Program

February 25-March 1, 2019
Spokane Convention Center
Spokane, Washington, U.S.

Sponsored by the SIAM Activity Group on Computational Science and Engineering (CSE)

This activity group fosters collaboration and interaction among applied mathematicians, computer scientists, domain scientists and engineers in those areas of research related to the theory, development, and use of computational technologies for the solution of important problems in science and engineering. The activity group promotes computational science and engineering as an academic discipline and promotes simulation as a mode of scientific discovery on the same level as theory and experiment.

SIAM Events Mobile App

Scan the QR code with any QR reader and download the TripBuilder EventMobile™ app to your iPhone, iPad, iTouch or Android mobile device.

You can also visit *www.tripbuildermedia.com/apps/siamevents*

Society for Industrial and Applied Mathematics
3600 Market Street, 6th Floor
Philadelphia, PA 19104-2688 U.S.
Telephone: +1-215-382-9800 Fax: +1-215-386-7999
Conference E-mail: meetings@siam.org
Conference Web: www.siam.org/meetings/
Membership and Customer Service:
(800) 447-7426 (U.S. & Canada) or +1-215-382-9800 (worldwide)
https://www.siam.org/conferences/CM/Main/cse19

Table of Contents

Organizing Committee Co-Chairs

Jeffrey Hittinger
Lawrence Livermore National Laboratory, U.S.

Luke Olson
University of Illinois at Urbana-Champaign, U.S.

Suzanne Shontz
University of Kansas, U.S.

Organizing Committee

Katherine J. Evans
Oak Ridge National Laboratory, U.S.

Hans Johansen
Lawrence Berkeley National Laboratory, U.S.

Matthew G. Knepley
University at Buffalo, U.S.

Tamara G. Kolda
Sandia National Laboratories, U.S.

Alison Marsden
Stanford University, U.S.

Marie E. Rognes
Simula Research Laboratory, Norway

Robert Scheichl
University of Heidelberg, Germany

Andrea Walther
Paderborn University, Germany

Stefan Wild
Argonne National Laboratory, U.S.

Rebecca Willett
University of Chicago, U.S.

Rio Yokota
Tokyo Institute of Technology, Japan

Conference Themes

Computational science and machine learning
Statistical modeling, methods, and computation
Multiscale, multiphysics, and multilevel methods
High performance software: packages and design
Algorithms at extreme scales
Tensor computations
High-order methods, novel discretizations, and scalable solvers
Data science, analytics, and visualization
Applications in science, engineering, and industry
Biological and biomedical computations
Scientific simulation and uncertainty
Numerical optimization: methods and applications
Reduced order modeling
Emerging trends in CS&E education and training

Spokane Convention Center

Unless otherwise specified, all sessions take place at the convention center.

334 W Spokane Falls Blvd

Spokane, WA 99201

SIAM Registration Desk

The SIAM registration desk is in the Ballroom Foyer of the Spokane Convention Center. It is open during the following hours:

Sunday, February 24
3:00 p.m. - 7:30 p.m.

Monday, February 25
7:30 a.m. - 4:00 p.m.

Tuesday, February 26
7:30 a.m. - 4:00 p.m.

Wednesday, February 27
8:00 a.m. - 4:00 p.m.

Thursday, February 28
7:45 a.m. - 4:30 p.m.

Friday, March 1
8:00 a.m. - 12:00 p.m.

Child Care

Visit Spokane recommends The Little Red Schoolhouse (509-456-3776) for attendees interested in child care services. Attendees are responsible for making their own child care arrangements.

Corporate Members and Affiliates

SIAM corporate members provide their employees with knowledge about, access to, and contacts in the applied mathematics and computational sciences community through their membership benefits. Corporate membership is more than just a bundle of tangible products and services; it is an expression of support for SIAM and its programs. SIAM is pleased to acknowledge its corporate members and sponsors. In recognition of their support, non-member attendees who are employed by the following organizations are entitled to the SIAM member registration rate.

Corporate/Institutional Members

The Aerospace Corporation
Air Force Office of Scientific Research
Amazon
Argonne National Laboratory
Bechtel Marine Propulsion Laboratory
The Boeing Company
CEA/DAM
Cirrus Logic
Department of National Defence (DND/CSEC)
DSTO- Defence Science and Technology Organisation, Edinburgh
Exxon Mobil
IDA Center for Communications Research, La Jolla
IDA Institute for Defense Analyses, Princeton
IDA Institute for Defense Analyses, Bowie, Maryland
Lawrence Berkeley National Laboratory

Lawrence Livermore National Labs
Lockheed Martin Maritime Systems & Sensors
Los Alamos National Laboratory
Max-Planck-Institute
Mentor Graphics
National Institute of Standards and Technology (NIST)
National Security Agency
Oak Ridge National Laboratory
Sandia National Laboratories
Schlumberger
Simons Foundation
United States Department of Energy
U.S. Army Corps of Engineers, Engineer Research and Development Center

List current as of December 2018.

Funding Agencies

SIAM and the conference organizing committee wish to extend their thanks and appreciation to the U.S. National Science Foundation and the Department of Energy for their support of this conference.

Join SIAM and save!

Leading the applied mathematics community . . .

SIAM members save up to $140 on full registration for the 2019 SIAM Conference on Computational Science and Engineering! Join your peers in supporting the premier professional society for applied mathematicians and computational scientists. SIAM members receive subscriptions to *SIAM Review*, *SIAM News* and *SIAM Unwrapped*, and enjoy substantial discounts on SIAM books, journal subscriptions, and conference registrations.

If you are not a SIAM member and paid the *Non-Member* rate to attend, you can apply the difference of $140 between what you paid and what a member paid towards a SIAM membership. Contact SIAM Customer Service for details or join at the conference registration desk.

If you are a SIAM member, it only costs $15 to join the SIAM Activity Group on Computational Science and Engineering (SIAG/CSE). As a SIAG/CSE member, you are eligible for an additional $15 discount on this conference, so if you paid the SIAM member rate to attend the conference, you might be eligible for a free SIAG/CSE membership. Check at the registration desk.

Students who paid the *Student Non-Member* Rate will be automatically enrolled as SIAM Student Members. Please go to *my.siam.org* to update your education and contact information in your profile. If you attend a SIAM Academic Member Institution or are part of a SIAM Student Chapter you will be able to renew next year for free.

Join onsite at the registration desk, go to *https://www.siam.org/Membership/Join-SIAM* to join online or download an application form, or contact SIAM Customer Service:

Telephone: +1-215-382-9800 (worldwide); or 800-447-7426 (U.S. and Canada only)

Fax: +1-215-386-7999

Email: membership@siam.org

Postal mail: Society for Industrial and Applied Mathematics, 3600 Market Street, 6th floor, Philadelphia, PA 19104-2688 U.S.

Standard Audio/Visual Set-Up in Meeting Rooms

SIAM does not provide computers for any speaker. When giving an electronic presentation, speakers must provide their own computers. SIAM is not responsible for the safety and security of speakers' computers.

A data (LCD) projector and screen will be provided in all technical session meeting rooms. The data projectors support both VGA and HDMI connections. Presenters requiring an alternate connection must provide their own adaptor.

If you have questions regarding availability of equipment in the meeting room of your presentation, please see a SIAM staff member at the registration desk.

Internet Access

Complimentary wireless internet access is available in the Spokane Convention Center. In addition, a limited number of computers with internet access will be available during registration hours.

Registration Fee Includes

Admission to all technical sessions
Business Meeting (open to SIAG/CSE members)
Coffee breaks daily
Room set-ups and audio/visual equipment
Poster Sessions
Welcome Reception

In addition, the following events are available to attendees at no additional costs. These events are subsidized by SIAM and are not covered by the registration fees:

Career Fair
Industry Reception

Job Postings

Please check at the SIAM registration desk regarding the availability of job postings or visit *jobs.siam.org*.

Important Notice to Poster Presenters

Poster boards and push pins will be available to presenters beginning 12:00 p.m. on Tuesday, February 26. Materials must be posted by 4:30 p.m. on Tuesday, February 7. Presenters are requested to display their posters from 4:30 p.m. on Tuesday, February 26 through 6:50 p.m. on Wednesday, February 27. Presenters should plan to stand by their poster during the session they've been assigned. If you are unsure which session you have been assigned, please refer to the speaker index in this book or check the online program. **Posters must be removed by 7:00 p.m. on Wednesday, February 27.**

Poster Blitzes and Plenary Poster Sessions including Minisymposteria

There are two plenary poster sessions, each include Minisymposteria. Minisymposteria are collections of three or more posters by different presenters grouped around a central theme.

Poster sessions will take place in Riverside Hall D on Tuesday, February 26 and Wednesday, February 27, 4:50 p.m. – 6:40 p.m. Poster Blitzes will precede each poster session at 4:10 p.m. in Ballroom 100BC.

Presenters have been requested to display their posters during both poster sessions from 4:30 p.m. on Tuesday, February 26 through 6:50 p.m. on Wednesday, February 27.

These sessions provide a great opportunity to network and are a great prelude to dinner in Spokane!

SIAM Books and Journals

Display copies of books and complimentary copies of journals are available on site. SIAM books are available at a discounted price during the conference. The books booth will be staffed from 9:00 a.m. – 5:00 p.m. Monday through Thursday, and 9:00 a.m. - 11:30 a.m. on Friday. If a SIAM books representative is temporarily away from the booth, completed order forms and payment (credit cards are preferred) may be taken to the SIAM registration desk. The books table will close at 11:30 a.m. on Friday.

Table Top Displays

Arm Ltd
Association for Women in Mathematics
Cambridge University Press
ICIAM 2019
National Academies Board on Mathematical Sciences and Analytics
Princeton University Press
Resume Doctor
SIAM
Springer Nature
Students @SC19

Thank you to our Welcome Reception sponsor!

2019 Conference Bag Sponsor

Supporter

Name Badges

A space for emergency contact information is provided on the back of your name badge. Help us help you in the event of an emergency!

Comments?

Comments about SIAM meetings are encouraged! Please send to:

Cynthia Phillips, SIAM Vice President for Programs (vpp@siam.org).

Get-togethers

Welcome Reception
Sunday, February 24
5:45 p.m. - 7:45 p.m

Poster Sessions
Tuesday, February 26
4:50 p.m. - 6:50 p.m.
Wednesday, February 27
4:50 p.m. - 6:50 p.m.

Business Meeting
(open to SIAG/CSE members)
Wednesday, February 27
7:00 p.m. – 8:00 p.m.

Complimentary beer and wine will be served.

Statement on Inclusiveness

As a professional society, SIAM is committed to providing an inclusive climate that encourages the open expression and exchange of ideas, that is free from all forms of discrimination, harassment, and retaliation, and that is welcoming and comfortable to all members and to those who participate in its activities. In pursuit of that commitment, SIAM is dedicated to the philosophy of equality of opportunity and treatment for all participants regardless of gender, gender identity or expression, sexual orientation, race, color, national or ethnic origin, religion or religious belief, age, marital status, disabilities, veteran status, field of expertise, or any other reason not related to scientific merit. This philosophy extends from SIAM conferences, to its publications, and to its governing structures and bodies. We expect all members of SIAM and participants in SIAM activities to work towards this commitment.

Please Note

SIAM is not responsible for the safety and security of attendees' computers. Do not leave your laptop computers unattended. Please remember to turn off your cell phones, tablets, etc. during sessions.

Recording of Presentations

Audio and video recording of presentations at SIAM meetings is prohibited without the written permission of the presenter and SIAM.

Social Media

SIAM is promoting the use of social media, such as Facebook and Twitter, to enhance scientific discussion at its meetings and enable attendees to connect with each other prior to, during and after conferences. If you are tweeting about a conference, please use the designated hashtag to enable other attendees to keep up with the Twitter conversation and to allow better archiving of our conference discussions. The hashtag for this meeting is #SIAMCSE19.

SIAM's Twitter handle is @TheSIAMNews.

Changes to the Printed Program

The printed program was current at the time of printing, however, please review the online program schedule (http://meetings.siam.org/program.cfm?CONFCODE=CS19) or use the mobile app for up-to-date information.

Award Announcements

On Thursday, February 28, 8:00-8:15 a.m. the award announcements will take place. The following prizes will be acknowledged and announced: SIAM/ACM Prize in Computational Science and Engineering, SIAM SIAG CSE Best Paper Prize, SIAM CSE19 Poster prizes, SIAM James H. Wilkinson Prize for Numerical Software, and the Bavarian Graduate Computational Engineering (BGCE) Student Paper Prize.

SIAM Events Mobile App

Scan the QR code with any QR reader and download the TripBuilder EventMobile™ app to your iPhone, iPad, iTouch or Android mobile device.

You can also visit *www.tripbuildermedia.com/apps/siam*

SIAM Activity Group on Computational Science and Engineering (SIAG/CSE)

www.siam.org/activity/cse

ACTIVITIES INCLUDE

- Special sessions at SIAM Annual Meetings
- Biennial conference
- Wiki

BENEFITS OF SIAG/CSE MEMBERSHIP

- Additional $15 discount on registration at the SIAM Conference on Computational Science and Engineering (excludes student)
- Electronic communications about recent developments in your specialty
- Eligibility for candidacy for SIAG/CSE office
- Participation in the selection of SIAG/CSE officers

ELIGIBILITY

- Be a current SIAM member

COST

- $15 per year
- Student members can join two activity groups for free!

2019-20 SIAG/CSE OFFICERS

Chair:	Karen Devine, Sandia National Laboratories
Vice-Chair:	Suzanne Shontz, University of Kansas
Program Director:	Stefan Wild, Argonne National Laboratory
Secretary:	Judith Hill, Oak Ridge National Laboratory

TO JOIN

SIAG/CSE:	*my.siam.org/forms/join_siag.htm*
SIAM:	*siam.org/joinsiam*

A great way to get involved!

Collaborate and interact with mathematicians and applied scientists whose work involves computational science and engineering.

Graduating? Earning your final degree?

"Awarding a PhD is just the beginning in our professions. Needless to say, interacting with the pillars of our community is a must in preparing ourselves for future career goals. Being able to remain a SIAM member thanks to the discounted rate in the years following the PhD degree allows one to have more chances for these interactions."

– Necibe Tuncer, Florida Atlantic University

SIAM offers a reduced rate membership for individuals early in their careers

Individuals who have graduated or received their final degree within the last five years are eligible for a reduced rate SIAM membership with the same benefits as a regular member—for less!

SIAM offers early career membership at 50% off the price of regular membership for the first three years after receiving a final degree and 25% off for the fourth and fifth years.

If you graduated this year, you can remain a member for just 22 cents a day!

If you are a SIAM student member, notify SIAM that you would like to renew next year as an early career member.

Starting a career can be a little unsettling. SIAM can help you through the rough spots in lots of ways:

☑ Develop your career
SIAM has numerous resources for finding jobs, including the SIAM job board, professional ads in *SIAM News*, and various activity group electronic mailing lists. Plus, SIAM's science policy electronic mailing list can keep you informed about funding opportunities and issues affecting applied mathematics and computational science.

☑ Networking Opportunities
SIAM membership helps graduates make the transition from completing their education to building a career. With a wealth of resources, SIAM will support your journey as you grow professionally. Plus, you can make a difference to your profession by participating in activity groups, presenting at SIAM conferences, and volunteering to serve on SIAM committees.

☑ Discounts
You'll receive generous discounts on SIAM conference registrations, books, and journals. Take advantage of the specially-reduced conference fee for the SIAM Annual Meeting available only to SIAM early career members.

☑ Keep up-to-date on what's happening in the field
SIAM members receive *SIAM Review*, a quarterly publication providing an overview of applied mathematics, in print as well as in electronic format. You'll also receive *SIAM News*, the news journal of the applied mathematics community.

☑ Additional Benefits
As an early career member, these additional benefits are available to you:

- Vote, hold office, and serve on SIAM committees.
- Nominate two students for free membership.
- Nominate eligible colleagues for the SIAM Fellows program and begin to accumulate the years of membership that will qualify you to be nominated as a SIAM Fellow.

For more information about becoming an early career member, contact Membership Manager Tim Fest at fest@siam.org.

siam | Society for Industrial and Applied Mathematics
3600 Market Street, 6th Floor, Philadelphia, PA 19104-2688 U.S.
+1-215-382-9800 · Fax: +1-215-386-7999 · membership@siam.org · www.siam.org

12/18

Invited Plenary Speakers

**** All Plenary Lectures are scheduled take place in Ballroom 100BC****

Monday, February 25

8:30 a.m. - 9:15 a.m.

IP1 Modelling 100 Percent Renewable Electricity

Michael C. Ferris, *University of Wisconsin, Madison, U.S.*

1:00 p.m. – 1:45 p.m.

IP2 Deep Learning for Inverse Problems - Some Recent Approaches

Carola-Bibiane Schönlieb, *University of Cambridge, United Kingdom*

Tuesday, February 26

8:30 a.m. - 9:15 a.m.

IP3 Reduced Order Methods for PDEs: State of the Art and Perspectives with Applications in Industry, Medicine and Environmental Sciences

Gianluigi Rozza, *SISSA, International School for Advanced Studies, Trieste, Italy*

1:00 p.m. - 1:45 p.m.

IP4 Fluid-Structure Interaction in Medicine and Biology: Methods, Models, and Applications

Boyce E. Griffith, *University of North Carolina at Chapel Hill, U.S.*

Wednesday, February 27

8:30 a.m. - 9:15 a.m.

IP5 Communication Avoiding: The Past Decade and the New Challenges

Laura Grigori, *Inria Paris, France*

1:00 p.m. - 1:45 p.m.

IP6 Stochastic Gradient Descent, in Theory and Practice

Rachel Ward, *University of Texas at Austin, U.S.*

Thursday, February 28

(see next page for prize lectures scheduled for Thursday morning)

1:00 p.m. - 1:45 p.m.

IP7 Physical, Numerical, and Computational Challenges in Modelling Oceans for Climate

Alistair Adcroft, *Princeton University and NOAA-GFDL, U.S.*

Friday, March 1

8:30 a.m. - 9:15 a.m.

IP8 Role of Tensors in Machine Learning

Anima Anandkumar, *Amazon and California Institute of Technology, U.S*

Invited Plenary Speakers – Prize Lectures

**** All Plenary Lectures are scheduled take place in Ballroom 100BC ****

Wednesday, February 27

11:30 a.m. - 12:00 p.m.

SP1 SIAG CSE Best Paper Prize Lecture

Recipient To Be Announced

Thursday, February 28

8:15 a.m. - 8:45 a.m.

SP2 SIAG/CSE Early Career Prize

Data-Driven Discovery and Control of Complex Systems: Uncovering Interpretable and Generalizable Nonlinear Models

Steven Brunton, *University of Washington, U.S.*

8:45 a.m. - 9:15 a.m.

SP3 SIAM/ACM Prize in Computational Science and Engineering

The Singular Value Decomposition: Anatomy of an Algorithm, Optimizing for Performance

Jack J. Dongarra, *University of Tennessee and Oak Ridge National Laboratory, U.S.*

11:30 a.m. – 12:00 p.m.

SP4 James H. Wilkinson Prize for Numerical Software

Solving the Two Language Problem in Scientific Computing and Machine Learning with Julia

Jeffrey Bezanson, Stefan Karpinski, and Viral Shah, *Julia Computing, Inc., U.S.*

Minitutorials

***** All Minitutorials will take place in room 300D*****

Tuesday, February 26

9:45 a.m. - 11:25 a.m.

MT1 Introduction To PETSc

Jed Brown, *University of Colorado Boulder, U.S.*

Tobin Isaac, *Georgia Institute of Technology, U.S.*

Matthew G. Knepley, *State University of New York at Buffalo, U.S.*

Wednesday, February 27

9:45 a.m. - 11:25 a.m.

MT2 Solving Differential Equations with the libMesh Finite Element Library

Paul Bauman, *State University of New York at Buffalo, U.S.*

John W. Peterson, *Idaho National Laboratory, U.S.*

Roy Stogner, *University of Texas at Austin, U.S.*

Thursday, February 28

9:45 a.m. - 11:25 a.m.

MT3 MOOSE: Enabling Multiphysics

Derek R. Gaston, *Idaho National Laboratory, U.S.*

Andrew Slaughter, *Idaho National Laboratory, U.S.*

Friday, March 1

9:45 a.m. - 11:25 a.m.

MT4 Firedrake: Automated High Performance Finite Element Simulation

Thomas H. Gibson, *Imperial College London, United Kingdom*

David Ham, *Imperial College London, United Kingdom*

Lawrence Mitchell, *Durham University, United Kingdom*

Featured Minisymposia

Selected by the CSE19 Organizing Committee, these minisymposia provide overview presentations on fundamental advances in fields related to the conference themes. Each featured minisymposium will begin with a broadly accessible introduction to the topic, followed by presentations about important advances in the area. These sessions are a great opportunity for conference attendees to broaden their horizons and learn about important advances in areas with which they may not be fully familiar.

Monday, February 25

MS1 Mathematical Advances in Deep Learning - Part I of II

MS8 Uncertainty Quantification and Data Assimilation - Part I of II

MS35 Mathematical Advances in Deep Learning - Part II of II

MS42 Uncertainty Quantification and Data Assimilation - Part II of II

MS78 Optimization with Coupled PDEs in Multiphysics Applications

MS95 Performance Portability and Numerical Libraries: Challenges and Opportunities for Sustainable Science - Part I of II

MS101 Scientific Data Visualization Platforms Facilitating New Paradigms

Tuesday, February 26

MS118 Hydrodynamics at Small Scales: Fluctuating Hydrodynamics - Part I of II

MS119 Model Reduction for Problems with Strong Convection, Sharp Gradients, and Discontinuities - Part I of II

MS127 Design and Usability of High-performance PDE Software Engines and Frameworks - Part I of II

MS129 Performance Portability and Numerical Libraries: Challenges and Opportunities for Sustainable Science - Part II of II

MS152 Hydrodynamics at Small Scales: Fluctuating Hydrodynamics - Part II of II

MS153 Model Reduction for Problems with Strong Convection, Sharp Gradients, and Discontinuities - Part II of II

MS161 Design and Usability of High-performance PDE Software Engines and Frameworks - Part II of II

Wednesday, February 27

MS195 CSE Education and Workforce - Part I of II

MS229 CSE Education and Workforce - Part II of II

Thursday, February 28

MS239 Multiphysics: Extensible, Composable Algorithms and Software - Part I of II

MS260 Applications of Data Assimilation in Science, Engineering, and Industry - Part I of II

MS272 Multiphysics: Extensible, Composable Algorithms and Software - Part II of II

MS294 Applications of Data Assimilation in Science, Engineering, and Industry - Part II of II

Friday, March 1

MS314 Optimal Experimental Design for Bayesian Inverse Problems

MS353 Mathematics of Energy Materials - Part I of II

MS363 Summation-by-Parts: A Framework for the Development and Analysis of Modern Numerical Methods - Part I of II

MS385 Mathematics of Energy Materials - Part II of II

MS395 Summation-by-Parts: A Framework for the Development and Analysis of Modern Numerical Methods - Part II of II

Career Fair

Careers in Business, Industry, and Government

Tuesday, February 26

Job Fair Session I: 9:45 a.m. – 11:45 a.m.

Job Fair Session II: 2:00 p.m. – 4:00 p.m.

Room: *Ballroom 100A*

The career fair will feature representatives from nonacademic employers from industry and government. These representatives will be prepared to discuss with you the opportunities for internships, postdoctoral appointments and full-time jobs at their organizations.

The career fair will feature morning and afternoon sessions during which you will have the opportunity to speak with the representatives of the participating organizations.

To maximize the value of your discussion, SIAM recommends that you take time beforehand to learn about the organizations that interest you.

SIAM is pleased to share the list of organizations that will participate in the 2019 Career Fair. This list is current at time of printing. The most up to date list of participants can be found at *https://www.siam.org/conferences/CM/P/CF/cse19-career-fair.*

- **Argonne National Laboratory**
- **Department of Energy Computational Science Graduate Fellowship (DOE CSGF)**
- **ExxonMobil Upstream Research Company**
- **Los Alamos National Laboratory**
- **MathWorks**
- **NASA Langley Research Center**
- **Naval Nuclear Laboratory**
- **Nokia Bell Labs**
- **Oak Ridge National Laboratory**
- **Pacific Northwest National Laboratory**
- **Raytheon**
- **Sandia National Laboratories**
- **U.S. Army Research and Development Center (ERDC)**

Student Days

Student Days at the 2019 SIAM Conference on Computational Science and Engineering are for students and about students. Organizers have got it all covered this year, with activities and sessions where students can meet with both peers and professionals in their field, participate in a career fair, attend an information session on hot areas for jobs and research, and network with SIAM Student Chapters from all over the world.

Goals

Organized by the SIAM Education Committee (Chaired by Suzanne Weekes, Worcester Polytechnic Inst), Student Days are designed to encourage student participation in SIAM, to help students learn more about applied mathematics and computational science as both fields of study and as careers, and to provide a forum for emerging mathematicians to learn about their field from the professionals who know the answers. Organizers also hope to encourage those in the learning community to establish new student chapters of SIAM and to promote interaction between students and SIAM leadership.

Events and Happenings

Student Days sessions include presentations and posters by student chapter. In addition, students can attend plenary sessions from the CSE19 schedule. Other activities that will enhance the meeting for students include a special orientation prior to Sunday's welcome reception, Tutorials for Students, an Early Career Panel, and Prize sessions. Other activities that will be of interest to students include the career fair, evening Welcome Reception on Sunday, and the poster sessions on Tuesday and Wednesday.

Sunday, February 24

4:45 p.m. - 5:45 p.m. Student Days: Student Orientation
5:45 p.m. - 7:45 p.m. Welcome Reception

Monday, February 25

9:45 a.m. - 11:25 a.m. MS4 Student Days: Student Chapter Presentations - Part I of II
9:45 a.m. - 11:25 a.m. MS16 Tutorials for Students: Accessible Introductions to Active Research Areas
11:30 a.m. - 12:30 p.m. PD2 Early Career Panel
2:15 p.m. - 3:55 p.m. MS52 Tutorials for Students: Accessible Introductions to Active Research Areas - Part II of II
4:10 p.m. - 5:50 p.m. MS73 Student Days: Student Chapter Presentations - Part II of II

Tuesday, February 26

7:00 a.m. - 8:15 a.m. Student Days: Chapter Breakfast with SIAM Leadership (by invitation)
9:45 a.m. - 11:25 a.m. MS130 Computational Engineering (BGCE) Student Paper Prize - Part I of II
9:45 a.m. - 11:45 a.m. Career Fair
11:30 a.m. - 12:30 p.m. PD Mentor-Mentee Mixer
2:00 p.m. - 4:00 p.m. Career Fair
2:15 p.m. - 3:55 p.m. MS164 Computational Engineering (BGCE) Student Paper Prize - Part II of II
4:50 p.m. – 6:50 p.m.PP1 Poster Session
6:30 p.m. - 8:00 p.m. Industry Reception
7:00 p.m. - 8:00 p.m. Book Panel & Reception

Wednesday, February 27

9:45 a.m. - 11:25 a.m. MS173 Student Days: Undergraduate Presentations - Part I of II
4:50 p.m. - 6:50 p.m. PP204 Minisymposterium: Student Days - Student Chapter Posters
4:50 p.m. - 6:50 p.m. PP205 Minisymposterium: Student Days - Undergraduate Posters

Thursday, February 28

8:00 a.m. - 8:15 a.m. Awards Announcement
8:15 a.m. - 8:45 a.m. SIAM Activity Group on Computational Science and Engineering Early Career Prize
8:45 a.m. - 9:15 a.m. SIAM/ACM Prize in Computational Science and Engineering

Workshops

Workshop Celebrating Diversity (WCD)

This annual event provides a chance for students to listen to technical talks presented by minority graduate students. The workshop is intended to accomplish several goals:

- To send a clear, explicit message of enthusiastic welcome and support from SIAM to members of under-represented groups. The workshop is deliberately held as part of a regular SIAM meeting so that the participants can combine the experiences of attending a regular scientific meeting and a special occasion dedicated to them.
- To bring together a mixture of people from different levels of age and professional experience, ranging from undergraduate students to senior scientists.
- To provide an opportunity for minority graduate students to present their research.
- To provide an informal, comfortable setting (a lunch) where all the students can meet applied and computational mathematicians with a wide variety of jobs in academia, national laboratories, industry, and government.

All registered attendees are invited to attend Workshop Celebrating Diversity sessions.

The following sessions are part of the Workshop Celebrating Diversity.

MS135 WCD Workshop **MS169 WCD Workshop**
MS203 WCD Workshop **MS237 WCD Workshop**

Check the online program for an update on these sessions.

SIAM Workshop Celebrating Diversity and Sustainable Horizons Institute Broader Engagement Program are hosting a joint lunch (by invitation) on Wednesday, February 27, 11:30 a.m.-1:00 p.m. Those participating in the programs are encouraged to attend.

Association for Women in Mathematics (AWM) Workshop

The Workshop for Women Graduate Students and Recent P.h.D.'s is held in conjunction with the SIAM Annual Meeting. All registered attendees are invited to attend the AWM Workshop technical sessions. See https://awm-math.org/meetings/awm-siam/ for additional information.

The following sessions are part of the AWM Workshop.
MS3 AWM Workshop: Data Science and Mathematics - Part I of II
MS37 AWM Workshop: Data Science and Mathematics - Part II of II
MS147 AWM Workshop Panel: Perspectives and Advice from Women in Research
PP101 AWM Workshop: Poster Session

Sustainable Horizons Institute Broader Engagement and Mentoring Program

The goals of the program are to bring a diverse group of students, faculty, and professional to the conference, engage them in conference activities and the community, and engage the CSE community in learning about and supporting diversity and inclusion.

The following sessions are part of this program.
MS13 BE: Mentor Protégé
MS48 BE: Best Practices for CSE Diversity and Inclusion
MS82 BE: Lightning Talks
MS114 BE: Broader Engagement Technical Research - Part I of II
MS115 BE: Securing Extreme-Scale Scientific Computing
MS148 BE: Broader Engagement Technical Research - Part II of II
MS149 BE: Social Modeling through Game Design, A STEAM Workshop
MS182 BE: GPU Accelerated Computing on Summit
MS215 BE: Parallel Algorithm Design
MS250 BE: Advances in Computational Drug Discovery - Part I of II
MS283 BE: Advances in Computational Drug Discovery - Part II of II
MS317 BE: Overcoming Workplace Challenges Panel: Trials, Tribulations, and Triumphs
MS349 BE: Wrap-up Session

SIAM Workshop Celebrating Diversity and Sustainable Horizons Institute Broader Engagement Program are hosting a joint lunch (by invitation) on Wednesday, February 27, 11:30 a.m.-1:00 p.m. Those participating in the programs are encouraged to attend.

Program Schedule

February 25-March 1, 2019
Spokane Convention Center
Spokane, Washington, U.S.

Sunday, February 24

Registration
3:00 p.m.-7:30 p.m.

Room: Ballroom Foyer

Broader Engagement (BE): Welcome and Orientation
3:45 p.m.-4:45 p.m.

Room: 302A

Student Days: Student Orientation
4:45 p.m.-5:45 p.m.

Room: Conference Theater

Welcome Reception


5:45 p.m.-7:45 p.m.

Room: Ballroom Foyer

Monday, February 25

Data Science Book Series Editorial Board Meeting
7:00 a.m.-8:15 a.m.

Room: Davenport Hotel -- Meeting Room 11

Registration
7:30 a.m.-4:00 p.m.

Room: Ballroom Foyer

Welcome Remarks
8:15 a.m.-8:30 a.m.

Room: Ballroom 100BC

Monday, February 25

IP1

Modelling 100 Percent Renewable Electricity

8:30 a.m.-9:15 a.m.

Room: Ballroom 100BC

Chair: Suzanne M. Shontz, University of Kansas, U.S.

In the past few decades, power grids across the world have become dependent on markets that aim to efficiently match supply with demand at all times via a variety of pricing and auction mechanisms. These markets are based on models that capture interactions between producers, transmission and consumers. The recent explosion in the use of renewable supply such as wind, solar and hydro has led to increased volatility in this system at a number of different time scales. Nevertheless, various countries have announced target dates for a 100% renewable electricity system. Such targets require long-term investment planning, medium-term storage management (including batteries and pumped storage), as well as a short-term analysis of demand-response, involuntary load curtailment and transmission congestion. We develop models that aim to ensure enough generation capacity for the long term under various constraints related to environmental concerns, and consider the recovery of costs for this enhanced infrastructure. We demonstrate how risk can impose significant costs on the system that are not modeled in the context of socially optimal power system markets and highlight the use of contracts to reduce or recover these costs. We discuss how models can be used to compare several 100% renewable investment scenarios, and consider the performance in terms of investment and fuel costs, emissions, electricity prices, and lost-load.

Michael C. Ferris
University of Wisconsin, Madison, U.S.

Coffee Break
9:15 a.m.-9:45 a.m.

Room: Ballroom Foyer

Monday, February 25

MS1

Mathematical Advances in Deep Learning - Part I of II

9:45 a.m.-11:25 a.m.

Room: Ballroom 100BC

For Part 2 see MS35
Featured Minisymposium

Training neural networks with many hidden layers, aka Deep Learning, has become a booming area of machine learning. Despite enormous success in applications including computer vision and speech recognition, several critical challenges in designing, training, and understanding deep networks persist. This minisymposium presents recent advances toward rigorous mathematical underpinnings of deep learning and more efficient training. Speakers will report on recent success enabled by applying techniques from partial differential equations, optimization, optimal control, inverse problems, uncertainty quantification, to the deep learning problem.

Organizer: Lars Ruthotto

Emory University, U.S.

Organizer: Eldad Haber

University of British Columbia, Canada

9:45-10:05 Adversarial Regularizers in Inverse Problem

Sebastian Lunz and Carola-Bibiane Schönlieb, University of Cambridge, United Kingdom; Ozan Öktem, KTH Royal Institute of Technology, Sweden

10:10-10:30 Simple and Effective PDE Based Approaches to Deep Learning

Stanley J. Osher and Bao Wang, University of California, Los Angeles, U.S.; Li Zhen, Hong Kong University of Science and Technology, Hong Kong; Wei Zhu, University of California, Los Angeles, U.S.; Zuoqiang Shi, Tsinghua University, China; Adam M. Oberman, McGill University, Canada; Pratik A. Chaudhari, Penghang Yin, Xiyang Luo, Minh Pham, and Alex Lin, University of California, Los Angeles, U.S.

10:35-10:55 Simultaneous Parallel-in-layer Optimization for Training of Deep Residual Networks

Stefanie Guenther, Technische Universität Kaiserslautern, Germany; Lars Ruthotto, Emory University, U.S.; Jacob B. Schroder, Lawrence Livermore National Laboratory, U.S.; Nicolas R. Gauger, Technische Universität Kaiserslautern, Germany

11:00-11:20 Neural Ordinary Differential Equations

Ricky T. Q. Chen, David Duvenaud, Yulia Rubanova, and Jesse Bettencourt, University of Toronto, Canada

continued in next column

Monday, February 25

MS2

Scientific Software: Practices, Concerns, and Solution Strategies - Part I of II

9:45 a.m.-11:25 a.m.

Room: Conference Theater

For Part 2 see MS36

Software is an increasingly important component of scientific research. It would not be an exaggeration to say that almost all forms of scientific research today involve software, whether to compute, to manage, or to analyze. The kinds of software used and the associated issues, which range from technical to sociological, are as diverse as the scientific communities themselves. This minisymposium reflects this rich diversity, including presentations from software developers, users, and scientific project managers from various parts of the world and from various science communities. The common thread among these topics is that many overlapping subsets of scientific software issues are of interest to various participants. The minisymposium will, therefore, be a snapshot of the current state of practices, concerns and solution strategies in the world of scientific software.

Organizer: Anshu Dubey

Argonne National Laboratory, U.S.

9:45-10:05 Toward Automatic Generation of Scientific Software Artifacts

Spencer Smith, McMaster University, Canada

10:10-10:30 Runtime Data Analysis for CSE Applications

Alvaro Coutinho, Universidade Federal de Rio de Janeiro, Brazil; Marta Mattoso, Federal University of Rio de Janeiro, Brazil; José Camata, COPPE/Universidade Federal do Rio e Janeiro, Brazil; Vitor Souza, Linda Gsenhues, and Renan Souza, Federal University of Rio de Janeiro, Brazil

continued on next page

10:35-10:55 Development and Integration Workflows for Large Complex Distributed CSE Software Efforts

Roscoe A. Bartlett, Sandia National Laboratories, U.S.

11:00-11:20 How to Assure Quality of Software Preservation Early in a Project Life Cycle and Ongoing Efforts

Sandra Gesing, Natalie K. Meyers, Rick Johnson, and Zheng Wang, University of Notre Dame, U.S.

Monday, February 25

MS3

AWM Workshop: Data Science and Mathematics - Part I of II

9:45 a.m.-11:25 a.m.

Room: 102A

For Part 2 see MS37

Organizer: Deanna Needell
University of California, Los Angeles, U.S.

Organizer: Giseon Heo
University of Alberta, Canada

9:45-10:05 Stochastic Methods for High Dimensional Data with Multiple Measurement Vectors

Natalie Durgin, Spiceworks, U.S.; Rachel Grotheer, Goucher College, U.S.; Chenxi Huang, Yale University, U.S.; Shuang Li, Colorado School of Mines, U.S.; *Anna Ma*, University of California, San Diego, U.S.; Deanna Needell and Jing Qin, University of California, Los Angeles, U.S.

10:10-10:30 Compressed Anomaly Detection with Multiple Mixed Observations

Natalie Durgin, Spiceworks, U.S.; Rachel Grotheer, Goucher College, U.S.; Chenxi Huang, Yale University, U.S.; Shuang Li, Colorado School of Mines, U.S.; Anna Ma, University of California, San Diego, U.S.; Deanna Needell and Jing Qin, University of California, Los Angeles, U.S.

10:35-10:55 Heuristic Framework for Multi-scale Testing of the Multi-manifold Hypothesis

Karamatou Yacoubou Djima, Amherst College, U.S.; Ness Linda, Rutgers University, U.S.; Patricia Medina, Worcester Polytechnic Institute, U.S.; Melanie Weber, Princeton University, U.S.

11:00-11:20 Path-based Spectral Clustering: Guarantees, Robustness to Outliers, and Fast Algorithms

Anna Little, Michigan State University, U.S.; James Murphy, Tufts University, U.S.; Mauro Maggioni, Johns Hopkins University, U.S.

Monday, February 25

MS4

Student Days: Student Chapter Presentations - Part I of II

9:45 a.m.-11:25 a.m.

Room: 102B

For Part 2 see MS73

Presentations given by students from SIAM Student Chapters.

See online program for an update on this session.

Monday, February 25

MS5

Recent Advances in Modeling and Numerical Analysis for Electronic Structure Calculations - Part I of II

9:45 a.m.-11:25 a.m.

Room: 102C

For Part 2 see MS39

The goal of this minisymposium is to bring experts working on new modeling approaches and algorithmic contributions for electronic structure problems in materials. For many cutting edge applications, materials have to been studied at the quantum level to predict their properties, such as transport (conduction), magnetic or optical response. Systems that are of interest to materials scientists are typically of increasing complexity and scales, leading to new challenges from a mathematical and computational perspective for the computation and characterization of the electronic structure of materials. One recent example among others is that of devices formed of multiple layers of recently discovered 2D materials such as graphene, boron nitride, etc. and which lack periodicity. The first talk of the session will give an introduction to electronic structure computations and the other invited speakers will present their latest developments in the field.

Organizer: Paul Cazeaux

University of Kansas, U.S.

9:45-10:05 Recent Developments of Quantum Embedding Theories

Lin Lin, University of California, Berkeley and Lawrence Berkeley National Laboratory, U.S.

10:10-10:30 Permutation-invariant Interatomic Potentials Based on Ab-initio Data

Geneviève Dusson, University of Warwick, United Kingdom

10:35-10:55 Electronic Structure Methods for 2D Materials

Matthias Maier, Texas A&M University, U.S.; Mitchell Luskin, University of Minnesota, U.S.

11:00-11:20 Electronic Structure of Incommensurate Materials using Momentum Space

Daniel Massatt, University of Minnesota, U.S.

continued in next column

Monday, February 25

MS6

Roundoff Error in High-Performance Implementations of CG/ Lanczos-type Solvers - Part I of II

9:45 a.m.-11:25 a.m.

Room: 102D

For Part 2 see MS40

The conjugate gradient and Lanczos algorithms are widely used methods for solving Hermitian positive definite linear systems, computing eigenvalues of Hermitian matrices and, more generally, for computing the product of a function of a Hermitian matrix with a given vector. It is well-known that the methods may behave differently in finite precision arithmetic than they would in exact arithmetic, which means that whenever different implementations are considered -- such as implementations that make better use of parallelism -- one must be careful about the effects on actual performance and accuracy. This minisymposium brings together researchers to talk about different mathematically equivalent CG/Lanczos implementations and their behavior in finite precision arithmetic.

Organizer: Erin C. Carson

Charles University, Czech Republic

Organizer: Anne Greenbaum

University of Washington, U.S.

9:45-10:05 Why Worry about Rounding Errors in Hermitian Krylov Subspace Methods?

Anne Greenbaum, University of Washington, U.S.

10:10-10:30 Computing Matrix Functions with the Finite Precision Lanczos Method: New Convergence Bounds and Limitations

Christopher Musco, Massachusetts Institute of Technology, U.S.

continued on next page

10:35-10:55 Characterizing Parallel Scalability of Asynchronous Krylov Solvers

Kevin DeWeese, University of Washington, U.S.

11:00-11:20 Spectral Clustering of Graph Vertex Subsets via Krylov Subspace Model Reduction

Alexander V. Mamonov, University of Houston, U.S.; Vladimir Druskin, Worcester Polytechnic Institute, U.S.; Mikhail Zaslavsky, Schlumberger-Doll Research, U.S.

Monday, February 25

MS7

Numerical Linear Algebra for Machine Learning - Part I of II

9:45 a.m.-11:25 a.m.

Room: 111A

For Part 2 see MS41

Machine learning algorithms such as deep neural networks with convolution and fully connected layers, recurrent neural networks, support vector machines, and kernel machines rely on linear algebra for matrix-vector and matrix-matrix products. Many other algorithms for "big data processing" rely on low-rank decomposition such as the thin SVD, rank-revealing QR factorizations, and randomized projection and sampling algorithms. In this MS, we will explore the connections between machine learning algorithms and numerical linear algebra, including techniques such as the fast multipole method, hierarchical matrices, structured matrices, tensor representations, tensor-trains, low-rank and sparse representations, and numerical algorithms for large-scale eigenvalue and singular value calculations.

Organizer: Leopold Cambier
Stanford University, U.S.

Organizer: Eric F. Darve
Stanford University, U.S.

9:45-10:05 Kronecker Factorization for Second Order Optimization in Deep Learning

Rio Yokota, Tokyo Institute of Technology, Japan

10:10-10:30 Using Deep Learning for the Solution of Partial Differential Equations

Ziyi Yang, Stanford University, U.S.

10:35-10:55 Fast Approximations of Hessians

George M. Turkiyyah, American University of Beirut, Lebanon and King Abdullah University of Science & Technology (KAUST), Saudi Arabia; Wajih Halim Boukaram and David E. Keyes, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

11:00-11:20 A Novel Compression Algorithm for the Construction of Hierarchically Semi-separable Matrices for Kernel Matrices

Gustavo Chavez, Lawrence Berkeley National Laboratory, U.S.; Rebrova Elizaveta, University of California, Los Angeles, U.S.; Pieter Ghysels and Sherry Li, Lawrence Berkeley National Laboratory, U.S.

continued in next column

Monday, February 25

MS8

Uncertainty Quantification and Data Assimilation - Part I of II

9:45 a.m.-11:25 a.m.

Room: 111B

For Part 2 see MS42
Featured Minisymposium

Data assimilation is the primary means of prediction and parameter estimation in weather forecasting, geophysics, fluid dynamics, and finance. Standard methods in filtering and Bayesian inference are particularly appealing considering the rapid advances in data acquisition in recent years. However, incorporating observation data with forward models introduces both model and measurement uncertainty into the prediction. As a result, these methods may require a large number of forward model evaluations to converge, which can be computationally prohibitive for high-dimensional problems. This minisymposium invites speakers working at the forefront of data assimilation and uncertainty quantification to discuss recent challenges, innovations, and strategies in these areas.

Organizer: Krithika Manohar
California Institute of Technology, U.S.

9:45-10:05 Scalable Bayesian Inference for Inverse Problems
Cheng Zhang, Fred Hutchinson Cancer Research Center, U.S.; *Babak Shahbaba* and Hongkai Zhao, University of California, Irvine, U.S.

10:10-10:30 Computer Model Calibration and Uncertainty Quantification for Binary Black Hole Formation
Derek Bingham, Simon Fraser University, Canada

10:35-10:55 Derivative Informed MCMC Methods for Subsurface Models with Faults
Jeonghun J. Lee, Baylor University, U.S.; Omar Ghattas and Tan Bui, University of Texas at Austin, U.S.; Umberto Villa, Washington University, St. Louis, U.S.

11:00-11:20 Tensorised Approximations for Intractable Probability Distributions
Tiangang Cui, Monash University, Australia; Sergey Dolgov, University of Bath, United Kingdom; Robert Scheichl, Universität Heidelberg, Germany

continued in next column

Monday, February 25

MS9

Avoiding Big-data: In-situ Visualisation and Analysis for Large-scale CFD Simulations - Part I of II

9:45 a.m.-11:25 a.m.

Room: 111C

For Part 2 see MS43
Modern-day CFD, especially time-resolved simulations of turbulent flow, can readily produce unmanageable amounts of data. Disk I/O and fundamental storage capacity preclude writing such data to magnetic disk for a posteriori analysis. Therefore, new technologies are required that can visualise and analyse these data on-the-fly 'in-situ' as they are generated. This minisymposia will explore the state-of-the-art in this area, including the development of frameworks for in-situ visualisation, data compression, feature identification, and feature tracking.

Organizer: Marius K. Koch
Imperial College London, United Kingdom

Organizer: Peter E. Vincent
Imperial College London, United Kingdom

Organizer: Paul Kelly
Imperial College London, United Kingdom

Organizer: Freddie Witherden
Imperial College London, United Kingdom

9:45-10:05 GPU Accelerated Feature Extraction on Unstructured Grids using Contour Trees
Marius K. Koch, Paul Kelly, and Peter E. Vincent, Imperial College London, United Kingdom

10:10-10:30 In Situ Analysis and Visualization Overview
Patrick O'Leary, Kitware, Inc., U.S.

10:35-10:55 The Development and use of in Situ Visualization and Analysis Approaches for the U.S. Exascale Program
James Ahrens, Los Alamos National Laboratory, U.S.

11:00-11:20 In-transit Data Analytics for Large-scale CFD
Steve Legensky, Intelligent Light, U.S.

Monday, February 25

MS10

Model Reduction, Adaptivity, and High Dimensionality in Uncertainty Quantification - Part I of II

9:45 a.m.-11:25 a.m.

Room: 300A

For Part 2 see MS44

Uncertainty Quantification (UQ) is essential for credible simulation-aided prediction and design. However, many computational strategies for both forward and inverse UQ are intractable when applied to models with a moderate or large number of parameters. Exploiting structure of the model input-output map and the goals of the overall analysis are needed to improve tractability. This minisymposium will focus on goal-oriented techniques whose aim is to reduce the cost of uncertainty quantification while simultaneously maintaining or improving accuracy. Areas of interest include dimension reduction, sparsity, approximations with non-tensorial measures, adaptivity, multi-level methods, and multi-fidelity methods. This minisymposium will bring together leading researchers from the applied and computational mathematics and engineering community to discuss and collaborate on novel theoretical and computational advances, and to discuss future directions for research.

Organizer: Akil Narayan

University of Utah, U.S.

Organizer: John D. Jakeman

Sandia National Laboratories, U.S.

9:45-10:05 A Multi-fidelity Spectral Technique for Uncertainty Quantification

Alireza Doostan and Felix Newberry, University of Colorado Boulder, U.S.

10:10-10:30 Energy Norm Regularized Sparse Simultaneous Reconstruction of Solutions to Parameterized PDEs

Nick Dexter, Simon Fraser University, Canada; Hoang A. Tran, Oak Ridge National Laboratory, U.S.; Clayton G. Webster, University of Tennessee and Oak Ridge National Laboratory, U.S.

10:35-10:55 Adaptive Multi-index Collication for Balancing Statistical and Deterministic Errors

John D. Jakeman, Michael S. Eldred, and Gianluca Geraci, Sandia National Laboratories, U.S.; Alex Gorodetsky, University of Michigan, U.S.

11:00-11:20 Sparse Grid Approximation of Elliptic PDEs with Lognormal Diffusion Coefficient

Lorenzo Tamellini, Istituto di Matematica Applicata e Tecnologie Informatiche-CNR, Italy

continued in next column

Monday, February 25

MS11

Data Enabled Modeling and Discovery in Science and Engineering - Part I of II

9:45 a.m.-11:25 a.m.

Room: 300B

For Part 2 see MS45

There is a growing amount of data generated from instruments and computational simulations associated with the scientific enterprise. The volume of data has driven an interest in scientific machine learning as a guided way to determine relationships and incorporate data into models. This presents a number of challenges unique to scientific machine learning of relevance to the applied mathematics community. For instance, developing scalable algorithms to run learning algorithms on leadership class platforms, incorporation of physical constraints into the model, dealing with heterogeneous data, and building trust through verified and validated results for the learning models. This minisymposium will bring together researchers developing algorithms and learning approaches to overcome these challenging issues. In addition, the discussion will include a motivation of the need for scientific machine learning while highlighting some additional challenges.

Organizer: Eric C. Cyr

Sandia National Laboratories, U.S.

Organizer: Paris Perdikaris

University of Pennsylvania, U.S.

Organizer: Prasanna Balaprakash

Argonne National Laboratory, U.S.

9:45-10:05 PDE-inspired Deep Neural Networks for Machine Learning and Inverse Problems

Lars Ruthotto, Emory University, U.S.

continued on next page

Monday, February 25

MS11

Data Enabled Modeling and Discovery in Science and Engineering - Part I of II

continued

10:10-10:30 Energy Minimizing Deep Neural Networks for Solving High-dimensional Stochastic Partial Differential Equations

Sharmila Karumuri, Ilias Bilionis, and Jitesh Panchal, Purdue University, U.S.

10:35-10:55 Towards Scalable Scientific Machine Learning: Motivation and Approaches

Eric C. Cyr, Sandia National Laboratories, U.S.

11:00-11:20 Data-driven Approximation of High-dimensional Flow Maps and Probability Distributions using Deep Neural Networks

Tenavi Nakamura-Zimmerer, Daniele Venturi, and Qi Gong, University of California, Santa Cruz, U.S.

Monday, February 25

MS12

Task-based Programming for Scientific Computing: Runtime Support - Part I of II

9:45 a.m.-11:25 a.m.

Room: 300C

For Part 2 see MS46

As the architecture of modern supercomputing systems is becoming increasingly complex and heterogeneous, runtime systems have recently emerged as a solid alternative to traditional parallel programming paradigms and tools because they allow for the delevopment of efficient and scalable code which is portable across a wide range of architectures thanks to the use of a modular approach and architecture-agnostic parallel programming paradigms. Numerous runtimes exist which propose different programming interfaces and features, this minisymposium presents recent advances and developments in some of the best known and used runtime systems.

Organizer: Alfredo Buttari
CNRS, France

Organizer: Emmanuel Agullo
Inria, France

Organizer: Hatem Ltaief
King Abdullah University of Science \\& Technology (KAUST), Saudi Arabia

9:45-10:05 Task-based Programming, from Sequential-looking Source Code to Exascale: What a Runtime System Can Help You With

Samuel Thibault, University of Bordeaux, France

10:10-10:30 Automatic Task-based Parallelization of Python Codes

Cristian Ramon-Cortes, Ramon Amela, and Jorge Ejarque, Barcelona Supercomputing Center, Spain; Philippe Clauss, INRIA and Université de Strasbourg, France; Rosa M. Badia, Barcelona Supercomputing Center, Spain

10:35-10:55 Task-based Programming for Scientific Computing: Runtime Supports

Thomas Sterling, Indiana University, U.S.

11:00-11:20 Abstraction Layer for Standardizing APIs of Task-based Engines

Rabab Omairy, Hatem Ltaief, and David E. Keyes, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

continued in next column

Monday, February 25

MS13

BE: Mentor Protégé

9:45 a.m.-11:25 a.m.

Room: 300D

Organizer: Mary Ann E. Leung
Sustainable Horizons Institute, U.S.

Monday, February 25

MS14

Emerging Trends for Structure Preserving Model Order Reduction - Part I of II

9:45 a.m.-11:25 a.m.

Room: 302B

For Part 2 see MS49

Model Order Reduction (MOR) is an indispensable tool for multi-query simulations and (faster than) real-time estimation and control, in particular for multiscale and multiphysics problems. Standard MOR approaches, such as, the Reduced Basis Method, the Proper Orthogonal Decomposition, etc. in conjunction with the (Discrete) Empirical Interpolation Method, fail to strike an acceptable balance between fidelity and complexity to obtain computationally compact representations of several phenomena of interest (for instance, multiple-transport phenomena). The limitations of standard MOR techniques instigate a drive to develop novel and automated techniques tailored for identification of underlying embedding (structure), and, preservation of key (physical) model properties (such as, positivity, evolving discontinuous front, etc.) and system theoretic properties (for instance, Port-Hamiltonian structure, propagation delays, stability, etc.) of the associated (parametric) reduced-order models, which are guaranteed to offer robust approximation with lowest possible reduced-order representation. In this minisymposium, we will discuss recent advances and perspectives in the field of structure and property preserving (discretization and) MOR frameworks to deal with (large-scale) linear and non-linear dynamical systems.

Organizer: Harshit Bansal
Eindhoven University of Technology, Netherlands

9:45-10:05 Perspective on Structure Preserving Model Order Reduction
Wil Schilders, Eindhoven University of Technology, Netherlands

10:10-10:30 Structure Preserving Reduced Order Models for Hamiltonian Problems with a Non-linear Poisson Structure
Cecilia Pagliantini and Jan S. Hesthaven, École Polytechnique Fédérale de Lausanne, Switzerland

10:35-10:55 Novel Structure Preserving Model Reduction Schemes
Igor Pontes Duff, Max Planck Institute for Dynamics of Complex Systems, Germany; Pawan Goyal and Peter Benner, Max Planck Institute for Dynamics of Complex Technical Systems, Germany

11:00-11:20 On Conservation of Dissipative Relations in Model Order and Complexity-reduced Systems
Bjoern Liljegren-Sailer and Nicole Marheineke, Trier University, Germany

continued in next column

Monday, February 25

MS15

Fast Methods for High-Frequency Wave Propagation - Part I of II

9:45 a.m.-11:25 a.m.

Room: 303A

For Part 2 see MS50

High-frequency wave equations, e.g., variable-coefficient Helmholtz equations, govern many large-scale electromagnetic, acoustic, and elastodynamic phenomena in inhomogeneous media. Their numerical solution remains a grand challenge due to difficulties in convergence, dispersion, and computational efficiency in the high-frequency regime. This minisymposium reports recent progress in low-complexity solvers including directional methods, sweeping preconditioners, method of polarized traces, as well as reliable asymptotic methods such as geometrical optics, Gaussian beam methods, and Huygens sweeping methods.

Organizer: Leonardo Zepeda-Nunez

Lawrence Berkeley National Laboratory, U.S.

9:45-10:05 Redemption of the High Frequencies: Scalable Algorithms for 3D Helmholtz

Laurent Demanet and Matthias Taus, Massachusetts Institute of Technology, U.S.; Leonardo Zepeda-Núñez, Lawrence Berkeley National Laboratory, U.S.; Russell Hewett, Total E&P, U.S.; Adrien Scheuer, Université Catholique de Louvain, Belgium

10:10-10:30 Hadamard-Babich Ansatz for Point-source Maxwell's Equations in Inhomogeneous Media

Jianliang Qian, Michigan State University, U.S.

10:35-10:55 Parallel-in-time Multiscale Algorithm for Wave Equations

Yen-Hsi Richard Tsai, University of Texas at Austin, U.S.

11:00-11:20 Fast Huygens Sweeping (FHS) Methods for Highly Oscillatory Phenomena

Shingyu Leung, Hong Kong University of Science and Technology, Hong Kong

Monday, February 25

MS16

Tutorials for Students: Accessible Introductions to Active Research Areas

9:45 a.m.-11:25 a.m.

Room: 303B

For Part 2 see MS51

This set of tutorials provides accessible introductions for students to areas of active research in computational science and engineering. It features four 45 minute tutorials which represent various areas of active research that will recur in talks throughout the conference. It serves to give students from various areas of computational science and engineering the necessary background to benefit from other talks throughout the conference, and future conferences as these active fields continue to grow. In particular, these four tutorials will touch on aspects of data driven modeling, integrating data in computational models, and algorithmic methodologies.

Organizer: Joseph L. Hart

North Carolina State University, U.S.

Organizer: Kathleen Kavanagh

Clarkson University, U.S.

9:45-10:30 Inverse Problems: Integrating Data with Models under Uncertainty

Noemi Petra, University of California, Merced, U.S.

10:35-11:20 Global Sensitivity Measures: Theory to Practice

Manav Vohra, Vanderbilt University, U.S.

Monday, February 25

MS17

Advances in Phase Retrieval: Theory and Applications - Part I of II

9:45 a.m.-11:25 a.m.

Room: 201B

For Part 2 see MS52

Phase retrieval is the process of recovering the phase of a complex signal given only its Fourier magnitude. This problem is known to be non-convex and plays a vital role in various scientific applications. Improved phase recovery methods and its applications are therefore an active area of research. This minisymposium reports on recent advances in phase retrieval ranging from algorithm advancement, experimental design, to application in high-energy x-ray science.

Organizer: Selin Aslan

Argonne National Laboratory, U.S.

Organizer: Zichao Di

Argonne National Laboratory, U.S.

9:45-10:05 Joint Ptycho-tomography Reconstruction Through Alternating Direction Method of Multipliers

Selin Aslan, Doga Gürsoy, and Sven Leyffer, Argonne National Laboratory, U.S.

10:10-10:30 New Methods for Partially Coherent Ptychography

Oliver Cossairt and Yudong Yao, Northwestern University, U.S.

10:35-10:55 Hyperspectral Ptychography and Beyond

Huibin Chang, Tianjin Normal University; Pablo Enfedaque, Hari krishnan, and *Stefano Marchesini*, Lawrence Berkeley National Laboratory, U.S.

11:00-11:20 Relaxation Algorithms for Nonsmooth Phase Retrieval

Aleksandr Aravkin and Peng Zheng, University of Washington, U.S.

Monday, February 25

MS18

Computational and Numerical Methods in Electronics - Part I of II

9:45 a.m.-11:25 a.m.

Room: 201C

For Part 2 see MS53

Research in electronics faces two challenges which are unfortunately often considered separate from each other: First, the rapid technological progress requires short simulation times for electronic devices with increasing geometric complexity. Such short simulation times can only be achieved through the development of new parallel algorithms since the serial performance of modern computing hardware stagnates. Second, sound numerical methods are required to accurately describe physical processes such as electronic transport at the discrete level. For example, important physical properties such as positivity and charge conservation need to be reflected in the numerics. This minisymposium - with the aim to bridge the gap between both challenges - has been initially founded with the previous SIAM Conference on Computational Science and Engineering in 2017. The continuation in 2019 will allow the minisymposium to establish itself as a series which will in turn enable to provide a continued presence as an interdisciplinary platform for physicists, numerical analysts, and computational scientists within the area of electronics. There is currently no other venue providing such a focused forum further underlining the importance of this minisymposium series.

Organizer: Karl Rupp
Technische Universität Wien, Austria

Organizer: Daniel Brinkman
San Jose State University, U.S.

Organizer: Patricio Farrell
Hamburg University of Technology, Germany

Organizer: Nella Rotundo
Weierstrass Institute, Germany

Organizer: Josef Weinbub
Technische Universität Wien, Austria

9:45-10:05 Conservation Laws in Diffusive Chemical Reaction Networks for Semiconductor Defects

Daniel Brinkman, San Jose State University, U.S.; Christian Ringhofer, Arizona State University, U.S.

10:10-10:30 High Order Sparse Grid Discontinuous Galerkin Methods for Nonlinear PDEs

Yingda Cheng, Michigan State University, U.S.

10:35-10:55 Highly Accurate Quadrature-based Scharfetter-Gummel Schemes for Charge Transport in Degenerate Semiconductors

Patricio Farrell, Hamburg University of Technology, Germany

11:00-11:20 Modeling Negative Capacitance and Inductive Loop in Perovskite Solar Cells

Evelyne Knapp, ZHAW Zurich University of Applied Sciences, Switzerland; Martin Neukom, Fluxim AG, Switzerland; Beat Ruhstaller, ZHAW Zurich University of Applied Sciences, Switzerland

continued in next column

Monday, February 25

MS19

Fast Solvers for Inverse Problems with PDEs - Part I of II

9:45 a.m.-11:25 a.m.

Room: 202A

For Part 2 see MS54

We present recent advances in the development of fast solvers and effective numerical schemes for the solution of inverse problems with a particular focus on large-scale PDE-constrained optimization problems. Key algorithmic challenges include computational complexity, memory consumption, and a vast number unknowns as well as model and data uncertainties. We showcase state-of-the-art techniques in scientific computing to tackle these challenges.

Organizer: James L. Herring
University of Houston, U.S.

Organizer: Andreas Mang
University of Houston, U.S.

Organizer: George Biros
University of Texas at Austin, U.S.

9:45-10:05 Adaptive Product-convolution Approximation for Hessians, Interface Schur Complements, and other Locally Translation-invariant Operators

Nick Alger, University of Texas at Austin, U.S.; Vishwas Rao, Argonne National Laboratory, U.S.; Aaron Myers, Tan Bui-Thanh, and Omar Ghattas, University of Texas at Austin, U.S.

10:10-10:30 Automatic Determination of Regularization Parameters in Krylov Methods for Large Scale Inverse Problems

Wim I. Vanroose and Nick Schenkels, University of Antwerp, Belgium

continued on next page

Monday, February 25

MS19

Fast Solvers for Inverse Problems with PDEs - Part I of II

continued

10:35-10:55 Accelerating Reconstruction Algorithms for Fluorescence Optical Tomography

Matthias Schlottbom, University of Twente, Netherlands; Herbert Egger, Technische Universität Darmstadt, Germany

11:00-11:20 Reduced Order Modeling for Time-dependent Optimal Control Problems with Variable Initial Values

Matthias Heinkenschloss, Rice University, U.S.

Monday, February 25

MS20

Moment Methods in Kinetic Theory - Part I of II

9:45 a.m.-11:25 a.m.

Room: 202B

For Part 2 see MS55

Kinetic theory enables modeling of numerous physical processes such as radiative transfer or rarefied gas flows in the transition or kinetic regime for moderate to large Knudsen numbers. Here the Knudsen number relates the mean free path of gas particles to the reference length of the system. Instead of tracking individual particles, a velocity distribution function is used to model the dynamics of the ensemble. Solutions of the underlying Boltzmann equation or variants thereof are often computed using moment methods. To that extend, a set of moments of the distribution function, including e.g. physical quantities like density, momentum, energy, heat flux, stress tensor and other higher moments, is chosen to describe the flow properties. The resulting moment equations then describe the evolution of these moments. In this minisymposium we want to show recent advances and applications of moment methods that highlight the applicability and accuracy of moment methods in comparison to standard fluid dynamics approaches and particle-based models. This includes the ongoing search for efficient yet accurate closures for moment equations as well as numerical methods ensuring asymptotic stability. On the other hand, recent developments of hyperbolic moment models shall be discussed together with first results for adaptive moment models and predictive hierarchical simulations. At the end, we want to exemplify various applications of moment methods, e.g. for polydisperse sprays.

Organizer: Julian Koellermeier
Freie Universität Berlin, Germany

Organizer: Yuwei Fan
Stanford University, U.S.

9:45-10:05 Filtered Hyperbolic Moment Equations for Improved Convergence

Julian Koellermeier, Freie Universität Berlin, Germany

10:10-10:30 Numerical Study of the Maximum Entropy Moment System

Yuwei Fan, Stanford University, U.S.

10:35-10:55 Residual-based Model-adaptivity with Moment Equations

Manuel Torrilhon, RWTH Aachen University, Germany

11:00-11:20 Hermite Spectral Method for Boltzmann Equation

Zhenning Cai, National University of Singapore, Singapore; Zhicheng Hu, Nanjing University of Aeronautics and Astronautics, China; Yanli Wang, Peking University, China

continued in next column

Monday, February 25

MS21

Recent Advances in High Order Lagrangian/ALE Methods - Part I of II

9:45 a.m.-11:25 a.m.

Room: 202C

For Part 2 see MS56

The numerical approximation of the Euler equations of gas dynamics in a moving frame is a common approach for solving many multiphysics problems involving e.g. large deformations, strong shocks and interactions of multiple materials. In Lagrangian methods, the mesh is moving with the fluid velocity, therefore they are well-suited for accurate resolution of material interfaces. On the other hand, multidimensional Lagrangian meshes tend to tangle so that the mesh elements become invalid, and in general cannot represent large deformation. This difficulty can be resolved in ALE (Arbitrary-Eulerian-Lagrangian) methods, which assume that the mesh moves independently on the flow and therefore offer additional flexibility and accuracy. ALE techniques can also be directly implemented in numerical schemes, for example to take into account moving objects, in fluid structure interaction, etc. In both cases, the question of the mesh quality is a central question. The aim of this minisymposium is to bring together the researchers working in the fields of Lagrangian hydrodynamics and ALE methods to discuss the state-of-the-art of single- and multi-material hydrodynamic simulations. We are particularly interested in recent advances in high order methods, such as high order finite volume (WENO, ADER), discontinuous Galerkin, high order finite elements, residual distribution methods, as applied to Lagrangian or ALE description of the flow.

Organizer: Svetlana Tokareva

Los Alamos National Laboratory, U.S.

Organizer: Remi Abgrall

Universität Zürich, Switzerland

9:45-10:05 A High Order Multidimensional Residual Distribution Scheme for Lagrangian Hydrodynamics

Svetlana Tokareva, Los Alamos National Laboratory, U.S.; Remi Abgrall, Universität Zürich, Switzerland; Konstantin Lipnikov and Nathaniel Morgan, Los Alamos National Laboratory, U.S.

10:10-10:30 A Multi-dimensional Discontinuous Galerkin Method for Solid Dynamics

Donald E. Burton, Nathaniel Morgan, Xiaodong Liu, Michael Berry, and Evan Lieberman, Los Alamos National Laboratory, U.S.

10:35-10:55 An Interpolation-free Adaptive ALE Approach with Multi-step Time Schemes

Barbara Re, University of Zurich, Switzerland; Alberto Guardone, Politecnico di Milano, Italy; Cecile Dobrzynski, Inria Bordeaux Sud-Ouest, France

11:00-11:20 Matrix-free Advection-based Remap Algorithms for Lagrangian/ALE Methods

Hennes Hajduk, Technische Universität Dortmund, Germany

continued in next column

Monday, February 25

MS22

Use of C++ in Computational Science Libraries and Applications - Part I of II

9:45 a.m.-11:25 a.m.

Room: 203

For Part 2 see MS57

The number of C++ libraries and applications in computational science have grown over the past years. We will be presenting some of them and the advancements in the C++ standard that have contributed to this growth. The talks in this minisymposium will focus on the relevant aspects of the C++ standard and accompanying libraries. Additionally, practical aspects and implementation methods will be shown to deal with a wide range of computational science codes on many modern platforms in high performance manner.

Organizer: Piotr Luszczek

University of Tennessee, U.S.

Organizer: Mark Hoemmen

Sandia National Laboratories, U.S.

9:45-10:05 Interfacing Dense Linear Algebra Libraries in C++

Piotr Luszczek, University of Tennessee, Knoxville, U.S.

10:10-10:30 Kokkos Libraries and Applications

David S. Hollman, Sandia National Laboratories, U.S.

10:35-10:55 Benchmarking Modern C++ Abstraction Penalty

Andrew Lumsdaine, University of Washington, U.S.

11:00-11:20 Modern C++ in Computation Science

Daniel Sunderland, Sandia National Laboratories, U.S.

Monday, February 25

MS23

Finite Element Methods on Unfitted Meshes - Part I of II

9:45 a.m.-11:25 a.m.

Room: 205

For Part 2 see MS58

Finite element method (FEM) solving partial differential equations has been largely focusing on the equations and overlooking the coupling of the domain or interface with complex or evolving geometry. For problems with difficult geometries, classical FEMs that employ fitted meshes often require highly non-trivial and costly steps to generate qualified meshes, especially in the three dimensions. It is then extremely advantageous to employ meshes that do not necessarily fit the physical geometry. This minisymposium focuses on most recent techniques to design stable and robust FEMs that employ the unfitted meshes.

Organizer: Cuiyu He

University College London, United Kingdom

9:45-10:05 A Posteriori Error Estimation for Cut Finite Element Method

Cuiyu He, University College London, United Kingdom

10:10-10:30 A Finite Element Method For PDEs in Time-dependent Domains

Maxim A. Olshanskii, University of Houston, U.S.

10:35-10:55 An Interface-unfitted Finite Element Method for Elliptic Interface Optimal Control Problems

Xiaoping Xie, Sichuan University, China

11:00-11:20 Hybrid High-order Methods on Unfitted Meshes

Guillaume Delay, ENPC and INRIA, France

Monday, February 25

MS24

Applications of the AMReX Block Structured Adaptive Mesh Refinement Framework - Part I of II

9:45 a.m.-11:25 a.m.

Room: 206A

For Part 2 see MS59

Multiphysics PDE simulation at scale can be a challenging task given rapidly evolving supercomputing architectures. As part of the DOE Exascale Computing Project (ECP), researchers at LBNL, NRL, and ANL have publicly released the AMReX software framework for massively parallel block-structured adaptive mesh refinement (AMR) applications. The core libraries include linear solvers, particle support, embedded boundary geometry representation, profiling tools, and hybrid parallelism. AMReX is the basis for many mature scientific applications, including several supported by ECP and SciDAC (e.g., combustion, carbon capture/storage, accelerator physics, and astrophysics, to name a few). In this minisymposium, 8 computational scientists from a wide range of fields will showcase their latest code innovations and scientific achievements using the AMReX software framework. Additionally, each will discuss many aspects of their problem, as listed below. Who should attend this minisymposium? -Researchers interested in multiphysics PDE simulation techniques at scale. This includes all aspects of the problem from model equation development, numerical discretizations, implementation in HPC frameworks, utilization of supercomputing resources, and analysis of data. -Researchers in fields represented by the speakers, including combustion, materials science, multiphase flow, astrophysics, stochastic PDEs modeling fluctuating hydrodynamics, and accelerator physics.

Organizer: Andy J. Nonaka

Lawrence Berkeley National Laboratory, U.S.

9:45-10:05 Overview of the AMReX Framework - Capabilities and Scientific Applications

Andy J. Nonaka, Lawrence Berkeley National Laboratory, U.S.

10:10-10:30 Modeling of Interface-driven Microstructure Evolution in Metals using the Multiphase Field Method

Brandon S. Runnels, University of Colorado, Colorado Springs, U.S.

10:35-10:55 MFIX-Exa: Using AMReX to Modernize a Legacy Code for Fossil Energy Applications

Jordan Musser, National Energy Technology Laboratory, U.S.; Madhava Syamlal, Fluent, Inc., U.S.; Ann S. Almgren and John B. Bell, Lawrence Berkeley National Laboratory, U.S.

11:00-11:20 Hierarchical Fidelity to Meet the Evolving Needs of Basic and Applied Combustion Simulations

Ray W. Grout, National Renewable Energy Laboratory, U.S.

continued in next column

Monday, February 25

MS25

Discovering and Exploiting Low-dimensional Structures in Computational Models - Part I of II

9:45 a.m.-11:25 a.m.

Room: 206B

For Part 2 see MS60

High-dimensional spaces of input parameters and observable output data are commonly encountered in the modeling of practical problems. Unfortunately, many of the most reliable and well understood methods we rely upon to numerically solve models, perform optimization, and quantify uncertainties suffer from some form of the "curse of dimensionality." Recently, there is great interest in developing methods that identify and exploit lower-dimensional structures in high-dimensional spaces. This session focuses on recent advances in methodologies and approaches for discovering and exploiting low-dimensional structures in the input and output spaces of models.

Organizer: Steven A. Mattis
Technische Universität München, Germany

Organizer: Troy Butler
University of Colorado, Denver, U.S.

Organizer: Steven A. Mattis
Technische Universität München, Germany

9:45-10:05 Accelerating Prediction under Uncertainty with Dimension Reduction

Steven A. Mattis and Barbara Wohlmuth, Technische Universität München, Germany

10:10-10:30 Data-informed Subspace Identification using a Consistent (Data-oriented) Bayesian Method

Tan Bui-Thanh, University of Texas at Austin, U.S.

10:35-10:55 Scalable Approximations for the Consistent Bayes Method

Brad Marvin, University of Texas at Austin, U.S.; Timothy Wildey, Sandia National Laboratories, U.S.; Tan Bui-Thanh, University of Texas at Austin, U.S.

11:00-11:20 Dimensionality Reduction for Subsurface Flow Models

Daniel O'Malley, Los Alamos National Laboratory, U.S.

Monday, February 25

MS26

Parallel-in-Time Integration Techniques - Part I of II

9:45 a.m.-11:25 a.m.

Room: 206C

For Part 2 see MS61

For the efficient use of massively parallel supercomputers for applications in science and engineering new parallel algorithms are required. The idea of parallelizing in the time dimension and, thus, adding an additional layer of parallelism to a numerical algorithm has become an increasing popular way to tackle evolutionary problems. This minisymposium features talks on both theoretical and computational aspects of parallel-in-time integration. Research topics ranging from algorithm development and mathematical analysis to software implementation and exploitation of parallel-in-time integration techniques in realistic applications are discussed.

Organizer: Stephanie Friedhoff
University of Wuppertal, Germany

Organizer: Debasmita Samaddar
UK Atomic Energy Authority, United Kingdom

9:45-10:05 Convergence Analysis for Parallel-in-Time Solution of Hyperbolic Systems

Stephanie Friedhoff, University of Wuppertal, Germany; Hans De Sterck, Monash University, Australia; Alexander Howse, University of Waterloo, Canada; Scott Maclachlan, Memorial University, Newfoundland, Canada

10:10-10:30 A Relaxation Approach to Parallel-in-Time

Benjamin W. Ong, Michigan Technological University, U.S.; Felix Kwok, Hong Kong Baptist University, Hong Kong

10:35-10:55 Parallel Multigrid in Time for Hyperbolic Problems

Robert D. Falgout, Lawrence Livermore National Laboratory, U.S.

11:00-11:20 Space-time Reduction Methods for Linear PDEs

Oliver A. Krzysik and Hans De Sterck, Monash University, Australia; Scott Maclachlan, Memorial University, Newfoundland, Canada; Alexander Howse, University of Waterloo, Canada; Stephanie Friedhoff, University of Wuppertal, Germany

continued in next column

Monday, February 25

MS27

Modeling, Analyzing and Computing of Fractional Partial Differential Equations - Part I of II

9:45 a.m.-11:25 a.m.

Room: 206D

For Part 2 see MS62

The minisymposium covers modeling analyzing, and computing of fractional partial differential equations (PDEs) in physics, mechanics, engineering, and biology etc. The interdisciplinary talks stimulate research at the crossroads of computational science and engineering, and applied mathematics. The minisymposium reports novel numerical techniques for fractional PDEs and their high-performance computer implementations. Works that investigate the analytical properties of the solutions of systems consisting of PDEs with derivatives of fractional orders are also of interest to the symposium. Physical modeling of real-life phenomena employing fractional derivatives is also a topic covered by this symposium. Moreover, the interplay between modeling, analysis and simulation of fractional systems is a topic of high interest to this event.

Organizer: Abdul M. Khaliq

Middle Tennessee State University, U.S.

Organizer: Khaled Furati

King Fahd University of Petroleum and Minerals, Saudi Arabia

9:45-10:05 Source Determination for a Two-parameter Anomalous Diffusion with Local Time Datum

Khaled Furati, King Fahd University of Petroleum and Minerals, Saudi Arabia

10:10-10:30 Structure-preserving and Numerically Efficient Methods for Fractional Wave Equations

Jorge E. Macias-Diaz, Universidad Autonoma de Aguascalientes, Mexico

10:35-10:55 Sparse Approximation of Non Local Operator Equations via Compressed Sensing

Viktor Reshniak, Oak Ridge National Laboratory, U.S.

11:00-11:20 Parallel Algorithms for Nonlinear Time-space Fractional PDEs

Toheeb A. Biala, Middle Tennessee State University, U.S.

Monday, February 25

MS28

FASTMath Tools and Technologies - Part I of II

9:45 a.m.-11:25 a.m.

Room: 207

For Part 2 see MS63

The FASTMath (Frameworks, Algorithms and Scalable Technologies for Mathematics) Institute is a R&D project funded by the SciDAC Program at the U.S. Department of Energy (DOE). The goal of FASTMath is to develop and deploy scalable mathematical algorithms and software tools for reliable simulation of complex physical phenomena and collaborating with DOE domain scientists to ensure the usefulness and applicability of the work in the project. The focus of FASTMath is strongly driven by the requirements of DOE application scientists who require fast, accurate, and robust forward simulation along with the ability to efficiently perform ensembles of simulations in optimization or uncertainty quantification studies. This minisymposium will present work from the eight core areas of FASTMath: structured meshes, unstructured meshes, time integrators, linear solvers, eigensolvers, numerical optimization, data analytics, and uncertain quantification.

Organizer: Esmond G. Ng

Lawrence Berkeley National Laboratory, U.S.

9:45-10:05 AMReX and Applications

Ann S. Almgren, Lawrence Berkeley National Laboratory, U.S.

10:10-10:30 Unstructured Mesh Support for Particle-in-Cell Simulations

Mark S. Shephard, Rensselaer Polytechnic Institute, U.S.; Eisung Yoon, Ulsan National Institute of Science and Technology, South Korea; Agnieszka Truszkowska, Gopan Perumpilly, E. Seegyoung Seol, Onkar Sahni, Cameron Smith, Gerrett Diamond, and William R. Tobin, Rensselaer Polytechnic Institute, U.S.

continued in next column

continued on next page

10:35-10:55 Communication and Synchronization Reduction in Sparse Factorization and Triangular Solution

Xiaoye S. Li, Lawrence Berkeley National Laboratory, U.S.

11:00-11:20 Solving Nonlinear Eigenvalue Problems with Padé Approximate Linearization

Osni A. Marques, Lawrence Berkeley National Laboratory, U.S.; Zhaojun Bai, University of California, Davis, U.S.

Monday, February 25

MS29

Computational Methods for Kinetic Models of Plasma - Part I of II

9:45 a.m.-11:25 a.m.

Room: 401A

For Part 2 see MS64

Statistical mechanics provides a mathematical modeling framework in which a collection of "particles" (e.g., electrons, nucleons, atoms, or molecules) is represented by a probability density function (PDF). This representation allows one to model the influence of microscopic dynamics on larger scales, without explicitly computing the detailed dynamics on the microscopic scale. Models in this framework are referred to as "kinetic models", and arise in many application areas, including in rarefied gas dynamics, nuclear reactor modeling, and plasma physics. Transport phenomena refers to the evolution and redistribution of macroscopic quantities such as mass, momentum, energy, and heat flux. This minisymposium addresses recent advances in computational methods for the simulation of statistical mechanical models and the resulting transport phenomena. Our goal is to showcase a diverse array of application areas, models, and numerical techniques.

Organizer: James A. Rossmanith

Iowa State University, U.S.

9:45-10:05 Regionally Implicit Discontinuous Galerkin Methods for Solving the Relativistic Vlasov-Maxwell System

Pierson Guthrey, Michigan State University, U.S.; James A. Rossmanith, Iowa State University, U.S.

10:10-10:30 Solving the Continuum Kinetic Equations for Plasmas using the Discontinuous Galerkin Method

Bhuvana Srinivasan and Petr Cagas, Virginia Tech, U.S.; Ammar Hakim, Princeton Plasma Physics Laboratory, U.S.

10:35-10:55 Comparison of Boltzmann and Fokker-Planck Collision Operators for Moderately Coupled Plasmas

Jeffrey Haack, Los Alamos National Laboratory, U.S.

11:00-11:20 Verification Test Cases of Grid-based Direct Kinetic Modeling Framework for Plasma Flows

Ken Hara, Texas A&M University, U.S.; Kyle Hanquist, University of Michigan, U.S.

continued in next column

Monday, February 25

MS30

Graph and Combinatorial Algorithms for Enabling Exascale Applications - Part I of II

9:45 a.m.-11:25 a.m.

Room: 401B

For Part 2 see MS65

Combinatorial algorithms in general and graph algorithms in particular play a critical enabling role in numerous scientific applications. The low computation to communication ratios and irregular memory accesses of these algorithms makes them some of the hardest algorithmic kernels to implement on parallel systems. The talks in this minisymposium will consider: (i) exascale applications that drive the selection of combinatorial kernels and integration of software tools; (ii) combinatorial (graph) kernels such as graph traversals, graph matching, graph coloring, graph clustering, and graph partitioning, that play a crucial enabling role in the chosen application areas; and (iii) software frameworks for efficient implementation of the algorithms on hierarchical distributed-memory architectures that are representative of potential exascale platforms.

Organizer: Alex Pothen

Purdue University, U.S.

Organizer: Mahantesh Halappanavar

Pacific Northwest National Laboratory, U.S.

9:45-10:05 A Scheduling Problem Motivated by Cybersecurity and Adaptive Machine Learning

Ojas Parekh and *Cynthia Phillips*, Sandia National Laboratories, U.S.; Vladlena Powers, Nouri Sakr, and Cliff Stein, Columbia University, U.S.

10:10-10:30 A Scalable Algorithm for Data Anonymization

Arif Khan, Pacific Northwest National Laboratory, U.S.

10:35-10:55 Parallel Algorithms for Degree Constrained Subgraphs with Applications

S M Ferdous and Alex Pothen, Purdue University, U.S.

11:00-11:20 Construction of Massive Protein Similarity Networks using Distributed Suffix Arrays

Saliya P. Ekanayake, Lawrence Berkeley National Laboratory, U.S.

Monday, February 25

MS31

Numerical Methods for Earth System Modeling - Part I of II

9:45 a.m.-11:25 a.m.

Room: 401C

For Part 2 see MS66

Earth system models are complex multiscale, multiphysics systems that require advanced discretization schemes, temporal integration methods, and algebraic solvers to obtain accurate and stable solutions on today's high-performance computing systems. This minisymposium brings together researchers to present state of the art methods in these areas applied to Earth system modeling. It also focuses on real-world challenges and opportunities encountered while implementing these schemes on current and upcoming computing systems.

Organizer: Christopher J. Vogl

Lawrence Livermore National Laboratory, U.S.

Organizer: David J. Gardner

Lawrence Livermore National Laboratory, U.S.

9:45-10:05 High-order Finite-volume Weno Discretizations that Effectively use Accelerated HPC for Atmospheric Dynamics

Matthew R. Norman, Oak Ridge National Laboratory, U.S.

10:10-10:30 Hybrid-mixed Finite Element Solvers for Compressible Atmospheric Equations

Thomas H. Gibson, Imperial College London, United Kingdom

continued in next column

continued on next page

10:35-10:55 Evaluation of Additive Runge-Kutta Methods for the E3SM Non-hydrostatic Dycore

Christopher J. Vogl, Lawrence Livermore National Laboratory, U.S.; Daniel R. Reynolds, Southern Methodist University, U.S.; Paul Ullrich, University of California, Davis, U.S.; Andrew J. Steyer, Sandia National Laboratories, U.S.; Carol S. Woodward and David J. Gardner, Lawrence Livermore National Laboratory, U.S.

11:00-11:20 The E3SM Non-hydrostatic Dynamical Core at Cloud Resolving Resolution - DCMIP 2016 Super Cell Thunderstorm

Jorge E. Guerra, University of California, Davis, U.S.

Monday, February 25

MS32

Recent Advances in Error Estimation for Complex Problems - Part I of II

9:45 a.m.-11:25 a.m.

Room: 402A

For Part 2 see MS67

The accurate simulation of complex systems which incorporate multiple physical processes interacting across a range of spatial and temporal scales demands sophisticated multi-discretization numerical methods. These include operator decomposition and domain decomposition, as well as specialized time discretizations such as explicit, semi-implicit, operating splitting and parallel-in-time methods. Each of these discretization choices introduces a new source of error and the various sources of discretization error interact in non-trivial ways. Accurate error estimation is essential for building confidence in numerical solutions and is a key component of uncertainty quantification for multiphysics simulations. This minisymposium will feature recent advances in a posteriori error estimation for simulation and uncertainty quantification and the discretization choices necessary to meet constraints on accuracy.

Organizer: Jehanzeb H. Chaudhry

University of New Mexico, U.S.

Organizer: Simon Tavener

Colorado State University, U.S.

9:45-10:05 Error Control for Hybridized and Embedded Discontinuous Galerkin Methods

Krzysztof Fidkowski and Guodong Chen, University of Michigan, U.S.

10:10-10:30 Variational Formulation and Error Estimation for Explicit Time Integrators

Judit Muñoz-Matute and David Pardo, University of the Basque Country, Spain; Victor M. Calo, Curtin University of Technology, Perth, Australia; Elisabete Alberdi, University of the Basque Country, Spain

10:35-10:55 Error Estimation for Multiplicative Schwarz Domain Decomposition Method

Jehanzeb H. Chaudhry, University of New Mexico, U.S.; Simon Tavener and Don Estep, Colorado State University, U.S.

11:00-11:20 An A Posteriori Error Estimation for Numerical Solution of Richards Equation

Victor E. Ginting, University of Wyoming, U.S.

continued in next column

Monday, February 25

MS33

Recent Progresses in Data Analysis - Part I of II

9:45 a.m.-11:25 a.m.

Room: 402B

For Part 2 see MS68

Data analysis plays a critical role in various applications and its techniques are being actively developed today. The proposed minisymposium will bring together researchers in data science, particularly in the areas of topological data analysis and optimization. In this minisymposium we will try to present both the theoretical and practical developments recently made in those areas and provide an opportunity for researchers to have conversation on bridging those areas.

Organizer: Jae-Hun Jung

State University of New York at Buffalo, U.S.

Organizer: Guohui Song

Clarkson University, U.S.

9:45-10:05 Topological Data Analysis of Vascular Disease

Jae-Hun Jung, Megan Johnson, and John Nicponski, State University of New York at Buffalo, U.S.

10:10-10:30 Stitch-fix for Mapper

Bala Krishnamoorthy and Nathaniel Saul, Washington State University, U.S.; Bei Wang, University of Utah, U.S.

10:35-10:55 Community Tree Persistence: A Topological Method for Community Structure Analysis in Dynamic Networks

Ashis G. Banerjee, Wei Guo, Ruqian Chen, and Yen-Chi Chen, University of Washington, U.S.

11:00-11:20 R Package TDA for Statistical Inference on Topological Data Analysis

Jisu Kim, Carnegie Mellon University, U.S.

Monday, February 25

MS34

Methods for Large-scale Risk-averse Optimization - Part I of II

9:45 a.m.-11:25 a.m.

Room: 402C

For Part 2 see MS69

Many science and engineering applications require, e.g., the control or design of a physical system modeled by partial differential equations (PDEs) with uncertain inputs. For such problems, determining optimal solutions that are resilient to uncertainty is critical. This session focuses on stochastic optimization problems motivated by a range of applications including power networks, financial markets, petroleum engineering and structural design. In particular, we target mathematical approaches for risk aversion and probabilistic computations with an emphasis on efficient numerical methods.

Organizer: Drew P. Kouri

Sandia National Laboratories, U.S.

Organizer: Harbir Antil

George Mason University, U.S.

9:45-10:05 How to Supplement Risk Regulations to Avoid Industrial Catastrophes

Stan Uryasev and Giorgi Pertaia, University of Florida, U.S.

10:10-10:30 Chance Constrained Optimization in Banach Spaces with Application to PDE Systems

Abebe Geletu, Armin Hoffmann, and Pu Li, Technische Universität Ilmenau, Germany

10:35-10:55 A Primal-dual Algorithm for Risk Minimization with Application to PDE-constrained Optimization

Drew P. Kouri, Sandia National Laboratories, U.S.

11:00-11:20 Scalable Quasi-Newton Methods for the Optimization of Engineering Systems under Uncertainty using High-performance Computing

Cosmin G. Petra, Lawrence Livermore National Laboratory, U.S.

Monday, February 25

CP1

Modeling and Uncertainty Quantification I

9:45 a.m.-11:25 a.m.

Room: 301

Chair: Carley R. Walker, University of Southern Mississippi, U.S.

9:45-10:00 Parameter Identification for a Viscoplastic Model with Damage and Effect of Conditions on Results using Bayesian Approaches

Ehsan Adeli, Technische Universität Braunschweig, Germany

10:05-10:20 Surrogate Model Design for Non-smooth and Computationally Expensive Radiation Detection Models

Jared A. Cook, North Carolina State University, U.S.

10:25-10:40 Sampling Input Parameters of Stochastic Models using Consistent Pull-back Measures

Tian Yu Yen, University of Colorado, Denver, U.S.

10:45-11:00 Parameter Inference and Uncertainty Quantification in Simulating Blood Flow in Coronary Arteries

Minglang Yin, Alireza Yazdani, and George Karniadakis, Brown University, U.S.

11:05-11:20 Removal of Multiplicative Poisson Noise using Variance Based Joint Sparsity Recovery

Carley R. Walker, University of Southern Mississippi, U.S.

Monday, February 25

CP2

Computational Fluid Mechanics I

9:45 a.m.-11:25 a.m.

Room: 201A

Chair: Fady M. Najjar, Lawrence Livermore National Laboratory, U.S.

9:45-10:00 Equilibrium Statistics and Entropy Computation for Vortex Filaments on a Cubic Lattice: A Modified Pivot Algorithm

Pavel Belik, Aleksandr Lukanen, and Robert Laskowski, Augsburg University; Douglas Dokken, Kurt Scholz, and Mikhail M. Shvartsman, University of St. Thomas, U.S.

10:05-10:20 Reference Mapping: Application in Eulerian Hydrocodes

Serge B. Ndanou, Gary Dilts, and Thomas Masser, Los Alamos National Laboratory, U.S.

10:25-10:40 Future Proofing a Massive Parallel Framework for Real-time Flood Prediction

Christoph M. Ertl, Bobby Minola Ginting, and Ralf-Peter Mundani, Technische Universität München, Germany

10:45-11:00 Visual Exploration of Flood Simulations in Cave-based Environments

Ralf-Peter Mundani, Mario Silaci, and Christoph M. Ertl, Technische Universität München, Germany

11:05-11:20 Exploring the Predictibility of Random Forests for the Sedov-Von Neumann-Taylor Blast Wave Solution

Mokbel Karam, University of Utah, U.S.; *Fady M. Najjar* and Ming Jiang, Lawrence Livermore National Laboratory, U.S.; James C. Sutherland and Tony Saad, University of Utah, U.S.

Monday, February 25

Fundamental Algorithms Book Series Editorial Board Meeting

11:30 a.m.-1:00 p.m.

Room: Davenport Hotel -- Meeting Room 11

Lunch Break

11:30 a.m.-1:00 p.m.

Attendees on their own

Monday, February 25

PD1

Future Directions of Research Funding Programs

11:30 a.m.-12:30 p.m.

Room: 111A

Chair: Hans Johansen, Lawrence Berkeley National Laboratory, U.S.

Chair: Robert Scheichl, Universität Heidelberg, Germany

Chair: Rebecca Willett, University of Chicago, U.S.

Chair: Rio Yokota, Tokyo Institute of Technology, Japan

What opportunities are available for obtaining research funding? What makes a research proposal stand out? How can you build a research program that is attractive to the funding agencies? How can you conduct your research to make the biggest impact and increase your chances of future funding? We address all these questions and more as a part of this panel discussion.

Vipin Chaudhary
National Science Foundation, U.S.

Jan Hesthaven
École Polytechnique Fédérale de Lausanne, Switzerland

David Keyes
KAUST, Saudia Arabia and Columbia University, U.S.

Steven Lee
Lawrence Livermore National Laboratory, U.S.

Lee Seversky
Air Force Research Laboratory, U.S.

Monday, February 25

PD2

Early Career Panel

11:30 a.m.-12:30 p.m.

Room: Conference Theater

Chair: Katherine J. Evans, Oak Ridge National Laboratory, U.S.

Chair: Stefan Wild, Argonne National Laboratory, U.S.

Chair: Alison Marsden, Stanford University, U.S.

The goal of this panel is to provide new scientists useful information and guidance about different career choices available to them for a career. All career paths, for example academia, national labs, and industry, have unique attributes and challenges. Several panelists with diverse experiences will be on hand to provide some background and real-world experiences as they have progressed in their multi-faceted careers.

Ilene Carpenter
Cray, Inc., U.S.

Jed Brown
University of Colorado, Boulder, U.S.

Aydin Buluc
Lawrence Berkeley National Laboratory, U.S.

Monday, February 25

IP2

Deep Learning for Inverse Problems - Some Recent Approaches

1:00 p.m.-1:45 p.m.

Room: Ballroom 100BC

Chair: Katherine J. Evans, Oak Ridge National Laboratory, U.S.

In this talk we discuss the idea of data-driven regularisers for inverse imaging problems. We are in particular interested in the combination of model-based and purely data-driven image processing approaches. In this context we will make a journey from "shallow" learning for computing optimal parameters for variational regularisation models by bilevel optimization to the investigation of different approaches that use deep neural networks for solving inverse imaging problems. Alongside all approaches that are being discussed, their numerical solution and available solution guarantees will be stated.

Carola-Bibiane Schönlieb
University of Cambridge, United Kingdom

Coffee Break

1:45 p.m.-2:15 p.m.

Room: Ballroom Foyer

Monday, February 25

MS35

Mathematical Advances in Deep Learning - Part II of II

2:15 p.m.-3:55 p.m.

Room: Ballroom 100BC

For Part 1 see MS1
Featured Minisymposium

Training neural networks with many hidden layers, aka Deep Learning, has become a booming area of machine learning. Despite enormous success in applications including computer vision and speech recognition, several critical challenges in designing, training, and understanding deep networks persist. This minisymposium presents recent advances toward rigorous mathematical underpinnings of deep learning and more efficient training. Speakers will report on recent success enabled by applying techniques from partial differential equations, optimization, optimal control, inverse problems, uncertainty quantification, to the deep learning problem.

Organizer: Lars Ruthotto
Emory University, U.S.

Organizer: Eldad Haber
University of British Columbia, Canada

2:15-2:35 A Mean-field Optimal Control Formulation of Deep Learning
Jiequn Han, Princeton University, U.S.

2:40-3:00 Structured Models for Convolutional Neural Networks
Xiaojin Tan and Eldad Haber, University of British Columbia, Canada

3:05-3:25 Extensions and Algorithms for Neural Net Training
Peng Zheng and Aleksandr Aravkin, University of Washington, U.S.

3:30-3:50 Deep Inversion - Latent Space Analysis of Autoencoders for Inverse Problems
Yoeri Boink, Stephan A. Van Gils, Srirang Manohar, and *Christoph Brune*, University of Twente, Netherlands

Monday, February 25

MS36

Scientific Software: Practices, Concerns, and Solution Strategies - Part II of II

2:15 p.m.-3:55 p.m.

Room: Conference Theater

For Part 1 see MS2

Software is an increasingly important component of scientific research. It would not be an exaggeration to say that almost all forms of scientific research today involve software, whether to compute, to manage, or to analyze. The kinds of software used and the associated issues, which range from technical to sociological, are as diverse as the scientific communities themselves. This minisymposium reflects this rich diversity, including presentations from software developers, users, and scientific project managers from various parts of the world and from various science communities. The common thread among these topics is that many overlapping subsets of scientific software issues are of interest to various participants. The minisymposium will, therefore, be a snapshot of the current state of practices, concerns and solution strategies in the world of scientific software.

Organizer: Anshu Dubey

Argonne National Laboratory, U.S.

2:15-2:35 Software Engineering Guidelines for Scientists - A Practical Handout for the Developing Researcher

Carina S. Haupt, German Aerospace Center (DLR), Germany

2:40-3:00 Symbiosis between Software and Scientific Career

Matthew J. Turk, University of Illinois, U.S.

3:05-3:25 Role of Requirements in Scientific Software

Jared O'Neal, Argonne National Laboratory, U.S.

3:30-3:50 Repository Analysis of Open-source and Scientific Software Development Projects

Boyana Norris and Kanika Sood, University of Oregon, U.S.; Anshu Dubey and Lois Curfman McInnes, Argonne National Laboratory, U.S.

Monday, February 25

MS37

AWM Workshop: Data Science and Mathematics - Part II of II

2:15 p.m.-3:55 p.m.

Room: 102A

For Part 1 see MS3

Session description to follow.

Organizer: Deanna Needell

University of California, Los Angeles, U.S.

Organizer: Giseon Heo

University of Alberta, Canada

2:15-2:35 Fast and Efficient Distributed Matrix-vector Multiplication using Rateless Fountain Codes

Gauri Joshi, Carnegie Mellon University, U.S.

2:40-3:00 On the Service Capacity Region of Content Access from Coded Storage

Sarah Anderson, University of St. Thomas, U.S.

3:05-3:25 Data Analysis and Automated Sleep Apnea Diagnosis

Kathryn Leonard, Occidental College, U.S.

3:30-3:50 Automated Obstructive Sleep Apnea Diagnosis through High-dimensional Multi-Source Data Analysis

Xu Wang, Wilfrid Laurier University, Canada

Monday, February 25

MS39

Recent Advances in Modeling and Numerical Analysis for Electronic Structure Calculations - Part II of II

2:15 p.m.-3:55 p.m.

Room: 102C

For Part 1 see MS5

The goal of this minisymposium is to bring experts working on new modeling approaches and algorithmic contributions for electronic structure problems in materials. For many cutting edge applications, materials have to been studied at the quantum level to predict their properties, such as transport (conduction), magnetic or optical response. Systems that are of interest to materials scientists are typically of increasing complexity and scales, leading to new challenges from a mathematical and computational perspective for the computation and characterization of the electronic structure of materials. One recent example among others is that of devices formed of multiple layers of recently discovered 2D materials such as graphene, boron nitride, etc. and which lack periodicity. The first talk of the session will give an introduction to electronic structure computations and the other invited speakers will present their latest developments in the field.

Organizer: Paul Cazeaux

University of Kansas, U.S.

2:15-2:35 Transport Coefficients of Thermally Disordered Materials

Emil Prodan, Yeshiva University, U.S.

2:40-3:00 Exploiting Sparsity in the Chebyshev Coefficients for Faster Conductivity Calculations

Simon Etter, University of Warwick, United Kingdom

continued on next page

Monday, February 25

MS39

Recent Advances in Modeling and Numerical Analysis for Electronic Structure Calculations - Part II of II

continued

3:05-3:25 Practical Computation of the Noncommutative Kubo Formula for the Conductivity in Incommensurate Multilayers

Paul Cazeaux, University of Kansas, U.S.

3:30-3:50 A Convolutional Neural Network to Determine Effective Transport Properties for Electrochemical Lithium-ion Battery Models

Ross M. Weber, Svyatoslav Korneev, Simona Onori, and Ilenia Battiato, Stanford University, U.S.

Monday, February 25

MS40

Roundoff Error in High-Performance Implementations of CG/ Lanczos-type Solvers - Part II of II

2:15 p.m.-3:55 p.m.

Room: 102D

For Part 1 see MS6

The conjugate gradient and Lanczos algorithms are widely used methods for solving Hermitian positive definite linear systems, computing eigenvalues of Hermitian matrices and, more generally, for computing the product of a function of a Hermitian matrix with a given vector. It is well-known that the methods may behave differently in finite precision arithmetic than they would in exact arithmetic, which means that whenever different implementations are considered -- such as implementations that make better use of parallelism -- one must be careful about the effects on actual performance and accuracy. This minisymposium brings together researchers to talk about different mathematically equivalent CG/Lanczos implementations and their behavior in finite precision arithmetic.

Organizer: Anne Greenbaum
University of Washington, U.S.

Organizer: Erin C. Carson
Charles University, Czech Republic

2:15-2:35 Numerical Analysis of the Maximal Attainable Accuracy in Communication-hiding Pipelined Conjugate Gradients

Siegfried Cools and Jeffrey Cornelis, University of Antwerp, Belgium; Emmanuel Agullo, Inria, France; Emrullah Fatih Yetkin, Kadir Has University, Turkey; Luc Giraud, Inria, France; Wim Vanroose, Antwerp University, Belgium

2:40-3:00 The s-Step Conjugate Gradient Method in Finite Precision

Erin C. Carson, Charles University, Czech Republic

3:05-3:25 On the Convergence Rate of Different Variants of the Conjugate Gradient Algorithm

Hexuan Liu, University of Washington, U.S.

3:30-3:50 Exploiting Multiprecision in Krylov Subspace Methods

David Titley-Peloquin, McGill University, Canada

continued in next column

Monday, February 25

MS41

Numerical Linear Algebra for Machine Learning - Part II of II

2:15 p.m.-3:55 p.m.

Room: 111A

For Part 1 see MS7

Machine learning algorithms such as deep neural networks with convolution and fully connected layers, recurrent neural networks, support vector machines, and kernel machines rely on linear algebra for matrix-vector and matrix-matrix products. Many other algorithms for "big data processing" rely on low-rank decomposition such as the thin SVD, rank-revealing QR factorizations, and randomized projection and sampling algorithms. In this MS, we will explore the connections between machine learning algorithms and numerical linear algebra, including techniques such as the fast multipole method, hierarchical matrices, structured matrices, tensor representations, tensor-trains, low-rank and sparse representations, and numerical algorithms for large-scale eigenvalue and singular value calculations.

Organizer: Leopold Cambier
Stanford University, U.S.

Organizer: Eric F. Darve
Stanford University, U.S.

2:15-2:35 Online Optimization for Continuous Hyperparameters with Applications to Traffic Flow Prediction
Hongyuan Zhan, Pennsylvania State University, U.S.

2:40-3:00 Randomized Algorithms to Compute Low-rank Approximation for Data Analysis and Machine Learning Applications
Ichitaro Yamazaki, University of Tennessee, Knoxville, U.S.; Yuechao Lu, Osaka University, Japan; Michael Mahoney, University of California, Berkeley, U.S.; Jack J. Dongarra, University of Tennessee and Oak Ridge National Laboratory, U.S.

3:05-3:25 Randomized Sparse PCA using the Variable Projection Method
N. Benjamin Erichson, University of California, Berkeley, U.S.

3:30-3:50 Scalable Hierarchical Matrix Methods in Machine Learning
George Biros, University of Texas at Austin, U.S.

continued in next column

Monday, February 25

MS42

Uncertainty Quantification and Data Assimilation - Part II of II

2:15 p.m.-3:55 p.m.

Room: 111B

For Part 1 see MS8
Featured Minisymposium

Data assimilation is the primary means of prediction and parameter estimation in weather forecasting, geophysics, fluid dynamics, and finance. Standard methods in filtering and Bayesian inference are particularly appealing considering the rapid advances in data acquisition in recent years. However, incorporating observation data with forward models introduces both model and measurement uncertainty into the prediction. As a result, these methods may require a large number of forward model evaluations to converge, which can be computationally prohibitive for high-dimensional problems. This minisymposium invites speakers working at the forefront of data assimilation and uncertainty quantification to discuss recent challenges, innovations, and strategies in these areas.

Organizer: Krithika Manohar
California Institute of Technology, U.S.

2:15-2:35 Interdisciplinary Data Assimilation and Bayesian Learning
Pierre F. Lermusiaux, Massachusetts Institute of Technology, U.S.

2:40-3:00 EnKF-Based Estimators for Aerodynamics
Tim Colonius and Andre da Silva, California Institute of Technology, U.S.; Darwin Darakananda and Jeff D. Eldredge, University of California, Los Angeles, U.S.

3:05-3:25 Ensemble Kalman Inversion: A Derivative-free Technique for Machine Learning Tasks
Nikola Kovachki and Andrew Stuart, California Institute of Technology, U.S.

3:30-3:50 Greedy Sensor Placement for Controlling High-dimensional Dynamics
Krithika Manohar, California Institute of Technology, U.S.; J. Nathan Kutz and Steven Brunton, University of Washington, U.S.

Monday, February 25

MS43

Avoiding Big-data: In-situ Visualisation and Analysis for Large-scale CFD Simulations - Part II of II

2:15 p.m.-3:55 p.m.

Room: 111C

For Part 1 see MS9

Modern-day CFD, especially time-resolved simulations of turbulent flow, can readily produce unmanageable amounts of data. Disk I/O and fundamental storage capacity preclude writing such data to magnetic disk for a posteriori analysis. Therefore, new technologies are required that can visualise and analyse these data on-the-fly 'in-situ' as they are generated. This minisymposia will explore the state-of-the-art in this area, including the development of frameworks for in-situ visualisation, data compression, feature identification, and feature tracking.

Organizer: Marius K. Koch
Imperial College London, United Kingdom

Organizer: Peter E. Vincent
Imperial College London, United Kingdom

Organizer: Paul Kelly
Imperial College London, United Kingdom

Organizer: Freddie Witherden
Imperial College London, United Kingdom

2:15-2:35 Online Compression of High-order CFD Solutions using Machine Learning

Freddie Witherden, Imperial College London, United Kingdom; Jeremy Morton, Stanford University, U.S.; Liqian Peng and Kevin T. Carlberg, Sandia National Laboratories, U.S.; Mykel Kochenderfer and Antony Jameson, Stanford University, U.S.

2:40-3:00 In-situ Data Visualization Approaches in the SU2 Analysis and Design Framework

Scott Imlay, Tecplot, Inc., U.S.; Juan J. Alonso, Stanford University, U.S.; Thomas D. Economon, Bosch Research And Technology Center, U.S.

3:05-3:25 Siac Filtering for Visualization and Feature Detection

Ashok Jallepalli and Robert M. Kirby, University of Utah, U.S.

3:30-3:50 Advanced in Situ Analysis of High Order Dg Solutions with Application to Scale-resolving Simulations of Complex Turbulent Flows

Michel Rasquin and Koen Hillewaert, CENAERO, Belgium

continued in next column

Monday, February 25

MS44

Model Reduction, Adaptivity, and High Dimensionality in Uncertainty Quantification - Part II of II

2:15 p.m.-3:55 p.m.

Room: 300A

For Part 1 see MS10

Uncertainty Quantification (UQ) is essential for credible simulation-aided prediction and design. However, many computational strategies for both forward and inverse UQ are intractable when applied to models with a moderate or large number of parameters. Exploiting structure of the model input-output map and the goals of the overall analysis are needed to improve tractability. This minisymposium will focus on goal-oriented techniques whose aim is to reduce the cost of uncertainty quantification while simultaneously maintaining or improving accuracy. Areas of interest include dimension reduction, sparsity, approximations with non-tensorial measures, adaptivity, multi-level methods, and multi-fidelity methods. This minisymposium will bring together leading researchers from the applied and computational mathematics and engineering community to discuss and collaborate on novel theoretical and computational advances, and to discuss future directions for research.

Organizer: John D. Jakeman
Sandia National Laboratories, U.S.

Organizer: Akil Narayan
University of Utah, U.S.

2:15-2:35 Adaptive Construction of Quadrature Rules for Bayesian Inference and Prediction

Laurent van den Bos and Benjamin Sanderse, Centrum voor Wiskunde en Informatica (CWI), Netherlands

continued on next page

2:40-3:00 Stability-preserving Model Order Reduction for Random Dynamical Systems
Roland Pulch, University of Greifswald, Germany

3:05-3:25 A Multi-Level Optimization Based Monte-Carlo Sampler for Large Scale Inverse Problems
Chuntao Chen, Monash University, Australia

3:30-3:50 Large-scale Uncertainty Propagation via Overlapping Domain Decomposition
Kevin T. Carlberg, Sandia National Laboratories, U.S.; Sofia Guzzetti, Emory University, U.S.; *Mohammad Khalil* and Khachik Sargsyan, Sandia National Laboratories, U.S.

Monday, February 25

MS45

Data Enabled Modeling and Discovery in Science and Engineering - Part II of II

2:15 p.m.-3:55 p.m.

Room: 300B

For Part 1 see MS11

There is a growing amount of data generated from instruments and computational simulations associated with the scientific enterprise. The volume of data has driven an interest in scientific machine learning as a guided way to determine relationships and incorporate data into models. This presents a number of challenges unique to scientific machine learning of relevance to the applied mathematics community. For instance, developing scalable algorithms to run learning algorithms on leadership class platforms, incorporation of physical constraints into the model, dealing with heterogeneous data, and building trust through verified and validated results for the learning models. This minisymposium will bring together researchers developing algorithms and learning approaches to overcome these challenging issues. In addition, the discussion will include a motivation of the need for scientific machine learning while highlighting some additional challenges.

Organizer: Prasanna Balaprakash
Argonne National Laboratory, U.S.

Organizer: Eric C. Cyr
Sandia National Laboratories, U.S.

Organizer: Paris Perdikaris
University of Pennsylvania, U.S.

2:15-2:35 DeepHyper: Scalable Asynchronous Hyperparameter Search for Deep Learning
Prasanna Balaprakash, Argonne National Laboratory, U.S.

2:40-3:00 Data-Driven Multi-scale Decompositions for Forecasting and Model Discovery
Daniel Dylewsky and J. Nathan Kutz, University of Washington, U.S.; Molei Tao, Georgia Institute of Technology, U.S.

3:05-3:25 Differential Equations with Unknown Constitutive Relations as Recurrent Neural Networks
Tobias Hagge and Panos Stinis, Pacific Northwest National Laboratory, U.S.; Enoch Yeung, University of California, Santa Barbara, U.S.; Alexander Tartakovsky, Pacific Northwest National Laboratory, U.S.

3:30-3:50 Data-driven Modeling and Analysis of Non-stationary Fluid Flows
Kazuki Maeda and Steven Brunton, University of Washington, U.S.

continued in next column

Monday, February 25

MS46

Task-based Programming for Scientific Computing: Runtime Support - Part II of II

2:15 p.m.-3:55 p.m.

Room: 300C

For Part 1 see MS12

As the architecture of modern supercomputing systems is becoming increasingly complex and heterogeneous, runtime systems have recently emerged as a solid alternative to traditional parallel programming paradigms and tools because they allow for the delevopment of efficient and scalable code which is portable across a wide range of architectures thanks to the use of a modular approach and architecture-agnostic parallel programming paradigms. Numerous runtimes exist which propose different programming interfaces and features, this minisymposium presents recent advances and developments in some of the best known and used runtime systems.

Organizer: Alfredo Buttari

CNRS, France

Organizer: Emmanuel Agullo

Inria, France

Organizer: Hatem Ltaief

King Abdullah University of Science & Technology (KAUST), Saudi Arabia

2:15-2:35 Performance and Implementation of UPC++

Steven Hofmeyr, Lawrence Berkeley National Laboratory, U.S.

2:40-3:00 Dynamic Algorithms with PaRSEC

George Bosilca, Thomas Herault, Aurélien Bouteiller, and Damien Genet, University of Tennessee, Knoxville, U.S.; Jack J. Dongarra, University of Tennessee and Oak Ridge National Laboratory, U.S.

3:05-3:25 SPETABARU: A Task-based Runtime System with Speculative Execution Capability

Berenger Bramas, Inria, France

3:30-3:50 How to Adjust an OpenMP Runtime to Better Reach Hardware Performance on HPC Applications: libKOMP's Experiences?

Thierry Gautier and *Philippe Virouleau*, Inria, France

continued in next column

Monday, February 25

MS47

Advances in Computational Methods for Data Assimilation - Part I of II

2:15 p.m.-3:55 p.m.

Room: 300D

For Part 2 see MS87

Data assimilation is a mathematical procedure that estimates the state of some hidden system based on observation data. It is one of the most important research topics in computational science and engineering and it has many practical applications. A main effort to solve the data assimilation problem is the simulation based computational methods, which is usually combined by forward simulations and inverse inferences. With the progress in uncertainty quantification and the advances in large scale computing, computational methods for data assimilation attract more and more attentions. In this minisymposium, we present recent research achievements in computational methods for data assimilation, and we explore both the forward simulations and the inverse inferences for the data assimilation problem.

Organizer: Feng Bao

Florida State University, U.S.

Organizer: Xuemin Tu

University of Kansas, U.S.

2:15-2:35 Data Assimilation Subject to Practical Imperfections

Henry Huang, Pacific Northwest National Laboratory, U.S.

2:40-3:00 A Direct Filter Method for Parameter Estimation

Feng Bao, Florida State University, U.S.

3:05-3:25 An Adaptive Multi-fidelity PC Based Iterative Ensemble Kalman Method for Inverse Problem

Tao Zhou, Chinese Academy of Sciences, China

3:30-3:50 An Improved Particle Filter and its Application for Tracking in a Wireless Sensor Network

Kai Kang, National Institutes of Health, U.S.

Monday, February 25

MS48

BE: Best Practices for CSE Diversity and Inclusion

2:15 p.m.-3:55 p.m.

Room: 302A

Developing the capacity for innovation and productivity in the global marketplace relies on the ability to develop an outstanding science and technology (S&T) workforce. Recent economic projections call for one million more science, technology, engineering, and mathematics (STEM) professionals over the next decade in the United States (US). Yet, less than 40 percent of entering college students who intend to major in STEM disciplines actually complete a STEM degree. Disparities in the representation in by women and minorities compound the challenge, and is more pronounced in computational science and engineering (CSE) fields. STEM and CSE fields have been historically dominated by a white, male population; however, data indicates that there has been movement towards diversification in academics and employment. Yet, the underrepresentation of women and minorities continues to be problematic, especially in CSE. This session will include a series of presentations on programs and activities aimed at promoting diversity and creating more inclusive environments. The intention is that attendees will learn about best practices that can be employed in a variety of settings and can participate in engaging discussion on the topic.

Organizer: Mary Ann E. Leung
Sustainable Horizons Institute, U.S.

2:15-2:35 Sustainable Research Pathways: Building Connections Across Communities to Diversify the National Laboratory Workforce
Mary Ann E. Leung, Sustainable Horizons Institute, U.S.

2:40-3:00 Creating and Sustaining a Research Network for Women in Mathematics of Materials
Malena I. Espanol, University of Akron, U.S.

3:05-3:25 Effective Workplace Communication
Rebecca J. Hartman-Baker, Lawrence Berkeley National Laboratory, U.S.

3:30-3:50 Diversifying the Workforce at Lawrence Livermore National Laboratory
Tony Baylis, Lawrence Livermore National Laboratory, U.S.

Monday, February 25

MS49

Emerging Trends for Structure Preserving Model Order Reduction - Part II of II

2:15 p.m.-3:55 p.m.

Room: 302B

For Part 1 see MS14
Model Order Reduction (MOR) is an indispensable tool for multi-query simulations and (faster than) real-time estimation and control, in particular for multiscale and multiphysics problems. Standard MOR approaches, such as, the Reduced Basis Method, the Proper Orthogonal Decomposition, etc. in conjunction with the (Discrete) Empirical Interpolation Method, fail to strike an acceptable balance between fidelity and complexity to obtain computationally compact representations of several phenomena of interest (for instance, multiple-transport phenomena). The limitations of standard MOR techniques instigate a drive to develop novel and automated techniques tailored for identification of underlying embedding (structure), and, preservation of key (physical) model properties (such as, positivity, evolving discontinuous front, etc.) and system theoretic properties (for instance, Port-Hamiltonian structure, propagation delays, stability, etc.) of the associated (parametric) reduced-order models, which are guaranteed to offer robust approximation with lowest possible reduced-order representation. In this minisymposium, we will discuss recent advances and perspectives in the field of structure and property preserving (discretization and) MOR frameworks to deal with (large-scale) linear and non-linear dynamical systems.

Organizer: Harshit Bansal
Eindhoven University of Technology, Netherlands

continued in next column

continued on next page

Monday, February 25

MS49

Emerging Trends for Structure Preserving Model Order Reduction - Part II of II

continued

2:15-2:35 Randomized Residual-based Error Estimators for Parametrized Equations

Kathrin Smetana, University of Twente, Netherlands; Olivier Zahm, Inria Grenoble, France; Anthony T. Patera, Massachusetts Institute of Technology, U.S.

2:40-3:00 From Data to Structured Reduced-order Models

Benjamin Unger, Technische Universität Berlin, Germany

3:05-3:25 Operator Splitting and Model Reduction

Klajdi Sinani, Serkan Gugercin, Christopher A. Beattie, and Jeff Borggaard, Virginia Tech, U.S.

3:30-3:50 Model Reduction of Synchronized Lur'e Networks

Xiaodong Cheng and *Jacquelien M.A. Scherpen*, University of Groningen, Netherlands

Monday, February 25

MS50

Fast Methods for High-Frequency Wave Propagation - Part II of II

2:15 p.m.-3:55 p.m.

Room: 303A

For Part 1 see MS15

High-frequency wave equations, e.g., variable-coefficient Helmholtz equations, govern many large-scale electromagnetic, acoustic, and elastodynamic phenomena in inhomogeneous media. Their numerical solution remains a grand challenge due to difficulties in convergence, dispersion, and computational efficiency in the high-frequency regime. This minisymposium reports recent progress in low-complexity solvers including directional methods, sweeping preconditioners, method of polarized traces, as well as reliable asymptotic methods such as geometrical optics, Gaussian beam methods, and Huygens sweeping methods.

Organizer: Leonardo Zepeda-Nunez

Lawrence Berkeley National Laboratory, U.S.

2:15-2:35 Frozen Gaussian Approximation for 3-D Elastic Wave Propagation and Seismic Tomography

James Hateley, University of California, Santa Barbara, U.S.; Lihui Chai, National Sun Yat-Sen University, Taiwan; Ping Tong, Nanyang Technological University, Singapore; *Xu Yang*, University of California, Santa Barbara, U.S.

2:40-3:00 Recent Progress on Solving High Frequency Wave Equations

Lexing Ying, Stanford University, U.S.

3:05-3:25 Learning Dominant Wave Directions For High-frequency Helmholtz Equations

Jun Fang, University of California, Irvine, U.S.; Jianliang Qian, Michigan State University, U.S.; *Hongkai Zhao*, University of California, Irvine, U.S.; Leonardo Zepeda, Lawrence Berkeley National Laboratory, U.S.

3:30-3:50 High-order Galerkin Method for Helmholtz and Laplace Problems on Multiple Open Arcs

Carlos Jerez-Hanckes, Pontificia Universidad Católica de Chile, Chile

continued in next column

Monday, February 25

MS51

Tutorials for Students: Accessible Introductions to Active Research Areas - Part II of II

2:15 p.m.-3:55 p.m.

Room: 303B

For Part 1 see MS16

This set of tutorials provides accessible introductions for students to areas of active research in compu- tational science and engineering. It features four 45 minute tutorials which represent various areas of active research that will recur in talks throughout the conference. It serves to give students from various areas of computational science and engineering the necessary background to benet from other talks throughout the conference, and future conferences as these active elds continue to grow. In particular, these four tutorials will touch on aspects of data driven modeling, integrating data in computational models, and algorithmic methodologies.

Organizer: Joseph L. Hart
North Carolina State University, U.S.

Organizer: Kathleen Kavanagh
Clarkson University, U.S.

2:15-3:00 Mathematical Optimization: Modern Algorithms for Applications in Data Science

Ekkehard W. Sachs, Universität Trier, Germany

3:05-3:50 Introduction to Neural Networks and Deep Learning

Jeff Humphreys, Brigham Young University, U.S.

Monday, February 25

MS52

Advances in Phase Retrieval: Theory and Applications - Part II of II

2:15 p.m.-3:55 p.m.

Room: 201B

For Part 1 see MS17

Phase retrieval is the process of recovering the phase of a complex signal given only its Fourier magnitude. This problem is known to be non-convex and plays a vital role in various scientific applications. Improved phase recovery methods and its applications are therefore an active area of research. This minisymposium reports on recent advances in phase retrieval ranging from algorithm advancement, experimental design, to application in high-energy x-ray science.

Organizer: Selin Aslan
Argonne National Laboratory, U.S.

Organizer: Zichao Di
Argonne National Laboratory, U.S.

2:15-2:35 Advances in Ptychography Applications: Building Tools Towards Collaborative Development

Pablo Enfedaque and Stefano Marchesini, Lawrence Berkeley National Laboratory, U.S.; Huibin Chang, Tianjin Normal University

2:40-3:00 Ptychography: An Algorithmic Playground

Andrew M. Maiden, University of Sheffield, United Kingdom

3:05-3:25 Sparse Dictionary Learning Methods for Coherent X-ray Diffractive Imaging

Ashish Tripathi, Brendt Wohlberg, and Richard Sandberg, Los Alamos National Laboratory, U.S.

3:30-3:50 On the use of Automatic Differentiation as a Generic Phase Retrieval Framework

Saugat Kandel and Ming Du, Northwestern University, U.S.; Siddharth Maddali, Argonne National Laboratory, U.S.; Marc Allain, Institute Fresnel, France; Chris Jacobsen, Stephan O. Hruszkewycz, and Youssef Nashed, Argonne National Laboratory, U.S.

Monday, February 25

MS53

Computational and Numerical Methods in Electronics - Part II of II

2:15 p.m.-3:55 p.m.

Room: 201C

For Part 1 see MS18

Research in electronics faces two challenges which are unfortunately often considered separate from each other: First, the rapid technological progress requires short simulation times for electronic devices with increasing geometric complexity. Such short simulation times can only be achieved through the development of new parallel algorithms since the serial performance of modern computing hardware stagnates. Second, sound numerical methods are required to accurately describe physical processes such as electronic transport at the discrete level. For example, important physical properties such as positivity and charge conservation need to be reflected in the numerics. This minisymposium - with the aim to bridge the gap between both challenges - has been initially founded with the previous SIAM Conference on Computational Science and Engineering in 2017. The continuation in 2019 will allow the minisymposium to establish itself as a series which will in turn enable to provide a continued presence as an interdisciplinary platform for physicists, numerical analysts, and computational scientists within the area of electronics. There is currently no other venue providing such a focused forum further underlining the importance of this minisymposium series.

continued on next page

Monday, February 25

MS53

Computational and Numerical Methods in Electronics - Part II of II

continued

Organizer: Karl Rupp
Technische Universität Wien, Austria

Organizer: Daniel Brinkman
San Jose State University, U.S.

Organizer: Patricio Farrell
Hamburg University of Technology, Germany

Organizer: Nella Rotundo
Weierstrass Institute, Germany

Organizer: Josef Weinbub
Technische Universität Wien, Austria

2:15-2:35 Multi-scaled Simulations and Modelling of Novel Electronic Devices
Vihar Georgiev, University of Glasgow, Scotland, UK

2:40-3:00 GPU vs CPU - Comparative Performance Evaluation of Parallel Linear Solvers for TCAD
Karl Rupp, Technische Universität Wien, Austria

3:05-3:25 Algorithms for Large-scale Ab Initio Quantum Transport Simulations
Mathieu Luisier, Fabian Ducry, Sascha Brück, M. Hossein Bani-Hashemian, and Mauro Calderara, ETH Zürich, Switzerland

3:30-3:50 Recent Advances in High Performance Process TCAD
Georgios Diamantopoulos, Paul Manstetten, Lukas Gnam, Vito Simonka, Luiz Felipe Aguinsky, Michael Quell, Alexander Toifl, Andreas Hössinger, and *Josef Weinbub*, Technische Universität Wien, Austria

Monday, February 25

MS54

Fast Solvers for Inverse Problems with PDEs - Part II of II

2:15 p.m.-3:55 p.m.

Room: 202A

For Part 1 see MS19

We present recent advances in the development of fast solvers and effective numerical schemes for the solution of inverse problems with a particular focus on large-scale PDE-constrained optimization problems. Key algorithmic challenges include computational complexity, memory consumption, and a vast number unknowns as well as model and data uncertainties. We showcase state-of-the-art techniques in scientific computing to tackle these challenges.

Organizer: James L. Herring
University of Houston, U.S.

Organizer: Andreas Mang
University of Houston, U.S.

Organizer: George Biros
University of Texas at Austin, U.S.

2:15-2:35 Efficient Marginalization-based MCMC Approaches for Hierarchical Bayesian Inverse Problems
Arvind Saibaba, North Carolina State University, U.S.; Johnathan M. Bardsley, University of Montana, U.S.; Andrew Brown, Clemson University, U.S.; Alen Alexanderian, North Carolina State University, U.S.

2:40-3:00 Title Not Available
James G. Nagy, Emory University, U.S.

3:05-3:25 Title Not Available
Georg Stadler, Courant Institute of Mathematical Sciences, New York University, U.S.

3:30-3:50 LAP— A Fast Solver for Coupled Imaging Problems
James L. Herring, University of Houston, U.S.

Monday, February 25

MS55

Moment Methods in Kinetic Theory - Part II of II

2:15 p.m.-3:55 p.m.

Room: 202B

For Part 1 see MS20

Kinetic theory enables modeling of numerous physical processes such as radiative transfer or rarefied gas flows in the transition or kinetic regime for moderate to large Knudsen numbers. Here the Knudsen number relates the mean free path of gas particles to the reference length of the system. Instead of tracking individual particles, a velocity distribution function is used to model the dynamics of the ensemble. Solutions of the underlying Boltzmann equation or variants thereof are often computed using moment methods. To that extend, a set of moments of the distribution function, including e.g. physical quantities like density, momentum, energy, heat flux, stress tensor and other higher moments, is chosen to describe the flow properties. The resulting moment equations then describe the evolution of these moments. In this minisymposium we want to show recent advances and applications of moment methods that highlight the applicability and accuracy of moment methods in comparison to standard fluid dynamics approaches and particle-based models. This includes the ongoing search for efficient yet accurate closures for moment equations as well as numerical methods ensuring asymptotic stability. On the other hand, recent developments of hyperbolic moment models shall be discussed together with first results for adaptive moment models and predictive hierarchical simulations. At the end, we want to exemplify various applications of moment methods, e.g. for polydisperse sprays.

continued on next page

Organizer: Julian Koellermeier
Freie Universität Berlin, Germany

Organizer: Yuwei Fan
Stanford University, U.S.

2:15-2:35 Moment Methods for Polydisperse Sprays with Evaporation

James McDonald and Mathieu Giroux, University of Ottawa, Canada; Lucian Ivan, Canadian Nuclear Laboratories, Canada

2:40-3:00 Nonlinear Moment Model for Radiative Transfer Equation

Ruo Li, Peking University, China

3:05-3:25 Unified Gas Kinetic Particle Method for Multiscale Photon Transport

Weiming Li, Beijing Computational Science Research Center and Hong Kong University of Science & Technology, China; Chang Liu, Hong Kong University of Science and Technology, Hong Kong; Yajun Zhu, Northwestern Polytechnical University, China; Jiwei Zhang, ; Kun Xu, Hong Kong University of Science and Technology, Hong Kong

3:30-3:50 Asymptotic-preserving Lax-Wendroff Discontinuous Galerkin Schemes for Quasi-exponential Moment-closure Approximations of Kinetic Models

James A. Rossmanith, Iowa State University, U.S.

Monday, February 25

MS56

Recent Advances in High Order Lagrangian/ALE Methods - Part II of II

2:15 p.m.-3:55 p.m.

Room: 202C

For Part 1 see MS21

The numerical approximation of the Euler equations of gas dynamics in a moving frame is a common approach for solving many multiphysics problems involving e.g. large deformations, strong shocks and interactions of multiple materials. In Lagrangian methods, the mesh is moving with the fluid velocity, therefore they are well-suited for accurate resolution of material interfaces. On the other hand, multidimensional Lagrangian meshes tend to tangle so that the mesh elements become invalid, and in general cannot represent large deformation. This difficulty can be resolved in ALE (Arbitrary-Eulerian-Lagrangian) methods, which assume that the mesh moves independently on the flow and therefore offer additional flexibility and accuracy. ALE techniques can also be directly implemented in numerical schemes, for example to take into account moving objects, in fluid structure interaction, etc. In both cases, the question of the mesh quality is a central question. The aim of this minisymposium is to bring together the researchers working in the fields of Lagrangian hydrodynamics and ALE methods to discuss the state-of-the-art of single- and multi-material hydrodynamic simulations. We are particularly interested in recent advances in high order methods, such as high order finite volume (WENO, ADER), discontinuous Galerkin, high order finite elements, residual distribution methods, as applied to Lagrangian or ALE description of the flow.

Organizer: Svetlana Tokareva
Los Alamos National Laboratory, U.S.

Organizer: Remi Abgrall
Universität Zürich, Switzerland

2:15-2:35 Using Virtual Element Method for Data Remap on Polytopal Meshes

Konstantin Lipnikov and Nathaniel Morgan, Los Alamos National Laboratory, U.S.

2:40-3:00 Invariant Domains Preserving ALE Approximation of Hyperbolic Systems with Continuous Finite Elements

Jean-Luc Guermond and Bojan Popov, Texas A&M University, U.S.; *Laura Saavedra*, Polytechnic University of Madrid, Spain

3:05-3:25 A New High Order Cell-centered ALE Method with Uniform Accuracy in Both Space and Time

Yibing Chen, Institute of Applied Physics and Computational Mathematics Beijing, China

3:30-3:50 Exploration of a High-order Lagrangian Discontinuous Galerkin Method on Quadratic Triangular Meshes

Xiaodong Liu, Nathaniel Morgan, and Donald E. Burton, Los Alamos National Laboratory, U.S.

continued in next column

Monday, February 25

MS57

Use of C++ in Computational Science Libraries and Applications - Part II of II

2:15 p.m.-3:55 p.m.

Room: 203

For Part 1 see MS22

The number of C++ libraries and applications in computational science have grown over the past years. We will be presenting some of them and the advancements in the C++ standard that have contributed to this growth. The talks in this minisymposium will focus on the relevant aspects of the C++ standard and accompanying libraries. Additionally, practical aspects and implementation methods will be shown to deal with a wide range of computational science codes on many modern platforms in high performance manner. This will include HPC aspects of modern C++ and its impact of expresivity of the code. The talks will also include interfacing to modern Fortran code and how the integration between C++ and Fortran may be made easier or even completely automated.

Organizer: Piotr Luszczek

University of Tennessee, U.S.

Organizer: Mark Hoemmen

Sandia National Laboratories, U.S.

Organizer: Heike Jagode

University of Tennessee, U.S.

Organizer: Damien Genet

University of Tennessee, Knoxville, U.S.

2:15-2:35 Data Flow Graph Programming for High-performance Scientific Computing in C++

Edward F. Valeev, Virginia Tech, U.S.

2:40-3:00 The Simulation Development Environment (SDE): A C++ Framework for Reusable Computational Chemistry

Ryan M. Richard, Iowa State University, U.S.

3:05-3:25 Automated Fortran-C++ Bindings for Scientific Applications

Seth Johnson, Oak Ridge National Laboratory, U.S.

Monday, February 25

MS58

Finite Element Methods on Unfitted Meshes - Part II of II

2:15 p.m.-3:55 p.m.

Room: 205

For Part 1 see MS23

Finite element method (FEM) solving partial differential equations has been largely focusing on the equations and overlooking the coupling of the domain or interface with complex or evolving geometry. For problems with difficult geometries, classical FEMs that employ fitted meshes often require highly non-trivial and costly steps to generate qualified meshes, especially in the three dimensions. It is then extremely advantageous to employ meshes that do not necessarily fit the physical geometry. This minisymposium focuses on most recent techniques to design stable and robust FEMs that employ the unfitted meshes.

Organizer: Cuiyu He

University College London, United Kingdom

2:15-2:35 An Unfitted Interface Penalty Finite Element Method for Elliptic Interface Problems

Peiqi Huang, Nanjing Forestry University, China

2:40-3:00 A Conforming Enriched Finite Element Method for Interface Problems

Jinru Chen, Nanjing Normal University, China

3:05-3:25 Unconditionally Stable Cutfem for Dynamic Interfaces in a Fluid Structure Interaction Problem

Sarkis Marcus, Worcester Polytechnic Institute, U.S.

3:30-3:50 Numerical Study of Lateral Phase Separation in Two-component Bio-membranes with TraceFEM

Vladimir Yushutin, University of Houston, U.S.

Monday, February 25

MS59

Applications of the AMReX Block Structured Adaptive Mesh Refinement Framework - Part II of II

2:15 p.m.-3:55 p.m.

Room: 206A

For Part 1 see MS24

Multiphysics PDE simulation at scale can be a challenging task given rapidly evolving supercomputing architectures. As part of the DOE Exascale Computing Project (ECP), researchers at LBNL, NRL, and ANL have publicly released the AMReX software framework for massively parallel block-structured adaptive mesh refinement (AMR) applications. The core libraries include linear solvers, particle support, embedded boundary geometry representation, profiling tools, and hybrid parallelism. AMReX is the basis for many mature scientific applications, including several supported by ECP and SciDAC (e.g., combustion, carbon capture/storage, accelerator physics, and astrophysics, to name a few). In this minisymposium, 8 computational scientists from a wide range of fields will showcase their latest code innovations and scientific achievements using the AMReX software framework. Additionally, each will discuss many aspects of their problem, as listed below. Who should attend this minisymposium? -Researchers interested in multiphysics PDE simulation techniques at scale. This includes all aspects of the problem from model equation development, numerical discretizations, implementation in HPC frameworks, utilization of supercomputing resources, and analysis of data. -Researchers in fields represented by the speakers, including combustion, materials science, multiphase flow, astrophysics, stochastic PDEs modeling fluctuating hydrodynamics, and accelerator physics.

continued on next page

Organizer: Andy J. Nonaka
Lawrence Berkeley National Laboratory, U.S.

2:15-2:35 Castro: A Compressible Astrophysical Hydrodynamics Code
Maria G. Barrios Sazo, Stony Brook University, U.S.; Ann S. Almgren, Lawrence Berkeley National Laboratory, U.S.; Max Katz, Nvidia Corporation, U.S.; Donald E. Willcox and Weiqun Zhang, Lawrence Berkeley National Laboratory, U.S.; Michael Zingale, Stony Brook University, U.S.

2:40-3:00 Solving Low Mach Number Stellar Hydrodynamics using MAESTROeX
Duoming Fan, Lawrence Berkeley National Laboratory, U.S.

3:05-3:25 WarpX: Exascale Modeling of Advanced Particle Accelerators with AMReX
Maxence Thevenet, Rémi Lehe, Andrew Myers, and Weiqun Zhang, Lawrence Berkeley National Laboratory, U.S.; Mark Hogan, SLAC National Accelerator Laboratory, U.S.; Cho-Kuen Ng, Stanford Linear Accelerator Center, U.S.; Olga Shapoval and Jaehong Park, Lawrence Berkeley National Laboratory, U.S.; David Grote, Lawrence Livermore National Laboratory, U.S.; John B. Bell, Lawrence Berkeley National Laboratory, U.S.; Lixin Ge, Stanford Linear Accelerator Center, U.S.; Ann S. Almgren and Jean-Luc Vay, Lawrence Berkeley National Laboratory, U.S.

3:30-3:50 Fluctuating Hydrodynamics Simulations of Reactive Electrolyte Solutions
Changho Kim, University of California, Merced, U.S.; Andrew Nonaka and John B. Bell, Lawrence Berkeley National Laboratory, U.S.; Alejandro Garcia, San Jose State University, U.S.; Aleksandar Donev, Courant Institute of Mathematical Sciences, New York University, U.S.

Monday, February 25

MS60

Discovering and Exploiting Low-dimensional Structures in Computational Models - Part II of II

2:15 p.m.-3:55 p.m.

Room: 206B

For Part 1 see MS25

High-dimensional spaces of input parameters and observable output data are commonly encountered in the modeling of practical problems. Unfortunately, many of the most reliable and well understood methods we rely upon to numerically solve models, perform optimization, and quantify uncertainties suffer from some form of the "curse of dimensionality." Recently, there is great interest in developing methods that identify and exploit lower-dimensional structures in high-dimensional spaces. This session focuses on recent advances in methodologies and approaches for discovering and exploiting low-dimensional structures in the input and output spaces of models.

Organizer: Steven A. Mattis
Technische Universität München, Germany

Organizer: Troy Butler
University of Colorado, Denver, U.S.

2:15-2:35 Exploiting Low-dimensional Structure to Efficiently Perform Stochastic Inference for Prediction
Timothy Wildey, Sandia National Laboratories, U.S.; Troy Butler, University of Colorado, Denver, U.S.; John D. Jakeman, Sandia National Laboratories, U.S.

2:40-3:00 Lifting Nonlinear Systems: More Structure, More Opportunities for ROM?
Boris Kramer, Massachusetts Institute of Technology, U.S.; Karen E. Willcox, University of Texas at Austin, U.S.

3:05-3:25 Scalable Online Nonlinear Goal-oriented Inference with Physics-Informed Maps
Harriet Li, Massachusetts Institute of Technology, U.S.; Karen E. Willcox, University of Texas at Austin, U.S.

3:30-3:50 Exploiting Low-dimensional Structure for Multifidelity Simulations
Akil Narayan, University of Utah, U.S.

Monday, February 25

MS61

Parallel-in-Time Integration Techniques - Part II of II

2:15 p.m.-3:55 p.m.

Room: 206C

For Part 1 see MS26

For the efficient use of massively parallel supercomputers for applications in science and engineering new parallel algorithms are required. The idea of parallelizing in the time dimension and, thus, adding an additional layer of parallelism to a numerical algorithm has become an increasing popular way to tackle evolutionary problems. This minisymposium features talks on both theoretical and computational aspects of parallel-in-time integration. Research topics ranging from algorithm development and mathematical analysis to software implementation and exploitation of parallel-in-time integration techniques in realistic applications are discussed.

Organizer: Stephanie Friedhoff
University of Wuppertal, Germany

Organizer: Debasmita Samaddar
UK Atomic Energy Authority, United Kingdom

2:15-2:35 pySDC: A Parallel Framework for Spectral Deferred Corrections
Robert Speck, Jülich Supercomputing Centre, Germany

2:40-3:00 Parallel-in-Time Simulation of the Schrödinger Equation
Hannah Rittich, Forschungszentrum Jülich, Germany

3:05-3:25 Adaptive Parareal: Introducing Time-grid Adaptivity to Improve the Stability of the Parareal Algorithm
Federico Danieli, University of Oxford, United Kingdom

3:30-3:50 Multigrid Based Strategies for the Solution of Non-linear Space-time Problems
Pietro Benedusi, Rolf Krause, and Patrick Zulian, Università della Svizzera italiana, Switzerland

Monday, February 25

MS62

Modeling, Analyzing and Computing of Fractional Partial Differential Equations - Part II of II

2:15 p.m.-3:55 p.m.

Room: 206D

For Part 1 see MS27

The minisymposium covers modeling analyzing, and computing of fractional partial differential equations (PDEs) in physics, mechanics, engineering, and biology etc. The interdisciplinary talks stimulate research at the crossroads of computational science and engineering, and applied mathematics. The minisymposium reports novel numerical techniques for fractional PDEs and their high-performance computer implementations. Works that investigate the analytical properties of the solutions of systems consisting of PDEs with derivatives of fractional orders are also of interest to the symposium. Physical modeling of real-life phenomena employing fractional derivatives is also a topic covered by this symposium. Moreover, the interplay between modeling, analysis and simulation of fractional systems is a topic of high interest to this event.

Organizer: Abdul M. Khaliq

Middle Tennessee State University, U.S.

Organizer: Khaled Furati

King Fahd University of Petroleum and Minerals, Saudi Arabia

2:15-2:35 Modeling and Finite Difference Computing of Fractional Convection Equation

Changpin Li, Shanghai University, China

2:40-3:00 Anomalous Diffusion, Dilation, and Erosion in Image Processing

Andreas Kleefeld, Jülich Supercomputing Centre, Germany

3:05-3:25 Numerical Methods for Multidimensional Fractional Reaction-diffusion Systems in Developmental Biology

Ashlin Harris and Abdul M. Khaliq, Middle Tennessee State University, U.S.

3:30-3:50 Inverse Source Problems for Fractional Differential Equations with Nonlocal Boundary Conditions

Salman A. Malik, COMSAT University Islamabad, Pakistan

continued in next column

Monday, February 25

MS63

FASTMath Tools and Technologies - Part II of II

2:15 p.m.-3:55 p.m.

Room: 207

For Part 1 see MS28

The FASTMath (Frameworks, Algorithms and Scalable Technologies for Mathematics) Institute is a R&D project funded by the SciDAC Program at the U.S. Department of Energy (DOE). The goal of FASTMath is to develop and deploy scalable mathematical algorithms and software tools for reliable simulation of complex physical phenomena and collaborating with DOE domain scientists to ensure the usefulness and applicability of the work in the project. The focus of FASTMath is strongly driven by the requirements of DOE application scientists who require fast, accurate, and robust forward simulation along with the ability to efficiently perform ensembles of simulations in optimization or uncertainty quantification studies. This minisymposium will present work from the eight core areas of FASTMath: structured meshes, unstructured meshes, time integrators, linear solvers, eigensolvers, numerical optimization, data analytics, and uncertain quantification.

Organizer: Esmond G. Ng

Lawrence Berkeley National Laboratory, U.S.

2:15-2:35 Nonlinear Optimization using the Toolkit for Advanced Optimization

Adam Denchfield, University of Illinois at Chicago, U.S.; Alp Dener and *Todd Munson*, Argonne National Laboratory, U.S.

continued on next page

2:40-3:00 Uncertainty Quantification in Computations of Large Scale Physical Systems

Habib N. Najm, Sandia National Laboratories, U.S.

3:05-3:25 UQTk, a FASTMath C++/ Python Toolkit for Uncertainty Quantification: Overview and Applications

Bert Debussschere, Khachik Sargsyan, and Cosmin Safta, Sandia National Laboratories, U.S.

3:30-3:50 Sparse Representations of Data for HPC

Rick Archibald, Oak Ridge National Laboratory, U.S.

Monday, February 25

MS64

Computational Methods for Kinetic Models of Plasma - Part II of II

2:15 p.m.-3:55 p.m.

Room: 401A

For Part 1 see MS29

Statistical mechanics provides a mathematical modeling framework in which a collection of "particles" (e.g., electrons, nucleons, atoms, or molecules) is represented by a probability density function (PDF). This representation allows one to model the influence of microscopic dynamics on larger scales, without explicitly computing the detailed dynamics on the microscopic scale. Models in this framework are referred to as "kinetic models", and arise in many application areas, including in rarefied gas dynamics, nuclear reactor modeling, and plasma physics. Transport phenomena refers to the evolution and redistribution of macroscopic quantities such as mass, momentum, energy, and heat flux. This minisymposium addresses recent advances in computational methods for the simulation of statistical mechanical models and the resulting transport phenomena. Our goal is to showcase a diverse array of application areas, models, and numerical techniques.

Organizer: James A. Rossmanith

Iowa State University, U.S.

2:15-2:35 Implicit, Conservative Multiscale Algorithms for Hybrid Kinetic (Particle) Ion/Fluid Electron Models

Luis Chacon, Adam Stanier, and Guangye Chen, Los Alamos National Laboratory, U.S.

2:40-3:00 Computational Methods for Multiscale Levy-Fokker-Planck Equation

Li Wang, State University of New York, Buffalo, U.S.

3:05-3:25 A Hybrid Discontinuous Galerkin Method for 2D-2V Vlasov-Poisson Problems with Geometry

David C. Seal, United States Naval Academy, U.S.

3:30-3:50 Simulation of Interpenetrating Plasmas in 2D using the Grid Based Continuum Code LOKI

Jeffrey W. Banks, Rensselaer Polytechnic Institute, U.S.; Richard Berger, Lawrence Livermore National Laboratory, U.S.; Stephan Brunner, École Polytechnique Fédérale de Lausanne, Switzerland; Thomas Chapman, Andris Dimits, and Debojyoti Ghosh, Lawrence Livermore National Laboratory, U.S.

continued in next column

Monday, February 25

MS65

Graph and Combinatorial Algorithms for Enabling Exascale Applications - Part II of II

2:15 p.m.-3:55 p.m.

Room: 401B

For Part 1 see MS30

Combinatorial algorithms in general and graph algorithms in particular play a critical enabling role in numerous scientific applications. The low computation to communication ratios and irregular memory accesses of these algorithms makes them some of the hardest algorithmic kernels to implement on parallel systems. The talks in this minisymposium will consider: (i) exascale applications that drive the selection of combinatorial kernels and integration of software tools; (ii) combinatorial (graph) kernels such as graph traversals, graph matching, graph coloring, graph clustering, and graph partitioning, that play a crucial enabling role in the chosen application areas; and (iii) software frameworks for efficient implementation of the algorithms on hierarchical distributed-memory architectures that are representative of potential exascale platforms.

Organizer: Alex Pothen

Purdue University, U.S.

Organizer: Mahantesh Halappanavar

Pacific Northwest National Laboratory, U.S.

2:15-2:35 Partitioning Strategies for Exascale Computing

Erik G. Boman, Karen D. Devine, Siva Rajamanickam, and Michael W. Wolf, Sandia National Laboratories, U.S.

2:40-3:00 Graph Algorithms in Scalable Implementations of Computational Chemistry Methods

Sriram Krishnamoorthy, Pacific Northwest National Laboratory, U.S.

3:05-3:25 Accelerating Biological Network Clustering with Sparse Matrix Kernels on Large-scale System

Oguz Selvitopi, Lawrence Berkeley National Laboratory, U.S.; Ariful Azad, Indiana University, U.S.; Aydin Buluc, Lawrence Berkeley National Laboratory, U.S.

3:30-3:50 Parallelism in Deep Learning through Linear Algebra

Yu-Hong Yeung and Alex Pothen, Purdue University, U.S.; Maxim Naumov, NVIDIA, U.S.; Mahantesh Halappanavar, Pacific Northwest National Laboratory, U.S.

continued in next column

Monday, February 25

MS66

Numerical Methods for Earth System Modeling - Part II of II

2:15 p.m.-3:55 p.m.

Room: 401C

For Part 1 see MS31

Earth system models are complex multiscale, multiphysics systems that require advanced discretization schemes, temporal integration methods, and algebraic solvers to obtain accurate and stable solutions on today's high-performance computing systems. This minisymposium brings together researchers to present state of the art methods in these areas applied to Earth system modeling. It also focuses on real-world challenges and opportunities encountered while implementing these schemes on current and upcoming computing systems.

Organizer: Christopher J. Vogl

Lawrence Livermore National Laboratory, U.S.

Organizer: David J. Gardner

Lawrence Livermore National Laboratory, U.S.

2:15-2:35 Diagnosing Physically-simulated Mixing in MPAS-O and E3SM via Online Lagrangian Particle Tracking with LIGHT

Phillip J. Wolfram, Todd Ringler, Luke Van Roekel, Mark R. Petersen, and Mathew Maltrud, Los Alamos National Laboratory, U.S.; Riley Brady, University of Colorado Boulder, U.S.; Jon Wolfe, Adrian Turner, Philip Jones, and Xylar Asay-Davis, Los Alamos National Laboratory, U.S.

2:40-3:00 Modeling Antarctic Ice Sheet Dynamics using Adaptive Mesh Refinement

Daniel Martin, Lawrence Berkeley National Laboratory, U.S.

3:05-3:25 Development of Terrestrial Dynamical Core for Earth System Models

Nathan Collier, Oak Ridge National Laboratory, U.S.

Monday, February 25

MS67

Recent Advances in Error Estimation for Complex Problems - Part II of II

2:15 p.m.-3:55 p.m.

Room: 402A

For Part 1 see MS32

The accurate simulation of complex systems which incorporate multiple physical processes interacting across a range of spatial and temporal scales demands sophisticated multi-discretization numerical methods. These include operator decomposition and domain decomposition, as well as specialized time discretizations such as explicit, semi-implicit, operating splitting and parallel-in-time methods. Each of these discretization choices introduces a new source of error and the various sources of discretization error interact in non-trivial ways. Accurate error estimation is essential for building confidence in numerical solutions and is a key component of uncertainty quantification for multiphysics simulations. This minisymposium will feature recent advances in a posteriori error estimation for simulation and uncertainty quantification and the discretization choices necessary to meet constraints on accuracy.

Organizer: Jehanzeb H. Chaudhry

University of New Mexico, U.S.

Organizer: Simon Tavener

Colorado State University, U.S.

2:15-2:35 On a Goal-oriented Formulation for the PGD Method

Serge Prudhomme, Kenan Kergrene, and Marc Laforest, École Polytechnique de Montréal, Canada; Ludovic Chamoin, ENS Cachan, France

2:40-3:00 Estimating and Controlling "Communication" Errors in Multiphysics Problems

Simon Tavener and Don Estep, Colorado State University, U.S.

3:05-3:25 An Analysis of the Approximations and Errors in a Boundary Integral Solution of the Poisson-Boltzmann Equation

Christopher D. Cooper, Universidad Técnica Federico Santa María, Chile

3:30-3:50 Error Estimation for Computation of Ionic Current in Biomolecular Channels

Cairn Overturf and Jehanzeb H. Chaudhry, University of New Mexico, U.S.

continued in next column

Monday, February 25

MS68

Recent Progresses in Data Analysis - Part II of II

2:15 p.m.-3:55 p.m.

Room: 402B

For Part 1 see MS33

Data analysis plays a critical role in various applications and its techniques are being actively developed today. The proposed minisymposium will bring together researchers in data science, particularly in the areas of topological data analysis and optimization. In this minisymposium we will try to present both the theoretical and practical developments recently made in those areas and provide an opportunity for researchers to have conversation on bridging those areas.

Organizer: Jae-Hun Jung

State University of New York at Buffalo, U.S.

Organizer: Guohui Song

Clarkson University, U.S.

2:15-2:35 A Distributed ℓ^1 Minimization Algorithm

Guohui Song, Clarkson University, U.S.

2:40-3:00 Stability and Generalization for SGD in Pairwise Learning

Yiming Ying, State University of New York, Albany, U.S.

3:05-3:25 Complexity Bounds of First-order Methods for Large-scale Optimization

Yuyuan Ouyang, Clemson University, U.S.

3:30-3:50 Overcomplete Sparse Tensor Decomposition

Ashley Prater-Bennette, Air Force Research Laboratory, U.S.; Kenneth Carr, Northeastern University, U.S.

Monday, February 25

MS69

Methods for Large-scale Risk-averse Optimization - Part II of II

2:15 p.m.-3:55 p.m.

Room: 402C

For Part 1 see MS34

Many science and engineering applications require, e.g., the control or design of a physical system modeled by partial differential equations (PDEs) with uncertain inputs. For such problems, determining optimal solutions that are resilient to uncertainty is critical. This session focuses on stochastic optimization problems motivated by a range of applications including power networks, financial markets, petroleum engineering and structural design. In particular, we target mathematical approaches for risk aversion and probabilistic computations with an emphasis on efficient numerical methods.

Organizer: Drew P. Kouri
Sandia National Laboratories, U.S.

Organizer: Harbir Antil
George Mason University, U.S.

2:15-2:35 Risk-Averse Control of Nonlocal Pdes

Harbir Antil, George Mason University, U.S.; Drew P. Kouri, Sandia National Laboratories, U.S.; Johannes Pfefferer, Technische Universität München, Germany

2:40-3:00 Low-Rank Solvers for Optimal Control Problems Constrained by PDEs with Uncertain Inputs

Akwum Onwunta, University of Maryland, Baltimore County, U.S.

3:05-3:25 Adaptive Local Reduced Basis Approach for Stochastic Optimization with Nonlinear PDEs

Wilkins Aquino and Zilong Zou, Duke University, U.S.; Drew P. Kouri, Sandia National Laboratories, U.S.

3:30-3:50 Measuring Uncertainty and Performance of Machine Learning Algorithms with Buffered Probability of Exceedance

Matthew Norton, Naval Postgraduate School, U.S.

continued in next column

Monday, February 25

CP3

Modeling and Uncertainty Quantification II

2:15 p.m.-3:55 p.m.

Room: 201A

Chair: Veronika S. Vasylkivska, AECOM Corporation, U.S.

2:15-2:30 An Adaptive Sampling Approach for Surrogate Modeling of Expensive Computer Experiments

Ashwin Renganathan, Georgia Institute of Technology, U.S.

2:35-2:50 Characterising Failure in Hypersonic Vehicles using Unified Bayesian Networks

Sean R. Breckling and Daniele E. Schiavazzi, University of Notre Dame, U.S.

2:55-3:10 Long-term Risk Analysis of a Geologic CO2 Storage Project During the Post-injection Period

Veronika S. Vasylkivska, AECOM Corporation, U.S.; Gregory Lackey, Oak Ridge Institute for Science and Education, U.S.; Seth King, AECOM Corporation, U.S.; Andrew Wentworth, Oak Ridge Institute for Science and Education, U.S.; Nicolas Huerta, National Energy Technology Laboratory, U.S.; Christopher Creason and Jennifer DiGiulio, AECOM Corporation, U.S.; Ya-Mei Yang, Oak Ridge Institute for Science and Education, U.S.; Robert Dilmore, National Energy Technology Laboratory, U.S.

3:15-3:30 Coupling Scalable Sampling Methods for Uncertainty Quantification and Propagation

James Warner, Patrick Leser, William Leser, and Geoffrey Bomarito, NASA Langley Research Center, U.S.

3:35-3:50 Combining Physical Models and Gaussian Processes to Control Risk-sensitive Systems with Limited Data

Ehecatl A. Del Rio Chanona, Imperial College London, United Kingdom

Monday, February 25

CP4

Fluid-Structure Interactions

2:15 p.m.-3:55 p.m.

Room: 301

Chair: Yuexia Lin, Harvard University, U.S.

2:15-2:30 Fluid-Structure Interaction Analysis of Transcatheter Aortic Valve Replacement in a Patient with Bicuspid Aortic Valve

Andres Caballero, Georgia Institute of Technology, U.S.

2:35-2:50 Immersed Interface Method for Modeling Fluid-structure Interaction of Flexible Faceted Surfaces

Ebrahim (amin) M. Kolahdouz, University of North Carolina at Chapel Hill, U.S.; Amneet P.S. Bhalla, San Diego State University, U.S.; Kenneth I. Aycock and Brent A. Craven, U. S. Food and Drug Administration, U.S; Boyce E. Griffith, University of North Carolina, Chapel Hill, U.S.

2:55-3:10 Reference Map Technique: A Fully Eulerian Method for Fluid-structure Interactions

Yuexia Lin and Nicholas Derr, Harvard University, U.S.; Ken Kamrin, Massachusetts Institute of Technology, U.S.; Chris H. Rycroft, Harvard University, U.S.

3:15-3:30 Volumetric Stabilization of the Immersed Boundary Method

Ben Vadala-Roth, Simone Rossi, and Boyce E. Griffith, University of North Carolina at Chapel Hill, U.S.

3:35-3:50 An Inverse Problem Formulation of the Immersed Boundary Method

Jianfeng Yan and Jason E. Hicken, Rensselaer Polytechnic Institute, U.S.

Intermission

3:55 p.m.-4:10 p.m.

Monday, February 25

MS70

Machine Learning in Computational Science

4:10 p.m.-5:50 p.m.

Room: Ballroom 100BC

Machine learning have contributed significantly to recent advances in image and signal processing, pattern recognition, recommendation systems, natural language processing and machine translation. A large number of these machine learning concepts could be adopted to solve a wide range of problems in computational science. This minisymposium covers recent developments in machine learning algorithms for inverse problems, uncertainty quantification (UQ), uncertainty propagation, multi-fidelity models, reduced order models, surrogate models, model selection and calibration, among others. Target applications include (but not restricted to) numerical weather prediction, subsurface flow, flow in porous media.

Organizer: Vishwas Rao

Argonne National Laboratory, U.S.

Organizer: Ahmed Attia

Argonne National Laboratory, U.S.

Organizer: Adrian Sandu

Virginia Tech, U.S.

4:10-4:30 A Learning-based Approach for Data Assimilation

Azam Moosavi, Tesla, U.S.; Ahmed Attia, Argonne National Laboratory, U.S.; Adrian Sandu, Virginia Tech, U.S.

4:35-4:55 Bayesian Model Selection, Calibration and Uncertainty Quantification for Thermodynamic Properties

Noah Paulson, Elise Jennings, and Marius Stan, Argonne National Laboratory, U.S.

5:00-5:20 Inferring Black Box Functions under a Limited Computational Budget

Piyush Pandita and Ilias Bilionis, Purdue University, U.S.; Jesper Kristensen, Cornell University, U.S.; Jitesh Panchal, Purdue University, U.S.

5:25-5:45 Statistical Learning Approaches for Automatic Parameter Tuning to Increase Parallel Performance in Large CFD Simulations

Zachary Cooper, Virginia Tech, U.S.; Azam Moosavi, Tesla, U.S.; Adrian Sandu and Danesh Tafti, Virginia Tech, U.S.; Vishwas Rao, Argonne National Laboratory, U.S.

continued in next column

Monday, February 25

MS71

Scientific Machine Learning

4:10 p.m.-5:50 p.m.

Room: Conference Theater

This minisymposium will provide research highlights related to the Department of Energy (DOE) Advanced Scientific Computing Research (ASCR) "Scientific Machine Learning Basic Research Needs" workshop held in Jan-Feb 2018. The ASCR workshop considered the status, recent trends, and broad use of machine learning in scientific computing. Although use of such scientific machine learning (SciML) methods is increasing rapidly, there are several opportunities, barriers, and potential for high scientific impact through fundamental advances in the mathematical, statistical, and computational research foundations for SciML. Participants in the ASCR workshop identified three foundational priority research directions focused on domain-aware, interpretable, and robust SciML as well as three capabilities-focused priority research directions centered on data-intensive, inner-loop, and outer-loop SciML. This minisymposium will give a brief overview of the workshop outcomes and present recent research related to these priority research directions.

Organizer: Nathan Baker

Pacific Northwest National Laboratory, U.S.

4:10-4:30 Projection-based Model Reduction: Formulations for Physics-based Machine Learning

Renee Swischuk and Laura Mainini, Massachusetts Institute of Technology, U.S.; Benjamin Peherstorfer, Courant Institute of Mathematical Sciences, New York University, U.S.; *Karen E. Willcox*, University of Texas at Austin, U.S.

4:35-4:55 Learning Parameters and Constitutive Relationships with Physics Informed Deep Neural Networks

Alexander Tartakovsky, Pacific Northwest National Laboratory, U.S.

5:00-5:20 No Equations, No Variables, No Parameters, No Space, No Time: Data and the Modeling of Complex Systems

Yannis Kevrekidis, Johns Hopkins University, U.S.

5:25-5:45 Machine-learning Error Models for Approximate Solutions to Parameterized Systems of Nonlinear Equations

Brian A. Freno and Kevin T. Carlberg, Sandia National Laboratories, U.S.

Monday, February 25

MS72

Recent Progress in Coordinate-wise Descent Methods

4:10 p.m.-5:50 p.m.

Room: 102A

Coordinate-wise descent methods (CDMs) are a class of optimization methods which minimize the objective function along coordinate directions. The history of CDM dates to the early development of optimization. Recently the raise of super high dimensional problems in data sciences and quantum physics revitalizes CDM, which takes advantages of CDM's fast convergence property and asynchronous parallel scalability. Currently CDMs are still under fast development. New coordinate-pick strategies are proposed, and momentum acceleration is also combined with CDMs to further accelerate the convergence. CDMs are also tailored for different applications. For example, leading eigenvalue problem which is important in data science and quantum physics are tackled by CDMs recently. Doubly greedy primal dual coordinate descent algorithm is developed to solve sparse empirical risk minimization. Some asynchronous parallel block coordinate methods are used to solve convex optimization problems with non-separable linear constraint. CDMs are also used in image processing, signal processing and many more fields.

Organizer: Zhe Wang

Duke University, U.S.

Organizer: Yingzhou Li

Duke University, U.S.

4:10-4:30 Coordinate-Wise Descent Methods for Leading Eigenvalue Problems

Yingzhou Li, Jianfeng Lu, and *Zhe Wang*, Duke University, U.S.

continued in next column

continued on next page

4:35-4:55 Efficient Coordinate-wise Leading Eigenvector Computation
Jialei Wang, University of Chicago, U.S.; *Weiran Wang*, Amazon, U.S.; Dan Garber, Technion Israel Institute of Technology, Israel; Nathan Srebro, Toyota Technological Institute, U.S.

5:00-5:20 Primal-dual Block Coordinate Update Methods for Multi-block Structured Affinely Constrained Problems
Yangyang Xu, Rensselaer Polytechnic Institute, U.S.

5:25-5:45 Recent Advances in Primal-dual Coordinate Methods for ERM
Qi Lei, University of Texas at Austin, U.S.; Ian En-Hsu Yen, Carnegie Mellon University, U.S.; Chao-yuan Wu and Inderjit S. Dhillon, University of Texas at Austin, U.S.; Pradeep Ravikumar, Carnegie Mellon University, U.S.

Monday, February 25

MS73

Student Days: Student Chapter Presentations - Part II of II

4:10 p.m.-5:50 p.m.

Room: 102B

For Part 1 see MS4
Presentations given by students from SIAM Student Chapters.

See online program for an update on this session.

Monday, February 25

MS74

Transmission Eigenvalue Spectrum for Electromagnetic and Elastic Scatterers

4:10 p.m.-5:50 p.m.

Room: 102C

Electromagnetic and elastic interior transmission eigenvalues naturally appear in the theory of reconstructing electromagnetic and elastic scatterers from the knowledge of far-field data. Additionally, the spectrum itself might be useful to obtain qualitative information of the scatterer which is a desired task in non-destructive testing. But the numerical calculation of such interior transmission eigenvalues is very challenging due to the fact that the corresponding transmission problem is neither self-adjoint nor elliptic. The goal of this minisymposium is to highlight new developments and novel algorithm. It aims to bring together mathematicians and scientists working on these methods to exchange ideas and share new results.

Organizer: Andreas Kleefeld
Jülich Supercomputing Centre, Germany

Organizer: Lukas Pieronek
Jülich Supercomputing Centre, Germany

4:10-4:30 Method of Fundamental Solution for Computing Elastic Interior Transmission Eigenvalues
Lukas Pieronek, Jülich Supercomputing Centre, Germany

4:35-4:55 Transmission Eigenvalues for Materials with a Conductive Boundary Condition
Isaac Harris, Purdue University, U.S.

5:00-5:20 Determination of Electromagnetic Bloch Modes in a Medium with Frequency-dependent Coefficients
Shixu Meng, University of Michigan, U.S.

5:25-5:45 Scattering by Obstacles with Periodic Material Properties and the Effect of Boundary Correctors
Shari Moskow, Drexel University, U.S.

Monday, February 25

MS75

Unfitted Discretization Methods - Part I of II

4:10 p.m.-5:50 p.m.

Room: 102D

For Part 2 see MS107

The use of unfitted finite element (or isogeometric) techniques is of main interest for different reasons. On one hand, it allows one to eliminate the need to have body-fitted meshes, simplifying the pre-processing step, and having a closer interaction with the CAD data, e.g., in shape or topology optimization frameworks. On the other hand, extended finite element-like techniques are appealing to track interfaces in multi-phase and multi-physics applications. However, the use of unfitted finite element techniques have some drawbacks, e.g., the small cut cell problem and the resulting ill-conditioning of the problem at hand, the imposition of Dirichlet boundary conditions, etc. In this minisymposium, we are interested in works dealing with unfitted finite element simulations in a broad sense, including stabilization techniques for well-posedness, high-order frameworks, applications with advanced discretization spaces, robust and parallel preconditioning techniques, application to PDEs in manifolds and moving domains, surface extraction and cut cell integration techniques, adaptivity, etc., but also the use of these techniques in complex applications, e.g., additive manufacturing and bio applications, and software and implementation aspects.

Organizer: Guglielmo Scovazzi

Duke University, U.S.

Organizer: Xianyi Zeng

University of Texas at El Paso, U.S.

Organizer: Santiago Badia

Universitat Politecnica de Catalunya, Spain

4:10-4:30 Coupling Multi-material Flows and Structures using Embedded Methods on ALE Grids

Xianyi Zeng, University of Texas at El Paso, U.S.

4:35-4:55 Large Scale Unfitted Finite Elements for Fluid and Solid Mechanics Problems

Santiago Badia, Universitat Politecnica de Catalunya, Spain; Alberto F Martin and Francesc Verdugo, CIMNE, Spain

5:00-5:20 Unfitted FEM/DEM Methods for Modeling Variably Saturated Granular Media

Chris Kees, Yong Yang, and Manuel Quezada De Luna, U.S. Army Engineer Research and Development Center, U.S.

5:25-5:45 The Shifted Interface Method for Multiphase Flows Computations

Mehdi Khalloufi, Kangan Li, and Guglielmo Scovazzi, Duke University, U.S.

continued in next column

Monday, February 25

MS76

Multigrid Solvers for Partially Structured Meshes and for Advanced Architectures

4:10 p.m.-5:50 p.m.

Room: 111A

Multigrid methods are an important class of iterative solvers and preconditioners as they provide fast convergence rates with O(n) complexity on specific classes of problems such as elliptic problems with uniform coefficients. While their implementation on distributed memory system has been shown to scale well both weakly and strongly, the advent of new mixed architectures requires a redesign of some core kernels of these methods. This minisymposium aims at exploring methods leveraging fully or partially structured grids: fully structured grids, block-structured grids, mixed structured/unstructured grids or nested grids. Exploiting the regularity and the predictable indexing of the data leads to new more parallel and scalable algorithms. This minisymposium also focuses on the parallel shared memory implementations of these algorithms for mixed architectures as it is an increasingly important aspect of high performance computing on recent architectures aiming at scalability at exascale.

Organizer: Luc Berger-Vergiat

Sandia National Laboratories, U.S.

Organizer: Ray S. Tuminaro

Sandia National Laboratories, U.S.

4:10-4:30 An Algebraic Multigrid Method Tailored to Semi-structured Grids

Ray S. Tuminaro, Sandia National Laboratories, U.S.; Matthias Mayr, Universität der Bundeswehr München, Germany; Luc Berger-Vergiat, Sandia National Laboratories, U.S.; Peter Ohm, Tufts University, U.S.

continued on next page

4:35-4:55 Robust Structure-exploiting Multilevel Solvers

J. David Moulton, Los Alamos National Laboratory, U.S.; Andrew Reisner and Luke Olson, University of Illinois at Urbana-Champaign, U.S.

5:00-5:20 Black Box Multigrid for An Exact Penalty Formulation of Resistive MHD

Ari Rappaport, John Shadid, and Ray S. Tuminaro, Sandia National Laboratories, U.S.; Jehanzeb H. Chaudhry, University of New Mexico, U.S.; Luc Berger-Vergiat, Sandia National Laboratories, U.S.

5:25-5:45 HyTeG: A High Performance Multigrid Framework on Hybrid Meshes

Dominik Thoennes and Nils Kohl, University of Erlangen-Nuernberg, Germany; Daniel Drzisga, Technische Universität München, Germany; Dominik Bartuschat, University of Erlangen-Nuernberg, Germany; Ulrich J. Ruede, University of Erlangen-Nuremberg, Germany

Monday, February 25

MS77

Multirate Time Stepping Methods

4:10 p.m.-5:50 p.m.

Room: 111B

Time-stepping methods are numerical methods for the time evolution of ODEs, DAEs, and evolutionary PDEs. Many applications require specialized time-stepping methods in order to enhance efficiency and expose parallelism at the time stepping level. Multirate time stepping methods are integration methods that can accommodate different time steps for different solution components or different sub-domains so that the time step is locally adapted to satisfy stability and consistency requirements. The talks in this minisymposium will describe recent novel developments in multirate time-stepping methods and demonstrate their benefits in practical applications.

Organizer: Michael Guenther

Bergische Universität Wuppertal, Germany

4:10-4:30 Multirate Schemes for Differential-Algebraic Equations

Michael Guenther, Bergische Universität Wuppertal, Germany

4:35-4:55 Multi-rate Time Integration on Overset Meshes

Cory Mikida, University of Illinois at Urbana-Champaign, U.S.; Andreas Kloeckner, University of Illinois, U.S.; Dan Bodony, University of Illinois at Urbana-Champaign, U.S.

5:00-5:20 Multirate Parareal with Application in Electrical Engineering

Iryna Kulchytska-Ruchka, Idoia Garcia, and Sebastian Schoeps, Technische Universität Darmstadt, Germany

5:25-5:45 Coupling Multirate and Parallel-in-time Approaches

Andreas Bartel, Bergische Universität Wuppertal, Germany

Monday, February 25

MS78

Optimization with Coupled PDEs in Multiphysics Applications

4:10 p.m.-5:50 p.m.

Room: 111C

Featured Minisymposium

In the context of optimization for multiphysics problems, one has to deal with coupled PDEs as state constraints. Nowadays, several coupled PDE problems can be solved by the use of existing solvers provided by open source software. Some of these software packages already incorporate the coupling between several disciplines, such as the multiphysics suites SU2 or Kratos. However, the minisymposium discusses how to transfer coupled simulations to design and optimization tools. In particular, we discuss the setup of appropriate optimization methodologies making use of coupled adjoints.

Organizer: Nicolas R. Gauger

Technische Universität Kaiserslautern, Germany

4:10-4:30 Topology Optimization of Heat Exchanging Components and Networks

Maarten Blommaert, Martine Baelmans, Geert Buckinx, Wouter Dekeyser, Paul Lacko, Bart Peremans, and Yang Zhou, KU Leuven, Belgium

4:35-4:55 Coupled Discrete Adjoint Computations for CFD-related Multiphysics Problems

Ole Burghardt, Tim Albring, and Nicolas R. Gauger, Technische Universität Kaiserslautern, Germany

5:00-5:20 A Modified Search Direction Method with Weakly Imposed Karush-Kuhn-Tucker Conditions using Singular-value Decomposition for Large Constrained Optimization Problems

Long Chen, Technische Universität München, Germany

5:25-5:45 Numerical Optimization for Aeroacoustics

Beckett Zhou, Tim Albring, and Nicolas R. Gauger, Technische Universität Kaiserslautern, Germany

Monday, February 25

MS79

Recent Developments in Numerical Methods for PDEs and their Applications - Part I of II

4:10 p.m.-5:50 p.m.

Room: 300A

For Part 2 see MS233

Numerical Methods for partial differential equations and their analysis are important and challenging topics in applied and computational mathematics. This minisymposium is focused on recent developments in numerical methods for PDEs, including new developments in finite element methods and relevant applications. The goal of this minisymposium is to bring together leading researchers in the field of numerical methods to discuss and disseminate the latest results and envisage future challenges in traditional and new areas of science. The topics of the minisymposium cover a broad range of numerical methods, including but not limited to finite element methods, finite difference methods, discontinuous Galerkin methods, weak Galerkin methods. A wide range of application fields will also be covered, such as MHD equations, Helmholtz equations, equations of poroelasticity, Alfeld splits and gradient flows.

Organizer: Chunmei Wang

Texas Tech University, U.S.

4:10-4:30 Primal-dual Weak Galerkin Finite Element Methods

Junping Wang, National Science Foundation, U.S.

4:35-4:55 A New Spectral Method for Elliptic Problems in General Domains

Yiqi Gu and Jie Shen, Purdue University, U.S.

5:00-5:20 A Charge-conservative Mixed Finite Element Method for Inductionless MHD Equations

Weiying Zheng, Chinese Academy of Sciences, China

5:25-5:45 Exact Smooth Piecewise Polynomial Sequences on Alfeld Splits

Michael J. Neilan, University of Pittsburgh, U.S.; Johnny Guzman and Guosheng Fu, Brown University, U.S.

Monday, February 25

MS80

Applications of Tensor Decompositions

4:10 p.m.-5:50 p.m.

Room:300B

Tensors are a natural way to represent a multi-relational dataset. Given the data as a tensor, a tensor decomposition serves as a promising analytical tool for mining that data to uncover hidden structure among the data's relations. This minisymposium explores applications of tensor decompositions in data analytics, across areas spanning signal processing, cybersecurity, machine learning, and beyond. The talks will consider additional issues including algorithmic scalability, parallelism, and software stacks, among others.

Organizer: Richard Vuduc

Georgia Institute of Technology, U.S.

4:10-4:30 Applications of Tensor Decompositions at Amazon

Anima Anandkumar, Amazon and California Institute of Technology, U.S.

4:35-4:55 Tensor Methods and Recommender Systems

Ivan Oseledets, Skolkovo Institute of Science and Technology, Russia

5:00-5:20 Tensor Factorizations with Applications to Pattern and Topic Detection

Evangelos Papalexakis, University of California, Riverside, U.S.

5:25-5:45 Detecting Malware using Tensor Decompositions

Charles Nicholas, University of Maryland, Baltimore County, U.S.

Monday, February 25

MS81

A Posteriori Error Estimation for Various Adaptive Finite Element Methods

4:10 p.m.-5:50 p.m.

Room: 300C

A posterior error estimation plays a crucial role in the adaptive finite element methods (FEM). Regarding as one of the most effective approaches to solve the large scale and grand challenging problems, the adaptive numerical methods aim to obtain a solution within prescribed accuracy using (almost) the least amount of computational costs. This minisymposium will bring together researchers to discuss recent advances in developing and analyzing a posteriori error estimation for a variety of finite element methods, including classical FEM, least-squares FEM, virtual element method, and immersed finite element method (a FEM with unfitted meshes).

Organizer: Xu Zhang

Mississippi State University, U.S.

4:10-4:30 Adaptive Least-squares Finite Element Methods for Linear Transport Equations

Shun Zhang, City University of Hong Kong, Hong Kong

4:35-4:55 Hybrid a Posteriori Error Estimators for Finite Element Approximations

Difeng Cai, Purdue University, U.S.

5:00-5:20 A Posteriori Error Estimates for the Virtual Element Method

Shuhao Cao, University of California, Irvine, U.S.

5:25-5:45 Title Not Available

Xu Zhang, Mississippi State University, U.S.

Monday, February 25

MS82

BE: Lightning Talks

4:10 p.m.-5:50 p.m.

Room: 300D

Organizer: Mary Ann E. Leung
Sustainable Horizons Institute, U.S.

Monday, February 25

MS83

Adjoints of Nonlinear Dynamics in CFD Applications

4:10 p.m.-5:50 p.m.

Room: 302A

Large Scale Computational Fluid Dynamics (CFD) applications typically pivot around highly complex physical phenomena, are nonlinear and even chaotic in nature, and overall computationally expensive. Subsequently, simulations of fluid motion demand state of the art solutions from fields such as numerical analysis, high performance computing, mathematical modeling, to just name a few. A CFD simulation may be supplemented with adjoint capabilities in order to address many of the computational challenges, but also to aid in uncovering new physical insights, as well as to tune solver parameters. This addition, however, brings about an extra layer of difficulties, such as increased data intensity, high computational costs, poor convergence rates in adjoint-based optimization routines etc. The present minisymposium seeks to gather and disseminate knowledge on the challenges of adjoint solvers, and their resolutions, in the context of nonlinear fluid dynamics. The focus is in particular on efficient computational strategies in renowned CFD codes and libraries they build on, but also applications that benefit from adjoint implementations.

Organizer: Oana Marin
Argonne National Laboratory, U.S.

Organizer: Barry Smith
Argonne National Laboratory, U.S.

4:10-4:30 Large-scale PDE Constrained Optimization from the Perspective of a Library (PETSc)
Barry F. Smith, Oana Marin, and Emil M. Constantinescu, Argonne National Laboratory, U.S.

4:35-4:55 A Unsteady Discrete Adjoint Solver for RANS and Hybrid-RANS/LES
Nicolas R. Gauger, Max Sagebaum, and Tim Albring, Technische Universität Kaiserslautern, Germany

5:00-5:20 Obtaining Discrete Adjoint Sensitivities using Algorithmic Differentiation
Markus Towara, RWTH Aachen University, Germany

5:25-5:45 Adjoint-based Mesh Optimisation and Optimal Initial Conditions in a Spectral Element Framework
Philipp Schlatter and Nicolas Offermans, KTH Royal Institute of Technology, Sweden; Oana Marin, Argonne National Laboratory, U.S.

continued in next column

Monday, February 25

MS84

Advances and Applications of Numerical Methods for Free Boundaries

4:10 p.m.-5:50 p.m.

Room: 302B

This minisymposium will focus on novel numerical methods for free boundary value problems and their applications to the fields of materials science and fluid dynamics. Numerical methods that seek to impose sharp boundary conditions will be emphasized as well as methods on adaptive grids and strategies for parallel environment.

Organizer: Frederic Gibou

University of California, Santa Barbara, U.S.

4:10-4:30 Free Boundary Solvers on Octree Grids and in Parallel Environments

Frederic Gibou, University of California, Santa Barbara, U.S.

4:35-4:55 Pressure Based Compressible Solver for Two-phase Flows with Arbitrary Equation of State and Entropic Effects

Sebastien Tanguy, Universite de Toulouse, France

5:00-5:20 Parallel Octree Simulations of Incompressible Flows Coupled with Soluble Surfactant in a Level-set Framework

Fernando Temprano-Coleto, University of California, Santa Barbara, U.S.

5:25-5:45 A Multi-scale and Multi-physics Tissue Simulation Schema to Assess Emergent Rules of Complex Systems

Pouria Mistani, University of California, Santa Barbara, U.S.

Monday, February 25

MS85

Developments in Algebraic Multigrid for Nonsymmetric and Hyperbolic Problems

4:10 p.m.-5:50 p.m.

Room: 303A

Algebraic multigrid (AMG) is a relatively mature field of research in the applied mathematics community. However, despite several decades of research, highly nonsymmetric and hyperbolic-type problems remain difficult for AMG solvers, both in theory and in practice. With the continuing growth and evolution of parallel computer architectures, extra care must also be taken to develop algorithms that are amenable for highly parallel environments. This symposium highlights recent developments in the field of nonsymmetric AMG, including two unique theoretical frameworks that provide sufficient conditions for convergence-in-norm of nonsymmetric AMG, a new nonsymmetric AMG solver based on approximate ideal restriction (AIR), and the application of new nonsymmetric AMG methods to nonlinear hyperbolic PDEs.

Organizer: Ben Southworth

University of Colorado Boulder, U.S.

4:10-4:30 Nonsymmetric Algebraic Multigrid Based on Approximate Ideal Restriction (AIR)

Ben Southworth and Thomas Manteuffel, University of Colorado Boulder, U.S.; John Ruge, University of Colorado, U.S.

4:35-4:55 Parallel High-order Space-time AIR for Hyperbolic PDEs

Hans De Sterck and Oliver A. Krzysik, Monash University, Australia; Scott Maclachlan, Memorial University, Newfoundland, Canada; Ben Southworth, University of Colorado Boulder, U.S.

5:00-5:20 Algebraic Multigrid for Hypersonic Simulations

Luc Berger-Vergiat, Ray S. Tuminaro, Jonathan J. Hu, and Chris Siefert, Sandia National Laboratories, U.S.; Michael Brazell and Dimitri Mavriplis, University of Wyoming, U.S.

5:25-5:45 A Convergence Theory for Algebraic Multigrid (AMG) Applied to Nonsymmetric Systems

Thomas Manteuffel and Ben Southworth, University of Colorado Boulder, U.S.

continued in next column

Monday, February 25

MS86

Hardware-aware Algorithms and Numerics for Heterogeneous Supercomputers

4:10 p.m.-5:25 p.m.

Room: 303B

The advent of heterogeneous supercomputers - consisting of a blend of CPUs and GPUs - has promised significant gains in speedup and scalability over their homogeneous CPU-only counterparts. However, to effectively utilize them, heterogeneous architectures require a paradigm shift in how we develop numerical methods and algorithms: while in the past, physics was the main driver for algorithms development, today, both physics and hardware need to be accounted for. For example, hybrid computing requires numerical methods that increase arithmetic intensity and algorithms that reduce data transfers between hardware and across compute nodes. The proposed session aims at bringing together some of the most recent research in this area with specific applications in evolutionary differential equations such as those encountered in fluids, combustion, gas turbines, atmospheric flows, and biological transport. Topics will focus numerical methods and algorithms that aim at increasing arithmetic intensity and include, but not limited to, scalable linear solvers on GPUs, alternative techniques to global linear solvers, block-Implicit methods, the use of machine learning to solve linear and nonlinear systems of equations, asynchronous algorithms, high-order methods, and others.

Organizer: Tony Saad
University of Utah, U.S.

Organizer: James C. Sutherland
University of Utah, U.S.

4:10-4:30 Solving Evolutionary Differential Equations on Heterogeneous Architectures

Tony Saad and James C. Sutherland, University of Utah, U.S.

4:35-4:55 Software Library for Accelerating Chemical Kinetics on Hybrid Architectures: The SLACKHA Project

Kyle E. Niemeyer, Andrew Alferman, and Emily F. Klee, Oregon State University, U.S.; Nicholas Curtis and Chih-Jen Sung, University of Connecticut, U.S.

5:00-5:20 Towards a Performance Portable Multi-fluid Plasma Capability

Roger Pawlowski, Eric C. Cyr, Edward G. Phillips, Matthew Bettencourt, Eric Phipps, Christian Trott, Paul Lin, and John Shadid, Sandia National Laboratories, U.S.

Monday, February 25

MS87

Advances in Computational Methods for Data Assimilation - Part II of II

4:10 p.m.-5:50 p.m.

Room: 201B

For Part 1 see MS47

Data assimilation is a mathematical procedure that estimates the state of some hidden system based on observation data. It is one of the most important research topics in computational science and engineering and it has many practical applications. A main effort to solve the data assimilation problem is the simulation based computational methods, which is usually combined by forward simulations and inverse inferences. With the progress in uncertainty quantification and the advances in large scale computing, computational methods for data assimilation attract more and more attentions. In this minisymposium, we present recent research achievements in computational methods for data assimilation, and we explore both the forward simulations and the inverse inferences for the data assimilation problem.

Organizer: Feng Bao
Florida State University, U.S.

Organizer: Xuemin Tu
University of Kansas, U.S.

4:10-4:30 Model Reduction and Basis Adaptation for Uncertainty Quantification of Time-dependent PDEs

Panos Stinis, Pacific Northwest National Laboratory, U.S.; Ramakrishna Tipireddy, University of Southern California, U.S.; Alexandre M. Tartakovsky and Jing Li, Pacific Northwest National Laboratory, U.S.

4:35-4:55 Ensemble Kernel-based Learning for Analytic Continuation

Xuping Xie, Oak Ridge National Laboratory, U.S.

continued in next column

continued on next page

Monday, February 25

MS87

Advances in Computational Methods for Data Assimilation - Part II of II

continued

5:00-5:20 Ensemble Kalman Filter with Localization

Xin T. Tong, National University of Singapore, Singapore

5:25-5:45 The Commutation Error for Large Eddy Simulation Reduced Order Models (LES-ROMs)

Birgul Koc, Muhammad Mohebujjaman, Changhong Mou, and Traian Iliescu, Virginia Tech, U.S.

Monday, February 25

MS88

Performance Portability through Source-to-source Code Transformations

4:10 p.m.-5:50 p.m.

Room: 202B

While compilers are reasonably effective at optimizing codes, getting the highest performance often requires creating special versions of performance-critical codes which depend on multiple factors including the specific machine architecture, problem parameters, and interactions with other components of the application. In this minisymposia, we will discuss the use of source-to-source transformations to allow the developer to work at a higher level of abstraction and exploit state-of-the-art techniques to create and improve system- and problem-specific code to achieve high performance across a range of systems. The evaluation of such systems, the level of abstraction and the developer participation in the process are also important topics to be discussed.

Organizer: William D. Gropp

University of Illinois at Urbana-Champaign, U.S.

4:10-4:30 Designing for High Performance

Tze Meng, Carnegie Mellon University, U.S.

4:35-4:55 Can Autotuning Compilers become Mainstream?

Mary Hall, University of Utah, U.S.

5:00-5:20 Transformation-based Code Optimization for Finite Element Methods

Kaushik Kulkarni, University of Illinois at Urbana-Champaign, U.S.; Andreas Kloeckner, University of Illinois, U.S.

5:25-5:45 A DSL for Program Optimization

Thiago Teixeira, University of Illinois at Urbana-Champaign, U.S.

Monday, February 25

MS89

Recent Advances in Accurate and Verified Numerical Computations

4:10 p.m.-5:50 p.m.

Room: 202C

This minisymposium is devoted to accurate and verified numerical computations related to self-validating methods and computer-assisted proofs. Since verified numerical computations enable us to rigorously solving various problems including linear algebra, nonlinear equations, integration, differentiation, and so forth, in finite precision arithmetic, they have become increasingly important in wide range of science and engineering. The main objective of the minisymposium is to discuss several recent topics on accurate and verified numerical computations and related numerical methods.

Organizer: Takeshi Ogita

Tokyo Woman's Christian University, Japan

4:10-4:30 Iterative Refinement for Singular Value Decomposition

Takeshi Ogita, Tokyo Woman's Christian University, Japan; Kensuke Aishima, Hosei University, Japan

4:35-4:55 Reproducibility and Performance of the Feltor Code on Parallel Architectures

Roman Iakymchuk, KTH Royal Institute of Technology, Sweden; Matthias Wiesenberger, Technical University of Denmark, Denmark; Stef Graillat, University Pierre and Marie Curie (UPMC), France

5:00-5:20 Verified Algorithm for High Order Differentiation Based on Hyper-dual Numbers

Naoya Yamanaka, Meisei University, Japan

5:25-5:45 Thin QR Decomposition using LU Factors and its Refinement

Takeshi Terao and Katsuhisa Ozaki, Shibaura Institute of Technology, Japan; Takeshi Ogita, Tokyo Woman's Christian University, Japan

Monday, February 25

MS90

Spectral Deferred Correction Methods for Temporal Integration

4:10 p.m.-5:50 p.m.

Room: 205

Spectral deferred corrections methods are a class of iterative temporal integration methods for differential equations. One feature of SDC methods that is often exploited is that it is straightforward to build higher- order integrators that allow additive splitting of the equation. SDC methods are also the key component of the parallel in time algorithm PFASST. This minisymposium will highlight new ways in which SDC methods are being leveraged to construct novel time integration strategies for computationally challenging applications.

Organizer: Michael Minion

Lawrence Berkeley National Laboratory, U.S.

4:10-4:30 Improved Coupling of Hydrodynamics and Nuclear Burning in Astrophysical Flows using SDC

Michael Zingale, Stony Brook University, U.S.

4:35-4:55 Tracking Fast Ions in Fusion Reactors using Spectral Deferred Corrections

Krasymyr Tretiak, University of Leeds, United Kingdom; James Buchanan, Culham Centre For Fusion Energy, United Kingdom; Debasmita Samaddar, UK Atomic Energy Authority, United Kingdom; Rob Akers, Culham Centre For Fusion Energy, United Kingdom; *Daniel Ruprecht*, University of Leeds, United Kingdom

5:00-5:20 Parallel-in-time Integration of the Shallow-water Equations on the Rotating Sphere

Francois P. Hamon, Total Exploration & Production, U.S.; Martin Schreiber, University of Exeter, United Kingdom; Michael Minion, Lawrence Berkeley National Laboratory, U.S.

5:25-5:45 Adaptive Parallel-in-time PDE-constrained Optimization

Sebastian Götschel, Zuse Institute Berlin, Germany; Michael Minion, Lawrence Berkeley National Laboratory, U.S.

Monday, February 25

MS91

Large-scale PDE-Constrained Optimization Algorithms and Applications

4:10 p.m.-5:50 p.m.

Room: 206A

Numerical optimization is a powerful methodology for the control and/or design of many engineering systems governed by partial differential equations (PDEs). Many data assimilation problems can also be formulated in this framework, where the objective function and/or the constraints depend on the solution of PDEs. The computational cost of these PDE solutions, however, necessitate additional considerations in the optimization algorithms. This is especially true for problems that feature large numbers of variables, some of which may also be stochastic or integer-valued in nature. In this minisymposium, we will highlight ongoing research efforts in the efficient solution of such large-scale PDE-constrained optimization problems.

Organizer: Alp Dener

Argonne National Laboratory, U.S.

Organizer: Todd Munson

Argonne National Laboratory, U.S.

4:10-4:30 Accelerating Limited-memory Quasi-Newton Convergence for Optimization of PDE-governed Systems

Alp Dener and Todd Munson, Argonne National Laboratory, U.S.

4:35-4:55 An Acceleration Framework for Parameter Estimation using Implicit Sampling and Adaptive Reduced-order Models

Robert Baraldi, North Carolina State University, U.S.; Matthew J. Zahr, University of California, Berkeley and Lawrence Berkeley National Laboratory, U.S.

5:00-5:20 Fast Solvers for Mixed-integer PDE-constrained Optimization

Sven Leyffer, Argonne National Laboratory, U.S.

5:25-5:45 Scalable Optimization of Discrete-time PDE Systems with Inequality Constraints

Denis Ridzal and Eric C. Cyr, Sandia National Laboratories, U.S.

Monday, February 25

MS92

Latest Advances in Viscoplastic CFD: Discretisation, Simulation and Applications

4:10 p.m.-5:50 p.m.

Room: 206B

When a nondifferentiable inclusion problem in function spaces arises jointly with the nonlinearity and differential-algebraic structure of the incompressible Navier-Stokes equations, then the resulting challenges continue to be formidable. This is the case for flow problems of viscoplastic fluids, gel-like materials with a stress-dependent behaviour that switches between viscous flow and plastic creep. As of today, simulations of viscoplastic flows in 3D or realistically complex 2D configurations remain largely infeasible. Thus, state-of-the-art viscoplastic solvers perform comparably to Newtonian solvers of the past century. The invited speakers contribute new strategies for efficiency gains and highly accurate solutions, as required in important applications. They discuss the key ingredients of viscoplastic solvers: (i) properties of discretisation schemes (accuracy, stability, robustness, discrete conservation laws, sparsity, amenability to multigrid solvers, mesh-independent convergence), (ii) algorithmic treatment of the nonsmoothness (regularisation techniques vs genuinely nonsmooth methods) and (iii) convergence acceleration (preconditioners, Newton-type methods, inertial algorithms, parallelisation). Speakers working on applications address the resolution of problem-inherent features at very different length scales: boundary layers, internal layers up to the extreme case of a lower-dimensional manifold: the yield surface that separates subdomains of flow and creep.

continued on next page

Monday, February 25

MS92

Latest Advances in Viscoplastic CFD: Discretisation, Simulation and Applications

continued

Organizer: Timm Treskatis
University of British Columbia, Canada

4:10-4:30 Generalized Newton Methods for the Numerical Simulation of Viscoplastic Fluids
Sergio González-Andrade, Escuela Politécnica Nacional del Ecuador, Ecuador

4:35-4:55 A Multigrid Optimization Algorithm for the Numerical Solution of Classical Viscoplastic Models
Sofía López-Ordóñez and Sergio González-Andrade, Escuela Politécnica Nacional del Ecuador, Ecuador

5:00-5:20 The Viscoplastic Navier-Stokes Equations: Fast Defect Correction, Mimetic Discretisation and Applications in Bioprinting
Timm Treskatis, University of British Columbia, Canada

5:25-5:45 Flow of a Yield-stress Fluid over a Cavity: Experimental and Numerical Investigation of an Oldroyd's Boundary Layer
Paul Vigneaux, UMPA - CNRS UMR 5669 ENS de LYON, France

Monday, February 25

MS93

Nonlinear Reduced Order Modeling of Realistic Engineering Fluid Flows

4:10 p.m.-5:50 p.m.

Room: 206C

This minisymposium presents important applications of nonlinear reduced order modeling in realistic fluid flows. The speakers will present their work on reduced order modeling of engineering fluid problems, such as wind turbine wake, fluid structure interaction, parametrized compressible flow simulations, and numerical analysis. Data-driven reduced order modeling of fluids plays an important role in those applications. Advantages of their model and potential future research directions will be discussed in the minisymposium.

Organizer: Xuping Xie
Oak Ridge National Laboratory, U.S.

Organizer: Mithu Debnath
National Renewable Energy Laboratory, U.S.

4:10-4:30 Modulation in Wind Turbine Wakes Identified with Modal Decomposition Techniques
Mithu Debnath, National Renewable Energy Laboratory, U.S.

4:35-4:55 Structure-preserving Galerkin POD Reduced-order Modeling of Hamiltonian Systems
Zhu Wang, University of South Carolina, U.S.

5:00-5:20 Domain Decomposition Least-Squares Petrov--Galerkin for Nonlinear Model Reduction
Chi K. Hoang and Kevin T. Carlberg, Sandia National Laboratories, U.S.

5:25-5:45 Reduced Order Modeling of a Fluid Structure Interaction Problem under Embedded
Lei Lei and Charbel Farhat, Stanford University, U.S.

Monday, February 25

MS94

Asynchronous Iterative Methods

4:10 p.m.-5:50 p.m.

Room: 206D

Classical synchronous iterative methods alternate between local computation and boundary data exchange. In asynchronous iterative methods this dependency is relaxed and processing units are allowed to use whatever data is available at the beginning of a computation phase. Originally called `Chaotic Relaxation' for fixed-point iterations, asynchronous iterative methods are used in various areas of high-performance computing and numerical optimization. In this minisymposium, recent research is presented both on the theory and implementation of asynchronous iterative methods.

Organizer: Christian Glusa
Sandia National Laboratories, U.S.

Organizer: Daniel B. Szyld
Temple University, U.S.

4:10-4:30 Asynchronous Optimized Schwarz Methods for the Solution of PDEs on Bounded Domains
Daniel B. Szyld, Temple University, U.S.; Federeic Magoules, CentraleSupélec, France; Jose Garay, Louisiana State University, U.S.; Mireille El Haddad, Saint Joseph University, Lebanon

4:35-4:55 Asynchronous Subgradient-push
Mike Rabbat, Facebook Research, U.S.

5:00-5:20 An Asynchronous, Decentralized Solution Framework for the Large Scale Unit Commitment Problem
Paritosh P. Ramanan, Georgia Institute of Technology, U.S.

5:25-5:45 Asynchronous Two-level Domain Decomposition Solvers
Christian Glusa, Sandia National Laboratories, U.S.

Monday, February 25

MS95

Performance Portability and Numerical Libraries: Challenges and Opportunities for Sustainable Science - Part I of II

4:10 p.m.-5:25 p.m.

Room: 207

For Part 2 see MS129
Featured Minisymposium

A software component is said to be performance portable if it performs well across a range of architectures and problem configurations with modest development and maintenance effort. There are many programming tools that strive to assist with portability, but many challenges remain, especially balancing vectorization with memory locality across disparate architectures and managing composition across module boundaries where developers must eschew assumptions about environment and context. Optimized numerical libraries provide an opportunity to encapsulate any necessary specialization in a form that can be leveraged by many applications, but innovative interfaces and implementations are needed to serve diverse user requirements and to control complexity when porting and optimizing for an ever-growing range of architectures. This session brings together exemplars, promising techniques, and lessons learned in developing robust performance-portable applications and numerical libraries for computational science and engineering

Organizer: Jed Brown
University of Colorado Boulder, U.S.

Organizer: Rebecca J. Hartman-Baker
Lawrence Berkeley National Laboratory, U.S.

4:10-4:30 Performance Optimisation of Finite Element Assembly in Firedrake
Tianjiao Sun, Imperial College London, United Kingdom; Lawrence Mitchell, Durham University, United Kingdom; David Ham and Paul Kelly, Imperial College London, United Kingdom

4:35-4:55 Kokkos: Insights from 4 Years of Application Adoption Work
Christian Trott, Sandia National Laboratories, U.S.

5:00-5:20 Chapel's Language-based Approach to Performance Portability
Bradford L. Chamberlain, Cray, Inc., U.S.

continued in next column

Monday, February 25

MS96

Computationally Efficient Derivative-Free Local and Global Optimization

4:10 p.m.-5:50 p.m.

Room: 401A

There are many important nonlinear optimization problems for which derivatives are not available or they are too expensive to compute numerically. This minisymposium will discuss recent advances in computationally efficient derivative-free algorithms that can solve these problems. This includes global optimization methods that find the best of multiple local minima using a global surrogate to reduce the number of evaluations of an expensive objective function required to find an accurate solution. Trust region methods are used for local optimization and in one of the papers as a part of global search. Surrogates, including trust region approximation, can reduce the number of function evaluations to obtain an accurate solution and are hence especially important for expensive objective functions such as those requiring solution of a simulation model. The minisymposium includes the discussion of sequential and parallel algorithms that are based on decomposition in order to expand the scope of derivative-free optimization to larger-scale problems than what was possible in the past. Finally, the discussion includes some real-world applications, including tuning of machine learning models and design of nanoparticle synthesis.

Organizer: Nikolaos Sahinidis
Carnegie Mellon University, U.S.

Organizer: Christine Shoemaker
National University of Singapore, Singapore

continued on next page

Monday, February 25

MS96

Computationally Efficient Derivative-Free Local and Global Optimization

continued

4:10-4:30 Improving Surrogate Global Optimization Accuracy with RBF Trust Regions

Christine Shoemaker, Limeng Liu, and Jin Yi, National University of Singapore, Singapore

4:35-4:55 Decomposition in Global Derivative-free Optimization

Kaiwen Ma and Nikolaos Sahinidis, Carnegie Mellon University, U.S.; Satyajith Amaran and Scott Bury, The Dow Chemical Company, U.S.

5:00-5:20 Targeted Design of Nanoparticle Synthesis using Derivative-free Optimization

Aekaansh Verma, Stanford University, U.S.; Stefan Wild, Argonne National Laboratory, U.S.

5:25-5:45 Tuning Machine Learning Models with Derivative-free Optimization

Steven Gardner, Oleg Golovidov, Joshua Griffin, Patrick Koch, and Scott Pope, SAS Institute, Inc., U.S.

Monday, February 25

MS97

Computational Challenges in the ECP Energy Applications

4:10 p.m.-5:50 p.m.

Room: 401B

The Exascale Computing Project (ECP) is aimed at ensuring that science applications are ready and able to fully take advantage of the unique challenges associated with next-generation computing architectures. This session explores recent developments within the Energy portfolio of the Applications Development focus area of ECP. These talks will describe new approaches for enabling simulation of complex physical systems using large-scale computing environments. These projects are supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of two U.S. Department of Energy organizations (Office of Science and the National Nuclear Security Administration).

Organizer: Steven Hamilton

Oak Ridge National Laboratory, U.S.

4:10-4:30 Coupled Monte Carlo Radiation Transport and CFD for Nuclear Reactor Modeling

Steven Hamilton and Thomas Evans, Oak Ridge National Laboratory, U.S.

4:35-4:55 Novel Computational Methods for High-fidelity Modeling of Plasma-based Particle Accelerators

Olga Shapoval, Jean-Luc Vay, Rémi Lehe, and Maxence Thevenet, Lawrence Berkeley National Laboratory, U.S.; Henri Vincenti, CEA, France

5:00-5:20 Multiphase Flow Modeling with MFIX-Exa

Michele Rosso, Lawrence Berkeley National Laboratory, U.S.

5:25-5:45 ExaWind: Multi-scale Challenges of Massively Parallel Blade-resolved Wind Turbine Simulations

Michael Sprague, Shreyas Ananthan, Stephen Thomas, Michael Lawson, and Katarzyna Swirydowicz, National Renewable Energy Laboratory, U.S.

Monday, February 25

MS98

Computational Advances for Large-scale Geophysical Data

4:10 p.m.-5:50 p.m.

Room: 401C

New geophysical acquisition technologies are enabling the widespread collection of continuous streams of data from an unprecedented number of sensors. Often, the same dataset may be used for multiple purposes: event detection, source inversion, subsurface imaging, and more. To fully exploit these data, new algorithms must be developed that scale both in terms of floating point operations and memory access. Recent algorithmic advances have included robust automated data pre-processing, fast data transforms and solvers, improved preconditioners for iterative methods, and new applications of artificial intelligence to improve efficiency. In this minisymposium, we welcome submissions related to algorithmic advances that improve the scalability of data analysis with respect to longer times (large-T) or many sensors (large-N), or that operate in a streaming paradigm (accessing data a small number of times) or a cloud computing environment. This minisymposium will include studies using open source data and novel data sources (e.g. smartphone/MEMS accelerometers, fiber optics, wireless nodes) and we encourage the presentation of new applications.

Organizer: Eileen R. Martin

Virginia Tech, U.S.

4:10-4:30 Fast Radon Transforms for Ambient Noise Interferometry

Eileen R. Martin, Virginia Tech, U.S.; Biondo Biondi, Stanford University, U.S.; Jonathan Ajo-Franklin, Lawrence Berkeley National Laboratory, U.S.

4:35-4:55 MyShake - Dense Global Seismic Network using Smartphones

Qingkai Kong and Richard Allen, University of California, Berkeley, U.S.

continued on next page

5:00-5:20 Wireless Sensing for the Energy Industry over the Internet of Things

Hadi Jamali-Rad, Shell International Exploration & Production B.V., Netherlands; Xander Campman, Shell Global Solutions International B.V., Rijswijk, Netherlands

5:25-5:45 Compressive Least-squares Migration with on-the-fly Fourier Transforms

Philipp A. Witte, Georgia Institute of Technology, U.S.; Mathias Louboutin, University of British Columbia, Canada; Fabio Luporini and Gerard J Gorman, Imperial College London, United Kingdom; Felix Herrmann, Georgia Institute of Technology, U.S.

Monday, February 25

MS99

Distributed-Memory Graph Analytics: Programming Models, Algorithms and Applications

4:10 p.m.-5:50 p.m.

Room: 402A

Graph analytics is central to many applications in computational science and engineering (CSE) and data science. This minisymposium will feature talks addressing recent developments in: programing models in support of distributed-memory graph analytics; design, implementation and performance evaluation of scalable distributed-memory graph algorithms; and applications in bioinformatics. Featured graph kernels include community identification and graph traversals.

Organizer: Assefaw Gebremedhin

Washington State University, U.S.

4:10-4:30 Challenges in MPI for Supporting Distributed-memory Graph Analytics

Pavan Balaji, Argonne National Laboratory, U.S.

4:35-4:55 Kronecker Graphs as Benchmarks for Distributed Graph Analytics

Geoff Sanders, Lawrence Livermore National Laboratory, U.S.; Trevor Steil, University of Minnesota, U.S.; Timothy La Fond, Roger Pearce, and Keita Iwabuchi, Lawrence Livermore National Laboratory, U.S.; Benjamin Priest, Dartmouth College, U.S.

5:00-5:20 Distributed-memory Graph Algorithms for Bioinformatics

Kamesh Madduri, Pennsylvania State University, U.S.

5:25-5:45 Scalable Graph Community Detection using the Louvain Method

Sayan Ghosh, Washington State University, U.S.; Mahantesh Halappanavar and Antonino Tumeo, Pacific Northwest National Laboratory, U.S.; Ananth Kalyanaraman and Assefaw Gebremedhin, Washington State University, U.S.

Monday, February 25

MS100

Scalable Adaptive Applications: Recent Developments

4:10 p.m.-5:50 p.m.

Room: 402B

We present selected applications of scalable adaptive mesh refinement (AMR). The focus is on new algorithms for important features and functionalities and on latest science applications that rely on such features. Even though several algorithmic approaches to AMR are widely and successfully used for the numerical solution of partial differential equations, it remains challenging to incorporate complex physics, multiple scales, and non-element-based methods, as well as other sources of non-locality and non-uniformity. In this minisymposion, we highlight how such generalizations impact performance and scalability. We discuss the particular requirements on algorithm and application design at the largest scale, how low- and high-level simulation layers depend on each other, and how synergy can be maximized.

Organizer: Carsten Burstedde

Universität Bonn, Germany

Organizer: Martin Berzins

University of Utah, U.S.

4:10-4:30 Non-local Parallel Algorithms for Data Assignment and Analysis

Carsten Burstedde, Universität Bonn, Germany

4:35-4:55 Using AMR for Scaling Globally Connected Radiation Calculations

Martin Berzins and Alan Humphrey, University of Utah, U.S.

5:00-5:20 ForestClaw: An Adaptive, Finite Volume, Cartesian Grid Solver Based on p4est

Donna Calhoun, Boise State University, U.S.; Carsten Burstedde, Universität Bonn, Germany

continued on next page

Monday, February 25

MS100

Scalable Adaptive Applications: Recent Developments

continued

5:25-5:45 Parallel Variational Transfer for Scalable Multiphysics Simulations

Marco Favino, Patrick Zulian, Rolf Krause, and *Maria Nestola*, Università della Svizzera italiana, Switzerland

Monday, February 25

MS101

Scientific Data Visualization Platforms Facilitating New Paradigms

4:10 p.m.-5:50 p.m.

Room:402C

Featured Minisymposium

The influx of digital data, unprecedented high-performance computing (HPC) resources, and advances in visualization algorithms are driving a paradigm shift in how scientists work with information in the modern computational landscape. This session brings together application scientists and developers leveraging the rapidly changing technological advances to facilitate scientific discovery. Examples include data analytics, web-based interfaces for working with and analyzing data, larger infrastructure CAVES, and portable VR platforms such as 3D TVs or the HTC VIVE. For the successful communication of research done with these emerging technologies, methods are needed to convey the insights and analysis gained without requiring the audience to replicate the technological environment. Movie recordings panning through 3D data and simulation output can be one way to convey the depth and 3D complexity, as can well-designed web-based applications that allow the user to parse data remotely. For facilitating scientific reproducibility in the era of HPC and big data, university sponsored streaming and storage sites, such as institutional repositories may provide a critical link in moving this new technology forward and redefining approaches to scientific research.

Organizer: Margarete Jadamec

State University of New York at Buffalo, U.S.

4:10-4:30 Using 3D Virtual Reality for Interactively Exploring Models and Big Data

Margarete Jadamec, State University of New York at Buffalo, U.S.; Oliver Kreylos and Burak Yikilmaz, University of California, Davis, U.S.; Karen Fischer, Brown University, U.S.; Benjamin Chang, University of Houston, U.S.

4:35-4:55 3D Visual Tour of Earth's Interior Based on Global Adjoint Tomography

Ebru Bozdag, Colorado School of Mines, U.S.; David Pugmire, Oak Ridge National Laboratory, U.S.; Cagdas Demirkan, Sebnem Duzgun, and Ergin Isleyen, Colorado School of Mines, U.S.; Wenjie Lei, Princeton University, U.S.; Ridvan Orsvuran, University of Cote d'Azur, France; Youyi Ruan, Nanjing University, China; Jeroen Tromp, Princeton University, U.S.

5:00-5:20 Creating Visualizations for Science Communication

Maria Weber, The University of Chicago, U.S.

5:25-5:45 Opening Your Data: Libraries and Repositories Help You Go Public

Karlen Chase, State University of New York at Buffalo, U.S.

continued in next column

Monday, February 25

CP5

Data Science Applications

4:10 p.m.-5:30 p.m.

Room: 201A

Chair: Joanna Slawinska, University of Wisconsin, Milwaukee, U.S.

4:10-4:25 A Graph-based Framework for Load Flow Analysis of Multi-carrier Energy Systems

Anne S. Markensteijn and Kees Vuik, Delft University of Technology, Netherlands

4:30-4:45 Basin-preserving Lossy Compression for Optimization of a Molecular Dynamics System

Michael A. Retzlaff, University of Maryland, Baltimore County, U.S.; Zichao (Wendy) Di, Argonne National Laboratory, U.S.

4:50-5:05 Data-driven Koopman Analysis of Tropical Climate Space-time Variability

Joanna Slawinska, University of Wisconsin, Milwaukee, U.S.; Dimitrios Giannakis, Courant Institute of Mathematical Sciences, New York University, U.S.

5:10-5:25 Study Type B Aortic Dissection using a Deconvolution-based Nonlinear Filter

Huijuan Xu, Georgia Institute of Technology, U.S.; Alessandro Veneziani, Emory University, U.S.; Davide Baroli, University of Luxembourg, Luxembourg

Monday, February 25

CP6

Computational Fluid Mechanics II

4:10 p.m.-5:50 p.m.

Room: 201C

Chair: Alexander Litvinenko, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

4:10-4:25 Regularized Single and Double Layer Integrals in 3D Stokes Flow

Svetlana Tlupova, Farmingdale State College, U.S.; J. Thomas Beale, Duke University, U.S.

4:30-4:45 Chaotic Dynamics and Self-organization in a Depth-averaged Model of Turbidity Currents

Ramanathan Vishnampet, University of Illinois at Urbana-Champaign, U.S.; Mrugesh Shringarpure, Huafei Sun, Amit Kumar, and David Hoyal, ExxonMobil Upstream Research Company, U.S.

4:50-5:05 A Single Fluid Flow Poroelastic Model for a Fractured Porous Medium with Transverse Isotropy

Pedro Aguilar-Gastelum, Martín Díaz-Viera, Manuel Coronado, Edscott Wilson-García, and Mario Vadillo-Sáenz, Instituto Mexicano del Petróleo, México

5:10-5:25 Efficient Simulations for Contamination of Groundwater Aquifers under Uncertainties

Alexander Litvinenko, Dmitry Logashenko, David E. Keyes, Gabriel Wittum, and Raul F. Tempone, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

5:30-5:45 Large-scale Stochastic Particle Systems with Hydrodynamic Interactions

Gaddiel Ouaknin, Stanford University, U.S.; Yu Su, Cornell University, U.S.; Roseanna Zia, Stanford University, U.S.

Monday, February 25

PD3

Strategies for Promoting Diversity and Inclusion within our Profession

6:00 p.m.-7:15 p.m.

Room: Conference Theater

Moderator: Lisa Fauci, Tulane University, U.S.

Chair: Tamara G. Kolda, Sandia National Laboratories, U.S.

Chair: Helen Moore, AstraZeneca, U.S.

Chair: Karen E. Willcox, University of Texas at Austin, U.S.

As a professional society, SIAM supports its members in taking on the biggest challenges in the field. This panel is focused on one such challenge: inadequate representation of under-represented groups such as women and minorities. Our panelists have carefully studied the issue from different viewpoints and implemented various solutions. Their remarks will be focused on concrete actions for improving our profession in this regard. Join the discussion as we address questions such as:

- What are the biggest obstacles that our field faces in achieving better diversity and inclusion?
- What is the most successful diversity and inclusion strategy/action you have personally led or experienced?
- What have you tried and found not to work?
- What are your three top recommendations for actions that departments / individuals can take?

The panel moderator is SIAM President, *Lisa Fauci, Tulane University, U.S.*

Carlos Castillo-Chavez
Arizona State University, U.S.

Rachel Kuske
University of British Columbia, Canada

William Massey
Princeton University, U.S.

Eve Riskin
University of Washington, U.S.

SINUM Editorial Board Meeting

7:30 p.m.-9:30 p.m.

Room: Davenport Hotel -- Meeting Room 10

Tuesday, February 26

Math in Industry Book Committee Meeting

7:00 a.m.-8:15 a.m.

Room: Davenport Hotel -- Meeting Room 11

Student Days: Chapter Breakfast with SIAM Leadership (by invitation)

7:00 a.m.-8:15 a.m.

Room: Davenport Hotel -- Cedar Ballroom

Registration

7:30 a.m.-4:00 p.m.

Room: Ballroom Foyer

IP3

Reduced Order Methods for PDEs: State of the Art and Perspectives with Applications in Industry, Medicine and Environmental Sciences

8:30 a.m.-9:15 a.m.

Room: Ballroom 100BC

Chair: Luke Olson, University of Illinois at Urbana-Champaign, U.S.

We provide the state of the art of Reduced Order Methods (ROM) for parametric Partial Differential Equations (PDEs), and we focus on some perspectives in their current trends and developments, with a special interest in parametric problems arising in Computational Fluid Dynamics (CFD). Efficient parametrisations (random inputs, geometry, physics) are very important to be able to properly address an offline-online decoupling of the computational procedures and to allow competitive performances. Current ROM developments in CFD include: a better use of stable high fidelity methods, considering also spectral element method and finite volume discretisations; efficient sampling techniques to reduce the number of the basis functions, retained as snapshots, as well as the dimension of online systems; the improvements of the certification of accuracy based on residual based error bounds and of the stability factors, as well as the the guarantee of the stability of the approximation. For nonlinear systems, also the investigation on bifurcations of parametric solutions are crucial and they may be obtained by a reduced eigenvalue analysis of the linearised operator. All the previous aspects are crucial to analyse in real time complex parametric industrial, environmental and biomedical flow problems. We will focus on few benchmarks, including fluid-structure interaction problems, flow control and shape optimization.

Gianluigi Rozza
SISSA, International School for Advanced Studies, Trieste, Italy

Coffee Break

9:15 a.m.-9:45 a.m.

Room: Ballroom Foyer

continued in next column

Tuesday, February 26

MT1

Introduction To PETSc

9:45 a.m.-11:25 a.m.

Room: 300D

This short tutorial will introduce the PETSc libraries for the parallel solution of nonlinear algebraic equations. We will cover linear and nonlinear algebraic solvers and timestepping. It will focus primarily on optimal multilevel solvers for PDE problems. Test problem construction will leverage parallel mesh manipulation capabilities in PETSc and the p4est package. We will also help participants with design choices for their particular research codes.

Jed Brown, University of Colorado Boulder, U.S.

Tobin Isaac, Georgia Institute of Technology, U.S.

Matthew G. Knepley, State University of New York at Buffalo, U.S.

Tuesday, February 26

MS102

Scalable Optimization: Paradigms, Tools, and Applications - Part I of II

9:45 a.m.-11:25 a.m.

Room: Ballroom 100BC

For Part 2 see MS136

This minisymposium seeks to bring together experts in different optimization domains (PDEs, DAEs, stochastic, global, data/learning, mixed-integer) to discuss computational scalability issues as well as strategies and tools to overcome them.

Organizer: Victor M. Zavala

University of Wisconsin, Madison, U.S.

Organizer: Mihai Anitescu

University of Chicago, U.S.

Organizer: Kibaek Kim

Argonne National Laboratory, U.S.

9:45-10:05 Using Machine Learning to Understand and Solve Power System Optimization Problems

Line A. Roald, University of Wisconsin, Madison, U.S.; Sidhant Misra, Los Alamos National Laboratory, U.S.; Yee Sian Ng, Massachusetts Institute of Technology, U.S.

10:10-10:30 Bilevel Mixed-binary Programming for Identifying Critical Contingency Events in AC Power Systems

Kibaek Kim, Brian C. Dandurand, and Sven Leyffer, Argonne National Laboratory, U.S.

10:35-10:55 A Scalable Global Optimization Algorithm for Stochastic Nonlinear Programs

Yankai Cao, University of British Columbia, Canada; Victor Zavala, University of Wisconsin, Madison, U.S.

11:00-11:20 A Framework for Scalable Multiperiod Security Constrained AC-based Optimal Power Flow

Michel Schanen, Argonne National Laboratory, U.S.; Cosmin G. Petra, Lawrence Livermore National Laboratory, U.S.; Mihai Anitescu and Francois Gilbert, Argonne National Laboratory, U.S.

Tuesday, February 26

MS103

Recent Advances in Multirate Time Integration for Multiphysics Problems

9:45 a.m.-11:25 a.m.

Room:Conference Theater

Many problems arising in computational science and engineering involve multiple physical processes, where complex interactions between these components can result in dynamics that evolve on time scales differing by several orders of magnitude. This poses significant challenges for accurate and efficient numerical simulations that attempt to capture the dynamic behavior of each component. In recent years, multirate time integration techniques, which use multiple time steps tuned to these disparate physical processes, have shown advantages in multiphysics simulations. The main aim of this minisymposium is to present some of the latest developments in multirate time integration; however a secondary goal is to bring together experts and young researchers working in this field, with practitioners who make use of such time integration techniques in science and engineering applications.

Organizer: Vu Thai Luan

Southern Methodist University, U.S.

Organizer: Daniel R. Reynolds

Southern Methodist University, U.S.

9:45-10:05 Efficient Multirate Methods from High Order

Jean Sexton, Lawrence Berkeley National Laboratory, U.S.; Daniel R. Reynolds, Southern Methodist University, U.S.

10:10-10:30 Coupling MOR and Multirate Techniques in Multiphysics Problems

Michael Günther, Bergische Universität Wuppertal, Germany

10:35-10:55 Design of High-order Multirate General Additive Runge Kutta Schemes

Arash Sarshar, Steven Roberts, and Adrian Sandu, Virginia Tech, U.S.

continued on next page

Tuesday, February 26

MS103

Recent Advances in Multirate Time Integration for Multiphysics Problems

continued

11:00-11:20 A New Class of High-order, Flexible, IMEX Multirate Integrators for Multiphysics Applications

Rujeko Chinomona, Vu Thai Luan, and Daniel R. Reynolds, Southern Methodist University, U.S.

Tuesday, February 26

MS104

Accelerating Electronic Structure Calculations with GPUs - Part I of II

9:45 a.m.-11:25 a.m.

Room: 102A

For Part 2 see MS138

Several electronic structure codes have been ported to GPUs in recent years. However, with recent hardware developments, particularly the architecture of some of the new top supercomputers in the world, taking advantage of GPU accelerators is becoming a more pressing need in this field. In addition, since First-Principles (quantum-based) molecular dynamics can be computationally quite expensive, being able to reduce time-to-solution can significantly enhance the accuracy of many simulations. Additionally larger problems and longer timescales can be explored. In this minisymposium, we will discuss how electronic structure codes, targeting various applications in materials sciences, chemistry and biology, are being adapted and ported to GPU-based architectures. The speakers will cover a broad range of algorithms, models, languages and implementations used today in a variety of codes, and will discuss how these developments are helping to push the computational limits in real applications.

Organizer: Jean-Luc Fattebert
Oak Ridge National Laboratory, U.S.

Organizer: Susan Mniszewski
Los Alamos National Laboratory, U.S.

9:45-10:05 Accelerating Electronic Structure Calculations using GPUs: From Hand-coded Loops to Relying on Third-party Libraries

Jean-Luc Fattebert, Oak Ridge National Laboratory, U.S.; Susan Mniszewski, Los Alamos National Laboratory, U.S.

10:10-10:30 Experiences Porting Quantum Espresso's PWscf Solver to GPUs with CUDA Fortran

Josh Romero, NVIDIA, U.S.

10:35-10:55 Large Scale Hybrid Density Functional Theory Calculations with a Systematic Basis Set using GPUs

Laura Radcliff, Imperial College London, United Kingdom; Augustin Degomme, CEA Grenoble; José Flores-Livas and Stefan Goedecker, University of Basel, Switzerland; Luigi Genovese, CEA, France

11:00-11:20 Accelerating Ab Initio Quantum Monte Carlo Calculations

Paul Kent, Oak Ridge National Laboratory, U.S.

continued in next column

Tuesday, February 26

MS105

Computational Challenges at the Frontiers of Numerical Relativity - Part I of II

9:45 a.m.-11:25 a.m.

Room: 102B

For Part 2 see MS139

In LIGO's first year of observing time, it has detected gravitational waves from five binary black hole mergers and detected the first neutron star binary inspiral. The analysis and interpretation of LIGO data requires careful comparison to model solutions of the gravitational two-body problem. Some black hole binaries, however, pose significant challenges to both analytic approximation techniques and numerical simulations. The theme for our minisymposium proposal is the computational challenges in numerical relativity and gravity research as we enter the era of regular gravitational wave detections and multi-messenger astronomy. This will include discussion of numerical relativity and gravitational wave searches, algorithms for new physics, and new algorithms for modern supercomputer architectures. The speakers represent computational scientists, physicists and data analysis researchers who are working on the next generation of codes for numerical relativity. This talks will be of broad interest and naturally combines expertise and knowledge that reaches beyond physics to computational and computer science.

Organizer: Hari Sundar
University of Utah, U.S.

Organizer: David Neilsen
Brigham Young University, U.S.

9:45-10:05 Modelling Relativity in Software: We Need More Mathematical Rigour in our Codes for Both Clarity and Performance
Erik Schnetter, Perimeter Institute, Canada

10:10-10:30 Senr/NRPy+: Numerical Relativity in Singular Curvilinear Coordinate Systems
Zach Etienne, West Virginia University, U.S.

10:35-10:55 SpECTRE: A Task-based Discontinuous Galerkin Code for Relativistic Astrophysics
Lawrence Kidder, Cornell University, U.S.

11:00-11:20 Numerical Relativity and the Challenge of Turbulent Post-mergers
Matthew Duez, Washington State University, U.S.

continued in next column

Tuesday, February 26

MS106

Computational Scalability and Complex Geometry in Integral Equation Methods - Part I of II

9:45 a.m.-11:25 a.m.

Room: 102C

For Part 2 see MS140

Due to recent advances in the underlying technology, integral equations have become highly competitive in attacking the most challenging of elliptic, globally-coupled PDE problems. However, compared to a standard off-the-shelf finite element methods, a larger amount of machinery is involved, including fast algorithms, preconditioners, and singular quadrature methods. To make this theoretically attractive family of methods practically usable, many of these components need to be readily available in software form. This minisymposium seeks to bring together practitioners and researchers with the goal of encouraging exchange of ideas and discovering opportunities for collaboration in this direction.

Organizer: Andreas Kloeckner
University of Illinois, U.S.

Organizer: Adrianna Gillman
Rice University, U.S.

Organizer: Timo Betcke
University College London, United Kingdom

9:45-10:05 High Order Layer Potential Evaluation in three Dimensions using Quadrature by Expansion
Manas N. Rachh, Simons Foundation and Flatiron Institute, U.S.

10:10-10:30 Scalable Integral Equation Methods for Problems with Moving Boundaries
Dhairya Malhotra and Antoine Cerfon, Courant Institute of Mathematical Sciences, New York University, U.S.; Michael O'Neil, New York University, U.S.

continued on next page

Tuesday, February 26

MS106

Computational Scalability and Complex Geometry in Integral Equation Methods - Part I of II

continued

10:35-10:55 Parallel Space-time Boundary Element Methods for the Heat Equation

Olaf Steinbach and Stefan Dohr, Technische Universität, Graz, Austria

11:00-11:20 Fast Algorithms with Error Bounds for Quadrature by Expansion

Matt Wala and Andreas Klöckner, University of Illinois, U.S.

Tuesday, February 26

MS107

Unfitted Discretization Methods - Part II of II

9:45 a.m.-11:25 a.m.

Room: 102D

For Part 1 see MS75

The use of unfitted finite element (or isogeometric) techniques is of main interest for different reasons. On one hand, it allows one to eliminate the need to have body-fitted meshes, simplifying the pre-processing step, and having a closer interaction with the CAD data, e.g., in shape or topology optimization frameworks. On the other hand, extended finite element-like techniques are appealing to track interfaces in multi-phase and multi-physics applications. However, the use of unfitted finite element techniques have some drawbacks, e.g., the small cut cell problem and the resulting ill-conditioning of the problem at hand, the imposition of Dirichlet boundary conditions, etc. In this minisymposium, we are interested in works dealing with unfitted finite element simulations in a broad sense, including stabilization techniques for well-posedness, high-order frameworks, applications with advanced discretization spaces, robust and parallel preconditioning techniques, application to PDEs in manifolds and moving domains, surface extraction and cut cell integration techniques, adaptivity, etc., but also the use of these techniques in complex applications, e.g., additive manufacturing and bio applications, and software and implementation aspects.

Organizer: Guglielmo Scovazzi
Duke University, U.S.

Organizer: Xianyi Zeng
University of Texas at El Paso, U.S.

Organizer: Santiago Badia
Universitat Politecnica de Catalunya, Spain

9:45-10:05 The Shifted Boundary Method: A New Framework for Embedded Domain Computations

Guglielmo Scovazzi, Alex Main, Nabil Atallah, Ting Song, Kangan Li, Oriol Colomés, and Leo Nouveau, Duke University, U.S.

10:10-10:30 The Shifted Boundary Method for Embedded Domain Computations: Application to Solid Mechanics

Nabil Atallah, Alex Main, and Guglielmo Scovazzi, Duke University, U.S.

10:30-10:55 Embedded Stabilized Methods for Free Surface Flow Problems

Oriol Colomes, Leo Nouveau, and Guglielmo Scovazzi, Duke University, U.S.

11:00-11:20 The Shifted Interface Method for Embedded Interface Computations

Kangan Li, Nabil Atallah, and Guglielmo Scovazzi, Duke University, U.S.

continued in next column

Tuesday, February 26

MS108

Physics Motivated Problems in Machine Learning - Part I of II

9:45 a.m.-11:25 a.m.

Room: 111A

For Part 2 see MS141

There has been considerable recent interest in using machine learning techniques to solve or assist in the solution of problems in the physical sciences. In this minisymposium we explore recent developments, particularly focusing on learning unknown dynamics or efficient representations of dynamics in a reduced basis. We will also investigate efforts to constrain learning algorithms to produce more physical and interpretable models.

Organizer: Samuel Rudy
University of Washington, U.S.

Organizer: J. Nathan Kutz
University of Washington, U.S.

Organizer: Steven Brunton
University of Washington, U.S.

9:45-10:05 Deep Learning of Dynamics and Signal-noise Decomposition

Samuel Rudy, J. Nathan Kutz, and Steven Brunton, University of Washington, U.S.

10:10-10:30 Nngp: Deep Neural Networks Combined with Gaussian Processes

George E. Karniadakis, Brown University, U.S.

10:35-10:55 A Sequential Sampling Strategy for Extreme Event Statistics in Nonlinear Dynamical Systems

Themistoklis Sapsis, Massachusetts Institute of Technology, U.S.; Mustafa A. Mohamad, Courant Institute of Mathematical Sciences, New York University, U.S.

11:00-11:20 Numerical Aspects for Approximating Governing Equations using Data

Dongbin Xiu and *Kailiang Wu*, Ohio State University, U.S.

Tuesday, February 26

MS109

Physics/Chemistry-aware Machine Learning - Part I of II

9:45 a.m.-11:25 a.m.

Room: 111B

For Part 2 see MS142

The explosive growth in machine learning and data-sciences has unavoidably affected computational modeling of physical systems by providing a powerful alternative to overcoming traditional challenges in chemical and physical simulations pertaining to the large disparity in scales,the extremely high-dimensionality of the descriptions involved and the gigantic, and often discrete, search spaces in the design of new materials. While a direct transfer and application of machine learning tools to data generated by physics/chemistry simulations and experiments has already produced promising results, it is becoming more and more apparent that in order to provide further breakthroughs a synergistic environment should be developed. The challenge that this minisymposium hopes to address is how one can integrate physical/chemical laws, which provide a potentially infinite source of data and encode invariances and symmetries that govern the behavior of physical systems, with the undisputed ability of un/semi/supervised learning tools for extracting patterns, learning relationships and effecting dimensionality reductions.

Organizer: Nicholas Zabaras
University of Notre Dame, U.S.

9:45-10:05 Neural Networks Learning Quantum Chemistry

Olexandr Isayev, University of North Carolina at Chapel Hill, U.S.

10:10-10:30 Enforcing Physical Constraints in Machine Learning with Application to Fluid Flows

Jinlong Wu and Yang Zeng, Virginia Tech, U.S.; Karthik Kashinath, Adrian Albert, and Mr Prabhat, Lawrence Berkeley National Laboratory, U.S.; Heng Xiao, Virginia Tech, U.S.

10:35-10:55 Spatio-temporal Tensor Super-resolution using Chemical Priors

Chandrajit Bajaj, University of Texas at Austin, U.S.; Pradeep Ravikumar, Carnegie Mellon University, U.S.

11:00-11:20 Physics-constrained Surrogates for Reduced-order Modeling and Uncertainty Quantification

Constantin Grigo and Phaedon-Stelios Koutsourelakis, Technische Universität München, Germany

continued in next column

Tuesday, February 26

MS110

Recent Advances in Machine Learning and Data-driven Methods for Physical Sciences - Part I of II

9:45 a.m.-11:25 a.m.

Room: 111C

For Part 2 see MS143

Despite the remarkable increase in computational power, most real-world systems are still too complex to simulate in full details. How to utilize available data to facilitate/accelerate the simulations becomes increasingly important in the recent years. Among all methods, and data-driven approaches incorporating machine learning techniques prove themselves as indispensable algorithmic tools for real-world applications. This minisymposium focuses on recent developments in algorithms and applications based on data-driven and machine learning approaches. Topics include data-driven surrogate construction, Bayesian parameters inferences and physics-informed machine learning retracts information of a significant amount real data and provides guides for the system design, decision-making, etc. Applications include problems from computational chemistry, geoscience and hydrology.

Organizer: Jing Li

Pacific Northwest National Laboratory, U.S.

Organizer: Xueyu Zhu

University of Iowa, U.S.

Organizer: Ling Guo

Shanghai Normal University, China

9:45-10:05 Approximate Bayesian Model Inversion for PDEs with Unknown Parameters and Constitutive Relations

David A. Barajas-Solano and Alexandre M. Tartakovsky, Pacific Northwest National Laboratory, U.S.

10:10-10:30 Physics Informed Deep Neural Networks in the Data Limited Regime

Carlos Ortiz-Marrero, Pacific Northwest National Laboratory, U.S.

10:35-10:55 Maximum Entropy Inference for Learning Model Parameters when Data is Unreported

Tiernan Casey, Sandia National Laboratories, U.S.

11:00-11:20 Multifideity Informed Neural Network in Reduced Order Modeling

Chuan Lu and Xueyu Zhu, University of Iowa, U.S.

continued in next column

Tuesday, February 26

MS111

Scalable Parallel Multigrid Solvers - Part I of II

9:45 a.m.-11:25 a.m.

Room: 300A

For Part 2 see MS144

Multigrid methods are a common tool for solving large-scale sparse linear systems, particularly those arising from discretized partial differential equations. However, parallel multigrid methods suffer from reduced scalability as a result of large costs associated with communication. A number of algorithmic and implementation optimizations reduce communication requirements and yield improved scalability. This session provides a variety of talks on state-of-the-art multigrid solvers and methods for improved scalability of parallel implementations with application that extend beyond the multigrid setting.

Organizer: Amanda Bienz

University of Illinois at Urbana-Champaign, U.S.

9:45-10:05 Node-aware Communication in Parallel AMG

Luke Olson, Amanda Bienz, and William D. Gropp, University of Illinois at Urbana-Champaign, U.S.

10:10-10:30 Scalable Structured Solvers on Emerging Architectures

Andrew Reisner and Luke Olson, University of Illinois at Urbana-Champaign, U.S.; David Moulton, Los Alamos National Laboratory, U.S.

10:35-10:55 Robust AMG Solvers for Linear Elasticity Problems

Ruipeng Li and Robert D. Falgout, Lawrence Livermore National Laboratory, U.S.; Victor A. Paludetto Magri, University of Padova, Italy

11:00-11:20 Large Scale Parallel Solution Methods for Electromagnetic Simulations

Jonathan J. Hu, Christian Glusa, and Paul Lin, Sandia National Laboratories, U.S.

Tuesday, February 26

MS112

Recent Advances in Model Reduction and Uncertainty Quantification - Part I of II

9:45 a.m.-11:25 a.m.

Room: 300B

For Part 2 see MS145

Despite unprecedented computational power, simulating high-dimensional numerical problems or large-scale models describing complex physical phenomena may require weeks on supercomputers. Reducing computational costs of these models while preserving high fidelity is essential for handling onerous many-query applications in computational science and engineering. This minisymposium will focus on recent mathematical and computational advances in reduced order modeling and its application in uncertainty quantification. Topics will include novel enhancements to the state-of-the-art reduced order modeling techniques in the linear and nonlinear settings, development in low-rank approximation methods for stochastic problems, etc. The speakers in this minisymposium will discuss work towards broader model order reduction and data-driven approaches that are necessary to support a wide range of physical, biological, and aerospace applications.

Organizer: Jiahua Jiang

Virginia Tech, U.S.

Organizer: Kookjin Lee

Sandia National Laboratories, U.S.

9:45-10:05 Low-rank Solution Methods for Stochastic Eigenvalue Problems

Tengfei Su and Howard C. Elman, University of Maryland, College Park, U.S.

10:10-10:30 A Low-rank Multigrid Method for the Stochastic Steady-state Diffusion Equation

Howard C. Elman and Tengfei Su, University of Maryland, College Park, U.S.

10:35-10:55 Reduced Basis Method: Recent Improvements on Efficiency and Robustness

Yanlai Chen, University of Massachusetts, Dartmouth, U.S.; Jiahua Jiang, Virginia Tech, U.S.; Akil Narayan, University of Utah, U.S.

11:00-11:20 Rank-adaptive Tensor Recovery Based Model Reduction for PDEs with High-dimensional Random Inputs

Kejun Tang and Qifeng Liao, ShanghaiTech University, China

continued in next column

Tuesday, February 26

MS113

Recent Advances in Computational Methods for High Dimensional Bayesian Inversion - Part I of II

9:45 a.m.-11:25 a.m.

Room: 300C

For Part 2 see MS146

In many practical Bayesian inverse problems, the parameters of interest are of very high dimensions. On the other hand, the relation between the parameters and the data is often described by computationally intensive mathematical models. Thus, developing efficient and accurate algorithms for such high dimensional problems poses a big challenge computationally. The purpose of this minisymposium is to bring researchers from different fields to discuss the recent advances in computational methods for such problems.

Organizer: Qifeng Liao

ShanghaiTech University, China

Organizer: Jinglai Li

Shanghai Jiao Tong University, China

9:45-10:05 Bayesian Inference for Structural Error Quantification

Khachik Sargsyan, Sandia National Laboratories, U.S.; Xun Huan, University of Michigan, U.S.; Habib N. Najm, Sandia National Laboratories, U.S.

10:10-10:30 An Approximate Empirical Bayesian Method for Large-Scale Linear-Gaussian Inverse Problems

Jinglai Li, Shanghai Jiao Tong University, China

10:35-10:55 Bayesian Identification of Discontinuous Fields with an Ensemble-based Variable Separation Multiscale Method

Lijian Jiang, Tongji University, China; Guang Lin, Purdue University, U.S.; *Na Ou*, Hunan University, China

11:00-11:20 Title Not Available

Xiang Zhou, City University of Hong Kong, Hong Kong

Tuesday, February 26

MS114

BE: Broader Engagement Technical Research - Part I of II

9:45 a.m.-11:25 a.m.

Room: 302A

For Part 2 see MS148

The Broader Engagement (BE) minisymposium brings together faculty, researchers, and students from a variety of CSE19 conference theme disciplines. The BE program offers a dynamic program aimed at promoting diversity and inclusion through travel grants to SIAM CSE conferences and opportunities for the CSE community to learn about and support diversity and inclusion. This minisymposium features presentations of research from BE funded participants and faculty funded through the Sustainable Research Pathways program organized by the Sustainable Horizons Institute and funded by Lawrence Berkeley National Laboratory.

Organizer: Mary Ann E. Leung
Sustainable Horizons Institute, U.S.

Organizer: Nitin Sukhija
Slippery Rock University, U.S.

9:45-10:05 Data Quality Challenges with Missing Values and Mixed Types in Joint Sequence Analysis
Alina Lazar, Youngstown State University, U.S.

10:10-10:30 Parallel High-resolution Compact Krylov-FFT-type Algorithms for Subsurface Scattering Problems
Yury Grazin, Ronald L. Gonzales, and Yun Teck Lee, Idaho State University, U.S.

10:35-10:55 Parallel High-resolution Compact Schemes for Thermal Convection Simulation
Ronald L. Gonzales and Yuiry Gryazin, Idaho State University, U.S.

11:00-11:20 Scalable Direct Solver for Compact Stencil Calculations on Rectangular Grids
Yun Teck Lee, Idaho State University, U.S.

Tuesday, February 26

MS115

BE: Securing Extreme-Scale Scientific Computing

9:45 a.m.-11:25 a.m.

Room: 302B

Securing High Performance Computing (HPC) systems from internal threats is becoming extremely challenging due to the number and the heterogeneous nature of their components and due to the complex interactions among them. Such systems are usually shared by a number of internal users that, even unintentionally, could introduce malicious software. Given the HPC Systems encompasses large amount of available resources, such, as CPU, memory and network, these systems are of particular interest to prospective adversaries, who could potentially exploit these systems in order to launch further attacks towards other resources in the HPC environments. Moreover, storing data generated by monitoring tools and security products for offline analysis has become unfeasible as well. Thus, protecting HPC systems has become extremely challenging for the scientists given the typical issues related to Big Data and complex and heterogeneous nature of these infrastructures. Through series of invited talks, this minisymposium aims to highlight the collaborative efforts that are underway to develop and deploy effective solutions for protecting the distributed cyberinfrastructure (CI) supporting scientific research workflows. The focus of this minisymposium is to encourage, educate and map the views and perspectives identified to a list of core challenges for developing action plans and potential collaborations for securing scientific fidelity of computing systems.

Organizer: Nitin Sukhija
Slippery Rock University, U.S.

9:45-10:05 Cybersecurity for Trustworthy Science: The NSF Cybersecurity Center of Excellence
Von Welch, Indiana University, U.S.

10:10-10:30 Minisymposterium: Leveraging Data Analytics Tools for Fraud Prevention
Blair William, University of Pittsburgh Medical Center, U.S.

10:35-10:55 Cyber Security in a High Performance Computing Environment
David Dampier, University of Texas, San Antonio, U.S.

11:00-11:20 Resilience Modeling and Analysis of Scheduling Scientific Applications
Nitin Sukhija, Slippery Rock University, U.S.

continued in next column

Tuesday, February 26

MS116

Advances in Unstructured Mesh Algorithms and their Applications - Part I of II

9:45 a.m.-11:25 a.m.

Room: 303A

For Part 2 see MS150

Unstructured meshes continue to be utilized in many computational science and engineering simulations. Recently, advancements have been made to unstructured mesh algorithms in several areas including optimization-based mesh quality improvement, hybrid volume meshing, multi-material remapping, and moving meshes to list just a few. In addition, meshes have been used in interesting applications involving fluid-structure interaction, and many others. This minisymposium will feature presentations on some of these advancements in meshing algorithms and their applications.

Organizer: Mike Stees
University of Kansas, U.S.

Organizer: Suzanne M. Shontz
University of Kansas, U.S.

9:45-10:05 A Vision for Research in Unstructured Mesh Generation

Suzanne M. Shontz, University of Kansas, U.S.

10:10-10:30 Data Remapping between Distributed General Unstructured Meshes

Gary Dilts, Rachel Ertl, Charles R. Ferenbaugh, Rao V. Garimella, Angela Herring, Eugene Kikinzon, Chris Malone, *Navamita Ray*, Daniel Shevitz, and Jan Velechovsky, Los Alamos National Laboratory, U.S.

10:35-10:55 Size Field Driven Advancing Layer and Hex Core Hybrid Volume Meshing

John P. Steinbrenner, Pointwise, Inc., U.S.

11:00-11:20 Combining Structured and Unstructured Meshing for Efficient and Robust 3D Fluid-structure Interaction Applications

Yongxing Wang, *Peter K. Jimack*, and Mark A. Walkley, University of Leeds, United Kingdom

Tuesday, February 26

MS117

Batched BLAS: API Standardization, Libraries, and Applications - Part I of II

9:45 a.m.-11:25 a.m.

Room: 303B

For Part 2 see MS151

The Basic Linear Algebra Subprograms (BLAS) are the most widely accepted standard in high performance dense matrix computation. The unified BLAS API, which is implemented by many vendors and research groups, enables performance portability for many applications across different architectures and platforms. The past few years have witnessed a continuously growing interest in optimizing BLAS for a batch of small independent problems, hence the name "Batched BLAS". Such interest is driven by numerous applications, including tensor contractions, sparse solvers, astrophysics, quantum chemistry, and many others. This minisymposium (MS) covers a wide range of ongoing research activities about the Batched BLAS. The MS features three main categories of research. The first one covers standardization of the Batched BLAS API, as well as the next-generation BLAS to support reproducibility and extended precisions. The second category features many libraries, from different vendors and research groups, that provide optimized Batched BLAS routines on different hardware architectures. The third category highlights several scientific applications where optimized Batched BLAS has a great impact as a critical building block. At the end of the MS, the audience will be aware of the latest developments in Batched BLAS. The MS also provides an excellent collaboration opportunity between the Batched BLAS research community and potential application developers.

Organizer: Ahmad Abdelfattah
University of Tennessee, Knoxville, U.S.

Organizer: Stanimire Tomov
University of Tennessee, Knoxville, U.S.

9:45-10:05 Standardization of the Batched BLAS

Sven J. Hammarling, University of Manchester, United Kingdom

10:10-10:30 CUDA Linear Algebra and High Performance Tensor Primitives

Cris Cecka, NVIDIA, U.S.

10:35-10:55 Batched Blas Moving Forward in MAGMA

Ahmad Abdelfattah, University of Tennessee, Knoxville, U.S.; Azzam Haidar, NVIDIA, U.S.; Stanimire Tomov, University of Tennessee, Knoxville, U.S.; Jack J. Dongarra, University of Tennessee and Oak Ridge National Laboratory, U.S.

11:00-11:20 Toward Fast Eigensolvers for Electronic Structure Calculations using Low-rank Approximations

Kadir Akbudak, Ali M. Charara, David E. Keyes, *Hatem Ltaief*, Aleksandr Mikhalev, and Dalal Sukkari, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

continued in next column

Tuesday, February 26

MS118

Hydrodynamics at Small Scales: Fluctuating Hydrodynamics - Part I of II

9:45 a.m.-11:25 a.m.

Room: 201B

For Part 2 see MS152
Featured Minisymposium

With the increased interest in nano- and micro-fluidics, as well as biological systems, it has become necessary to develop tools for hydrodynamic calculations at microscopic and mesoscopic scales. This minisymposium will focus on advances in multiscale numerical methods for simulating flows at mesoscopic scales. Coarse-grained models cover a broad range of time and length scales by incrementally sacrificing physical fidelity for computational efficiency. Of particular interest will be fluctuating hydrodynamics of complex fluids such as reactive mixtures, colloidal passive and active suspensions, and multi-phase fluids. Issues to be discussed will include the inclusion of thermal fluctuations in analytical and computational models, as well as applications in the physical sciences, biology, and engineering.

Organizer: Aleksandar Donev

Courant Institute of Mathematical Sciences, New York University, U.S.

9:45-10:05 Low Mach Number Fluctuating Hydrodynamics for Electrolytes

John B. Bell, Lawrence Berkeley National Laboratory, U.S.

10:10-10:30 Hybrid Discrete-continuous Dynamics Method for Chemically-powered Nanomotors

Nikolaos Voulgarakis, Washington State University, U.S.

10:35-10:55 Nano Hydrodynamics Near Walls

Jaime Arturo de la Torre and Diego Duque-Zumajo, Universidad Nacional de Educación a Distancia, Spain; Diego Camargo, Universidad Nacional de Colombia, Colombia; Pep Español, Universidad Nacional de Educación a Distancia, Spain

11:00-11:20 Fluctuating Hydrodynamics of Flow through Porous Membranes

Daniel R. Ladiges, Lawrence Berkeley National Laboratory, U.S.

continued in next column

Tuesday, February 26

MS119

Model Reduction for Problems with Strong Convection, Sharp Gradients, and Discontinuities - Part I of II

9:45 a.m.-11:25 a.m.

Room: 201C

For Part 2 see MS153
Featured Minisymposium

Model reduction exploits that many phenomena of interest in science and engineering can be approximated well by problem adapted low-dimensional subspaces of the high-dimensional solution spaces. Reduced bases are a prominent example, but also proper orthogonal decomposition and polynomial chaos can be seen from this perspective. However, if coherent structures (e.g., wave fronts or shocks) travel through the physical domain, then traditional model reduction techniques typically fail because these structures introduce high-dimensional features. This minisymposium highlights recent advances in model reduction methods that recover low-dimensionality in convection-dominated problems via transformations, adaptive bases, separation of transport, and other techniques.

Organizer: Maciej Balajewicz

University of Illinois at Urbana-Champaign, U.S.

Organizer: Benjamin Peherstorfer

Courant Institute of Mathematical Sciences, New York University, U.S.

Organizer: Gerrit Welper

University of Central Florida, U.S.

9:45-10:05 Dynamic Coupling of Full and Reduced Models via Randomized Online Basis Updates

Benjamin Peherstorfer, Courant Institute of Mathematical Sciences, New York University, U.S.

continued on next page

10:10-10:30 Reduced Basis Methods for Wave and Schrödinger Equations
Karsten Urban, University of Ulm, Germany

10:35-10:55 The Shifted Pod: A Data Based Model Reduction for Situations with Multiple Transports - Formulations and Application Aspects
Julius Reiss, Philipp Schulze, and Volker Mehrmann, Technische Universität Berlin, Germany

11:00-11:20 Model Reduction for Compressible Fluids: Application to Uncertainty Quantification
Remi Abgrall, Roxana Crisovan, and Davide Torlo, Universität Zürich, Switzerland; Svetlana Tokareva, Los Alamos National Laboratory, U.S.

Tuesday, February 26

MS120

Structure Preserving Techniques for Hyperbolic Systems - Part I of II

9:45 a.m.-11:25 a.m.

Room: 202A

For Part 2 see MS154

The generalization of the maximum principle to hyperbolic systems is the notion of convex invariant domains. The key property defining an invariant domain is that for any initial data in the domain in question, the (entropy satisfying) solution to the hyperbolic system remains in the domain. For instance positive density and positive internal energy are invariant properties for the compressible Euler equations; likewise, positive water height is an invariant property of the shallow water equations. Invariant domain preserving numerical methods are those that preserve invariant domains. Many first-order methods are known to be invariant-domain preserving, but extending theses techniques to higher-order is non-trivial. The difficult balance consists of being invariant, high-order accurate and to satisfy enough entropy inequalities for the numerical solution to converge to an entropy solution. These three constraints are difficult to achieve at the same time, but are essential to get robust and accurate methods. Invariant-domain preserving are important in any application where being 'in bounds' is essential. The objective of minisymposium is to gather specialists to present the latest developments on the above topic and share new ideas. The minisymposium will particularly focus on the approximation of hyperbolic systems.

Organizer: Jean-Luc Guermond
Texas A&M University, U.S.

Organizer: Xiangxiong Zhang
Purdue University, U.S.

9:45-10:05 High Order Accurate Bound-preserving Compact Finite Difference Schemes for Scalar Convection Diffusion Equations
Xiangxiong Zhang, Purdue University, U.S.

10:10-10:30 Entropy Stable and Well-balanced Discontinuous Galerkin Methods for the Shallow Water Equations
Yulong Xing, Ohio State University, U.S.

10:35-10:55 Hybrid Methods for Radiation Transport
Cory Hauck, Oak Ridge National Laboratory, U.S.

11:00-11:20 High-order Convex Limiting for Conservation Equations
Ignacio Thomas, Sandia National Laboratories, U.S.

continued in next column

Tuesday, February 26

MS121

Theoretical and Computational Advancements in Ice-Sheet Modeling - Part I of II

9:45 a.m.-11:25 a.m.

Room: 202B

For Part 2 see MS155

The demand for more accurate answers to questions regarding the future of the Earth's climate is a driving force for the development of Earth System Models (ESMs). Such questions include, for instance, assessment of climate sensitivity, and quantification of Sea Level Rise (SLR). To answer these questions, ESMs are pushing toward higher resolution, both in computational terms (the number of degrees of freedom), and in terms of the physical processes included in the model. When it comes to quantifying SLR, accurate and robust models for ice sheets dynamics are key to achieve reliable predictions. In this respect, advancements in both the computational and modeling aspects are equally important. In fact, discretization of the governing equations of ice-sheet dynamics usually lead to difficult mathematical problems, which require robust and efficient solvers. On the other hand, an accurate modeling of ice-sheets movement can involve physical processes at completely different time and space scales, and may involve physical quantities that are hard (if possible at all) to measure directly. This minisymposium will feature presentations aimed at improving the efficiency and reliability of ice-sheet components of ESMs. Topics of interest are: algorithmic advancements, next-generation HPC software development, model parametrizations and model extensions, multiphysics couplings, data assimilation and uncertainty quantification.

Organizer: Luca Bertagna
Sandia National Laboratories, U.S.

Organizer: Mauro Perego
Sandia National Laboratories, U.S.

Organizer: Jerry Watkins
Sandia National Laboratories, U.S.

Organizer: Irina K. Tezaur
Sandia National Laboratories, U.S.

Organizer: Daniel Martin
Lawrence Berkeley National Laboratory, U.S.

9:45-10:05 Performance Portability in Albany Land Ice using Kokkos

Jerry Watkins and Irina K. Tezaur, Sandia National Laboratories, U.S.

10:10-10:30 H-Div Conforming Methods for Geodynamic Stokes Flow

Tobin Isaac, Georgia Institute of Technology, U.S.

10:35-10:55 Parallel Graph Algorithms to Remove Degenerate Features from Ice Sheet Meshes, and Determine Biconnectivity

Ian Bogle, Rensselaer Polytechnic Institute, U.S.; Karen D. Devine, Mauro Perego, and Siva Rajamanickam, Sandia National Laboratories, U.S.; George M. Slota, Rensselaer Polytechnic Institute, U.S.

11:00-11:20 Inferring Past Temperature Changes in Greenland using the Unscented Transform

Jake Downs and Jesse Johnson, University of Montana, U.S.; Jason Briner, State University of New York at Buffalo, U.S.; NIcolas Young, Columbia University, U.S.; Joshua Cuzzone, California Institute of Technology, U.S.

Tuesday, February 26

MS122

High-order Solvers for Wave Problems - Part I of II

9:45 a.m.-11:25 a.m.

Room: 202C

For Part 2 see MS156

Significant progress in the computational simulation of waves has been made in recent years from consideration of efficient high-order wave solvers, including finite-difference approaches, Fourier-based methods, integral equation methods, Galerkin differences, and other high-order techniques. This minisymposium considers ideas that include the extension of Fourier methods to non-periodic problems, new time-stepping strategies beyond the standard method-of-lines, sparse matrix inversion, Galerkin difference approximations, high-order summation by parts schemes for seismic wave propagation and the use of robust upwind schemes for wave equations on overlapping structured grids. The resulting methods have provided important new mathematical tools applicable to a wide range of areas of science and engineering---including electronics, photonics,meta-materials, microscopy, remote sensing, imaging and communications, among many others: in some cases, the resulting techniques have provided solutions to previously intractable problems. This session will focus on several important topics that have recently attracted particular attention in the field, including modern finite-difference approaches, frequency-time hybrids, Fourier-based methods for general domains, fast high-order integral-equation methods, and Galerkin difference approximations. This minisymposium includes important recent contributions in the area and will help to disseminate and advance the field.

continued in next column

continued on next page

Organizer: Oscar P. Bruno
California Institute of Technology, U.S.

Organizer: William D. Henshaw
Rensselaer Polytechnic Institute, U.S.

9:45-10:05 High-order, Dispersionless, Spatio-Temporally Parallel "Fast-Hybrid" Wave Equation Solver at O(1) Sampling Cost

Thomas Anderson and Oscar P. Bruno, California Institute of Technology, U.S.; Mark Lyon, University of New Hampshire, U.S.

10:10-10:30 Higher-Order Implicit Methods for the Navier-Stokes Equations and Other Linear and Nonlinear Waves

Max Cubillos, Wright-Patterson Air Force Labs, U.S.

10:35-10:55 Approximating Eigenvalue Clusters using FEAST and Application to Beam Propagation

Jay Gopalikrishnan, Portland State University, U.S.

11:00-11:20 High-order Upwind Schemes for the Dispersive Maxwell Equations

William D. Henshaw, Rensselaer Polytechnic Institute, U.S.

Tuesday, February 26

MS123

Application of Fractional Calculus in Material Science and Engineering - Part I of II

9:45 a.m.-11:25 a.m.

Room: 203

For Part 2 see MS157

While the mathematics and computations of fractional calculus have matured in recent years, limited efforts have focused on material science and engineering applications. Fractional calculus has been shown as an effective operator in fractal media, viscoelasticity, and control designs. While the fractional integral has been used to describe the fractal structure of materials which leads to new thermodynamic relations, the fractional derivative could be used to describe viscoelasticity properties of materials, not only does the calibration of fractional models of viscoelasticity with experimental measurements provides less error compare to the integer one, but it also provides more information about the structure of a material. This area opens up an application of fractional calculus which may describe the multiscale thermomechanical material behavior of many polymers. On the other hand, fractional calculus has been used in the formulation of variational calculus, Euler-Lagrange equations and optimal control systems where the fractional derivative has been replaced with the integer derivative to create a new set of necessary conditions that must be satisfied by an optimal control law and its associated state-control equations. This minisymposium is designed to bring together the research scholars who are working in the application of fractional calculus in material science and engineering to present their results and exchange ideas to propose future research in this field.

Organizer: Somayeh Mashayekhi
Florida State University, U.S.

Organizer: William Oates
Florida State University, U.S.

9:45-10:05 Modeling Diffusion Signal Decay of Magnetic Resonance Images

Richard L. Magin, University of Illinois at Chicago, U.S.

10:10-10:30 Fractional Generalizations of Gradient Elasticity and Higher Order Diffusion

Elias Aifantis, Michigan Technological University, U.S.; Konstantinos Parisis, Aristotle University of Thessaloniki, Greece

10:35-10:55 A Physical Interpretation of Fractional Viscoelasticity Based on the Fractal Structure of Media

Somayeh Mashayekhi, M. Yousuff Hussaini, and William Oates, Florida State University, U.S.

11:00-11:20 Fractional Viscoelasticity of Auxetic Foams

Eugenia Stanisauskis, Florida State University, U.S.; Paul Miles, North Carolina State University, U.S.; William Oates, Florida State University, U.S.

continued in next column

Tuesday, February 26

MS124

Multiphysics Simulation with MOOSE - Part I of II

9:45 a.m.-11:25 a.m.

Room: 205

For Part 2 see MS158

The Multiphysics Object Oriented Simulation Environment (MOOSE) framework is Idaho National Laboratory's (INL) premier open source modeling and simulation tool. Under development since 2008, and available on GitHub since 2014, the framework provides a flexible, powerful, and analyst-friendly computational environment for the numerical solution of partial differential equations via the finite element method. The framework is written in C++, supports parallel 1, 2, and 3D implicit and explicit finite element simulations, and makes extensive use of the finite element building blocks provided by the libMesh finite element library, the solvers and preconditioners of PETSc, and the parallel communication routines of MPI. A major emphasis of the framework is facilitating multiphysics simulations by lowering the barriers of entry which inevitably arise when coupling different codes together, and codifying these coupling methods into a consistent programming interface. This minisymposium aims to bring together MOOSE users and developers to discuss the latest developments within the framework, as well as advances in both public and internally-developed MOOSE-based applications. Application studies are welcome from all areas of computational science, with a particular emphasis on multiphysics simulation, including algorithmic explorations as well as scientific and engineering studies enabled by MOOSE.

Organizer: John W. Peterson
Idaho National Laboratory, U.S.

Organizer: Cody J. Permann
Idaho National Laboratory, U.S.

9:45-10:05 An Overview of the Multiphysics Simulation Capabilities of MOOSE

Cody J. Permann, Idaho National Laboratory, U.S.

10:10-10:30 Porous Media Thermal Hydraulics Simulations of Pebble Bed Nuclear Reactors using Pronghorn

April Novak, University of California, Berkeley, U.S.; John Peterson and Ling Zou, Idaho National Laboratory, U.S.; Rachel Slaybaugh, University of California, Berkeley, U.S.; Richard Martineau, Idaho National Laboratory, U.S.

10:35-10:55 Moose Multiphysics Simulations Employing Physics-based, Transient Material Interfaces

James B. Tompkins, Texas A&M University, U.S.

11:00-11:20 Electromagnetics Simulations with Vector-valued Finite Elements in MOOSE

Casey T. Icenhour, North Carolina State University, U.S.; Alex Lindsay and Richard Martineau, Idaho National Laboratory, U.S.; Steven Shannon, North Carolina State University, U.S.

continued in next column

Tuesday, February 26

MS125

Computational Tools and Precision Medicine

9:45 a.m.-11:25 a.m.

Room: 206A

Precision Medicine challenges High Performance Computing (HPC). Challenges include analyzing large amounts of data that is instrumental in guiding precision medicine initiatives using rich computational resources. Precision Medicine also inspires HPC as the large computing resources can facilitating the acceleration of critical discoveries thus providing valuable insights and solutions. This symposium will highlight some of the on-going research, challenges and suitable solutions. The coordination of data generation and analysis cannot rely on manual, centralized approaches as it is predominately done today. One of the talks in this symposium will discuss combinations of machine learning and data analytics approach along with HPC techniques that can transition the runtime analysis of larger and larger MD trajectories towards the exascale era via in situ data analytics. The second talk will discuss advance computing capabilities, collaborations, and solutions for cancer. The third talk will share challenges for analysis and visualization of atomic-detail simulation of minimal cells on large-scale computing platforms and how such findings can create new opportunities for the study of the structure and function of large biomolecular complexes. The fourth talk will present knowledge gained by accelerating the prediction of the evaluation of large protein conformational ensembles that can lead to improved strategies for drug discovery.

Organizer: Sunita Chandrasekaran
University of Delaware, U.S.

9:45-10:05 Acceleration of Prediction of Chemical Shift Structures

Sunita Chandrasekaran and Juan Perilla, University of Delaware, U.S.

continued on next page

10:10-10:30 In Situ Data Analytics for Next Generation Molecular Dynamics Workflows

Michela Taufer, University of Tennessee, U.S.

10:35-10:55 Challenges for Analysis and Visualization of Atomic-detail Simulations of Minimal Cells

John E. Stone, University of Illinois at Urbana-Champaign, U.S.

11:00-11:20 Capabilities, Collaboration and Cancer: Co-design for Advanced Computing Solutions for Cancer

Eric Stahlberg, National Cancer Institute, U.S.

Tuesday, February 26

MS126

Data Assimilation and Optimization in Physiologic Modeling - Part I of II

9:45 a.m.-11:25 a.m.

Room: 206B

For Part 2 see MS160

Physics-based modeling for biological systems has been advanced over the past decade and plays an increasingly important role in understanding the underlying physiological mechanism and assisting clinical management of associated diseases. Although increased model complexity can better reveal detailed information of the biological system, it also poses great difficulties on how to calibrate the model to be patient-specific. Clinical measurement data (e.g., medical imaging) can be utilized to perform model calibration and parameter/hidden state estimation, which can be formularized as inverse problems. Therefore, proper optimization tools and data assimilation techniques are needed. However, solving these inverse problems is challenging due to large inter-patient variability and increased computational complexity. Moreover, large uncertainties from both clinical data and computational model also pose more difficulties. This minisymposium hopes to address the aforementioned challenges by reviewing and discussing recent advances in developing efficient optimization/data assimilation approaches and quantifying associated uncertainties for inverse problems in biomedical applications.

Organizer: Shawn Shadden

University of California, Berkeley, U.S.

Organizer: Jian-Xun Wang

University of Notre Dame, U.S.

9:45-10:05 Improving Convergence of Iterative Ensemble Kalman Filter by Preventing "Early Stopping" via a Resampling Approach

Jian-Xun Wang, University of Notre Dame, U.S.; Jiacheng Wu and Shawn Shadden, University of California, Berkeley, U.S.

10:10-10:30 Sensitivity Analysis, Model Reduction and Parameters Estimation for Thrombosis Modeling for CFD

Rodrigo Mendez, Simon Mendez, and Franck Nicoud, University of Montpellier, France; Didier Lucor, LIMSI-CNRS, France

10:35-10:55 Data Assimilation Methods for Defective Boundary Data Problems in Computational Hemodynamics

Alessandro Veneziani, Emory University, U.S.; Rodrigo M. Romarowski, Simone Morganti, and Ferdinando Auricchio, University of Pavia, Italy; Adrien Lefieux, Emory University, U.S.

11:00-11:20 Strategies for Handling Uncertainties in Model Parameters and Measurements for Cardiovascular Applications

Leif Hellevik and Lucas Omar Mueller, Norwegian University of Science and Technology, Norway

continued in next column

Tuesday, February 26

MS127

Design and Usability of High-performance PDE Software Engines and Frameworks - Part I of II

9:45 a.m.-11:25 a.m.

Room: 206C

For Part 2 see MS161
Featured Minisymposium

The combined complexity of the application areas being studied, the applied numerical methods, and the requirements of parallel hardware make the development of high performance PDE engines and frameworks a challenging area of research. In this minisymposium we bring together users and developers of several PDE software packages. We will discuss the requirements and expectations that users have of production-ready simulation software as well as the design decisions that go into the implementation of such software.

Organizer: Anne Reinarz
Technische Universität München, Germany

Organizer: David Ham
Imperial College London, United Kingdom

Organizer: Tobias Weinzierl
Durham University, United Kingdom

Organizer: Michael Bader
Technische Universität München, Germany

9:45-10:05 The Dividends of Investing in Computational Software Design
Anshu Dubey, Argonne National Laboratory, U.S.

10:10-10:30 Modularity of Lowlevel Forest-of-octree Libraries
Johannes Holke, German Aerospace Center (DLR), Germany; Carsten Burstedde, Universität Bonn, Germany

10:35-10:55 Exahype - An Exascale Engine for Solving Hyperbolic PDEs
Tobias Weinzierl, Durham University, United Kingdom

continued in next column

11:00-11:20 Firedrake: High Productivity High Performance Simulation
David Ham, Imperial College London, United Kingdom; Lawrence Mitchell, Durham University, United Kingdom; Thomas H. Gibson, Tianjiao Sun, and Miklós Homolya, Imperial College London, United Kingdom; Rob C. Kirby, Baylor University, U.S.; Paul Kelly, Imperial College London, United Kingdom

Tuesday, February 26

MS128

Exascale Applications with High-Order Methods - Part I of II

9:45 a.m.-11:25 a.m.

Room: 206D

For Part 2 see MS162
High-order discretizations have the potential to provide an optimal strategy for achieving high performance and delivering fast, efficient, and accurate simulations on next-generation architectures. This minisymposium will discuss the next-generation high-order discretization algorithm, based on finite/spectral element approaches that will enable a wide range of important scientific applications to run efficiently on future architecture.

Organizer: MiSun Min
Argonne National Laboratory, U.S.

Organizer: Paul Fischer
University of Illinois at Urbana-Champaign, U.S.

9:45-10:05 Bernstein-Bezier Techniques in High Order Finite Element Analysis
Mark Ainsworth and Shuai Jiang, Brown University, U.S.

10:10-10:30 Nektar++: Development of Large Scale, Spectral/Hp Element Simulation of Flow around Complex Geometries of Industrial Interest
Spencer Sherwin, Imperial College London, United Kingdom; David Moxey, University of Exeter, United Kingdom; Chris Cantwell, Imperial College London, United Kingdom; Mike Kirby, University of Utah, U.S.; Joaquim Peiro, Imperial College London, United Kingdom

10:35-10:55 A Performance Portable Implementation of the Spectral Elements Atmosphere Dynamical Core of E3SM
Luca Bertagna, Michael Deakin, Oksana Guba, Daniel Sunderland, Andrew Bradley, Irina K. Tezaur, Mark A. Taylor, and Andrew Salinger, Sandia National Laboratories, U.S.

11:00-11:20 Efficient Implementation of a High Order Control Volume Finite Element Scheme for Low-mach Flow
Robert Knaus, Sandia National Laboratories, U.S.

Tuesday, February 26

MS129

Performance Portability and Numerical Libraries: Challenges and Opportunities for Sustainable Science - Part II of II

9:45 a.m.-11:25 a.m.

Room: 207

For Part 1 see MS95
Featured Minisymposium

A software component is said to be performance portable if it performs well across a range of architectures and problem configurations with modest development and maintenance effort. There are many programming tools that strive to assist with portability, but many challenges remain, especially balancing vectorization with memory locality across disparate architectures and managing composition across module boundaries where developers must eschew assumptions about environment and context. Optimized numerical libraries provide an opportunity to encapsulate any necessary specialization in a form that can be leveraged by many applications, but innovative interfaces and implementations are needed to serve diverse user requirements and to control complexity when porting and optimizing for an ever-growing range of architectures. This session brings together exemplars, promising techniques, and lessons learned in developing robust performance-portable applications and numerical libraries for computational science and engineering

Organizer: Jed Brown
University of Colorado Boulder, U.S.

Organizer: Sunita Chandrasekaran
University of Delaware, U.S.

Organizer: Rebecca J. Hartman-Baker
Lawrence Berkeley National Laboratory, U.S.

9:45-10:05 Library Interface Design and Performance Portability
Jed Brown, Jeremy Thompson, and Valeria Barra, University of Colorado Boulder, U.S.

10:10-10:30 BLIS: A Case Study in Performance and Portability
Devangi N. Parikh and Field G. Van Zee, University of Texas at Austin, U.S.

10:35-10:55 Raptor: Parallel Algebraic Multigrid with Node-aware Communication
Amanda Bienz, Luke Olson, and William D. Gropp, University of Illinois at Urbana-Champaign, U.S.

11:00-11:20 Performance Portability of Communication Optimal GMRES
Katarzyna Swirydowicz and Stephen Thomas, National Renewable Energy Laboratory, U.S.; Julien Langou, University of Colorado, Denver, U.S.; Shreyas Ananthan, National Renewable Energy Laboratory, U.S.; Ulrike Meier Yang, Lawrence Livermore National Laboratory, U.S.

continued in next column

Tuesday, February 26

MS130

Computational Engineering (BGCE) Student Paper Prize - Part I of II

9:45 a.m.-11:25 a.m.

Room: 401A

For Part 2 see MS164
The 7th Bavarian Graduate School in Computational Engineering (BGCE) Student Paper Prize will be awarded at the 2019 SIAM CS&E Conference for outstanding student work in the field of Computational Science and Engineering. Eligible for the prize will be undergraduate and graduate students prior to receiving their PhD. Candidates are required to summarize their work in a short paper of at most 4 pages. The prize finalists will present their work in this minisymposium. The prize award announcement will be scheduled at one of the last days of the conference.

Organizer: Tobias Neckel
Technische Universität München, Germany

Organizer: Hans-Joachim Bungartz
Technische Universität München, Germany

Organizer: Dietmar Fey
Universität Erlangen-Nürnberg, Germany

Organizer: Alexander Ditter
Universität Erlangen-Nürnberg, Germany

Speakers To Be Announced

See online program for an update on this session.

Tuesday, February 26

MS131

Homogenization and Reduced Order Modelling for Wave Equations - Part I of II

9:45 a.m.-11:25 a.m.

Room: 401B

For Part 2 see MS165

Wave propagation in highly heterogeneous media is very important in applications ranging from geophysics to complex engineered structures such as metamaterials. It is also very challenging computationally as short waves induced by the medium do not decay rapidly in space-time, and thus must be modeled, while the medium variability itself leads to dispersive effects for longer waves. Speakers in this minisymposium will discuss a variety of techniques for treating these challenging problems.

Organizer: Thomas M. Hagstrom

Southern Methodist University, U.S.

Organizer: Daniel Appelo

University of Colorado Boulder, U.S.

9:45-10:05 Numerical Homogenization for Waves using Hermite Methods

Daniel Appelo, University of Colorado Boulder, U.S.; Thomas M. Hagstrom, Southern Methodist University, U.S.

10:10-10:30 Long Time Homogenization of the Wave Equation in Locally Periodic Media

Timothée Pouchon, École Polytechnique Fédérale de Lausanne, Switzerland

10:35-10:55 A Backscattering Model Based on Corrector Theory of Homogenization for the Random Helmholtz Equation

Olivier Pinaud, Colorado State University, U.S.

11:00-11:20 Sparse Data-driven Reduced Order Models and Inverse Scattering

Vladimir Druskin, Worcester Polytechnic Institute, U.S.; Liliana Borcea, University of Michigan, U.S.; Alexander V. Mamonov, University of Houston, U.S.; Mikhail Zaslavsky, Schlumberger-Doll Research, U.S.

Tuesday, February 26

MS132

Latest Advances in Topology and Shape Optimization - Part I of II

9:45 a.m.-11:25 a.m.

Room: 401C

For Part 2 see MS166

Topology and shape optimization are design methodologies to obtain the optimal geometry of engineering structures to minimize a cost function while satisfying design constraints. These methods are well established design tools that have found industrial application in recent years to design materials with specific mechanical, thermal, acoustic properties, etc. where they can provide non-intuitive and more optimal designs than the traditional trial and error can do. However, many challenges are still unresolved as these methods are expanded to larger scale problems with more complex physical phenomena, requiring new numerical methods, optimization formulations and physical modeling. This minisymposium aims to bring together researchers whose works are intended to extend the possibilities in topology optimization towards larger problems and new physics.

Organizer: Miguel Salazar de Troya

Lawrence Livermore National Laboratory, U.S.

9:45-10:05 A Stress Constrained Level Set Topology Optimization Method with the Cut Finite Element Method

Miguel Salazar de Troya, Lawrence Livermore National Laboratory, U.S.; Daniel Tortorelli, University of Illinois at Urbana-Champaign, U.S.

10:10-10:30 Algorithmic Differentiation for Shape Optimization Problems in the High Level Finite Element Framework FEniCS

Jørgen Dokken, Simon W. Funke, and Sebastian Kenji Mitusch, Simula Research Laboratory, Norway; Stephan Schmidt, Universität Würzburg, Germany

10:35-10:55 Latest Advances in Topology and Shape Optimization

H. Alicia Kim, Sandilya Kambampati, and Zongliang Du, University of California, San Diego, U.S.

11:00-11:20 H-Div Based ZZ Error Indicator for Adaptive Mesh Refinement in Topology Optimization

Kenneth Swartz, University of Illinois at Urbana-Champaign and Lawrence Livermore National Laboratory, U.S.; Youngsoo Choi, Jun Kudo, Miguel Salazar de Troya, and Dan White, Lawrence Livermore National Laboratory, U.S.; Dan Tortorelli, University of Illinois at Urbana-Champaign, U.S.

continued in next column

Tuesday, February 26

MS133

Quantitative Image Analysis - Part I of II

9:45 a.m.-11:25 a.m.

Room: 402A

For Part 2 see MS167

Image processing is ubiquitous in today's world. However, it is often the case that only a qualitative approach is needed, for example segmenting images on a smart phone app, in which aesthetics are usually valued over precision. For scientific imaging, however, quantitative techniques are necessary, and a wide variety of approaches have been developed based on the specific applications at hand. From a computational standpoint, interesting applications may involve running queries on large data sets, for instance those found in astronomy and agriculture, or inverse problems, such as those found in tomography and computer vision. This minisymposium will highlight several of the current techniques used in scientific imaging and discuss the computational challenges involved with each approach. A special focus will be on how these challenges can inform other areas of computational science and engineering.

Organizer: Laramie Paxton

Washington State University, U.S.

Organizer: Matthew Sottile

Noddle.io, U.S.

9:45-10:05 Challenges in Quantitative Image Analysis for Science and Engineering

Matthew Sottile, Noddle.io, U.S.

10:10-10:30 Pulling 3D Cloud Structure Out of Passive Overhead Imagery: Observational and Computational Challenges

Anthony B. Davis, California Institute of Technology, U.S.; Quentin Libois, Universite de Toulouse, France; Nicolas Ferlay, Universite de Lille 1, France; Aviad Levis and Yoav Schechner, Technion Israel Institute of Technology, Israel

10:35-10:55 Camera-Based, Mobile Disease Surveillance using Convolutional Neural Networks

Altansuren Tumurbaatar, Washington State University, U.S.; Lauren Charles, Pacific Northwest National Laboratory, U.S.; Saiteja Abbu, State University of New York, Albany, U.S.; Olatunde Madandola, Georgia Southern University, U.S.; Liangyu Tan, Iowa State University, U.S.

11:00-11:20 Phase Transformation Models as Regularizers for Electron Microscope Imagery

Jeff Simmons, Air Force Research Laboratory, U.S.; Amir Koushyar Ziabari, Purdue University, U.S.; Jeffery Rickman, Lehigh University, U.S.; Lawrence Drummy, Air Force Research Laboratory, U.S.; Charles Bouman, Purdue University, U.S.

continued in next column

Tuesday, February 26

MS134

Reproducibility in Network Algorithms - Part I of II

9:45 a.m.-11:25 a.m.

Room: 402B

For Part 2 see MS168

The growing use of network analysis in diverse critical applications, such as drug design, recommendation systems and criminal justice, has highlighted the importance of creating reproducible results. In this minisymposium, we focus on the challenges in achieving reproducibility in network analysis in the context of (i) modeling networks from raw data and (ii) in analyzing the networks, particularly dynamic ones that change with time. The first session will focus on methods for creating networks from data, and how modeling choices affect the structure of the network. In the second session, we will focus on the analysis algorithms and how variations in their design or implementation choices can change the results.

Organizer: Sanjukta Bhowmick

University of Nebraska, Omaha, U.S.

Organizer: Boyana Norris

University of Oregon, U.S.

9:45-10:05 Graph Structure Inference: Methodology and Applications

Ivan Brugere and Tanya Y. Berger-Wolf, University of Ilinois at Chicago, U.S.

10:10-10:30 The Core Resilience of Networks

Sucheta Soundarajan, Syracuse University, U.S.

10:35-10:55 Constructing a Graph from user Implicit Feedback

Haozhe Zhang, Iowa State University, U.S.; Yan Qu and *Shad Kirmani*, eBay Research Labs, U.S.

11:00-11:20 The Pursuit of Reproducibility through Well Planned and Executed Evidence Trails

Eric Stephan, Pacific Northwest National Laboratory, U.S.

Tuesday, February 26

MS135

WCD Workshop - 1 of 4

9:45 a.m.-11:25 a.m.

Room: 402C

For Part 2 see MS169
Part of the SIAM Workshop Celebrating Diversity

Organizer: Talea Mayo
University of Central Florida, U.S.

Organizer: Shelby Wilson
Morehouse College, U.S.

Organizer: Rebecca E. Morrison
University of Colorado Boulder, U.S.

9:45-10:05 Multi-scale Modeling of Paracrine PDGF-driven Glioma Growth and Invasion
Susan Massey, Mayo Clinic, U.S.

10:10-10:30 Accuracy and Performance Tradeoffs for Robust Porous Media Discretization with Non-orthogonal Grids and Tensor Coefficients
Reagan Cronin and Jed Brown, University of Colorado Boulder, U.S.

10:35-10:55 An Uncertainty Representation for Model Inadequacy in a Field-scale Contaminant Transport Model
Teresa Portone and Robert D. Moser, University of Texas at Austin, U.S.

11:00-11:20 Computation of Koopman Modes in Chaotic Systems
Nisha Chandramoorthy, Massachusetts Institute of Technology, U.S.

Tuesday, February 26

CP7

Multilevel Methods

9:45 a.m.-11:25 a.m.

Room: 301

Chair: Ana Maria Soane, US Naval Academy, U.S.

9:45-10:00 aSP-AMG: Adaptive Smoothing and Prolongation Algebraic MultiGrid
Victor A. Paludetto Magri, University of Padova, Italy; Andrea Franceschini, Stanford University, U.S.; Carlo Janna, University of Padova, Italy

10:05-10:20 Multigrid Preconditioning of Linear Systems Arising in the Semismooth Newton Solution of Distributed Optimal Control of Elliptic Equations with State Constraints
Jyoti Saraswat, Thomas More College, U.S.; Andrei Draganescu, University of Maryland, Baltimore County, U.S.

10:25-10:40 Multigrid Methods for Optimal Control Problems Constrained by Partial Differential Equations with Random Data
Ana Maria Soane, US Naval Academy, U.S.

10:45-11:00 Improved Convergence of Highly Parallel Additive Multigrid with Scaled Correction
Syam Vangara and Siva Nadarajah, McGill University, Canada

Tuesday, February 26

CP8

Computational Electromagnetics

9:45 a.m.-11:25 a.m.

Room: 201A

Chair: Zecheng Gan, University of Michigan, U.S.

9:45-10:00 Simulating Nonlinear Faraday Waves on a Cylinder
Saad Qadeer, University of North Carolina at Chapel Hill, U.S.

10:05-10:20 Effects of Induced Magnetic Field and Slip Velocity on Transient Rotating Magnetohydrodynamic Electroosmotic Micropumps in a Slowly Varying Periodically Micro-channel
Mohammed Abdulhameed, Federal Polytechnic Bauchi, Nigeria

10:25-10:40 Fast Algorithms for the Self-assembly of Charged Dielectric Spheres
Zecheng Gan, University of Michigan, U.S.; Weihua Geng, Southern Methodist University, U.S.; Robert Krasny, University of Michigan, U.S.

10:45-11:00 Modeling Diffusion Signal Decay of Magnetic Resonance Images
Richard Magin, University of Illinois, Chicago, U.S.

Tuesday, February 26

Career Fair

9:45 a.m.-11:45 a.m.

Room: Ballroom 100A

CSE Editorial Board Meeting

11:30 a.m.-1:00 p.m.

Room: Davenport Hotel -- Meeting Room 11

Lunch Break

11:30 a.m.-1:00 p.m.

Attendees on their own

Tuesday, February 26

PD4

Mentor-Mentee Mixer

11:30 a.m.-12:30 p.m.

Room: 111A

Chair: Alison Marsden, Stanford University, U.S.

Chair: Katherine J. Evans, Oak Ridge National Laboratory, U.S.

Chair: Stefan Wild, Argonne National Laboratory, U.S.

The goal of this mixer is to provide a casual and interactive opportunity for students and junior scientists to interact and build their network with more seniors SIAM scientists. It is also provides an opportunity for senior scientists to learn first-hand about all the great next generation science underway by the newest members of the community.

Tuesday, February 26

PD5

Mid-Career Panel

11:30 a.m.-12:30 p.m.

Room: Conference Theater

Chair: Stefan Wild, Argonne National Laboratory, U.S.

Chair: Alison Marsden, Stanford University, U.S.

Chair: Katherine J. Evans, Oak Ridge National Laboratory, U.S.

The goal of this panel is to provide information and guidance to mid-career scientists looking at the next level of their career. Choices such as line management, project management, scientific impact are all considerations. Several panelists with diverse experiences will be on hand to discuss their experiences with these options and can provide some context to navigate the often overlooked transition from early career to senior scientist.

Misha Kilmer
Tufts University, U.S.

Sven Leyffer
Argonne National Laboratory, U.S.

Lois Curfman McInnes
Argonne National Laboratory, U.S.

Tuesday, February 26

IP4

Fluid-Structure Interaction in Medicine and Biology: Methods, Models, and Applications

1:00 p.m.-1:45 p.m.

Room: Ballroom 100BC

Chair: Matthew G. Knepley, State University of New York at Buffalo, U.S.

Fluid-structure interaction (FSI) is ubiquitous in nature and occurs at molecular to environmental scales, from the writhing of DNA in nucleoplasm, to the beating of cilia and flagella and the projection of lamellipodia and bleb-like protrusions by motile cells, to the flow of blood in the heart, to swimming fish and flying birds and insects, to the dispersal of seeds and pollen in the wind. This talk will describe numerical methods and computational infrastructure for FSI, focusing on extensions of the immersed boundary (IB) method for fluid-structure interaction and applications of these methods to various models in medicine and biology. Different approaches are needed for FSI involving rigid and elastic structures, but both can be addressed within the framework of the IB method. I will discuss IB methods for FSI with prescribed structural kinematics and methods for FSI involving flexible bodies that use nonlinear structural dynamics formulations. I also will describe new extensions of these methods that aim to achieve higher-order accuracy for applications involving realistic biological and physiological geometries. I will survey applications of these IB methods in biology and medicine, including flagellar mechanics, aquatic locomotion and neuro-mechanical feedback, and esophageal transport. I will also detail ongoing work to develop IB models of the heart and its valves and applications to cardiovascular medical devices.

Boyce E. Griffith
University of North Carolina at Chapel Hill, U.S.

Tuesday, February 26

Coffee Break

1:45 p.m.-2:15 p.m.

Room: Ballroom Foyer

Career Fair

2:00 p.m.-4:00 p.m.

Room: Ballroom 100A

Tuesday, February 26

MS136

Scalable Optimization: Paradigms, Tools, and Applications - Part II of II

2:15 p.m.-3:55 p.m.

Room: Ballroom 100BC

For Part 1 see MS102

This minisymposium seeks to bring together experts in different optimization domains (PDEs, DAEs, stochastic, global, data/learning, mixed-integer) to discuss computational scalability issues as well as strategies and tools to overcome them.

Organizer: Victor M. Zavala
University of Wisconsin, Madison, U.S.

Organizer: Mihai Anitescu
University of Chicago, U.S.

Organizer: Kibaek Kim
Argonne National Laboratory, U.S.

2:15-2:35 Parallel Multigrid Reduction in Time (MGRIT) Applied to Power Grid Simulations

Carol S. Woodward and Robert D. Falgout, Lawrence Livermore National Laboratory, U.S.; Matthieu Lecouvez, CEA, France; Jacob B. Schroder, University of New Mexico, U.S.; Philip Top, Lawrence Livermore National Laboratory, U.S.

2:40-3:00 Chordal Decompositions in Semidefinite Programming: Trading Stability for Scalability

Arvind Raghunathan, Mitsubishi Electric Research Laboratories, U.S.; Larry Biegler, Carnegie Mellon University, U.S.

3:05-3:25 Structured Modeling and Decomposition Methods in Pyomo

Bethany Nicholson, Carl Laird, and John D. Siirola, Sandia National Laboratories, U.S.

3:30-3:50 Scalable Gaussian Process Computations using Hierarchical Matrices

Christopher Geoga, Rutgers University, U.S.; Mihai Anitescu, Argonne National Laboratory, U.S.; Michael Stein, University of Chicago, U.S.

Tuesday, February 26

MS137

Toward Software Ecosystems for Computational Science and Engineering

2:15 p.m.-3:55 p.m.

Room: Conference Theater

Software---cross-cutting technology that connects advances in mathematics, computer science, and domain-specific science and engineering---is a cornerstone of long-term collaboration and progress in computational science and engineering (CSE). As we leverage unprecedented high-performance computing resources to work toward predictive science, software complexity is increasing due to multiphysics and multiscale modeling, the coupling of simulations and data analytics, and the demand for greater reproducibility and sustainability, all in the midst of disruptive architectural changes. Applications increasingly require the combined use of independent software packages, whose development teams have diverse sponsors, priorities, software engineering expertise, and processes for development and release. The developers of open-source scientific software are increasingly encouraging community contributions and considering more effective strategies for connections among complementary packages. This speakers in this minisymposium will discuss work toward broader software interoperability and scientific software ecosystems needed to support next-generation CSE.

Organizer: Lois Curfmann McInnes
Argonne National Laboratory, U.S.

Organizer: James Willenbring
Sandia National Laboratories, U.S.

2:15-2:35 xSDK: Foundations of a Numerical Software Ecosystem for High-performance CSE

Ulrike M. Yang, Lawrence Livermore National Laboratory, U.S.

2:40-3:00 DUNE Collaborating via Interfaces

Christian Engwer, University of Münster, Germany

3:05-3:25 An Innovative Method for Integration of Simulation/Data/Learning in the Exascale/Post-Moore Era

Kengo Nakajima, University of Tokyo, Japan; Takeshi Iwashita, Hokkaido University, Japan; Takashi Shimokawabe, University of Tokyo, Japan; Takahiro Katagiri, Nagoya University, Japan; Hisashi Yashiro, RIKEN Advanced Institute for Computational Science, Japan; Hiroya Matsuba, RIKEN, Japan

3:30-3:50 Software Development Kits: A Software Integration Strategy for CSE

James Willenbring, Sandia National Laboratories, U.S.

continued in next column

Tuesday, February 26

MS138

Accelerating Electronic Structure Calculations with GPUs - Part II of II

2:15 p.m.-3:55 p.m.

Room: 102A

For Part 1 see MS104

Several electronic structure codes have been ported to GPUs in recent years. However, with recent hardware developments, particularly the architecture of some of the new top supercomputers in the world, taking advantage of GPU accelerators is becoming a more pressing need in this field. In addition, since First-Principles (quantum-based) molecular dynamics can be computationally quite expensive, being able to reduce time-to-solution can significantly enhance the accuracy of many simulations. Additionally larger problems and longer timescales can be explored. In this minisymposium, we will discuss how electronic structure codes, targeting various applications in materials sciences, chemistry and biology, are being adapted and ported to GPU-based architectures. The speakers will cover a broad range of algorithms, models, languages and implementations used today in a variety of codes, and will discuss how these developments are helping to push the computational limits in real applications.

Organizer: Jean-Luc Fattebert
Oak Ridge National Laboratory, U.S.

Organizer: Susan Mniszewski
Los Alamos National Laboratory, U.S.

2:15-2:35 Petaflop Scale Electronic Structure Calculations with RMG and GPU Accelerators

Emil Briggs, Wenchang Lu, and Jerzy Bernholc, North Carolina State University, U.S.

2:40-3:00 GPU-accelerated Real Space Electronic Structure Theory on HPC Resources

William P. Huhn, Bjoern Lange, and Victor Yu, Duke University, U.S.; Mina Yoon, Oak Ridge National Laboratory, U.S.; Volker Blum, Duke University, U.S.

continued on next page

Tuesday, February 26

MS138

Accelerating Electronic Structure Calculations with GPUs - Part II of II

continued

3:05-3:25 Development of Plane-wave and O(N) Methods in Nwchemex for Emerging Exascale Machines

Eric J. Bylaska, Pacific Northwest National Laboratory, U.S.

3:30-3:50 Large Scale Hybrid Functional Plane Wave Pseudopotential Density Functional Theory Calculation on GPU Cluster

Weile Jia, University of California, Berkeley, U.S.

Tuesday, February 26

MS139

Computational Challenges at the Frontiers of Numerical Relativity - Part II of II

2:15 p.m.-3:55 p.m.

Room: 102B

For Part 1 see MS105

In LIGO's first year of observing time, it has detected gravitational waves from five binary black hole mergers and detected the first neutron star binary inspiral. The analysis and interpretation of LIGO data requires careful comparison to model solutions of the gravitational two-body problem. Some black hole binaries, however, pose significant challenges to both analytic approximation techniques and numerical simulations. The theme for our minisymposium proposal is the computational challenges in numerical relativity and gravity research as we enter the era of regular gravitational wave detections and multi-messenger astronomy. This will include discussion of numerical relativity and gravitational wave searches, algorithms for new physics, and new algorithms for modern supercomputer architectures. The speakers represent computational scientists, physicists and data analysis researchers who are working on the next generation of codes for numerical relativity. This talks will be of broad interest and naturally combines expertise and knowledge that reaches beyond physics to computational and computer science.

Organizer: Hari Sundar
University of Utah, U.S.

Organizer: David Neilsen
Brigham Young University, U.S.

2:15-2:35 Numerical Relativity in the Age of Gravitational Wave Observations

Deirdre Shoemaker, Georgia Institute of Technology, U.S.

2:40-3:00 On the Development and use of ADER-DG Methods in Computational Relativistic Astrophysics

Luciano Rezzolla, Max Planck Institute for Gravitation Physics, Potsdam, Germany

3:05-3:25 Numerical Relativity, Artificial Intelligence, Big Data Analytics and High Performance Computing: An Emergent Transdisciplinary Paradigm to Realize Multimessenger Astrophysics

Elliu Huerta, University of Illinois at Urbana-Champaign, U.S.

3:30-3:50 Numerical Relativity at the Extremes

Yosef Zlochower, Rochester Institute of Technology, U.S.

continued in next column

Tuesday, February 26

MS140

Computational Scalability and Complex Geometry in Integral Equation Methods - Part II of II

2:15 p.m.-3:55 p.m.

Room: 102C

For Part 1 see MS106

Due to recent advances in the underlying technology, integral equations have become highly competitive in attacking the most challenging of elliptic, globally-coupled PDE problems. However, compared to a standard off-the-shelf finite element methods, a larger amount of machinery is involved, including fast algorithms, preconditioners, and singular quadrature methods. To make this theoretically attractive family of methods practically usable, many of these components need to be readily available in software form. This minisymposium seeks to bring together practitioners and researchers with the goal of encouraging exchange of ideas and discovering opportunities for collaboration in this direction.

Organizer: Andreas Kloeckner
University of Illinois, U.S.

Organizer: Adrianna Gillman
Rice University, U.S.

Organizer: Timo Betcke
University College London, United Kingdom

2:15-2:35 A Fast Direct Solver for Multilayered Quasi-periodic Scattering
Adrianna Gillman, Rice University, U.S.

2:40-3:00 Adaptive BEM with Inexact PCG Solver Yields Almost Optimal Computational Costs
Stefan Schimanko, Technische Universitaet Wien, Austria

3:05-3:25 An Integral Equation Method for the Cahn-Hilliard Equation in the Wetting Problem
Shidong Jiang, New Jersey Institute of Technology, U.S.; Andreas Kloeckner, University of Illinois, U.S.; Xiao-Ping Wang, Hong Kong University of Science and Technology, Hong Kong; *Xiaoyu Wei*, University of Illinois at Urbana-Champaign, U.S.

3:30-3:50 Symbolic Computation for Layer Potential Evaluation with Quadrature by Expansion
Andreas Kloeckner, University of Illinois, U.S.

continued in next column

Tuesday, February 26

MS141

Physics Motivated Problems in Machine Learning - Part II of II

2:15 p.m.-3:55 p.m.

Room: 111A

For Part 1 see MS108

There has been considerable recent interest in using machine learning techniques to solve or assist in the solution of problems in the physical sciences. In this minisymposium we explore recent developments, particularly focusing on learning unknown dynamics or efficient representations of dynamics in a reduced basis. We will also investigate efforts to constrain learning algorithms to produce more physical and interpretable models.

Organizer: Samuel Rudy
University of Washington, U.S.

Organizer: J. Nathan Kutz
University of Washington, U.S.

Organizer: Steven Brunton
University of Washington, U.S.

2:15-2:35 Transfer Learning and Generalization in Physical Problems
Karthik Duraisamy, University of Michigan, Ann Arbor, U.S.

2:40-3:00 Hidden Physics Models: Machine Learning of Non-linear Partial Differential Equations
Maziar Raissi, Brown University, U.S.

3:05-3:25 Deep Learning Algorithm for the Data-driven Simulation of Noisy Dynamical System
Kyongmin Yeo, IBM T.J. Watson Research Center, U.S.

3:30-3:50 Physics-informed Neural Networks
Paris Perdikaris, University of Pennsylvania, U.S.

Tuesday, February 26

MS142

Physics/Chemistry-aware Machine Learning - Part II of II

2:15 p.m.-3:55 p.m.

Room: 111B

For Part 1 see MS109

The explosive growth in machine learning and data-sciences has unavoidably affected computational modeling of physical systems by providing a powerful alternative to overcoming traditional challenges in chemical and physical simulations pertaining to the large disparity in scales,the extremely high-dimensionality of the descriptions involved and the gigantic, and often discrete, search spaces in the design of new materials. While a direct transfer and application of machine learning tools to data generated by physics/chemistry simulations and experiments has already produced promising results, it is becoming more and more apparent that in order to provide further breakthroughs a synergistic environment should be developed. The challenge that this minisymposium hopes to address is how one can integrate physical/chemical laws, which provide a potentially infinite source of data and encode invariances and symmetries that govern the behavior of physical systems, with the undisputed ability of un/semi/supervised learning tools for extracting patterns, learning relationships and effecting dimensionality reductions.

Organizer: Nicholas Zabaras

University of Notre Dame, U.S.

2:15-2:35 An Information-theoretic Approach to Fusion of Multi-scale Data and Models

Hannah Lu, Francesca Boso, and Daniel M. Tartakovsky, Stanford University, U.S.

2:40-3:00 Neural Style Transfer for Enriching Turbulent Flow

Ryan King, Peter Graf, and Michael Sprague, National Renewable Energy Laboratory, U.S.

3:05-3:25 Classical and Quantum Machine Learning with Tensor Networks

Miles Stoudenmire, Simons Foundation and Flatiron Institute, U.S.

3:30-3:50 Learning Atomic Force Fields with Compositional Neural Networks

Risi Kondor, University of Chicago, U.S.

continued in next column

Tuesday, February 26

MS143

Recent Advances in Machine Learning and Data-driven Methods for Physical Sciences - Part II of II

2:15 p.m.-3:55 p.m.

Room: 111C

For Part 1 see MS110

Despite the remarkable increase in computational power, most real-world systems are still too complex to simulate in full details. How to utilize available data to facilitate/accelerate the simulations becomes increasingly important in the recent years. Among all methods, and data-driven approaches incorporating machine learning techniques prove themselves as indispensable algorithmic tools for real-world applications. This minisymposium focuses on recent developments in algorithms and applications based on data-driven and machine learning approaches. Topics include data-driven surrogate construction, Bayesian parameters inferences and physics-informed machine learning retracts information of a significant amount real data and provides guides for the system design, decision-making, etc. Applications include problems from computational chemistry, geoscience and hydrology.

Organizer: Jing Li

Pacific Northwest National Laboratory, U.S.

Organizer: Xueyu Zhu

University of Iowa, U.S.

Organizer: Ling Guo

Shanghai Normal University, China

2:15-2:35 Data-driven Approach of Quantifying Uncertainty in Complex Systems with Arbitrary Randomness

Jing Li, Huan Lei, Peiyuan Gao, and Nathan Baker, Pacific Northwest National Laboratory, U.S.

continued on next page

2:40-3:00 Physics-informed GANs for Stochastic Differential Equations

Liu Yang, Dongkun Zhang, and George Em Karniadakis, Brown University, U.S.

3:05-3:25 Coupling of Non-intrusive and Conventional RB Models for Large-scale Structural Problems

Zhenying Zhang, Mengwu Guo, and Jan S. Hesthaven, École Polytechnique Fédérale de Lausanne, Switzerland

3:30-3:50 Reducing Parameter Space for Neural Network Training

Tong Qin, Ling Zhou, and Dongbin Xiu, Ohio State University, U.S.

Tuesday, February 26

MS144

Scalable Parallel Multigrid Solvers - Part II of II

2:15 p.m.-3:55 p.m.

Room: 300A

For Part 1 see MS111

Multigrid methods are a common tool for solving large-scale sparse linear systems, particularly those arising from discretized partial differential equations. However, parallel multigrid methods suffer from reduced scalability as a result of large costs associated with communication. A number of algorithmic and implementation optimizations reduce communication requirements and yield improved scalability. This session provides a variety of talks on state-of-the-art multigrid solvers and methods for improved scalability of parallel implementations with application that extend beyond the multigrid setting.

Organizer: Amanda Bienz

University of Illinois at Urbana-Champaign, U.S.

2:15-2:35 Multigrid Solver Integrated Tree-Based Adaptive Refinement ({SITAR}) in {PETS}c

Mark Adams, Lawrence Berkeley National Laboratory, U.S.; Matthew G. Knepley, State University of New York at Buffalo, U.S.; Tobin Isaac, Georgia Institute of Technology, U.S.; Johann Rudi, Argonne National Laboratory, U.S.

2:40-3:00 Parallel Performance of Algebraic Multigrid with Domain Decomposition

Wayne Mitchell, Heidelberg University, Germany

3:05-3:25 Multilevel Convergence Theory for Multigrid Reduction in Time (MGRIT)

Robert D. Falgout, Lawrence Livermore National Laboratory, U.S.; *Andreas Hessenthaler*, University of Stuttgart, Germany; Oliver Röhrle, Institut für Mechanik (Bauwesen), Germany; Jacob B. Schroder, University of New Mexico, U.S.; Ben Southworth, University of Colorado Boulder, U.S.

3:30-3:50 Performance of Parallel Approximate Ideal Restriction Multigrid (pair) with Applications in Transport

Joshua Hanophy, Texas A&M University, U.S.; Ben Southworth, University of Colorado Boulder, U.S.; Ruipeng Li, Lawrence Livermore National Laboratory, U.S.

continued in next column

Tuesday, February 26

MS145

Recent Advances in Model Reduction and Uncertainty Quantification - Part II of II

2:15 p.m.-3:55 p.m.

Room: 300B

For Part 1 see MS112

Despite unprecedented computational power, simulating high-dimensional numerical problems or large-scale models describing complex physical phenomena may require weeks on supercomputers. Reducing computational costs of these models while preserving high fidelity is essential for handling onerous many-query applications in computational science and engineering. This minisymposium will focus on recent mathematical and computational advances in reduced order modeling and its application in uncertainty quantification. Topics will include novel enhancements to the state-of-the-art reduced order modeling techniques in the linear and nonlinear settings, development in low-rank approximation methods for stochastic problems, etc. The speakers in this minisymposium will discuss work towards broader model order reduction and data-driven approaches that are necessary to support a wide range of physical, biological, and aerospace applications.

Organizer: Jiahua Jiang

Virginia Tech, U.S.

Organizer: Kookjin Lee

Sandia National Laboratories, U.S.

2:15-2:35 A Physics-informed Gaussian Process Regression Method for Data Model Convergence

Xiu Yang, Alexandre M. Tartakovsky, and Guzel Tartakovsky, Pacific Northwest National Laboratory, U.S.

2:40-3:00 A Stein Variational Newton Method for Optimal Experiment Design Problems

Keyi Wu, Peng Chen, and Omar Ghattas, University of Texas at Austin, U.S.

3:05-3:25 A Riemannian View on Active Subspaces

Zach Grey and Paul Constantine, University of Colorado Boulder, U.S.

3:30-3:50 Reduced Models for Uncertainty Quantification in the Cardiovascular Network via Domain Decomposition

Sofia Guzzetti, Emory University, U.S.; Alonso M. Alvarez and Pablo J. Blanco, Laboratorio Nacional de Computacao Cientifica, Brazil; Kevin T. Carlberg, Sandia National Laboratories, U.S.; Alessandro Veneziani, Emory University, U.S.

continued in next column

Tuesday, February 26

MS146

Recent Advances in Computational Methods for High Dimensional Bayesian Inversion - Part II of II

2:15 p.m.-3:55 p.m.

Room: 300C

For Part 1 see MS113

In many practical Bayesian inverse problems, the parameters of interest are of very high dimensions. On the other hand, the relation between the parameters and the data is often described by computationally intensive mathematical models. Thus, developing efficient and accurate algorithms for such high dimensional problems poses a big challenge computationally. The purpose of this minisymposium is to bring researchers from different fields to discuss the recent advances in computational methods for such problems.

Organizer: Qifeng Liao

ShanghaiTech University, China

Organizer: Jinglai Li

Shanghai Jiao Tong University, China

2:15-2:35 Trimmed Ensemble Kalman Filter for Nonlinear and Non-Gaussian Data Assimilation Problems

Guang Lin, Purdue University, U.S.

2:40-3:00 An Adaptive Reduced Basis Anova Method for High-dimensional Bayesian Inverse Problems

Qifeng Liao, ShanghaiTech University, China; Jinglai Li, Shanghai Jiao Tong University, China

3:05-3:25 An Efficient Multiscale Finite Element Method for Multiscale Elliptic PDEs with Random Coefficients

Zhiwen Zhang, University of Hong Kong, Hong Kong; Eric Chung, The Chinese University of Hong Kong, Hong Kong; Yalchin Efendiev, Texas A&M University, U.S.; Wing Tat Leung, University of Texas at Austin, U.S.

3:30-3:50 Uncertainty Quantification with non-Gaussian Correlated Uncertainties, and Applications to Electronic and Photonic Circuits

Chunfeng Cui and *Zheng Zhang*, University of California, Santa Barbara, U.S.

Tuesday, February 26

MS147

AWM Workshop Panel: Perspectives and Advice from Women in Research

2:15 p.m.-3:55 p.m.

Room: 300D

Leading women researchers from academia and industry will discuss their career paths, including challenges they have faced. Most of the panel time will be dedicated to questions from the audience and informal discussion about the issues raised. We welcome anyone interested in hearing about women's experiences navigating a research career.

Organizer: Malena I. Espanol
University of Akron, U.S.

See online program for an update on this session.

Tuesday, February 26

MS148

BE: Broader Engagement Technical Research - Part II of II

2:15 p.m.-3:55 p.m.

Room: 302A

For Part 1 see MS114

The Broader Engagement (BE) minisymposium brings together faculty, researchers, and students from a variety of CSE19 conference theme disciplines. The BE program offers a dynamic program aimed at promoting diversity and inclusion through travel grants to SIAM CSE conferences and opportunities for the CSE community to learn about and support diversity and inclusion. This minisymposium features presentations of research from BE funded participants and faculty funded through the Sustainable Research Pathways program organized by the Sustainable Horizons Institute and funded by Lawrence Berkeley National Laboratory.

Organizer: Mary Ann E. Leung
Sustainable Horizons Institute, U.S.

Organizer: Nitin Sukhija
Slippery Rock University, U.S.

2:15-2:35 Computational Tools for Adaptive Resolution Molecular Dynamics Simulation: Application to Biological Molecules

Masa Watanabe, Branden Sanchez, and Zachary Douglas, University of Saint Mary, Kansas, U.S.; Yu-Hang Tang and Bert de Jong, Lawrence Berkeley National Laboratory, U.S.; Silvia N. Crivelli, Lawrence Berkeley National Laboratory and University of California, Davis, U.S.

2:40-3:00 A Deep Learning Approach to Protein Structure Prediction

Wilson Rivera, University of Puerto Rico, Mayaguez, Puerto Rico; Dan Rosa de Jesus and Julian Cuevas, Lawrence Berkeley National Laboratory, U.S. and University of Puerto Rico at Mayaguez, P.R.; Silvia N. Crivelli, Lawrence Berkeley National Laboratory and University of California, Davis, U.S.

3:05-3:25 Machine Learning Enabled Suicide Prevention Research using ICU Patient Data

Rafael Zamora, Shirley Wang, Cheng Ding, Shahzeb Khan, and Ryan Kingery, Lawrence Berkeley National Laboratory, U.S.; *Xinlian Liu*, Hood College, U.S.; Silvia N. Crivelli, Lawrence Berkeley National Laboratory and University of California, Davis, U.S.

3:30-3:50 Learning From Protein Structures using Graph Convolutional Networks

Rafael Zamora-Resendiz, Hood College and Lawrence Berkeley National Laboratory, U.S.

continued in next column

Tuesday, February 26

MS149

BE: Social Modeling through Game Design, A STEAM Workshop

2:15 p.m.-3:55 p.m.

Room: 302B

Organizer: Mary Ann E. Leung
Sustainable Horizons Institute, U.S.

Tuesday, February 26

MS150

Advances in Unstructured Mesh Algorithms and their Applications - Part II of II

2:15 p.m.-3:55 p.m.

Room: 303A

For Part 1 see MS116

Unstructured meshes continue to be utilized in many computational science and engineering simulations. Recently, advancements have been made to unstructured mesh algorithms in several areas including optimization-based mesh quality improvement, hybrid volume meshing, multi-material remapping, and moving meshes to list just a few. In addition, meshes have been used in interesting applications involving fluid-structure interaction, and many others. This minisymposium will feature presentations on some of these advancements in meshing algorithms and their applications.

Organizer: Mike Stees
University of Kansas, U.S.

Organizer: Suzanne M. Shontz
University of Kansas, U.S.

2:15-2:35 High-order Mesh Untangling Based on Angles

Mike Stees and Suzanne M. Shontz, University of Kansas, U.S.

2:40-3:00 Simulation-driven Optimization of High-order Meshes by the Target-matrix Optimization Paradigm

Veselin Dobrev, Lawrence Livermore National Laboratory, U.S.; Patrick Knupp, Dihedral LLC, U.S.; Tzanio Kolev, Lawrence Livermore National Laboratory, U.S.; Ketan Mittal, University of Illinois at Urbana-Champaign, U.S.; *Vladimir Tomov*, Lawrence Livermore National Laboratory, U.S.

3:05-3:25 Optimization-based PDE-constrained Discontinuity Tracking with High-order Curved Unstructured Meshes

Andrew Shi and Per-Olof Persson, University of California, Berkeley, U.S.

3:30-3:50 An Efficient Parallel Algorithm for Multiphysics Simulations on 3D Unstructured Meshes

Fande Kong, Derek R. Gaston, John W. Peterson, Cody J. Permann, Andrew Slaughter, Alex Lindsay, and Richard Martineau, Idaho National Laboratory, U.S.

continued in next column

Tuesday, February 26

MS151

Batched BLAS: API Standardization, Libraries, and Applications - Part II of II

2:15 p.m.-3:55 p.m.

Room: 303B

For Part 1 see MS117

The Basic Linear Algebra Subprograms (BLAS) are the most widely accepted standard in high performance dense matrix computation. The unified BLAS API, which is implemented by many vendors and research groups, enables performance portability for many applications across different architectures and platforms. The past few years have witnessed a continuously growing interest in optimizing BLAS for a batch of small independent problems, hence the name "Batched BLAS". Such interest is driven by numerous applications, including tensor contractions, sparse solvers, astrophysics, quantum chemistry, and many others. This minisymposium (MS) covers a wide range of ongoing research activities about the Batched BLAS. The MS features three main categories of research. The first one covers standardization of the Batched BLAS API, as well as the next-generation BLAS to support reproducibility and extended precisions. The second category features many libraries, from different vendors and research groups, that provide optimized Batched BLAS routines on different hardware architectures. The third category highlights several scientific applications where optimized Batched BLAS has a great impact as a critical building block. At the end of the MS, the audience will be aware of the latest developments in Batched BLAS. The MS also provides an excellent collaboration opportunity between the Batched BLAS research community and potential application developers.

Organizer: Ahmad Abdelfattah

University of Tennessee, Knoxville, U.S.

Organizer: Stanimire Tomov

University of Tennessee, Knoxville, U.S.

2:15-2:35 A Proposal for Next-generation BLAS

Mark Gates, University of Tennessee, U.S.; James W. Demmel, University of California, Berkeley, U.S.; Greg Henry, Intel Corporation, U.S.; Xiaoye S. Li, Lawrence Berkeley National Laboratory, U.S.; Jason Riedy, Georgia Institute of Technology, U.S.; Peter Tang, Intel Corporation, U.S.

2:40-3:00 Adventures in Batched Linear Algebra in Intel{\ textregistered} Math Kernel Library

Sarah Knepper, Intel Corporation, U.S.; Peter A. Caday, Rice University, U.S.; Kazushige Goto, Louise Huot, Mesut Meterelliyoz, Arthur Mitrano, and Shane Story, Intel Corporation, U.S.

3:05-3:25 Batched Linear Algebra in Kokkos Kernels

Siva Rajamanickam, Sandia National Laboratories, U.S.

3:30-3:50 AcroTensor: Flexible Tensor Contractions on GPUs

Aaron Fisher and Tzanio Kolev, Lawrence Livermore National Laboratory, U.S.; Johann Dahm, IBM Research, U.S.

continued in next column

Tuesday, February 26

MS152

Hydrodynamics at Small Scales: Fluctuating Hydrodynamics - Part II of II

2:15 p.m.-3:55 p.m.

Room: 201B

For Part 1 see MS118
Featured Minisymposium

With the increased interest in nano- and micro-fluidics, as well as biological systems, it has become necessary to develop tools for hydrodynamic calculations at microscopic and mesoscopic scales. This minisymposium will focus on advances in multiscale numerical methods for simulating flows at mesoscopic scales. Coarse-grained models cover a broad range of time and length scales by incrementally sacrificing physical fidelity for computational efficiency. Of particular interest will be fluctuating hydrodynamics of complex fluids such as reactive mixtures, colloidal passive and active suspensions, and multi-phase fluids. Issues to be discussed will include the inclusion of thermal fluctuations in analytical and computational models, as well as applications in the physical sciences, biology, and engineering.

Organizer: Aleksandar Donev

Courant Institute of Mathematical Sciences, New York University, U.S.

2:15-2:35 Hydrodynamic Fluctuations in Quasi-two Dimensional Diffusion

Aleksandar Donev, Courant Institute of Mathematical Sciences, New York University, U.S.; *R. Delgado-Buscalioni*, Universidad Autonoma de Madrid, Spain

2:40-3:00 Fluctuating Hydrodynamics of Janus Particles Assembled via Long-range Attractive Hydrophobic Interactions

Yuan-Nan Young, New Jersey Institute of Technology, U.S.

3:05-3:25 Towards an Irreversible Thermodynamics for Active Systems

Katherine Klymko, Lawrence Berkeley National Laboratory and University of California Berkeley, U.S.

continued on next page

Tuesday, February 26

MS152

Hydrodynamics at Small Scales: Fluctuating Hydrodynamics - Part II of II

continued

3:30-3:50 Effects of Subscale Fluctuations on Predictions of Hot Spot Formation in Energetic Materials
Joseph Bakarji, Stanford University, U.S.

Tuesday, February 26

MS153

Model Reduction for Problems with Strong Convection, Sharp Gradients, and Discontinuities - Part II of II

2:15 p.m.-3:55 p.m.

Room: 201C

For Part 1 see MS119
Featured Minisymposium

Model reduction exploits that many phenomena of interest in science and engineering can be approximated well by problem adapted low-dimensional subspaces of the high-dimensional solution spaces. Reduced bases are a prominent example, but also proper orthogonal decomposition and polynomial chaos can be seen from this perspective. However, if coherent structures (e.g., wave fronts or shocks) travel through the physical domain, then traditional model reduction techniques typically fail because these structures introduce high-dimensional features. This minisymposium highlights recent advances in model reduction methods that recover low-dimensionality in convection-dominated problems via transformations, adaptive bases, separation of transport, and other techniques.

Organizer: Maciej Balajewicz
University of Illinois at Urbana-Champaign, U.S.

Organizer: Benjamin Peherstorfer
Courant Institute of Mathematical Sciences, New York University, U.S.

Organizer: Gerrit Welper
University of Central Florida, U.S.

2:15-2:35 Model Reduction of Multi-dimensional Hyperbolic Conservation Laws
Donsub Rim and Kyle T. Mandli, Columbia University, U.S.

2:40-3:00 Transported Snapshot Model Order Reduction for Parametric, Steady-state Fluid Flows Containing Parameter Dependent Shocks
Nirmal Nair and Maciej Balajewicz, University of Illinois at Urbana-Champaign, U.S.

3:05-3:25 Sparsity Promoting Acceleration for Reduced Order Modeling of Multi-scale Transport Problems
Jiayang Xu, Cheng Huang, and Karthik Duraisamy, University of Michigan, Ann Arbor, U.S.

3:30-3:50 Interpolation of Solutions of Parametric Hyperbolic PDEs with Changing Shock Structure
Gerrit Welper, University of Central Florida, U.S.

continued in next column

Tuesday, February 26

MS154

Structure Preserving Techniques for Hyperbolic Systems - Part II of II

2:15 p.m.-3:55 p.m.

Room: 202A

For Part 1 see MS120

The generalization of the maximum principle to hyperbolic systems is the notion of convex invariant domains. The key property defining an invariant domain is that for any initial data in the domain in question, the (entropy satisfying) solution to the hyperbolic system remains in the domain. For instance positive density and positive internal energy are invariant properties for the compressible Euler equations; likewise, positive water height is an invariant property of the shallow water equations. Invariant domain preserving numerical methods are those that preserve invariant domains. Many first-order methods are known to be invariant-domain preserving, but extending theses techniques to higher-order is non-trivial. The difficult balance consists of being invariant, high-order accurate and to satisfy enough entropy inequalities for the numerical solution to converge to an entropy solution. These three constraints are difficult to achieve at the same time, but are essential to get robust and accurate methods. Invariant-domain preserving are important in any application where being 'in bounds' is essential. The objective of minisymposium is to gather specialists to present the latest developments on the above topic and share new ideas. The minisymposium will particularly focus on the approximation of hyperbolic systems.

Organizer: Jean-Luc Guermond

Texas A&M University, U.S.

2:15-2:35 Convex Limiting

Jean-Luc Guermond, Ignacio Thomas, and Bojan Popov, Texas A&M University, U.S.; Murtazo Nazarov, Uppsala University, Sweden

2:40-3:00 Asymptotic Preserving Schemes for Euler Equations of Gas Dynamics and Shallow Water Equations with Coriolis Forces

Alexander Kurganov, Tulane University, U.S.

3:05-3:25 A New Approach for Designing Moving-water Equilibria Preserving Schemes for the Shallow Water Equations

Alina Chertock, North Carolina State University, U.S.

3:30-3:50 Invariant-region-preserving DG Methods for Multi-dimensional Hyperbolic Conservation Law Systems, with an Application to Compressible Euler Equations

Hailiang Liu, Iowa State University, U.S.

continued in next column

Tuesday, February 26

MS155

Theoretical and Computational Advancements in Ice-Sheet Modeling - Part II of II

2:15 p.m.-3:55 p.m.

Room: 202B

For Part 1 see MS121

The demand for more accurate answers to questions regarding the future of the Earth's climate is a driving force for the development of Earth System Models (ESMs). Such questions include, for instance, assessment of climate sensitivity, and quantification of Sea Level Rise (SLR). To answer these questions, ESMs are pushing toward higher resolution, both in computational terms (the number of degrees of freedom), and in terms of the physical processes included in the model. When it comes to quantifying SLR, accurate and robust models for ice sheets dynamics are key to achieve reliable predictions. In this respect, advancements in both the computational and modeling aspects are equally important. In fact, discretization of the governing equations of ice-sheet dynamics usually lead to difficult mathematical problems, which require robust and efficient solvers. On the other hand, an accurate modeling of ice-sheets movement can involve physical processes at completely different time and space scales, and may involve physical quantities that are hard (if possible at all) to measure directly. This minisymposium will feature presentations aimed at improving the efficiency and reliability of ice-sheet components of ESMs. Topics of interest are: algorithmic advancements, next-generation HPC software development, model parametrizations and model extensions, multiphysics couplings, data assimilation and uncertainty quantification.

Organizer: Luca Bertagna

Sandia National Laboratories, U.S.

Organizer: Mauro Perego

Sandia National Laboratories, U.S.

continued on next page

Tuesday, February 26

MS155

Theoretical and Computational Advancements in Ice-Sheet Modeling - Part II of II

continued

Organizer: Jerry Watkins
Sandia National Laboratories, U.S.

Organizer: Irina K. Tezaur
Sandia National Laboratories, U.S.

Organizer: Daniel Martin
Lawrence Berkeley National Laboratory, U.S.

2:15-2:35 The Case for Large Ice Sheet Model Ensembles from Theory and Practice

Alexander Robel and Helene Seroussi, California Institute of Technology, U.S.; Gerard Roe, University of Washington, U.S.

2:40-3:00 Large Ensemble Modelling of Greenland's Contribution to Sea Level Rise

Andy Aschwanden, Mark Fahnestock, and Martin Truffer, University of Alaska, Fairbanks, U.S.; *Douglas Brinkerhoff*, University of Montana, U.S.; Regine Hock and Constantine Khroulev, University of Alaska, Fairbanks, U.S.; Ruth Mottram, Danish Meteorological Institute, Denmark; Abbas Khan, DTU Space, Denmark

3:05-3:25 Reconciling Observations of Driving Stress with Observations of Velocity

Jesse Johnson, University of Montana, U.S.

3:30-3:50 Modeling Thermally Activated Sliding in Ice Sheet Flow

Elisa Mantelli, Stanford University, U.S.; Christian Schoof, University of British Columbia, Canada

Tuesday, February 26

MS156

High-order Solvers for Wave Problems - Part II of II

2:15 p.m.-3:55 p.m.

Room: 202C

For Part 1 see MS122

Significant progress in the computational simulation of waves has been made in recent years from consideration of efficient high-order wave solvers, including finite-difference approaches,Fourier-based methods, integral equation methods, Galerkin differences, and other high-order techniques. This minisymposium considers ideas that include the extension of Fourier methods to non-periodic problems, new time-stepping strategies beyond the standard method-of-lines, sparse matrix inversion, Galerkin difference approximations, high-order summation by parts schemes for seismic wave propagation and the use of robust upwind schemes for wave equations on overlapping structured grids. The resulting methods have provided important new mathematical tools applicable to a wide range of areas of science and engineering---including electronics, photonics,meta-materials, microscopy, remote sensing, imaging and communications, among many others: in some cases, the resulting techniques have provided solutions to previously intractable problems. This session will focus on several important topics that have recently attracted particular attention in the field, including modern finite-difference approaches, frequency-time hybrids, Fourier-based methods for general domains, fast high-order integral-equation methods, and Galerkin difference approximations. This minisymposium includes important recent contributions in the area and will help to disseminate and advance the field.

Organizer: Oscar P. Bruno
California Institute of Technology, U.S.

Organizer: William D. Henshaw
Rensselaer Polytechnic Institute, U.S.

2:15-2:35 Galerkin Difference Approximations for Biharmonic Equations

John Jacangelo, Rensselaer Polytechnic Institute, U.S.

2:40-3:00 Fast, Higher-order Direct/ Iterative Hybrid Solver for Scattering by Inhomogeneous Media – with Application to High-frequency and Discontinuous Refractivity Problems

Ambuj Pandey and OSCAR P Bruno, California Institute of Technology, U.S.

3:05-3:25 Mesh Refinement Interfaces with Hanging Nodes for the Elastic Wave Equation

N. Anders Petersson, Lawrence Livermore National Laboratory, U.S.; Siyang Wang, Chalmers University of Technology, Sweden, and University of Gothenburg, Sweden; Bjorn Sjogreen, Lawrence Livermore National Laboratory, U.S.

3:30-3:50 High Order Modeling of Visco-elastic Attenuation in Seismic Wave Computation

Bjorn Sjogreen and N. Anders Petersson, Lawrence Livermore National Laboratory, U.S.

continued in next column

Tuesday, February 26

MS157

Application of Fractional Calculus in Material Science and Engineering - Part II of II

2:15 p.m.-3:55 p.m.

Room: 203

For Part 1 see MS123

While the mathematics and computations of fractional calculus have matured in recent years, limited efforts have focused on material science and engineering applications. Fractional calculus has been shown as an effective operator in fractal media, viscoelasticity, and control designs. While the fractional integral has been used to describe the fractal structure of materials which leads to new thermodynamic relations, the fractional derivative could be used to describe viscoelasticity properties of materials, not only does the calibration of fractional models of viscoelasticity with experimental measurements provides less error compare to the integer one, but it also provides more information about the structure of a material. This area opens up an application of fractional calculus which may describe the multiscale thermomechanical material behavior of many polymers. On the other hand, fractional calculus has been used in the formulation of variational calculus, Euler-Lagrange equations and optimal control systems where the fractional derivative has been replaced with the integer derivative to create a new set of necessary conditions that must be satisfied by an optimal control law and its associated state-control equations. This minisymposium is designed to bring together the research scholars who are working in the application of fractional calculus in material science and engineering to present their results and exchange ideas to propose future research in this field.

Organizer: Somayeh Mashayekhi
Florida State University, U.S.

Organizer: William Oates
Florida State University, U.S.

2:15-2:35 Numerical Techniques to Approximate Fractional-order Nonlinear Viscoelasticity in Soft Elastomers

Paul Miles, Graham Pash, and *Ralph Smith*, North Carolina State University, U.S.; William Oates, Florida State University, U.S.

2:40-3:00 Fractional Modeling of Visco-elasto-plastic Materials

Jorge L. Suzuki and Mohsen Zayernouri, Michigan State University, U.S.

3:05-3:25 Nonlocal Operators in Dynamics and Control: Theories and Applications

Arman Dabiri, University of Arizona, U.S.

3:30-3:50 The Zolotarev Fractional Derivative: Modeling Anomalous Diffusion and Dissipative Wave Propagation

James F. Kelly, Mark Meerschaert, and Robert J. McGough2, Michigan State University, U.S.

continued in next column

Tuesday, February 26

MS158

Multiphysics Simulation with MOOSE - Part II of II

2:15 p.m.-3:55 p.m.

Room: 205

For Part 1 see MS124

The Multiphysics Object Oriented Simulation Environment (MOOSE) framework is Idaho National Laboratory's (INL) premier open source modeling and simulation tool. Under development since 2008, and available on GitHub since 2014, the framework provides a flexible, powerful, and analyst-friendly computational environment for the numerical solution of partial differential equations via the finite element method. The framework is written in C++, supports parallel 1, 2, and 3D implicit and explicit finite element simulations, and makes extensive use of the finite element building blocks provided by the libMesh finite element library, the solvers and preconditioners of PETSc, and the parallel communication routines of MPI. A major emphasis of the framework is facilitating multiphysics simulations by lowering the barriers of entry which inevitably arise when coupling different codes together, and codifying these coupling methods into a consistent programming interface. This minisymposium aims to bring together MOOSE users and developers to discuss the latest developments within the framework, as well as advances in both public and internally-developed MOOSE-based applications. Application studies are welcome from all areas of computational science, with a particular emphasis on multiphysics simulation, including algorithmic explorations as well as scientific and engineering studies enabled by MOOSE.

Organizer: John W. Peterson
Idaho National Laboratory, U.S.

Organizer: Cody J. Permann
Idaho National Laboratory, U.S.

continued on next page

Tuesday, February 26

MS158

Multiphysics Simulation with MOOSE - Part II of II

continued

2:15-2:35 Implementation of a State-based Peridynamic Analysis in MOOSE Framework

Deepak K. Behera, University of Arizona, U.S.; Hailong Chen, Idaho National Laboratory and University of Kentucky, U.S.; Benjamin Spencer, Idaho National Laboratory, U.S.; Erdogan Madenci, University of Arizona, U.S.

2:40-3:00 Multiphysics Analysis of Pebble Bed Reactors in MOOSE

Sebastian Schunert and Daniel Schwen, Idaho National Laboratory, U.S.

3:05-3:25 Parallel Phasefield Simulation with MOOSE and Marmot

Daniel Schwen and Larry K. Aagesen, Idaho National Laboratory, U.S.

3:30-3:50 Simulation of Molten Salt Reactors with Moltres

Andrei Rykhlevskii, University of Illinois at Urbana-Champaign, U.S.; Alexander Lindsay, Idaho National Laboratory, U.S.; Kathryn Huff, University of Illinois at Urbana-Champaign, U.S.

Tuesday, February 26

MS159

Data Assimilation in Fluid Models

2:15 p.m.-3:55 p.m.

Room: 206A

This minisymposium groups talks related to data assimilation and applications to fluid problems requiring fast solutions to complex computations. Multiple approaches to working with data assimilation problems is made through methods such as graph network reduction and machine learning. These networks are used to model underground transport of contaminant particles to predict future concentrations and flows underground. We explore using graph networks to find quantities of interest and the inversion problem central to relating these quantities back to physical system properties. In a related application, the Kalman filter (EnKF) is implemented to to determine the key model parameters that have the most impact on material deformation and failure fracture propagation. The application of the Azouani-Olson-Titi (AOT) algorithm is used to examine how perturbations of fluid viscosity affect the convergence of the AOT algorithm to the 2D incompressible Navier-Stokes equations. A new AOT algorithm with continuous data assimilation is based on feedback control from a PDE and weighted least square interpolation is in development. Testing for this algorithm is being done on multiple PDEs, including the 1D Kuramoto-Sivashinsky equation and 2D shallow water equations.

Organizer: Jaime A. Lopez-Merizalde

Tulane University, U.S.

2:15-2:35 Model Reduction using Graph Networks

Jaime A. Lopez-Merizalde, Tulane University, U.S.

2:40-3:00 Estimating Key Parameters for a Fracture Propagation Model using Data Assimilation

Humberto C. Godinez, Los Alamos National Laboratory, U.S.

3:05-3:25 An Analysis of Parameter Sensitivity in Continuous Data Assimilation of 2D Incompressible Navier-Stokes Equations

Elizabeth Carlson, University of Nebraska, U.S.

3:30-3:50 Continuous Data Assimilation from Scattered Spatial Observations in Time Dependent PDEs

Tong Wu and James Hyman, Tulane University, U.S.; Humberto C. Godinez and Vitaliy Gyrya, Los Alamos National Laboratory, U.S.

continued in next column

Tuesday, February 26

MS160

Data Assimilation and Optimization in Physiologic Modeling - Part II of II

2:15 p.m.-3:55 p.m.

Room: 206B

For Part 1 see MS126

Physics-based modeling for biological systems has been advanced over the past decade and plays an increasingly important role in understanding the underlying physiological mechanism and assisting clinical management of associated diseases. Although increased model complexity can better reveal detailed information of the biological system, it also poses great difficulties on how to calibrate the model to be patient-specific. Clinical measurement data (e.g., medical imaging) can be utilized to perform model calibration and parameter/hidden state estimation, which can be formularized as inverse problems. Therefore, proper optimization tools and data assimilation techniques are needed. However, solving these inverse problems is challenging due to large inter-patient variability and increased computational complexity. Moreover, large uncertainties from both clinical data and computational model also pose more difficulties. This minisymposium hopes to address the aforementioned challenges by reviewing and discussing recent advances in developing efficient optimization/data assimilation approaches and quantifying associated uncertainties for inverse problems in biomedical applications.

Organizer: Jian-Xun Wang

University of Notre Dame, U.S.

Organizer: Shawn Shadden

University of California, Berkeley, U.S.

2:15-2:35 Adding Constraints to Bayesian Inverse Problems

Jiacheng Wu, University of California, Berkeley, U.S.; Jian-Xun Wang, University of Notre Dame, U.S.; Shawn Shadden, University of California, Berkeley, U.S.

2:40-3:00 Data Assimilation on Lumped Parameter Models for Diastolic Heart Failure

Karlyn Harrod, University of Notre Dame, U.S.; Alison Marsden, Stanford University, U.S.; Daniele E. Schiavazzi, University of Notre Dame, U.S.

3:05-3:25 Multilevel Multifidelity Approaches for Uncertainty Quantification in Cardiovascular Modeling

Casey M. Fleeter, Stanford University, U.S.; Gianluca Geraci, Sandia National Laboratories, U.S.; Daniele E. Schiavazzi, University of Notre Dame, U.S.; Andrew Kahn, University of California, San Diego, U.S.; Alison Marsden, Stanford University, U.S.

3:30-3:50 Shape Based Effective Estimation for Cardiovascular Models

Philippe Moireau, Inria, France

continued in next column

Tuesday, February 26

MS161

Design and Usability of High-performance PDE Software Engines and Frameworks - Part II of II

2:15 p.m.-3:55 p.m.

Room: 206C

For Part 1 see MS127
Featured Minisymposium

The combined complexity of the application areas being studied, the applied numerical methods, and the requirements of parallel hardware make the development of high performance PDE engines and frameworks a challenging area of research. In this minisymposium we bring together users and developers of several PDE software packages. We will discuss the requirements and expectations that users have of production-ready simulation software as well as the design decisions that go into the implementation of such software.

Organizer: Anne Reinarz

Technische Universität München, Germany

Organizer: David Ham

Imperial College London, United Kingdom

Organizer: Tobias Weinzierl

Durham University, United Kingdom

Organizer: Michael Bader

Technische Universität München, Germany

2:15-2:35 Supporting Advanced Finite Element Methods in the Deal.II library

Wolfgang Bangerth, Colorado State University, U.S.

2:40-3:00 The SUNDIALS Suite of Time Integrators and Nonlinear Solvers: Preparing for Exascale Computing

David J. Gardner and Carol S. Woodward, Lawrence Livermore National Laboratory, U.S.; Daniel R. Reynolds, Southern Methodist University, U.S.; Alan Hindmarsh, Slaven Peles, and Cody J. Balos, Lawrence Livermore National Laboratory, U.S.

continued on next page

Tuesday, February 26

MS161

Design and Usability of High-performance PDE Software Engines and Frameworks - Part II of II

continued

3:05-3:25 Balancing the Numerical and Parallel Performance for Reservoir Simulations

Andreas Thune and *Xing Cai*, Simula Research Laboratory, Norway

3:30-3:50 Walberla: A General Purpose Software Framework for Massively Parallel Simulations

Christoph Rettinger, Martin Bauer, Sebastian Eibl, Christian Godenschwager, Florian Schornbaum, Harald Köstler, and Ulrich Rüde, University of Erlangen-Nuremberg, Germany

Tuesday, February 26

MS162

Exascale Applications with High-Order Methods - Part II of II

2:15 p.m.-3:55 p.m.

Room: 206D

For Part 1 see MS128

High-order discretizations have the potential to provide an optimal strategy for achieving high performance and delivering fast, efficient, and accurate simulations on next-generation architectures. This minisymposium will discuss the next-generation high-order discretization algorithm, based on finite/spectral element approaches that will enable a wide range of important scientific applications to run efficiently on future architecture.

Organizer: MiSun Min

Argonne National Laboratory, U.S.

Organizer: Paul Fischer

University of Illinois at Urbana-Champaign, U.S.

2:15-2:35 High-order Steady State Solvers for Multiphysics Applications

MiSun Min, Argonne National Laboratory, U.S.

2:40-3:00 Geometry, Meshing and Mesh Adaptation of High-order Curved Meshes

Morteza Hakimi, Kazem Kamran, *Onkar Sahni*, and Mark S. Shephard, Rensselaer Polytechnic Institute, U.S.; Blair Downie, Rocco Nastasia, Saurabh Tendulkar, and Mark Beall, Simmetrix Inc, U.S.

3:05-3:25 Low-memory Implicit Solvers for the Sparse Line-DG Method with Kronecker-SVD Preconditioning

Per-Olof Persson and Will Pazner, University of California, Berkeley, U.S.

3:30-3:50 Recent Developments in Spectral Element Methods for Turbulent Flows

Paul Fischer, University of Illinois at Urbana-Champaign, U.S.

Tuesday, February 26

MS163

Advances In Nonlinearly Stable Methods

2:15 p.m.-3:30 p.m.

Room: 207

High-order nonlinearly stable methods have recently attracted much attention because of their ability to provide stronger stability estimates for numerical solutions of nonlinear partial differential equations, such as the Navier-Stokes equations, MHD equations, shallow water equations, etc. The main objective of this minisymposium is to bring together experts in nonlinearly stable methods to discuss innovative approaches, unsolved problems, and future directions for constructing high-order methods that mimic key stability properties of the governing nonlinear PDEs including entropy stability, dissipation of kinetic energy, and others.

Organizer: Nail Yamaleev

Old Dominion University, U.S.

2:15-2:35 High-order Entropy Stable Spectral Collocation Schemes for the Navier-Stokes Equations on Moving Deforming Grids

Nail Yamaleev, Old Dominion University, U.S.; David C. Del Rey Fernandez, NASA Langley Research Center and National Institute of Aerospace, U.S.; Jialin Lou, Old Dominion University, U.S.; Mark H. Carpenter, NASA Langley Research Center, U.S.

2:40-3:00 Conservative Split form Summation-by-parts Discretization of the Euler Equations on Unstructured Grids

Siavosh Shadpey, University of Toronto Institute for Aerospace Studies, Canada; David Zingg, University of Toronto, Canada

3:05-3:25 H/P Adaptive Entropy Stable Methods for the Compressible Navier-Stokes Equations: Curvilinear Grids

Mark H. Carpenter, NASA Langley Research Center, U.S.; David C. Del Rey Fernandez, NASA Langley Research Center and National Institute of Aerospace, U.S.; Matteo Parsani, King Abdullah University of Science & Technology (KAUST), Saudi Arabia; Travis Fisher, Sandia National Laboratories, U.S.; Jared Crean and Jason E. Hicken, Rensselaer Polytechnic Institute, U.S.

Tuesday, February 26

MS164

Computational Engineering (BGCE) Student Paper Prize - Part II of II

2:15 p.m.-3:55 p.m.

Room: 401A

For Part 1 see MS130

The 7th Bavarian Graduate School in Computational Engineering (BGCE) Student Paper Prize will be awarded at the 2019 SIAM CS&E Conference for outstanding student work in the field of Computational Science and Engineering. Eligible for the prize will be undergraduate and graduate students prior to receiving their PhD. Candidates are required to summarize their work in a short paper of at most 4 pages. The prize finalists will present their work in this minisymposium. The prize award announcement will be scheduled at one of the last days of the conference.

Organizer: Tobias Neckel
Technische Universität München, Germany

Organizer: Hans-Joachim Bungartz
Technische Universität München, Germany

Organizer: Dietmar Fey
Universität Erlangen-Nürnberg, Germany

Organizer: Alexander Ditter
Universität Erlangen-Nürnberg, Germany

See online program for an update on this session.

Tuesday, February 26

MS165

Homogenization and Reduced Order Modelling for Wave Equations - Part II of II

2:15 p.m.-3:55 p.m.

Room: 401B

For Part 1 see MS131

Wave propagation in highly heterogeneous media is very important in applications ranging from geophysics to complex engineered structures such as metamaterials. It is also very challenging computationally as short waves induced by the medium do not decay rapidly in space-time, and thus must be modeled, while the medium variability itself leads to dispersive effects for longer waves. Speakers in this minisymposium will discuss a variety of techniques for treating these challenging problems.

Organizer: Thomas M. Hagstrom
Southern Methodist University, U.S.

Organizer: Daniel Appelo
University of Colorado Boulder, U.S.

2:15-2:35 Efficient Density Estimation in Noisy PDEs and Uncertainty Propagation

Amir Sagiv, Gadi Fibich, and Adi Ditkowski, Tel Aviv University, Israel

2:40-3:00 Multiscale Model Reduction for Wave Equations

Eric Chung, The Chinese University of Hong Kong, Hong Kong

3:05-3:25 Rational Krylov Subspaces and Phase-preconditioning for Model Reduction of Wave Equations

Jorn Zimmerling, University of Michigan, U.S.; Rob F. Remis, Delft University of Technology, Netherlands; Vladimir Druskin, Worcester Polytechnic Institute, U.S.; Mikhail Zaslavsky, Schlumberger-Doll Research, U.S.

3:30-3:50 Hyperbolic Homogenization

Thomas M. Hagstrom, Southern Methodist University, U.S.

Tuesday, February 26

MS166

Latest Advances in Topology and Shape Optimization - Part II of II

2:15 p.m.-3:30 p.m.

Room: 401C

For Part 1 see MS132

Topology and shape optimization are design methodologies to obtain the optimal geometry of engineering structures to minimize a cost function while satisfying design constraints. These methods are well established design tools that have found industrial application in recent years to design materials with specific mechanical, thermal, acoustic properties, etc. where they can provide non-intuitive and more optimal designs than the traditional trial and error can do. However, many challenges are still unresolved as these methods are expanded to larger scale problems with more complex physical phenomena, requiring new numerical methods, optimization formulations and physical modeling. This minisymposium aims to bring together researchers whose works are intended to extend the possibilities in topology optimization towards larger problems and new physics.

Organizer: Miguel Salazar de Troya
Lawrence Livermore National Laboratory, U.S.

2:15-2:35 Topology Optimization of Structures Subject to Snapping Behaviour

Mathias Wallin, Lund University, Sweden; Daniel Tortorelli, University of Illinois at Urbana-Champaign, U.S.

2:40-3:00 A Density Gradient Approach to Topology Optimization under Design-dependent Boundary Loading

Cunfu Wang and *Xiaoping Qian*, University of Wisconsin, Madison, U.S.

3:05-3:25 Ultra High Resolution Topology Optimization: Brute Force or Smart Discretizations?

Niels Aage, Technical University of Denmark, Denmark

Tuesday, February 26

MS167

Quantitative Image Analysis - Part II of II

2:15 p.m.-3:55 p.m.

Room: 402A

For Part 1 see MS133

Image processing is ubiquitous in today's world. However, it is often the case that only a qualitative approach is needed, for example segmenting images on a smart phone app, in which aesthetics are usually valued over precision. For scientific imaging, however, quantitative techniques are necessary, and a wide variety of approaches have been developed based on the specific applications at hand. From a computational standpoint, interesting applications may involve running queries on large data sets, for instance those found in astronomy and agriculture, or inverse problems, such as those found in tomography and computer vision. This minisymposium will highlight several of the current techniques used in scientific imaging and discuss the computational challenges involved with each approach. A special focus will be on how these challenges can inform other areas of computational science and engineering.

Organizer: Laramie Paxton
Washington State University, U.S.

Organizer: Matthew Sottile
Noddle.io, U.S.

2:15-2:35 A High Performance Algorithm for X-ray Image Deblurring
Jesse Adams, Nevada National Security Site, U.S.

2:40-3:00 Recent Advances in Lifting Factorization for Discrete Wavelet Transforms
Christopher M. Brislawn, Unaffiliated

3:05-3:25 Quantifying Spatially Varying Blur in X-ray Radiographic Imaging Systems
Kevin Joyce, Nevada National Security Site, U.S.

3:30-3:50 Liver Segmentation using Graph Cuts, Tissue Intensity Means, and B-spline Interpolation
Yufeng Cao, *Laramie Paxton*, Kevin R. Vixie, and Yuan Wang, Washington State University, U.S.

Tuesday, February 26

MS168

Reproducibility in Network Algorithms - Part II of II

2:15 p.m.-3:55 p.m.

Room: 402B

For Part 1 see MS134

The growing use of network analysis in diverse critical applications, such as drug design, recommendation systems and criminal justice, has highlighted the importance of creating reproducible results. In this minisymposium, we focus on the challenges in achieving reproducibility in network analysis in the context of (i) modeling networks from raw data and (ii) in analyzing the networks, particularly dynamic ones that change with time. The first session will focus on methods for creating networks from data, and how modeling choices affect the structure of the network. In the second session, we will focus on the analysis algorithms and how variations in their design or implementation choices can change the results.

Organizer: Sanjukta Bhowmick
University of Nebraska, Omaha, U.S.

Organizer: Boyana Norris
University of Oregon, U.S.

2:15-2:35 Dynamic Graph Clustering with Variable Time Resolution
Timothy La Fond, Lawrence Livermore National Laboratory, U.S.

2:40-3:00 Parallel Dynamic Networks
Sriram Srinivasan, University of Nebraska, Omaha, U.S.

3:05-3:25 Reproducibility in Parallel Graph Algorithms
Samuel D. Pollard, University of Oregon, U.S.

3:30-3:50 Discussion on Reproducibility of Network Algorithms
Sanjukta Bhowmick, University of Nebraska, Omaha, U.S.

Tuesday, February 26

MS169

WCD Workshop - 2 of 4

2:15 p.m.-3:55 p.m.

Room:402C

For Part 1 see MS135
For Part 3 see MS203
Part of the SIAM Workshop Celebrating Diversity

Organizer: Talea Mayo
University of Central Florida, U.S.

Organizer: Shelby Wilson
Morehouse College, U.S.

Organizer: Jamol Pender
Cornell University, U.S.

See online program for an update on this session.

Tuesday, February 26

CP9

Numerical Linear Algebra I

2:15 p.m.-3:55 p.m.

Room: 102D

Chair: Brian Weston, Lawrence Livermore National Laboratory, U.S.

2:15-2:30 Multigrid Method for Incompressible Flow Simulation on Arbitrarily Complex Domain using Cartesian Mesh
Gwangsoo Go and Hyung Taek Ahn, University of Ulsan, South Korea

2:35-2:50 Multigrid Preconditioning for Space-time Distributed Optimal Control of Parabolic Equations
Mona Hajghassem, University of Baltimore, U.S.; Andrei Draganescu, University of Maryland, Baltimore County, U.S.

2:55-3:10 Effect of Matrix Reordering on the Performance of Asynchronous Preconditioners for Fluid Flow Problems
Aditya Kashi and Siva Nadarajah, McGill University, Canada

3:15-3:30 Matrices, Moments, Quadrature and PDEs
James V. Lambers, University of Southern Mississippi, U.S.

3:35-3:50 P-multigrid Block Reduction Preconditioning for a Fully-implicit High-order Reconstructed Discontinuous Galerkin Flow Solver
Brian Weston and Robert Nourgaliev, Lawrence Livermore National Laboratory, U.S.

Tuesday, February 26

CP10

Mathematical Optimization I

2:15 p.m.-3:55 p.m.

Room: 201A

Chair: Andrei Draganescu, University of Maryland, Baltimore County, U.S.

2:15-2:30 Multigrid Preconditioners for Optimization-based Domain Decomposition of Elliptic Equations
Andrei Draganescu and Michael A. Retzlaff, University of Maryland, Baltimore County, U.S.

2:35-2:50 Duality and the Level-set Method for Convex Optimization
Ron Estrin, Stanford University, U.S.; Michael P. Friedlander, University of British Columbia, Canada

2:55-3:10 Null Space Gradient Flows for Constrained Optimization with Applications to Shape Optimization
Florian Feppon, CMAP, Ecole Polytechnique, France

3:15-3:30 A Hybrid Genetic Algorithm with Local Search Approach for Laminated Composite Structures Optimization
Leila Gharsalli, Centre National d'Etudes Spatiales (CNES), France

3:35-3:50 Optimization of Programmable Elastic Materials to Produce Targeted Deformed States
Gareth W. Jones, University of Manchester, United Kingdom

Intermission

3:55 p.m.-4:10 p.m.

Poster Blitz

4:10 p.m.-4:50 p.m.

Room: Ballroom 100BC

Tuesday, February 26

PP1

General Posters

4:50 p.m.-6:50 p.m.

Room: Riverside Hall D

Design Exploration of Fuel Injectors Based on Reduced Order Modeling
Nissrine Akkari, Fabien Casenave, and Christian Rey, SAFRAN, France; Vincent R. Moureau, CORIA, France

Parallel Implementation of a Monolithic Li-ion Battery Model using Fenics
Jeffery M. Allen, Justin Chang, Francois Usseglio-Viretta, and Peter Graf, National Renewable Energy Laboratory, U.S.

An Edge-preserving Method for Joint Bayesian Inversion with Non-Gaussian Priors
Ilona Ambartsumyan, Tan Bui-Thanh, Omar Ghattas, and Eldar Khattatov, University of Texas at Austin, U.S.

CFD Simulation of Premixed Combustion in a Polygon Engine
Kevin R. Anderson and Christian Mendez, California State Polytechnic University, Pomona, U.S.

Feasibility Study of Multilevel Schur Complement Methods for Sparse Linear Systems
Gregory S. Bolet and *Joshua Booth*, Franklin & Marshall College, U.S.

Discontinuous-Galerkin Galerkin-differences for the Wave Equation in Second-order Form
Benjamin B. Buckner and Jeff Banks, Rensselaer Polytechnic Institute, U.S.; T Hagstrom and K Juhnke, Southern Methodist University, U.S.

Towards the Understanding of Multirate Schemes on Adaptive Mesh Refinement Grids
Valentin Dallerit, University of California, Merced, U.S.; John Loffeld, Lawrence Livermore National Laboratory, U.S.

A Kronecker Product Implementation of Density Matrix Renormalization Group
Eduardo F. D'Azevedo, Wael R. Elwasif, Arghya Chatterjee, and Gonzalo Alvarez, Oak Ridge National Laboratory, U.S.

continued on next page

Tuesday, February 26

PP1

General Posters

continued

Lorenz vs. "Boids"

Jorge Diaz-Castro, University of Puerto Rico, Puerto Rico

Sensitivity-driven Dimension-adaptive Sparse Stochastic Approximations in Linear Gyrokinetics

Ionut-Gabriel Farcas, Technische Universität München, Germany; Tobias Goerler, Max Planck Institute for Plasma Physics, Germany; Hans-Joachim Bungartz, Technische Universität München, Germany; Frank Jenko, Max Planck Institute for Plasma Physics, Germany; Tobias Neckel, Technische Universität München, Germany

Tissue Geometry May Govern Lung Branching Mode Selection

Uduak Z. George, San Diego State University, U.S.; Sharon Lubkin, North Carolina State University, U.S.

A Fixed Mesh Method With Immersed Finite Elements for Solving Interface Inverse Problems

Ruchi Guo and Tao Lin, Virginia Tech, U.S.; Yanping Lin, Hong Kong Polytechnic University, China

Performing Derivative-free Optimization using Active Subspaces and Random Walks

Jordan R. Hall, University of Colorado, Denver, U.S.

An Isofrequency Remapping Scheme for Harmonic Balance Methods

Andy Huang, Sandia National Laboratories, U.S.

Contour Integration and Moments for the Solution of Large Eigenproblems

Sarah Huber, Universität Wuppertal, Germany; Yasunori Futamura, University of Tsukuba, Japan; Martin Galgon and Bruno Lang, Universität Wuppertal, Germany; Tetsuya Sakurai, University of Tsukuba, Japan

3D Deep Learning in High Resolution Point Clouds using the Grid-Octree Data Structure

Ivana Jovanovic, Technische Universität München, Germany

A Two-stage Method for Spectral-Spatial Classification of Hyperspectral Images

Raymond H. Chan, Chinese University of Hong Kong, Hong Kong; *Kelvin Kan*, Emory University, U.S.; Mila Nikolova, CMLA, CNRS, ENS Cachan, France; Robert Plemmons, Wake Forest University, U.S.

Efficient Black Box Optimization for Thin Film Solar Cell Design using Transfer Learning

Mine Kaya and Shima Hajimirza, Texas A&M University, U.S.

Vector Potentials for Incompressible Flows in Bounded Domains

Matthias Kirchhart, RWTH Aachen University, Germany; Yutaka Kobayashi, Keio University, Japan

Computational Study of Non-isothermal Non-Newtonian Biofluid Flows in the Influence of Magnetic Fields for Microfluidic Applications

Ilias Konstantinou, Newcastle University, United Kingdom; Laszlo Konozsy, Cranfield University, United Kingdom

NUMO – A New Non-hydrostatic Ocean Model for Fjord Circulation and Ice-sheet/Ocean Interaction

Michal Kopera, Boise State University, U.S.; Wieslaw Maslowski and Frank Giraldo, Naval Postgraduate School, U.S.

Linking Mechanistic Infectious Disease Models to Genomic Surveillance Data Reveals Transmission Properties

Albert Lee, Institute for Disease Modeling, U.S.

A High-order Cross-platform Incompressible Navier-Stokes Solver via Artificial Compressibility with Application to Submarine Hydrodynamics

Niki A. Loppi, Freddie Witherden, and Peter E. Vincent, Imperial College London, United Kingdom

Fast Manifold Updates for Non-stationary Data Streams

Anthony Marcich and Benjamin W. Ong, Michigan Technological University, U.S.

Calibration, Propagation, and Validation of Model Discrepancy Across Experiments

Kathryn Maupin and Laura Swiler, Sandia National Laboratories, U.S.

A Low-communication Method to Solve Poisson's Equation on Locally-structured Grids

Peter Mccorquodale, Phillip Colella, and Brian Van Straalen, Lawrence Berkeley National Laboratory, U.S.; Christos Kavouklis, Lawrence Livermore National Laboratory, U.S.

Machine Learning of a Dynamical Systems Model of Aircraft Environmental Control System Pneumatics

Kirsten Meeker, US Naval Air Warfare Center, U.S.

The Theoretical and Computational Analysis of the Non-conserved Interactive Driven Diffusive Systems

Tripti Midha, Indian Institute of Technology Ropar, India

The Impact of Noise on Krylov Methods

Hannah M. Morgan, Argonne National Laboratory, U.S.; Patrick Sanan, Università della Svizzera italiana, Switzerland; Matthew G. Knepley, State University of New York at Buffalo, U.S.

Benchmarking as an Aid to Selecting the Right Hardware Architecture and Numerical Method

Benson K. Muite, University of Tartu, Estonia

Refined Isogeometric Analysis (rIGA): A Fluid Flow Application

Daniel Garcia, Basque Center for Applied Mathematics, Spain; David Pardo, University of the Basque Country, Spain; Victor Calo, King Abdullah University of Science & Technology (KAUST), Saudi Arabia; *Judit Muñoz-Matute*, University of the Basque Country, Spain

Stability of Traveling Wave Solutions of Nonlinear Conservation Laws for Image Processing

Jeungeun Park and Tong Li, University of Iowa, U.S.

LFA Lab - Flexible Local Fourier Analysis Library

Hannah Rittich, Forschungszentrum Jülich, Germany

Improved Newton Linearization for L^1-Norm-Type Minimization with Application to Viscoplastic Fluid Solvers

Johann Rudi, Argonne National Laboratory, U.S.; Georg Stadler, Courant Institute of Mathematical Sciences, New York University, U.S.; Omar Ghattas, University of Texas at Austin, U.S.

continued in next column *continued in next column* *continued on next page*

Adjoint Based H-(m)odel Adaptive Scheme for Linearised Boltzmann Equation

Neeraj Sarna and Manuel Torrilhon, RWTH Aachen University, Germany

On Composable Block Solvers and Performance Spectrum Model for the Four-Field Double Porosity/ permeability Model

Mohammad Sarraf Joshaghani, University of Houston, U.S.; Justin Chang, Rice University, U.S.; Kalyana Nakshatrala, University of Houston, U.S.; Matthew G. Knepley, State University of New York at Buffalo, U.S.

An Immersed Boundary, Fourier Pseudospectral Method for Advection-Diffusion-Reaction in Incompressible Flow

Joseph Schiff, California State University, Long Beach, U.S.

Communication-avoiding Mesh Multiplication for Large-scale Unstructured Meshes

Romulo M. Silva and Benaia Lima, Universidade Federal de Rio de Janeiro, Brazil; José Camata, COPPE/Universidade Federal do Rio e Janeiro, Brazil; Renato Elias and Alvaro Coutinho, Universidade Federal de Rio de Janeiro, Brazil

Performance Analytics for Computational Experiments (PACE)

Sarat Sreepathi, Oak Ridge National Laboratory, U.S.; Zachary Mitchell, Pellissippi State Community College, U.S.; Gaurab Kc, University of Tennessee, Knoxville, U.S.

Physics Constrained Machine Learning

Renee Swischuk, Massachusetts Institute of Technology, U.S.; Karen E. Willcox, University of Texas at Austin, U.S.

A Computationally Inexpensive Approach to Set Partitioning with Applications to Optimal Planning for Large-scale Robot Swarms

Shadi Tasdighi Kalat, Worcester Polytechnic Institute, U.S.

Algorithms for Multi-model Coupled Problems in Precice

Benjamin Uekermann, Technische Universität München, Germany; Harald Van Brummelen, Eindhoven University of Technology, Netherlands; Hans-Joachim Bungartz and Gerasimos Chourdakis, Technische Universität München, Germany; Miriam Mehl, Universität Stuttgart, Germany; Benjamin Rüth, Technische Universität München, Germany

Predicting Optimal Error Bounded Lossy Compression Configuration for Sampled Data

Robert R. Underwood, Jon Calhoun, and Amy Apon, Clemson University, U.S.

A Dynamical Systems Approach to Transforming Disparate Timescales in Data Driven Equation-free Modeling of Disease Dynamics

Mollie Van Gordon and Joshua Proctor, Institute for Disease Modeling, U.S.

A Variational Approach to Design a Numerical Scheme for N-Fluid Flow

Thibaud Vazquez-Gonzalez and Antoine Llor, CEA, France

Using Bayesian Inference to Evaluate Rare Event Probabilities

Siddhant Wahal and Biros George, University of Texas at Austin, U.S.

Locally-implicit Discontinuous Galerkin Schemes with Limiters that Guarantee Moment-Invertibility for Hyperbolic Quadrature-Based Moment Closures

Christine Wiersma and James A. Rossmanith, Iowa State University, U.S.

Physics-informed Machine Learning for Data-driven Turbulence Modeling

Jinlong Wu and Carlos Michelen-Strofer, Virginia Tech, U.S.; Jianxun Wang, University of Notre Dame, U.S.; Heng Xiao, Virginia Tech, U.S.

Fast Multipole Method for Non-oscillating Kernel Based on Cartesian Tensor and Differential Algebra

He Zhang, Thomas Jefferson National Accelerator Facility, U.S.; He Huang and Li-Shi Luo, Old Dominion University, U.S.

Tuesday, February 26

PP101

Minisymposterium: AWM Workshop

4:50 p.m.-6:50 p.m.

Room: Riverside Hall D

Malena I. Espanol, University of Akron, U.S.

Cindy Grimm, Oregon State University, U.S.

Posters in this session are part of the Association of Women in Mathematics (AWM) Workshop.

Minisymposterium: Randomized Least Squares Regression: Combining Model and Algorithm-induced Uncertainties

Jocelyn Chi and Ilse Ipsen, North Carolina State University, U.S.

Minisymposterium: Finite Difference Moving Mesh Methods for PDEs on Curved Domains

Kelsey Dipietro and Alan E. Lindsay, University of Notre Dame, U.S.

Minisymposterium: Classification of Vascular Disease Based on the Persistence Diagram using Topological Data Analysis (TDA) of Vascular Data

Megan Johnson, State University of New York at Buffalo, U.S.

Minisymposterium: A Randomized, Inexact, Newton-based Approach for Quantitative Photo-acoustic Tomography

Katrina Petroske, North Carolina State University, U.S.

Minisymposterium: Fast Algorithms for Cosmic Microwave Background Radiation Data on Healpix Points

Kathryn P. Drake and Grady B. Wright, Boise State University, U.S.

Minisymposterium: Efficient Denoising of High Resolution Color Digital Images Utilizing Krylov Subspace Spectral Methods

Eva Comino, University of Southern Mississippi, U.S.

Minisymposterium: A New Theory for Movement of Surface Meshes

Avary Kolasinski and Weizhang Huang, University of Kansas, U.S.

continued in next column

continued on next page

Tuesday, February 26

PP101

Minisymposterium: AWM Workshop

continued

Minisymposterium: Modeling the Effects of Macrophages on Bone Fracture Healing
Imelda Trejo, University of Texas at Arlington, U.S.

Minisymposterium: Fractional Derivatives and Laplacians in One and Two-sided Weighted Sobolev Spaces
Pablo Stinga and *Mary Vaughan*, Iowa State University, U.S.

Minisymposterium: Troll Hunter: Understanding Swarm Behavior in Social Networks
Chrysm Watson Ross, University of New Mexico, U.S.

Minisymposterium: AWM Selection of the Regularization Parameter in the Ambrosio - Tortorelli Approach to Image Segmentation
Yufei Yu, University of Kansas, U.S.

Minisymposterium: Computationally Efficient Fast Optimization over Time-varying Directed Graphs
Fakhteh Saadatniaki, Tufts University, U.S.

Tuesday, February 26

PP102

Minisymposterium: Broader Engagement

4:50 p.m.-6:50 p.m.

Room: Riverside Hall D

Mary Ann E. Leung, Sustainable Horizons Institute, U.S.

Minisymposterium: Joint Sequence Analysis Challenges: How to Handle Missing Values and Mixed Variable Types
Alexandra Ballow, Youngstown State University, U.S.

Minisymposterium: Deployment of Automatic Differentiation Package in R
Chitrak Banerjee, Michigan State University, U.S.

Minisymposterium: Capsule Networks for Protein Structure Classifciation and Prediction
Julian Cuevas, Lawrence Berkeley National Laboratory, U.S. and University of Puerto Rico at Mayaguez, P.R.

Minisymposterium: Finite Difference Moving Mesh for Nonlinear PDEs in Curved Domains
Kelsey Dipietro, University of Notre Dame, U.S.

Minisymposterium: Gaussian Processing for Coarse-grained Potential Development
Zachary Douglas, University of Saint Mary, Kansas, U.S.

Minisymposterium: Hybridizable Discontinuous Methods for Flow and Transport in Porous Media
Maurice Fabien, Rice University, U.S.

Minisymposterium: In Situ Analysis with Apache Spark
Aparna S. Gollakota, Loyola University Chicago, U.S.

Minisymposterium: Utilizing Graphics Techniques as a Means of Calculating Light Absorption in Staple Crops
Colleen Heinemann, University of Illinois at Urbana-Champaign, U.S.

Minisymposterium: Checkpoints Compression for Adjoint Computation
Kai-yuan Hou, Northwestern University, U.S.

Minisymposterium: CFD Study of Varying Pebble Diameters in Pebble Bed Nuclear Reactors
Andrew Jones, Boise State University, U.S.

Minisymposterium: Finite-difference Time-domain Method for AZO/ZnO Multilayered 1D Structures
Priscilla Kelly, San Diego State University, U.S.

Minisymposterium: Test Harness
Bhavya Kumaran, Case Western Reserve University, U.S.

Minisymposterium: Network Traffic Performance Prediction with Multivariate Clusters in Time Windows
Tyler Leibengood, Youngstown State University, U.S.

Minisymposterium: Modification and Application of a Method for Studying Stability of High-speed Boundary Layers
Kevin Luna, University of Arizona, U.S.

Minisymposterium: Computational Tools for the Reconstruction of Atmospheric Aerosols via Spectroscopy
Aimee Maurais, Virginia Tech, U.S.

Minisymposterium: New Gene Editing Capabilities Against Vector-borne Infectious Diseases
Pedro X. Medina, University of Puerto Rico at Arecibo, Puerto Rico

Minisymposterium: PDE-informed Covariance for Model-form Uncertainty
Carlos Michelen-Strofer, Virginia Tech, U.S.

Minisymposterium: Numerical Study of Impact of Interparticle Interactions on Particle Dynamics in Lab Generated Spark Discharge Plasma
Pooja Rao, University of Illinois at Urbana-Champaign, U.S.

Minisymposterium: Capsule Networks for Protein Structure Classification
Dan Rosa, University of Puerto Rico, Puerto Rico

Minisymposterium: In Situ Performance Analysis of HPC Network Simulations
Caitlin Ross, Rensselaer Polytechnic Institute, U.S.

continued in next column

continued on next page

Minisymposterium: Molecular Communication in Biological Cells: Foundational Study and Development of Computational Techniques
Zahmeeth Sayed Sakkaff, University of Nebraska, Lincoln, U.S.

Minisymposterium: A Comparative Analysis of Parallel Louvain Algorithms for Community Detection
Naw Safrin Sattar, University of New Orleans, U.S.

Minisymposterium: Collective I/O using RAM Area Network (RAN)
Sergio Servantez, Illinois Institute of Technology, U.S.

Minisymposterium: Analyzing and Evaluating Resilience of Scheduling Scientific Applications
Sonny R Sevin, Slippery Rock University of Pennsylvania, U.S.

Minisymposterium: Quantum Local Search for Graph Community Detection
Ruslan Shaydulin, Clemson University, U.S.

Minisymposterium: Blow-up Behavior of Conservation Law with Spatially Varying Flux
Choah Shin, Oregon State University, U.S.

Minisymposterium: Utilization of the Polymerase Chain Reaction and DNA Barcoding Method in Bioinformatics for the Identification of Unknown Fish Species
Kaayla Tippins, Jarvis Christian College, U.S.

Minisymposterium: Using Deep Learning on Medical Data for Suicide Prevention
Xiange "Shirley" Wang, Hood College, U.S.

Minisymposterium: Locally-implicit Discontinuous Galerkin Schemes with Limiters that Guarantee Moment-invertibility for Hyperbolic Quadrature-based Moment Closures
Christine Wiersma, Iowa State University, U.S.

Minisymposterium: Application of the Ensemble Kalman Filter in Tsunami Wavefield Reconstruction
Yuyun Yang, Stanford University, U.S.

Tuesday, February 26

PP103

Minisymposterium: Software Productivity and Sustainability for CSE and Data Science

4:50 p.m.-6:50 p.m.

Room: Riverside Hall D

Lois Curfman McInnes, Argonne National Laboratory, U.S.

Software is the key crosscutting technology that enables advances in mathematics, computer science, and domain-specific science and engineering to achieve robust simulations and analysis for predictive science, engineering, and other research fields. While software is becoming more complex due to multiphysics and multiscale modeling, the coupling of data analytics, and disruptive changes in computer hardware (due to increases in typical system scale and heterogeneity, including GPUs and additional alternative architectures), software itself has not traditionally received focused attention in the CSE community or been rewarded by that community. The presenters in this minisymposterium will address work that addresses growing technical and social challenges in software productivity, quality, and sustainability, and thereby helps software fulfill its critical role as a cornerstone of long-term CSE collaboration. Having a minisymposterium for these topics provides a natural gathering point during poster sessions for informal conversation.

Minisymposterium: Ginkgo: Designing a Single-Node Linear Operator Framework for High Performance Computing
Hartwig Anzt, University of Tennessee, U.S.; Goran Flegar, Universitat Jaume I, Spain; *Terry Cojean* and Pratik Nayak, Karlsruhe Institute of Technology, Germany; Enrique Quintana-Orti, Universidad Jaume I, Spain

Minisymposterium: hIPPYlib: An Extensible Software Framework for Large-scale Inverse Problems
Olalekan Babaniyi, University of California, Merced, U.S.; Omar Ghattas, University of Texas at Austin, U.S.; Noemi Petra, University of California, Merced, U.S.; Umberto Villa, Washington University, St. Louis, U.S.

Minisymposterium: Shaping a Sustainable SUNDIALS: Applying Software Sustainability Practices to a CSE Library
Cody J. Balos and David J. Gardner, Lawrence Livermore National Laboratory, U.S.; Daniel R. Reynolds, Southern Methodist University, U.S.; Carol S. Woodward and Slaven Peles, Lawrence Livermore National Laboratory, U.S.; Alan Hindmarsh, Lawrence Berkeley National Laboratory, U.S.

Minisymposterium: Outreach for Better Scientific Software
David E. Bernholdt, Oak Ridge National Laboratory, U.S.

Minisymposterium: Using Software Best Practices to Advance Research Across the Aviation Industry
Stuart Bowman, MITRE Corporation, U.S.

Minisymposterium: Software Engineering for Research Software
Jeffrey C. Carver, University of Alabama, U.S.

Minisymposterium: Slate: Developing Sustainable Linear Algebra Software for Exascale
Jamie M. Finney, Jakub Kurzak, Gerald Ragghianti, and Asim YarKhan, University of Tennessee, Knoxville, U.S.; Mark Gates and Piotr Luszczek, University of Tennessee, U.S.; Jack J. Dongarra, University of Tennessee and Oak Ridge National Laboratory, U.S.

Minisymposterium: A Survey of Development Practices in High Performance Computing Applications
Alexander Grannan, Argonne National Laboratory, U.S.; Kanika Sood, Nisansa de Silva, and Boyana Norris, University of Oregon, U.S.; Anshu Dubey, Argonne National Laboratory, U.S.

Minisymposterium: A Look at PFLOTRAN's Cloud-based Continuous Integration
Glenn Hammond, Sandia National Laboratories, U.S.; Gautam Bisht, Lawrence Berkeley National Laboratory, U.S.

continued in next column

continued on next page

Tuesday, February 26

PP103

Minisymposterium: Software Productivity and Sustainability for CSE and Data Science

continued

Minisymposterium: LibEnsemble + PETSc/TAO Sustaining a Library for Dynamic Ensemble-based Computations

Stephen Hudson, Argonne National Laboratory, U.S.; David S. Bindel, Cornell University, U.S.; Jeffrey Larson, Barry F. Smith, and Stefan Wild, Argonne National Laboratory, U.S.

Minisymposterium: Sustaining Student Software

Kathryn Huff, University of Illinois at Urbana-Champaign, U.S.

Minisymposterium: Productive and Sustainable Python Workflows in ParSL

Daniel S. Katz, University of Illinois at Urbana-Champaign, U.S.; Yadu Babuji, Kyle Chard, and Ben Clifford, University of Chicago, U.S.; Ian Foster, Argonne National Laboratory, U.S.; Lukasz Lacinski, University of Chicago, U.S.; Connor Pigg, University of Illinois, U.S.; Michael Wilde, Parallel Works, U.S.; Anna Woodard, University of Chicago, U.S.

Minisymposterium: The NWChemEX Simulation Development Environment - A General Computational Chemistry Software Framework

Ryan Richard, Iowa State University, U.S.; Kristopher Keipert, Argonne National Laboratory, U.S.; Thom Dunning, University of Washington, U.S.; Robert Harrison, Stony Brook University, U.S.; Theresa Windus, Iowa State University, U.S.

Minisymposterium: Improving the Development Workflow of the SETSM Photogrammetry Software

Samuel Khuvis and Judith Gardiner, Ohio Supercomputer Center, U.S.; Ian Howat, Ohio State University, U.S.; Caleb Lehman, Ohio Supercomputer Center, U.S.; Myoung-Jong Noh, Ohio State University, U.S.; Karen Tomko, Ohio Supercomputer Center, U.S.

Minisymposterium: Automated Performance Analysis with PETSc

Matthew G. Knepley, State University of New York at Buffalo, U.S.; Justin Chang, National Renewable Energy Laboratory, U.S.; Albert Cowie, State University of New York at Buffalo, U.S.

Minisymposterium: Managing Software Development Requirements with user Stories

Osni A. Marques, Lawrence Berkeley National Laboratory, U.S.

Minismymposterium: Progress in CSE Software Ecosystems

Lois Curfman McInnes, Argonne National Laboratory, U.S.

Minisymposterium: The Journal of Open Source Software

Kyle E. Niemeyer, Oregon State University, U.S.; Arfon Smith, Space Telescope Science Institute, U.S.; Lorena Barba, George Washington University, U.S.; Jed Brown, University of Colorado Boulder, U.S.; Jason Clark, Montana State University, U.S.; George Githinji, Wellcome Trust, United Kingdom; Melissa Gymrek, University of California, San Diego, U.S.; Lindsey Heagy, University of British Columbia, Canada; Kathryn Huff, University of Illinois at Urbana-Champaign, U.S.

Minisymposterium: How to Professionally Develop Reusable Scientific Software – and When Not To

Vyas Ramasubramani, University of Michigan, U.S.; Carl S. Adorf, University of Michigan, Ann Arbor, U.S.; Joshua Anderson and Sharon Glozter, University of Michigan, U.S.

Minisymposterium: Exahype: An Exascale Hyperbolic PDE Engine

Anne Reinarz and Michael Bader, Technische Universität München, Germany

Minisymposterium: Changing Mindsets for Large-scale Modelling

Chris Richardson, Garth Wells, and Nathan Sime, University of Cambridge, United Kingdom

Minisymposterium: Increasing Software Testing Coverage and Portability with Spack

Jon Rood and Shreyas Ananthan, National Renewable Energy Laboratory, U.S.

Minisymposterium: Modernizing the Scientific Software Approach for the Fusion Analysis Code Transp

Jai Sachdev, Marina Gorelenkova, Xingqiu Yuan, Joshua Breslau, and Francesca Poli, Princeton Plasma Physics Laboratory, U.S.

Minisymposterium: No Need for Excuses: Applying Software Engineering Principles to Facilitate Scientific Software Documentation

Spencer Smith, McMaster University, Canada

Minisymposterium: Analyzing Open-source Scientific Software Projects

Kanika Sood, Nisansa Silva, and Boyana Norris, University of Oregon, U.S.; Anshu Dubey and Lois Curfman McInnes, Argonne National Laboratory, U.S.

Minisymposterium: Reproducible Computational Scientific Workflows with Signac

Carl S. Adorf, University of Michigan, Ann Arbor, U.S.; Vyas Ramasubramani, University of Michigan, U.S.; Bradley Dice and Paul Dodd, University of Michigan, Ann Arbor, U.S.; Sharon C. Glotzer, University of Michigan, U.S.

Minisymposterium: A Software Productivity and Sustainability Case Study: Multithreaded Requests to Cloud Services for Intelligent Address Standardization

Aleksei Sorokin, Illinois Institute of Technology, U.S.; Andy Liu, Illinois Mathematics and Science Academy, U.S.; Sou-Cheng T. Choi, Illinois Institute of Technology, U.S. and Allstate Insurance Corporation, U.S.

Minisymposterium: IDEAS PSIP in Practice: Adopting Continuous Integration for Exascale MD

Richard J. Zamora, Argonne National Laboratory, U.S.; Christoph Junghans and *David Moulton*, Los Alamos National Laboratory, U.S.

Industry Reception

6:30 p.m.-8:00 p.m.

Room: Ballroom Foyer

continued in next column

continued in next column

Tuesday, February 26

PD6

Thinking of Writing a Book?

7:00 p.m.-8:00 p.m.

Room: Conference Theater

Chair: Nicholas J. Higham, University of Manchester, United Kingdom

Ever thought about writing a book? Ever wondered just what that entails? This session brings together successful authors and publishing staff to discuss the process. Topics of interest will include:

• Why and when you should consider writing a book

• A step-by-step description of the process, from initial idea to published book

• How to choose a publisher

• The author/publisher relationship – who does what

• Pitfalls to avoid

• Tips from successful authors Even if you aren't currently thinking of becoming an author, this session promises to be lively and engaging!

Daniela Calvetti
Case Western Reserve University, U.S.

Paul Constantine
Sandia National Laboratories, U.S.

Elizabeth Greenspan
SIAM, U.S.

Jeffrey Humpherys
Brigham Young University, U.S.

SISC Editorial Board Meeting

7:30 p.m.-9:30 p.m.

Room: Davenport Hotel -- Meeting Room 10

Wednesday, February 27

MAC Committee Meeting

7:00 a.m.-8:30 a.m.

Room: Davenport Hotel -- Meeting Room 12

Book Committee Meeting

7:00 a.m.-8:15 a.m.

Room: Davenport Hotel -- Meeting Room 11

Registration

8:00 a.m.-4:00 p.m.

Room: Ballroom Foyer

Wednesday, February 27

IP5

Communication Avoiding: From Linear Algebra to High Order Tensors

8:30 a.m.-9:15 a.m.

Room: Ballroom 100BC

Chair: Tamara G. Kolda, Sandia National Laboratories, U.S.

In this talk I will review 10 years of research in designing robust algorithms for linear algebra that reduce or even minimize in some cases data movement between processors, hence enhancing scalability while preserving accuracy. I will focus on two major operations in linear algebra. The first one, ubiquitous in scientific computing, involves solving large sparse linear systems of equations arising from the discretization of PDEs with strongly heterogeneous coefficients. In particular, I will discuss here enlarged Krylov subspace solvers and algebraic preconditioners that are robust, in the sense that they allow to bound the condition number of the preconditioned matrix. The second operation, ubiquitous in the data analysis applications of scientific computing, involves the compression of large volumes of data while preserving information. I will present a unified view of deterministic and randomized algorithms for computing the low rank approximation of a matrix that allow to attain a given precision with low communication cost. I will discuss the efficiency of the proposed methods in the context of several different real life applications. I will conclude the talk with tensors in high dimensions presenting some recent results while emphasizing many remaining open questions.

Laura Grigori
Inria Paris, France

Coffee Break

9:15 a.m.-9:45 a.m.

Room: Ballroom Foyer

Wednesday, February 27

MT2

Solving Differential Equations with the libMesh Finite Element Library

9:45 a.m.-11:25 a.m.

Room: 300D

This tutorial provides a hands-on introduction to the libMesh finite element library. Attendees will be exposed to the scope of the applications studied by current libMesh developers and users, an overview of the organization of the library, instructions for building the library on a variety of platforms, and examples of constructing application codes based on the library. It is expected that attendees are familiar with the Linux operating system and are familiar writing and building C++ applications.

11:25-11:25 Solving Differential Equations with the libMesh Finite Element Library

Paul Bauman, State University of New York at Buffalo, U.S.

11:25-11:25 Solving Differential Equations with the libMesh Finite Element Library

Roy Stogner, University of Texas at Austin, U.S.

11:25-11:25 Solving Differential Equations with the libMesh Finite Element Library

John W. Peterson, Idaho National Laboratory, U.S.

Wednesday, February 27

MS170

Exploiting Model Hierarchies, Sparsity and Low-Rank Structure of Large-scale Bayesian Computation - Part I of II

9:45 a.m.-11:25 a.m.

Room: Ballroom 100BC

For Part 2 see MS204

One of the central tasks in modern computational science and engineering is to quantify the uncertainty associated with the computational results. This can be broadly classified as Bayesian Computation. Frequently these tasks involve expensive or intractable numerical models, embedded in likelihood functions that involve high-dimensional parameters and/or high-dimensional data sets. The output is typically the posterior (post data) distribution of parameters or quantities of interest, or statistics of those parameters or quantities. Markov chain Monte Carlo, sequential Monte Carlo, and other posterior exploration schemes require repeated evaluations of such models, in principle over high-dimensional spaces. In this setting, standard algorithms quickly become intractable. Methods for identifying and exploiting model hierarchies, low-rank structure and sparsity in the representation of the posterior distribution are becoming essential for solving these otherwise intractable Bayesian inference problems. This minisymposium will bring together researchers working on the forefront of structure-exploiting methods intended to accelerate such large-scale Bayesian computation.

Organizer: Tiangang Cui
Monash University, Australia

Organizer: Robert Scheichl
Universität Heidelberg, Germany

9:45-10:05 Efficient Sampling from High-dimensional Distributions using Low-rank Tensor Surrogates

Robert Scheichl, Universität Heidelberg, Germany; Karim Anaya-Izquierdo and Sergey Dolgov, University of Bath, United Kingdom; Colin Fox, University of Otago, New Zealand

10:10-10:30 Layers of Low-rank Couplings for Large-scale Bayesian Inference

Daniele Bigoni, Massachusetts Institute of Technology, U.S.; Olivier Zahm, Inria Grenoble, France; Alessio Spantini and Youssef M. Marzouk, Massachusetts Institute of Technology, U.S.

10:35-10:55 Multi-Index and Multi-level Markov Chain Monte Carlo in MUQ2

Linus Seelinger, Ole Klein, and Robert Scheichl, Universität Heidelberg, Germany

11:00-11:20 Use of the Bayesian Approximation Error Approach to Account for Model Discrepancy: The Robin Problem Revisited

Ruanui Nicholson, University of Auckland, New Zealand; Noemi Petra, University of California, Merced, U.S.; Jari Kaipio, University of Auckland, New Zealand

continued in next column

Wednesday, February 27

MS171

Progress and Challenges in Extreme Scale Computing and Data - Part I of II

9:45 a.m.-11:25 a.m.

Room: Conference Theater

For Part 2 see MS205

Extreme scale computing efforts have resulted in numerous advances for multicore, manycore and accelerator based scalable systems. In addition, large-scale applications must increasingly deal with data management and analysis as a first-class concern. Therefore, new applications often have to manage distributed and parallel computing, and have to manage workflows of different tasks (computing, data analytics, machine learning, visualization,…). In this minisymposium we present some of the latest works in scalable algorithms, programming paradigms, and libraries for next generation computing platforms. Furthermore, we discuss efforts to better incorporate data science concerns as a principle component of our scientific workflows.

Organizer: Serge Petiton

University of Lille, France

Organizer: Michael A. Heroux

Sandia National Laboratories, U.S.

Organizer: Kengo Nakajima

University of Tokyo, Japan

9:45-10:05 Combining Extreme Computing and Big Data for Future Machine Learning

Serge Petiton, University of Lille, France; Kesheng Wu and Osni A. Marques, Lawrence Berkeley National Laboratory, U.S.

10:10-10:30 A Scalable Randomized Singular Value Decomposition with Multiple Sketches for Big Data Analysics

Weichung Wang, National Taiwan University, Taiwan

10:35-10:55 Multidisciplinary High Performance Data Analysis Forum: summary and recommendations

Michel Dayde and Pierre-Henri Cros, University of Toulouse, France; Osni A. Marques, Lawrence Berkeley National Laboratory, U.S.

11:00-11:20 Obtaining Performance from a Julia-Implementation of Trilinos Data Librairies

Neil Lindquist, St. John's University, U.S.

continued in next column

Wednesday, February 27

MS172

Advanced HPC Trends for Oil and Gas Applications - Part I of II

9:45 a.m.-11:25 a.m.

Room: 102A

For Part 2 see MS206

For decades, advances in high performance computing have led to more accurate, safer, and faster oil and gas exploration and production processes. Applications range from seismic imaging and reservoir simulations to seismic interpretation and digital rock physics. This evolution is paramount today to further develop production by exploring increasingly complex reservoirs and by enhancing the recovery ratios from existing fields. In fact, these requirements translate into larger volumes of data to process and models of higher fidelity in terms of physical formulation and space and time resolutions. To cope with these challenges, the oil and gas industry has to adapt constantly to a changing technology landscape in terms of algorithms, platforms, software, and tools. The move towards less synchrony at all system levels, the widening gap between cost of IO versus floating-point operations, and the recent advent of Deep Learning are a few of these important changes the oil and gas software industry has to consider in moving forward with exascale. This minisymposium is an opportunity to discuss today's and future trends defining oil and gas HPC applications design.

Organizer: Rached Abdelkhalak

King Abdullah University of Science \\\\& Technology (KAUST), Saudi Arabia

Organizer: Gabriel Fabien-Ouellet

Polytechnique Montreal, Canada

continued on next page

Wednesday, February 27

MS172

Advanced HPC Trends for Oil and Gas Applications - Part I of II

continued

9:45-10:05 Leveraging Seismic Modeling with Mixed Precision Arithmetic

Rached Abdelkhalak, Hatem Ltaief, and Kadir Akbudak, King Abdullah University of Science & Technology (KAUST), Saudi Arabia; Vincent Etienne and Thierry Tonellot, Saudi Aramco Oil Company, Saudi Arabia; David E. Keyes, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

10:10-10:30 Evaluation of Compression Algorithms and Memory Hierarchies Impact Within Time Reversal Algorithms

Philippe Thierry, Intel Corporation, France; sunny gogar, Mathieu Gontier, Alberto Villarreal, Yvan Zavarzin, Sergey Khlystov, and Leonardo Borges, Intel Corporation, U.S.; Essam Algizawy, Brightskies Inc, Egypt

10:35-10:55 High-performance Asynchronous Execution of the Reverse Time Migration for the Oil and Gas Industry

Amani Alonazi and Hatem Ltaief, King Abdullah University of Science & Technology (KAUST), Saudi Arabia; Issam Said, NVIDIA, U.S.; Rached Abdelkhalak and David E. Keyes, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

11:00-11:20 The ARM HPC Ecosystem for Oil and Gas Applications

Fabrice Dupros and Conrad Hillairet, ARM, France

Wednesday, February 27

MS173

Student Days: Undergraduate Presentations

9:45 a.m.-11:25 a.m.

Room: 102B

Talks presented by undergraduates. Part of "Student Days". 12/11/18 there were two sessions earmarked for this but only one was needed. Part II was unlinked and deleted. CEY

Organizer: Kathleen Kavanagh

Clarkson University, U.S.

9:45-10:00 Analysis of Equity Markets: A Graph-theory Approach

Joshua Abrams, José Celaya-Alcalá, Drew Baldwin, Ryan Gonda, and Zhaoren Chen, University of Arizona, U.S.

10:05-10:20 Calibration of the Ross Recovery Theorem to Real-world Data, and Tests of its Practical Value

Ling Lan and Zhengxu Li, Courant Institute of Mathematical Sciences, New York University, U.S.

10:25-10:40 Multiresolution Methods for Convolutional Neural Networks

Evan Scope Crafts, Emory University, U.S.

10:45-11:00 Reducing Computation Time of Image Deblurring Algorithms Using Wavelets

Jonathan Wittmer, University of Akron, U.S.

11:05-11:20 Evaluating Tangling in High-order Meshes with Tangent Vectors

Myra Dotzel and Mike Stees, University of Kansas, U.S.

Wednesday, February 27

MS174

Advances and Applications in Numerical Methods for Interfacial Flows - Part I of II

9:45 a.m.-11:25 a.m.

Room: 102C

For Part 2 see MS207

Flows with multiple components separated by fluid interfaces are ubiquitous in nature and engineering. In the past tens of years, vast progress has been made in the numerical simulation of interfacial flows and many outstanding methods have been developed, eg., the volume-of-fluid, level-set, front-tracking, and phase-field methods. However, challenges, such as the coupling between fluid interface and complex rheology, moving contact line, and accurate and robust discretization, still remain. This minisymposium aims to bring together experts in this area to discuss new developments across different methods and identify new application areas.

Organizer: Pengtao Yue

Virginia Tech, U.S.

Organizer: Shahriar Afkhami

New Jersey Institute of Technology, U.S.

9:45-10:05 Breakup of Finite-size Liquid Filaments: Transition from No-breakup to Breakup Including Substrate Effects

Shahriar Afkhami, New Jersey Institute of Technology, U.S.

10:10-10:30 A Moment-of-Fluid Method for Diffusion Equations on Irregular Domains with Application to Problems in Multi-material Systems

Yang Liu and Mark Sussman, Florida State University, U.S.; Yongsheng Lian, University of Louisville, U.S.; M. Yousuff Hussaini, Florida State University, U.S.

10:35-10:55 Numerical Modeling and Simulation of Parachute Inflation by Tracking an Immersed Elastic Interface in Incompressible Flow

Xiaolin Li, Tengbo Yang, and Brandon Balletine, State University of New York, Stony Brook, U.S.

11:00-11:20 Dynamics of Particles Trapped at Complex Fluid Interfaces

Mingfeng Qiu, James Feng, and *Jean-Christophe Loudet*, University of British Columbia, Canada

Wednesday, February 27

MS175

Numerical Methods for Multi-Material Fluid Flows - Part I of II

9:45 a.m.-11:25 a.m.

Room: 102D

For Part 2 see MS208

The minisymposia aims to bring together researchers from universities and research laboratories to discuss the state-of-the-art in multi-material hydrodynamics simulations. It focuses on relevant numerical methods, on the analysis of such methods and on the modeling of complex multi-material flows. Topice include Lagrangian Hydrodynamics, Arbitrary Lagrangian Eulerian (ALE) Methods, Eulerian Methods, Mesh Generation Methods, including Mesh Adaptation, Interface Reconstruction Methods, Data Transfer Between Meshes—Remapping, Advanced Discretization Methods, High-order Methods.

Organizer: Mikhail Shashkov

Los Alamos National Laboratory, U.S.

9:45-10:05 A Sub Cell Dynamics Closure Model for Solids with Void Opening and Closure

Andrew J. Barlow, Atomic Weapons Establishment, United Kingdom; Misha Shashkov, Los Alamos National Laboratory, U.S.; Matej Klima, Czech Technical University, Czech Republic

10:10-10:30 The Distribution-Based Remapping of the Nodal Mass and Momentum Between Arbitrary Meshes for Staggered Arbitrary Lagrangian-Eulerian Hydrodynamics

Mikhail Shashkov and Mack Kenamond, Los Alamos National Laboratory, U.S.

10:35-10:55 3D Remapping on General Polyedral Mesh: Overlays, Polynomial Reconstructions and a Posteriori Limiting

Ghislain Blanchard, ONERA, France; *Raphael Loubere*, Bordeaux University, France; Pavel Vachal, Czech Technical University, Prague, Czech Republic

11:00-11:20 Optimizing Artificial Viscosity using Machine Learning

Jason Albright, Mikhail Shashkov, and Nathan Urban, Los Alamos National Laboratory, U.S.

Wednesday, February 27

MS176

Machine Learning Strategies for Computer Simulation of Physical Systems - Part I of II

9:45 a.m.-11:25 a.m.

Room: 111A

For Part 2 see MS209

The bidirectional relation between computational models and data necessitates commensurate advances in machine learning algorithms and tools. This minisymposium aims to go beyond the parametric calibration of models and to address questions that pertain to the discovery of new models from computationally or experimentally-generated data. This need arises often in several instances of computational modeling as for example in uncertainty quantification where the construction of reduced-order or surrogate models is imperative in order to overcome the huge number of calls to highly-complex and expensive models. Furthermore, similar problems appear in the context of multiscale simulations where high-resolution, computationally-generated data over relatively small spatial domains and time horizons must be used in order to construct effective models of the system's behavior over larger spatiotemporal scales. In both of the aforementioned cases, significant challenges relate to high-dimensionality of inputs and outputs which usually exceed those encountered in typical machine learning applications, the discovery of meaningful and physically-interpretable collective variables and structural features of the models as well as the incorporation of physical constraints and invariances whenever these are available.

Organizer: Phaedon S. Koutsourelakis

Technische Universität München, Germany

9:45-10:05 Physics-informed Learning for Multiscale Systems

Nicholas Zabaras, University of Notre Dame, U.S.; Phaedon-Stelios Koutsourelakis, Technische Universität München, Germany; Paris Perdikaris, University of Pennsylvania, U.S.; Jianxun Wang, University of Notre Dame, U.S.

10:10-10:30 Discovering Unknown Physics and Enforcing Known Symmetries and Constraints with Machine Learning,

Steven Brunton, University of Washington, U.S.

10:35-10:55 LES Wall Modeling via Physics-informed Neural Networks

Xiang Yang and Suhaib Zafar, Pennsylvania State University, U.S.; Jian-Xun Wang, University of Notre Dame, U.S.; Heng Xiao, Virginia Tech, U.S.

11:00-11:20 A Machine Learning Guided Stochastic Modelling of Biomolecule Systems

Lei Wu, Peking University, China; Huan Lei, Pacific Northwest National Laboratory, U.S.

continued in next column

Wednesday, February 27

MS177

Machine Learning Approaches for the Sciences and Engineering: Recent Developments - Part I of II

9:45 a.m.-11:25 a.m.

Room: 111B

For Part 2 see MS210

This session brings together researchers and practioners to present recent results on emerging data-driven approaches and machine learning methods, with an emphasis on their adaptation for use in the sciences and engineering. There have been many recent successes in areas, such as Natural Language Processing (NLP) and Image Processing (IP). However, problems encountered in the sciences and engineering disciplines often have rather different requirements and structure than those of NLP/IP. For adoption in these new problem domains, further mathematical and computational developments are needed for effective use of machine learning methods. In this session, we bring together researchers working on specific applications that would benefit from such data-driven/ML approaches, as well as, investigators working on fundamental new developments in emerging areas of machine learning.

Organizer: Ben J. Gross

University of California, Santa Barbara, U.S.

Organizer: Paul J. Atzberger

University of California, Santa Barbara, U.S.

9:45-10:05 Machine Learning Methods for Fluid Dynamics on Manifolds: Physics-Informed Gaussian Process Regression (PI-GPR) and PIV

Ben J. Gross and Paul J. Atzberger, University of California, Santa Barbara, U.S.

10:10-10:30 Enforcing Constraints for Interpolation and Extrapolation in Generative Adversarial Networks

Panos Stinis, Pacific Northwest National Laboratory, U.S.

10:35-10:55 Model Compression and Fast Prediction for LSTM

Cho-Jui Hsieh, University of California, Davis, U.S.

11:00-11:20 Data-driven Discovery of Nonlinear Dynamics

Kathleen Champion, J. Nathan Kutz, and Steven Brunton, University of Washington, U.S.

continued in next column

Wednesday, February 27

MS178

Recent Advances in Multilevel Solvers - Part I of II

9:45 a.m.-11:25 a.m.

Room: 111C

For Part 2 see MS211

For many years, multilevel solvers have played a central role in many areas of computer simulation, including fluid dynamics, solid mechanics, computer graphics, geophysical exploration, and more. As a result, development and analysis of multilevel solvers remains an important area of research in computational science and engineering. In this minisymposium, we bring together speakers focused on several aspects of multilevel solvers, including developing algorithms for coupled multiphysics problems, analysis of algorithms, and graph problems.

Organizer: James H. Adler

Tufts University, U.S.

Organizer: Scott Maclachlan

Memorial University, Newfoundland, Canada

Organizer: Xiaozhe Hu

Tufts University, U.S.

9:45-10:05 Robust Multilevel Preconditioners for a New Stabilized Discretization of the Poroelastic Equation

Peter Ohm, James H. Adler, and Xiaozhe Hu, Tufts University, U.S.; Ludmil Zikatanov, Pennsylvania State University, U.S.; Carmen Rodrigo and Francisco José Gaspar, University of Zaragoza, Spain

10:10-10:30 An Adaptive Unsmoothed Aggregation Multigrid Method Based on Path Cover

Junyuan Lin, Tufts University, U.S.

10:35-10:55 Enrichment for Multilevel Solvers

Matthew Knepley, State University of New York at Buffalo, U.S.

11:00-11:20 Auxiliary Space Preconditioning for Mixed Finite Element Discretizations of Richards Equation

Juan Batista, Pennsylvania State University, U.S.

Wednesday, February 27

MS179

High-order Finite Element Methods for Complex and Multiphysics Applications - Part I of II

9:45 a.m.-11:25 a.m.

Room: 300A

For Part 2 see MS212

A current challenge before the applied mathematics and computational science community is the solution of complex and multiphysics systems. Science and engineering applications for these systems are highly-nonlinear and have strongly coupled physical mechanisms that interact on a wide range of length- and time-scales and require robust and efficient high-resolution numerical approximations. For this reason, developing robust, and accurate methods that can effectively use parallel computation at extreme scales is critical. In this context numerical discretizations and solvers for practical multiphysics simulations must be: (1) High-order accurate in space and time; (2) Stable; (3) Conservative; (4) Having minimum degrees of freedom for implicit solution approaches; (5) Well suited for unstructured meshes; (6) Well suited for hp-adaptivity; (7) Well suited for applications with disparate temporal and spatial scales; and (8) Well suited for fine-grain parallelism. This minisymposium focuses on the latest developments in high(er) order finite element, and related methods and associated numerical methods. The speakers in this minisymposium will address theoretical/numerical and computational issues that are critical to developing approaches with these desired properties. Applications will include aerodynamics, magnetohydrodynamics, plasma physics, subsurface flows, geophysical flows, etc.

Organizer: Sriramkrishnan Muralikrishnan
University of Texas at Austin, U.S.

Organizer: Shinhoo Kang
University of Texas at Austin, U.S.

Organizer: Tan Bui
University of Texas at Austin, U.S.

9:45-10:05 High-order Hybridized Discontinuous Galerkin Method and a Scalable Solver for Incompressible Magnetohydrodynamics

Sriramkrishnan Muralikrishnan and Stephen Shannon, University of Texas at Austin, U.S.; Tim Wildey, Sandia National Laboratories, U.S.; Tan Bui, University of Texas at Austin, U.S.; John N. Shadid, Sandia National Laboratories, U.S.

10:10-10:30 High-order Discontinuous Galerkin Discretization for Multi-physics Simulations of Plasmas

Petr Cagas, Virginia Tech, U.S.; Ammar Hakim, Princeton Plasma Physics Laboratory, U.S.; James Juno, University of Maryland, U.S.; Bhuvana Srinivasan, Virginia Tech, U.S.

10:35-10:55 Output-based Mesh Optimization for Embedded Discontinuous Galerkin Methods

Guodong Chen and Krzysztof Fidkowski, University of Michigan, U.S.

11:00-11:20 A High Order Locally Field-aligned Discontinuous Galerkin Method for Anisotropic Problems in Plasma Physics

Benedict Dingfelder, Florian Hindenlang, and Eric Sonnendrücker, Max Planck Institute for Plasma Physics, Germany

Wednesday, February 27

MS180

Multiphysics and Multiscale Problems in Computational Science and Engineering - Part I of II

9:45 a.m.-11:25 a.m.

Room: 300B

For Part 2 see MS213

Multiphysics and multiscale problems are challenging algorithmically and taxing computationally. Since the interface or boundary is constantly changing nonlinearly, accurate representation in a stable manner is not trivial, not to mention nonlinearities in the governing equations. Hence, high order and high-fidelity discretization methods combined with accurate and stable coupling strategies are needed, and analysis tools as well as acceleration (stabilization) techniques are of critical importance. In this minisymposium, we will review a number of new ideas for developing efficient and reliable coupling algorithms and share interesting applications in various engineering and science fields. The topics of interest include and are not limited to: fluid–structure interaction, atomistic-to-continuum coupling problems, nonlocal-to-nonlocal/local coupling problems, quantum mechanics, micro-mechanics of materials, concurrent coupling approaches, multiphysics/ multiscale coupling methods.

Organizer: Xingjie Li
University of North Carolina, Charlotte, U.S.

9:45-10:05 An Immersogeometric FSI Framework for the Modeling and Simulation of Transcatheter Heart Valves

Ming-Chen Hsu, Michael C. H. Wu, and Heather Muchowski, Iowa State University, U.S.

10:10-10:30 Computational Methods for Fluid–structure Interaction with Applications

Padmanabhan Seshaiyer, George Mason University, U.S.

continued in next column

continued on next page

Wednesday, February 27

MS180

Multiphysics and Multiscale Problems in Computational Science and Engineering - Part I of II

continued

10:35-10:55 The Dynamic Augmented Lagrangian Method: Application to Immersogeometric Heart Valve Analysis and Implementation using {FEniCS}

David Kamensky, Brown University, U.S.; Ming-Chen Hsu, Iowa State University, U.S.; Yue Yu, Lehigh University, U.S.; Yuri Bazilevs, Brown University, U.S.

11:00-11:20 Coupled Peridynamic and Subsurface Flow Models for Environmentally-assisted Crack Growth in Concrete Structures

David Littlewood, Jessica Rimsza, Jennifer Frederick, and Reese Jones, Sandia National Laboratories, U.S.

Wednesday, February 27

MS181

Novel Computational Algorithms for Future Computing Platforms - Part I of II

9:45 a.m.-11:25 a.m.

Room: 300C

For Part 2 see MS218

In the early 2000s, due to constraints on economical heat dissipation, clock speeds of single-core CPUs could no longer be increased, which marked the adoption of multi-core CPUs, together with a paradigm shift to algorithms specifically designed for multi-core architectures. About 15 years into this current architectural cycle and on its way to exascale performance, the computing industry finds itself at the confluence of technical difficulties that cast doubt on its ability to sustain this architectural model beyond the exascale capability. These difficulties are driving the hardware industry to develop application-specific chips and to look beyond silicon-based chips (e.g., quantum computing, physical annealing, neuromorphics, etc.), with a continued emphasis on raw processing power and emerging concerns about energy efficiency. Hardware specialization will likely redefine the way computational algorithms are developed over the next two decades for a wide range of important applications: large-scale PDE-based problems (CFD, wave propagation, subsurface modeling, etc.), artificial intelligence, computational chemistry, and optimization problems, to name a few. The pressure to decrease time to solution or improve simulation fidelity, or both, for these applications will continue unabated. This minisymposium provides a forum for sharing innovative ideas on algorithm development for leveraging future computing platforms.

Organizer: Dimitar Trenev
ExxonMobil Research and Engineering

Organizer: Laurent White
ExxonMobil Research and Engineering

9:45-10:05 Scientific Computing in a Changing Landscape

Laurent White and Dimitar Trenev, ExxonMobil Research and Engineering

10:10-10:30 Adapting Scientific Software and Designing Algorithms for Next Generation GPU Computing Platforms

John Stone, University of Illinois at Urbana-Champaign, U.S.

10:35-10:55 Neural-inspired Computing at Sandia Labs – Enabling and Performing Advanced Computation

Craig Vineyard, Sandia National Laboratories, U.S.

11:00-11:20 Fast Recovery from Node Failures for the Parallel Preconditioned Conjugate Gradient Method

Markus Levonyak, Carlos Pachajoa, and Wilfried N. Gansterer, University of Vienna, Austria

continued in next column

Wednesday, February 27

MS182

BE: GPU Accelerated Computing on Summit

9:45 a.m.-11:25 a.m.

Room: 301

Organizer: Mary Ann E. Leung
Sustainable Horizons Institute, U.S.

Wednesday, February 27

MS183

Mathematical Methods for Control and Optimization of Large-Scale Energy Networks - Part I of II

9:45 a.m.-11:25 a.m.

Room: 302A

For Part 2 see MS216

Recent developments in the energy market require new mathematical models and suitable algorithms for the efficient usage of the existing networks. For instance, in the case of electrical power networks, the increase of renewable energy sources and the rise of the electric car lead to new challenges for the mathematical model and require optimization of the power network. If additionally storage units such as batteries are considered, this becomes an even more challenging problem. Moreover, the integration of renewable energy sources such as natural gas, photovoltaics, and wind energy into the power grids poses new challenges from both theoretical and practical sides. Starting from these real-world problems, the mathematical aspects of this minisymposium concentrate on methods for large-scale energy networks and, in particular, address optimization and stability analysis, model predictive control, model-order reduction, uncertainty quantification, and related topics on energy networks. The abstract setting allows for consideration of applications arising from both, gas and power networks.

Organizer: Manuel Baumann
Max Planck Institute for Dynamics of Complex Technical Systems, Germany

Organizer: Yue Qiu
Max Planck Institute for Dynamics of Complex Technical Systems, Germany

9:45-10:05 Model-order Reduction for the Optimization of Distribution Grids with Battery Control

Manuel Baumann and Sara Grundel, Max Planck Institute for Dynamics of Complex Technical Systems, Germany

continued in next column

10:10-10:30 Newton-Krylov Methods for Solving the AC Load Flow Equations

Domenico Lahaye, Technical University of Delft, Netherlands

10:35-10:55 Hierarchical Modeling and Simulation of Power Networks

Riccardo Morandin, Technische Universität, Berlin, Germany

11:00-11:20 Four Mathematical Formulations of the (Optimal) Power Flow Problem and Their Impact on the Performance of Solution Methods

Baljinnyam Sereeter, Kees Vuik, and Cees Witteveen, Delft University of Technology, Netherlands

Wednesday, February 27

MS184

Entropy-stable Formulations for Numerical Solution of Conservation Laws - Part I of II

9:45 a.m.-11:25 a.m.

Room: 302B

For Part 2 see MS217

Entropy stability presents an elegant route to improve the robustness of numerical solutions of compressible flows. The demand for high-order-accurate numerical discretizations and considerations of complex physics such as those involving sharp gradients and multi-component flows has resulted in a renewed interest in entropy stable discretizations. Many issues, such as the interplay between accuracy and robustness, the potential role of energy-preserving schemes, extensions to multi-species flows, and interaction with sub-grid-scale models are currently being explored. This minisymposium will address the related mathematical and numerical issues as well as applications in fluid dynamics.

Organizer: Scott Murman

NASA Ames Research Center, U.S.

9:45-10:05 Entropy Stability and the Computation of Entropic Measure-valued Solutions

Eitan Tadmor, University of Maryland, U.S.

10:10-10:30 A Fully-Discrete Kinetic Energy Preserving and Entropy Conservative Scheme for Compressible Flows

Deep Ray, École Polytechnique Fédérale de Lausanne, Switzerland; Praveen Chandrashekar, TIFR Centre, Bangalore, India

10:35-10:55 Entropy Stable Schemes for Multicomponent Flows

Ayoub Gouasmi, University of Michigan, U.S.; Karthik Duraisamy, University of Michigan, Ann Arbor, U.S.; Scott Murman, NASA Ames Research Center, U.S.

11:00-11:20 Entropy Stable Discretizations of Compressible Flows in Thermochemical Non-equilibrium

Michael Hansen and Travis Fisher, Sandia National Laboratories, U.S.

Wednesday, February 27

MS185

High-order PDE Methods on CPU Architectures with Wide SIMD Units - Part I of II

9:45 a.m.-11:25 a.m.

Room: 303A

For Part 2 see MS222

In Summer 2016 the 2nd generation Intel Xeon Phi Processor was released and since 2017 the Intel Scalable processor is available. Both are the computational heart of many current and future supercomputing installations. Examples are the DoE-machines Trinity-II, Cori-II and Theta, the NSF-machine Stampede-2, Oakforest-PACS in Japan, or Europe's SuperMUC-NG. These processors combine the computational power of accelerator-based machines with the traditional approach of homogeneous high performance computing. From an application viewpoint, many research articles in the last years proved that high-order methods are able to the computational power of modern supercomputers. However, it is a challenging engineering problem to enable complex software packages on manycore architectures. This mini-symposium brings together researchers working on efficient implementations of scalable high-order solvers. Of particular interest are applications which are known to require exascale computing resources in future, such as computational fluid dynamics, electro-magnetics and seismic simulations, as well as weather forecasting. Due to the broad spectrum of applications, this MS aims at identifying common algorithmic patterns of efficient high-order methods. Additionally, the comparison of different numerical approaches (e.g. CG vs. DG) will allow to identify best practices when leveraging Intel Xeon Phi and Intel Xeon Scalable systems at scale.

Organizer: Alexander Heinecke

Intel Corporation, U.S.

Organizer: Alexander Breuer

University of California, San Diego, U.S.

Organizer: David Moxey

University of Exeter, United Kingdom

9:45-10:05 Vectorisation for High-order Simplicial Elements

David Moxey, University of Exeter, United Kingdom; Roman Amici and Mike Kirby, University of Utah, U.S.

10:10-10:30 Fluxo: An Open Source Parallel Split-Form DG Solver for Compressible Navier-Stokes and Resistive MHD Equations

Florian Hindenlang and Eric Sonnendrücker, Max Planck Institute for Plasma Physics, Germany; Gregor Gassner and Andrew R. Winters, University of Cologne, Germany

10:35-10:55 Performance Portable Implementation of High Order DG Methods through Code Generation

Dominic Kempf, René Heß, Steffen Müthing, and Peter Bastian, Universität Heidelberg, Germany

11:00-11:20 SIMD Vectorization for High-order Matrix-free Finite Element Computations in CFD

Peter Munch, Niklas Fehn, Wolfgang A Wall, and Martin Kronbichler, Technische Universität München, Germany

continued in next column

Wednesday, February 27

MS186

Advances in Analyzing Floating-point Errors in Computational Science - Part I of II

9:45 a.m.-11:25 a.m.

Room: 303B

For Part 2 see MS219

Modern error analysis of floating-point computations faces a double challenge with the rise of large-scale, low-precision computations. On the one hand, larger and larger problems are being solved routinely; on the other hand, the use of half-precision arithmetic (fp16) is becoming increasingly attractive. These two trends raise questions about the validity and sharpness of traditional rounding error bounds. In this minisymposium we will discuss recent advances in the analysis of errors in floating-point computations, and in particular some new approaches that have emerged to meet the challenges described above.

Organizer: Pierre Blanchard

University of Manchester, United Kingdom

Organizer: Nicholas J. Higham

University of Manchester, United Kingdom

9:45-10:05 A New Approach to Probabilistic Roundoff Error Analysis

Nicholas J. Higham and *Theo Mary*, University of Manchester, United Kingdom

10:10-10:30 Probabilistic Error Analysis for Inner Products

Ilse Ipsen, North Carolina State University, U.S.

10:35-10:55 Stochastic Analysis and Correction of Floating Point Errors in Monte Carlo Simulations

Oliver Sheridan-Methven and Mike Giles, University of Oxford, United Kingdom

11:00-11:20 Probabilistic Numerics and Applications to Roundoff Error

Jon Cockayne, University of Warwick, United Kingdom; Chris Oates, Newcastle University, United Kingdom; Tim Sullivan, University of Warwick, United Kingdom; Mark Girolami, Imperial College London, United Kingdom; Han Cheng Lie, Free University of Berlin, Germany

Wednesday, February 27

MS187

Computational Methods for Linear Kinetic Transport Equations - Part I of II

9:45 a.m.-11:25 a.m.

Room: 201B

For Part 2 see MS220

Statistical mechanics provides a mathematical modeling framework in which a collection of "particles" (e.g., electrons, nucleons, atoms, or molecules) is represented by a probability density function (PDF). This representation allows one to model the influence of microscopic dynamics on larger scales, without explicitly computing the detailed dynamics on the microscopic scale. Models in this framework are referred to as "kinetic models", and arise in many application areas, including in rarefied gas dynamics, nuclear reactor modeling, and plasma physics. Transport phenomena refers to the evolution and redistribution of macroscopic quantities such as mass, momentum, energy, and heat flux. This minisymposium addresses recent advances in computational methods for the simulation of statistical mechanical models and the resulting transport phenomena. Our goal is to showcase a diverse array of application areas, models, and numerical techniques.

Organizer: Andrew J. Christlieb

Michigan State University, U.S.

9:45-10:05 Hamilton Jacobi Formulation of Vlasov Poisson and a Method of Lines Transpose Approach

Andrew J. Christlieb and William A. Sands, Michigan State University, U.S.

10:10-10:30 Assessment of the Lagrange Discrete Ordinates Equations for Three-dimensional Neutral Particle Transport

Kelly L. Rowland, National Energy Research Scientific Computing Center, U.S.; Cory Ahrens, Los Alamos National Laboratory, U.S.; Steven Hamilton, Oak Ridge National Laboratory, U.S.; Rachel Slaybaugh, University of California, Berkeley, U.S.

10:35-10:55 Piecewise P_N Approximations for Multidimensional Radiative Transfer

Minwoo Shin and James A. Rossmanith, Iowa State University, U.S.

11:00-11:20 Neutron Kinetics in Liquid-fueled Nuclear Reactors

Kathryn Huff, University of Illinois at Urbana-Champaign, U.S.

continued in next column

Wednesday, February 27

MS188

Reduced Order Modeling for Parametric CFD Problems- Part I of II

9:45 a.m.-11:25 a.m.

Room: 201C

For Part 2 see MS221

Large-scale computing is recurrent in several contexts such as computational mechanics or fluid dynamics, due to the high computational complexity in solving parametric and/or stochastic systems, based, for instance, on partial differential equations. This often leads to an unaffordable computational burden, especially when dealing with real-world applications, real-time or multi-query computing. In order to lessen this computational burden, reduced-order modeling techniques (e.g., reduced basis, proper orthogonal and generalized decomposition, empirical interpolation, hierarchical model reduction) play a crucial role as they aim to capture the most important features of the problem at hand without giving up accuracy. This minisymposium focuses on the development and application of reduced-order modeling techniques in computational fluid dynamics for direct and inverse modeling, and for control, optimization and design purposes. The actual goal is twofold, on one hand to provide the state of the art on the most recently developed techniques; on the other hand to identify the new research directions and perspectives on the scientific panorama.

Organizer: Annalisa Quaini
University of Houston, U.S.

Organizer: Gianluigi Rozza
SISSA, International School for Advanced Studies, Trieste, Italy

Organizer: Simona Perotto
Politecnico di Milano, Italy

9:45-10:05 A Localized Reduced-order Modeling Approach for PDEs with Bifurcating Solutions

Martin Hess, SISSA, International School for Advanced Studies, Trieste, Italy; Alessandro Alla, Pontifical Catholic University of Rio de Janeiro, Brazil; *Annalisa Quaini*, University of Houston, U.S.; Gianluigi Rozza, SISSA, International School for Advanced Studies, Trieste, Italy; Max Gunzburger, Florida State University, U.S.

10:10-10:30 An Approximated Minimum Residual Reduced Basis Method for the Parameterized Incompressible Navier-Stokes Equations

Tommaso Taddei, INRIA Bordeaux, France

10:35-10:55 Reduced Order Methods for Parametric Optimal Flow Control in Coronary Bypass Grafts: Patients Specific Data Assimilation and Geometrical Reconstruction

Zakia Zainib, Francesco Ballarin, and Gianluigi Rozza, SISSA, International School for Advanced Studies, Trieste, Italy

11:00-11:20 HIGAMod Approximation for the Stokes Equations in Patient-specific Geometries

Yves Barbosa, MOX, Politecnico di Milano, Italy

Wednesday, February 27

MS189

Advances in Quasi-Monte Carlo Methods - Part I of II

9:45 a.m.-11:25 a.m.

Room: 202B

For Part 2 see MS223

Quasi-Monte Carlo (qMC) methods are the algorithms of choice for evaluating high dimensional integrals. Recent advances in qMC include the development of open source software, applications to uncertainty quantification, and multivariate statistical inference. This minisymposium highlights some of these advances.

Organizer: Fred J. Hickernell
Illinois Institute of Technology, U.S.

Organizer: Michael B. Giles
University of Oxford, United Kingdom

9:45-10:05 The Start of Community Supported Quasi-Monte Carlo Software

Fred J. Hickernell, Illinois Institute of Technology, U.S.

10:10-10:30 QMC Methods for Elliptic PDEs Driven by White Noise

Matteo Croci and Mike Giles, University of Oxford, United Kingdom; Marie E. Rognes, Simula Research Laboratory, Norway; Patrick E. Farrell, University of Oxford, United Kingdom

10:35-10:55 Computing Multivariate Normal Variance Mixture Distributions with Quasi-Monte Carlo Methods

Erik Hintz, Marius Hofert, and Christiane Lemieux, University of Waterloo, Canada

11:00-11:20 Goodness-of-fit Testing of Copulas using Quasi-Monte Carlo Methods

Yiran Chen, Florida State University, U.S.

continued in next column

Wednesday, February 27

MS190

Fast and Accurate Integral Methods for Highly Oscillatory Phenomena - Part I of II

9:45 a.m.-11:25 a.m.

Room: 202C

For Part 2 see MS224

Integral methods are useful tools in applied science and engineering. In particular, they are an important topic for large-scale scientific computing. Many challenges remain open and attract much attention especially in the high-frequency regime. This minisymposium focuses on recent advances in integral equations and integral transforms for highly oscillatory phenomena, including new formulations for high-frequency wave propagation, efficient and accurate discretizations, novel fast algorithms and their implementation based on locally rank-structured matrices and non-oscillatory phase functions, with applications in various imaging science and computational electromagnetism.

Organizer: Haizhao Yang

National University of Singapore, Singapore

Organizer: Yang Liu

Lawrence Berkeley National Laboratory, U.S.

9:45-10:05 Boundary Element Methods for High-intensity Focused Ultrasound Simulations

Timo Betcke, Pierre Gelat, Seyyed Reza Haqshenas, and Nader Saffari, University College London, United Kingdom

10:10-10:30 The Windowed Green Function Method and Other Specialized High-frequency Techniques in Computational Electromagnetism

Oscar P. Bruno, California Institute of Technology, U.S.

10:35-10:55 On Accurate Evaluation of Oscillatory Integrals and Functions

Gregory Beylkin, University of Colorado Boulder, U.S.

11:00-11:20 Generalized Plane Waves for the 3D Convected Helmholtz Equation

Lise-Marie Imbert-Gerard, University of Maryland, Baltimore, U.S.; Guillaume Sylvand, Airbus Group Innovations, France

continued in next column

Wednesday, February 27

MS191

New Challenges and Opportunities for Model Order Reduction - Part I of II

9:45 a.m.-11:25 a.m.

Room: 203

For Part 2 see MS225

In order to facilitate a fast simulation response, model order reduction techniques construct a low-dimensional, problem-adapted subspace from solutions at various time instances, different parameter values, or different boundary conditions. Examples comprise the Reduced Basis method, the Proper Orthogonal Decomposition, tensor-based approaches, or localized model order reduction approaches. In this minisymposium we want to discuss challenges and opportunities for model order reduction methods that arise for instance when applying model order reduction to data assimilation, parameter estimation and large input-parameter spaces, incorporating ideas originally developed in data science, or large-scale optimization.

Organizer: Laura Iapichino

Technische Universiteit Eindhoven, The Netherlands

Organizer: Kathrin Smetana

University of Twente, Netherlands

Organizer: Tommaso Taddei

INRIA Bordeaux, France

9:45-10:05 Reduced-order Models: Convergence Between Data and Simulation

Angelo Iollo, Institut de Mathématiques de Bordeaux, France; Michel Bergmann, Inria Bordeaux Sud-Ouest, France; Andrea Ferrero, Politecnico di Torino, Italy; Sebastien Riffaud, Inria and University of Bordeaux, France; Edoardo Lombardi, Angela Scardigli, and Haysam Telib, Optimad, Italy

continued on next page

Wednesday, February 27

MS191

New Challenges and Opportunities for Model Order Reduction - Part I of II

continued

10:10-10:30 Certified Reduced Basis Methods for Variational Data Assimilation

Sebastien J. Boyaval, École des Ponts ParisTech, France; Martin Grepl, Mark Kaercher, Nicole Nellesen, and *Karen Veroy*, RWTH Aachen University, Germany

10:35-10:55 Greedy Controllability of Reduced-order Linear Dynamical Systems

Giulia Fabrini, University of Konstanz, Germany; Laura Iapichino, Technische Universiteit Eindhoven, The Netherlands; Martin Lazar, University of Dubrovnik, Croatia; Stefan Volkwein, University of Konstanz, Germany

11:00-11:20 Adaptive Multiscale and Asynchronous Optimization Methods for Large-scale PDE Paramteter Estimation

Samy Wu Fung and Lars Ruthotto, Emory University, U.S.

Wednesday, February 27

MS192

Computational Strategies for High-dimensional Data Assimilation and Bayesian Inverse Problems - Part I of II

9:45 a.m.-11:25 a.m.

Room: 205

For Part 2 see MS226

Bayesian inverse problems, in which one estimates parameters of a model based on noisy data, and data assimilation problems, in which one sequentially estimates the state and parameters of a dynamical model using noisy observations, occur throughout science and engineering. Practical problems are often characterized by nonlinear numerical models, non-Gaussian distributions, model error, and extremely high dimensionality. In this session we present and discuss computational methods for solving inference problems in this setting, with an emphasis on Markov chain Monte Carlo, importance sampling, and ensemble methods for high-dimensional posterior distributions.

Organizer: Matthias Morzfeld

University of Arizona, U.S.

Organizer: Xin T. Tong

National University of Singapore, Singapore

Organizer: Youssef M. Marzouk

Massachusetts Institute of Technology, U.S.

9:45-10:05 Bayesian Inference with Nonparametric Likelihood Functions

Shixiao W. Jiang and John Harlim, Pennsylvania State University, U.S.

10:10-10:30 When is High-dimensional State Estimation Hard or Easy?

Chris Snyder, National Center for Atmospheric Research, U.S.

10:35-10:55 Accounting for Model Error in ENKF

Xuemin Tu, University of Kansas, U.S.; Alexandre Chorin, University of California, Berkeley, U.S.; Fei Lu, Johns Hopkins University, U.S.

11:00-11:20 A Unified Framework for Transportation Particle Filters

Peter Jan van Leeuwen, University of Reading, United Kingdom and Colorado State University, U.S.

Wednesday, February 27

MS193

Data-driven Methods in Fluid Dynamics - Part I of II

9:45 a.m.-11:25 a.m.

Room: 206A

For Part 2 see MS227

Data-fueled modeling and control of complex systems is currently undergoing a revolution, driven by the confluence of big data, advanced algorithms in machine learning, and modern computational power. Over the recent years, there has been an increasing focus on data-driven methods in fluid dynamics, e.g. for extraction of coherent flow structures, reduced-order modeling, control design in complex cost landscapes, sensor and actuator placement, and full-scale design optimization. This symposium explores the challenges and state-of-the-art innovations in data-driven methods for characterization, modeling, and control in challenging fluid flow problems, that include techniques such as dimensionality reduction, machine learning, operator-theoretic approaches, and control. In this minisymposium, experts will demonstrate achievements in the field, discuss open problems and challenges and inspire future directions.

Organizer: Kazuki Maeda

University of Washington, U.S.

Organizer: Eurika Kaiser

University of Washington, U.S.

9:45-10:05 Flow Reconstruction using Manifold Learning

Lionel Mathelin, CNRS, France; Srikanth Derebail Muralidhar and Berengere Podvin, LIMSI-CNRS, France

continued on next page

10:10-10:30 Optimal Actuator Selection for Airfoil Separation Control using Empirical Data

Maziar S. Hemati and Debraj Bhattacharjee, University of Minnesota, U.S.; Bjoern Klose and Gustaaf Jacobs, San Diego State University, U.S.

10:35-10:55 Physically-informed Bayesian Learning of Linear Embeddings for Fluid Dynamics Problems

Shaowu Pan and Karthik Duraisamy, University of Michigan, Ann Arbor, U.S.

11:00-11:20 Using Machine Learning for the Development of Closure Models for Multiphase Flows

Gretar Tryggvason, Johns Hopkins University, U.S.

Wednesday, February 27

MS194

Efficient Solvers for Coupled Multiphysics Problems - Part I of II

9:45 a.m.-11:25 a.m.

Room: 206B

For Part 2 see MS228

Modeling the simultaneous occurrence of different physical processes is an essential requirement for reliable simulation in many computational fluid and solid mechanics applications. The interaction between tightly coupled thermal, mechanical, chemical and flow processes can play a critical role. The discretization of these problems gives rise to linearized equations with an inherent block structure reflecting the underlying physical couplings. The development of efficient solution strategies, either monolithic or sequential, for such problems is an active field of research. This minisymposium aims to join scientists working on multi-physics problems, both from a purely algebraic and a more physics-based viewpoint, with the objective of sharing experiences and new ideas arising from a wide spectrum of applications.

Organizer: Massimiliano Ferronato

University of Padova, Italy

Organizer: Nicola Castelletto

Lawrence Livermore National Laboratory, U.S.

Organizer: Joshua A. White

Lawrence Livermore National Laboratory, U.S.

9:45-10:05 A Relaxed Physical Factorization Preconditioner for Mixed Finite Element Coupled Poromechanics

Matteo Frigo, University of Padova, Italy; Nicola Castelletto, Lawrence Livermore National Laboratory, U.S.; Massimiliano Ferronato, University of Padova, Italy

10:10-10:30 Robust Preconditioners for the Biot's Model

James H. Adler, Tufts University, U.S.; Francisco José Gaspar, University of Zaragoza, Spain; *Xiaozhe Hu* and Peter Ohm, Tufts University, U.S.; Carmen Rodrigo, University of Zaragoza, Spain; Ludmil Zikatanov, Pennsylvania State University, U.S.

10:35-10:55 Multiscale Computation of Multiphase Fluid Dynamics at the Pore Scale

Yashar Mehmani and Hamdi Tchelepi, Stanford University, U.S.

11:00-11:20 Flexible Monolithic Solvers for Dynamic Poroelastic Wave Propagation

Uwe Koecher, Helmut-Schmidt-University, Germany

continued in next column

Wednesday, February 27

MS195

CSE Education and Workforce - Part I of II

9:45 a.m.-11:25 a.m.

Room: 206C

For Part 2 see MS229
Featured Minisymposium

This minisymposium will highlight recent advances and future challenges in CSE education and workforce. Academic speakers will discuss new and thriving interdisciplinary programs in undergraduate and graduate CSE. Industry and National Lab speakers will discuss perspectives on CSE workforce needs from various industry and government sectors. Each talk slot will include time for an active discussion with minisymposium attendees.

Organizer: Serkan Gugercin
Virginia Tech, U.S.

Organizer: Karen E. Willcox
University of Texas at Austin, U.S.

9:45-10:05 CMDA: Math Modeling, Data Science, and HPC in Undergraduate Education
Serkan Gugercin, Virginia Tech, U.S.

10:10-10:30 A New Undergraduate Engineering Degree in Computational Engineering at UT Austin
Clint Dawson and Noel Clemens, University of Texas at Austin, U.S.

10:35-10:55 Applied and Computational Mathematics: BYU's New Degree for 21st Century Discovery and Innovation
Jeffrey Humpherys, Brigham Young University, U.S.

11:00-11:20 Experiences in CSE Graduate Education: 12 Years of AICES
Karen Veroy, RWTH Aachen University, Germany

Wednesday, February 27

MS196

Recent Developments of Numerical Methods for Hyperbolic and Parabolic Equations - Part I of II

9:45 a.m.-11:25 a.m.

Room: 206D

For Part 2 see MS230
This minisymposium is to bring people together to discuss the recent advances and exchange ideas in the algorithm design of high-order numerical methods for hyperbolic and parabolic equations and other high-order partial different equations, including the implementation, numerical analysis. In the minisymposium, the speakers will apply those high-order numerical methods to computational fluid, biology and physics, etc. This minisymposium is a good opportunity for people to discuss with researchers from different areas, and explore more applications and future research collaborations. We expect 8 speakers to present in this minisymposium.

Organizer: Yang Yang
Michigan Technological University, U.S.

9:45-10:05 Higher Order Globally Constraint-preserving FVTD and DGTD Schemes for Time-dependent Computational Electrodynamics
Dinshaw Balsara, Notre Dame University, U.S.

10:10-10:30 Local Discontinuous Galerkin Method for Convection-diffusion Equations on Overlapping Meshes
Jie Du, Tsinghua University, P. R. China; *Yang Yang*, Michigan Technological University, U.S.; Eric Chung, The Chinese University of Hong Kong, Hong Kong

10:35-10:55 High-order Bound-preserving Discontinuous Galerkin Methods for Compressible Miscible Displacements in Porous Media on Triangular Meshes
Nattaporn Chuenjarern, Yang Yang, and Ziyao Xu, Michigan Technological University, U.S.

11:00-11:20 An Alternative Formulation of Discontinous Galerkin Schemes for Solving Hamilton-Jacobi Equations
Wei Guo, Texas Tech University, U.S.

continued in next column

Wednesday, February 27

MS197

Task-based Programming for Scientific Computing: Linear Algebra Applications - Part I of II

9:45 a.m.-11:25 a.m.

Room: 207

For Part 2 see MS231

The evolution of modern architectures has led to the development of numerous new programming models. These models aims at simplifying the work of the programmer to help with performance portability among various architectures, but also in enabling more asynchronism in applications. Among all solutions task-based runtimes systems and language extensions provide a large spectrum of features to the programmer to adapt to various applications. In this mini-symposium, we propose to study a set of modern linear algebra solvers that exploits these expressive models to reach high performance, and or express new algorithms.

Organizer: Mathieu Faverge
Bordeaux INP, Inria, LaBRI, France

Organizer: Mathias Jacquelin
Lawrence Berkeley National Laboratory, U.S.

9:45-10:05 Hierarchical Algorithms on Hierarchical Architectures

Kadir Akbudak, Ali M. Charara, *David E. Keyes*, Hatem Ltaief, and Aleksandr Mikhalev, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

10:10-10:30 Exploiting Nested Task-based Parallelism in the Factorization of Hierarchical Matrices

Rocio Carratalá-Sáez and Enrique S. Quintana-Ortí, Universidad Jaume I, Spain

10:35-10:55 Simulation of a Sparse Direct Solver on Heterogeneous Systems using Starpu and Simgrid

Emmanuel Agullo, Inria, France; *Alfredo Buttari*, CNRS, France; Ian Masliah, INRIA Bordeaux, France

11:00-11:20 Exploiting Parameterized Task-graph in Sparse Direct Solvers

Mathieu Faverge, Bordeaux INP, Inria, LaBRI, France; Gregoire Pichon, Inria, France; Pierre Ramet, Université de Bordeaux, Inria, LaBRI, France

Wednesday, February 27

MS198

Surrogate Modeling and Data Compression for Exascale Applications - Part I of II

9:45 a.m.-11:25 a.m.

Room: 401A

For Part 2 see MS232

High performance computing systems are expected to reach the exascale in the near future, requiring a new generation of simulation tools and introducing new challenges arising from high-dimensional, extremely large data sets and the highly-distributed nature of exascale systems. This minisymposium will emphasize surrogate modeling and data compression techniques applicable to exascale applications and requiring scaling across many thousands of cores or to large datasets. Talks that address topics in data compression, dimension reduction, in situ visualization, surrogate modeling, or contain applications targeting exascale computing resources are encouraged.

Organizer: Matthew J. Reynolds
National Renewable Energy Laboratory, U.S.

Organizer: Ryan King
National Renewable Energy Laboratory, U.S.

Organizer: Alireza Doostan
University of Colorado Boulder, U.S.

9:45-10:05 Scenario Selection for Reliable Operations and Planning of Large Power Grids

Matthew J. Reynolds, National Renewable Energy Laboratory, U.S.

10:10-10:30 Exploiting Ridge Structure in Chance Constrained Design under Uncertainty

Jeffrey M. Hokanson and Paul Constantine, University of Colorado Boulder, U.S.

10:35-10:55 Data Compression Based on Orthogonal Transforms, Impact on Post-processing of Large Data Sets

Oana Marin, Argonne National Laboratory, U.S.

11:00-11:20 Low-rank Tucker Decomposition of Large Tensors using TensorSketch

Osman Asif Malik and Stephen Becker, University of Colorado Boulder, U.S.

Wednesday, February 27

MS199

Recent Developments in Theory and Implementation of Subspace Correction Methods

9:45 a.m.-11:25 a.m.

Room:401B

The method of subspace corrections offers a very general framework in which to analyse solvers, preconditioners, and smoothers for a wide class of problems. Variants of this method have been successfully used to demonstrate robust solvers for the Navier-Stokes equations, Maxwell equations, and many others. Despite this, reusable, robust, library interfaces that take advantage of these techniques are in short supply. The challenge in designing such libraries is to balance the diverse requirements of individual, problem-specific, methods with the goal of providing clean, usable interfaces. The aim of this minisymposium is to bring together theory and practice of advanced multilevel and domain decomposition solvers in order to facilitate discussion and development of usable, robust solvers for wide classes of difficult problems.

Organizer: Florian Wechsung
University of Oxford, United Kingdom

Organizer: Lawrence Mitchell
Durham University, United Kingdom

9:45-10:05 A New Configurable Preconditioner for Additive Subspace Correction Methods in PETSc

Lawrence Mitchell, Durham University, United Kingdom; Patrick E. Farrell, University of Oxford, United Kingdom; Rob C. Kirby, Baylor University, U.S.; Matthew G. Knepley, State University of New York at Buffalo, U.S.

10:10-10:30 Nested Solvers for Improved Scalability of Domain Decomposition Methods

Pierre Jolivet, CNRS, France

continued on next page

Wednesday, February 27

MS199

Recent Developments in Theory and Implementation of Subspace Correction Methods

continued

10:35-10:55 BDDC and FETI-DP Methods in PETSc
Stefano Zampini, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

11:00-11:20 A Scalable Solver for the Stationary Navier-Stokes Equations in 3D Based on Augmented Lagrangian and Subspace Correction Methods
Florian Wechsung, University of Oxford, United Kingdom

Wednesday, February 27

MS200

Sparse Function Approximations: Theory and Applications - Part I of II

9:45 a.m.-11:25 a.m.

Room: 401C

For Part 2 see MS234

Learning non-linear systems from noisy, limited, and/or dependent data is an important task for data analysis across various many scientific fields including statistic, engineering, computer science, mathematics, and many more. One of the major paradigms is to learn an unknown generating function from a set of input-output pairs, which can be rephrased as finding an appropriate approximation to a high-dimensional function. Without additional restrictions or structures, this learning task is ill-posed. This minisymposium addresses some recent developments of several learning techniques, which leverage structures such as sparsity, randomness, dimensionality, etc, in order to make the learning well-posed. The presentations will cover analysis based on orthogonal polynomials, sampling, approximation theory, probability theory, as well as novel computational techniques.

Organizer: Giang Tran
University of Waterloo, Canada

Organizer: Rachel Ward
University of Texas at Austin, U.S.

Organizer: Hayden Schaeffer
Carnegie Mellon University, U.S.

9:45-10:05 Sparse Recovery and Outlier Detection for Dependent Data
Giang Tran, University of Waterloo, Canada; Rachel Ward, University of Texas at Austin, U.S.; Hayden Schaeffer, Carnegie Mellon University, U.S.

10:10-10:30 Sparse Polynomial Approximation on Irregular Domains
Ben Adcock, Simon Fraser University, Canada; Daan Huybrechs, KU Leuven, Belgium; Mohsen Seifi, Simon Fraser University, Canada

10:35-10:55 Learning Functions with Low Complexity in High Dimensions
Wenjing Liao, Georgia Institute of Technology, U.S.; Mauro Maggioni and Stefano Vigogna, Johns Hopkins University, U.S.

11:00-11:20 Analysis of Sparse Recovery for Legendre Expansions using Envelope Bound
Hoang A. Tran, Oak Ridge National Laboratory, U.S.; Clayton G. Webster, University of Tennessee and Oak Ridge National Laboratory, U.S.

continued in next column

Wednesday, February 27

MS201

Theoretical and Computational Aspects in Nonlocal and Material Science Modeling - Part I of II

9:45 a.m.-11:25 a.m.

Room: 402A

For Part 2 see MS333

The mathematics of material science encompasses a variety of aspects, from formulating models that describe material properties and phenomena, to theoretical investigations, and numerical simulations. A variety of models describing local and nonlocal effects have been shown successful in applications, their investigation and implementation in real-world applications is ongoing attracting the interest of a growing community. Nonlocal models such as anomalous diffusion and peridynamics can incorporate temporal and spatial nonlocal behaviors due to memory effects and long-range forces. For problems where these effects cannot be neglected, such descriptions are more accurate than PDEs and have recently gained popularity. Mathematical analysis and numerical solutions for nonlocal problems remain challenges, requiring new approaches to design robust, efficient and predictive numerical methods. The goal of this minisymposium is to bring together researchers working on different applications in material science to present and discuss novel approaches related to theoretic nonlocal models and present approaches for their numerical solution.

Organizer: Marta D'Elia

Sandia National Laboratories, U.S.

Organizer: Petronela Radu

University of Nebraska, Lincoln, U.S.

9:45-10:05 Theoretical and Numerical Aspects for Doubly Nonlocal Problems on Bounded Domains

Petronela Radu, University of Nebraska, Lincoln, U.S.

10:10-10:30 Investigation of Dispersion Relation of the 1D Quasi-nonlocal Coupling for Nonlocal and Local Mechanics

Helen Li, University of North Carolina, Charlotte, U.S.

10:35-10:55 Theoretical Aspects of Nonlocal Helmholtz Decomposition of a Vector Field

Marta D'Elia, Sandia National Laboratories, U.S.; *Cynthia Flores*, California State University, Channel Islands, U.S.; Xingjie Li, University of North Carolina, Charlotte, U.S.; Petronela Radu, University of Nebraska, Lincoln, U.S.; Yue Yu, Lehigh University, U.S.

11:00-11:20 Numerical Aspects of Nonlocal Helmholtz Decomposition of a Vector Field

Yue Yu, Lehigh University, U.S.; Marta D'Elia, Sandia National Laboratories, U.S.; Cynthia Flores, California State University, Channel Islands, U.S.; Xingjie Li, University of North Carolina, Charlotte, U.S.; Petronela Radu, University of Nebraska, Lincoln, U.S.

continued in next column

Wednesday, February 27

MS202

Nonlocal Models in Computational Science and Engineering - Part I of II

9:45 a.m.-11:25 a.m.

Room: 402B

For Part 2 see MS236

Nonlocal models provide a new framework to overcome limitations and challenges present in classical PDE-based models. For instance, peridynamics, a nonlocal extension of classical continuum mechanics, admits discontinuous solutions and naturally describes material failure and damage. Similarly, nonlocal and fractional diffusion models can represent anomalous diffusion and heat transfer. Furthermore, nonlocal models introduce length scales, which can be used for multiscale modeling. Recent years have witnessed a tremendous advance in modeling, mathematical analysis, and computational practice for nonlocal problems. This minisymposium focuses on recent developments in peridynamics, nonlocal and fractional diffusion, and other related nonlocal models. The topics of interest include and are not limited to: peridynamics, nonlocal and fractional mass diffusion and heat transfer, discretization and time-stepping methods for nonlocal models, nonlocal boundary conditions, multiscale methods and multi-physics modeling with nonlocal models, nonlocal modeling of heterogeneous systems, engineering and scientific applications with nonlocal models.

Organizer: Yue Yu

Lehigh University, U.S.

Organizer: Pablo Seleson

Oak Ridge National Laboratory, U.S.

9:45-10:05 Foundations of Material Anisotropy in Classical Elasticity and Peridynamics

Pablo Seleson and Jeremy Trageser, Oak Ridge National Laboratory, U.S.; Max Gunzburger, Florida State University, U.S.

continued on next page

Wednesday, February 27

MS202

Nonlocal Models in Computational Science and Engineering - Part I of II

continued

10:10-10:30 Subsurface Applications for Peridynamics
Michael L. Parks, Sandia National Laboratories, U.S.

10:35-10:55 Finite Element and Finite Difference Modeling and Convergence for State Based Peridynamic Fracture
Robert P. Lipton, Louisiana State University, U.S.

11:00-11:20 Asymptotically Compatible Foundations for Nonlocal Mechanics
Nathaniel Trask, Sandia National Laboratories, U.S.

Wednesday, February 27

MS203

Generalizable Machine Learning and Clustering Methods: Applications in Bias, Biology, and Imaging

9:45 a.m.-11:25 a.m.

Room: 402C

For Part 2 see MS169
For Part 4 see MS237
Part of the SIAM Workshop Celebrating Diversity

As the implementation of machine learning methods become more commonplace to answer questions in science and society, the need for generalizable methods have increased. At the same, these methods should be approached with appropriate domain knowledge of the particular research question and understanding of any underlying bias. The talks in this session will present advances in these methods in the context of bias, biology, and imaging applications. The first talk will address optimization and machine learning methods to detect genomic variation in related individuals. The second communication will incorporate data from intensive care units (ICU) to assess bias in mortality prediction. The next presentation will use topological data analysis to improve predictions from ECG signals. The final talk will develop a new density-based clustering method in the context of image segmentation.

Organizer: Mario Banuelos
California State University, Fresno, U.S.

Organizer: Talea Mayo
University of Central Florida, U.S.

Organizer: Shelby Wilson
Morehouse College, U.S.

9:45-10:05 Predicting Genomic Variation in Error-prone Data Regimes
Mario Banuelos, California State University, Fresno, U.S.

10:10-10:30 Assessing Bias in the Prediction of Mortality in the ICU
Lara Reichmann, University of San Francisco, U.S.

10:35-10:55 Classification of Short ECG Readings: TDA Informed Machine Learning
David T. Uminsky, University of San Francisco, U.S.

11:00-11:20 A Consistent Density-based Clustering Algorithm and its Application to Microstructure Image Segmentation
Marilyn Vazquez, ICERM, U.S.

continued in next column

Wednesday, February 27

CP11

Numerical Linear Algebra II

9:45 a.m.-11:25 a.m.

Room: 201A

Chair: Christophe Audouze, University of Toronto, Canada

9:45-10:00 Sparse Low-rank Separated Representation Models for Regression

Christophe Audouze and Prasanth B. Nair, University of Toronto, Canada

10:05-10:20 High Performance Qr Factorization of Ill-conditioned Matrices Based on the Cholesky Qr Algorithm

Takeshi Fukaya, Hokkaido University, Japan; Ramaseshan Kannan, Arup, Manchester, United Kingdom; Yuji Nakatsukasa, National Institute of Informatics, Japan; Yusaku Yamamoto, University of Electro-Communications, Japan; Yuka Yanagisawa, Waseda University, Japan

10:25-10:40 Communication-optimal Qr Decomposition Using Task-based Parallelism

Heather Pacella and Gianluca Iaccarino, Stanford University, U.S.

10:45-11:00 Comparing Randomized and Deterministic Approaches for Computing Low-rank Approximations of Large Matrices

Carlos Pachajoa, Viktoria Mayer, and Wilfried N. Gansterer, University of Vienna, Austria

11:05-11:20 A Partial Singular Value Decomposition Algorithm and its Implementation on Distributed-memory Systems.

David E. Keyes and Hatem Ltaief, King Abdullah University of Science & Technology (KAUST), Saudi Arabia; Yuji Nakatsukasa, University of Oxford, United Kingdom; *Dalal Sukkari*, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

Wednesday, February 27

CP12

Mathematical Optimization II

9:45 a.m.-11:25 a.m.

Room: 202A

Chair: Silke Glas, University of Ulm, Germany

9:45-10:00 Scalable Incremental Nonconvex Optimization Approach for Phase Retrieval from Minimal Measurements

Ji Li, Beijing Computational Science Research Center, China; Jian-Feng Cai, Hong Kong University of Science and Technology, Hong Kong; Hongkai Zhao, University of California, Irvine, U.S.

10:05-10:20 Advances in Greedy Optimization Algorithm in Locating Mfs Nodes and Collocation Points

Aidan Doyle, Anton Nelson, and Matthew Yuan, United States Military Academy, U.S.; Noah DeMoes and Bryce Wilkins, Massachusetts Institute of Technology, U.S.; Kameron Grubaugh, Unaffiliated; Theodore Hromadka, United States Military Academy, U.S.

10:25-10:40 Efficient Multilevel Methods for the Optimal Control of Large-scale Stochastic Partial Differential Equations.

Sumaya Alzuhairy, Andrei Draganescu, and Bedrich Sousedik, University of Maryland, Baltimore County, U.S.

10:45-11:00 Model Reduction for Hamilton-Jacobi-Bellman Equations Resulting from Intraday Trading of Electricity

Silke Glas and Karsten Urban, University of Ulm, Germany

11:05-11:20 Robust Residual-based and Residual-free Greedy Algorithms for Reduced Basis Methods

Jiahua Jiang, Virginia Tech, U.S.; Yanlai Chen, University of Massachusetts, Dartmouth, U.S.; Akil Narayan, University of Utah, U.S.

Wednesday, February 27

SIAM Workshop Celebrating Diversity and Broader Engagement Joint Lunch (by invitation)

11:30 a.m.-1:00 p.m.

Room: Davenport Hotel -- Cedar Ballroom

SP1 SIAG CSE Best Paper Prize Lecture

11:30 a.m.-12:00 p.m.

Room: Ballroom 100BC

Recipient to be announced.

Lunch Break

12:00 p.m.-1:00 p.m.

Wednesday, February 27

IP6

Stochastic Gradient Descent, in Theory and Practice

1:00 p.m.-1:45 p.m.

Room:Ballroom 100BC

Chair: Stefan Wild, Argonne National Laboratory, U.S.

Stochastic Gradient Descent (SGD) is a first-order stochastic optimization algorithm which is the algorithm of choice behind powerful deep learning architectures which are becoming increasingly omnipresent in society. Still, there remains is a wide gap between the setting where SGD theoretical guarantees and the setting where SGD is most effective and useful in practice. We discuss recent theoretical results which make progress towards closing this gap. First, we present the first theoretical guarantees for "adaptive learning rate" SGD algorithms such as AdaGrad which are used widely in practice, as they make SGD less sensitive to choice of hyperparameters such as the step-size, but which are challenging to analyze theoretically as they are non-linear stochastic dynamical systems. Second, we provide new guarantees for the convergence of SGD to global minimizers of certain non-convex optimization problems via a novel analysis for products of independent random matrices. We conclude by discussing several open problems.

Rachel Ward

University of Texas at Austin, U.S.

Coffee Break

1:45 p.m.-2:15 p.m.

Room:Ballroom Foyer

Wednesday, February 27

MS204

Exploiting Model Hierarchies, Sparsity and Low-rank Structure of Large-scale Bayesian Computation - Part II of II

2:15 p.m.-3:55 p.m.

Room:Ballroom 100BC

For Part 1 see MS170

One of the central tasks in modern computational science and engineering is to quantify the uncertainty associated with the computational results. This can be broadly classified as Bayesian Computation. Frequently these tasks involve expensive or intractable numerical models, embedded in likelihood functions that involve high-dimensional parameters and/or high-dimensional data sets. The output is typically the posterior (post data) distribution of parameters or quantities of interest, or statistics of those parameters or quantities. Markov chain Monte Carlo, sequential Monte Carlo, and other posterior exploration schemes require repeated evaluations of such models, in principle over high-dimensional spaces. In this setting, standard algorithms quickly become intractable. Methods for identifying and exploiting model hierarchies, low-rank structure and sparsity in the representation of the posterior distribution are becoming essential for solving these otherwise intractable Bayesian inference problems. This minisymposium will bring together researchers working on the forefront of structure-exploiting methods intended to accelerate such large-scale Bayesian computation.

Organizer: Tiangang Cui
Monash University, Australia

Organizer: Robert Scheichl
Universität Heidelberg, Germany

2:15-2:35 Consistent Parameter Distributions in High-dimensional Spaces: Built One Dimension at a Time

Troy Butler, University of Colorado, Denver, U.S.

2:40-3:00 Low-rank Structure in Optimization-based Sampling

Zheng Wang and Youssef M. Marzouk, Massachusetts Institute of Technology, U.S.; Tiangang Cui, Monash University, Australia

3:05-3:25 Mitigating the Cost of PDE-constrained Bayesian Inverse Problems Using Dimensionality Reduction and Machine Learning

Sheroze Sheriffdeen, Institute for Computational Engineering and Sciences, U.S.

3:30-3:50 Hierarchical Off-diagonal Low-rank Approximation for Hessians in Bayesian Inference

Tucker Hartland, University of California, Merced, U.S.; Ruanui Nicholson, University of Auckland, New Zealand; Noemi Petra, University of California, Merced, U.S.

continued in next column

Wednesday, February 27

MS205

Progress and Challenges in Extreme Scale Computing and Data - Part II of II

2:15 p.m.-3:55 p.m.

Room:Conference Theater

For Part 1 see MS171

Extreme scale computing efforts have resulted in numerous advances for multicore, manycore and accelerator based scalable systems. In addition, large-scale applications must increasingly deal with data management and analysis as a first-class concern. Therefore, new applications often have to manage distributed and parallel computing, and have to manage workflows of different tasks (computing, data analytics, machine learning, visualization,…). In this minisymposium we present some of the latest works in scalable algorithms, programming paradigms, and libraries for next generation computing platforms. Furthermore, we discuss efforts to better incorporate data science concerns as a principle component of our scientific workflows.

Organizer: Serge Petiton

University of Lille, France

Organizer: Michael A. Heroux

Sandia National Laboratories, U.S.

Organizer: Kengo Nakajima

University of Tokyo, Japan

2:15-2:35 Recent Trends and Challenges for High Performance Sparse Linear Algebra

Michael A. Heroux, Sandia National Laboratories, U.S.

2:40-3:00 A Parallel Time Integration Approach without Re-discretization in Time Direction

Akihiro Fujii and Teruo Tanaka, Kogakuin University, Japan; Takeshi Iwashita, Hokkaido University, Japan

3:05-3:25 On Performances Portability: Illustration on Cross-section Calculations for Monte Carlo Neutron Transport

Christophe Calvin, CEA Saclay, France; Emeric Brun, Fausto Malvagi, and Tao Chang, CEA, DEN, SRMP, France

3:30-3:50 H-Matrix Framework for Many-core Processors

Tetsuya Hoshino and Akihiro Ida, University of Tokyo, Japan

continued in next column

Wednesday, February 27

MS206

Advanced HPC Trends for Oil and Gas Applications - Part II of II

2:15 p.m.-3:55 p.m.

Room:102A

For Part 1 see MS172

For decades, advances in high performance computing have led to more accurate, safer, and faster oil and gas exploration and production processes. Applications range from seismic imaging and reservoir simulations to seismic interpretation and digital rock physics. This evolution is paramount today to further develop production by exploring increasingly complex reservoirs and by enhancing the recovery ratios from existing fields. In fact, these requirements translate into larger volumes of data to process and models of higher fidelity in terms of physical formulation and space and time resolutions. To cope with these challenges, the oil and gas industry has to adapt constantly to a changing technology landscape in terms of algorithms, platforms, software, and tools. The move towards less synchrony at all system levels, the widening gap between cost of IO versus floating-point operations, and the recent advent of Deep Learning are a few of these important changes the oil and gas software industry has to consider in moving forward with exascale. This minisymposium is an opportunity to discuss today's and future trends defining oil and gas HPC applications design.

Organizer: Rached Abdelkhalak

King Abdullah University of Science & Technology (KAUST), Saudi Arabia

Organizer: Gabriel Fabien-Ouellet

Polytechnique Montreal, Canada

continued on next page

Wednesday, February 27

MS206

Advanced HPC Trends for Oil and Gas Applications - Part II of II

continued

2:15-2:35 Smoothing Data Movement Between Ram and Storage for Reverse Time Migration

Tariq Alturkestani, King Abdullah University of Science and Technology (KAUST), Saudia Arabia and University of Pittsburgh, U.S.; Thierry Tonellot, Saudi Aramco Oil Company, Saudi Arabia; Hatem Ltaief, King Abdullah University of Science & Technology (KAUST), Saudi Arabia; Vincent Etienne, Saudi Aramco Oil Company, Saudi Arabia; David E. Keyes, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

2:40-3:00 Automation of Velocity Analysis with Recurrent Neural Networks

Gabriel Fabien-Ouellet, Polytechnique Montreal, Canada; Rahul Sarkar, Stanford University, U.S.

3:05-3:25 Deep Learning in an Industrial Context: an Introduction to the NVIDIA Ecosystem

Issam Said, NVIDIA, U.S.

3:30-3:50 Fossil Segmentation and Classification using 3D U-NET Based Architecture

Amine Kerkeni, InstaDeep, Tunisia; Youssef Ben Dhia and Karim Beguir, InstaDeep, U.K.; Noomane Keskes and Maxime Bellay, Total, France

Wednesday, February 27

MS207

Advances and Applications in Numerical Methods for Interfacial Flows - Part II of II

2:15 p.m.-3:55 p.m.

Room: 102C

For Part 1 see MS174

Flows with multiple components separated by fluid interfaces are ubiquitous in nature and engineering. In the past tens of years, vast progress has been made in the numerical simulation of interfacial flows and many outstanding methods have been developed, eg., the volume-of-fluid, level-set, front-tracking, and phase-field methods. However, challenges, such as the coupling between fluid interface and complex rheology, moving contact line, and accurate and robust discretization, still remain. This minisymposium aims to bring together experts in this area to discuss new developments across different methods and identify new application areas.

Organizer: Pengtao Yue
Virginia Tech, U.S.

Organizer: Shahriar Afkhami
New Jersey Institute of Technology, U.S.

2:15-2:35 Stabilized IEQ/SAV Approach for Gradient Flow System with Strong Spatial Anisotropy

Xiaofeng Yang, University of South Carolina, U.S.

2:40-3:00 Modelling Moving Contact Lines in Multi-phase Fluids

Weiqing Ren, National University of Singapore and IHPC, Singapore

3:05-3:25 An Interface-preserving Level-set Method for Interfacial Flows with Contact Lines

Jiaqi Zhang and Pengtao Yue, Virginia Tech, U.S.

3:30-3:50 Domain Decomposition Solvers and Preconditioners for the Implicit Closest Point Method

Ian May, Simon Fraser University, Canada; Ronald Haynes, Memorial University, Newfoundland, Canada; *Steven Ruuth*, Simon Fraser University, Canada

Wednesday, February 27

MS208

Numerical Methods for Multi-Material Fluid Flows - Part II of II

2:15 p.m.-3:55 p.m.

Room: 102D

For Part 1 see MS175

The minisymposia aims to bring together researchers from universities and research laboratories to discuss the state-of-the-art in multi-material hydrodynamics simulations. It focuses on relevant numerical methods, on the analysis of such methods and on the modeling of complex multi-material flows. Topice include Lagrangian Hydrodynamics, Arbitrary Lagrangian Eulerian (ALE) Methods, Eulerian Methods, Mesh Generation Methods, including Mesh Adaptation, Interface Reconstruction Methods, Data Transfer Between Meshes—Remapping, Advanced Discretization Methods, High-order Methods.

Organizer: Mikhail Shashkov
Los Alamos National Laboratory, U.S.

2:15-2:35 On a High-order Lagrangian Discontinuous Galerkin Hydrodynamic Method

Nathaniel Morgan, Xiaodong Liu, and Donald E. Burton, Los Alamos National Laboratory, U.S.

2:40-3:00 Gradient-based Nonlinear Stabilization for Finite Element Discretizations of Conservation Laws

Dmitri Kuzmin, Technische Universität Dortmund, Germany

3:05-3:25 High Order Direct ALE Schemes on Moving Voronoi Meshes with Changing Connectivity

Elena Gaburro, University of Trento, Italy; Walter Boscheri, Free University of Bolzen-Bolzano, Italy; Michael Dumbser, University of Trento, Italy; Christian F. Klingenberg, Wurzburg University, Germany

3:30-3:50 High-order Matrix-free Hyper-viscosity for Arbitrary Lagrangian-Eulerian Hydrodynamics

Robert Rieben, Lawrence Livermore National Laboratory, U.S.; Pedro David Bello-Maldonado, University of Illinois at Urbana-Champaign, U.S.

Wednesday, February 27

MS209

Machine Learning Strategies for Computer Simulation of Physical Systems - Part II of II

2:15 p.m.-3:55 p.m.

Room:111A

For Part 1 see MS176

The bidirectional relation between computational models and data necessitates commensurate advances in machine learning algorithms and tools. This minisymposium aims to go beyond the parametric calibration of models and to address questions that pertain to the discovery of new models from computationally or experimentally-generated data. This need arises often in several instances of computational modeling as for example in uncertainty quantification where the construction of reduced-order or surrogate models is imperative in order to overcome the huge number of calls to highly-complex and expensive models. Furthermore, similar problems appear in the context of multiscale simulations where high-resolution, computationally-generated data over relatively small spatial domains and time horizons must be used in order to construct effective models of the system's behavior over larger spatiotemporal scales. In both of the aforementioned cases, significant challenges relate to high-dimensionality of inputs and outputs which usually exceed those encountered in typical machine learning applications, the discovery of meaningful and physically-interpretable collective variables and structural features of the models as well as the incorporation of physical constraints and invariances whenever these are available.

Organizer: Phaedon S. Koutsourelakis

Technische Universität München, Germany

2:15-2:35 Learning Coarse-grained Models from Molecular Dynamics

Markus Schoeberl, Technische Universität München, Germany; Nicholas Zabaras, University of Notre Dame, U.S.; Phaedon-Stelios Koutsourelakis, Technische Universität München, Germany

2:40-3:00 Learning Strategies for Generating Low-rank Representations of Physical Systems

Alex Gorodetsky, University of Michigan, U.S.

3:05-3:25 Theory-guided Data Science: A New Paradigm for Scientific Discovery from Data

Anuj Karpathe, University of Minnesota, U.S.

3:30-3:50 Physics-aware and Sparse Coarse-graining of Multiscale Dynamics

Sebastian Kaltenbach and Phaedon-Stelios Koutsourelakis, Technische Universität München, Germany

continued in next column

Wednesday, February 27

MS210

Machine Learning Approaches for the Sciences and Engineering: Recent Developments - Part II of II

2:15 p.m.-3:55 p.m.

Room: 111B

For Part 1 see MS177

This session brings together researchers and practitioners to present recent results on emerging data-driven approaches and machine learning methods, with an emphasis on their adaptation for use in the sciences and engineering. There have been many recent successes in areas, such as Natural Language Processing (NLP) and Image Processing (IP). However, problems encountered in the sciences and engineering disciplines often have rather different requirements and structure than those of NLP/IP. For adoption in these new problem domains, further mathematical and computational developments are needed for effective use of machine learning methods. In this session, we bring together researchers working on specific applications that would benefit from such data-driven/ML approaches, as well as, investigators working on fundamental new developments in emerging areas of machine learning.

Organizer: Ben J. Gross

University of California, Santa Barbara, U.S.

Organizer: Paul J. Atzberger

University of California, Santa Barbara, U.S.

2:15-2:30 Training Domain Expert Neural Networks with Weak Supervised Learning

Garrett Goh, Pacific Northwest National Laboratory, U.S.

2:35-2:50 Neural Network Surrogates of PDE-based Dynamical Systems, Application to Ice Sheet Dynamics

Mauro Perego, Sandia National Laboratories, U.S.

2:55-3:10 Interpretability and Representation of Data in Recurrent Neural Networks

Eli Shlizerman, University of Washington, Seattle, U.S.

3:35-3:50 Multilevel Adaptive Reduction of Data for Large Scale Scientific Simulation

Ben Whitney, Brown University, U.S.

Wednesday, February 27

MS211

Recent Advances in Multilevel Solvers - Part II of II

2:15 p.m.-3:55 p.m.

Room: 111C

For Part 1 see MS178

For many years, multilevel solvers have played a central role in many areas of computer simulation, including fluid dynamics, solid mechanics, computer graphics, geophysical exploration, and more. As a result, development and analysis of multilevel solvers remains an important area of research in computational science and engineering. In this minisymposium, we bring together speakers focused on several aspects of multilevel solvers, including developing algorithms for coupled multiphysics problems, analysis of algorithms, and graph problems.

Organizer: James H. Adler
Tufts University, U.S.

Organizer: Scott Maclachlan
Memorial University, Newfoundland, Canada

Organizer: Xiaozhe Hu
Tufts University, U.S.

2:15-2:35 Computing the Eigenvalues of the Dirichlet-to-Neumann Map for Indefinite Helmholtz Equation
Yangqingxiang Wu, Pennsylvania State University, U.S.

2:40-3:00 Randomized Solvers
John Urschel, Massachusetts Institute of Technology, U.S.

3:05-3:25 Efficient Operator-coarsening Multigrid Schemes for Local Discontinuous Galerkin Methods
Daniel Fortunato and Chris H. Rycroft, Harvard University, U.S.; Robert Saye, Lawrence Berkeley National Laboratory, U.S.

3:30-3:50 Local Fourier Analysis of BDDC-like Algorithms
Yunhui He, Memorial University, Newfoundland, Canada; Jed Brown, University of Colorado Boulder, U.S.; Scott Maclachlan, Memorial University, Newfoundland, Canada

Wednesday, February 27

MS212

High-order Finite Element Methods for Complex and Multiphysics Applications - Part II of II

2:15 p.m.-3:55 p.m.

Room: 300A

For Part 1 see MS179

A current challenge before the applied mathematics and computational science community is the solution of complex and multiphysics systems. Science and engineering applications for these systems are highly-nonlinear and have strongly coupled physical mechanisms that interact on a wide range of length- and time-scales and require robust and efficient high-resolution numerical approximations. For this reason, developing robust, and accurate methods that can effectively use parallel computation at extreme scales is critical. In this context numerical discretizations and solvers for practical multiphysics simulations must be: (1) High-order accurate in space and time; (2) Stable; (3) Conservative; (4) Having minimum degrees of freedom for implicit solution approaches; (5) Well suited for unstructured meshes; (6) Well suited for hp-adaptivity; (7) Well suited for applications with disparate temporal and spatial scales; and (8)Well suited for fine-grain parallelism. This minisymposium focuses on the latest developments in high(er) order finite element, and related methods and associated numerical methods. The speakers in this minisymposium will address theoretical/numerical and computational issues that are critical to developing approaches with these desired properties. Applications will include aerodynamics, magnetohydrodynamics, plasma physics, subsurface flows, geophysical flows, etc.

Organizer: Sriramkrishnan Muralikrishnan
University of Texas at Austin, U.S.

Organizer: Shinhoo Kang
University of Texas at Austin, U.S.

Organizer: Tan Bui
University of Texas at Austin, U.S.

2:15-2:35 Wind Turbine Simulation with High-order Hybridized Discontinuous Galerkin Methods
Shinhoo Kang and Tan Bui-Thanh, University of Texas at Austin, U.S.

2:40-3:00 Enabling Large-scale Wind Farm Simulations Using An Hp-Adaptive Discontinuous-Galerkin Method and Overset Grids
Michael Brazell and Dimitri Mavriplis, University of Wyoming, U.S.

3:05-3:25 A New Adaptive Block ILU Preconditioning Technique with Application in High-order Simulations of Three-dimensional Aircrafts
Behzad Ahrabi and Dimitri Mavriplis, University of Wyoming, U.S.

3:30-3:50 Nonhydrostatic Atmospheric Models using Dynamically Adaptive Continuous and Discontinuous Galerkin Methods
Jeremy E. Kozdon and Francis X. Giraldo, Naval Postgraduate School, U.S.; Michal Kopera, Boise State University, U.S.; Lucas Wilcox, Naval Postgraduate School, U.S.

continued in next column

Wednesday, February 27

MS213

Multiphysics and Multiscale Problems in Computational Science and Engineering - Part II of II

2:15 p.m.-3:55 p.m.

Room: 300B

For Part 1 see MS180

Multiphysics and multiscale problems are challenging algorithmically and taxing computationally. Since the interface or boundary is constantly changing nonlinearly, accurate representation in a stable manner is not trivial, not to mention nonlinearities in the governing equations. Hence, high order and high-fidelity discretization methods combined with accurate and stable coupling strategies are needed, and analysis tools as well as acceleration (stabilization) techniques are of critical importance. In this minisymposium, we will review a number of new ideas for developing efficient and reliable coupling algorithms and share interesting applications in various engineering and science fields. The topics of interest include and are not limited to: fluid–structure interaction, atomistic-to-continuum coupling problems, nonlocal-to-nonlocal/local coupling problems, quantum mechanics, micro-mechanics of materials, concurrent coupling approaches, multiphysics/ multiscale coupling methods.

Organizer: Xingjie Li

University of North Carolina, Charlotte, U.S.

2:15-2:35 A Neumann-type Boundary Condition for Peridynamics and its Coupling with Fluids

Yue Yu and *Huaiqian You*, Lehigh University, U.S.; Nathaniel Trask, Sandia National Laboratories, U.S.; Ming-Chen Hsu, Iowa State University, U.S.

2:40-3:00 Understanding the Tunable Adhesion and Wetting of Wrinkled Surfaces via a Lattice Model

Teng Zhang, Syracuse University, U.S.

3:05-3:25 Analysis of An Atomistic Model for Anti-plane Fracture with Sequential Coupling

Maciej Buze, University of Warwick, United Kingdom

3:30-3:50 Large-scale Embedding Quantum Simulations of Twisted Graphene Layers.

Xiaojie Wu, University of California, Berkeley, U.S.

continued in next column

Wednesday, February 27

MS214

Tensor-based Methods for Scientific and Engineering Computing

2:15 p.m.-3:55 p.m.

Room: 300C

Tensor-based methods are power tools in modern scientific and engineering computing, especially for high-dimensional and large data problems. Once a dataset has been tensorized, tensor decomposition methods such as the tensor-train (TT) or high-order singular value decomposition (HOSVD) can be used to form extremely compact representations, much smaller than the original data size. By limiting the rank in a tensor decomposition, it is also possible to find an accurate approximation of the original data, which again reduces the size of the representation. In many practical problems, extracting essential information from a huge data set has the potential to help enhancing the new scientific discovery or improving the physical model. Similar to the singular value decomposition (SVD), principal component analysis (PCA) or proper orthogonal decomposition (POD), the tensor decomposition methods not only provide a concise representation of the data but also help to discovery critical properties of the physical systems. This minisymposium will bring in researchers working on numerical analysis, algorithm development and applications using tensor-based methods from different areas to leverage the enhancement of the tensor methods in scientific and engineering computing.

Organizer: Xiu Yang

Pacific Northwest National Laboratory, U.S.

Organizer: Zheng Zhang

University of California, Santa Barbara, U.S.

2:15-2:35 Low-rank Tensor Methods for Optimal Control of Fractional PDEs

Gennadij Heidel, Universität Trier, Germany; Venera Khoromskaia, Max Planck Institute for Mathematics in the Sciences, Germany; Boris Khoromskij, Max Planck Institut Leipzig, Germany; Volker Schultz, Universität Trier, Germany

continued on next page

Wednesday, February 27

MS214

Tensor-based Methods for Scientific and Engineering Computing

continued

2:40-3:00 Optimal Solutions to Complex Tensor Approximation Problems almost Always Exist

Lek-Heng Lam and Yang Qi, University of Chicago, U.S.; Mateusz Michalek, Max Planck Institute for Mathematics in the Sciences, Germany and Polish Academy of Sciences, Poland

3:05-3:25 The Space-time Problem with Model Reduction for Traveling Waves

Ariana Mendible, Steven Brunton, and J. Nathan Kutz, University of Washington, U.S.

3:30-3:50 Fast and Flexible Bayesian Inference for Low-rank Tensor Problems

Cole Hawkins and Zheng Zhang, University of California, Santa Barbara, U.S.

Wednesday, February 27

MS215

BE: Parallel Algorithm Design

2:15 p.m.-3:55 p.m.

Room: 301

Organizer: Mary Ann E. Leung
Sustainable Horizons Institute, U.S.

Wednesday, February 27

MS216

Mathematical Methods for Control and Optimization of Large-scale Energy Networks - Part II of II

2:15 p.m.-3:55 p.m.

Room: 302A

For Part 1 see MS183

Recent developments in the energy market require new mathematical models and suitable algorithms for the efficient usage of the existing networks. For instance, in the case of electrical power networks, the increase of renewable energy sources and the rise of the electric car lead to new challenges for the mathematical model and require optimization of the power network. If additionally storage units such as batteries are considered, this becomes an even more challenging problem. Moreover, the integration of renewable energy sources such as natural gas, photovoltaics, and wind energy into the power grids poses new challenges from both theoretical and practical sides. Starting from these real-world problems, the mathematical aspects of this minisymposium concentrate on methods for large-scale energy networks and, in particular, address optimization and stability analysis, model predictive control, model-order reduction, uncertainty quantification, and related topics on energy networks. The abstract setting allows for consideration of applications arising from both, gas and power networks.

Organizer: Manuel Baumann
Max Planck Institute for Dynamics of Complex Technical Systems, Germany

Organizer: Yue Qiu
Max Planck Institute for Dynamics of Complex Technical Systems, Germany

continued on next page

2:15-2:35 Fast Numerical Methods for Gas Network Modeling and Simulation
Yue Qiu and Sara Grundel, Max Planck Institute for Dynamics of Complex Technical Systems, Germany; Martin Stoll, Chemnitz University of Technology, Germany; Peter Benner, Max Planck Institute for Dynamics of Complex Technical Systems, Germany

2:40-3:00 Hyperbolic Stochastic Galerkin Formulations to Model Fluctuations in Supply Networks
Stephan Gerster and Michael Herty, RWTH Aachen University, Germany

3:05-3:25 Rigorous MINLP Optimization of Power-to-methane Processes
Jennifer Uebbing, Max Planck Institute for Dynamics of Complex Technical Systems, Germany; Sebastian Sager, Universität Magdeburg, Germany; Kai Sundmacher, Max Planck Institute for Dynamics of Complex Technical Systems, Germany

Wednesday, February 27

MS217

Entropy-stable Formulations for Numerical Solution of Conservation Laws - Part II of II

2:15 p.m.-3:55 p.m.

Room: 302B

For Part 1 see MS184

Entropy stability presents an elegant route to improve the robustness of numerical solutions of compressible flows. The demand for high-order-accurate numerical discretizations and considerations of complex physics such as those involving sharp gradients and multi-component flows has resulted in a renewed interest in entropy stable discretizations. Many issues, such as the interplay between accuracy and robustness, the potential role of energy-preserving schemes, extensions to multi-species flows, and interaction with sub-grid-scale models are currently being explored. This minisymposium will address the related mathematical and numerical issues as well as applications in fluid dynamics.

Organizer: Scott Murman
NASA Ames Research Center, U.S.

2:15-2:35 On the Convergence of Entropy Stable Schemes to Statistical Solutions of Hyperbolic Systems of Conservation Laws
Siddhartha Mishra, ETH Zürich, Switzerland

2:40-3:00 Convex Limiting for Nonlinear Hyperbolic Systems with Source Terms
Bojan Popov, Texas A&M University, U.S.

3:05-3:25 Fully Discrete Entropy Stable and Kinetic Energy Preserving Discontinuous Galerkin Methods
Gregor Gassner and *Andrew R. Winters*, University of Cologne, Germany; David Flad, Universität Stuttgart, Germany; Gero Schnuecke and Lucas Friedrich, University of Cologne, Germany

3:30-3:50 Entropy-stable, High-order Discretizations based on Continuous Summation-by-parts Operators
Jason E. Hicken, Rensselaer Polytechnic Institute, U.S.

Wednesday, February 27

MS218

Novel Computational Algorithms for Future Computing Platforms - Part II of II

2:15 p.m.-3:55 p.m.

Room: 303A

For Part 1 see MS181

In the early 2000s, due to constraints on economical heat dissipation, clock speeds of single-core CPUs could no longer be increased, which marked the adoption of multi-core CPUs, together with a paradigm shift to algorithms specifically designed for multi-core architectures. About 15 years into this current architectural cycle and on its way to exascale performance, the computing industry finds itself at the confluence of technical difficulties that cast doubt on its ability to sustain this architectural model beyond the exascale capability. These difficulties are driving the hardware industry to develop application-specific chips and to look beyond silicon-based chips (e.g., quantum computing, physical annealing, neuromorphics, etc.), with a continued emphasis on raw processing power and emerging concerns about energy efficiency. Hardware specialization will likely redefine the way computational algorithms are developed over the next two decades for a wide range of important applications: large-scale PDE-based problems (CFD, wave propagation, subsurface modeling, etc.), artificial intelligence, computational chemistry, and optimization problems, to name a few. The pressure to decrease time to solution or improve simulation fidelity, or both, for these applications will continue unabated. This minisymposium provides a forum for sharing innovative ideas on algorithm development for leveraging future computing platforms.

Organizer: Dimitar Trenev
ExxonMobil Research and Engineering

Organizer: Laurent White
ExxonMobil Research and Engineering

continued on next page

Wednesday, February 27

MS218

Novel Computational Algorithms for Future Computing Platforms - Part II of II

continued

2:15-2:35 Shannon & Shrödinger: Statistical Algorithms and Architectures
Aaron Lott, D-Wave Systems, Inc., Canada

2:40-3:00 Non-von Neumann Computing using Networks of Optical Parametric Oscillators
Peter McMahon, Stanford University, U.S.

3:05-3:25 Quantum Algorithms for Approximate Optimization
Stuart Hadfield, NASA, U.S.

3:30-3:50 Leveraging Future Computing Platforms: An Open Discussion
Dimitar Trenev, ExxonMobil Research and Engineering

Wednesday, February 27

MS219

Advances in Analyzing Floating-point Errors in Computational Science - Part II of II

2:15 p.m.-3:55 p.m.

Room: 303B

For Part 1 see MS186

Modern error analysis of floating-point computations faces a double challenge with the rise of large-scale, low-precision computations. On the one hand, larger and larger problems are being solved routinely; on the other hand, the use of half-precision arithmetic (fp16) is becoming increasingly attractive. These two trends raise questions about the validity and sharpness of traditional rounding error bounds. In this minisymposium we will discuss recent advances in the analysis of errors in floating-point computations, and in particular some new approaches that have emerged to meet the challenges described above.

Organizer: Pierre Blanchard
University of Manchester, United Kingdom

Organizer: Nicholas J. Higham
University of Manchester, United Kingdom

Organizer: Theo Mary
University of Manchester, United Kingdom

2:15-2:35 Algorithm Based Error Analysis for Mixed Precision Matrix Factorizations.
Pierre Blanchard, University of Manchester, United Kingdom

2:40-3:00 Approximate Arithmetic - A Hardware Perspective
George Constantinides, Imperial College London, United Kingdom

3:05-3:25 Automated Backward Error Analysis
Zhoulai Fu, IT University of Copenhagen, Denmark

3:30-3:50 Bayesian Analysis of the Effects of Lower Precision Arithmetic in Inverse Problems
Daniela Calvetti, Case Western Reserve University, U.S.

Wednesday, February 27

MS220

Computational Methods for Linear Kinetic Transport Equations - Part II of II

2:15 p.m.-3:55 p.m.

Room: 201B

For Part 1 see MS187

Statistical mechanics provides a mathematical modeling framework in which a collection of "particles" (e.g., electrons, nucleons, atoms, or molecules) is represented by a probability density function (PDF). This representation allows one to model the influence of microscopic dynamics on larger scales, without explicitly computing the detailed dynamics on the microscopic scale. Models in this framework are referred to as "kinetic models", and arise in many application areas, including in rarefied gas dynamics, nuclear reactor modeling, and plasma physics. Transport phenomena refers to the evolution and redistribution of macroscopic quantities such as mass, momentum, energy, and heat flux. This minisymposium addresses recent advances in computational methods for the simulation of statistical mechanical models and the resulting transport phenomena. Our goal is to showcase a diverse array of application areas, models, and numerical techniques.

Organizer: Andrew J. Christlieb
Michigan State University, U.S.

2:15-2:35 A Third Order Accurate Wave Propagation Algorithm for Hyperbolic Partial Differential Equations
Christiane Helzel, Heinrich-Heine Universitaet Duesseldorf, Germany

2:40-3:00 A Second-order Asymptotic-preserving and Positivity-preserving Exponential Runge-Kutta Method for a Class of Stiff Kinetic Equations
Jingwei Hu, Purdue University, U.S.; Ruiwen Shu, University of Wisconsin-Madison, U.S.

continued on next page

3:05-3:25 A Positive Asymptotic Preserving Scheme for Linear Kinetic Transport Equations

Paul Laiu, Oak Ridge National Laboratory, U.S.; Martin Frank, Karlsruhe Institute of Technology, Germany; Cory Hauck, Oak Ridge National Laboratory, U.S.

3:30-3:50 Efficient and Highly Accurate Semi-Lagrangian Discontinuous Galerkin Method for Convection-diffusion Problems

Jingmei Qiu, University of Delaware, U.S.

Wednesday, February 27

MS221

Reduced Order Modeling for Parametric CFD Problems - Part II of II

2:15 p.m.-3:55 p.m.

Room: 201C

For Part 1 see MS188

Large-scale computing is recurrent in several contexts such as computational mechanics or fluid dynamics, due to the high computational complexity in solving parametric and/or stochastic systems, based, for instance, on partial differential equations. This often leads to an unaffordable computational burden, especially when dealing with real-world applications, real-time or multi-query computing. In order to lessen this computational burden, reduced-order modeling techniques (e.g., reduced basis, proper orthogonal and generalized decomposition, empirical interpolation, hierarchical model reduction) play a crucial role as they aim to capture the most important features of the problem at hand without giving up accuracy. This minisymposium focuses on the development and application of reduced-order modeling techniques in computational fluid dynamics for direct and inverse modeling, and for control, optimization and design purposes. The actual goal is twofold, on one hand to provide the state of the art on the most recently developed techniques; on the other hand to identify the new research directions and perspectives on the scientific panorama.

Organizer: Annalisa Quaini
University of Houston, U.S.

Organizer: Gianluigi Rozza
SISSA, International School for Advanced Studies, Trieste, Italy

Organizer: Simona Perotto
Politecnico di Milano, Italy

2:15-2:35 Goal-oriented Model Reduction of Parametrized Nonlinear PDEs in Aerodynamics

Masayuki Yano, University of Toronto, Canada

2:40-3:00 Learning a Hierarchy of Context-aware Low-fidelity Models for Adaptive Multifidelity Monte Carlo Estimation

Ionut-Gabriel Farcas, Technische Universität München, Germany; Benjamin Peherstorfer, Courant Institute of Mathematical Sciences, New York University, U.S.

3:05-3:25 Sampling and Clustering on the Pod-Grassmann Manifold

Michel Bergmann, Inria Bordeaux Sud-Ouest, France; Angelo Iollo, Institut de Mathématiques de Bordeaux, France

3:30-3:50 PGD-based Computational Vademecums for Parameterized Flows

Matteo Giacomini, Universitat Politecnica de Catalunya, Spain; Ruben Sevilla, Swansea University, United Kingdom; *Antonio Huerta*, Universitat Politecnica de Catalunya, Spain

continued in next column

Wednesday, February 27

MS222

High-order PDE Methods on CPU Architectures with Wide SIMD Units - Part II of II

2:15 p.m.-3:55 p.m.

Room: 202A

For Part 1 see MS185

In Summer 2016 the 2nd generation Intel Xeon Phi Processor was released and since 2017 the Intel Scalable processor is available. Both are the computational heart of many current and future supercomputing installations. Examples are the DoE-machines Trinity-II, Cori-II and Theta, the NSF-machine Stampede-2, Oakforest-PACS in Japan, or Europe's SuperMUC-NG. These processors combine the computational power of accelerator-based machines with the traditional approach of homogeneous high performance computing. From an application viewpoint, many research articles in the last years proved that high-order methods are able to the computational power of modern supercomputers. However, it is a challenging engineering problem to enable complex software packages on manycore architectures. This minisymposium brings together researchers working on efficient implementations of scalable high-order solvers. Of particular interest are applications which are known to require exascale computing resources in future, such as computational fluid dynamics, electro-magnetics and seismic simulations, as well as weather forecasting. Due to the broad spectrum of applications, this MS aims at identifying common algorithmic patterns of efficient high-order methods. Additionally, the comparison of different numerical approaches (e.g. CG vs. DG) will allow to identify best practices when leveraging Intel Xeon Phi and Intel Xeon Scalable systems at scale.

Organizer: Alexander Heinecke
Intel Corporation, U.S.

Organizer: Alexander Breuer
University of California, San Diego, U.S.

Organizer: David Moxey
University of Exeter, United Kingdom

2:15-2:35 Fused Earthquake Simulations on Deep Learning Hardware

Alexander Breuer, University of California, San Diego, U.S.; Alexander Heinecke, Intel Corporation, U.S.; Yifeng Cui, San Diego Supercomputer Center, U.S.

2:40-3:00 Experiences with Code Generation for High-order ADER-DG Schemes in SeisSol

Carsten Uphoff and Michael Bader, Technische Universität München, Germany

3:05-3:25 Efficient Implementation of Communication-avoiding One-step Ader-Dg Schemes for Nonlinear Hyperbolic Systems of Balance Laws Using AVX 512 Instructions

Francesco Fambri, University of Trento, Italy

3:30-3:50 Optimizing Performance for Portable Generic Finite Element Interfaces

Jeremy Thompson, University of Colorado Boulder, U.S.

Wednesday, February 27

MS223

Advances in Quasi-Monte Carlo Methods - Part II of II

2:15 p.m.-3:30 p.m.

Room: 202B

For Part 1 see MS189

Quasi-Monte Carlo (qMC) methods are the algorithms of choice for evaluating high dimensional integrals. Recent advances in qMC include the development of open source software, applications to uncertainty quantification, and multivariate statistical inference. This minisymposium highlights some of these advances.

Organizer: Fred J. Hickernell
Illinois Institute of Technology, U.S.

Organizer: Mike Giles
University of Oxford, United Kingdom

2:15-2:35 Mc and Qmc Algorithms for the Calculation of Dempster-Shafer Belief Functions

Nima Salehy, Florida State University, U.S.

2:40-3:00 An Adaptive Quasi-Monte Carlo Method for Bayesian Inference with User-specified Error Tolerance

Kan Zhang, Illinois Institute of Technology, U.S.

3:05-3:25 Automatic Bayesian Cubature with Quasi-Monte Carlo Sampling

Jagadeeswaran R, Illinois Institute of Technology, U.S.

continued in next column

Wednesday, February 27

MS224

Fast and Accurate Integral Methods for Highly Oscillatory Phenomena - Part II of II

2:15 p.m.-3:55 p.m.

Room: 202C

For Part 1 see MS190

Integral methods are useful tools in applied science and engineering. In particular, they are an important topic for large-scale scientific computing. Many challenges remain open and attract much attention especially in the high-frequency regime. This minisymposium focuses on recent advances in integral equations and integral transforms for highly oscillatory phenomena, including new formulations for high-frequency wave propagation, efficient and accurate discretizations, novel fast algorithms and their implementation based on locally rank-structured matrices and non-oscillatory phase functions, with applications in various imaging science and computational electromagnetism.

Organizer: Haizhao Yang
National University of Singapore, Singapore

Organizer: Yang Liu
Lawrence Berkeley National Laboratory, U.S.

2:15-2:35 Fast Algebras on Optimized Butterfly Structures
Yang Liu, Lawrence Berkeley National Laboratory, U.S.

2:40-3:00 Recent Advances in Butterfly Algorithms
Michael O'Neil, New York University, U.S.

3:05-3:25 Applying Integral Transforms using Nonoscilatory Phase Functions
James Bremer, University of California, Davis, U.S.

3:30-3:50 A Unified Framework for Oscillatory Integral Transforms: When to use Nufft Or Butterfly Algorithms?
Haizhao Yang, National University of Singapore, Singapore

Wednesday, February 27

MS225

New Challenges and Opportunities for Model Order Reduction - Part II of II

2:15 p.m.-3:55 p.m.

Room: 203

For Part 1 see MS191

In order to facilitate a fast simulation response, model order reduction techniques construct a low-dimensional, problem-adapted subspace from solutions at various time instances, different parameter values, or different boundary conditions. Examples comprise the Reduced Basis method, the Proper Orthogonal Decomposition, tensor-based approaches, or localized model order reduction approaches. In this minisymposium we want to discuss challenges and opportunities for model order reduction methods that arise for instance when applying model order reduction to data assimilation, parameter estimation and large input-parameter spaces, incorporating ideas originally developed in data science, or large-scale optimization.

Organizer: Laura Iapichino
Technische Universiteit Eindhoven, The Netherlands

Organizer: Kathrin Smetana
University of Twente, Netherlands

Organizer: Tommaso Taddei
INRIA Bordeaux, France

2:15-2:35 Highly Effective A-Posteriori Error Estimation for Approximation of Nonlinear Problems
Bernard Haasdonk, Andreas Schmidt, and Dominik Wittwar, Universität Stuttgart, Germany

2:40-3:00 Some New Results on PBDW, Application to Time-dependent Problems
Yvon Maday, Université Pierre et Marie Curie, France and Brown University, U.S.

3:05-3:25 Model Reduction of Systems of Hyperbolic Conservation Laws
Kyle T. Mandli and Donsub Rim, Columbia University, U.S.

3:30-3:50 Dictionary-based Reduced Basis Method via Random Sketching
Oleg Balabanov, Ecole Centrale de Nantes, France and Polytechnic University of Catalonia, Spain; Anthony Nouy, Ecole Centrale de Nantes, France

Wednesday, February 27

MS226

Computational Strategies for High-dimensional Data Assimilation and Bayesian Inverse Problems - Part II of II

2:15 p.m.-3:55 p.m.

Room: 205

For Part 1 see MS192

Bayesian inverse problems, in which one estimates parameters of a model based on noisy data, and data assimilation problems, in which one sequentially estimates the state and parameters of a dynamical model using noisy observations, occur throughout science and engineering. Practical problems are often characterized by nonlinear numerical models, non-Gaussian distributions, model error, and extremely high dimensionality. In this session we present and discuss computational methods for solving inference problems in this setting, with an emphasis on Markov chain Monte Carlo, importance sampling, and ensemble methods for high-dimensional posterior distributions.

Organizer: Matthias Morzfeld
University of Arizona, U.S.

Organizer: Xin T. Tong
National University of Singapore, Singapore

Organizer: Youssef M. Marzouk
Massachusetts Institute of Technology, U.S.

2:15-2:35 New Strategies for Inference with Multilevel Monte Carlo
Kody Law, University of Manchester, United Kingdom

2:40-3:00 Localization for MCMC - Sampling High Dimensional Posterior Distributions with Local Structure
Matthias Morzfeld, University of Arizona, U.S.; Xin T. Tong, National University of Singapore, Singapore; Youssef M. Marzouk, Massachusetts Institute of Technology, U.S.

continued on next page

Wednesday, February 27

MS226

Computational Strategies for High-dimensional Data Assimilation and Bayesian Inverse Problems - Part II of II

continued

3:05-3:25 Data Assimilation for Stochastic Advection by Lie Transport Euler SPDE

Colin J. Cotter, Dan Crisan, Darryl D. Holm, *Wei Pan*, and Igor Shevchenko, Imperial College London, United Kingdom

3:30-3:50 Nonlinear Filtering with Local Couplings

Youssef M. Marzouk, Alessio Spantini, and Ricardo Baptista, Massachusetts Institute of Technology, U.S.

Wednesday, February 27

MS227

Data-driven Methods in Fluid Dynamics - Part II of II

2:15 p.m.-3:55 p.m.

Room: 206A

For Part 1 see MS193

Data-fueled modeling and control of complex systems is currently undergoing a revolution, driven by the confluence of big data, advanced algorithms in machine learning, and modern computational power. Over the recent years, there has been an increasing focus on data-driven methods in fluid dynamics, e.g. for extraction of coherent flow structures, reduced-order modeling, control design in complex cost landscapes, sensor and actuator placement, and full-scale design optimization. This symposium explores the challenges and state-of-the-art innovations in data-driven methods for characterization, modeling, and control in challenging fluid flow problems, that include techniques such as dimensionality reduction, machine learning, operator-theoretic approaches, and control. In this minisymposium, experts will demonstrate achievements in the field, discuss open problems and challenges and inspire future directions.

Organizer: Eurika Kaiser
University of Washington, U.S.

Organizer: Kazuki Maeda
University of Washington, U.S.

2:15-2:35 Feedback Control of Nonlinear PDEs using Data-efficient Reduced Order Models Based on the Koopman Operator

Sebastian Peitz, University of Paderborn, Germany; Stefan Klus, Freie Universität Berlin, Germany

2:40-3:00 Data-driven Feedback Control Strategies for Unsteady Flows

Aditya G. Nair and Chi-An Yeh, Florida State University, U.S.; Eurika Kaiser, University of Washington, U.S.; Bernd Noack, Technische Universität Braunschweig, Germany; Steven Brunton, University of Washington, U.S.; Kunihiko Taira, University of California, Los Angeles, U.S.

3:05-3:25 Identifying and Modeling Coherent Structures in Turbulent Flows using Spectral Proper Orthogonal Decomposition and Resolvent Analysis

Aaron Towne, University of Michigan, U.S.

3:30-3:50 Modal Decomposition for Fluid-structure Interaction, and its Application to Flag Flapping

Andres Goza, University of Illinois at Urbana-Champaign, U.S.; Tim Colonius, California Institute of Technology, U.S.

continued in next column

Wednesday, February 27

MS228

Efficient Solvers for Coupled Multiphysics Problems - Part II of II

2:15 p.m.-3:55 p.m.

Room: 206B

For Part 1 see MS194

Modeling the simultaneous occurrence of different physical processes is an essential requirement for reliable simulation in many computational fluid and solid mechanics applications. The interaction between tightly coupled thermal, mechanical, chemical and flow processes can play a critical role. The discretization of these problems gives rise to linearized equations with an inherent block structure reflecting the underlying physical couplings. The development of efficient solution strategies, either monolithic or sequential, for such problems is an active field of research. This minisymposium aims to join scientists working on multi-physics problems, both from a purely algebraic and a more physics-based viewpoint, with the objective of sharing experiences and new ideas arising from a wide spectrum of applications.

Organizer: Massimiliano Ferronato

University of Padova, Italy

Organizer: Nicola Castelletto

Lawrence Livermore National Laboratory, U.S.

Organizer: Joshua A. White

Lawrence Livermore National Laboratory, U.S.

2:15-2:35 A Coupled Multi-physics Preconditioning Technique Based on Block Sparse Approximate Inverses

Andrea Franceschini, Stanford University, U.S.; Nicola Castelletto, Lawrence Livermore National Laboratory, U.S.; Carlo Janna and Massimiliano Ferronato, University of Padova, Italy; Hamdi Tchelepi, Stanford University, U.S.

2:40-3:00 A Parallel Multigrid Reduction Framework for Coupled Multiphase Poromechanics

Quan M. Bui and Daniel Osei-Kuffuor, Lawrence Livermore National Laboratory, U.S.

3:05-3:25 Block Preconditioners for Incompressible Magnetohydrodynamics Problems

Michael P. Wathen, *Chen Greif*, and Dominik Schoetzau, University of British Columbia, Canada

3:30-3:50 Multiphysics Simulations with PFLOTRAN and PETSc on Novel Manycore Computer Architectures

Richard T. Mills, Argonne National Laboratory, U.S.

Wednesday, February 27

MS229

CSE Education and Workforce - Part II of II

2:15 p.m.-3:55 p.m.

Room: 206C

For Part 1 see MS195
Featured Minisymposium

This minisymposium will highlight recent advances and future challenges in CSE education and workforce. Academic speakers will discuss new and thriving interdisciplinary programs in undergraduate and graduate CSE. Industry and National Lab speakers will discuss perspectives on CSE workforce needs from various industry and government sectors. Each talk slot will include time for an active discussion with minisymposium attendees.

Organizer: Serkan Gugercin

Virginia Tech, U.S.

Organizer: Karen E. Willcox

University of Texas at Austin, U.S.

2:15-2:35 Past, Present and Future of CSE Education at ICES

Karen E. Willcox, University of Texas at Austin, U.S.

2:40-3:00 Industrial Strength Computing in a Python Universe

Thomas A. Grandine, The Boeing Company, U.S.

3:05-3:25 Perspectives from a DOE Lab on CSE Education and Workforce Needs

Tamara G. Kolda, Sandia National Laboratories, U.S.

3:30-3:50 High Performance Computing at BP

Keith Gray, BP Exploration, United Kingdom

continued in next column

Wednesday, February 27

MS230

Recent Developments of Numerical Methods for Hyperbolic and Parabolic Equations - Part II of II

2:15 p.m.-3:55 p.m.

Room: 206D

For Part 1 see MS196

This minisymposium is to bring people together to discuss the recent advances and exchange ideas in the algorithm design of high-order numerical methods for hyperbolic and parabolic equations and other high-order partial different equations, including the implementation, numerical analysis. In the minisymposium, the speakers will apply those high-order numerical methods to computational fluid, biology and physics, etc. This minisymposium is a good opportunity for people to discuss with researchers from different areas, and explore more applications and future research collaborations. We expect 8 speakers to present in this minisymposium.

Organizer: Yang Yang

Michigan Technological University, U.S.

2:15-2:35 Adaptive Radial Basis Function Methods for PDEs

Jiaxi Gu and Jae-Hun Jung, State University of New York at Buffalo, U.S.

2:40-3:00 An Implicit Sparse Grid Discontinuous Galerkin Method for High Dimensional Reaction-diffusion Equations

Yuan Liu, Mississippi State University, U.S.; Yingda Cheng, Michigan State University, U.S.; Shanqin Chen, Indiana University, U.S.; Yong-Tao Zhang, University of Notre Dame, U.S.

3:05-3:25 Accessing Superconvergence for Accurate Multi-resolution Analysis of Discontinuous Galerkin Approximations

Jennifer K. Ryan, University of East Anglia, United Kingdom and Heinrich-Heine University, Germany

3:30-3:50 Symmetric Direct Discontinuous Galerkin Method for Elliptic Interface Problems

Jue Yan, Iowa State University, U.S.

Wednesday, February 27

MS231

Task-based Programming for Scientific Computing: Linear Algebra Applications - Part II of II

2:15 p.m.-3:55 p.m.

Room: 207

For Part 1 see MS197

The evolution of modern architectures has led to the development of numerous new programming models. These models aims at simplifying the work of the programmer to help with performance portability among various architectures, but also in enabling more asynchronism in applications. Among all solutions task-based runtimes systems and language extensions provide a large spectrum of features to the programmer to adapt to various applications. In this minisymposium, we propose to study a set of modern linear algebra solvers that exploits these expressive models to reach high performance, and or express new algorithms.

Organizer: Mathieu Faverge

Bordeaux INP, Inria, LaBRI, France

Organizer: Mathias Jacquelin

Lawrence Berkeley National Laboratory, U.S.

2:15-2:35 Slate: Multilevel Tasking in Dense Linear Algebra Libraries

Asim YarKhan, University of Tennessee, Knoxville, U.S.; Mark Gates, University of Tennessee, U.S.; Piotr Luszczek and Jakub Kurzak, University of Tennessee, Knoxville, U.S.; Jack J. Dongarra, University of Tennessee and Oak Ridge National Laboratory, U.S.

2:40-3:00 Task-based Sparse Direct Solver for Symmetric Indefinite Systems

Florent Lopez, Rutherford Appleton Laboratory, United Kingdom; Iain Duff, Science & Technology Facilities Council, United Kingdom and CERFACS, Toulouse, France

3:05-3:25 Preconditioning using Rank-structured Sparse Matrix Factorization

Pieter Ghysels, Lawrence Berkeley National Laboratory, U.S.

3:30-3:50 A 2D Task-based Fan-both Factorization Algorithm for Sparse Symmetric Matrices

Mathias Jacquelin and Esmond G. Ng, Lawrence Berkeley National Laboratory, U.S.

continued in next column

Wednesday, February 27

MS232

Surrogate Modeling and Data Compression for Exascale Applications - Part II of II

2:15 p.m.-3:55 p.m.

Room: 401A

For Part 1 see MS198

High performance computing systems are expected to reach the exascale in the near future, requiring a new generation of simulation tools and introducing new challenges arising from high-dimensional, extremely large data sets and the highly-distributed nature of exascale systems. This minisymposium will emphasize surrogate modeling and data compression techniques applicable to exascale applications and requiring scaling across many thousands of cores or to large datasets. Talks that address topics in data compression, dimension reduction, in situ visualization, surrogate modeling, or contain applications targeting exascale computing resources are encouraged.

Organizer: Matthew J. Reynolds

National Renewable Energy Laboratory, U.S.

Organizer: Ryan King

National Renewable Energy Laboratory, U.S.

Organizer: Alireza Doostan

University of Colorado Boulder, U.S.

2:15-2:35 Data Reconstruction for Computational Fluid Dynamics using Deep Neural Networks

Marc Henry de Frahan, Ryan King, and Ray W. Grout, National Renewable Energy Laboratory, U.S.

2:40-3:00 Dimensional Analysis and Reduced Order Models for Large Scale Simulations of Particle-laden Turbulence in a Radiation Environment

Zachary del Rosario, Lluis Jofre, and Gianluca Iaccarino, Stanford University, U.S.

3:05-3:25 Pass-efficient Matrix Algorithms for Lossy Data Compression

Alec M. Dunton, ; Lluis Jofre and Gianluca Iaccarino, Stanford University, U.S.; Alireza Doostan, University of Colorado Boulder, U.S.

3:30-3:50 A Low-rank, Bi-fidelity Approximation for Linear Bayesian Inference

Hillary Fairbanks, Lawrence Livermore National Laboratory, U.S.; Alireza Doostan, University of Colorado Boulder, U.S.

continued in next column

Wednesday, February 27

MS233

Recent Developments in Numerical Methods for PDEs and their Applications - Part II of II

2:15 p.m.-3:55 p.m.

Room: 401B

For Part 1 see MS79

Numerical Methods for partial differential equations and their analysis are important and challenging topics in applied and computational mathematics. This minisymposium is focused on recent developments in numerical methods for PDEs, including new developments in finite element methods and relevant applications. The goal of this minisymposium is to bring together leading researchers in the field of numerical methods to discuss and disseminate the latest results and envisage future challenges in traditional and new areas of science. The topics of the minisymposium cover a broad range of numerical methods, including but not limited to finite element methods, finite difference methods, discontinuous Galerkin methods, weak Galerkin methods. A wide range of application fields will also be covered, such as MHD equations, Helmholtz equations, equations of poroelasticity, Alfeld splits and gradient flows.

Organizer: Chunmei Wang

Texas Tech University, U.S.

Organizer: Jun Zou

The Chinese University of Hong Kong, Hong Kong

2:15-2:35 On the Equations of Poroelasticity

Amnon J Meir, Southern Methodist University, U.S.

2:40-3:00 A Posteriori Error Estimators of Recovery Type

Zhiqiang Cai, Purdue University, U.S.

3:05-3:25 Adaptive FEM for Helmholtz Equation with High Wave Number

Haijun Wu, Nanjing University, China

3:30-3:50 Adaptive Finite Element Methods for Solving Inverse Problems of PDEs

Jun Zou, The Chinese University of Hong Kong, Hong Kong

Wednesday, February 27

MS234

Sparse Function Approximations: Theory and Applications - Part II of II

2:15 p.m.-3:55 p.m.

Room: 401C

For Part 1 see MS200

Learning non-linear systems from noisy, limited, and/or dependent data is an important task for data analysis across various many scientific fields including statistic, engineering, computer science, mathematics, and many more. One of the major paradigms is to learn an unknown generating function from a set of input-output pairs, which can be rephrased as finding an appropriate approximation to a high-dimensional function. Without additional restrictions or structures, this learning task is ill-posed. This minisymposium addresses some recent developments of several learning techniques, which leverage structures such as sparsity, randomness, dimensionality, etc, in order to make the learning well-posed. The presentations will cover analysis based on orthogonal polynomials, sampling, approximation theory, probability theory, as well as novel computational techniques.

Organizer: Giang Tran

University of Waterloo, Canada

Organizer: Rachel Ward

University of Texas at Austin, U.S.

Organizer: Hayden Schaeffer

Carnegie Mellon University, U.S.

2:15-2:35 Learning Large-scale Sparse Graphical Models: Theory, Algorithm, and Applications

Somayeh Sojoudi, Salar Fattahi, and Richard Zhang, University of California, Berkeley, U.S.

2:40-3:00 Ridge Approximations and Active Subspaces

Paul Constantine and Jeffrey M. Hokanson, University of Colorado Boulder, U.S.; Rachel Ward, University of Texas at Austin, U.S.; Armin Eftekhari, Alan Turing Institute, United Kingdom

3:05-3:25 Learning Interaction Laws in Agent Dynamics from Observations

Sui Tang, Fei Lu, and Mauro Maggioni, Johns Hopkins University, U.S.

3:30-3:50 Identifying and Elimination of Data Corruption

Dongbin Xiu and Jun Hou, Ohio State University, U.S.

continued in next column

Wednesday, February 27

MS235

Data Science and Analytics in Industry

2:15 p.m.-3:55 p.m.

Room: 402A

The role of data science and analytics in industry is opening up many new opportunities for applied mathematicians. This minisymposium will feature speakers from different industries who will share their jobs in data science and analytics, the challenges the companies have and how applied mathematicians can make a difference, and how they got to their current positions.

Organizer: Malena I. Espanol

University of Akron, U.S.

2:15-2:35 Data Analytics is Imperative for Business Success – Why are Companies Struggling?

Mario R. Garzia, Diga Group, U.S.

2:40-3:00 The Technical and Organizational Challenges of Data Science

Catherine Micek, 3M, U.S.

3:05-3:25 Applications of Data Science and Analytics in the Defense Industry and Challenges

Veronica Bloom, Northrop Grumman Corp., U.S.

3:30-3:50 Being Rational: Weighted Adaptive Approximation of Scientific Data

Youssef Nashed, Tom Peterka, Vijay Mahadevan, and Iulian Grindeanu, Argonne National Laboratory, U.S.

Wednesday, February 27

MS236

Nonlocal Models in Computational Science and Engineering - Part II of II

2:15 p.m.-3:55 p.m.

Room: 402B

For Part 1 see MS202

Nonlocal models provide a new framework to overcome limitations and challenges present in classical PDE-based models. For instance, peridynamics, a nonlocal extension of classical continuum mechanics, admits discontinuous solutions and naturally describes material failure and damage. Similarly, nonlocal and fractional diffusion models can represent anomalous diffusion and heat transfer. Furthermore, nonlocal models introduce length scales, which can be used for multiscale modeling. Recent years have witnessed a tremendous advance in modeling, mathematical analysis, and computational practice for nonlocal problems. This minisymposium focuses on recent developments in peridynamics, nonlocal and fractional diffusion, and other related nonlocal models. The topics of interest include and are not limited to: peridynamics, nonlocal and fractional mass diffusion and heat transfer, discretization and time-stepping methods for nonlocal models, nonlocal boundary conditions, multiscale methods and multi-physics modeling with nonlocal models, nonlocal modeling of heterogeneous systems, engineering and scientific applications with nonlocal models.

Organizer: Yue Yu
Lehigh University, U.S.

Organizer: Pablo Seleson
Oak Ridge National Laboratory, U.S.

2:15-2:35 Improving the Efficiency of FEM Discretizations for Nonlocal Problems: a New Concept of Nonlocal Neighborhoods

Marta D'Elia, Sandia National Laboratories, U.S.; Christian Vollmann, Universität Trier, Germany; Max Gunzburger, Florida State University, U.S.; Volker Schultz, Universität Trier, Germany

2:40-3:00 Finite Element Approximation of Fractional Minimal Surfaces

Juan Pablo Borthagaray, Ricardo Nochetto, and Wenbo Li, University of Maryland, U.S.

3:05-3:25 Finite Element Approximation of An Obstacle Problem for a Class of Integro-differential Operators

Abner J. Salgado, University of Tennessee, U.S.

3:30-3:50 A Fractional Calculus Framework for the Prediction of Material Failure

Mohsen Zayernouri, Michigan State University, U.S.

continued in next column

Wednesday, February 27

MS237

WCD Workshop - 4 of 4

2:15 p.m.-3:55 p.m.

Room: 402C

For Part 3 see MS203
Part of the SIAM Workshop Celebrating Diversity

Organizer: Talea Mayo
University of Central Florida, U.S.

Organizer: Shelby Wilson
Morehouse College, U.S.

See online program for an update on this session.

Wednesday, February 27

CP13

Numerical PDEs IV

2:15 p.m.-3:55 p.m.

Room: 201A

Chair: Joseph V. Koebbe, Utah State University, U.S.

2:15-2:30 An Overlapping Multiscale Mixed Method: Formulation and Iterative Downscaling

Het Y. Mankad, Abdullah Al Mamun, and Felipe Pereira, University of Texas at Dallas, U.S.; Fabricio Sousa, University of Sao Paulo, Brazil

2:35-2:50 Recovery-based Discontinuous Galerkin Method for the Cahn-Hilliard Equation

Fanchen He, Philip Johnson, and Eric Johnsen, University of Michigan, U.S.

2:55-3:10 Homogenization Wavelet Reconstruction for Partial Equations with Elliptic Terms

Joseph V. Koebbe, Utah State University, U.S.; Abibat Lasisi, Virginia Military Institute, U.S.

3:15-3:30 High-order Stochastic Semi-Lagrangian Method to Solve Transport Equations

Hareshram Natarajan and Gustaaf Jacobs, San Diego State University, U.S.

3:35-3:50 Fast Finite Difference Schemes for the Time-fractional Diffusion Equation with a Weak Singularity at the Initial Time

Jinye Shen, Zhizhong Sun, and Rui Du, Southeast University, China

Intermission

3:55 p.m.-4:10 p.m.

Poster Blitz

4:10 p.m.-4:50 p.m.

Room: Ballroom 100BC

Wednesday, February 27

PP2

General Posters

4:50 p.m.-6:50 p.m.

Room:Riverside Hall D

To Be Determined

Padè Time Stepping Method of Rational Form for Reaction Diffusion Equations

Said Algarni, King Fahd University of Petroleum and Minerals, Saudi Arabia

Modeling, Simulation, Testing and Correlation of a Solar Photovolatic / Battery Storage System

Kevin R. Anderson and Wael Yassine, California State Polytechnic University, Pomona, U.S.

Sparse Approximate Matrix Multiplication in a Fully Recursive Distributed Task-based Parallel Framework

Anton G. Artemov, Uppsala University, Sweden

A Double Core Tensor Factorization for Heterogeneous Data

Davoud Ataee Tarzanagh and George Michailidis, University of Florida, U.S.

Randomized Sub-sampled Methods for Matrix Approximation

Andrew T. Azzam, Benjamin W. Ong, and Allan Struthers, Michigan Technological University, U.S.

Index-aware MOR for Nonlinear DAEs Arising from Gas Transport Networks

Nicodemus Banagaaya, Sara Grundel, and Peter Benner, Max Planck Institute for Dynamics of Complex Technical Systems, Germany

A Stabilized, Hypoelastic Constitutive Correspondence Framework for Peridynamics

Masoud Behzadinasab and John Foster, University of Texas at Austin, U.S.

Anomalous Sub- and Super-diffusion in Image Processing

Toheeb A. Biala and Abdul Khaliq, Middle Tennessee State University, U.S.

Blocked Multigrid Methods for Structured Matrices

Matthias Bolten, University of Wuppertal, Germany

Exploring the Variability and Error Distributions of Lossy Compression Algorithms

Jon Calhoun, Tasmia Reza, and Pavlo Triantafyllides, Clemson University, U.S.

A Scalable Petsc Implementation of the 3-Phase Unbalanced Ac Power Flow Solver

Justin Chang, Gord Stephen, and Dheepak Krishnamurthy, National Renewable Energy Laboratory, U.S.

Sensitivity Analysis in Particle-in-Cell Methods

Seung Whan Chung, University of Illinois at Urbana-Champaign, U.S.; Stephen D. Bond and Eric C. Cyr, Sandia National Laboratories, U.S.; Jonathan B. Freund, University of Illinois at Urbana-Champaign, U.S.

Pipelined Krylov Subspace Methods with Improved Parallel Scalability

Jeffrey Cornelis, University of Antwerp, Belgium; Pieter Ghysels, Lawrence Berkeley National Laboratory, U.S.; Siegfried Cools and Wim Vanroose, University of Antwerp, Belgium

Exponential Affine Solutions for Correlated Stochastic Process

Allan J. Da Silva and Jack Baczynski, National Laboratory for Scientific Computing, Brazil; José Vicente, BCB/ IBMEC, Brazil

An EXCMG Accelerated Multiscale Multigrid Computation for 3D Poisson Equation

Ruxin Dai, University of Wisconsin-River Falls, U.S.

Accelerated Boundary Integral Treecodes

Satyen V. Dhamankar and Benjamin W. Ong, Michigan Technological University, U.S.

Registry Effects in Carbon Nanostructures

Malena I. Espanol, Dmitry Golovaty, and J. Patrick Wilber, University of Akron, U.S.

Towards Precision Medicine: Simulation Based Parameter Estimation for Drug Metabolism

Yoko Franchetti and Thomas Nolin, University of Pittsburgh, U.S.; Franz Franchetti, Carnegie Mellon University, U.S.

Vesicle Adhesion in Constricted Geometries

Ashley R. Gannon, Florida State University, U.S.

continued in next column

continued on next page

Quantitative Engineering Analysis Curriculum

John B. Geddes, Franklin W. Olin College of Engineering, U.S.

In-situ Visualization for a Viscoplastic Column Collapse

Linda Gesenhues, José Camata, and Fernando A. Rochinha, COPPE/Universidade Federal do Rio e Janeiro, Brazil; Alvaro Coutinho, Universidade Federal de Rio de Janeiro, Brazil

High-order Accurate Vlasov Discretizations in 2D+2V

Andre Gianesini Odu and Jeffrey W. Banks, Rensselaer Polytechnic Institute, U.S.

Elasticity in Space and Time

Silke Glas and Karsten Urban, University of Ulm, Germany

Numerical Simulation of Basal Entrainment of a Viscoplastic Fluid

Malú Grave, Linda Gesenhues, and José Camata, COPPE/Universidade Federal do Rio e Janeiro, Brazil; Alvaro Coutinho, Universidade Federal de Rio de Janeiro, Brazil

Analysis of a Coupled Reaction-diffusion Model for Tumour Induced Angiogenesis in Breast Cancer Tissue

Phebe Mawuena A. Havor, Kwame Nkrumah University of Science and Technology, Ghana

Dynamics of Disease Models with Self-diffusion: A Study of Cholera in Ghana

Phebe Mawuena A. Havor and Anas Musah, Kwame Nkrumah University of Science and Technology, Ghana

A High-resolution Godunov Method with HLLC Riemann Solver for Two-phase Modeling of Practical Condensed-phase Explosives

Michael P. Hennessey, Donald W. Schwendeman, and Ashwani K. Kapila, Rensselaer Polytechnic Institute, U.S.

High-order Runge-Kutta Discontinuous Galerkin Methods with Local Time-stepping for Conservation Laws

Thi-Thao-Phuong Hoang, Auburn University, U.S.

Fully Kinetic PIC-DSMC Simulations to Study Backflow of Ion Thruster Plasma Plumes using Hybrid MPI-CUDA Paradigm

Revathi Jambunathan and Deborah Levin, University of Illinois, Urbana-Champaign, U.S.

CFD and Heat Transport Study of Varying Pebble Diameters in Pebble Bed Nuclear Reactors

Andrew Jones, Boise State University, U.S.; Tien Yee and Eduardo Farfan, Kennesaw State University, U.S.

Validation of Assumptions in MILP Scheme for Determining Optimal Orificing in Fast Reactors using SAM

Chris Keckler, University of California, Berkeley, U.S.; Yves Robert, INSA Lyon, France; Massimiliano Fratoni and Ehud Greenspan, University of California, Berkeley, U.S.

Bayesian Inversion of Fault Properties in Two-phase Flow in Fractured Media

Eldar Khattatov, Omar Ghattas, Tan Bui-Thanh, and Ilona Ambartsumyan, University of Texas at Austin, U.S.; Umberto Villa, Washington University, St. Louis, U.S.

Flexible Parallel Mesh Partitioning Strategies for Large-scale Multiphysics Simulations

Fande Kong, Derek R. Gaston, John Peterson, Cody J. Permann, and Richard Martineau, Idaho National Laboratory, U.S.

Machine Learning Algorithm to Support Interpretability and Accuracy Simultaneously

Boris Kovalerchuk and Nathan Newhous, Central Washington University, U.S.

Algorithm Design at Scale: Porting Parallel FFT-based Fortran Simulations to GPUs

Anuva Kulkarni, Carnegie Mellon University, U.S.; Jelena Kovacevic, New York University, U.S.; Franz Franchetti, Carnegie Mellon University, U.S.

A Fourth-order Accurate Multigrid Solver for Overset Grids

Chang Liu and William Henshaw, Rensselaer Polytechnic Institute, U.S.

Experiences Porting a Large-scale Deterministic Neutron Transport Application to Heterogeneous Architectures

John Loffeld, Teresa S. Bailey, Peter Brown, Adam Kunen, and Bujar Tagani, Lawrence Livermore National Laboratory, U.S.

Meshfree Semi-Lagrangian Schemes for Advection on Surfaces: Polyharmonic Splines Augmented with Polynomials

Daniel Malmuth and Grady B. Wright, Boise State University, U.S.; Varun Shankar, University of Utah, U.S.

Cluster Extraction using Compressive Sensing

Daniel Mckenzie and Ming-Jun Lai, University of Georgia, U.S.

A Time Adaptive Multirate Dirichlet-Neumann Waveform Relaxation Method for Heterogeneous Coupled Heat Equations

Azahar Monge and Philipp Birken, Lund University, Sweden

Sparse Fourier Interpolation in Computational Chemistry

Zachary B. Morrow, ELENA Jakubikova, Tim Kelley, and Chang Liu, North Carolina State University, U.S.; Miroslav Stoyanov, Oak Ridge National Laboratory, U.S.

A Highly Scalable, Fault-tolerant Implementation of the Sparse Grid Combination Technique

Michael Obersteiner and Hans-Joachim Bungartz, Technische Universität München, Germany; Dirk Pflüger, Universität Stuttgart, Germany

An Efficient p-Multigrid Preconditioner for an Implicit Dg Compressible Flow Solver

Yu Pan, Zhenguo Yan, Peiro Joaquim, and Spencer Sherwin, Imperial College London, United Kingdom

Randomized Methods for Recompression of Low-rank Tensors

Lana Perisa and Daniel Kressner, École Polytechnique Fédérale de Lausanne, Switzerland

A Continuous Model of Discrete Scientific Data

Tom Peterka, Youssef Nashed, Iulian Grindeanu, and Vijay Mahadevan, Argonne National Laboratory, U.S.; Raine Yeh and Xavier M. Tricoche, Purdue University, U.S.

Using Push-forward Measures for Parameter Identification under Uncertainty

Michael Pilosov, University of Colorado, Denver, U.S.

Multi-level and Multi-index Monte Carlo Discontinuous Galerkin Methods for Uncertainty Quantification of Nonlinear Hyperbolic Problems

Stanislav Y. Polishchuk, Hans De Sterck, and Tiangang Cui, Monash University, Australia

Multigrid Preconditioning for Optimization-based Domain Decomposition of the Helmholtz Equation

Michael A. Retzlaff and Andrei Draganescu, University of Maryland, Baltimore County, U.S.

continued in next column

continued in next column

continued on next page

Wednesday, February 27

PP2

General Posters

continued

Reconstruction of Sparse Sums with a Generalized Music Algorithm

Jacob D. Rezac, National Institute of Standards and Technology, U.S.

Different Measure Approximations for Efficient Constrained Multi-objective Optimization under Uncertainty

Mickaël Rivier and Pietro M. Congedo, Inria Bordeaux Sud-Ouest, France

Modeling of Low Salinity Water Flooding in Naturally Fractured Core Samples

Carlos A. Romano, Universidad Nacional Autónoma de México, Mexico; Martin Diaz, Instituto Mexicano del Petróleo, México

Communication Lower Bounds for Computing a Matricized-tensor Times Khatri-Rao Product

Kathryn Rouse, Inmar; Grey Ballard, Wake Forest University, U.S.; Nicholas Knight, University of California, Merced, U.S.

A Kernel-based Approach for Solving the Hamilton-Jacobi Formulation of the Vlasov-Poisson System

William A. Sands, Andrew J. Christlieb, and Yan Jiang, Michigan State University, U.S.

Mathematical Modeling of an Application Specific Processor Architecture with Power Optimization

Vijayalakshmi Saravanan and Anpalagan A, Ryerson University, Canada

Sparse Grid Density Estimation Techniques for Clustering-based Collocation

Paul Cristian Sarbu and Hans-Joachim Bungartz, Technische Universität München, Germany

ETD Algorithm for Fractional Model of Pressure Distribution in Fractured Rocks

Ibrahim O. Sarumi and Khaled Furati, King Fahd University of Petroleum and Minerals, Saudi Arabia; Abdul Khaliq, Middle Tennessee State University, U.S.

Deep Learning - Neural Networks

Alexandru Savoiu, Stanford University, U.S.

MPI-OpenMP Load Balanced Simulation of Inhomogeneous Particle Systems in ls1 Mardyn at Extreme Scale

Steffen Seckler and Nikola Tchipev, Technische Universität München, Germany; Matthias Heinen, University of Paderborn, Germany; Fabio A. Gratl and Hans-Joachim Bungartz, Technische Universität München, Germany; Philipp Neumann, University of Hamburg, Germany

A Comparison of RBF-FD Methods for Solving Partial Differential Equations on Surfaces

Sage B. Shaw and Grady B. Wright, Boise State University, U.S.; Varun Shankar, University of Utah, U.S.

A Weighted Essentially Non-oscillatory Forward Semi-Lagrangian Scheme for Vlasov-Poisson Systems

David Sirajuddin and William Hitchon, University of Wisconsin, Madison, U.S.

Comparative Study of Exponential Time Integrators for Chemical Kinetic Problems

Jared Stewart, University of California, Merced, U.S.

Anisotropy in Two-dimensional and Plane Elasticity Bond-based Linear Peridynamics

Jeremy Trageser and Pablo Seleson, Oak Ridge National Laboratory, U.S.

NFFT-based Fast Summation for the Graph Laplacian of Fully Connected Networks

Toni Volkmer, Dominik Alfke, Daniel Potts, and Martin Stoll, Chemnitz University of Technology, Germany

Primal-dual Weak Galerkin Finite Element Methods for PDEs

Chunmei Wang, Texas Tech University, U.S.

Sharp Uniform in Time Error Estimate on a Stochastic Structure-preserving Lagrangian Method and Computation of Effective Diffusivity in 3D Chaotic Flows

Zhongjian Wang, University of Hong Kong, Hong Kong; Jack Xin, University of California, Irvine, U.S.; Zhiwen Zhang, The University of Hong Kong, Hong Kong

User Interfaces and Practical Examples of Opendiel

Kwai L. Wong, University of Tennessee and Oak Ridge National Laboratory, U.S.; Chung Ng and Efosa Asemota, Morehouse College, U.S.; Frank Betancourt, University of Tennessee, U.S.; Quindell Marshall, Morehouse College, U.S.; Zac Trzil, University of Tennessee, U.S.

Development of Efficient Preconditioners for Newton-Krylov Method in Spectral/hp Element, DG Compressible Flow Simulations

Zhen-Guo Yan, Yu Pan, Joaquim Peiro, and Spencer Sherwin, Imperial College London, United Kingdom

Spectral Velocity Discretization of the Vlasov Equation using Generalized Hermite Functions

Anna Yurova and Katharina Kormann, Max Planck Institute for Plasma Physics, Germany; Caroline Lasser, Technische Universität München, Germany

Forward Stability of ResNet and its Variants

Linan Zhang and Hayden Schaeffer, Carnegie Mellon University, U.S.

Exploration of Numerical Precision in Deep Neural Networks

Yunkai Zhang, University of California, Santa Barbara, U.S.; Yu Ma, University of California, Berkeley, U.S.; Zhaoqi Li, University of Washington, U.S.; Catalina Vajiac, University of Notre Dame, U.S.

A Novel Method to Average Images by Averaging Diffeomorphisms

Zicong Zhou and Guojun Liao, University of Texas at Arlington, U.S.

High Degree Discontinuous Petrov-Galerkin Immersed Finite Element Methods using Fictitious Elements for Elliptic Interface Problems

Qiao Zhuang and Ruchi Guo, Virginia Tech, U.S.

Monolithic Simulation of Binary Alloy Phase-change Coupled to Thermal and Compositional Convection

Alexander G. Zimmerman and Julia Kowalski, RWTH Aachen University, Germany

Forecasting Volcanic Eruption Time using a Stochastic Enhancement of the Failure Forecast Method

Abani K. Patra, State University of New York, Buffalo, U.S.

continued in next column

continued in next column

continued on next page

Understanding the Structure of Germanium Sulfide (GeS)

Miguel A. Castro, Inter-American University of Puerto Rico, Bayamón Campus, P.R.; Betul Pamuk and Darrell Schlom, Cornell University, U.S.

Density-equalizing Reference Map with Applications

Gary Choi and Chris H. Rycroft, Harvard University, U.S.

On the less than Reasonable Effectiveness of Floating Point Arithmetic in the Modelling of Physical Systems: the Case of the Generalised Bernoulli Map

Peter Coveney, University College London, United Kingdom

Coupling an Explicit Variable Density Projection Method to Finite Rate Kinetics

Josh T. Mcconnell and James C. Sutherland, University of Utah, U.S.

Wednesday, February 27

PP201

Minisymposterium: AMReX: Software and Applications

4:50 p.m.-6:50 p.m.

Room: Riverside Hall D

Andrew Meyers, Lawrence Berkeley National Laboratory, U.S.

AMReX is an ECP-funded, open source software framework for developing block-structured adaptive mesh refinement (AMR) applications. AMReX supports a variety of multiphysics application codes spanning a wide range of scientific problems, including combustion research, astrophysics and cosmology, accelerator physics, and carbon capture and sequestration. In this minisymposterium session, we highlight a few of the recent scientific and software developments featuring AMReX and its application codes, with a particular focus on 1) new application domains and 2) extending AMReX to work on hydrid GPU/CPU systems.

Minisymposterium: SedonaEx: A Monte Carlo Radiation Transfer Code for Astrophysical Events

Donald E. Willcox and Ann S. Almgren, Lawrence Berkeley National Laboratory, U.S.; Daniel Kasen, University of California, Berkeley, U.S.; Andrew Myers and Weiqun Zhang, Lawrence Berkeley National Laboratory, U.S.

Minisymposterium: Building a Microphysics Cloud Model with AMReX and PADDI

Sara Nasab, University of California, Santa Cruz, U.S.

Minisymposterium: Overview of Amrex - a New Framework for Block-structured Adaptive Mesh Refinement Calculations

Andrew Myers, Ann S. Almgren, John B. Bell, Marcus Day, Brian Friesen, Kevin N. Gott, Andy J. Nonaka, Steven Reeves, and Weiqun Zhang, Lawrence Berkeley National Laboratory, U.S.

Minisymposterium: Application of Implicit Methods for Complex Fluid Flow

Guy C. Moore and Andrew Nonaka, Lawrence Berkeley National Laboratory, U.S.

Minisymposterium: Performance Study of GPU Offloading via CUDA, OpenACC, and OpenMP in AMReX

Yuexia Lin, Harvard University, U.S.; Ann S. Almgren, Brian Friesen, and Andrew Myers, Lawrence Berkeley National Laboratory, U.S.

Minisymposterium: Modeling the Chemistry and Hydrodynamics of Micro-swimmers

Matea Alvarado, University of California, Merced, U.S.; Johannes P. Blaschke, Lawrence Berkeley National Laboratory, U.S.

Minisymposterium: Modeling Electro-kinetic Flows with Fluctuating Hydrodynamics

Sean P. Carney, The University of Texas at Austin, U.S.; John B. Bell and Andy J. Nonaka, Lawrence Berkeley National Laboratory, U.S.; Alejandro Garcia, San Jose State University, U.S.

Minisymposterium: An Overview of GPU Strategies for Porting Amrex-Based Applications to Next-generation HPC Systems

Kevin N. Gott, Lawrence Berkeley National Laboratory, U.S.

continued in next column

Wednesday, February 27

PP202

Minisymposterium: Clawpack and ForestClaw Software

4:50 p.m.-6:50 p.m.

Room: Riverside Hall D

Randall LeVeque, University of Washington, U.S.

Clawpack is an open source software package for solving time-dependent hyperbolic partial differential equations with adaptive mesh refinement capabilities (AMR). The posters in this session present some new developments to accelerate AMR using GPUs, to the ForestClaw package using p4est for the AMR, and toward the development of better absorbing boundary conditions for wave propagation problems.

Minisymposterium: Accelerating Wave-propagation Algorithms with Adaptive Mesh Refinement Using the Graphics Processing Unit (GPU)

Xinsheng Qin, Randall LeVeque, and Michael Motley, University of Washington, U.S.

Minisymposterium: A Massively Parallel Solver for Poisson's Equation on Block Structured Cartesian Grids

Scott Aiton, Donna Calhoun, and Grady B. Wright, Boise State University, U.S.

Minisymposterium: Evaluation of Shallow Water Models for Tsunami Prediction with a Near-field Seismic Source

Christopher J. Vogl, Lawrence Livermore National Laboratory, U.S.; Randall LeVeque, University of Washington, U.S.

Minisymposterium: Absorbing Outgoing Waves with Coordinate Mappings in Clawpack

Randall LeVeque, University of Washington, U.S.; Christopher J. Vogl, Lawrence Livermore National Laboratory, U.S.

Wednesday, February 27

PP203

Minisymposterium: Multi-level and Multi-fidelity Monte Carlo Methods for Uncertainty Quantification

4:50 p.m.-6:50 p.m.

Room: Riverside Hall D

Panagiotis Tsilifis, École Polytechnique Fédérale de Lausanne, Switzerland

Minisymposterium: Markov Chain Monte Carlo Methods for Seismic Source Inversion

Juan Pablo Madrigal Cianci and Fabio Nobile, École Polytechnique Fédérale de Lausanne, Switzerland; Raul F. Tempone, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

Minisymposterium: Kullback-Leibler Algorithms for Bayesian Inverse Problems using Multi-Level Monte Carlo Methods

Panagiotis Tsilifis, Juan Pablo Madrigal Cianci, and Fabio Nobile, École Polytechnique Fédérale de Lausanne, Switzerland

Minisymposterium: Some Recent Advancements on Multifidelity Monte Carlo Techniques for Uncertainty Quantification

Gianluca Geraci, Sandia National Laboratories, U.S.; Alex A. Gorodetsky, University of Michigan, U.S.; Michael S. Eldred and John D. Jakeman, Sandia National Laboratories, U.S.

Minisymposterium: Computation of Electromagnetic Fields Scattered From Objects of Uncertain Shapes Using Multilevel Monte Carlo

Alexander Litvinenko and *Raul F. Tempone*, King Abdullah University of Science & Technology (KAUST), Saudi Arabia; Abdulkadir Yucel, Nanyang Technological University, Singapore; Hakan Bagci, King Abdullah University of Science & Technology (KAUST), Saudi Arabia; Jesper Oppelstrup, KTH Royal Institute of Technology, Sweden; Erik Michielssen, University of Michigan, U.S.

Wednesday, February 27

PP204

Minisymposterium: Student Days - Student Chapter Posters

4:50 p.m.-6:50 p.m.

Room: Riverside Hall D

Kathleen Kavanagh, Clarkson University, U.S.

Poster presentations from SIAM Student Chapters.

Minisymposterium: Computational Developments for the Bayesian Conjugate Gradient Method

Timothy Ried, North Carolina State University, U.S.

See online program for additional posters in this session.

Wednesday, February 27

PP205

Minisymposterium: Student Days - Undergraduate Posters

4:50 p.m.-6:50 p.m.

See online program for an update on this session.

Business Meeting

7:00 p.m.-8:00 p.m.

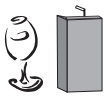

Room: Ballroom 100BC

Complimentary beer and wine will be served.

Thursday, February 28

Registration

7:45 a.m.-4:30 p.m.

Room: Ballroom Foyer

Awards Announcement

8:00 a.m.-8:15 a.m.

Room: Ballroom 100BC

Thursday, February 28

SP2

SIAG/CSE Early Career Prize: Data-Driven Discovery and Control of Complex Systems: Uncovering Interpretable and Generalizable Nonlinear Models

8:15 a.m.-8:45 a.m.

Room: Ballroom 100BC

Chair: To Be Determined

Accurate and efficient reduced-order models are essential to understand, predict, estimate, and control complex, multiscale, and nonlinear dynamical systems. These models should ideally be generalizable, interpretable, and based on limited training data. This work develops a general framework to discover the governing equations underlying a dynamical system simply from data measurements, leveraging advances in sparsity-promoting techniques and machine learning. The resulting models are parsimonious, balancing model complexity with descriptive ability while avoiding overfitting. This perspective, combining dynamical systems with machine learning and sparse sensing, is explored with the overarching goal of real-time closed-loop feedback control of unsteady fluid systems. We will discuss how to enforce known constraints, such as energy conserving quadratic nonlinearities in incompressible fluids, to "bake in" known physics. Next, we will demonstrate that higher-order nonlinearities can approximate the effect of truncated modes, resulting in more accurate models of lower order than Galerkin projection. Finally, we will discuss the use of intrinsic measurement coordinates to build nonlinear models, circumventing the well-known issue

continued on next page

Thursday, February 28

SP2

SIAG/CSE Early Career Prize: Data-Driven Discovery and Control of Complex Systems: Uncovering Interpretable and Generalizable Nonlinear Models continued

of continuous mode deformation associated with methods based on the proper orthogonal decomposition. This approach is demonstrated on several relevant systems with low-dimensional dynamics.

Steven Brunton
University of Washington, U.S.

Thursday, February 28

SP3

SIAM/ACM Prize in Computational Science and Engineering: The Singular Value Decomposition: Anatomy of an Algorithm, Optimizing for Performance

8:45 a.m.-9:15 a.m.

Room: Ballroom 100BC

Chair: To Be Determined

The computation of the singular value decomposition, or SVD, has a long history with many improvements over the years, both in its implementations and algorithmically. Here, we look at the evolution of SVD algorithms for dense matrices, discussing the motivation and performance impacts of changes. We will look at the changes over time by testing various historical and current implementations on a common, modern multicore machine and a distributed computing platform. We show that algorithmic and implementation improvements have increased the speed of the SVD by several orders of magnitude, while using up to 40 times less energy.

Jack J. Dongarra
University of Tennessee and Oak Ridge National Laboratory, U.S.

Coffee Break

9:15 a.m.-9:45 a.m.

Room: Ballroom Foyer

Thursday, February 28

MT3

MOOSE: Enabling Multiphysics

9:45 a.m.-11:25 a.m.

Room: 300D

The Multiphysics Object Oriented Simulation Environment (MOOSE) is an open-source, modular system for creating simulation tools where highly coupled systems of equations need to be solved simultaneously in massively-parallel environments. This tutorial will introduce the core systems within MOOSE to begin building a custom application to solve your unique and challenging problems.

Derek R. Gaston, Idaho National Laboratory, U.S.

Andrew Slaughter, Idaho National Laboratory, U.S.

Thursday, February 28

MS238

Theory and Application of Surrogate Models for Bayesian Inverse Problems - Part I of II

9:45 a.m.-11:25 a.m.

Room: Ballroom 100BC

For Part 2 see MS271

Today the availability of exascale computing resources enables high-fidelity simulations of complex processes in science and engineering. Moreover, it is becoming standard to include and quantify uncertainties in these processes. However, the forward and inverse uncertainty quantification with computationally demanding models is still infeasible in many situations. Expensive models arise for instance in hydrology, oncology, and multiphysics simulations such as fluid structure interaction. In practice, such models can be replaced by computationally cheap surrogates based on full model runs. Surrogates enable and accelerate uncertainty quantification. This minisymposium is dedicated to surrogate modelling for inverse uncertainty quantification. We highlight recent theoretical advances from mathematical and computational sciences, for example, adaptive strategies for surrogate construction, surrogate error analysis, and multifidelity surrogates, together with results from practical estimation problems. Surrogate models include Gaussian processes, sparse grids, reduced bases, (piecewise) polynomials, and low-rank tensors.

Organizer: Jonas Latz

Technische Universität München, Germany

Organizer: Claudia Schillings

Universitaet Mannheim, Germany

Organizer: Elisabeth Ullmann

Technische Universität München, Germany

9:45-10:05 Fast Sampling of Parameterised Gaussian Random Fields

Jonas Latz and Elisabeth Ullmann, Technische Universität München, Germany

10:10-10:30 Adaptive Multi-fidelity Polynomial Chaos Approach to Bayesian Inference in Inverse Problems

Yan Liang, Southeast University, China; Tao Zhou, Chinese Academy of Sciences, China

10:35-10:55 Dimension Adaptive Sparse Quadrature and Sparse Polynomial Parametrized Transport Maps for High Dimensional Bayesian Integration

Joshua Chen and Peng Chen, University of Texas at Austin, U.S.; Daniele Bigoni and Youssef M. Marzouk, Massachusetts Institute of Technology, U.S.; Omar Ghattas, University of Texas at Austin, U.S.

11:00-11:20 Marginal MCMC with Noisy Local Approximations

Andrew D. Davis and Youssef M. Marzouk, Massachusetts Institute of Technology, U.S.; Natesh Pillai, Harvard University, U.S.; Aaron Smith, University of Ottawa, Canada

continued in next column

Thursday, February 28

MS239

Multiphysics: Extensible, Composable Algorithms and Software - Part I of II

9:45 a.m.-11:25 a.m.

Room: Conference Theater

For Part 2 see MS272
Featured Minisymposium

A major thrust for more than two decades has been the development of flexible and extensible methodologies for tackling the most challenging scientific and engineering problems. These problems naturally exhibit multiscale and multiphysics characteristics which strain algorithmic foundations and software capabilities, assumptions, and optimization while targeting contemporary and emerging hardware. In this environment, it is not possible for a single person to understand and customize components for each class of target problem, thus algorithmic and software robustness, composability, and extensibility is paramount. Training, automation, and quality control become increasingly challenging as users and contributors from diverse backgrounds attempt to conduct increasingly complex simulations and analyses. This session brings together exemplars, promising techniques, and lessons learned in developing algorithms, software, and communities for computational multiphysics.

Organizer: Paul Bauman

State University of New York at Buffalo, U.S.

9:45-10:05 Composable Multiphysics: From Solvers to Quantities of Interest

Paul Bauman, State University of New York at Buffalo, U.S.

10:10-10:30 Multiscale Modeling of Bloodflow in the Human Vasculature

Amanda Randles, Duke University, U.S.; John Gounley, Oak Ridge National Laboratory, U.S.

continued on next page

Thursday, February 28

MS239

Multiphysics: Extensible, Composable Algorithms and Software - Part I of II

continued

10:35-10:55 preCICE -- a Comprehensive Coupling Library for Large-scale Surface-coupled Multi-physics Problems

Miriam Mehl, Universität Stuttgart, Germany

11:00-11:20 Scalable Solution of Implicit / Imex Fe Continuum Plasma Physics Models

John Shadid, Roger Pawlowski, Edward G. Phillips, Paul Lin, Sidafa Conde, Ray S. Tuminaro, and Sibusiso Mabuza, Sandia National Laboratories, U.S.; Luis Chacon, Los Alamos National Laboratory, U.S.

Thursday, February 28

MS240

Boundary Integral Methods for Particulate Flows - Part I of II

9:45 a.m.-11:25 a.m.

Room: 102A

For Part 2 see MS273

This minisymposium will focus on recent advances in the development of integral equation methods for simulating low-Re viscous flow of bubbles, drops, vesicles, rigid particles and other particulate suspensions. The accurate simulation of such complex fluid flows is challenging because of strong nonlinearities, time-dependent geometries, close interactions, and multiple length and time scales. Integral equation methods offer a suitable and powerful framework able to address these challenges while providing high-order accuracy.

Organizer: Lukas Bystricky
Florida State University, U.S.

Organizer: Chiara Sorgentone
KTH Royal Institute of Technology, Sweden

9:45-10:05 Numerical Simulations of 3D Surfactant-covered Drop Electrohydrodynamics

Chiara Sorgentone and Anna-Karin Tornberg, KTH Royal Institute of Technology, Sweden

10:10-10:30 Transport in Viscous Erosion

Bryan D. Quaife and Matthew N. Moore, Florida State University, U.S.

10:35-10:55 Optimal Design of Deterministic Lateral Displacement Devices for Cell Sorting

Gokberk Kabacaoglu, University of Texas, U.S.; George Biros, University of Texas at Austin, U.S.

11:00-11:20 Numerical Simulations of Two-phase Flow of Electrolyte Drops

Michael Siegel, New Jersey Institute of Technology, U.S.

Thursday, February 28

MS241

Close Evaluation of Layer Potentials: Advanced Numerical Methods and Applications

9:45 a.m.-11:25 a.m.

Room: 102B

Boundary integral equation methods are useful for solving boundary value problems for linear, elliptic, partial differential equations. Rather than solving the partial differential equation directly, one represents the solution as a layer potential, an integral operator applied to a density over the boundary. The close evaluation problem arises in boundary integral methods for numerically computing the solution for evaluation points near the boundary. It has been established that accuracy of standard numerical methods greatly suffer from the close evaluation problem. This session will include recent advances in numerical methods for boundary integral equation methods near boundaries and application of these for methods for a diverse range of problem. In particular it will cover electromagnetic wave propagation and Stokes flow. Topics may include development and analysis of numerical methods for boundary integral equations, asymptotic analysis, applications to electromagnetic wave propagation or scattering in biological tissue or exotic metamaterials, fluid-structure or multiphase flow problems, and more.

Organizer: Camille Carvalho
University of California, Merced, U.S.

Organizer: Shilpa Khatri
University of California, Merced, U.S.

9:45-10:05 Asymptotic Analysis for the Close Evaluation Problem

Camille Carvalho, Shilpa Khatri, and *Arnold D. Kim*, University of California, Merced, U.S.

continued on next page

10:10-10:30 An Efficient Method for Nearly Singular Line Integrals in Two and Three Dimensions
Ludvig af Klinteberg, Simon Fraser University, Canada

10:35-10:55 Close Interactions of Particles in Stokes Flow
Anna-Karin Tornberg, KTH Royal Institute of Technology, Sweden

11:00-11:20 A Fredholm Operator Approach to Clamped Plate Problems
Travis Askham, New Jersey Institute of Technology, U.S.

Thursday, February 28

MS242

Recent Advances in Models and Numerical Methods for Multiphase Problems - Part I of II

9:45 a.m.-11:25 a.m.

Room:102C

For Part 2 see MS275

Multiphase problems arise in a wide range of computational science and engineering applications. They are ubiquitous in the universe. They exist in material systems consisting of coexistent, multiple physical phases (gas, liquid, solid) and/or a single physical phase but multiple components (i.e., oil/water, polymer/polymer/solvent). How to develop the thermodynamically consistent description of the multiphase phenomena? How to develop appropriate numerical strategies and computational methods to efficiently solve the multiphase problems are pressing issues in the research community. Phase field has emerged as an effective modeling and computational tool to deal with multiple phase phenomena recently. How to exploit the efficiency of the formulation and computational flexibility and how to couple this efficient computational strategy to other methods for multiphase problems will be the central theme of the symposium. In this minisymposium, experts in modeling and computation are invited to share their recent advances in studying multiphase problems and relevant applications. It serves as a platform for advancing modeling and numerical strategies in the related field.

Organizer: Jia Zhao
Utah State University, U.S.

Organizer: Qi Wang
University of South Carolina, U.S. and Beijing Computational Science Research Center, China

9:45-10:05 Compressible Multiphase Fluid Flows
Qi Wang, University of South Carolina, U.S. and Beijing Computational Science Research Center, China

10:10-10:30 Thermal Effects in Transport of Ionic Solutions
Chun Liu, Illinois Institute of Technology, U.S.

10:35-10:55 Approximation of Contact Angle Hysteresis by using Onsager Principle
Xianmin Xu, Chinese Academy of Sciences, China

11:00-11:20 Cell Motility Dependence on Adhesive Wetting
Yanxiang Zhao, George Washington University, U.S.

continued in next column

Thursday, February 28

MS243

Reduced Order Models for Fluids: Achievements and Open Problems - Part I of II

9:45 a.m.-11:25 a.m.

Room:102D

For Part 2 see MS276

This minisymposium aims at giving a survey of recent developments in the reduced order modeling of fluid flows. Computational modeling, numerical analysis, and applications to realistic engineering and geophysical flow problems will be covered in this minisymposium. Both achievements and open problems in the reduced order modeling of fluid flows will be discussed.

Organizer: Traian Iliescu

Virginia Tech, U.S.

Organizer: Michael Schneier

University of Pittsburgh, U.S.

9:45-10:05 Model Reduction of Dynamical Systems on Nonlinear Manifolds using Deep Convolutional Autoencoders

Kevin T. Carlberg and Kookjin Lee, Sandia National Laboratories, U.S.

10:10-10:30 Reduced-order Model for the BGK Equation based on Pod and Optimal Transport

Florian Bernard, Inria and University of Bordeaux, France; Angelo Iollo, Institut de Mathématiques de Bordeaux, France; *Sebastien Riffaud*, Inria and University of Bordeaux, France

10:35-10:55 Physically-constrained Data-driven Correction for Reduced Order Modeling of Fluid Flows

Muhammad Mohebujjaman, Virginia Tech, U.S.; Leo Rebholz, Clemson University, U.S.; Traian Iliescu, Virginia Tech, U.S.; Xuping Xie, Oak Ridge National Laboratory, U.S.

11:00-11:20 Reduced Stochastic Models for Complex Fluids through Computational Information Geometry

Sorin Mitran, University of North Carolina, Chapel Hill, U.S.

Thursday, February 28

MS244

Theory and Algorithms for Improved Performance of Machine Learning in Scientific Applications - Part I of II

9:45 a.m.-11:25 a.m.

Room: 111A

For Part 2 see MS277

Machine learning (ML) and artificial intelligence (AI) have had unprecedented success in Silicon Valley, and for big data problems that have emerged in the last decade. However, the application of ML to scientific discovery requires an entirely new focus on algorithmic development and mathematical rigor, aimed at producing reliable, confident, and reproducible results, customized to specific type (plentiful or sparse) and form (labeled or unlabeled, static or online, deterministic or noisy) of observational data. An overarching goal of this symposium is to amalgamate international experts in numerical analysis, approximation theory, uncertainty quantification, and data science, to systematically investigate: the design of architectures/models which accurately capture the complexities of the data, with robust estimates of confidence in predictions on defined domains; and fast and scalable algorithms to fit the proposed models to data, with a theory that explains the convergence and success of these techniques.

Organizer: Guannan Zhang

Oak Ridge National Laboratory, U.S.

Organizer: Clayton G. Webster

University of Tennessee and Oak Ridge National Laboratory, U.S.

9:45-10:05 Approximability Models and Optimal System Identification

Simon Foucart, Texas A&M University, U.S.

10:10-10:30 Large-scale Classification Techniques with Applications to Medical Data

Deanna Needell, University of California, Los Angeles, U.S.

10:35-10:55 An Iterative Method for Classification of Binary Data

Denali Molitor, University of California, Los Angeles, U.S.

11:00-11:20 Deep Learning for Quasi-optimal Polynomial Approximations of High-dimensional Functions

Clayton G. Webster, University of Tennessee and Oak Ridge National Laboratory, U.S.

continued in next column

Thursday, February 28

MS245

Models with Superior Reasons for Complex Big Data - Part I of II

9:45 a.m.-11:25 a.m.

Room: 111B

For Part 2 see MS278

This minisymposium concerns state-of-the-art computational methods for solving real-world problems involving complex high-dimensional big-volume data stemming from heterogenous sources, such as high-frequency data from finance or complex claims records in insurance, spatio-temporal sensory data from mobile connected devices, semi- or fully autonomous transportation, and emerging Internet-of-Things (IoT) applications. We are particularly interested in building highly accurate learning and prediction models with excellent decision reasons: overall variable importance at the model level down to personalized key features at the data record level, which demonstrate superior model intelligence and hence inspires confidence and trust. Enabling software tools and hardware architecture will be highlighted along with key development principles. This minisymposium is partially sponsored by American Family Insurance and eMALI. IO Ltd.

Organizer: Sou-Cheng T. Choi

Illinois Institute of Technology, U.S. and Allstate Insurance Corporation, U.S.

Organizer: Luisa Polania

Target Corporation, U.S.

Organizer: Teja Kanchinadam

American Family Insurance, U.S.

Organizer: Lawrence K.H. Ma

Hong Kong Blockchain Society, Hong Kong

9:45-10:05 Real-time Prediction of Traffic Speed During Traffic Incidents

Sou-Cheng T. Choi, Illinois Institute of Technology, U.S. and Allstate Insurance Corporation, U.S.; Rory Hartong-Redden, Tim Gibson, and Sunil Chintakindi, Allstate Insurance Corporation, U.S.

10:10-10:30 Dealing with Uncertainties in Investment Robo-advice

Shi Yu, Vanguard, U.S.

10:35-10:55 Using Discriminative Graphical Models for Insurance Recommender Systems

Teja Kanchinadam, American Family Insurance, U.S.

11:00-11:20 Mapping Advanced Knowledge: the Lessons from Economic-climate and Macro-financial Modeling

Victor Zhorin, Independent Researcher

continued in next column

Thursday, February 28

MS246

Machine Learning Methods in Computational Fluid Dynamics - Part I of II

9:45 a.m.-11:25 a.m.

Room: 111C

For Part 2 see MS279

Machine learning (ML) methods and in particular deep learning via neural networks have generated significant interest in the last years. This interest reaches beyond the ML community itself into other fields of science and engineering. Since these methods can provide approximations to general functions by being coaxed to learn optimally from data without a-priori assumptions, they are particularly attractive for the generation of subspace models. In the area of computational fluid dynamics, research into how ML methods can enhance current capabilities is an active research topic. Deep learning methods have been shown to provide accurate shock-capturing sensors, improve RANS turbulence models, provide approximate deconvolutions of a coarse-scale flow fields and reconstruct the exact LES stresses terms to generate data-driven closure models. With these encouraging successes, a number of questions arise: What did the models actually learn? How can one incorporate physical constraints into the learning process? What is the confidence in the model predictions? And what are the limits of ML and can we quantify them? This symposium invites contributions that discuss how to enhance CFD methods by machine learning algorithms. The generation of subspace models, e.g. for RANS and LES, is of particular interest, but the minisymposium is also open to other topics, e.g. learning intelligent flow sensors, data compression and feature extraction, implementation examples etc.

continued on next page

Thursday, February 28

MS246

Machine Learning Methods in Computational Fluid Dynamics - Part I of II

continued

Organizer: Andrea D. Beck
Universität Stuttgart, Germany

Organizer: Romit Maulik
Oklahoma State University, U.S.

Organizer: Omer San
Oklahoma State University, U.S.

9:45-10:05 Data-driven Filter Estimation for the Sub-grid Modelling of Kraichnan Turbulence

Romit Maulik and Omer San, Oklahoma State University, U.S.; Adil Rasheed, SINTEF Digital, Norway; Prakash Vedula, University of Oklahoma, U.S.

10:10-10:30 Deep Neural Networks for Data-driven Turbulence Models

Andrea D. Beck and David Flad, Universität Stuttgart, Germany; Claus-Dieter Munz, Institut fuer Aerodynamik und Gasdynamik (IAG), Germany

10:35-10:55 Deep Learning for Data-driven Modeling and Analysis of Fluid Flows

Balaji Jayaraman, Shivakanth Chary, and Chen Lu, Oklahoma State University, U.S.

11:00-11:20 Convolutional Neural Networks for Direct Prediction of Aerodynamic Flowfields

Saakaar Bhatnagar, University of Michigan, U.S.; Shaowu Pan, University of Michigan, Ann Arbor, U.S.

Thursday, February 28

MS247

Recent Advances in PDE-constrained Optimization under Uncertainty - Part I of II

9:45 a.m.-11:25 a.m.

Room: 300A

For Part 2 see MS280

Over the past two decades, significant strides have been made in theory and algorithms for optimization of systems governed by PDEs. These advances in PDE-constrained optimization have enabled many applications of optimal design and control to complex deterministic systems in computational science and engineering. A major remaining challenge for PDE-constrained optimization is accounting for uncertainty in the PDE models. Many PDE models are characterized by uncertain parameters due to lack of knowledge or intrinsic variability of the inputs, including initial or boundary conditions, sources, coefficients, or geometry. It is important to incorporate this uncertainty in the optimization problem to make the optimal solution more robust and reliable. Indeed, recent years have witnessed rapid growth in research on PDE-constrained optimization under uncertainty. In this minisymposium, leading experts will present their work on advances in the development, analysis, and application of methods for PDE-constrained optimization under uncertainty. In particular, the following themes will be emphasized: (1) development and analysis of scalable algorithms to solve optimization problems under high- or infinite-dimensional uncertainty; (2) investigation of different risk measures and constraints; (3) application to problems governed by more challenging models, including multiphysics, multiscale, and fractional PDE problems.

Organizer: Peng Chen
University of Texas at Austin, U.S.

Organizer: Omar Ghattas
University of Texas at Austin, U.S.

9:45-10:05 Breaking the Curse of Dimensionality for PDE-constrained Optimization under High-dimensional Uncertainty

Peng Chen and Omar Ghattas, University of Texas at Austin, U.S.

10:10-10:30 Adaptive Multifidelity Contour Estimation for Reliability-based Design Optimization

Anirban Chaudhuri and Alexandre Marques, Massachusetts Institute of Technology, U.S.; Karen E. Willcox, University of Texas at Austin, U.S.

10:35-10:55 Robust Optimization of PDE-constrained Problems using Second-order Models

Philip Kolvenbach and Stefan Ulbrich, Technische Universität Darmstadt, Germany

11:00-11:20 Sparse Optimal Control of PDEs under Uncertainty

Chen Li, New York University, U.S.; Georg Stadler, Courant Institute of Mathematical Sciences, New York University, U.S.

continued in next column

Thursday, February 28

MS248

High Performance Sparse Matrix, Tensor, and Graph Kernels - Part I of II

9:45 a.m.-11:25 a.m.

Room: 300B

For Part 2 see MS281

High performance computational kernels are critical to various real-world applications, such as scientific computing, machine learning, social networks, and healthcare analytics, to name a few. Data from many applications are sparse, which means most of the entries are meaningless or missing values, represented as zeros. In this proposed minisymposium, we will discuss state-of-the-art studies of sparse computational kernels in sparse matrix, sparse tensor, and graph problems. We will discuss their interactions, challenges, solutions, and future directions. Our topics cover these sparse kernels from different views: sparse kernels as benchmarks to compare different computer architectures, new sparse data structures and optimization methods for these algorithms, and the applications of these kernels. Our topics also cover the optimization techniques from different domains: algorithms, compilers, runtime systems, and computer architectures. We expect this minisymposium will explore high-performance sparse algorithms and their interactions among different kernels.

Organizer: Jiajia Li

Pacific Northwest National Laboratory, U.S.

Organizer: Xu Liu

College of William & Mary, U.S.

9:45-10:05 Gluon: A Communication Optimizing Framework for Distributed Heterogeneous Graph Analytics

Roshan Dathathri and Keshav Pingali, University of Texas at Austin, U.S.

10:10-10:30 A Model Driven Sparse Gemm on GPUs

Xin He and Guangming Tan, Chinese Academy of Sciences, China

10:35-10:55 Enabling On-the-fly Storage Format Prediction and Optimization for SpMV

Weijie Zhou and Xipeng Shen, North Carolina State University, U.S.

11:00-11:20 CVR: Efficient SpMV Vectorization on X86 Processors

Xu Liu, College of William & Mary, U.S.; Biwei Xie, Chinese Academy of Sciences, P.R. of China

continued in next column

Thursday, February 28

MS249

Tensor Based Methods in Scientific Computing and Data Science - Part I of II

9:45 a.m.-11:25 a.m.

Room: 300C

For Part 2 see MS282

Tensor-based methods are becoming increasingly important because of their ability to extract and exploit inherent multidimensional structure. However, several computational challenges persist that prevent the use of tensor decompositions for large-scale problems. This minisymposium explores advances in numerical methods and applications involving tensor decompositions in scientific computing and data sciences.

Organizer: Arvind Saibaba

North Carolina State University, U.S.

Organizer: Misha E. Kilmer

Tufts University, U.S.

9:45-10:05 Stochastic Gradient Descent for Large-scale Generalized CP Tensor Decomposition

Tamara G. Kolda, Sandia National Laboratories, U.S.; David Hong, University of Michigan, U.S.; Jed A. Duersch, Sandia National Laboratories, U.S.

10:10-10:30 Stable Tensor Neural Networks

Elizabeth Newman, Tufts University, U.S.; Lior Horesh, IBM Research, U.S.; Haim Avron, Tel Aviv University, Israel; Misha E. Kilmer, Tufts University, U.S.

10:35-10:55 Tensor Canonical Correlation Analysis - A Probabilistic Method and Estimation Algorithm

Eric Chi, North Carolina State University, U.S.; Hua Zhou, University of California, Los Angeles, U.S.; Eun Jeong Min, University of Pennsylvania, U.S.

11:00-11:20 Eigenvector-based Centrality Measures in Multilayer Networks -- Theory and Computation

Francesca Arrigo, University of Strathclyde, United Kingdom; Francesco Tudisco, University of Padua, Italy; Antoine Gautier, Saarland University, Germany

Thursday, February 28

MS250

BE: Advances in Computational Drug Discovery - Part I of II

9:45 a.m.-11:25 a.m.

Room: 302A

For Part 2 see MS283

Drug discovery and development continues to be an expensive, lengthy, and sometimes fatal process. With the large chemical space of potential drugs and the number of proteins and observable genetic variation of these proteins, an experimental and clinical understanding of all potential interactions is not feasible, making accurate computational studies extremely relevant and important. The massive computational power of leadership class computing and advances in machine learning provide exciting opportunities to overcome hurdles in drug discovery and development. Drugs are known to be promiscuous; they interact with more than one protein. Some estimates predict that every protein interacts with tens to hundreds of different known drugs. Accurate predictions of how drugs interact with the entire proteome would establish computational polypharmacological networks. These networks would provide insight on drug repurposing, side-effect prediction, and the development of more efficacious drugs. This session focuses on combining traditional biomolecular simulations with biomedical big data sources using machine learning techniques to improve drug binding predictions to enable the development of accurate polypharmacological networks.

Organizer: Sally R. Ellingson

University of Kentucky, U.S.

9:45-10:05 Opportunities for Polypharmacology Using Machine Learning

Sally R. Ellingson, University of Kentucky, U.S.

10:10-10:30 Accelerating Therapeutics for Opportunities in Medicine

Jonathan Allen, Lawrence Livermore National Laboratory, U.S.

10:35-10:55 Feature Selection in Biomolecular Models

Julie Mitchell, Oak Ridge National Laboratory, U.S.

11:00-11:20 Leveraging Large Scale Molecular Dynamics Simulations and Deep Learning for Binding Activity Models

William D. Jones, Lawrence Livermore National Laboratory, U.S.

continued in next column

Thursday, February 28

MS251

Algorithms and Software for Nonlinear Eigenvalue Problems - Part I of II

9:45 a.m.-11:25 a.m.

Room: 302B

For Part 2 see MS284

Eigenvalue problems arise in many fields of science and engineering and their mathematical properties and numerical solution methods for standard, linear eigenvalue problems are well understood. Recently, due to breakthrough applications, e.g., electronic structure calculations for 2D materials, we observe emerging challenges in the nonlinear eigenvalue problems (NLEVPs) $T(\lambda) x=0$, which exhibits nonlinearity in the eigenvalue parameter. Although, the nonlinear eigenvalue problems received a lot of attention from the numerical linear algebra community during the last decade, the majority of the work has been focused on polynomial eigenvalue problems. The minisymposium will provide a forum for discussing the novel algorithmic contributions and latest software developments for the general nonlinear eigenvalue problems involving nonlinear functions such as exponential, rational, and irrational ones. Presentations will focus on Newton-based techniques, Krylov subspace methods applied to linearizations, and contour integration and rational filtering methods and their role in large-scale science and engineering simulations.

Organizer: Roel Van Beeumen

Lawrence Berkeley National Laboratory, U.S.

Organizer: Agnieszka Miedlar

University of Kansas, U.S.

9:45-10:05 NLEVP: A Collection of Nonlinear Eigenvalue Problems

Francoise Tisseur and Gian Maria Negri Porzio, University of Manchester, United Kingdom

continued on next page

10:10-10:30 NEP-PACK: A Julia Package for Nonlinear Eigenproblems
Elias Jarlebring, KTH Royal Institute of Technology, Sweden

10:35-10:55 The Rational Krylov Toolbox
Steven Elsworth and Stefan Guettel, University of Manchester, United Kingdom

11:00-11:20 {CORK}\texttt{++}: Compact Rational Krylov Methods for Nonlinear Eigenvalue Problems
Roel Van Beeumen, Lawrence Berkeley National Laboratory, U.S.

Thursday, February 28

MS252

Preconditioning for High-order Matrix-free PDE Operators - Part I of II

9:45 a.m.-11:25 a.m.

Room: 303A

For Part 2 see MS285

This minisymposium explores various approaches to preconditioning operators when no explicit matrix representation is easily available. An ongoing move in high-performance computing toward next-generation high-concurrency hardware makes the higher arithmetic intensity of high polynomial order discretizations much more attractive, but explicit matrix assembly and application of such high-order operators rapidly becomes a performance bottleneck. A partial solution to this issue is representing the operator in an unassembled or matrix-free setting, taking advantage for example of tensor representations of a finite element basis. In the unassembled setting, however, it is not at all clear how to precondition the operator, especially in situations where a geometric mesh hierarchy is not available. In this minisymposium, we consider approaches such as preconditioning a high-order method with a lower-order method, unassembled algebraic multigrid preconditioners, and related preconditioning methods.

Organizer: Andrew T. Barker
Lawrence Livermore National Laboratory, U.S.

Organizer: Per-Olof Persson
University of California, Berkeley, U.S.

9:45-10:05 Algebraic Multigrid-based Approaches for Preconditioning Unassembled Operators in H(curl)
Andrew T. Barker and Tzanio Kolev, Lawrence Livermore National Laboratory, U.S.

10:10-10:30 Approximate Fast Diagonalization for the Spectral Element Method
Pablo Brubeck and Paul Fischer, University of Illinois at Urbana-Champaign, U.S.

10:35-10:55 Efficient Solvers and Tensor-product Preconditioners for the Implicit Time Integration of Discontinuous Galerkin Methods
Will Pazner and Per-Olof Persson, University of California, Berkeley, U.S.

11:00-11:20 Acceleration of Tensor-product Operations for High-order Finite Element Methods
Tim Warburton, Virginia Tech, U.S.

continued in next column

Thursday, February 28

MS253

Mitigating Communication Costs Using Variable Precision Computing Techniques - Part I of II

9:45 a.m.-11:25 a.m.

Room: 303B

For Part 2 see MS286

Communication costs, including global data movement between processors and local movement within a processor, are becoming significantly more expensive compared to arithmetic operations. It is now of high importance to develop algorithms with communications costs in mind. Traditionally, to guard against floating-point round off errors, we store data in a 64-bit double precision representation even when very few significant digits are needed. However, many scientific simulations have traditional error polluting the 64-bit representation - truncation error, iteration error, and floating-point round off error. Consequently, since many of the bits represent error, we are wasting valuable resources. Thus, a simple way to reduce communication costs is to reduce the size of the data; either using fewer bits, i.e., in a lower precision, or by using a lossy compression algorithm. Mixed-precision or lower precision algorithms have had a renewed interest not only to mitigate communication costs, but due to the fact that GPU's can perform half precision arithmetic approximately 16 times as fast as double precision arithmetic. In the past few years, lossy compression algorithms have been specifically designed for scientific data, sacrificing a small amount precision and additional computations for reduced bandwidth. To this end, this minisymposium will highlight recent algorithmic advances in variable precision computing.

Organizer: Alyson Fox
Lawrence Livermore National Laboratory, U.S.

Organizer: Jeffrey A. Hittinger
Lawrence Livermore National Laboratory, U.S.

Organizer: James D. Diffenderfer
University of Florida, U.S.

9:45-10:05 Error Analysis of ZFP Compression for Floating-point Data

James D. Diffenderfer, University of Florida, U.S.

10:10-10:30 Stability Analysis of Inline ZFP Compression for Floating-point Data in Iterative Methods

Alyson Fox, Lawrence Livermore National Laboratory, U.S.

10:35-10:55 Error Analysis of ZFP Compression in Multigrid Methods

Avary Kolasinski, University of Kansas, U.S.

11:00-11:20 Beyond IEEE: Next-generation Floating-point Formats

Peter Lindstrom, Scott Lloyd, and Jeffrey A. Hittinger, Lawrence Livermore National Laboratory, U.S.

continued in next column

Thursday, February 28

MS254

Mori-Zwanzig Formulation, Implementation and Applications - Part I of II

9:45 a.m.-11:25 a.m.

Room: 201C

For Part 2 see MS288

The Mori-Zwanzig formalism allows reducing the number of variables in large systems of coupled equations. For differential equations, the reduced equations model the effect of the unresolved variables, leading to a Markovian, memory and fluctuating terms. This formalism can be a starting point for multiscale and meso-scale modeling, based on first-principles calculations. It can also be used to derive popular reduced order models as appropriate limiting cases. In addition to formulating better reduced models, there is a growing need for efficient simulation of those reduced models. The minisymposium will address the related mathematical and numerical issues as well as applications to materials, fluid mechanics, soft matter, biology and uncertainty quantification.

Organizer: Xiantao Li
Pennsylvania State University, U.S.

Organizer: Panos Stinis
Pacific Northwest National Laboratory, U.S.

Organizer: Daniele Venturi
University of California, Santa Cruz, U.S.

Organizer: Karthik Duraisamy
University of Michigan, Ann Arbor, U.S.

9:45-10:05 Coarse-graining Strategies for Multi-scale Partial Differential Equations using the Mori-Zwanzig Approach

Aniruddhe Pradhan, The University of Michigan, Ann Arbor, U.S.; Karthik Duraisamy, University of Michigan, Ann Arbor, U.S.

continued on next page

10:10-10:30 Coarse-graining the Overdamped Langevin Equation for a Potential with Entropic Barriers via the Mori-Zwanzig Formalism
Thomas Hudson, University of Warwick, United Kingdom

10:35-10:55 The Derivation of Fluctuating Heat Conduction Models from Molecular Dynamics
Weiqi Chu and Xiantao Li, Pennsylvania State University, U.S.

11:00-11:20 Data-driven Modelling and Parameterization of Generalized Langevin Equation
Huan Lei, Pacific Northwest National Laboratory, U.S.

Thursday, February 28

MS255

Underwater Sensing and Signal Processing - Part I of II

9:45 a.m.-11:25 a.m.

Room:202A

For Part 2 see MS289

The earth's vast oceans are largely unexplored, but at the same time are of great importance for economic, ecological and military objectives. Advances in sensing technology (e.g. vector sensors), platform capabilities (e.g. Unmanned Underwater Vehicles) and computational resources create an exciting environment for new signal processing algorithms and sensing opportunities. Many of the problems involve computational challenges such as prediction of long-range acoustic propagation through a changing, multiple-scale environment, and simultaneous processing of data from hundreds or thousands of sensors. This minisymposium will identify current directions of research and methods being developed to solve problems, including their advantages and shortcomings.

Organizer: Margaret Cheney
Colorado State University, U.S.

Organizer: Kevin P. Bongiovanni
Raytheon BBN Technologies, U.S.

9:45-10:05 Problems in Underwater Acoustics
Kevin P. Bongiovanni, Raytheon BBN Technologies, U.S.

10:10-10:30 Shallow Water Inversion with Sequential Filtering and Linearization
Eliza (Z.-H.) Michalopoulou, New Jersey Institute of Technology, U.S.

10:35-10:55 Localization in a Random Medium Using Lucky Techniques
Ivars Kirsteins, Naval Undersea Warfare Center, U.S.

11:00-11:20 Maximal Invariants and the Singular Value Decomposition in Detection Theory
Louis Scharf, Colorado State University, U.S.

Thursday, February 28

MS256

Industrial Eigensolution Technology: Advances and Challenges - Part I of II

9:45 a.m.-11:25 a.m.

Room:202B

For Part 2 see MS274

Numerical eigenvalue problems and eigensolution technology are one of the centerpiece of industrial scale modeling and simulations for acoustics and structural analysis among many others. In this minisymposium, speakers from industry and academic will present recent advancements in eigenvalue algorithms, numerical analysis and industrial strength high-performance software, and discuss emerging needs and challenges.

Organizer: Zhaojun Bai
University of California, Davis, U.S.

Organizer: Roger Grimes
Livermore Software Technology Corporation, U.S.

9:45-10:05 Eigenvalue Problems in Mechanical Engineering
Roger Grimes, Livermore Software Technology Corporation, U.S.

10:10-10:30 FISS: A Scalable Parallel Eigenvalue Solver for Large-scale Structural Analysis
Yasunori Futamura and *Tetsuya Sakurai*, University of Tsukuba, Japan; Sakujiro Hatazawa, MSC Software Corporation, U.S.; Tei Yamato, Mitsubishi Heavy Industries, Ltd, Japan

10:35-10:55 Substructuring Eigensolver Capabilities in Siemens PLM Software with Applications in Dynamic Analyses
Ben-Shan Liao, Siemens PLM Software, U.S.

11:00-11:20 Nonlinear Eigenvalue Problems: Applications and Challenges from Mechanical Engineering
Karl Meerbergen, KU Leuven, Belgium

Thursday, February 28

MS257

Towards Digital Twins for Industrial Applications - Part I of II

9:45 a.m.-11:25 a.m.

Room: 202C

For Part 2 see MS291

The future development of industry and society exhibits strongly increasing complexity, at the same time ever-shorter innovation cycles. In addition, digitisation and the internet of things have led to an explosion of data and information. Without novel computational tools and paradigms we will not be able to manage these challenges. For almost all domains of science and engineering and in most industrial sectors, a multitude of commercial and open source software for modelling, simulation and optimization (MSO) based on mathematical models is available. At the same time increasingly large amounts of process and product data are available and strong artificial intelligence solutions have been developed to exploit these. All this is fostered by computers becoming more and more powerful. These developments lead to the vision that holistic approaches can be achieved that combine all these developments. A complete industrial product or process in its whole life cycle can be accompanied by a virtual representation, often called digital twin that allows design optimization, process control, lifecycle management, predictive maintenance, risk analysis and many other features. Digital twins are so important to business today, that they were named one of Gartner's Top 10 Strategic Technology Trends for 2017. They are becoming a business imperative, covering the entire lifecycle of an asset or process and forming the foundation for connected products and services.

Organizer: Wil Schilders

Eindhoven University of Technology, Netherlands

9:45-10:05 Towards Digital Twins: Merging Data and Model-based Approaches

Dirk Hartmann, Siemens AG, Germany

10:10-10:30 PDE-informed Surrogate Modeling for Fluid-mechanics Applications

Yous van Halder, Centrum voor Wiskunde en Informatica (CWI), Netherlands

10:35-10:55 The GANESO Software (Gas Network Simulation and Optimization)

Julio González Diaz and Alfredo Bermúdez de Castro, University of Santiago de Compostela, Spain

11:00-11:20 The European Network Eu-maths-in and its Efforts for Digital Twins

Volker Mehrmann, Technische Universität Berlin, Germany; Wil Schilders, Eindhoven University of Technology, Netherlands

continued in next column

Thursday, February 28

MS258

Showcase of Research Supported by the DOE Computational Science Graduate Fellowship - Part I of II

9:45 a.m.-11:25 a.m.

Room: 203

For Part 2 see MS292

Since 1991 the DOE computational science graduate fellowship (CSGF) program has supported the training of exceptional researchers across a range of science and engineering fields. In this minisymposium, CSGF fellows and alumni present their latest research accomplishments. The sessions will cover a wide range of computational science application areas.

Organizer: David I. Ketcheson

King Abdullah University of Science & Technology (KAUST), Saudi Arabia

Organizer: Mayya Tokman

University of California, Merced, U.S.

9:45-10:05 Incorporating Computational Science and Engineering into Traditional STEM Education

Claire Ralph, California Institute of Technology, U.S.

10:10-10:30 OpenKIM: Reliable Interatomic Models for Multiscale Simulations

Ryan S. Elliott, University of Minnesota, U.S.

10:35-10:55 Asynchronous Finite Element Simulation of Coastal Inundation

Maximilian Bremer, University of Texas at Austin, U.S.; Hartmut Kaiser, Louisiana State University, U.S.; Clint Dawson, University of Texas at Austin, U.S.

11:00-11:20 Computing in the Kitchen Sink: Computational Stability of a Circular Hydraulic Jump

David I. Ketcheson, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

Thursday, February 28

MS259

Advances in Scalable Iterative, Multigrid and Direct Sparse Linear Solvers

9:45 a.m.-11:25 a.m.

Room: 205

Iterative, multigrid and direct solver algorithms for large sparse linear systems present different and unique challenges for scaling to very high core counts. Krylov methods rely on matrix-vector products and global reductions. The number of global reductions can be reduced, however, maintaining numerical stability and accuracy is a concern. Communication avoiding Krylov algorithms and pipelined variants have been developed to hide these latencies and extend strong scaling. For direct solvers, the introduction of new fill and load-balancing are important for scaling, in addition to possible matrix partitioning and re-ordering techniques. Multigrid methods, either as a preconditioner or stand-alone solver, rely on coarse-grid correction together with relaxation and some smoothers may not exhibit sufficient fine-grain parallelism. The smoother algorithms work local to a sub-domain and can be a hybrid combination of Jacobi and Gauss-Seidel. Asynchronous relaxation schemes are a more scalable alternative. In addition to coarse-grain MPI parallelism, fine-grain thread based parallelism is also possible but may require different approaches and trade-offs such as higher iteration counts for an overall reduction in the time to solution. We have brought together experts in each of these areas who have pioneered new approaches to improving the parallel scalability of sparse solvers.

Organizer: Katarzyna Swirydowicz

National Renewable Energy Laboratory, U.S.

Organizer: Stephen Thomas

National Renewable Energy Laboratory, U.S.

9:45-10:05 Asynchronous Multigrid Methods

Edmond Chow and Jordi Wolfson-Pou, Georgia Institute of Technology, U.S.

10:10-10:30 Communication-avoiding and Pipelined Krylov Solvers in Trilinos

Mark Hoemmen, Sandia National Laboratories, U.S.; Ichitaro Yamazaki, University of Tennessee, Knoxville, U.S.; Erik G. Boman, Sandia National Laboratories, U.S.; Jack J. Dongarra, University of Tennessee and Oak Ridge National Laboratory, U.S.

10:35-10:55 Recent Developments in the Symbolic Phase of Sparse Direct Methods

Mathias Jacquelin and *Esmond G. Ng*, Lawrence Berkeley National Laboratory, U.S.; Barry Peyton, Dalton State College, U.S.

11:00-11:20 Low Synchronization GMRES Algorithms

Stephen Thomas and Katarzyna Swirydowicz, National Renewable Energy Laboratory, U.S.; Julien Langou, University of Colorado, Denver, U.S.; Shreyas Ananthan, National Renewable Energy Laboratory, U.S.; Yang Ulrike, Lawrence Livermore National Laboratory, U.S.

continued in next column

Thursday, February 28

MS260

Applications of Data Assimilation in Science, Engineering, and Industry - Part I of II

9:45 a.m.-11:25 a.m.

Room: 206A

For Part 2 see MS294
Featured Minisymposium

These days data is generated at unprecedented levels in Science, Engineering, and Industry. It is becoming increasingly important to use this data in a smart way in order to stay competitive. In many cases, this may involve inversion or assimilation of data into a complex model, for example derived from physical laws, or it may involve inversion of data alone, for example to learn a complex model in a context where no first principles model exists. This minisymposium aims to explore recent novel applications of methodology and algorithms for data assimilation and inference in Science, Engineering, and Industry.

Organizer: Amr El-Bakry

ExxonMobil Upstream Research Company, U.S.

Organizer: Kody Law

University of Manchester, United Kingdom

9:45-10:05 Assimilating Data for Adaptive Multi-modal Sensing in Nuclear Security

Aaron B. Luttman, National Security Technologies, LLC, U.S.; Derek Constantino, Nevada National Security Site, U.S.; Christopher Stewart and Bethany Goldblum, University of California, Berkeley, U.S.; Tori Hoff, Embry-Riddle Aeronautical University, U.S.; Margaret Hoeller, Illinois Institute of Technology, U.S.

10:10-10:30 Multilevel Markov Chain Monte Carlo Methods for Seismic Source Inversion

Juan Pablo Madrigal Cianci and Fabio Nobile, École Polytechnique Fédérale de Lausanne, Switzerland; Raul F. Tempone, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

continued on next page

Thursday, February 28

MS260

Applications of Data Assimilation in Science, Engineering, and Industry - Part I of II

continued

10:35-10:55 Using Large-eddy Simulations and Feature-based Data Assimilation to Estimate Environmental Variables of Simplified Cloud and Rain Models

Spencer C. Lunderman, University of Arizona, U.S.

11:00-11:20 Inverting a Diverse Set of Spatially Sparse Diagnostic Measurements to Reconstruct a Complete 3D Picture of the Plasma Equilibrium in Fusion Devices

Mark Cianciosa, Oak Ridge National Laboratory, U.S.; J.D Hanson, Auburn University, U.S.; E.C. Howell, Tech-X Corporation, U.S.; S.K. Seal and S.P. Hirshman, Oak Ridge National Laboratory, U.S.

Thursday, February 28

MS261

Computational Advances and Challenges in Data-enabled Life Sciences

9:45 a.m.-11:25 a.m.

Room: 206B

Life sciences, broadly put, have been at the center of the recent data revolution. In particular, the use of devices connected to the internet including sensors, on-field cameras, and personal medical monitors have become increasingly widespread. The data collected by such devices in a wide variety of domains such as crop sciences and medicine have enabled a new level of fundamental understanding and discoveries. Concurrently, the analysis of large-scale complex datasets has also emerged to be an active area of research, spurring new advances in mathematics and computational sciences. This minisymposium will bring together experts and practitioners from three domains, namely computer scientists, mathematicians, and biologists, from both academia and industry, to discuss the advances, challenges and opportunities that uniquely arise at the intersection of computational and life sciences. In particular, the minisymposium will provide a platform to present an overview of the main research challenges stemming out of high-throughput data collection in life sciences, and certain key algorithmic and mathematical advancements including in the areas of topological data analysis, graph and visual analytics, and IoT tools, which are paving the way for tackling those research challenges. The overarching goal is to get the audience excited about the outstanding challenges in this domain and engage in transdisciplinary efforts toward scientific discovery and advancement.

Organizer: Bala Krishnamoorthy
Washington State University, U.S.

Organizer: Ananth Kalyanaraman
Washington State University, U.S.

9:45-10:05 Computational Advances and Challenges in Data-enabled Life Sciences: An Overview

Ananth Kalyanaraman and Bala Krishnamoorthy, Washington State University, U.S.; Patrick Schnable, Iowa State University, U.S.; Eric Lofgren, Washington State University, U.S.

10:10-10:30 Computational Challenges and Opportunities in Crop Phenotyping

Patrick Schnable, Iowa State University, U.S.

10:35-10:55 Persistent Homology in Metagenomics

Laxmi Parida, IBM T.J. Watson Research Center, U.S.

11:00-11:20 Empowering Farmers with Affordable Digital Agriculture Solutions

Ranveer Chandra, Microsoft Research, U.S.

continued in next column

Thursday, February 28

MS262

Modelling with Fractional PDEs: Statistical and Numerical Analysis - Part I of II

9:45 a.m.-11:25 a.m.

Room: 206C

For Part 2 see MS296

Many nonlocal phenomena can be effectively modeled with fractional calculus. Fractional models often require estimation of orders of fractional integral and/or differentiation and also require efficient numerical methods for accurate solutions. In this minisymposium, we present various topics in fractional modeling with PDEs using statistical and numerical tools.

Organizer: Zhongqiang Zhang

Worcester Polytechnic Institute, U.S.

Organizer: Hong Wang

University of South Carolina, U.S.

9:45-10:05 Lagrangian Solution of Fractional Advection-dispersion Equations

Yong Zhang, University of Alabama, U.S

10:10-10:30 Machine Learning of Linear Fractional Differential Equations

Mamikon Gulian and Maziar Raissi, Brown University, U.S.; Paris Perdikaris, University of Pennsylvania, U.S.; George E. Karniadakis, Brown University, U.S.

10:35-10:55 Hybrid Methods for Steady Fractional Advection Diffusion Reaction Equations

Yuchen Dong and Zhongqiang Zhang, Worcester Polytechnic Institute, U.S.

11:00-11:20 The Variable-order Fractional Differential Equation Driven by the White Noise: Mathematical and Numerical Analysis

Xiangcheng Zhen, University of South Carolina, U.S.

Thursday, February 28

MS263

Multi-scale Modeling and Computation of Complex and Active Fluids - Part I of II

9:45 a.m.-11:25 a.m.

Room: 206D

For Part 2 see MS297

Active suspensions, which can consist of micro-swimmers, vesicles, filaments, microtubule-motor networks or other types of interacting particles immersed in liquid, can exhibit complex dynamics and patterns. Capturing the interplays between the particles and the surrounding fluid in continuum or agent-based models presents significant challenges as hydrodynamic interactions happen on short and long scales. Added difficulties arise when such suspensions interact with boundaries of non-trivial shape. This minisymposium's goal is to explore recent advances, especially on the computational front, on this emerging and expanding topic.

Organizer: Enkeleida Lushi

New Jersey Institute of Technology, U.S.

9:45-10:05 Torque-dipolar Micro-swimmers: Modeling, Circling Behavior and Large-scale Simulations

Enkeleida Lushi, New Jersey Institute of Technology, U.S.

10:10-10:30 Rigid Particle Brownian Dynamics without Green's Functions

Brennan Sprinkle, Courant Institute of Mathematical Sciences, New York University, U.S.; Amneet P.S. Bhalla, San Diego State University, U.S.; Aleksandar Donev, Courant Institute of Mathematical Sciences, New York University, U.S.; Neelesh Patankar, Northwestern University, U.S.

10:35-10:55 Simulating Suspensions of Passive and Active Filaments

Adam Townsend, Simon Scholler, Timothy Westwood, and Eric Keaveny, Imperial College London, United Kingdom

11:00-11:20 Nonequilibrium Dynamics of Vesicles: Comparison with Experiments

Afsoun Falavarjani and *David Salac*, State University of New York at Buffalo, U.S.

Thursday, February 28

MS264

Exascale Software for High-order Methods - Part I of II

9:45 a.m.-11:25 a.m.

Room: 207

For Part 2 see MS298

This minisymposium focuses on discussing high-order discretization kernels and lightweight portable libraries that are critical for performance and exascale applications. Topics will include efficient next-generation discretization algorithm and software of arbitrarily high order to impact the design of exascale architectures, system and application software for improved portability and performance of the high-order methods.

Organizer: Tzanio Kolev

Lawrence Livermore National Laboratory, U.S.

9:45-10:05 Efficient Discretizations for Exascale Applications

Tzanio Kolev, Lawrence Livermore National Laboratory, U.S.

10:10-10:30 High-performance Implementation of High-order Discontinuous Galerkin Methods

Peter Bastian, Universität Heidelberg, Germany

10:35-10:55 High-performance Matrix-free Finite Element Programming with the DEAL.II Library

Martin Kronbichler, Technische Universität München, Germany

11:00-11:20 From the Variational Formulation to High Performance Kernels in NGsolve

Joachim Schoeberl, Technische Universitaet Wien, Germany

Thursday, February 28

MS265

Dynamics with Inherent Noise: Stochastic Modelling and Simulation - Part I of II

9:45 a.m.-11:25 a.m.

Room: 401A

For Part 2 see MS299

Inherent noise is ubiquitous in complex systems such as physics, chemistry, engineering and system biology. Numerical simulations based on stochastic models provide an important tool to understand the influence of noise and the dynamic properties of these systems beyond equilibrium. Synergy of stochastic modelling and numerical solutions techniques often leads to novel ideas and promote applications of stochastic models and solvers. In this minisymposium, we focus on both stochastic modelling and numerical methods with emphasis on the interaction of the-state-of-art computational techniques with applications in modelling dynamic process of complex systems. We invite speakers from both communities and expect them to have fruitful discussion. The speakers will address stochastic modelling problems and numerical techniques to solve stochastic equations arising in various applications. Specific topics includes stochastic dynamics modelled by Markov processes with applications to biology and chemical reaction systems, numerical techniques such as singular perturbation methods, surrogate model methods, long time integration of nonlinear SDE, model reduction methods, etc..

Organizer: Huan Lei

Pacific Northwest National Laboratory, U.S.

Organizer: Zhongqiang Zhang

Worcester Polytechnic Institute, U.S.

9:45-10:05 Minimum Action Method for Systems with Constant Time Delays

Xiaoliang Wan, Louisiana State University, U.S.

10:10-10:30 Efficient Numerical Methods for the Generalized Langevin Equation

Matthias Sachs, Duke University, U.S.; Benedict Leimkuhler, University of Edinburgh, United Kingdom

10:35-10:55 High-order Methods for Second-order Stochastic Differential Equations with Rough Noises

Wanrong Cao, Southeast University, China; Zhongqiang Zhang, Worcester Polytechnic Institute, U.S.

11:00-11:20 Finite Element Methods for the Stochastic Cahn-Hilliard Equation with Gradient-type Noise

Yi Zhang, University of North Carolina at Greensboro, U.S

continued in next column

Thursday, February 28

MS266

Inverse Problems in Medical Imaging - Part I of II

9:45 a.m.-11:25 a.m.

Room: 401B

For Part 2 see MS300

Image formation is a nonlinear inverse problem in many emerging medical imaging modalities. Examples include electrical impedance tomography, magnetic particle imaging, photoacoustic tomography and regularized sparse-data X-ray tomography. Typically some part of the imaging process is ill-posed, or extremely sensitive to modelling errors and measurement noise. Regularized nonlinear mathematical reconstruction algorithms are needed for robust and reliable imaging. This session presents results where new nonlinear mathematical techniques are combined with machine learning to yield results with unprecedented image quality.

Organizer: Melody Alsaker

Gonzaga University, U.S.

Organizer: Tatiana A. Bubba

University of Helsinki, Finland

Organizer: Samuli Siltanen

University of Helsinki, Finland

9:45-10:05 Advancements in Spatial Resolution of D-Bar EIT Reconstructions for Human Thoracic Imaging

Melody Alsaker, Gonzaga University, U.S.

10:10-10:30 Dynamic Inverse Problems in Medical Imaging

Christina Brandt, Universitat Hamburg, Germany

10:35-10:55 Limited Angle Tomography: Inpanting in Phase Space by Deep Learning

Tatiana A. Bubba, University of Helsinki, Finland

11:00-11:20 Learned Model-based Reconstructions for Accelerated Limited-view Photoacoustic Tomography

Andreas Hauptmann, University College London, United Kingdom

Thursday, February 28

MS267

Model Reduction and Reduced-order Modeling of Dynamical Systems - Part I of II

9:45 a.m.-11:25 a.m.

Room: 401C

For Part 2 see MS301

Large-scale dynamical systems are used to describe real-world processes in almost every area of computational sciences and engineering, ranging from structural dynamics to nano-electronics to fluid dynamics. It is long since known that in many-query scenarios, where the same model is to be simulated for varying control functions, source terms, and/or parameters, reduced-order surrogates can be used to significantly accelerate the engineering design, control and optimization processes for the underlying problem. In this minisymposium, we will investigate two facets of generating reduced-order surrogate models: first, using "classical" model order reduction in order to obtain a compact representation of a large-scale model arising, e.g., from a finite element discretization of an unsteady PDE, and second, computing a low-dimensional surrogate from available (experimental or simulation) data. Ideally, both approaches are used complementary, and we hope to bring together researchers from both sides in order to exchange ideas for potential hybrid methods using the best from both worlds!

Organizer: Peter Benner

Max Planck Institute for Dynamics of Complex Technical Systems, Germany

Organizer: Athanasios C. Antoulas

Rice University, U.S.

9:45-10:05 System-theoretic Model Order Reduction for Classes of Nonlinear Systems

Peter Benner, Max Planck Institute for Dynamics of Complex Technical Systems, Germany

10:10-10:30 H-Infinity Balanced Truncation for Feedback Control of Flow Problems

Peter Benner, Jan Heiland, and *Steffen W. R. Werner*, Max Planck Institute for Dynamics of Complex Technical Systems, Germany

10:35-10:55 A New Perspective on ADI-based Balanced Trunction

Heike Fassbender and Christian Bertram, Technische Universitaet Braunschweig, Germany

11:00-11:20 Interpolation-based $\mathcal{H}_2$ Optimal Approximation for Quadratic-bilinear Systems

Xingang Cao, Joseph Maubach, Siep Weiland, and Wil Schilders, Eindhoven University of Technology, Netherlands

continued in next column

Thursday, February 28

MS268

Novel Approaches for Design and Analysis of Chaotic Dynamical Systems - Part I of II

9:45 a.m.-11:25 a.m.

Room: 402A

For Part 2 see MS302

As access to high performance computational resources increases, the demand for high-fidelity numerical simulations will only increase. Many of these high-fidelity numerical simulations exhibit chaotic behavior, including climate models, scale-resolving computational fluid dynamics simulations, combustion models, and even road traffic models. The high sensitivity to small perturbations and seemingly random behavior of chaotic dynamical systems make conventional design and analysis methods ill-conditioned or computationally intractable, especially sensitivity analysis. This minisymposium will discuss novel approaches that will allow mathematicians and engineers to analyze, design, and control chaotic dynamical systems. Two classes of sensitivity analysis approaches for mean flow quantities will be discussed. The first class of methods control the exponential growth of unstable linearized perturbations and includes least squares shadowing and space-split sensitivity. The second class of methods involves linearizing about the mean flow and modeling the linearization of the closure terms. Approaches for predicting behavior that deviates from the mean, intermittent and extreme events, will be also discussed. Computationally efficient implementations of all of these algorithms and further strategies to minimize their costs will be discussed in the talks as well as how they can enable gradient-based optimization, error estimation, and mesh adaptation for chaotic dynamical systems.

continued on next page

Thursday, February 28

MS268

Novel Approaches for Design and Analysis of Chaotic Dynamical Systems - Part I of II

continued

Organizer: Patrick J. Blonigan
Sandia National Laboratories, U.S.

Organizer: Qiqi Wang
Massachusetts Institute of Technology, U.S.

Organizer: Nisha Chandramoorthy
Massachusetts Institute of Technology, U.S.

9:45-10:05 Adjoint-based Sensitivity Analysis and Error Estimation for Scale-resolving Turbulent Flow Simulations
Patrick J. Blonigan, Sandia National Laboratories, U.S.

10:10-10:30 Model Reduction and Error Control in Chaotic Flow Simulations
Yukiko S. Shimizu and Krzysztof Fidkowski, University of Michigan, U.S.

10:35-10:55 Grid-adaptation for Chaotic Multi-scale Simulations As a Verification-driven Inverse Problem
Johan Larsson and Siavash Toosi, University of Maryland, U.S.

11:00-11:20 Nilsas: Adjoint Sensitivity Analysis for Chaos via Computing Adjoint Shadowing Directions
Angxiu Ni, University of California, Berkeley, U.S.

Thursday, February 28

MS269

Resilience for Large Scale CSE Applications - Part I of II

9:45 a.m.-11:25 a.m.

Room: 402B

For Part 2 see MS303

Resilience is an imminent issue for large scale computing systems and their applications due to the reliability concerns that stem from the architectural complexity to meet the power and performance requirements together. Under such unreliable computing systems, it is essential to introduce failure mitigations at the runtime and application layers to compliment the resilience at the hardware/system levels. Today, the major resilience scheme is coordinated checkpoint and restart (CR) that involves global coordination of processes and threads. This global recovery model entails inherent scalability issues to handle a large class of errors and failures, and thus alternative approaches would improve the scalability and reliability of computing systems and applications together. In this minisymposium, we will discuss the reliability issues of the large scale computing systems, application/algorithm based resilience and programming model support to enhance the reliability of application program executions.

Organizer: Keita Teranishi
Sandia National Laboratories, U.S.

Organizer: Luc Giraud
Inria, France

Organizer: Christian Engelmann
Oak Ridge National Laboratory, U.S.

Organizer: George Bosilca
University of Tennessee, Knoxville, U.S.

9:45-10:05 On Soft Errors in the Conjugate Gradient: Sensitivity and Robust Numerical Detection
Emmanuel Agullo, Inria, France; Siegfried Cools, University of Antwerp, Belgium; Luc Giraud, Inria, France; Wim Vanroose, Antwerp University, Belgium; *Emrullah Fatih Yetkin*, Kadir Has University, Turkey

10:10-10:30 Correcting Uncorrectable Data Faults Using Mathematics and Application Context
Laura Monroe and Nathan A. DeBardeleben, Los Alamos National Laboratory, U.S.

10:35-10:55 Physics-based Checksums for Silent-error Detection in PDE Solvers
Maher Salloum, *Jackson Mayo*, and Rob Armstrong, Sandia National Laboratories, U.S.

11:00-11:20 Exascale Resilience Strategies for Transient Solvers
Chris Cantwell, Imperial College London, United Kingdom; Allan Nielsen, École Polytechnique Fédérale de Lausanne, Switzerland; David Moxey, University of Exeter, United Kingdom; Jan S. Hesthaven, École Polytechnique Fédérale de Lausanne, Switzerland; Spencer Sherwin, Imperial College London, United Kingdom

continued in next column

Thursday, February 28

MS270

State-of-the-art high-order Numerical Methods and Complex Fluid Simulations - Part I of II

9:45 a.m.-11:25 a.m.

Room: 402C

For Part 2 see MS304

Simulating compressible flows is challenging, due to the presence of discontinuities, e.g. shock and contact waves, and broad-band continuous flow scales. High-order and high-resolution schemes dedicate to resolving smooth flow scales accurately and to capturing the discontinuities simultaneously. In the past decades, many mathematical concepts, such as the artificial viscosity scheme, the total variation diminishing (TVD) scheme, the essentially non-oscillatory (ENO) scheme and the weighted essentially non-oscillatory (WENO) scheme, have been proposed and widely used in both academic and engineering simulations, as well as the compact scheme, flux reconstruction (FR) scheme, spectral volume/difference (SV/D) scheme etc. More recently, the high-order numerical methods pay a more and more important role in high-fidelity simulations, e.g. the large eddy simulations (LES) and the direct numerical simulations (DNS). This session focuses on the state-of-the-arts of high-order numerical methods and the complex fluid simulations, including but not limited to turbulence, multiphase flows, plasma, MHD flows, combustion etc. We aim to bridging the cutting-edge applied mathematics and the request of flow physics simulations.

Organizer: Lin Fu

Technische Universitaet Muenchen, Germany

Organizer: Nikolaus Adams

Technische Universitaet Muenchen, Germany

9:45-10:05 Reviewing of High-order Teno Schemes for Hyperbolic Conservation Laws

Lin Fu, Technische Universitaet Muenchen, Germany

10:10-10:30 High Resolution Schemes for Shock and Interface Capturing in Turbulent Flows

Sanjiva Lele, Akshay Subramaniam, and Man Long Wong, Stanford University, U.S.

10:35-10:55 Compact High-order Gas-kinetic Schemes for the Compressible Navier-Stokes Equations

Kun Xu, Hong Kong University of Science and Technology, Hong Kong

11:00-11:20 A Multi-species, Multi-fluid Model for Simulating Plasma Interpenetration

Debojyoti Ghosh, Thomas Chapman, Richard Berger, and Andris Dimits, Lawrence Livermore National Laboratory, U.S.; Jeff Banks, Rensselaer Polytechnic Institute, U.S.

continued in next column

Thursday, February 28

CP14

Time Integration Methods

9:45 a.m.-10:45 a.m.

Room: 201A

Chair: Hannah Rittich, Forschungszentrum Jülich, Germany

9:45-10:00 Padè Time Stepping Method of Rational Form for PDEs

Said Algarni, King Fahd University of Petroleum and Minerals, Saudi Arabia

10:05-10:20 Efficient Multistage Time Integrators for Incompressible Flows using Projection Methods

Mokbel Karam, University of Utah, U.S.; Mike Hansen, Sandia National Laboratories, U.S.; James C. Sutherland and Tony Saad, University of Utah, U.S.

10:25-10:40 Automatic Construction of Scalable Time-stepping Methods for Stiff PDEs

Vivian A. Montiforte and James V. Lambers, University of Southern Mississippi, U.S.

Thursday, February 28

CP15

Applications in CSE

9:45 a.m.-11:25 a.m.

Room: 201B

Chair: Charles Puelz, Courant Institute of Mathematical Sciences, New York University, U.S.

9:45-10:00 Hyperelastic Immersed Boundary Finite Element Model of the Human Heart

Charles Puelz, Courant Institute of Mathematical Sciences, New York University, U.S.; Margaret Anne Smith, Simone Rossi, and Boyce E. Griffith, University of North Carolina at Chapel Hill, U.S.

10:05-10:20 Computational Modeling for Blood Flow Through Stenotic Artery

Abdulsalam G. Ya'u, Abubakar Tafawa Balewa University Bauchi, Nigeria

10:25-10:40 Reducing Biogeochemical Models using Computational Singular Perturbation

Emily F. Klee and Kyle E. Niemeyer, Oregon State University, U.S.

10:45-11:00 Two-dimensional Local Fourier Image Reconstruction via Domain Decomposition Fourier Continuation Method

Ruonan Shi, Jae Hun Jung, and Ferdinand Schweser, State University of New York at Buffalo, U.S.

11:05-11:20 A Comparison of Some Recently Developed Langevin Integrators on a Course-grained Polymer Melt

Joshua Finkelstein, Temple University, U.S.

Thursday, February 28

SP4

James H. Wilkinson Prize for Numerical Software Prize Lecture: Solving the Two Language Problem in Scientific Computing and Machine Learning with Julia

11:30 a.m.-12:00 p.m.

Room: Ballroom 100BC

Chair: To Be Determined

Bridging cultures that have often been distant, Julia combines expertise from the diverse fields of computer science and computational science to create a new approach to numerical computing. Julia is designed to be easy and fast and questions notions generally held to be "laws of nature" by practitioners of numerical computing: 1. High-level dynamic programs have to be slow. 2. One must prototype in one language and then rewrite in another language for speed or deployment. 3. There are parts of a system appropriate for the programmer, and other parts that are best left untouched as they have been built by the experts. We introduce the Julia programming language and its design—a dance between specialization and abstraction. Specialization allows for custom treatment. Multiple dispatch, a technique from computer science, picks the right algorithm for the right circumstance. Abstraction, which is what good computation is really about, recognizes what remains the same after differences are stripped away. Abstractions in mathematics are captured as code through another technique from computer science, generic programming. Julia has been downloaded over 3,000,000 times. The Julia community has contributed almost 2,500 Julia packages, while also building connectivity to C, C++, Fortran, MPI, Python, R, Java and others. Julia shows that one can achieve machine performance without sacrificing human convenience.

Jeffrey Bezanson

Julia Computing, Inc.,, U.S.

Stefan Karpinski

Julia Computing, Inc.,, U.S.

Viral Shah

Julia Computing, Inc.,, U.S.

Lunch Break

12:00 p.m.-1:00 p.m.

Attendees on their own

continued in next column

Thursday, February 28

IP7

Physical, Numerical, and Computational Challenges in Modelling Oceans for Climate

1:00 p.m.-1:45 p.m.

Room: Ballroom 100BC

Chair: Jeffrey A. Hittinger, Lawrence Livermore National Laboratory, U.S.

The world's oceans play a critical role in determining the evolution of Earth's climate. While changes in surface climate and weather extremes are usually presented in the context of the atmosphere and "weather", the long time scales in the climate systems belong to the ice caps, land, and oceans. In the ocean, short spatial- and fast time-scale processes, such as wave breaking and turbulence, ultimately control the planetary scale circulation which evolves and equilibrates over millennia. These same processes make the oceans an important part of the climate system at long time-scales, but computationally challenging to simulate. Until recently most ocean climate models were non-eddy resolving; the dominant modes of variability, mesoscale eddies - analogous to weather systems, were not resolved but parameterized with lower order diffusive closures. Models from this prior era used numerical methods and algorithms that emphasized efficiency over accuracy. Contemporary computing resources allow the use of eddy-permitting ocean models for climate, along with short-duration eddy-resolving calculations. A 2-D turbulent eddy cascade introduces new constraints on ocean models. Modern computing architectures present new challenges for simulating the long time-scales and near adiabatic nature of the ocean interior. We review innovations in numerical and computational methods used in ocean modelling, and examine some of the unique challenges posed by the physics of the ocean.

Alistair Adcroft

Princeton University and NOAA-GFDL, U.S.

Coffee Break

1:45 p.m.-2:15 p.m.

Room: Ballroom Foyer

Thursday, February 28

MS271

Theory and Application of Surrogate Models for Bayesian Inverse Problems - Part II of II

2:15 p.m.-3:55 p.m.

Room: Ballroom 100BC

For Part 1 see MS238

Today the availability of exascale computing resources enables high-fidelity simulations of complex processes in science and engineering. Moreover, it is becoming standard to include and quantify uncertainties in these processes. However, the forward and inverse uncertainty quantification with computationally demanding models is still infeasible in many situations. Expensive models arise for instance in hydrology, oncology, and multiphysics simulations such as fluid structure interaction. In practice, such models can be replaced by computationally cheap surrogates based on full model runs. Surrogates enable and accelerate uncertainty quantification. This minisymposium is dedicated to surrogate modelling for inverse uncertainty quantification. We highlight recent theoretical advances from mathematical and computational sciences, for example, adaptive strategies for surrogate construction, surrogate error analysis, and multifidelity surrogates, together with results from practical estimation problems. Surrogate models include Gaussian processes, sparse grids, reduced bases, (piecewise) polynomials, and low-rank tensors.

Organizer: Jonas Latz

Technische Universität München, Germany

Organizer: Claudia Schillings

Universitaet Mannheim, Germany

Organizer: Elisabeth Ullmann

Technische Universität München, Germany

2:15-2:35 A New Approach for Bayesian Inversion using Sequential Piecewise Polynomial Chaos Expansion

Paul-Remo Wagner, Stefano Marelli, Christos Lataniotis, and Bruno Sudret, ETH Zürich, Switzerland

continued in next column

2:40-3:00 A Bayesian Framework for Assessing the Strength Distribution of Composite Structures with Random Defects

Anne Reinarz, Technische Universität München, Germany; Robert Scheichl, Universität Heidelberg, Germany; Tim J. Dodwell, University of Exeter, United Kingdom; Linus Seelinger, Universität Heidelberg, Germany

3:05-3:25 Uncertainty Quantification of Global Climate Models Using MCMC with Gaussian Process Emulators

Emmet Cleary, Tapio Schneider, and Andrew Stuart, California Institute of Technology, U.S.

3:30-3:50 Accelerating Physics-constrained Bayesian Inverse Problems Using Inaccurate Models and Data-driven Learning

Maximilian Koschade and Phaedon-Stelios Koutsourelakis, Technische Universität München, Germany

Thursday, February 28

MS272

Multiphysics: Extensible, Composable Algorithms and Software - Part II of II

2:15 p.m.-3:55 p.m.

Room: Conference Theater

For Part 1 see MS239
Featured Minisymposium

A major thrust for more than two decades has been the development of flexible and extensible methodologies for tackling the most challenging scientific and engineering problems. These problems naturally exhibit multiscale and multiphysics characteristics which strain algorithmic foundations and software capabilities, assumptions, and optimization while targeting contemporary and emerging hardware. In this environment, it is not possible for a single person to understand and customize components for each class of target problem, thus algorithmic and software robustness, composability, and extensibility is paramount. Training, automation, and quality control become increasingly challenging as users and contributors from diverse backgrounds attempt to conduct increasingly complex simulations and analyses. This session brings together exemplars, promising techniques, and lessons learned in developing algorithms, software, and communities for computational multiphysics.

Organizer: Paul Bauman

State University of New York at Buffalo, U.S.

2:15-2:35 Efficient Solver Composition with High-order Methods

Valeria Barra, Jed Brown, and Jeremy Thompson, University of Colorado Boulder, U.S.

2:40-3:00 Albany: A Trilinos-based Multi-physics Partial Differential Equation Research Tool Created using the Agile Components Code Development Strategy

Irina K. Tezaur, Andrew Salinger, Glen Hansen, and Daniel Ibanez, Sandia National Laboratories, U.S.

3:05-3:25 Collaborative Multiphysics, Multiscale Simulation using Moose

Derek R. Gaston, Cody J. Permann, David Andrs, John Peterson, Andrew Slaughter, Fande Kong, Robert Carlsen, Alex Lindsay, Richard Martineau, Brian Alger, and Jason Miller, Idaho National Laboratory, U.S.

3:30-3:50 Coupled Finite Element/ Boundary Element Methods for Fast High Intensity Focused Ultrasound Treatment Modelling

Timo Betcke, University College London, United Kingdom; Sam Groth and *Garth Wells*, University of Cambridge, United Kingdom

continued in next column

Thursday, February 28

MS273

Boundary Integral Methods for Particulate Flows - Part II of II

2:15 p.m.-3:30 p.m.

Room: 102A

For Part 1 see MS240

This minisymposium will focus on recent advances in the development of integral equation methods for simulating low-Re viscous flow of bubbles, drops, vesicles, rigid particles and other particulate suspensions. The accurate simulation of such complex fluid flows is challenging because of strong nonlinearities, time-dependent geometries, close interactions, and multiple length and time scales. Integral equation methods offer a suitable and powerful framework able to address these challenges while providing high-order accuracy.

Organizer: Lukas Bystricky

Florida State University, U.S.

Organizer: Chiara Sorgentone

KTH Royal Institute of Technology, Sweden

2:15-2:35 Contact-free Simulations of Rigid Particle Suspensions

Lukas Bystricky and Bryan D. Quaife, Florida State University, U.S.

2:40-3:00 Accurate Simulation of Surfactant-covered Droplets in Wall-confined Two-dimensional Stokes Flow

Sara Pålsson and Anna-Karin Tornberg, KTH Royal Institute of Technology, Sweden

3:05-3:25 Boundary Integral Computation via the Polynomial Atlas Method

William H. Mitchell, Macalester College, U.S.

Thursday, February 28

MS274

Industrial Eigensolution Technology: Advances and Challenges - Part II of II

2:15 p.m.-3:55 p.m.

Room: 102B

For Part 1 see MS256

Numerical eigenvalue problems and eigensolution technology are one of the centerpiece of industrial scale modeling and simulations for acoustics and structural analysis among many others. Speakers from industry and academic will present latest advancements in algorithms, numerical analysis and industrial strength high-performance software, and discuss emerging needs and challenges.

Organizer: Zhaojun Bai

University of California, Davis, U.S.

Organizer: Roger Grimes

Livermore Software Technology Corporation, U.S.

2:15-2:35 The Accelerated ACMS (Automated Component Modal Synthesis) Method: Extracting Parallel Efficiency through Shared Memory Processing

David Ehrlich and Peter Schartz, MSC Software Corporation, U.S.

2:40-3:00 Lanczos Method for Buckling Eigenvalue Problems of Singular Pencils

Chao-Ping Lin, University of California, Davis, U.S.; Huiqing Xie, East China University of Science and Technology, China; Zhaojun Bai, University of California, Davis, U.S.

3:05-3:25 Accelerating Commercial FEA Software through Advanced Computational Technologies

Vladimir Belsky, Mintae Kim, and Mikhail Belyi, Dassault Systèmes, U.S.

3:30-3:50 The State-of-the-art of Solvers for Quadratic Eigenvalue Problems

Zhaojun Bai, University of California, Davis, U.S.; Ding Lu, University of Geneva, Switzerland; Tianyi Lu and Yangfeng Su, Fudan University, China

Thursday, February 28

MS275

Recent Advances in Models and Numerical Methods for Multiphase Problems - Part II of II

2:15 p.m.-3:55 p.m.

Room: 102C

For Part 1 see MS242

Multiphase problems arise in a wide range of computational science and engineering applications. They are ubiquitous in the universe. They exist in material systems consisting of coexistent, multiple physical phases (gas, liquid, solid) and/or a single physical phase but multiple components (i.e., oil/water, polymer/polymer/solvent). How to develop the thermodynamically consistent description of the multiphase phenomena? How to develop appropriate numerical strategies and computational methods to efficiently solve the multiphase problems are pressing issues in the research community. Phase field has emerged as an effective modeling and computational tool to deal with multiple phase phenomena recently. How to exploit the efficiency of the formulation and computational flexibility and how to couple this efficient computational strategy to other methods for multiphase problems will be the central theme of the symposium. In this minisymposium, experts in modeling and computation are invited to share their recent advances in studying multiphase problems and relevant applications. It serves as a platform for advancing modeling and numerical strategies in the related field.

Organizer: Jia Zhao

Utah State University, U.S.

Organizer: Qi Wang

University of South Carolina, U.S. and Beijing Computational Science Research Center, China

2:15-2:35 The SAV Approach for Gradient Flows and Positivity Preserving Schemes for Poisson-Nernst-Planck Equations

Jie Shen, Purdue University, U.S.

2:40-3:00 Maximum Principle Preserving Exponential Time Differencing Schemes for the Nonlocal Allen-Cahn Equation

Qiang Du, Columbia University, U.S.; *Lili Ju*, University of South Carolina, U.S.; Xiao Li, Beijing Computational Science Research Center, China; Zhonghua Qiao, Hong Kong Polytechnic University, China

3:05-3:25 A Highly-accurate Numerical Method for the Solution of the Two-dimensional Allen-Cahn Equation in Non-convex Polygons

Ping Lin, University of Dundee, United Kingdom

3:30-3:50 A Time Adaptive Rescaling Method for Computing Interface Problems

Shuwang Li, Illinois Institute of Technology, U.S.

continued in next column

Thursday, February 28

MS276

Reduced Order Models for Fluids: Achievements and Open Problems - Part II of II

2:15 p.m.-3:55 p.m.

Room: 102D

For Part 1 see MS243

This minisymposium aims at giving a survey of recent developments in the reduced order modeling of fluid flows. Computational modeling, numerical analysis, and applications to realistic engineering and geophysical flow problems will be covered in this minisymposium. Both achievements and open problems in the reduced order modeling of fluid flows will be discussed.

Organizer: Traian Iliescu

Virginia Tech, U.S.

Organizer: Michael Schneier

University of Pittsburgh, U.S.

2:15-2:35 A Pod Based Artificial Compression Scheme for the Incompressible Navier-Stokes Equations

Michael Schneier, University of Pittsburgh, U.S.; Traian Iliescu, Virginia Tech, U.S.; William Layton, Victor P. Decaria, and Michael McLaughlin, University of Pittsburgh, U.S.

2:40-3:00 Incremental POD Mode Algorithm for Fluids with Easily Computable Error Bounds

John Singler, Missouri University of Science and Technology, U.S.

3:05-3:25 Online Adaptive Basis Refinement and Compression for Reduced-order Models

Philip Etter, Stanford University, U.S.; Kevin T. Carlberg, Sandia National Laboratories, U.S.

3:30-3:50 A Variational Approach to Closure of Nonlinear Dynamical Systems Based on Small-scale Parameterizations

Honghu Liu, Virginia Tech, U.S.; Mickael Chekroun and James McWilliams, University of California, Los Angeles, U.S.; Shouhong Wang, Indiana University, U.S.

Thursday, February 28

MS277

Theory and Algorithms for Improved Performance of Machine Learning in Scientific Applications - Part II of II

2:15 p.m.-3:30 p.m.

Room: 111A

For Part 1 see MS244

Machine learning (ML) and artificial intelligence (AI) have had unprecedented success in Silicon Valley, and for big data problems that have emerged in the last decade. However, the application of ML to scientific discovery requires an entirely new focus on algorithmic development and mathematical rigor, aimed at producing reliable, confident, and reproducible results, customized to specific type (plentiful or sparse) and form (labeled or unlabeled, static or online, deterministic or noisy) of observational data. An overarching goal of this symposium is to amalgamate international experts in numerical analysis, approximation theory, uncertainty quantification, and data science, to systematically investigate: the design of architectures/ models which accurately capture the complexities of the data, with robust estimates of confidence in predictions on defined domains; and fast and scalable algorithms to fit the proposed models to data, with a theory that explains the convergence and success of these techniques.

Organizer: Guannan Zhang

Oak Ridge National Laboratory, U.S.

Organizer: Clayton G. Webster

University of Tennessee and Oak Ridge National Laboratory, U.S.

2:15-2:35 Title Not Available

Rich Archibald, Oak Ridge National Laboratory, U.S.

2:40-3:00 Geometric Insights into Spectral Clustering by Graph Laplacian Embeddings

Franca Hoffmann, California Institute of Technology, U.S.; Nicolas Garcia Trillos, Brown University, U.S.; Bamdad Hosseini, California Institute of Technology, U.S.

3:05-3:25 Learning Reduced Order Model for Fluid Dynamics

Xuping Xie and *Guannan Zhang*, Oak Ridge National Laboratory, U.S.; Clayton G. Webster, University of Tennessee and Oak Ridge National Laboratory, U.S.

continued in next column

Thursday, February 28

MS278

Models with Superior Reasons for Complex Big Data - Part II of II

2:15 p.m.-3:55 p.m.

Room: 111B

For Part 1 see MS245

This minisymposium concerns state-of-the-art theoretical or computational methods for solving real-world problems involving complex high-dimensional big-volume data stemming from heterogenous sources, such as high-frequency data from finance, or spatio-temporal sensory data from mobile connected devices, semi- or fully autonomous transportation and emerging Internet-of-Things (IoT) applications. We are particularly interested in building highly accurate learning and prediction models with excellent decision reasoning: overall variable importance at the model level down to personalized reasons at the data record level, which demonstrates superior model intelligence and hence inspires confidence and trust. Enabling software tools and hardware architecture will be highlighted and emphasized along with key development principles. This minisymposium is partially sponsored by American Family Insurance and eMALI. IO Ltd.

Organizer: Sou-Cheng T. Choi

Illinois Institute of Technology, U.S. and Allstate Insurance Corporation, U.S.

Organizer: Luisa Polania

Target Corporation, U.S.

Organizer: Teja Kanchinadam

American Family Insurance, U.S.

Organizer: Lawrence K. H. Ma

Hong Kong Blockchain Society, Hong Kong

2:15-2:35 Non-IID Model Learning via Graph Neural Networks

Glenn Fung, American Family Insurance, U.S.

2:40-3:00 Understanding Climate-vegetation Interactions in Global Rainforests through a GP Model-tree Analysis

Anuradha Kodali, Allstate Corporation, U.S.; Kamalika Das, University of Maryland, Baltimore County, U.S.

3:05-3:25 Covariant Neural Networks for Learning Structured Data

Risi Kondor and *Brandon Anderson*, University of Chicago, U.S.

3:30-3:50 Contextually-sensitive Named Entity Disambiguation

Jeonghee Yi, Kyle Yan, Jingfei Du, Javad Dousti, and Bi Xue, Facebook, Inc., U.S.

continued in next column

Thursday, February 28

MS279

Machine Learning Methods in Computational Fluid Dynamics - Part II of II

2:15 p.m.-3:55 p.m.

Room: 111C

For Part 1 see MS246

Machine learning (ML) methods and in particular deep learning via neural networks have generated significant interest in the last years. This interest reaches beyond the ML community itself into other fields of science and engineering. Since these methods can provide approximations to general functions by being coaxed to learn optimally from data without a-priori assumptions, they are particularly attractive for the generation of subspace models. In the area of computational fluid dynamics, research into how ML methods can enhance current capabilities is an active research topic. Deep learning methods have been shown to provide accurate shock-capturing sensors, improve RANS turbulence models, provide approximate deconvolutions of a coarse-scale flow fields and reconstruct the exact LES stresses terms to generate data-driven closure models. With these encouraging successes, a number of questions arise: What did the models actually learn? How can one incorporate physical constraints into the learning process? What is the confidence in the model predictions? And what are the limits of ML and can we quantify them? This symposium invites contributions that discuss how to enhance CFD methods by machine learning algorithms. The generation of subspace models, e.g. for RANS and LES, is of particular interest, but the minisymposium is also open to other topics, e.g. learning intelligent flow sensors, data compression and feature extraction, implementation examples etc.

continued on next page

Thursday, February 28

MS279

Machine Learning Methods in Computational Fluid Dynamics - Part II of II

continued

Organizer: Andrea D. Beck
Universität Stuttgart, Germany

Organizer: Romit Maulik
Oklahoma State University, U.S.

Organizer: Omer San
Oklahoma State University, U.S.

2:15-2:35 Machine Learning Algorithms for Explicit Algebraic RANS Turbulent Closures: A Galilean Invariant Approach
Corrado Sotgiu and Bernhard Weigand, University of Stuttgart, Germany

2:40-3:00 Controlling Oscillations in High-order Accurate Methods Through Artificial Neural Networks
Jan S. Hesthaven, Niccolo Discacciati, and Deep Ray, École Polytechnique Fédérale de Lausanne, Switzerland; Jian Yu, Beihang University, China

3:05-3:25 Data-driven Discovery of Governing Physical Laws and Their Parametric Dependencies
J. Nathan Kutz, Samuel Rudy, and Steven Brunton, University of Washington, U.S.

3:30-3:50 Exploiting Sparsity for Modeling and Control of Dynamical Systems
Eurika Kaiser, Steven Brunton, and J. Nathan Kutz, University of Washington, U.S.; Bernd Noack, Technische Universität Braunschweig, Germany; Andreas Spohn, ENSMA, France

Thursday, February 28

MS280

Recent Advances in PDE-constrained Optimization under Uncertainty- Part II of II

2:15 p.m.-3:55 p.m.

Room: 300A

For Part 1 see MS247

Over the past two decades, significant strides have been made in theory and algorithms for optimization of systems governed by PDEs. These advances in PDE-constrained optimization have enabled many applications of optimal design and control to complex deterministic systems in computational science and engineering. A major remaining challenge for PDE-constrained optimization is accounting for uncertainty in the PDE models. Many PDE models are characterized by uncertain parameters due to lack of knowledge or intrinsic variability of the inputs, including initial or boundary conditions, sources, coefficients, or geometry. It is important to incorporate this uncertainty in the optimization problem to make the optimal solution more robust and reliable. Indeed, recent years have witnessed rapid growth in research on PDE-constrained optimization under uncertainty. In this minisymposium, leading experts will present their work on advances in the development, analysis, and application of methods for PDE-constrained optimization under uncertainty. In particular, the following themes will be emphasized: (1) development and analysis of scalable algorithms to solve optimization problems under high- or infinite-dimensional uncertainty; (2) investigation of different risk measures and constraints; (3) application to problems governed by more challenging models, including multiphysics, multiscale, and fractional PDE problems.

Organizer: Peng Chen
University of Texas at Austin, U.S.

2:15-2:35 Efficient PDE-constrained Optimization under Uncertainty using Adaptive Model Reduction and Sparse Grids
Matthew J. Zahr, University of California, Berkeley and Lawrence Berkeley National Laboratory, U.S.

2:40-3:00 An Adaptive Sample-based Approximation Approach for Stochastic Inverse Problems
Zilong Zou, Duke University, U.S.; Harbir Antil, George Mason University, U.S.; Sayan Mukherjee and Wilkins Aquino, Duke University, U.S.

3:05-3:25 Stochastic Approximation Techniques for PDE Constrained Optimal Control Problem under Uncertainty.
Matthieu Martin and Fabio Nobile, École Polytechnique Fédérale de Lausanne, Switzerland

3:30-3:50 Large-scale Optimization in Deep Learning for PDE Representation
Tom O'Leary-Roseberry, Joshua Chen, and Peng Chen, University of Texas at Austin, U.S.; Umberto Villa, Washington University, St. Louis, U.S.; Omar Ghattas, University of Texas at Austin, U.S.

continued in next column

Thursday, February 28

MS281

High Performance Sparse Matrix, Tensor, and Graph Kernels - Part II of II

2:15 p.m.-3:55 p.m.

Room: 300B

For Part 1 see MS248

High performance computational kernels are critical to various real-world applications, such as scientific computing, machine learning, social networks, and healthcare analytics, to name a few. Data from many applications are sparse, which means most of the entries are meaningless or missing values, represented as zeros. In this proposed minisymposium, we will discuss state-of-the-art studies of sparse computational kernels in sparse matrix, sparse tensor, and graph problems. We will discuss their interactions, challenges, solutions, and future directions. Our topics cover these sparse kernels from different views: sparse kernels as benchmarks to compare different computer architectures, new sparse data structures and optimization methods for these algorithms, and the applications of these kernels. Our topics also cover the optimization techniques from different domains: algorithms, compilers, runtime systems, and computer architectures. We expect this minisymposium will explore high-performance sparse algorithms and their interactions among different kernels.

Organizer: Xu Liu

College of William & Mary, U.S.

Organizer: Jiajia Li

Pacific Northwest National Laboratory, U.S.

2:15-2:35 Communication-avoiding Sparse Matrix Algorithms for Large Graph and Machine Learning Problems

Aydin Buluc, Lawrence Berkeley National Laboratory, U.S.

2:40-3:00 SPARTan: Scalable PARAFAC2 for Large and Sparse Data

Ioakeim Perros, Georgia Institute of Technology, U.S.; Evangelos Papalexakis, University of California, Riverside, U.S.; Fei Wang, Cornell University, U.S.; Rich Vuduc, Georgia Institute of Technology, U.S.; Elizabeth Searles, Children's Healthcare of Atlanta, U.S.; Michael Thompson, Unaffiliated; Jimeng Sun, Georgia Institute of Technology, U.S.

3:05-3:25 S-BLAS: Designing Highly Scalable Sparse-BLAS Kernels for Complex Modern HPC Architectures

Shuaiwen Song, Pacific Northwest National Laboratory, U.S.; *Jieyang Chen*, University of California, Riverside, U.S.

3:30-3:50 Sparse Tensor Algebra and its Relations to Matrix and Graph Problems

Jiajia Li, Pacific Northwest National Laboratory, U.S.; Jimeng Sun and Richard Vuduc, Georgia Institute of Technology, U.S.

continued in next column

Thursday, February 28

MS282

Tensor Based Methods in Scientific Computing and Data Science - Part II of II

2:15 p.m.-3:55 p.m.

Room: 300C

For Part 1 see MS249

Tensor-based methods are becoming increasingly important because of their ability to extract and exploit inherent multidimensional structure. However, several computational challenges persist that prevent the use of tensor decompositions for large-scale problems. This minisymposium explores advances in numerical methods and applications involving tensor decompositions in scientific computing and data sciences.

Organizer: Arvind Saibaba

North Carolina State University, U.S.

Organizer: Misha E. Kilmer

Tufts University, U.S.

2:15-2:35 Parallel Algorithms for Multiple Tensor-times-matrix Computation

Grey Ballard, Wake Forest University, U.S.

2:40-3:00 Randomized Pass-efficient Algorithms for Tucker Decompositions

Rachel Minster and Arvind Saibaba, North Carolina State University, U.S.

3:05-3:25 Tensor-tensor Products for Optimal Representation and Compression of Snapshot Data

Misha E. Kilmer, Tufts University, U.S.; Lior Horesh, IBM Research, U.S.; Haim Avron, Tel Aviv University, Israel; Elizabeth Newman, Tufts University, U.S.

3:30-3:50 Spacey Random Walks and Closely Related Tensor Eigenvector Problems for Data Analysis

David F. Gleich, Purdue University, U.S.; Austin Benson, Cornell University, U.S.

Thursday, February 28

MS283

BE: Advances in Computational Drug Discovery - Part II of II

2:15 p.m.-3:30 p.m.

Room: 302A

For Part 1 see MS250

Drug discovery and development continues to be an expensive, lengthy, and sometimes fatal process. With the large chemical space of potential drugs and the number of proteins and observable genetic variation of these proteins, an experimental and clinical understanding of all potential interactions is not feasible, making accurate computational studies extremely relevant and important. The massive computational power of leadership class computing and advances in machine learning provide exciting opportunities to overcome hurdles in drug discovery and development. Drugs are known to be promiscuous; they interact with more than one protein. Some estimates predict that every protein interacts with tens to hundreds of different known drugs. Accurate predictions of how drugs interact with the entire proteome would establish computational polypharmacological networks. These networks would provide insight on drug repurposing, side-effect prediction, and the development of more efficacious drugs. This session focuses on combining traditional biomolecular simulations with biomedical big data sources using machine learning techniques to improve drug binding predictions to enable the development of accurate polypharmacological networks.

Organizer: Sally R. Ellingson

University of Kentucky, U.S.

2:15-2:35 An Open Access Platform Empowering Research, Education, and the Crowdsourced Discovery of Medicine

Benjamin Samudio, Sierra College, U.S.

2:40-3:00 The Abstraction of 3-D Biomolecular Structure and Property Data for the Training of OpeNNdd Models

Bryce Kroencke, University of California, Davis, U.S.

3:05-3:25 The Development and Validation of OpeNNdd Network Architecture

Nicholas Pavini, American River College, U.S.

continued in next column

Thursday, February 28

MS284

Algorithms and Software for Nonlinear Eigenvalue Problems - Part II of II

2:15 p.m.-3:55 p.m.

Room: 302B

For Part 1 see MS251

Eigenvalue problems arise in many fields of science and engineering and their mathematical properties and numerical solution methods for standard, linear eigenvalue problems are well understood. Recently, due to breakthrough applications, e.g., electronic structure calculations for 2D materials, we observe emerging challenges in the nonlinear eigenvalue problems (NLEVPs) $T(\lambda) x=0$, which exhibits nonlinearity in the eigenvalue parameter. Although, the nonlinear eigenvalue problems received a lot of attention from the numerical linear algebra community during the last decade, the majority of the work has been focused on polynomial eigenvalue problems. The minisymposium will provide a forum for discussing the novel algorithmic contributions and latest software developments for the general nonlinear eigenvalue problems involving nonlinear functions such as exponential, rational, and irrational ones. Presentations will focus on Newton-based techniques, Krylov subspace methods applied to linearizations, and contour integration and rational filtering methods and their role in large-scale science and engineering simulations.

Organizer: Roel Van Beeumen

Lawrence Berkeley National Laboratory, U.S.

Organizer: Agnieszka Miedlar

University of Kansas, U.S.

2:15-2:35 Nonlinear Eigenvalue Localization for Damping Bounds

David S. Bindel, Cornell University, U.S.

continued on next page

2:40-3:00 Reliable Computation of Certain Exterior Eigenvalues of Large Polynomial Eigenproblems
Fei Xue, Clemson University, U.S.; Roel Van Beeumen, Lawrence Berkeley National Laboratory, U.S.

3:05-3:25 Automatic Rational Approximation and Linearization for Nonlinear Eigenvalue Problems
Javier Perez-Alvaro, University of Montana, U.S.

3:30-3:50 The Nonlinear FEAST Algorithm
Agnieszka Miedlar, University of Kansas, U.S.

Thursday, February 28

MS285

Preconditioning for High-order Matrix-free PDE Operators - Part II of II

2:15 p.m.-3:30 p.m.

Room: 303A

For Part 1 see MS252

This minisymposium explores various approaches to preconditioning operators when no explicit matrix representation is easily available. An ongoing move in high-performance computing toward next-generation high-concurrency hardware makes the higher arithmetic intensity of high polynomial order discretizations much more attractive, but explicit matrix assembly and application of such high-order operators rapidly becomes a performance bottleneck. A partial solution to this issue is representing the operator in an unassembled or matrix-free setting, taking advantage for example of tensor representations of a finite element basis. In the unassembled setting, however, it is not at all clear how to precondition the operator, especially in situations where a geometric mesh hierarchy is not available. In this minisymposium, we consider approaches such as preconditioning a high-order method with a lower-order method, unassembled algebraic multigrid preconditioners, and related preconditioning methods.

Organizer: Andrew T. Barker
Lawrence Livermore National Laboratory, U.S.

Organizer: Per-Olof Persson
University of California, Berkeley, U.S.

2:15-2:35 BDDC Preconditioners for Matrix-free Higher-order Methods
Clark R. Dohrmann, Sandia National Laboratories, U.S.; Stefano Zampini, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

2:40-3:00 Matrix-free Algebraic Multigrid Preconditioners
Bruno Turcksin, Wayne Joubert, Andrey Prokopenko, and Damien Lebrun-Grandie, Oak Ridge National Laboratory, U.S.

3:05-3:25 Uniform Preconditioner for the Mass Matrix for Hp-Fem on Triangles
Mark Ainsworth and *Shuai Jiang*, Brown University, U.S.

Thursday, February 28

MS286

Mitigating Communication Costs Using Variable Precision Computing Techniques - Part II of II

2:15 p.m.-3:55 p.m.

Room: 303B

For Part 1 see MS253

Communication costs, including global data movement between processors and local movement within a processor, are becoming significantly more expensive compared to arithmetic operations. It is now of high importance to develop algorithms with communications costs in mind. Traditionally, to guard against floating-point round off errors, we store data in a 64-bit double precision representation even when very few significant digits are needed. However, many scientific simulations have traditional error polluting the 64-bit representation - truncation error, iteration error, and floating-point round off error. Consequently, since many of the bits represent error, we are wasting valuable resources. Thus, a simple way to reduce communication costs is to reduce the size of the data; either using fewer bits, i.e., in a lower precision, or by using a lossy compression algorithm. Mixed-precision or lower precision algorithms have had a renewed interest not only to mitigate communication costs, but due to the fact that GPU's can perform half precision arithmetic approximately 16 times as fast as double precision arithmetic. In the past few years, lossy compression algorithms have been specifically designed for scientific data, sacrificing a small amount precision and additional computations for reduced bandwidth. To this end, this minisymposium will highlight recent algorithmic advances in variable precision computing.

Organizer: Alyson Fox
Lawrence Livermore National Laboratory, U.S.

continued on next page

Thursday, February 28

MS286

Mitigating Communication Costs Using Variable Precision Computing Techniques - Part II of II

continued

Organizer: Jeffrey A. Hittinger
Lawrence Livermore National Laboratory, U.S.

Organizer: James D. Diffenderfer
University of Florida, U.S.

2:15-2:35 Exploiting Half Precision Arithmetic in Solving $Ax = b$
Nicholas J. Higham, Srikara Pranesh, and Mawussi Zounon, University of Manchester, United Kingdom

2:40-3:00 Simulated Half-precision Implementation of Blocked QR Factorization and Graph Clustering Applications
Lucia Yang, University of Colorado Boulder, U.S.; Geoff Sanders, Lawrence Livermore National Laboratory, U.S.

3:05-3:25 Experiments with Mixed Precision Algorithms in Linear Algebra
Jack J. Dongarra, University of Tennessee and Oak Ridge National Laboratory, U.S.

3:30-3:50 Nonlinear Error Transport for Parabolic Problems in Mixed Precision
Dylan M. Copeland and Daniel Osei-Kuffuor, Lawrence Livermore National Laboratory, U.S.

Thursday, February 28

MS287

Advances in Rare Event Simulation for Complex Dynamical Systems - Part I of II

2:15 p.m.-3:55 p.m.

Room: 201B

For Part 2 see MS320

The study of rare events in dynamical systems is essential for a wide range of applications including the robust design of engineering systems, the prediction of extreme weather events, and the study of economic catastrophes. Over the last twenty years, different research communities have developed a variety of tools for studying rare events in dynamical systems. In these sessions, we showcase tools adapted from sampling, dynamical systems, and statistics including subset simulation, importance sampling, particle splitting, information theoretic bounds, and large deviations. These talks will show applications of these tools in fluid dynamics, aerospace systems, chemistry, atmospheric modeling, finance, and electrical systems. We will also explore common themes and challenges across these approaches.

Organizer: Benjamin J. Zhang
Massachusetts Institute of Technology, U.S.

Organizer: Tuhin Sahai
United Technologies Research Center, U.S.

2:15-2:35 A Generalized Theory of Rare Event Simulation for Stochastic Differential Equations
Benjamin J. Zhang, Massachusetts Institute of Technology, U.S.; Tuhin Sahai, United Technologies Research Center, U.S.; Youssef M. Marzouk, Massachusetts Institute of Technology, U.S.

2:40-3:00 Are Extreme Dissipation Events Predictable in Turbulent Fluid Flows?
Patrick J. Blonigan, Sandia National Laboratories, U.S.; *Mohammad Farazmand* and Themistoklis Sapsis, Massachusetts Institute of Technology, U.S.

3:05-3:25 Importance Sampling for Slow-fast Diffusions Based on Moderate Deviations
Konstantinos Spiliopoulos, Boston University, U.S.

3:30-3:50 Simplified Gentlest Ascent Dynamics for Saddle Point with Application to Allen-Cahn in Presence of Shear
Xiang Zhou and *Shuting Gu*, City University of Hong Kong, Hong Kong

continued in next column

Thursday, February 28

MS288

Mori-Zwanzig Formulation, Implementation and Applications - Part II of II

2:15 p.m.-3:55 p.m.

Room: 201C

For Part 1 see MS254

The Mori-Zwanzig formalism allows reducing the number of variables in large systems of coupled equations. For differential equations, the reduced equations model the effect of the unresolved variables, leading to a Markovian, memory and fluctuating terms. This formalism can be a starting point for multiscale and meso-scale modeling, based on first-principles calculations. It can also be used to derive popular reduced order models as appropriate limiting cases. In addition to formulating better reduced models, there is a growing need for efficient simulation of those reduced models. The minisymposium will address the related mathematical and numerical issues as well as applications to materials, fluid mechanics, soft matter, biology and uncertainty quantification.

Organizer: Xiantao Li
Pennsylvania State University, U.S.

Organizer: Panos Stinis
Pacific Northwest National Laboratory, U.S.

Organizer: Daniele Venturi
University of California, Santa Cruz, U.S.

Organizer: Karthik Duraisamy
University of Michigan, Ann Arbor, U.S.

2:15-2:35 On the Estimation of the Mori-Zwanzig Memory Integral
Yuanran Zhu, University of California, Santa Cruz, U.S.

2:40-3:00 Mori-Zwanzig Formalism and Discrete-time Modeling of Chaotic Dynamics
Kevin K. Lin, University of Arizona, U.S.; Fei Lu, Johns Hopkins University, U.S.; Alexandre Chorin, University of California, Berkeley, U.S.

3:05-3:25 Renormalized Reduced Order Models for Long Term Prediction
Jacob Price, University of Puget Sound, U.S.; Panos Stinis, Pacific Northwest National Laboratory, U.S.

3:30-3:50 Coarse-grained Modeling of Polymers in Solvent Via Mori-Zwanzig Formalism
Wenxiao Pan and Shu Wang, University of Wisconsin, Madison, U.S.

continued in next column

Thursday, February 28

MS289

Underwater Sensing and Signal Processing - Part II of II

2:15 p.m.-3:55 p.m.

Room: 202A

For Part 1 see MS255

The earth's vast oceans are largely unexplored, but at the same time are of great importance for economic, ecological and military objectives. Advances in sensing technology (e.g. vector sensors), platform capabilities (e.g. Unmanned Underwater Vehicles) and computational resources create an exciting environment for new signal processing algorithms and sensing opportunities. Many of the problems involve computational challenges such as prediction of long-range acoustic propagation through a changing, multiple-scale environment, and simultaneous processing of data from hundreds or thousands of sensors. This minisymposium will identify current directions of research and methods being developed to solve problems, including their advantages and shortcomings.

Organizer: Margaret Cheney
Colorado State University, U.S.

Organizer: Kevin P. Bongiovanni
Raytheon BBN Technologies, U.S.

Organizer: Ivars Kirsteins
Naval Undersea Warfare Center, U.S.

2:15-2:35 Multistatic Localization with Uncertain Data Association
Jamie Irza, Raytheon Company, U.S.

2:40-3:00 Identifying Acoustic Scattering Mechanisms for Tilted Solid Cylinders in Water by Comparing Ray Theory with Multi-domain Processing
Timothy Daniel and Philip Marston, Washington State University, U.S.

continued on next page

Thursday, February 28

MS289

Underwater Sensing and Signal Processing - Part II of II

continued

3:05-3:25 Modeling the Acoustic Response of Elastic Targets in a Layered Medium Using the Coupled Finite Element/Boundary Element Method

Ahmad Abawi, HLS Research, U.S.; Petr Krysl, University of California, San Diego, U.S.

3:30-3:50 Improving Sonar Training Systems using WaveQ3D

Sean Reilly, Raytheon BBN Technologies, U.S.

Thursday, February 28

MS290

Scalable Solvers for the Helmholtz Problem - Part I of II

2:15 p.m.-3:55 p.m.

Room: 202B

For Part 2 see MS322

Many wave phenomena can be described by the Helmholtz equation in the frequency domain approach. After discretization large linear systems have to be solved. To maintain a certain accuracy, the discretization relies on the number of grid points per wavelength, which should be as high as possible. Increasing the wavenumber, therefore leads to very large linear systems. Most solvers which are available have the property that the number of iterations increases at least linearly with the wave number. Only recently, various methods become available where the number of iterations is independent of the size of the wave number, and the work per iteration is approximately constant, so-called scalable methods. In this minisymposium various methods will be presented which are scalable when certain conditions are satisfied. Typical applications of these solvers are medical imaging, seismics, sonar etc.

Organizer: Kees Vuik
Delft University of Technology, Netherlands

Organizer: Reinhard Nabben
Institut für Mathematik, Germany

2:15-2:35 Field-of-values Analysis of a Two-level Shifted Laplace Preconditioner for Finite Element Helmholtz Problems

Luis Garcia Ramos, Technische Universitaet Berlin, Germany; Reinhard Nabben, Institut für Mathematik, Germany

2:40-3:00 Scalable Convergence using Deflation Preconditioning for the Helmholtz Equation

Vandana Dwarka, Delft University of Technology, Netherlands

3:05-3:25 A Hybrid Domain Decomposition and Shifted Laplacian Multigrid Preconditioner for the Elastic Helmholtz Equation

Eran Treister, Ben Gurion University Negev, Israel

3:30-3:50 Fast Contour Integral Preconditioner for Solving 3D High-frequency Helmholtz Equations

Yuanzhe Xi, Emory University, U.S.; Xiao Liu, Rice University, U.S.; Yousef Saad, University of Minnesota, U.S.; Maarten de Hoop, Rice University, U.S.

continued in next column

Thursday, February 28

MS291

Towards Digital Twins for Industrial Applications - Part II of II

2:15 p.m.-3:55 p.m.

Room: 202C

For Part 1 see MS257

The future development of industry and society exhibits strongly increasing complexity, at the same time ever-shorter innovation cycles. In addition, digitisation and the internet of things have led to an explosion of data and information. Without novel computational tools and paradigms we will not be able to manage these challenges. For almost all domains of science and engineering and in most industrial sectors, a multitude of commercial and open source software for modelling, simulation and optimization (MSO) based on mathematical models is available. At the same time increasingly large amounts of process and product data are available and strong artificial intelligence solutions have been developed to exploit these. All this is fostered by computers becoming more and more powerful. These developments lead to the vision that holistic approaches can be achieved that combine all these developments. A complete industrial product or process in its whole life cycle can be accompanied by a virtual representation, often called digital twin that allows design optimization, process control, lifecycle management, predictive maintenance, risk analysis and many other features. Digital twins are so important to business today, that they were named one of Gartner's Top 10 Strategic Technology Trends for 2017. They are becoming a business imperative, covering the entire lifecycle of an asset or process and forming the foundation for connected products and services.

Organizer: Wil Schilders

Eindhoven University of Technology, Netherlands

2:15-2:35 Towards Digital Twins for the Distribution Management in Energy Transport Networks

Caren Tischendorf, Humboldt University Berlin, Germany

2:40-3:00 HPC Oriented Cloud Infrastructure for Digital Twins

Zoltán Horváth, *Tamás Budai*, and Ákos Kovács, Széchenyi István University, Győr, Hungary

3:05-3:25 What Could Be Digital Twins for the Insurance Industry? Example of the Sensibility Analysis of the Solvency Capital Requirement

Veronique Maume-Deschamps and Kevin Elie dit Cosaque, University of Lyon 1, France

3:30-3:50 Towards Digital Tokamaks

Barry Koren, Eindhoven University of Technology, Netherlands

continued in next column

Thursday, February 28

MS292

Showcase of Research Supported by the DOE Computational Science Graduate Fellowship - Part II of II

2:15 p.m.-3:55 p.m.

Room: 203

For Part 1 see MS258

Since 1991 the DOE computational science graduate fellowship (CSGF) program has supported the training of exceptional researchers across a range of science and engineering fields. In this minisymposium, CSGF fellows and alumni present their latest research accomplishments. The sessions will cover a wide range of computational science application areas.

Organizer: David I. Ketcheson

King Abdullah University of Science & Technology (KAUST), Saudi Arabia

Organizer: Mayya Tokman

University of California, Merced, U.S.

2:15-2:35 Communication-avoiding Distributed-memory Algorithms for QR Factorization

Edward Hutter, University of Illinois, Urbana-Champaign, U.S.

2:40-3:00 Designing HPC Applications for Exascale Supercomputers and Beyond

Jeff R. Hammond, Intel Corporation, U.S.

3:05-3:25 Computer Security: The Science of the Impossible

Brent Kirkpatrick, Intrepid Net Computing, U.S.

3:30-3:50 Controlled Experimentation at Scale

Jordan Atlas, Microsoft Corporation, U.S.

Thursday, February 28

MS293

Partitioned and Adaptive Methods for Initial Value Problems - Part I of II

2:15 p.m.-3:55 p.m.

Room: 205

For Part 2 see MS325

Numerical methods for initial value problems often adapt or partition the method to suit the problem. For differential equations that have a separation into stiff and nonstiff terms, methods such as implicit-explicit methods and integrating factor methods exploit this structure by partitioning the method based on the structure of the differential equation. These methods allow larger stable time-steps than standard explicit methods at a smaller computational cost than traditional fully implicit methods. The analysis and efficient implementation of these methods requires careful treatment of the adaptivity or partitioning along with accuracy and stability. This minisymposium brings together speakers working on the application and analysis of and partitioned methods for initial value problems.

Organizer: Andrew J. Steyer

Sandia National Laboratories, U.S.

Organizer: Sidafa Conde

Sandia National Laboratories, U.S.

2:15-2:35 A Family of Second and Third Order Implicit-explicit Runge-Kutta Methods for Stiff Time-dependent Partial Differential Equations

Andrew J. Steyer, Sandia National Laboratories, U.S.; Christopher J. Vogl, Lawrence Livermore National Laboratory, U.S.; Mark A. Taylor and Oksana Guba, Sandia National Laboratories, U.S.

2:40-3:00 On the Development and Implementation of Optimized, High-order Time Integrators for Multi-physics Problems

Daniel R. Reynolds, Southern Methodist University, U.S.; David J. Gardner, Lawrence Livermore National Laboratory, U.S.

3:05-3:25 Devising Unconditionally Stable Multistep ImEx Schemes that Avoid Stiff Nonlinear Implicit Terms

David Shirokoff, New Jersey Institute of Technology, U.S.; Benjamin Seibold, Temple University, U.S.; Dong Zhou, California State University, Los Angeles, U.S.

3:30-3:50 Adaptive Krylov Based Time Integration Methods

Paul Tranquilli, Virginia Tech, U.S.

continued in next column

Thursday, February 28

MS294

Applications of Data Assimilation in Science, Engineering, and Industry - Part II of II

2:15 p.m.-3:55 p.m.

Room: 206A

For Part 1 see MS260
Featured Minisymposium

These days data is generated at unprecedented levels in Science, Engineering, and Industry. It is becoming increasingly important to use this data in a smart way in order to stay competitive. In many cases, this may involve inversion or assimilation of data into a complex model, for example derived from physical laws, or it may involve inversion of data alone, for example to learn a complex model in a context where no first principles model exists. This minisymposium aims to explore recent novel applications of methodology and algorithms for data assimilation and inference in Science, Engineering, and Industry.

Organizer: Amr El-Bakry

ExxonMobil Upstream Research Company, U.S.

Organizer: Kody Law

University of Manchester, United Kingdom

2:15-2:35 Big Data-based Analytics for Supercomputer Logs: Architecture and Application

Byung Hoon Park, Oak Ridge National Laboratory, U.S.

2:40-3:00 Multilevel Computational Design and Analysis of Logistics Networks

Quan Long, Nai-Yuan Chiang, *Thomas A. Frewen*, and Yiqing Lin, United Technologies Research Center, U.S.

continued on next page

3:05-3:25 Multilevel Ensemble Kalman Filter

Haakon Hoel and Raul F. Tempone, King Abdullah University of Science & Technology (KAUST), Saudi Arabia; Kody Law, University of Manchester, United Kingdom; Fabio Nobile, École Polytechnique Fédérale de Lausanne, Switzerland; Alexey Chernov, Carl von Ossietzky Universitaet Oldenburg, Germany

3:30-3:50 High-dimensional Bayesian Inversion with Generative and Invertible Neural Networks

Yinhao Zhu and Nicholas Zabaras, University of Notre Dame, U.S.

Thursday, February 28

MS295

Statistical Applications of Continuous and Discrete Transport - Part I of II

2:15 p.m.-3:55 p.m.

Room: 206B

For Part 2 see MS327

Transport between probability measures offers a rich toolbox of approaches for sampling and inference, and for comparing both discrete and continuous distributions. Deterministic couplings between probability measures can be induced by transport maps that minimize certain cost functionals; more generally, one may derive many tools and insights from the construction of optimal transport plans. Recent theoretical and computational advances in this area have enabled the solution to ubiquitous statistical problems, including density estimation, variational Bayesian inference (parametric and nonparametric), inference with intractable likelihoods, as well as nonlinear filtering and smoothing. Despite these successes, efficiently solving such problems remains challenging, especially in high dimensions with strongly non-Gaussian distributions, or when only limited information or few samples are available. This minisymposium will explore new algorithms and methodologies for computing and employing transport in the context of statistics and machine learning.

Organizer: Ricardo Baptista

Massachusetts Institute of Technology, U.S.

Organizer: Daniele Bigoni

Massachusetts Institute of Technology, U.S.

Organizer: Matthew Parno

US Army Corps of Engineers, U.S.

2:15-2:35 Layered Transport Maps for Conditional Sampling

Matthew Parno, US Army Corps of Engineers, U.S.

2:40-3:00 Scable Learning and Inference with Optimal Transport and Stein's Method

Qiang Liu, University of Texas at Austin, U.S.; Jian Peng, University of Illinois at Urbana-Champaign, U.S.; Jason Lee, University of Southern California, U.S.

3:05-3:25 Transport Maps and Outer Probability Measures for Bayesian Inference

Jeremie Houssineau, National University of Singapore, Singapore; Emmanuel Delande, University of Texas at Austin, U.S.; Junjie Wang, Harbin Institute of Technology, China

3:30-3:50 A Stein Variational Newton Method

Gianluca Detommaso, University of Bath, United Kingdom; Tiangang Cui, Monash University, Australia; Youssef M. Marzouk and Alessio Spantini, Massachusetts Institute of Technology, U.S.; Robert Scheichl, Universität Heidelberg, Germany

continued in next column

Thursday, February 28

MS296

Modelling with Fractional PDEs: Statistical and Numerical Analysis - Part II of II

2:15 p.m.-3:30 p.m.

Room: 206C

For Part 1 see MS262

Many nonlocal phenomena can be effectively modeled with fractional calculus. Fractional models often require estimation of orders of fractional integral and/or differentiation and also require efficient numerical methods for accurate solutions. In this minisymposium, we present various topics in fractional modeling with PDEs using statistical and numerical tools.

Organizer: Hong Wang

University of South Carolina, U.S.

2:15-2:35 Solving Fractional Partial Differential Equations using a Deep Neural Network

Hong Wang, University of South Carolina, U.S.

2:40-3:00 Lattice Boltzmann Model for Time Sub-diffusion Problem in Caputo Sense

Rui Du, Southeast University, China

3:05-3:25 A Fast Temporal Second-order Difference Scheme for the Fractional Wave Equation

Hong Sun, Jiangsu University of Science and Technology, China

3:30-3:50 Spectral Methods for Fractional Advection-diffusion-reaction Equations on Smooth Domains

Zhaopeng Hao, Marcus Sarkis, and Zhongqiang Zhang, Worcester Polytechnic Institute, U.S.

Thursday, February 28

MS297

Multi-scale Modeling and Computation of Complex and Active Fluids - Part II of II

2:15 p.m.-3:30 p.m.

Room: 206D

For Part 1 see MS263

Active suspensions, which can consist of micro-swimmers, vesicles, filaments, microtubule-motor networks or other types of interacting particles immersed in liquid, can exhibit complex dynamics and patterns. Capturing the interplays between the particles and the surrounding fluid in continuum or agent-based models presents significant challenges as hydrodynamic interactions happen on short and long scales. Added difficulties arise when such suspensions interact with boundaries of non-trivial shape. This minisymposium's goal is to explore recent advances, especially on the computational front, on this emerging and expanding topic.

Organizer: Enkeleida Lushi

New Jersey Institute of Technology, U.S.

2:15-2:35 Dynamics of Colloids Above a Bottom Wall Driven by Active Torques and Forces

Aleksandar Donev, Courant Institute of Mathematical Sciences, New York University, U.S.

2:40-3:00 A Coarse-grained Model for Beds of Elastic Fibers

David Stein, Simons Foundation, U.S.; Michael J. Shelley, Courant Institute of Mathematical Sciences, New York University, U.S.

3:05-3:25 Microswimmer Design: Pullers and Pushers from Bimetallic Rods

Florencio Balboa Usabiaga, Simons Foundation and Flatiron Institute, U.S.; Quentin Brosseau, New York University, U.S.; Enkeleida Lushi, New Jersey Institute of Technology, U.S.; Yang Wu, New York University, U.S.; Leif Ristroph and Jun Zhang, Courant Institute of Mathematical Sciences, New York University, U.S.; Michael D. Ward, New York University, U.S.; Michael J. Shelley, Courant Institute of Mathematical Sciences, New York University, U.S.

Thursday, February 28

MS298

Exascale Software for High-order Methods - Part II of II

2:15 p.m.-3:55 p.m.

Room: 207

For Part 1 see MS264

This minisymposium focuses on discussing high-order discretization kernels and lightweight portable libraries that are critical for performance and exascale applications. Topics will include efficient next-generation discretization algorithm and software of arbitrarily high order to impact the design of exascale architectures, system and application software for improved portability and performance of the high-order methods.

Organizer: Tzanio Kolev

Lawrence Livermore National Laboratory, U.S.

2:15-2:35 Efficient High-order Discretization with MFEM

Veselin Dobrev, Lawrence Livermore National Laboratory, U.S.

2:40-3:00 A Novel Approach to Using Fused Batched BLAS to Accelerate High-order Discretizations

Stanimire Tomov and Ahmad Abdelfattah, University of Tennessee, Knoxville, U.S.; Jack J. Dongarra, University of Tennessee and Oak Ridge National Laboratory, U.S.; Azzam Haidar, NVIDIA, U.S.

3:05-3:25 Code for Efficient Extensible Discretizations LibCEED

Thilina Rathnayake, University of Illinois, U.S.; Paul Fischer, University of Illinois at Urbana-Champaign, U.S.

3:30-3:50 Low-order Preconditioning High-order Triangular Finite Element Methods and Introducing LibParanumal

Ali Karakus, Virginia Tech, U.S.; Noel A. Chalmers, University of Waterloo, Canada; Tim Warburton, Virginia Tech, U.S.

Thursday, February 28

MS299

Dynamics with Inherent Noise: Stochastic Modelling and Simulation - Part II of II

2:15 p.m.-3:55 p.m.

Room: 401A

For Part 1 see MS265

Inherent noise is ubiquitous in complex systems such as physics, chemistry, engineering and system biology. Numerical simulations based on stochastic models provide an important tool to understand the influence of noise and the dynamic properties of these systems beyond equilibrium. Synergy of stochastic modelling and numerical solutions techniques often leads to novel ideas and promote applications of stochastic models and solvers. In this minisymposium, we focus on both stochastic modelling and numerical methods with emphasis on the interaction of the-state-of-art computational techniques with applications in modelling dynamic process of complex systems. We invite speakers from both communities and expect them to have fruitful discussion. The speakers will address stochastic modelling problems and numerical techniques to solve stochastic equations arising in various applications. Specific topics includes stochastic dynamics modelled by Markov processes with applications to biology and chemical reaction systems, numerical techniques such as singular perturbation methods, surrogate model methods, long time integration of nonlinear SDE, model reduction methods, etc..

Organizer: Huan Lei

Pacific Northwest National Laboratory, U.S.

2:15-2:35 Transitions as Rare Events in Stochastic Delayed Systems

Jiayu Zhai, University of Massachusetts, Amherst, U.S.; Xiaoliang Wan, Louisiana State University, U.S.

2:40-3:00 Homogenization for Generalized Langevin Equations

Soon Hoe Lim, KTH Royal Institute of Technology, Sweden

3:05-3:25 Computational Stochastic Dynamics

Jinqiao Duan, Illinois Institute of Technology, U.S.

3:30-3:50 Model-form Uncertainty Quantification for Predictive Modeling in Probabilistic Graphical Models

Jinchao Feng, University of Massachusetts, Amherst, U.S.; Joshua Lansford, University of Delaware, U.S.; Markos A. Katsoulakis, University of Massachusetts, Amherst, U.S.; Luc Rey-Bellet, University of Massachusetts, U.S.; Dionisios Vlachos, University of Delaware, U.S.

continued in next column

Thursday, February 28

MS300

Inverse Problems in Medical Imaging - Part II of II

2:15 p.m.-3:55 p.m.

Room: 401B

For Part 1 see MS266

Image formation is a nonlinear inverse problem in many emerging medical imaging modalities. Examples include electrical impedance tomography, magnetic particle imaging, photoacoustic tomography and regularized sparse-data X-ray tomography. Typically some part of the imaging process is ill-posed, or extremely sensitive to modelling errors and measurement noise. Regularized nonlinear mathematical reconstruction algorithms are needed for robust and reliable imaging. This session presents results where new nonlinear mathematical techniques are combined with machine learning to yield results with unprecedented image quality.

Organizer: Melody Alsaker

Gonzaga University, U.S.

Organizer: Tatiana A. Bubba

University of Helsinki, Finland

Organizer: Samuli Siltanen

University of Helsinki, Finland

2:15-2:35 Joint Reconstruction in Multienergy X-ray Tomography

Alexander Meaney, University of Helsinki, Finland; Jussi Toivanen, University of Eastern Finland, Finland; Samuli Siltanen, University of Helsinki, Finland; *Ville P. Kolehmainen*, University of Eastern Finland, Finland

2:40-3:00 Forward and Inverse Computation in Ultrasound Tomography

Jennifer L. Mueller, Colorado State University, U.S.

3:05-3:25 Classifying Stroke using Electrical Impedance Tomography

Samuli Siltanen, University of Helsinki, Finland

3:30-3:50 Debluring and Denoising of Signals with Uncertainty Quantification with Applications to Inverse Problems

Viktoria Taroudaki, Eastern Washington University, U.S.

Thursday, February 28

MS301

Model Reduction and Reduced-order Modeling of Dynamical Systems - Part II of II

2:15 p.m.-3:55 p.m.

Room: 401C

For Part 1 see MS267

Large-scale dynamical systems are used to describe real-world processes in almost every area of computational sciences and engineering, ranging from structural dynamics to nano-electronics to fluid dynamics. It is long since known that in many-query scenarios, where the same model is to be simulated for varying control functions, source terms, and/or parameters, reduced-order surrogates can be used to significantly accelerate the engineering design, control and optimization processes for the underlying problem. In this minisymposium, we will investigate two facets of generating reduced-order surrogate models: first, using "classical" model order reduction in order to obtain a compact representation of a large-scale model arising, e.g., from a finite element discretization of an unsteady PDE, and second, computing a low-dimensional surrogate from available (experimental or simulation) data. Ideally, both approaches are used complementary, and we hope to bring together researchers from both sides in order to exchange ideas for potential hybrid methods using the best from both worlds!

Organizer: Peter Benner

Max Planck Institute for Dynamics of Complex Technical Systems, Germany

Organizer: Athanasios C. Antoulas

Rice University, U.S.

2:15-2:35 Choosing Interpolation Points in Interpolatory Model Reduction

Athanasios C. Antoulas, Rice University, U.S.

2:40-3:00 Model Reduction of Linear and Nonlinear Semi-discretized Fluid Flow Problems from Data

Ion Victor Gosea, Max Planck Institute, Magdeburg, Germany

3:05-3:25 Learning Reduced-order Models in New Variables: Lifting and Operator Inference

Elizabeth Qian and Boris Kramer, Massachusetts Institute of Technology, U.S.; Karen E. Willcox, University of Texas at Austin, U.S.

3:30-3:50 Data-driven Strategies for Modeling Dissipative Systems

Christopher A. Beattie, Virginia Tech, U.S.

continued in next column

Thursday, February 28

MS302

Novel Approaches for Design and Analysis of Chaotic Dynamical Systems - Part II of II

2:15 p.m.-3:55 p.m.

Room: 402A

For Part 1 see MS268

As access to high performance computational resources increases, the demand for high-fidelity numerical simulations will only increase. Many of these high-fidelity numerical simulations exhibit chaotic behavior, including climate models, scale-resolving computational fluid dynamics simulations, combustion models, and even road traffic models. The high sensitivity to small perturbations and seemingly random behavior of chaotic dynamical systems make conventional design and analysis methods ill-conditioned or computationally intractable, especially sensitivity analysis. This minisymposium will discuss novel approaches that will allow mathematicians and engineers to analyze, design, and control chaotic dynamical systems. Two classes of sensitivity analysis approaches for mean flow quantities will be discussed. The first class of methods control the exponential growth of unstable linearized perturbations and includes least squares shadowing and space-split sensitivity. The second class of methods involves linearizing about the mean flow and modeling the linearization of the closure terms. Approaches for predicting behavior that deviates from the mean, intermittent and extreme events, will be also discussed. Computationally efficient implementations of all of these algorithms and further strategies to minimize their costs will be discussed in the talks as well as how they can enable gradient-based optimization, error estimation, and mesh adaptation for chaotic dynamical systems.

Organizer: Patrick J. Blonigan

Sandia National Laboratories, U.S.

continued on next page

Organizer: Qiqi Wang
Massachusetts Institute of Technology, U.S.

Organizer: Nisha Chandramoorthy
Massachusetts Institute of Technology, U.S.

2:15-2:35 Space-split Sensitivity (S3) Computation of Statistics in Chaotic Dynamical Systems
Nisha Chandramoorthy and Qiqi Wang, Massachusetts Institute of Technology, U.S.

2:40-3:00 Stability, Sensitivity and Optimization of Unsteady Thermoacoustic Systems
Luca Magri and Francisco Huhn, University of Cambridge, United Kingdom

3:05-3:25 Analysis of Rare Events using Transfer Operators and Community Clustering
Peter Schmid, Imperial College London, United Kingdom

3:30-3:50 Identifying Precursors to Intermittent or Extreme Events using Conditional Space-time Proper Orthogonal Decomposition
Oliver T. Schmidt, University of California, San Diego, U.S.; Peter Schmid, Imperial College London, United Kingdom

Thursday, February 28

MS303

Resilience for Large Scale CSE Applications - Part II of II

2:15 p.m.-3:55 p.m.

Room: 402B

For Part 1 see MS269

Resilience is an imminent issue for large scale computing systems and their applications due to the reliability concerns that stem from the architectural complexity to meet the power and performance requirements together. Under such unreliable computing systems, it is essential to introduce failure mitigations at the runtime and application layers to compliment the resilience at the hardware/system levels. Today, the major resilience scheme is coordinated checkpoint and restart (CR) that involves global coordination of processes and threads. This global recovery model entails inherent scalability issues to handle a large class of errors and failures, and thus alternative approaches would improve the scalability and reliability of computing systems and applications together. In this minisymposium, we will discuss the reliability issues of the large scale computing systems, application/algorithm based resilience and programming model support to enhance the reliability of application program executions.

Organizer: Keita Teranishi
Sandia National Laboratories, U.S.

Organizer: Luc Giraud
Inria, France

Organizer: Christian Engelmann
Oak Ridge National Laboratory, U.S.

Organizer: George Bosilca
University of Tennessee, Knoxville, U.S.

2:15-2:35 Resilience for Extreme Scale Systems: Understanding the Problem
Christian Engelmann, Oak Ridge National Laboratory, U.S.

2:40-3:00 Memory Errors on Modern Systems: An Analysis of Data over the Lifetime of Cielo
Scott Levy and Kurt Ferreira, Sandia National Laboratories, U.S.

3:05-3:25 Scalable, Efficient Fault Tolerance in Asynchronous Many Task (AMT) Programming Models
Keita Teranishi and Nicole Slattengren, Sandia National Laboratories, U.S.; Sri Raj Paul, Rice University, U.S.; *Hemanth Kolla*, Sandia National Laboratories, U.S.; Akihiro Hayashi, Rice University, U.S.; Jackson Mayo, Sandia National Laboratories, U.S.; Vivek Sarkar, Georgia Institute of Technology, U.S.

3:30-3:50 Title Not Available
Aurelien Bouteiller and George Bosilca, University of Tennessee, Knoxville, U.S.

continued in next column

Thursday, February 28

MS304

State-of-the-art High-order Numerical Methods and Complex Fluid Simulations - Part II of II

2:15 p.m.-3:55 p.m.

Room: 402C

For Part 1 see MS270

Simulating compressible flows is challenging, due to the presence of discontinuities, e.g. shock and contact waves, and broad-band continuous flow scales. High-order and high-resolution schemes dedicate to resolving smooth flow scales accurately and to capturing the discontinuities simultaneously. In the past decades, many mathematical concepts, such as the artificial viscosity scheme, the total variation diminishing (TVD) scheme, the essentially non-oscillatory (ENO) scheme and the weighted essentially non-oscillatory (WENO) scheme, have been proposed and widely used in both academic and engineering simulations, as well as the compact scheme, flux reconstruction (FR) scheme, spectral volume/difference (SV/D) scheme etc. More recently, the high-order numerical methods pay a more and more important role in high-fidelity simulations, e.g. the large eddy simulations (LES) and the direct numerical simulations (DNS). This session focuses on the state-of-the-arts of high-order numerical methods and the complex fluid simulations, including but not limited to turbulence, multiphase flows, plasma, MHD flows, combustion etc. We aim to bridging the cutting-edge applied mathematics and the request of flow physics simulations.

Organizer: Lin Fu
Technische Universitaet Muenchen, Germany

Organizer: Nikolaus Adams
Technische Universitaet Muenchen, Germany

2:15-2:35 Large-eddy Simulation of Separated Flows: The Role of Subgrid-scale Models
Zuoli Xiao and Rui Wang, Peking University, China

2:40-3:00 Compact Reconstruction Using Teno
James D. Baeder, University of Maryland, U.S.

3:05-3:25 A Fourth-order Adaptive Mesh Refinement Algorithm for the Multicomponent, Reacting Compressible Navier-Stokes Equations
Emmanuel Motheau, Matthew Emmett, Weiqun Zhang, Michael Minion, and John B. Bell, Lawrence Berkeley National Laboratory, U.S.

3:30-3:50 Time-accurate and Highly-stable Explicit Operators
Maxime Bassenne and Ali Mani, Stanford University, U.S.

continued in next column

Thursday, February 28

CP16

Numerical PDEs I

2:15 p.m.-3:55 p.m.

Room:301

Chair: Emine Celiker, University of Lincoln, United Kingdom

2:15-2:30 A Highly-accurate Finite Element Method for the Solution of the Dirichlet Problem of the Generalized Helmholtz Equation in Non-convex Polygons
Emine Celiker, University of Lincoln, United Kingdom; Ping Lin, University of Dundee, United Kingdom

2:35-2:50 An Deformation Based Adaptive Moving Grid Method
Yinlin Dong, University of Central Arkansas, U.S.

2:55-3:10 A Posteriori Error Estimates and Adaptive Mesh Refinement for the Stokes-Brinkman Problem
Kevin Williamson, University of Maryland, Baltimore County, U.S.; Pavel Burda, Czech Technical University, Prague, Czech Republic; Bedrich Sousedik, University of Maryland, Baltimore County, U.S.

3:15-3:30 A Communication Algorithm for the Patch-based Overlapping Grids Applications
Hong Guo, Institute of Applied Physics and Computational Mathematics, China

3:35-3:50 Generation of Structured Meshes for Shape Optimization
Camilla Hahn and Matthias Bolten, University of Wuppertal, Germany

Thursday, February 28

CP17

Computational Science Methods and Software

2:15 p.m.-3:55 p.m.

Room:201A

Chair: Andrey Prokopenko, Oak Ridge National Laboratory, U.S.

2:15-2:30 Recomputations and Communication Costs of Fast Matrix Multiplication with 2x2 Base Case
Roy Nissim and Oded Schwartz, Hebrew University of Jerusalem, Israel

2:35-2:50 Approximating the Generalized Singular Value Expansion
Matthew J. Roberts, Michigan Technological University, U.S.

2:55-3:10 Superlinear Convergence of the Methods of Quasi-Newton Method Based on Assumptions About the Initial Point
Sarah Nataj and Shaun Lui, University of Manitoba, Canada

3:15-3:30 A Visual Programming and Debugging Approach for Parallel Numerical Simulation Development
Li Liao and Aiqing Zhang, Institute of Applied Physics and Computational Mathematics, China

3:35-3:50 ArboX: A New Parallel Bounding Volume Hierarchy Implementation Using MPI+kokkos
Andrey Prokopenko, Damien Lebrun-Grandie, Bruno Turcksin, and Stuart Slattery, Oak Ridge National Laboratory, U.S.

Intermission

3:55 p.m.-4:10 p.m.

Thursday, February 28

MS305

Computing Tensor Decompositions

4:10 p.m.-5:50 p.m.

Room: Ballroom 100BC

Tensor decomposition -- representing a multiway tensor by a combination of simple rank-1 components -- is an important operation for processing and discovering features in complex data. But due to the size and/or irregularity of tensor data, tensor decomposition is challenging both algorithmically and computationally. Strategies that increase memory locality, reduce computation and memory costs, and exploit parallelism can enable efficient analysis of large- and extreme-scale tensors. This minisymposium presents a collection of numerical algorithms and implementation techniques that benefit tensor computations.

Organizer: Karen D. Devine
Sandia National Laboratories, U.S.

Organizer: Tamara G. Kolda
Sandia National Laboratories, U.S.

4:10-4:30 On Tensor Orderings for HiCOO
Bora Ucar, LIP-ENS Lyon, France

4:35-4:55 On the Decomposition of Tensors that are Given Implicitly
Nico Vervliet, Martijn Boussé, Ignat Domanov, and Lieven De Lathauwer, KU Leuven, Belgium

5:00-5:20 Tensor Decompositions on Emerging Manycore Hardware with Genten and Kokkos
Eric Phipps and Tamara G. Kolda, Sandia National Laboratories, U.S.

5:25-5:45 Parallel Sparse Tensor Decomposition with the Trilinos Parallel Linear Algebra Framework
Karen D. Devine, Sandia National Laboratories, U.S.

Thursday, February 28

MS306

Performance Portability of Scientific HPC Applications

4:10 p.m.-5:50 p.m.

Room:Conference Theater

Modernizing scientific codes to run performantly on multiple emerging architectures is a task of great complexity and challenge for computational scientists today. Today's developers face complex programming language features, use of portability abstractions, increases in hardware complexity, and a plethora of runtime systems on offer. Working at the intersection of these tools is not only difficult for developers, but also presents increasingly difficult problems to compilers. This minisymposium will address strategies that help ease computational science problems for large codes today.

Organizer: Geoff Womeldorff
Los Alamos National Laboratory, U.S.

4:10-4:30 Modernizing Compiler Design for Platform Portability
Patrick McCormick, Los Alamos National Laboratory, U.S.

4:35-4:55 Using Just-in-time Compilation to Optimize Scientific Codes
David F. Richards, Lawrence Livermore National Laboratory, U.S.

5:00-5:20 Enabling Application Optimizations with Unified Thread-based Runtime MPC
Julien C. Jaeger, CEA, France

5:25-5:45 Title Not Available
Simon D. Hammond, Sandia National Laboratories, U.S.

Thursday, February 28

MS307

Technology Transfer Using the Open-source Platform SU2

4:10 p.m.-5:50 p.m.

Room: 102A

SU2 is an open-source collection of software tools written in C++ and Python for the analysis of partial differential equations (PDEs) and PDE-constrained optimization problems on unstructured meshes with state-of-the-art numerical methods. Through the initiative of users and developers around the world, SU2 is now a well established tool in the computational sciences with wide applicability to aeronautical, automotive, naval, and renewable energy industries, to name a few. In this minisymposium we will present innovative methods from different groups in industry and academia which are in the stage of being transferred from research to their application to practical problems.

Organizer: Tim Albring

Technische Universität Kaiserslautern, Germany

Organizer: Thomas D. Economon

Bosch Research And Technology Center, U.S.

4:10-4:30 Parallel-in-time Integration in SU2

Tim Albring, Stefanie Guenther, and Nicolas R. Gauger, Technische Universität Kaiserslautern, Germany

4:35-4:55 Unsteady Optimization with Harmonic-balance Methods in SU2: Lessons Learned and Future Perspectives

Matteo Pini and Antonio Rubino, Delft University of Technology, Netherlands

5:00-5:20 Advances for Incompressible Flow Calculations with SU2

Thomas D. Economon, Bosch Research And Technology Center, U.S.

5:25-5:45 Wall-modeled LES Simulations using the SU2 DG-FEM Solver

Edwin van der Weide, University of Twente, Netherlands; *Juan J. Alonso*, Stanford University, U.S.

Thursday, February 28

MS308

Data-driven and Mathematical Model Reductions for Combustion System Simulation and Design

4:10 p.m.-5:50 p.m.

Room: 102B

Combustion systems are extremely computationally intensive to simulate. State of the art high-fidelity fluid-combustion simulations are typically large-eddy simulations of the Navier-Stokes equations with detailed chemical mechanisms composed of thousands of chemical species and reactions. Such simulations are intractable for all but the simplest problems and model reductions are necessary for any simulation of a physically relevant combustion system. Typically, such reductions neglect some subset of the physics that is not deemed necessary for the current problem. Physics reductions introduce error that is difficult to quantify as it is not fully know how much of the flow violates the assumptions of the reduction being used and for how long. Recent developments, presented in this session, use mathematical methods to select the most efficient reduced physics model to attain a target level of accuracy in each region of the flow. Such techniques enable the simulation of simple combustion devices. However, many problems of interest, remain intractable using high-fidelity simulations and, as such, further data-driven and mathematical model reductions, such as projection-based reduced basis methods, must be utilized to enable design iterations. The rest of this session discusses such methods, their pitfalls when applied to combustion problems, and concepts to overcome these roadblocks.

Organizer: Nathan L. Mundis

ERC Inc. and Air Force Research Labratory, U.S.

4:10-4:30 Combustion System Model Reduction Using Pod and Neural Networks

Nathan L. Mundis, ERC Inc. and Air Force Research Labratory, U.S.

4:35-4:55 Multi-fidelity Modeling of Rocket Combustion Instability

Cheng Huang and Karthik Duraisamy, University of Michigan, Ann Arbor, U.S.; Charles Merkle, Purdue University, U.S.

5:00-5:20 A Multi-fidelity Framework for Stability Analysis of Complex Dynamical Systems

Alexandre Marques, Massachusetts Institute of Technology, U.S.

5:25-5:45 Fidelity-adaptive Combustion Modeling for Turbulent Combustion

Matthias Ihme, Stanford University, U.S.

continued in next column

Thursday, February 28

MS309

Machine Learning for Inverse Problems and its Applications

4:10 p.m.-5:50 p.m.

Room: 111A

Recent advances in computer science and data analytics have brought machine-learning techniques to the forefront of computational inverse problems research. As a result, old questions are being addressed in new ways as techniques are being developed to exploit large volume and complex data sets. In computer-science domain, scientists have recently developed various novel machine-learning algorithms including deep learning, dictionary learning, etc. There have been several pioneering research works incorporating different machine-learning techniques to solve inverse problems and apply them to various applications. The goal of this session is to introduce those recent results to our inverse-problems community. We welcome researchers and students to participate in our minisymposium and join our discussions on these exciting new directions.

Organizer: Youzuo Lin

Los Alamos National Laboratory, U.S.

Organizer: Cristina Garcia-Cardona

Los Alamos National Laboratory, U.S.

4:10-4:30 Convolutional Dictionary Learning for Inverse Problems

Cristina Garcia-Cardona, Los Alamos National Laboratory, U.S.

4:35-4:55 Ultrasonic Tomographic Imaging Using Physics Constrained Learning Method

Weichang Li, Saudi Aramco Oil Company, Saudi Arabia

5:00-5:20 Linear Regression on Imbalanced Data: New Theoretical Insights

Shusen Wang, Stevens Institute of Technology, U.S.

5:25-5:45 InversionNet: Accurate and Efficient Seismic Waveform-inversion with Convolutional Neural Networks

Youzuo Lin, Los Alamos National Laboratory, U.S.

Thursday, February 28

MS310

The Intersection of Graph Algorithms and Machine Learning

4:10 p.m.-5:50 p.m.

Room:111B

We are witnessing large growth in the number of proposed graph-theoretic and machine learning solutions for a variety of problems. Claims and counterclaims are increasingly being made as to which approach is best. In some cases, old problems are recast in an alternate approach in the hope of finding a better solution; in other cases, an approach is chosen to solve a new problem without necessarily a sound theoretical basis for success. The conundrum is that both graph algorithms and machine learning can solve many real-world problems, and that their domains intersect, but the approaches are not equivalent. For example, both community detection algorithms and SVMs partition data into subsets of similar members, and Bayesian Networks are a probabilistic graphical model of random variables and conditional dependencies used to learn causal relationships. In reality, many analytic workloads require both approaches: graphs to understand relationships and organizational structures, and machine-learning methods to identify signature features. The talks in this minisymposium will: discuss the problem domains for graph methods, machine learning, and their intersections; discuss representations that rely on concepts from graph theory, and enable the formulation of machine learning problems; provide tractability analysis in terms of complexity, time-to-solution, problem size, and quality of solution; discuss integration of graphs and machine learning approaches.

Organizer: Antonino Tumeo

Pacific Northwest National Laboratory, U.S.

Organizer: Arif Khan

Pacific Northwest National Laboratory, U.S.

4:10-4:30 Representation Learning for Large Graphs

Nesreen Ahmed, Intel Labs, U.S.

4:35-4:55 Scaling Clustering Algorithms Using Graph Theoretic Approaches

Mostofa Patwary, Baidu, U.S.

5:00-5:20 Parallel Algorithms for Ensemble Clustering

Ariful Azad, Indiana University, U.S.

5:25-5:45 Asynchronous Communication-efficient Algorithm for Clustering Coefficient Computation

Flavio Vella, Free University of Bozen-Bolzano, Italy; Torsten Hoefler and Salvatore Di Girolamo, ETH Zürich, Switzerland

continued in next column

Thursday, February 28

MS311

Applications of Machine Learning in Fluid Mechanics

4:10 p.m.-5:50 p.m.

Room: 111C

Problems in fluid mechanics can be extremely difficult to solve, as they often involve strong nonlinear effects, bifurcations, and turbulence, among other technical challenges. Furthermore, since fluid flows are governed by partial differential equations, computational solutions often require discretizations with a million or more states, leading to extremely high-dimensional models and datasets. Driven by innovations in Silicon Valley, machine learning has been proven effective in solving a wide array of complex problems, many of which are inherently nonlinear, may not have any physical model to rely on, and can involve enormous datasets. As such, there is a great deal of interest in whether techniques developed to solve problems in Silicon Valley may be helpful in addressing challenges encountered in fluid mechanics. This minisymposium presents research demonstrating the promise of this approach. Among the methods that will be discussed are linear discriminant analysis, support vector machines, neural networks, and clustering, applied to a variety of problems in classifying and modeling vortex patterns in bio-inspired fluid flows.

Organizer: Jonathan H. Tu

NSWC Carderock

4:10-4:30 A Search for Physical Insight in Machine Learning Classification of Vortex Patterns

Jonathan H. Tu, NSWC Carderock

4:35-4:55 Classifying Flow Patterns using Neural Networks

Mohamad Alsalman, Brendan Colvert, and *Eva Kanso*, University of Southern California, U.S.

5:00-5:20 Clustering and Classifying Vortex Wakes by Dynamical Regimes

Mengying Wang and Maziar S. Hemati, University of Minnesota, U.S.

5:25-5:45 Deep Learning for Gust Detection from Available Wing Sensors in Unsteady Aerodynamics

Jeff D. Eldredge and Wei Hou, University of California, Los Angeles, U.S.

Thursday, February 28

MS312

Optimization Under Uncertainty Using Multifidelity and Derivative-free Approaches

4:10 p.m.-5:50 p.m.

Room: 300A

Optimization under uncertainty broadly describes optimization problems in which uncertain parameters enter the objective and constraints. In the derivative-free setting, the uncertain parameters and design parameters may be inputs to a (nonlinear) black-box model, whose outputs in turn define the optimization problem. In this context, one may seek optimal designs that are "robust" to parametric uncertainties by, for example, maximizing the expectation of the original objective or minimizing its variance. Other formulations include probabilistic (chance) constraints, or objective/constraints based on conditional-value-at-risk. Monte Carlo sampling can be used to evaluate these quantities, but at considerable expense; this is particularly challenging when the underlying models are complex and expensive to evaluate. Given limited sample sizes, another challenge is how to take the resulting noisy objective or constraint function estimates into account. This minisymposium is devoted to the latest developments and applications of optimization under uncertainty, with a particular focus on derivative-free and sampling-driven approaches. We are interested in robust or risk-sensitive formulations, the use of multilevel or multi-fidelity Monte Carlo methods, and the construction of noise-resilient surrogate models in trust region methods, among other relevant topics.

Organizer: Friedrich Menhorn

Technische Universität München, Germany

Organizer: Gianluca Geraci

Sandia National Laboratories, U.S.

Organizer: Michael S. Eldred

Sandia National Laboratories, U.S.

4:10-4:30 Derivative-free Multifidelity Design Optimization under Uncertainty of a Scramjet

Friedrich Menhorn, Technische Universität München, Germany; Michael S. Eldred and Gianluca Geraci, Sandia National Laboratories, U.S.; Youssef M. Marzouk, Massachusetts Institute of Technology, U.S.

4:35-4:55 Derivative-free Robust Optimization by Outer Approximations

Matt Menickelly and *Stefan Wild*, Argonne National Laboratory, U.S.

5:00-5:20 Bayesian Optimization with Expensive Integrands

Saul Toscano-Palmerin and Peter I. Frazier, Cornell University, U.S.

5:25-5:45 Bayesian Optimization of Combinatorial Structures

Matthias Poloczek, University of Arizona, U.S.; Ricardo Baptista, Massachusetts Institute of Technology, U.S.

continued in next column

Thursday, February 28

MS313

Parallel Numerical Linear Algebra for Future Extreme-scale Systems

4:10 p.m.-5:50 p.m.

Room: 300B

We describe recent developments in task-based algorithms for the solution of both dense and sparse linear systems and the solution of both symmetric and unsymmetric dense eigenproblems. By using directed acyclic graphs (DAGs) we are able to avoid synchronization and obtain high levels of parallelism. We first discuss the case of dense matrices. Our algorithms keep communication to a minimum, sometimes achieving a provably theoretical minimum. For the direct solution of symmetrically structured sparse systems, we use dense matrix kernels, showing how robust numerical pivoting can be accommodated for both symmetric indefinite and unsymmetric matrices without losing much parallelism. We compare the performance of the code for unsymmetric systems with another code we have developed that is based on a Markowitz threshold approach with a novel technique for getting a set of pivots that can be used in parallel. We study the scalability of an approach for block iterative methods that uses our parallel indefinite direct solver on the blocks. We study robust algebraic preconditioners that are highly scalable and examine a way to expand the Krylov space for better convergence while maintaining parallelism and reducing communication. We show the performance of prototype codes on a range of research and industrial applications. This minisymposium present research done in NLAFET, a Horizon 2020 FET-HPC project funded by the European Union.

Organizer: Iain Duff

Science & Technology Facilities Council, United Kingdom and CERFACS, Toulouse, France

4:10-4:30 Distributed Tasking in the PLASMA Numerical Library

Mawussi Zounon and Jakub Sistek, University of Manchester, United Kingdom; Jack J. Dongarra, University of Tennessee and Oak Ridge National Laboratory, U.S.

4:35-4:55 Solution of Sparse Unsymmetric Systems

Iain Duff, Science & Technology Facilities Council, United Kingdom and CERFACS, Toulouse, France; Florent Lopez and Stojce Nakov, Rutherford Appleton Laboratory, United Kingdom

5:00-5:20 The Scalability of Block Iterative Methods

Sebastien Cayrols, Rutherford Appleton Laboratory, United Kingdom; Iain Duff, Science & Technology Facilities Council, United Kingdom and CERFACS, Toulouse, France

5:25-5:45 Preconditioned Linear Solvers in Cosmic Microwave Background Data Analysis

Jan Papež, Inria Paris, France

continued in next column

Thursday, February 28

MS314

Optimal Experimental Design for Bayesian Inverse Problems

4:10 p.m.-5:50 p.m.

Room: 300C

Featured Minisymposium

In the context of Bayesian inverse problems, the goal of optimal experimental design (OED) is to find the optimal design of a data acquisition system (e.g., location of sensors, what quantities are measured and how often, what sources are used in each experiment), so that the uncertainty in the inferred parameter field—or in some quantity derived from this field—is minimized with respect to some criterion. OED for Bayesian inverse problems governed by partial differential equations (PDEs) is an extremely challenging problem. First, the parameter to be inferred is often a spatially correlated field and it leads, after discretization, to a high dimensional parameter space. Second, the forward PDE model is often complex and computationally expensive to solve. Third, the design space for the data acquisition system may be very large. And fourth, the Bayesian inverse problem—difficult as it is—is just a subproblem within the OED problem. This minisymposium explores advances in numerical methods for OED in a Bayesian framework that are scalable with respect to the parameter and data dimensions.

Organizer: Umberto Villa

Washington University, St. Louis, U.S.

Organizer: Omar Ghattas

University of Texas at Austin, U.S.

4:10-4:30 Scalable Methods for Bayesian Optimal Experimental Design Using Laplace Approximation

Umberto Villa, Washington University, St. Louis, U.S.; Omar Ghattas, University of Texas at Austin, U.S.

4:35-4:55 Optimal Bayesian Experimental Design Using Generalized Laplace Method

Quan Long, United Technologies Research Center, U.S.

continued on next page

Thursday, February 28

MS314

Optimal Experimental Design for Bayesian Inverse Problems

continued

5:00-5:20 Sum-up Rounding Approaches for Design of Experiments
Jing Yu, University of Chicago, U.S.; Mihai Anitescu, Argonne National Laboratory, U.S.

5:25-5:45 A Sparse Optimization Framework for Optimal Sensor Placement
Daniel Walter and Boris Vexler, Technische Universität München, Germany

Thursday, February 28

MS315

Numerical Methods and Regularization for Ill-posed and Ill-conditioned Problems

4:10 p.m.-5:50 p.m.

Room: 300D

Ill-posed problems often arise from inverse problems, integral equations of the first kind, and statistics. Inverse problems are usually to determine the internal structure of a physical system from the outside measured quantities of the system, or to determine the unknown input that gives rise to a measure output. Examples of inverse problems include electrical impedance tomography (EIT) that produces a picture of electric conductivity distribution inside a patient from the measurement of the corresponding voltages at the electrodes attached to the patient, and solving diffusion equations in negative time direction to detect the source of pollution from measurements. It is widely recognized that linear systems resulting from discretizing ill-posed problems are ill-conditioned. All ill-conditioned linear systems either have the singular values of the coefficient matrix decay gradually to zero with no particular gap in the spectrum or have large condition numbers. In this minisymposium, we focus on developing numerical methods for solving ill-posed problems that avoid introducing ill-conditioned discrete systems and numerical algorithms that regularize ill-conditioned discrete systems by introducing preconditioners that regularize the spectrum to reduce the large condition number. The aim of the methods and algorithms is to reduce the computational expense of the conventional methods so that large-scale computation for the ill-posed or ill-conditioned problems can be achieved.

Organizer: Long Lee
University of Wyoming, U.S.

Organizer: Man-Chung Yeung
University of Wyoming, U.S.

4:10-4:30 A Global Reconstruction Method for a Coefficient Inverse Problem with One Measurement
Dinh-Liem Nguyen, Kansas State University, U.S.

4:35-4:55 Numerical Methods for Direct Reconstruction Methods in Electrical Impedance Tomography
Peter Muller, Rowan University, U.S.

5:00-5:20 A Spectral Projection Preconditioner for Solving Ill-Conditioned Linear Systems
Man-Chung Yeung, University of Wyoming, U.S.

5:25-5:45 Constrained Inversion for Subsurface Imaging
Jodi Mead, Boise State University, U.S.

continued in next column

Thursday, February 28

MS316

Rational Approximation and its Applications

4:10 p.m.-5:50 p.m.

Room: 302A

Recent developments of algorithms generating low-cost, high-fidelity approximations, such as Adaptive Anderson-Antoulas (AAA), Vector Fitting (VF), and Least-Squares Rational Approximation (LSRA), have renewed interest in rational approximation for computational science and engineering applications. Recent work formulates both system identification and model reduction as a rational approximation problems, providing data-driven approaches instead of realization-dependent algorithms. Additionally, as rational functions often provide accurate approximations using fewer terms than polynomials, they can also solve differential equations with fewer degrees of freedom. This minisymposium explores recent advances exploiting these algorithms and their impact on system identification, model reduction, and function approximation.

Organizer: Jeffrey M. Hokanson
University of Colorado Boulder, U.S.

Organizer: Caleb C. Magruder
Tufts University, U.S.

4:10-4:30 Projected Nonlinear Least Squares for H2 Model Reduction
Jeffrey Hokanson, University of Colorado Boulder, U.S.; *Caleb C. Magruder*, Tufts University, U.S.

4:35-4:55 Quadrature Based Rational Ls for Transfer Function Approximation
Zlatko Drmac, University of Zagreb, Croatia; Christopher A. Beattie and Serkan Gugercin, Virginia Tech, U.S.

5:00-5:20 Solving the Laplace Equation using Rational Functions
Abinand Gopal, University of Oxford, United Kingdom

5:25-5:45 The AAA Algorithm for Parameterized Dynamical Systems
Andrea Carracedo Rodriguez and Serkan Gugercin, Virginia Tech, U.S.

continued in next column

Thursday, February 28

MS317

BE: Overcoming Workplace Challenges Panel: Trials, Tribulations, and Triumphs

4:10 p.m.-5:50 p.m.

Room: 302B

Navigating a lifelong career is not without challenges - even highly successful people struggle at some point along the way. This workshop, geared toward students and early career researchers, will connect the audience to a panel of experienced professionals with backgrounds in industry, academia, and national labs who will share the personal challenges they faced in their STEM careers and how they overcame them. In addition, the panelists will share resources and strategies for facing a variety of career challenges, such as not getting your dream job, unproductive research, struggling to create work/life balance, dealing with workplace biases and micro-aggressions, struggling with Imposter Syndrome, and more. Following the presentations, a moderated Q&A session will give the audience the opportunity to engage and interact with the panelists.

Organizer: Melissa Abdelbaky
Rutgers University, U.S.

Thursday, February 28

MS318

Large-scale Particulate Flows: Towards Multi-scale Modeling from µm to km

4:10 p.m.-5:50 p.m.

Room: 303A

Flows of many small solid particles immersed in a fluid are extremely common in astrophysical (planetary formation), geophysical (landslides, avalanches), biological (blood flow), and industrial processes (oil and pharmaceutical production). These flows exhibit aspects of all phases of matter --gas, liquid, and solid-- making it a challenging system to model. The coupling of characteristic length scales --from micron to km-- requires algorithms capable of bridging vast length scales, and resolving different coexisting physical regimes. Atomistic simulations, tracking individual particles, are computationally too expensive to scale up to realistic sizes. This calls for new hybrid models which make use of novel techniques such as adaptive mesh refinement, spectral deferred correction methods, or coarse graining. Here we will explore the physics of complex flows containing particles in order to identify coarse grained models, which in turn will inform the next-generation multi-scale simulation techniques. In this minisymposium we will bring together leading experts in granular physics, fluid-structure interaction, and hybrid computational methods.

Organizer: Johannes P. Blaschke
Lawrence Berkeley National Laboratory, U.S.

Organizer: Mazza Marco
Loughborough University, United Kingdom

4:10-4:30 CFD-DEM: Modeling the Small to Understand the Large
Christine Hrenya, University of Colorado Boulder, U.S.

continued on next page

Thursday, February 28

MS318

Large-scale Particulate Flows: Towards Multi-scale Modeling from µm to km

continued

4:35-4:55 Blood Flow Modelling in Science and Engineering - from Single Vessels to Networks

Timm Krueger, Albert Beardo Ricol, Nicholas Daly, and Miguel Bernabeu, University of Edinburgh, United Kingdom

5:00-5:20 A Hybrid Material Point and Discrete Element Method for Granular Media Modeling

Maytee Chantharayukhonthorn, Massachusetts Institute of Technology, U.S.; Breannan Smith, Columbia University, U.S.; Yonghao Yue, University of Tokyo, Japan; Peter Chen and Eitan Grinspun, Columbia University, U.S.; Ken Kamrin, Massachusetts Institute of Technology, U.S.

5:25-5:45 Modeling of Active Particles and Hydrodynamics: Density Heterogeneities at Multiple Scales

Marco G. Mazza, Loughborough University, United Kingdom

Thursday, February 28

MS319

Efficient Computational Methods for Molecular Dynamics

4:10 p.m.-5:50 p.m.

Room: 303B

Supercomputing architectures head towards the exascale era and undergo drastical changes on (rather) short time scales: hybrid host-accelerator systems have become "state-of-the-art" within the top 10 supercomputers; RISC architectures (such as ARM processors) are gaining more popularity; vector widths in processors keep increasing; etc. Due to the wide range of applicability of molecular dynamics (MD), including material science, chemistry or process engineering, big user communities are confronted with arising supercomputing challenges and corresponding changes in programming paradigms. Furthermore, there are application-driven challenges such as load balancing that need to be taken into account to efficiently exploit large-scale compute systems. In this minisymposium, challenges for large-scale MD applications are pointed out and potential solutions in terms of simulation methodology are discussed. The focus is put on strategies to exploit current supercomputing systems in terms of efficient algorithms, programming, and self-adaption of code (auto-tuning). This picture is complemented by discussions on the integration of new methodological approaches into community software, such as GROMACS, ESPREsSo, and ls1 mardyn, which will eventually enable end users to leverage today's and tomorrow's supercomputing systems for their applications.

Organizer: Philipp Neumann
University of Hamburg, Germany

Organizer: Steffen Seckler
Technische Universität München, Germany

4:10-4:30 Addressing Exascale Challenges for Molecular Dynamics in GROMACS

Mark Abraham, KTH Royal Institute of Technology, Sweden

4:35-4:55 Scalable Solvers for Dynamic Charge Models in Molecular Dynamics Simulations

H. Metin Aktulga, Kurt O'Hearn, and Abdullah Alperen, Michigan State University, U.S.

5:00-5:20 Leveraging Node-level Performance for Molecular Dynamics Through Auto-tuning

Fabio A. Gratl and Steffen Seckler, Technische Universität München, Germany; Philipp Neumann, University of Hamburg, Germany; Hans-Joachim Bungartz, Technische Universität München, Germany

5:25-5:45 Comparison of Load-balancing Techniques for MD Simulations with Dynamic Particle Binding

Steffen Hirschmann and Dirk Pflüger, Universität Stuttgart, Germany; Colin W. Glass, HLRS, Germany

continued in next column

Thursday, February 28

MS320

Advances in Rare Event Simulation for Complex Dynamical Systems - Part II of II

4:10 p.m.-5:50 p.m.

Room: 201B

For Part 1 see MS287

The study of rare events in dynamical systems is essential for a wide range of applications including the robust design of engineering systems, the prediction of extreme weather events, and the study of economic catastrophes. Over the last twenty years, different research communities have developed a variety of tools for studying rare events in dynamical systems. In these sessions, we showcase tools adapted from sampling, dynamical systems, and statistics including subset simulation, importance sampling, particle splitting, information theoretic bounds, and large deviations. These talks will show applications of these tools in fluid dynamics, aerospace systems, chemistry, atmospheric modeling, finance, and electrical systems. We will also explore common themes and challenges across these approaches.

Organizer: Benjamin J. Zhang
Massachusetts Institute of Technology, U.S.

Organizer: Tuhin Sahai
United Technologies Research Center, U.S.

4:10-4:30 From Catastrophe Principle to Strongly Efficient Importance Sampling for Heavy-tailed Rare Events

Chang-Han Rhee and Bohan Chen, Centrum Wiskunde & Informatica, Netherlands; Jose Blanchet, Stanford University, U.S.; Bert Zwart, Centrum Wiskunde & Informatica, Netherlands

4:35-4:55 Efficient PCE Representation for Estimation of Rare Events in High Dimensions

Iason Papaioannou, Max Ehre, and Daniel Straub, Technische Universität München, Germany

5:00-5:20 Optimal Sequential Monte Carlo Methods in Rare Event Sampling

Robert Webber and Jonathan Weare, University of Chicago, U.S.

5:25-5:45 Estimation of Failure Probabilities via Local MCMC Subset Approximations

Kenan Šehic and Mirza Karamehmedovic, Technical University of Denmark, Denmark; Youssef M. Marzouk, Massachusetts Institute of Technology, U.S.

continued in next column

Thursday, February 28

MS321

Fast Algorithms for Integral Equations and Their Applications

4:10 p.m.-5:50 p.m.

Room: 202A

The recent advances in integral equations and its fast numerical methods have provided useful tools for many applications ranging from nano-optics to medical imaging and geosciences. This minisymposium will discuss challenges in the formulation of the problem, cutting-edge fast algorithms and their efficient implementation, their applications in various fields. At the same time, it will provide opportunities to promote interdisciplinary research collaboration between computational scientists and other fields.

Organizer: Min Hyung Cho
University of Massachusetts, Lowell, U.S.

Organizer: Carlos Borges
University of Central Florida, U.S.

4:10-4:30 Domain Decomposition Based Preconditioners for Forward and Inverse Scattering Problems

Carlos Borges, University of Central Florida, U.S.; George Biros, University of Texas at Austin, U.S.

4:35-4:55 Efficient Boundary Integral Scheme for Threshold Dynamics and its Applications

Shidong Jiang, New Jersey Institute of Technology, U.S.; Dong Wang, University of Utah, U.S.; Xiao-Ping Wang, Hong Kong University of Science and Technology, Hong Kong

5:00-5:20 Fast Multipole Method in Layered Media

Bo Wang and Wei Cai, Southern Methodist University, U.S.; Min Hyung Cho, University of Massachusetts, Lowell, U.S.; Duan Chen, University of North Carolina, Charlotte, U.S.

5:25-5:45 A Fast Algorithmic Framework for Dense Rigid Body Suspensions in Stokes Flow

Eduardo Corona, University of Michigan, U.S.

Thursday, February 28

MS322

Scalable Solvers for the Helmholtz Problem - Part II of II

4:10 p.m.-5:50 p.m.

Room: 202B

For Part 1 see MS290

Many wave phenomena can be described by the Helmholtz equation in the frequency domain approach. After discretization large linear systems have to be solved. To maintain a certain accuracy, the discretization relies on the number of grid points per wavelength, which should be as high as possible. Increasing the wavenumber, therefore leads to very large linear systems. Most solvers which are available have the property that the number of iterations increases at least linearly with the wave number. Only recently, various methods become available where the number of iterations is independent of the size of the wave number, and the work per iteration is approximately constant, so-called scalable methods. In this minisymposium various methods will be presented which are scalable when certain conditions are satisfied. Typical applications of these solvers are medical imaging, seismics, sonar etc.

Organizer: Kees Vuik

Delft University of Technology, Netherlands

Organizer: Reinhard Nabben

Institut für Mathematik, Germany

4:10-4:30 Fast and Scalable Solvers for High-frequency Wave Propagation

Matthias Taus, Massachusetts Institute of Technology, U.S.; Leonardo Zepeda-Nunez, Lawrence Berkeley National Laboratory, U.S.; *Russell Hewett*, Total E&P, U.S.; Laurent Demanet, Massachusetts Institute of Technology, U.S.

4:35-4:55 Iterative Solvers for Modeling Ocean Waves Using Helmholtz-like Equation

Yogi Erlangga and Rysbergen Tabarek, Nazarbayev University, Kazakhstan

5:00-5:20 Domain Decomposition Preconditioning for Frequency Domain Wave Problems

Ivan G. Graham, University of Bath, United Kingdom

5:25-5:45 Bemfmm: An Extreme-scale Fmm-accelerated Solver Based on Boundary Element Method for the 3D Complex Helmholtz Equation

Noha Al-Harthi, Mustafa Abduljabbar, Mohammed A. Al Farhan, and Rui Chen, King Abdullah University of Science & Technology (KAUST), Saudi Arabia; Rio Yokota, Tokyo Institute of Technology, Japan; Hakan Bagci and David E. Keyes, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

continued in next column

Thursday, February 28

MS323

Novel Adaptation, Decomposition and Expansion Algorithms for a Class of Large Scale Models

4:10 p.m.-5:50 p.m.

Room: 202C

Large scale deterministic and stochastic models on complex geometries and high stochastic dimensions demand efficient computational techniques to avoid prohibitive cost to simulate the model. Development of such techniques crucially depend on the fundamental properties of the models, such as coercivity in elliptic systems and lack of coercivity in standard formulations to understand wave propagation in unbounded and heterogeneous media. We focus on recent developments in developing efficient algorithms for such models using spatial decomposition, expansions, and basis adaptation techniques.

Organizer: Mahadevan Ganesh

Colorado School of Mines, U.S.

4:10-4:30 An Overlapping Decomposition and Heterogeneity-homogeneity Adaptation for Wave Propagation in Unbounded Media

Mahadevan Ganesh, Colorado School of Mines, U.S.; Victor Dominguez, Universidad Pública de Navarra, Spain; Francisco J. Sayas, University of Delaware, U.S.

4:35-4:55 A Robust Solver for Elliptic PDEs in 3D Complex Geometries

Abtin Rahimian, University of Colorado Boulder, U.S.; Matthew J. Morse and Denis Zorin, Courant Institute of Mathematical Sciences, New York University, U.S.

5:00-5:20 An Efficient Domain Decomposition Method with Cross-point Treatment for Helmholtz Problems

Axel Modave, CNRS, France; Xavier L. Antoine, Université de Lorraine, France; Christophe Geuzaine, University of Liege, Belgium

5:25-5:45 Basis Adaptation and Domain Decomposition for a Class of Stochastic Models

Ramakrishna Tipireddy, University of Southern California, U.S.; Panos Stinis and Alexander Tartakovsky, Pacific Northwest National Laboratory, U.S.

Thursday, February 28

MS324

Flat Low-rank Matrix Formats: Potential and Limitations

4:10 p.m.-5:50 p.m.

Room: 203

In numerous scientific applications, matrices have been shown to possess off-diagonal blocks of low numerical rank. Among the different low-rank matrix formats proposed to exploit this property, this minisymposium focuses on flat formats, that is, formats with no or little hierarchy (as opposed to hierarchical formats such as H or HSS matrices). Flat formats include block low-rank (BLR) matrices, based on a 2D blocking of the matrix. The simplicity and regularity of this BLR format make it easier to take advantage of modern, possibly even heterogeneous, architectures. BLR solvers have been shown to achieve good performance on multicore and distributed architectures; however, this comes at the price of a higher theoretical complexity compared to the nearly linear complexity of hierarchical formats. In this minisymposium, we will discuss recent advances made on the BLR format. In addition, we will also present two recently proposed new flat formats, namely multilevel BLR (MBLR) and lattice-H matrices, which aim at achieving an improved compromise between asymptotic complexity and parallel performance. We hope this MS will bring together researchers from the field to discuss both the potential and limitations of flat formats compared to hierarchical ones.

Organizer: Theo Mary

University of Manchester, United Kingdom

Organizer: Gregoire Pichon

Inria, France

4:10-4:30 On the Inversion of BLR Matrices from Acoustic Fluid-structure Interactions

Cleve Ashcraft, Francois-Henry Rouet, and Clement Weisbecker, Livermore Software Technology Corporation, U.S.

4:35-4:55 Block Low-rank Algebraic Clustering for Sparse Direct Solvers

Gregoire Pichon, Inria, France; Eric F. Darve, Stanford University, U.S.; Mathieu Faverge, Bordeaux INP, Inria, LaBRI, France; Pierre Ramet, Université de Bordeaux, Inria, LaBRI, France; Jean Roman, Inria, France

5:00-5:20 Distributed Memory Lattice H-Matrix Factorization

Ichitaro Yamazaki, University of Tennessee, Knoxville, U.S.; *Akihiro Ida*, University of Tokyo, Japan; Rio Yokota, Tokyo Institute of Technology, Japan; Jack J. Dongarra, University of Tennessee and Oak Ridge National Laboratory, U.S.

5:25-5:45 Bridging the Gap between Flat and Hierarchical Low-rank Matrix Formats

Patrick Amestoy, ENSEEIHT-IRIT, France; Alfredo Buttari, CNRS, France; Jean-Yves L'Excellent, Inria-LIP-ENS Lyon, France; Theo Mary, University of Manchester, United Kingdom

Thursday, February 28

MS325

Partitioned and Adaptive Methods for Initial Value Problems - Part II of II

4:10 p.m.-5:50 p.m.

Room: 205

For Part 1 see MS293

Numerical methods for initial value problems often adapt or partition the method to suit the problem. For differential equations that have a separation into stiff and nonstiff terms, methods such as implicit-explicit methods and integrating factor methods exploit this structure by partitioning the method based on the structure of the differential equation. These methods allow larger stable time-steps than standard explicit methods at a smaller computational cost than traditional fully implicit methods. The analysis and efficient implementation of these methods requires careful treatment of the adaptivity or partitioning along with accuracy and stability. This minisymposium brings together speakers working on the application and analysis of and partitioned methods for initial value problems.

Organizer: Andrew J. Steyer

Sandia National Laboratories, U.S.

Organizer: Sidafa Conde

Sandia National Laboratories, U.S.

4:10-4:25 Strong Stability Preserving Integrating Factor Runge-Kutta Methods

Leah Isherwood, Sigal Gottlieb, and Zachary Grant, University of Massachusetts, Dartmouth, U.S.

4:30-4:45 Explicit Partitioned Methods for Interface Coupling

Kara Peterson, Pavel Bochev, and *Paul Kuberry*, Sandia National Laboratories, U.S.

5:00-5:20 Embedded Error Estimation and Adaptive Step-size Control for Optimal Explicit Strong Stability Preserving Runge–Kutta Methods

Sidafa Conde, Sandia National Laboratories, U.S.; Imre Fekete, Eötvös Loránd University, Hungary; John Shadid, Sandia National Laboratories, U.S.

continued in next column

continued on next page

Thursday, February 28

MS325

Partitioned and Adaptive Methods for Initial Value Problems - Part II of II

continued

5:30-5:45 Strong-stability-preserving Runge-Kutta Methods with Downwind-biased Operators

Yiannis Hadjimichael, Eötvös Loránd University, Hungary

Thursday, February 28

MS326

High-order Accurate Numerical Methods for Fluid-structure Interaction Problems

4:10 p.m.-5:50 p.m.

Room: 206A

Fluid-structure interaction (FSI) problems arise in many applications such as those in material science, physics, structural engineering and biomedicine, etc. The investigation of modeling those problems and the development of numerical approximations for their solution are very active areas of research. This minisymposium aims to bring researchers together to discuss the state-of-the-art methods and algorithms for solving FSI problems, and to share the numerical challenges in the various application areas

Organizer: Longfei Li

University of Louisiana at Lafayette

Organizer: Qi Tang

Los Alamos National Laboratory, U.S.

4:10-4:30 A Partitioned FSI Algorithm for Rigid Bodies and Incompressible Flow Based on Potentials

Qi Tang, Los Alamos National Laboratory, U.S.

4:35-4:55 A Stable Added-mass Partitioned (AMP) Algorithm for Elastic Solids and Incompressible Flows

Daniel A. Serino, Rensselaer Polytechnic Institute, U.S.

5:00-5:20 A Fourth-order Accurate Fractional-step Scheme for the Incompressible Navier-stokes Equations on Overlapping Grids

Donald W. Schwendeman, Rensselaer Polytechnic Institute, U.S.

5:25-5:45 Extending the Added-mass Partitioned (AMP) Scheme for Solving FSI Problems Coupling Incompressible Flows with Elastic Beams to 3D

Longfei Li, University of Louisiana at Lafayette

Thursday, February 28

MS327

Statistical Applications of Continuous and Discrete Transport - Part II of II

4:10 p.m.-5:50 p.m.

Room: 206B

For Part 1 see MS295

Transport between probability measures offers a rich toolbox of approaches for sampling and inference, and for comparing both discrete and continuous distributions. Deterministic couplings between probability measures can be induced by transport maps that minimize certain cost functionals; more generally, one may derive many tools and insights from the construction of optimal transport plans. Recent theoretical and computational advances in this area have enabled the solution to ubiquitous statistical problems, including density estimation, variational Bayesian inference (parametric and nonparametric), inference with intractable likelihoods, as well as nonlinear filtering and smoothing. Despite these successes, efficiently solving such problems remains challenging, especially in high dimensions with strongly non-Gaussian distributions, or when only limited information or few samples are available. This minisymposium will explore new algorithms and methodologies for computing and employing transport in the context of statistics and machine learning.

Organizer: Ricardo Baptista

Massachusetts Institute of Technology, U.S.

Organizer: Daniele Bigoni

Massachusetts Institute of Technology, U.S.

Organizer: Matthew Parno

US Army Corps of Engineers, U.S.

4:10-4:30 Accurate Estimation of High-dimensional Transport Maps for Bayesian Filtering

Ricardo Baptista and Youssef M. Marzouk, Massachusetts Institute of Technology, U.S.

continued on next page

4:35-4:55 Optimal Transport on Surfaces, Graphs, and Point Clouds

Justin Solomon, Sebastian Claici, and Edward Chien, Massachusetts Institute of Technology, U.S.; Gabriel Peyré, Ceremade CNRS-Université Paris-Dauphine; Marco Cuturi, ENSAE ParisTech, France

5:00-5:20 Langevin Monte Carlo and JKO Splitting

Espen Bernton, Harvard University, U.S.

5:25-5:45 Data Driven Schroedinger Bridge

Giulio Trigila, Baruch College, CUNY, U.S.; Esteban G. Tabak, Courant Institute of Mathematical Sciences, New York University, U.S.; Michele Pavon, University of Padova, Italy

Thursday, February 28

MS328

Integrated Mathematical and Computational Approaches to Interface Models and Data in Systems Biomedicine

4:10 p.m.-5:50 p.m.

Room: 206C

Complex biological systems comprise functional and structural networks interacting at different levels and scales i.e. metabolite and enzymatic networks, regulatory and co-expression networks, protein-protein interaction networks as well as signal transduction pathways that exhibit strong emergent properties. Diseases are caused when these networks are functionally or topologically perturbed due to a wide range of intrinsic and environmental cues. Integrated mathematical and computational approaches as they help decipher the interplay of these networks providing a comprehensive quantitative knowledge of the underlying functional mechanisms. The state-of-the art methods in computational systems biology allow for (a) bottom-up (microscale: from nodes to network) and (b) top-down (mesoscale: from network to nodes) modeling approaches. A bottom-up approach includes robust and computationally efficient graph-theoretic methods and functional genomics, but may not facilitate description of global system properties. On the other hand, the top-down approach uses deterministic and stochastic mathematical models coupled with parameter estimation techniques generally overcome the limitations of the bottom-up approach may be limited by model under-/over-fitting. The goal of this minisymposium is to bring together researchers working in computational modelling in order to showcase recent theories and tools in the overarching theme of these approaches in systems biology research.

Organizer: Jeyashree Krishnan
RWTH Aachen University, Germany

Organizer: Dimitris I. Kalogiros
University of Nottingham, United Kingdom

4:10-4:30 Modeling of Disease Progression in Cancer

Andreas Schuppert, RWTH Aachen, Germany

4:35-4:55 Healthcare by Design: How Computational Approaches Are Changing the Clinical Landscape

Himanshu Kaul, University of British Columbia, Canada; Yiannis Ventikos, University College London, United Kingdom; Rod Smallwood, University of Sheffield, United Kingdom; Chris Brightling, University of Leicester, United Kingdom

5:00-5:20 Combining in vitro Data and Mathematical Models to Understand Chemokine Biology: A Bayesian Parameter Inference Approach

Dimitris I. Kalogiros and Matthew Russell, University of Nottingham, United Kingdom; Willy Bonneuil and Jennifer Frattolin, Imperial College London, United Kingdom; Mikaila J. Bandara, Francesca F. Masci, and Robert Nibbs, University of Glasgow, Scotland, UK; James E. Moore Jr., Imperial College London, United Kingdom; Bindi S. Brook, University of Nottingham, United Kingdom

5:25-5:45 A Statistical Mechanical Perspective of Modeling Phase Transitions in Living Systems

Jeyashree Krishnan, RWTH Aachen University, Germany; Reza Torabi, University of Calgary, Canada; Edoardo Di Napoli and Andreas Schuppert, RWTH Aachen, Germany

continued in next column

Thursday, February 28

MS329

Recent Advances in Discontinuous Galerkin Methods for Partial Differential Equations

4:10 p.m.-5:50 p.m.

Room: 206D

This minisymposium provides scholars a platform to share and exchange their recent works and progresses on discontinuous Galerkin methods for partial differential equations. In this minisymposium, speakers will introduce the backgrounds and applications of several important PDE models in fluid dynamics and traffic flow theory. Speakers will show their ideas in designing DG schemes for those models as well as the numerical analysis and implementations. This minisymposium is supposed to provide a good opportunity for people to communicate with researchers from different concentrations and application fields and find potential academic cooperators for future. There will be 4 speakers to present their recent advances in this minisymposium.

Organizer: Ziyao Xu

Michigan Technological University, U.S.

4:10-4:30 High-order Bound-preserving Discontinuous Galerkin Method for Wormhole Propagation Model

Ziyao Xu, Michigan Technological University, U.S.

4:35-4:55 High-order Discontinuous Galerkin Schemes for Water Mass Transport on Curved Surface

Caleb D. Logemann, Iowa State University, U.S.

5:00-5:20 A Discontinuous Galerkin Method for Gradient Flows

Zheng Sun, Ohio State University, U.S.; Jose Carrillo, Imperial College London, United Kingdom; Chi-Wang Shu, Brown University, U.S.

5:25-5:45 A Discontinuous Galerkin Method for the Aw-rascle Traffic Flow Model on Networks

Joshua Buli, University of California, Riverside, U.S.; Yulong Xing, Ohio State University, U.S.

Thursday, February 28

MS330

Wave-based Imaging Meets Machine Learning

4:10 p.m.-5:50 p.m.

Room: 207

Developments in deep convolution neural networks (DCNNs) are driving exiting new approaches to problems in computational imaging. These include new ways to include learned prior knowledge via denoisers, auto-encoders, possibly in combination with unrolling of optimization schemes that are driven by local derivative information derived from a data objectives and physics-based forward models. While these approaches have resulted in major improvements in certain areas of computational imaging, these gains often rely on extensive training and on the availability of accurate (nonlinear) forward models, their derivatives and adjoints. During this mini symposium, we would like to explore possible ways forward to scale this technology and to make it less reliant on high-fidelity physics and/or numerics of the forward model.

Organizer: Tristan van Leeuwen

Utrecht University, The Netherlands

Organizer: Felix Herrmann

Georgia Institute of Technology, U.S.

4:10-4:30 Machine Learning in Seismic Imaging: from Low-fidelity to High-fidelity

Felix Herrmann, Ali Siahkoohi, and Mathias Louboutin, Georgia Institute of Technology, U.S.

4:35-4:55 Imaging for Inverse Scattering Using Learned Preconditioners

Hassan Mansour, Mitsubishi Electric Research Laboratories, U.S.; Ajinkya Kadu, Utrecht University, The Netherlands

5:00-5:20 Scaling Limits of Current Machine Learning Frameworks and How to Extend Them

Gerard J. Gorman, Imperial College London, United Kingdom

5:25-5:45 A Kernel Method for Full Waveform Inversion

Tristan van Leeuwen, Utrecht University, The Netherlands

Thursday, February 28

MS331

Recent Innovations in Restarting and Recycling Krylov Methods

4:10 p.m.-5:50 p.m.

Room: 401A

Scientists and engineers need to solve continually more difficult problems that require better and faster solvers for linear systems and eigenvalue problems. While Krylov subspace methods are simple in concept, they are often complicated in practice. In this minisymposium we highlight several important problems and recent innovations. We will focus mainly on the techniques of subspace recycling and polynomial preconditioning.

Organizer: Ronald Morgan

Baylor University, U.S.

Organizer: Eric De Sturler

Virginia Tech, U.S.

4:10-4:30 Using Krylov Subspace Information for Updating Preconditioners

Eric De Sturler, Virginia Tech, U.S.

4:35-4:55 Combining Preconditioner Updates with Krylov Subspace Recycling for Sequences of Linear Systems

Arielle K. Carr, Virginia Tech, U.S.

5:00-5:20 Polynomial Preconditioned GMRES to Reduce Communication in Parallel Computing

Jennifer A. Loe, Baylor University, U.S.; Heidi K. Thornquist and Erik G. Boman, Sandia National Laboratories, U.S.

5:25-5:45 Krylov Methods for Rank-one Updates of Linear Equations and Eigenvalues

Ron Morgan, Baylor University, U.S.; Mark Embree, Virginia Tech, U.S.

Thursday, February 28

MS332

SIMD Approaches for Achieving Performance and Portability on Emerging Computational Architectures

4:10 p.m.-5:50 p.m.

Room:401B

Single-Instruction Multiple-Data (SIMD) is a common technique being used by modern computer architectures to achieve high performance in a power-efficient manner by increasing the availability of fine-grained parallelism. Two notable examples are the use of vector instructions on modern CPU architectures and warps of concurrent threads on GPU architectures. However harnessing this fine-grained parallelism in scientific computing contexts can be quite challenging due to code complexity, sparsity and indirection. In this minisymposium, we explore approaches for leveraging fine-grained SIMD parallelism through the use of modern programming techniques such as operator overloading, discuss implications for performance and portability, and present applications of these ideas within several scientific computing contexts.

Organizer: Eric Phipps

Sandia National Laboratories, U.S.

Organizer: Siva Rajamanickam

Sandia National Laboratories, U.S.

4:10-4:30 Performance Portable SIMD Approach Implementing Block Line Solver For Coupled PDEs

Kyungjoo Kim, Sandia National Laboratories, U.S.

4:35-4:55 Ensemble Propagation for Efficient Uncertainty Quantification of Mechanical Contact Problems

Kim Liegeois and Romain Boman, Université de Liège, Belgium; Eric Phipps, Sandia National Laboratories, U.S.; Maarten Arnst, Université de Liège, Belgium

5:00-5:20 A Performance Portable SIMD Scalar Type for Effective Vectorization Across Heterogeneous Architectures

Damodar Sahasrabudhe, University of Utah, U.S.; Eric Phipps and Siva Rajamanickam, Sandia National Laboratories, U.S.; Martin Berzins, University of Utah, U.S.

5:25-5:45 Intel® Vector Math Libraries and Their Features, Usage, and Performance Improvements

Marius Cornea, Intel Corporation, U.S.

continued in next column

Thursday, February 28

MS333

Theoretical and Computational Aspects in Nonlocal and Material Science Modeling - Part II of II

4:10 p.m.-5:50 p.m.

Room: 401C

For Part 1 see MS201

The mathematics of material science encompasses a variety of aspects, from formulating models that describe material properties and phenomena, to theoretical investigations, and numerical simulations. A variety of models describing local and nonlocal effects have been shown successful in applications, their investigation and implementation in real-world applications is ongoing attracting the interest of a growing community. Nonlocal models such as anomalous diffusion and peridynamics can incorporate temporal and spatial nonlocal behaviors due to memory effects and long-range forces. For problems where these effects cannot be neglected, such descriptions are more accurate than PDEs and have recently gained popularity. Mathematical analysis and numerical solutions for nonlocal problems remain challenges, requiring new approaches to design robust, efficient and predictive numerical methods. The goal of this minisymposium is to bring together researchers working on different applications in material science to present and discuss novel approaches related to theoretic nonlocal models and present approaches for their numerical solution.

Organizer: Marta D'Elia

Sandia National Laboratories, U.S.

Organizer: Petronela Radu

University of Nebraska, Lincoln, U.S.

4:10-4:30 Boundary Conditions for Tempered Fractional Diffusion

Anna Lischke, Brown University, U.S.; James F. Kelly and Mark Meerschaert, Michigan State University, U.S.

continued on next page

Thursday, February 28

MS333

Theoretical and Computational Aspects in Nonlocal and Material Science Modeling - Part II of II

continued

4:35-4:55 Collocation Methods for the Integral Fractional Laplacian with Applications

Yanzhi Zhang, Missouri University of Science and Technology, U.S.

5:00-5:20 A Splitting Method for the Fast Computation of Nonlocal Diffusion and Peridynamics Models in Heterogeneous Media

Xiaochuan Tian, University of Texas at Austin, U.S.

5:25-5:45 Surface Energy Approach in Modeling of Nano-objects

Anna Zemlyanova, Kansas State University, U.S.

Thursday, February 28

MS334

Using Structure for Scalable Optimization

4:10 p.m.-5:50 p.m.

Room: 402A

The recent explosion of big-data applications has made it hard for traditional convex optimization algorithms to scale. However, in many real-world applications, problem structure significantly reduces the space of plausible solutions. In order to find these structurally-ideal solutions, often we employ convex regularizers, either explicitly or implicitly, for structured sparsity, such as low-rank or group-sparse applications. This minisymposium explores different ways in which incorporating problem structure may additionally facilitate the scalability of these algorithms. Examples include: designing methods with reduced per-iteration cost for structured optimization problems, evaluating the effectiveness of quick-and-dirty solutions on real-world problems, and improving the interpretability of favored heuristic approaches.

Organizer: Yifan Sun

University of British Columbia, Canada

Organizer: Michael P. Friedlander

University of British Columbia, Canada

4:10-4:30 Fast Methods for Low-rank Semidefinite Optimization

Yifan Sun, University of British Columbia, Canada

4:35-4:55 Certifying Accuracy and Uniqueness in Ill-conditioned Imaging Problems

Stephen Becker and James Folberth, University of Colorado Boulder, U.S.

5:00-5:20 Storage-optimal Convex Low Rank Semidefinite Programming

Madeleine R. Udell, Cornell University, U.S.

5:25-5:45 Manipulating Neural Networks with Inverse Problems

Tom Goldstein, University of Maryland, U.S.

Thursday, February 28

MS335

Black-box Optimization for Autotuning

4:10 p.m.-5:50 p.m.

Room: 402B

The challenges related to the optimization of modern applications targeting current and future computer architectures raise the need for systematic methods to handle this need, where humans are left out of the equation. Of particular interest to solve this problem is the use of autotuning. Given a problem and a parametrized computer program to solve it, the purpose of autotuning is to find the values of the parameters that will yield the best outcome of the program on the problem (computation time, memory or energy consumption, accuracy of the results, ...). In this minisymposium, the presentations will highlight new methodologies from the fields of black-box optimization and of of machine learning to target the autotuning needs. Several application examples will be targeting linear algebra related topics.

Organizer: Wissam M. Sid-Lakhdar

Lawrence Berkeley National Laboratory, U.S.

4:10-4:30 On the High Performance Implementation of Quaternionic Matrix Operations

David Williams-Young, Lawrence Berkeley National Laboratory, U.S.

4:35-4:55 From Matrix Multiplication to Deep Learning: Are We Tuning the Same Kernel?

Yaohung Tsai, University of Tennessee, Knoxville, U.S.

5:00-5:20 Asynchronous Model-based-search for Automatic Performance Tuning.

Paul D. Hovland, Argonne National Laboratory, U.S.

5:25-5:45 MultiVerse-Kriging: A Transfer and Multitask-learning Based Approach to Autotuning

Wissam M. Sid-Lakhdar, Lawrence Berkeley National Laboratory, U.S.

Thursday, February 28

MS336

High-performance Graph Algorithms

4:10 p.m.-5:50 p.m.

Room: 402C

Graph algorithms are critical for understanding and manipulating large sets of relational information from engineering, science, mathematics, and data analytics applications. Efficient execution of these procedures requires mapping novel spatial and temporal graph representations to the available programming models and hardware systems. This minisymposium continues the CSE17 symposium of the same name by bringing together researchers engaged in the development and application of high-performance parallel graph algorithms to discuss recent advances, challenges, and opportunities for collaboration.

Organizer: Cameron Smith
Rensselaer Polytechnic Institute, U.S.

Organizer: George M. Slota
Rensselaer Polytechnic Institute, U.S.

4:10-4:30 Diffusive Load Balancing of Hypergraphs for Partitioning Evolving Unstructured Mesh Applications
Gerrett Diamond and Mark S. Shephard, Rensselaer Polytechnic Institute, U.S.

4:35-4:55 New Models for Streaming Graphs
David A. Bader, Georgia Institute of Technology, U.S.

5:00-5:20 Local Algorithms for Hierarchical Dense Subgraph Discovery
A. Erdem Sariyuce, University at Buffalo, U.S.; Seshadhri C, University of California, Santa Cruz, U.S.; Ali Pinar, Sandia National Laboratories, U.S.

5:25-5:45 Fast Linear Algebra-based Triangle Analytics with Kokkos Kernels
Abdurrahman Yasar, Georgia Institute of Technology, U.S.; Michael Wolf and Jonathan W. Berry, Sandia National Laboratories, U.S.; Umit V. Catalyurek, Georgia Institute of Technology, U.S.

Thursday, February 28

CP18

Numerical PDEs II

4:10 p.m.-5:50 p.m.

Room: 102C

Chair: Timothy J. Barth, NASA Ames Research Center, U.S.

4:10-4:25 Hyperbolic Method for Incompressible Navier-Stokes Equations on Unstructured Grids
Hyung T. Ahn, University of Ulsan, South Korea

4:30-4:45 On the Calculation of Exact Cumulative Distribution Statistics for Burgers Equation
Timothy J. Barth, NASA Ames Research Center, U.S.; Jonas Sukys, Swiss Federal Institute of Aquatic Science and Technology, Switzerland

4:50-5:05 t Going to Infinity Limit of Solutions for Linear Advection Problems
Abhijit Biswas, Temple University, U.S.

5:10-5:25 Discontinuous Galerkin Methods for the Nonlinear Shallow Water Equations with Horizontal Temperature Gradients
Jolene Britton, University of California, Riverside, U.S.; Yulong Xing, Ohio State University, U.S.

5:30-5:45 Ssp and Positivity for Implicit Methods
Zoltan Horvath, Széchenyi István University, Györ, Hungary

Friday, March 1

Registration

8:00 a.m.-12:00 p.m.

Room: Ballroom Foyer

IP8

Role of Tensors in Machine Learning

8:30 a.m.-9:15 a.m.

Room:Ballroom 100BC

Chair: Rio Yokota, Tokyo Institute of Technology, Japan

Tensors are higher order extensions of matrices that can incorporate multiple modalities and encode higher-order relationships in data. After an introduction to tensor methods, I will present ways in which tensor methods can be used in deep learning, as well as in probabilistic modeling. I will show that tensor contractions, which are extensions of matrix products, provide high rates of compression in a variety of neural network models. I will also demonstrate the use of tensors for document categorization at scale through probabilistic topic models. These are available in a python library called Tensorly that provides a high-level API for tensor methods and deep tensorized architectures.

Anima Anandkumar
Amazon and California Institute of Technology, U.S.

Coffee Break

9:15 a.m.-9:45 a.m.

Room:Ballroom Foyer

Friday, March 1

MT4

Firedrake: Automated High Performance Finite Element Simulation

9:45 a.m.-11:25 a.m.

Room:300D

Firedrake is an automated simulation system. Users specify finite element problems mathematically using the Unified Form Language (UFL) embedded in Python. High performance parallel operator and residual assembly is automatically generated using advanced compiler technology. Firedrake integrates with the PETSc framework to provide a full suite of sophisticated linear and nonlinear solvers. In this hands-on Jupyter-based tutorial, you will have the chance to solve linear and nonlinear PDEs using Firedrake and try out some of its advanced features, including:
• Composable Schur complement, multilevel, and operator-based preconditioners
• Automated solution of time-dependend adjoint PDEs using dolfin-adjoint
• Hybridised DG and hybridised mixed methods.
• Automated exploitation of vectorised hardware.

David Ham, Imperial College London, United Kingdom
Lawrence Mitchell, Durham University, United Kingdom
Thomas H. Gibson, Imperial College London, United Kingdom
Tianjiao Sun, Imperial College London, United Kingdom

Friday, March 1

MS337

Tensor Decomposition for High Performance Data Analytics - Part I of II

9:45 a.m.-11:25 a.m.

Room: Ballroom 100BC

For Part 2 see MS370

Tensors are higher order generalization of matrices that map naturally to complex relationships found in the real world. In this regard, tensor decomposition has gained popularity in data analytics due its ability to identify hidden properties in large quantities of data. The talks in this minisymposium will explore state-of-the-art research into finding efficient and scalable solutions to decomposing tensors for data analysis on the latest processors and accelerators.

Organizer: Jee Choi

IBM T.J. Watson Research Center, U.S.

9:45-10:05 On Optimizing Distributed Tucker Decomposition for Sparse Tensors

Jee Choi, IBM T.J. Watson Research Center, U.S.

10:10-10:30 Tensor Computations: Efficiency Or Productivity?

Paolo Bientinesi, RWTH Aachen University, Germany

10:35-10:55 Accelerating Alternating Least Squares for Tensor Decomposition by Pairwise Perturbation

Edgar Solomonik, University of Illinois at Urbana-Champaign, U.S.; Linjian Ma, University of California, Berkeley, U.S.

11:00-11:20 HPC Formulations of Optimization Algorithms for Tensor Completion

Shaden Smith, University of Minnesota, U.S.

Friday, March 1

MS338

Software Productivity and Sustainability for CSE and Data Science - Part I of II

9:45 a.m.-11:25 a.m.

Room: Conference Theater

For Part 2 see MS371

Software is the key crosscutting technology that enables advances in mathematics, computer science, and domain-specific science and engineering to achieve robust simulations and analysis for predictive science, engineering, and other research fields. While software is becoming more complex due to multiphysics and multiscale modeling, the coupling of data analytics, and disruptive changes in computer hardware (due to increases in typical system scale and heterogeneity, including GPUs and additional alternative architectures), software itself has not traditionally received focused attention in the CSE community or been rewarded by that community. The presenters in this minisymposterium will address work that addresses growing technical and social challenges in software productivity, quality, and sustainability, and thereby helps software fulfill its critical role as a cornerstone of long-term CSE collaboration. Having a minisymposterium for these topics provides a natural gathering point during poster sessions for informal conversation.

Organizer: David E. Bernholdt

Oak Ridge National Laboratory, U.S.

Organizer: Daniel S. Katz

University of Illinois at Urbana-Champaign, U.S.

Organizer: Catherine Jones

Science and Technology Facilities Council, United Kingdom

9:45-10:05 Enabling a Culture of Developer Productivity and Software Sustainability

Elaine M. Raybourn, Sandia National Laboratories, U.S.

continued on next page

10:10-10:30 Planning to Make Research Software More Sustainable via a US Research Software Sustainability Institute (URSSI)
Daniel S. Katz, University of Illinois at Urbana-Champaign, U.S.; Jeffrey C. Carver, University of Alabama, U.S.; Sandra Gesing, University of Notre Dame, U.S.; Karthik Ram, University of California, Berkeley, U.S.; Nicholas Weber, University of Washington, U.S.

10:35-10:55 CIG Perspectives on 14 Years of Sustaining Software and their Communities in Geodynamics
Lorraine J. Hwang and Louise Kellogg, University of California, Davis, U.S.

11:00-11:20 Is It a Project Or a Business? Perspectives on the Consideration of Sustainability
Michael Zentner, Purdue University, U.S.

Friday, March 1

MS339

Recent Development of Numerical Methods for Optics and Plasmonics - Part I of II

9:45 a.m.-11:25 a.m.

Room: 102A

For Part 2 see MS372

The main purpose of this minisymposium is to bring together specialists in the fields of plasmonics and optics. Driven by spectacular advances in the design capabilities of materials at the nanoscale, there has been recent exponential growth in the fields of plasmonics and nano-optics. In particular, algorithm design, analysis, application and implementations of mathematical models in the area are now at the heart everyday technologies such as seismic imaging, underwater acoustics, and biomathematics; nevertheless some very important issues still remain open. This minisymposium is meant to be a platform to exchange ideas on these problems.

Organizer: YoungJoon Hong
San Diego State University, U.S.

Organizer: David P. Nicholls
University of Illinois, Chicago, U.S.

9:45-10:05 A High-order Perturbation of Surfaces Method for Vector Electromagnetic Scattering by Doubly Layered Periodic Crossed Gratings
David P. Nicholls, University of Illinois, Chicago, U.S.

10:10-10:30 Scattering by a Periodic Array of Subwavelength Slits I: Field Enhancement in the Diffraction Regime
Junshan Lin, Auburn University, U.S.

10:35-10:55 Numerical Methods for Wave Scattering from Multilayered Media
Min Hyung Cho, University of Massachusetts, Lowell, U.S.

11:00-11:20 Boundary Integral Equation Methods for Metasurface Scattering
Carlos Perez-Arancibia, Massachusetts Institute of Technology, U.S.; Raphael Pestourie, Harvard University, U.S.; Steven Johnson, Massachusetts Institute of Technology, U.S.

Friday, March 1

MS340

Neutrino Transport Methods in Astrophysics - Part I of II

9:45 a.m.-11:25 a.m.

Room: 102B

For Part 2 see MS373

Core-collapse supernovae (CCSNe) and compact binary mergers (CBMs) are cataclysmic astrophysical events responsible for heavy element synthesis and the emission of photon, neutrino, and gravitational wave (GW) signals. They have long been targets of instruments covering most of the electromagnetic (EM) spectrum, and more recently of neutrino and GW detectors. The recent detection of a merging binary neutron star pair in EM and GW channels marked a breakthrough for multi-messenger astronomy. Harvesting insights into the physical processes driving these events from the wealth of data relies heavily on sophisticated models requiring extreme-scale, high-fidelity computing. These models solve a coupled system of equations for self-gravity, magneto-hydrodynamics, and neutrino transport. In CCSNe and CBMs, lepton and four-momentum exchange between neutrinos and matter plays a major role in the dynamics. However, neutrino-matter interactions occur under non-equilibrium conditions and a kinetic description based on the Boltzmann transport equation is warranted. As a six-dimensional phase-space problem, our ability to model these astrophysical events with satisfactory realism relies on advances in multi-physics and multi-scale algorithms, novel discretization techniques, fast solvers, and sustainable software. The goal of this minisymposium is to bring together researchers working on topics of relevance to neutrino transport modeling to discuss recent work and exchange ideas.

Organizer: Eirik Endeve
Oak Ridge National Laboratory, U.S.

Organizer: Christian Cardall
Oak Ridge National Laboratory, U.S.

continued on next page

Friday, March 1

MS340

Neutrino Transport Methods in Astrophysics - Part I of II

9:45 a.m.-11:25 a.m.

continued

Organizer: Reuben Budiardja
University of Tennessee, U.S.

9:45-10:05 Neutrino Transport with Discontinuous Galerkin Methods
Eirik Endeve, Oak Ridge National Laboratory, U.S.

10:10-10:30 Numerical Methods of Fully Special and General Relativistic Boltzmann Neutrino Transport
Hiroki Nagakura, Princeton University, U.S.

10:35-10:55 Three-dimensional Monte Carlo Neutrino Transport in Neutron Star Mergers
Sherwood Richers, North Carolina State University, U.S.

Friday, March 1

MS341

Stochastic Modeling and Algorithms for Complex Physical Systems - Part I of II

9:45 a.m.-11:25 a.m.

Room: 102C

For Part 2 see MS374

This minisymposium focuses on the fundamental problem of how to approximate solutions of both forward and inverse stochastic systems. Predicting the behavior of complex physical processes relies on constructing and solving stochastic systems in high dimensional spaces, particularly in the case when the input data (coefficients, forcing terms, initial and boundary conditions, geometry) are affected by large amounts of uncertainties. We will highlight recent advances in theory and algorithms for stochastic systems, as well as their applications in uncertainty quantification, machine learning, computational fluid dynamics, statistical mechanics, etc. This minisymposium will bring together researchers from across the applied and computational mathematics communities to discuss recent advances, identify future direction, and promote new collaborations on stochastic methods.

Organizer: Miroslav Stoyanov
Oak Ridge National Laboratory, U.S.

Organizer: Tao Zhou
Chinese Academy of Sciences, China

9:45-10:05 Hypocoercivity Based Sensitivity and Numerical Analyses for Kinetic Equations with Uncertainty
Shi Jin, Shanghai Jiao Tong University, China, and the University of Wisconsin, U.S.

10:10-10:30 A Multi-fidelity Model Reduction Based on Variable-separation for Uncertainty Quantification
Lijian Jiang, Tongji University, China

10:35-10:55 Efficient Minimum Action Methods for Transition Path Computing in Noise-driven Non-gradient Systems
Haijun Yu, Institute of Computational Mathematics, China

11:00-11:20 Sparse Approximation of Data-driven Polynomial Chaos Expansions and their Applications in UQ
Ling Guo and Yongle Liu, Shanghai Normal University, China; Tao Zhou, Chinese Academy of Sciences, China

Friday, March 1

MS342

Low Mach Number AMR Combustion Simulations with PeleLM - Part I of II

9:45 a.m.-11:25 a.m.

Room: 102D

For Part 2 see MS375

PeleLM is a code for evolving chemically reacting low Mach number flows with block-structured adaptive mesh refinement (AMR). The code features an iterative time step that efficiently couples together advection, diffusion and chemical reactions, each potentially evolving with different time scales, all while evolving the numerically conservative discretization across multiple levels of mesh refinement on a manifold where the equation of state is always satified. PeleLM is being used to study the details of turbulence-chemistry interaction in both premixed and diffusion flame configurations, from small-scale laboratory experiments to large-scale pool fires and wildfires, while it is being simultaneously improved and extended to incorporate new physics, and to run efficiently on emerging computing hardware platforms. In this set of talks, we outline the current capabilities of PeleLM via several different combustion-specific applications, and discuss improvements to the code that are under development and supported by DOE's Exascale Computing Project. These improvements include arbitrary (CAD-based) 3D geometries, electric field effects on charged particles, and incorporation of real gas and closed-chamber pressurization effects. We also discuss further extensions to the PeleLM time stepping scheme to support high-order methods, and implementation strategies used in PeleLM for many-core and GPU-based large-scale computing platforms.

Organizer: Marcus Day
Lawrence Berkeley National Laboratory, U.S.

continued on next page

9:45-10:05 PeleImOverview: Existing and in-Progress Features and Plans for AMR Low Mach Number Combustion Simulation on the Desktop and at the Exascale

Marcus Day, Lawrence Berkeley National Laboratory, U.S.

10:10-10:30 Direct Numerical Simulation of Mixture Formation and Ignition under Multi-injection Diesel Relevant Conditions

Martin Rieth, Sandia National Laboratories, U.S.; Marcus Day, Lawrence Berkeley National Laboratory, U.S.; Jacqueline Chen, Sandia National Laboratories, U.S.

10:35-10:55 Turbulence-chemistry Interactions for Lean Premixed Flames

Debolina Dasgupta, Georgia Institute of Technology, U.S.; Marcus Day, Lawrence Berkeley National Laboratory, U.S.; Andrew Aspden, Newcastle University, United Kingdom; Timothy Lieuwen, Georgia Institute of Technology, U.S.

11:00-11:20 A Numerical Strategy for Unsteady Simulation of Low Mach Number Reacting Flows Subject to Electric Fields

Lucas Esclapez, Valentina Ricchiuti, and Marcus Day, Lawrence Berkeley National Laboratory, U.S.

Friday, March 1

MS343

Data-augmented Reduced-order Modeling: Operator Learning and Closure/error Modeling - Part I of II

9:45 a.m.-11:25 a.m.

Room: 111A

For Part 2 see MS376

As the availability of experimental and computational data has exploded in recent years, an important question has arisen in computational science and engineering: How can these data be exploited to improve the predictive capacity of dynamical system models? This minisymposium focuses on reduced-order models (ROMs) of dynamical systems, which are attractive due to their low simulation cost. In this setting, we explore three emerging classes of techniques that leverage the availability of experimental and computational data to construct and improve ROMs: (1) methods that learn low-dimensional dynamical-system operators from data (e.g., dynamic mode decomposition, operator inference), and (2) closure and error-modeling techniques (e.g., Mori–Zwanzig) for ROMs.

Organizer: Kevin T. Carlberg
Sandia National Laboratories, U.S.

Organizer: Boris Kramer
Massachusetts Institute of Technology, U.S.

9:45-10:05 Data-driven Dynamical Systems Modeling

Serkan Gugercin, Virginia Tech, U.S.

10:10-10:30 Operator Inference on Manifolds for Learning Physically Consistent Models from Data

Nihar Sawant, New York University, U.S.; Benjamin Peherstorfer, Courant Institute of Mathematical Sciences, New York University, U.S.

10:35-10:55 Data-driven Linear Parameter Varying Control of Wind Farms

Jennifer Annoni, National Renewable Energy Laboratory, U.S.

11:00-11:20 Nonintrusive Nonlinear Model Reduction via Machine-learning Approximations to Low-dimensional Operators

Zhe Bai, University of Washington, U.S.; Kevin T. Carlberg and Liqian Peng, Sandia National Laboratories, U.S.; Steven Brunton, University of Washington, U.S.

Friday, March 1

MS344

Inverse Problems in Machine Learning - Part I of II

9:45 a.m.-11:25 a.m.

Room: 111B

For Part 2 see MS377

Machine learning has been a fast growing area of research since the beginning of the 21st century due to the advent of computing resources and availability of data. Many of the questions of interest in machine learning can be formulated as inverse problems and vice versa. Therefore, it is not surprising that the developments in these two seemingly separate fields are in fact closely related. In this symposium we bring together researchers working at the intersection of machine learning and inverse problems to discuss recent advances, challenges and strategies in each area, aiming to promote cross-pollination and exchange of ideas that lead to new directions of research in both communities.

Organizer: Bamdad Hosseini
California Institute of Technology, U.S.

9:45-10:05 On the Information Capacity of Neural Networks

Gerald Friedland, University of California, Berkeley and Lawrence Livermore National Laboratory, U.S.

10:10-10:30 Bayesian Inference Using Adaptive Gaussian Process Models with Application to Cosmology

Timur Takhtaganov, Lawrence Livermore National Laboratory, U.S.; Zarija Lukic, Juliane Mueller, and Dmitriy Morozov, Lawrence Berkeley National Laboratory, U.S.

10:35-10:55 Learning-based Predictive Models: A New Approach to Integrating Large-scale Simulations and Experiments

Brian Spears, Lawrence Livermore National Laboratory, U.S.

11:00-11:20 Algorithmic Regularization for Machine Learning

Luigi Carratino, Universita' degli Studi di Genova, Italy; Alessandro Rudi, Inria, France; Lorenzo Rosasco, Instituto Italiano di Technologia, Italy and Massachusetts Institute of Technology, U.S.

Friday, March 1

MS345

Optimal Experimental Design for Inverse Problems - Part I of II

9:45 a.m.-11:25 a.m.

Room: 111C

For Part 2 see MS378

Good experimental data is critical for parameter inference in complex physical systems and enables subsequent robust decision making. The collection of useful experimental data, however, is arduous and expensive. It is therefore imperative to design experiments in an optimal manner, subject to physical or budgetary constraints while incorporating prior information. Leveraging simulation models, the computational cost of model-based optimal experimental design (OED) can be prohibitively expensive in many applications of interest. This minisymposium explores advances in numerical methods for OED of Inverse Problems.

Organizer: Xun Huan

University of Michigan, U.S.

Organizer: Jayanth Jagalur Mohan

Massachusetts Institute of Technology, U.S.

9:45-10:05 Large-scale Optimal Experimental Design for Bayesian Nonlinear Inverse Problems

Peng Chen and *Omar Ghattas*, University of Texas at Austin, U.S.; Umberto Villa, Washington University, St. Louis, U.S.; Keyi Wu, University of Texas at Austin, U.S.

10:10-10:30 A Sequential Greedy MM Framework for Optimal Experimental Design with Non-submodular Objectives

Jayanth Jagalur Mohan and Youssef M. Marzouk, Massachusetts Institute of Technology, U.S.

10:35-10:55 Efficient A-Optimal Experimental Design for Bayesian Linear Inverse Problems Using Randomized Estimators and Reweighted ℓ_1 Minimization

Elizabeth Herman, Alen Alexanderian, and Arvind Saibaba, North Carolina State University, U.S.

11:00-11:20 Optimal Experimental Design for Constraint Inverse Problems

Matthias Chung, Virginia Tech, U.S.; Lars Ruthotto, Emory University, U.S.; Julianne Chung, Virginia Tech, U.S.

continued in next column

Friday, March 1

MS346

Derivative-free and Global Optimization - Part I of II

9:45 a.m.-11:25 a.m.

Room: 300A

For Part 2 see MS379

Optimization problems arising in various science applications, such as engineering design, environmental applications, physics simulations, etc. require the development of new numerical optimization methods that efficiently and effectively find the optimal solutions. The objective functions are often multi-modal, computationally expensive, and black-box (their analytic descriptions and derivatives are not available). Often, the presence of stochasticity further increases the difficulty of solving these problems. New numerical optimization methods that do not rely on gradient information of the objective function and that are able to globally explore as well as locally search the parameter domain are therefore needed. The speakers in this minisymposium will present their recent developments of optimization methods that employ, among others, adaptive sampling methods, approximations of the objective functions, and the exploitation of any available problem structure to solve these difficult problems.

Organizer: Juliane Mueller

Lawrence Berkeley National Laboratory, U.S.

Organizer: Matt Menickelly

Argonne National Laboratory, U.S.

9:45-10:05 Derivative-free Optimization of Computationally Expensive Functions with Noisy Data

Juliane Mueller, Lawrence Berkeley National Laboratory, U.S.

10:10-10:30 Manifold Sampling for Optimization of Composite Nonsmooth Functions

Jeffrey Larson, Matt Menickelly, and Stefan Wild, Argonne National Laboratory, U.S.

10:35-10:55 Quantile Adaptive Search and Partition-based Algorithms for Global Optimization

Zelda B. Zabinsky and David Linz, University of Washington, U.S.

continued on next page

11:00-11:20 A Derivative-free Approach for Complex Problems

Danny D'Agostino, University of Rome, Italy; Matteo Diez, CNR-INSEAN, Italy; Giampaolo Liuzzi, CNR, Italy; Stefano Lucidi, Università di Roma "La Sapienza", Italy; Riccardo Pellegrini, CNR-INSEAN, National Research Council-Marine Technology Research Institute, Italy; *Francesco Rinaldi*, University of Padova, Italy; Andrea Serani, CNR-INSEAN, National Research Council-Marine Technology Research Institute, Italy

Friday, March 1

MS347

Advances in Global Sensitivity Analysis with Applications to Complex Systems - Part I of II

9:45 a.m.-11:25 a.m.

Room: 300B

For Part 2 see MS380

The traditional framework of global sensitivity analysis (GSA) considers scalar-valued quantities of interest (QoIs) that are functions of statistically independent inputs. The sensitivity of the QoIs to the inputs are determined using a number of classical methods such as Sobol' indices, Morris screening, or derivative based global sensitivity measures. To extend the applicability of GSA methods to broader classes of problems, such as problems with correlated inputs or problems with vectorial and functional outputs with high-dimensional parameters, new methods are needed. This minisymposium will highlight some of the recent advances in GSA and present recent generalizations and alternatives. Areas of particular interest are problems with dependent inputs, problems with vectorial or functional outputs, and applications of GSA in complex physical systems.

Organizer: Alen Alexanderian
North Carolina State University, U.S.

Organizer: Helen Cleaves
North Carolina State University, U.S.

Organizer: Pierre Gremaud
North Carolina State University, U.S.

9:45-10:05 Derivative Based Global Sensitivity Analysis and Dimension Reduction for Models with Function-valued Outputs

Helen Cleaves, Alen Alexanderian, Hayley Guy, and Ralph Smith, North Carolina State University, U.S.; Meilin Yu, University of Maryland, Baltimore County, U.S.

10:10-10:30 Exploiting Low-rank Structure for Sensitivity Analysis in Earth System Models

Cosmin Safta and Khachik Sargsyan, Sandia National Laboratories, U.S.; Daniel Ricciuto, Oak Ridge National Laboratory, U.S.; John D. Jakeman, Sandia National Laboratories, U.S.; Alex Gorodetsky, University of Michigan, U.S.

10:35-10:55 Global Sensitivity Analysis for Multilevel Scramjet Computations

Xun Huan, University of Michigan, U.S.; Cosmin Safta, Zachary Vane, Guilhem Lacaze, Joseph C. Oefelein, and Habib N. Najm, Sandia National Laboratories, U.S.

11:00-11:20 Probability-space Modeling for Global Sensitivity Analysis

Sankaran Mahadevan, Vanderbilt University, U.S.; Zhen Hu, University of Michigan, Dearborn, U.S.

continued in next column

Friday, March 1

MS348

Guidance from Early Applications of Quantum Computers

9:45 a.m.-11:25 a.m.

Room: 300C

Quantum computers (QCs) based on quantum annealing continue to grow rapidly in their delivered performance and the variety of applications being implemented on them. As the computational science community eagerly awaits fulfillment of the promise of dramatically higher computational performance, guidance from those early applications appears high value. This minisymposium aims to identify and disseminate successful methods for using these early QCs.

Organizer: Steve P. Reinhardt
D-Wave Systems, Inc., U.S.

Organizer: Scott Pakin
Los Alamos National Laboratory, U.S.

9:45-10:05 Logic Programming for Quantum Annealers
Scott Pakin, Los Alamos National Laboratory, U.S.

10:10-10:30 Applications of Quantum Computing in Computational Chemistry
Scott N. Genin, OTI Lumionics Inc., Canada; Ilya Ryabinkin and Artur Izmaylov, University of Toronto, Canada

10:35-10:55 Solving Industrial Network-optimization Problems on a Quantum Computer
Sheir Yarkoni, D-Wave Systems, Germany

11:00-11:20 Material Simulation on D-Wave
Michael Streif and *Florian Neukart*, Volkswagen, Germany

Friday, March 1

MS349

BE: Wrap-up Session

9:45 a.m.-11:25 a.m.

Room: 302A

Organizer: Mary Ann E. Leung
Sustainable Horizons Institute, U.S.

Friday, March 1

MS350

Multiscale and Domain Decomposition Approaches for PDEs with Rough Coefficients - Part I of II

9:45 a.m.-11:25 a.m.

Room: 302B

For Part 2 see MS382

Classical Lagrangian finite element methods are challenged by multi-scale problems due to the range of scales present in the problem. Over the last years, many discretization methods have been proposed to enable the accurate, efficient, and robust solution of these complex models. Examples include the Multi-scale Finite Element Method (and its variants, such as the Generalized Multi-scale Finite Element Method -- GMsFEM), the mixed multi-scale finite element, the Heterogeneous Multi-scale Method (HMM), multiscale methods based on a orthogonal decomposition of the space (such as the LOD), adaptive multi-scale methods, ... Such discretization methods often rely on approximation subspaces that incorporate specialized knowledge of the partial differential equation. Similar ideas appear in the context of preconditioners (or linear solvers) based on domain decomposition to build enhanced coarse spaces. The objective of this mini-symposium is to gather presentations where approximation subspaces incorporating specialized knowledge of the multi-scale partial differential equation are used either for the discretization method and/or for the linear solver.

Organizer: Ulrich Hetmaniuk
University of Washington, U.S.

Organizer: Frederic Legoll
Ecole Nationale des Ponts et Chaussées, France

9:45-10:05 Domain Decomposition Method for High Dimensional Stochastic Systems
Srikara Pranesh, University of Manchester, United Kingdom; Debraj Ghosh, Indian Institute of Science, Bangalore, India

continued on next page

10:10-10:30 The Schwarz Alternating Method for Concurrent Multiscale in Finite Deformation Solid Mechanics

Alejandro Mota and Irina K. Tezaur, Sandia National Laboratories, U.S.; Coleman Alleman, ; Gregory Phlipot, California Institute of Technology, U.S.

10:35-10:55 Robust Model Reduction for High-contrast Problems

Marcus Sarkis, Worcester Polytechnic Institute, U.S.; Alexandre L. Madureira, Laboratorio Nacional de Computacao Cientifica, Brazil

11:00-11:20 Reducing the Resonance Error in Numerical Homogenization Through Parabolic Micro Problems

Edoardo Paganoni, Assyr Abdulle, and Doghonay Arjmand, École Polytechnique Fédérale de Lausanne, Switzerland

Friday, March 1

MS351

Recent Advances in Numerical Methods for Multiphase Flow Problems - Part I of II

9:45 a.m.-11:25 a.m.

Room: 303A

For Part 2 see MS383

In nature and environment, rain, snow, fog, avalanches, mud slides, sediment transport, debris flows are all examples of multiphase flow where the behaviour of the phases are studied in different fields of natural science. Phase field method, or called diffusive interface approach, as a powerful numerical tool for simulating free interfacial motions, had been widely used in modeling and numerical approaches for multiphase flow problems. In this mini-symposium, the experts are expected to communicate the new idea, share their recent advances in the modeling, numerical Methods for, but not limited to, the phase field approach, as well as their applications in material science, fluid mechanics, etc.

Organizer: Xiaofeng Yang

University of South Carolina, U.S.

Organizer: Jia Zhao

Utah State University, U.S.

9:45-10:05 Phase-field Simulations of Contact Angle Hysteresis

Pengtao Yue, Virginia Tech, U.S.

10:10-10:30 Immersed Boundary Method in Computational Cardiology

Li Cai, Northwestern Polytechnical University, China

10:35-10:55 A Diffuse Domain Approach for Two Phase Flow in Complex Geometries

Zhenlin Guo, University of California, Irvine, U.S.

11:00-11:20 Thermodynamically Consistent Phase-field Models and Their Numerical Approximations

Jia Zhao, Utah State University, U.S.

Friday, March 1

MS352

Recent Developments in Model Order Reduction Methods - Part I of II

9:45 a.m.-11:25 a.m.

Room: 303B

For Part 2 see MS384

An increase in demands for simulations of complex systems has motivated active developments in model order reduction. We discuss several recent developments in model order reduction methods addressing challenges such as high dimensional parameter space, expensive offline phase. We also discuss how model reduction algorithms are developed in specific domains such as fluid control of quadratic systems, advection-dominated high-gradient structure systems, coupled flow-geomechanics problems, finite time Lyapunov exponents, and poroelastic media.

Organizer: Youngsoo Choi

Lawrence Livermore National Laboratory, U.S.

9:45-10:05 Accelerating Training Phase in Time-dependent Nonlinear Model Order Reduction

Youngsoo Choi, Lawrence Livermore National Laboratory, U.S.; Deshawn Coombs, Syracuse University, U.S.; Robert W. Anderson, Lawrence Livermore National Laboratory, U.S.

10:10-10:30 Data-driven Reduced Order Modeling for Parametrized Problems

Mengwu Guo and Jan S. Hesthaven, École Polytechnique Fédérale de Lausanne, Switzerland

10:35-10:55 Reduced-order Models for Advection-dominated Systems

Philipp Schulze, Technische Universität Berlin, Germany

11:00-11:20 Interpolatory Model Reduction and Control of Quadratic Systems

Jeff Borggaard, Virginia Tech, U.S.

Friday, March 1

MS353

Mathematics of Energy Materials - Part I of II

9:45 a.m.-11:25 a.m.

Room: 201C

For Part 2 see MS385
Featured Minisymposium

Dramatically improved performance of energy storage devices, e.g., batteries and capacitors, is a prerequisite for successful energy transition. Porous materials are a key component in most of such devices. This minisymposium deals with innovative approaches to modeling, simulations, and design of energy storage devices, including development of new porous materials.

Organizer: Ilenia Battiato
Stanford University, U.S.

Organizer: Daniel M. Tartakovsky
Stanford University, U.S.

Organizer: John H. Cushman
Purdue University, U.S.

9:45-10:05 Homogenization-informed Convolutional Neural Networks for Identification of Effective Properties of Energy Storage Devices
Ilenia Battiato, Stanford University, U.S.

10:10-10:30 Microstructurally-induced Properties and Degradation in Lithium-ion Batteries
Aniruddha Jana, Abhas Deva, Alfonso Campos, and Lucas D. Robins, Purdue University, U.S.; Ilenia Battiato, Stanford University, U.S.; *Edwin Garcia*, Purdue University, U.S.

10:35-10:55 Predictive Probabilistic Graphical Models for Energy Materials
Markos A. Katsoulakis, University of Massachusetts, Amherst, U.S.; Eric J. Hall, RWTH Aachen, Germany

11:00-11:20 Optimal Design of Nanoporous Materials for Electrical Storage Devices
Xuan Zhang, GEICO, U.S.; *Daniel M. Tartakovsky*, Stanford University, U.S.

Friday, March 1

MS354

Next Generation FFT Algorithms in Theory and Practice: Parallel Implementations, Sparse FFTs, and Applications - Part I of II

9:45 a.m.-11:25 a.m.

Room: 202A

For Part 2 see MS386
The fast Fourier Transform (FFT) is an algorithm used in a wide variety of applications, yet does not make optimal use of many current hardware platforms. Hardware utilization performance on its own does not however imply optimal problem solving. The purpose of this minisymposium is to enable exchange of information between people working on alternative FFT algorithms such as sparse and non uniform FFTs, to those working on FFT implementations, in particular for parallel hardware.

Organizer: Daisuke Takahashi
University of Tsukuba, Japan

Organizer: Mark Iwen
Michigan State University, U.S.

Organizer: Samar A. Aseeri
King Abdullah University of Science & Technology (KAUST), Saudi Arabia

Organizer: Benson K. Muite
University of Tartu, Estonia

9:45-10:05 Implementation of Parallel 3-D Real FFT with 2-D Decomposition on Intel Xeon Phi Clusters
Daisuke Takahashi, University of Tsukuba, Japan

10:10-10:30 Discrete Sparse Fourier Transforms: Faster Stable Implementations with Guarantees
Mark Iwen, Sami Merhi, and Ruochuan Zhang, Michigan State University, U.S.

10:35-10:55 Getting Best Performance of Memory Bandwidth Limited Algorithms with Intel MKL
Alexander Kalinkin, Intel Corporation, U.S.

11:00-11:20 Implementation of Sparse FFT with Structured Sparsity
Sina Bittens, University of Goettingen, Germany; Ruochuan Zhang and Mark Iwen, Michigan State University, U.S.

Friday, March 1

MS355

High-order Discretizations and Quadrature for Integral Equation Methods - Part I of II

9:45 a.m.-11:25 a.m.

Room: 202B

For Part 2 see MS387
Integral equation methods are a powerful set of techniques for simulating physical systems accurately and efficiently. They have been employed with great success in the design of devices in many fields such as audio equipment, microfluidic devices, photonic devices, and biomedical equipment. Recent years have seen tremendous improvement in the precision and complexity of devices which can be manufactured necessitating the development of new computationally-efficient high-precision numerical tools for complex geometries. In integral equation methods the solutions are typically represented as singular integrals, requiring efficient discretization and quadrature methods. Existing methods include skeletonization based approaches, generalized Gaussian quadrature, and quadrature by expansion. When coupled with fast algorithms the resulting discretizations yield solvers scaling linearly in the number of degrees of freedom. Two critical avenues of research are improving performance in complex geometries and in three dimensions, as well as developing automatically adaptive discretizations which resolve complexity in geometry and input data. The talks in this mini-symposium are focused on the construction of efficient Nystrom methods and associated high-order quadrature methods for surfaces in three dimensions; the design of automatically adaptive discretizations in moving geometries; and the detailed analysis of singularities of solutions near corners and edges.

Organizer: Manas N. Rachh
Simons Foundation and Flatiron Institute, U.S.

Organizer: Jeremy Hoskins
Yale University, U.S.

continued on next page

Organizer: Kirill Serkh
Yale University, U.S.

9:45-10:05 An Overview of Nystrom (and Not-So-Nystrom) High-order Surface Quadratures for Fast Solvers
Alex H. Barnett, Simons Foundation, U.S.

10:10-10:30 Integral Equation Methods for the Heat Equation in Moving Geometries
Jun Wang, Simons Foundation, U.S.; Leslie Greengard, Simons Foundation and Courant Institute of Mathematical Sciences, New York University, U.S.

10:35-10:55 An Adaptive Technique for 2D Boundary Integral Equation.
Yabin Zhang, Rice University, U.S.

11:00-11:20 An Efficient and High Order Accurate Solution Technique for Three Dimensional Elliptic Partial Differential Equations
Natalie N. Beams and Adrianna Gillman, Rice University, U.S.

Friday, March 1

MS356

Modeling Resource Utilization and Contention in HPC System-Application Interactions - Part I of II

9:45 a.m.-11:25 a.m.

Room: 202C

For Part 2 see MS388

Application performance, system throughput and stability, and procurement specifications can be improved by better understanding of applications' resource requirements and the performance impacts of resource contention within a site's workload. Production data can potentially provide the most accurate insights, however, quantification of resource utilization and contention impact may be difficult empirically because relationships between such data and performance are largely inferential and controlled conditions for experimentation are difficult to obtain in production. Conversely, theoretical models developed in the design phase of large-scale systems have missed large-scale contention scenarios and may not capture the range of behavior of time-dependent production workloads. Advances in instrumentation and large data analytical and visual techniques can improve our ability to develop data-driven models with the fidelity to capture events of interest. Efficient techniques for assessing near-term system behavior with respect to models can enable runtime responses to improve operations and be used for model feedback to improve confidence. We bring together experts to discuss current techniques and experiences in developing and validating models and characterizations of resource utilization and contention; exploring data to determine and affirm model dependencies; and extracting insights from models to dynamically effect improved performance and operations.

Organizer: James Brandt
Sandia National Laboratories, U.S.

Organizer: Ann Gentile
Sandia National Laboratories, U.S.

Organizer: Zbigniew Kalbarczyk
University of Illinois at Urbana-Champaign, U.S.

9:45-10:05 Extracting Actionable System-application Performance Factors
James Brandt and Ann Gentile, Sandia National Laboratories, U.S.; Jonathan Cook, New Mexico State University, U.S.

10:10-10:30 Holistic Performance Diagnosis for HPC Workflows
Karen L. Karavanic, Portland State University, U.S.

10:35-10:55 Performance Analysis of a Multi-application Plasma Surface Interactions Workflow
Philip C. Roth, Oak Ridge National Laboratory, U.S.; Ane Lasa, Sophie Blondel, and Timothy Younkin, University of Tennessee, Knoxville, U.S.; Jon Drobny, University of Illinois at Urbana-Champaign, U.S.; Mark Cianciosa and Wael R. Elwasif, Oak Ridge National Laboratory, U.S.; Davide Curreli, University of Illinois at Urbana-Champaign, U.S.; John Canik and David E. Bernholdt, Oak Ridge National Laboratory, U.S.; Brian Wirth, University of Tennessee, Knoxville, U.S.

11:00-11:20 Challenges and Opportunities in Capturing, Modeling, and Actively Managing Power and Energy Usage on HPC Systems
Martin Schulz, Technische Universitaet Muenchen, Germany

continued in next column

Friday, March 1

MS357

Divide and Conquer Strategies for Large-scale Eigenvalue Problems - Part I of II

9:45 a.m.-11:25 a.m.

Room: 203

For Part 2 see MS389

This minisymposium presents recent advances in large-scale eigenvalue computations based on divide-and-conquer strategies. These new methodologies aim to handle problems that are often found challenging for traditional algorithms, namely, extracting a large number of eigenvalues from a large and sparse matrix, or extracting eigenvalues that are deep inside the spectrum. The presented techniques include spectrum slicing methods, polynomial and rational filtering, algorithms based on contour integration, etc, along with Krylov subspace methods or subspace iterations. We bring together numerical analysts who study these techniques with different methodologies and also physicists who initiate an effort to assemble the different ideas to solve problems in real applications.

Organizer: Ruipeng Li

Lawrence Livermore National Laboratory, U.S.

Organizer: Yuanzhe Xi

Emory University, U.S.

Organizer: Edoardo A. Di Napoli

Jülich Supercomputing Centre, Germany

9:45-10:05 The EVSL Package for Symmetric Eigenvalue Problems

Yousef Saad, University of Minnesota, U.S.; Ruipeng Li, Lawrence Livermore National Laboratory, U.S.; Yuanzhe Xi, Emory University, U.S.

10:10-10:30 A Contour Intergral-based Stochastic Estimator for Eigenvalue Counts of Generalized Eigenproblems

Yasunori Futamura, Takahiro Yano, Akira Imakura, and Tetsuya Sakurai, University of Tsukuba, Japan

10:35-10:55 Beyond Automated Multilevel Substructuring: Domain Decomposition with Rational Filtering

Vasileios Kalantzis, University of Minnesota, U.S.; Yuanzhe Xi, Emory University, U.S.; Yousef Saad, University of Minnesota, U.S.

11:00-11:20 Lanczos Bidiagonalization, Orthogonal Polynomials, and Singular Value Distributions

Anthony P. Austin, Argonne National Laboratory, U.S.

continued in next column

Friday, March 1

MS358

Advances in Multi-method Time Discretizations of Evolutionary PDEs - Part I of II

9:45 a.m.-11:25 a.m.

Room:205

For Part 2 see MS390

Simulations of systems modeled by large time dependent partial differential equations faces challenges posed by the interaction of multiple physical processes and the existence of multiple dynamical scales. Time discretization is an essential ingredient of the solution process that determines the overall accuracy and stability of the solution, the computational efficiency, and the data dependencies in a parallel environment. This minisymposium highlights recent developments in time stepping methods for large evolutionary PDEs. Topics include, but are not restricted to, partitioned methods for multiphysics systems, multirate schemes, approaches based on a minimal amount of implicitness to achieve stability without sacrificing computational efficiency, methods that preserve special properties of the underlying system (such as monotonicity), and approaches that improve the scalable parallel performance in the context of large applications.

Organizer: Adrian Sandu

Virginia Tech, U.S.

Organizer: Michael Guenther

Bergische Universität Wuppertal, Germany

9:45-10:05 A Class of Multirate Infinitesimal GARK Schemes

Adrian Sandu, Virginia Tech, U.S.

10:10-10:30 On the Derivation of a New Class of Multirate Methods Based on Exponential Integrators

Vu Thai Luan, Southern Methodist University, U.S.

10:35-10:55 Exponential Propagation in Time Integration.

Mayya Tokman, University of California, Merced, U.S.

11:00-11:20 Partitioned Exponential Methods for Stiff Differential Equations

Mahesh Narayanamurthi and Adrian Sandu, Virginia Tech, U.S.

Friday, March 1

MS359

Computational Methods for Data Assimilation and Uncertainty Quantification - Part I of II

9:45 a.m.-11:25 a.m.

Room: 206A

For Part 2 see MS391

Model-based computer simulations about large-scale physical phenomena such as the atmosphere provide invaluable, yet uncertain, prediction of the evolution of the physical system. The model forecasts are corrected based on other available sources of information including measurements of reality, i.e. observations. Observational data are fused with model-based forecast, through the process of data assimilation, to estimate the trajectory or state of dynamical systems. This method has been extensively used in meteorology, hydrology, oceanography among other applications, to make accurate predictions about the state of these systems. The quality of the output of a data assimilation procedure is critically dependent on the data acquisition strategy, and thus an optimal experimental design is required for sensor placement. The focus of this minisymposium is on recent advances in computational methods for data assimilation, uncertainty quantification, sensor placement, and data acquisition.

Organizer: Ahmed Attia

Argonne National Laboratory, U.S.

Organizer: Vishwas Rao

Argonne National Laboratory, U.S.

9:45-10:05 An Optimal Experimental Design Framework for Adaptive Inflation and Covariance Localization for Ensemble Filters

Ahmed Attia and Emil M. Constantinescu, Argonne National Laboratory, U.S.

10:10-10:30 Goal Oriented Sensor Placement for Large-scale Inverse Problems Under Model Uncertainty

Alen Alexanderian, North Carolina State University, U.S.; Ahmed Attia, Argonne National Laboratory, U.S.; Noemi Petra, University of California, Merced, U.S.; Arvind Saibaba, North Carolina State University, U.S.; Georg Stadler, Courant Institute of Mathematical Sciences, New York University, U.S.

10:35-10:55 Adaptive Multivariate Schur Product Localization in Ensemble-based Filters

Andrey Popov, Virginia Polytechnic Institute and State University, U.S.; Adrian Sandu, Virginia Tech, U.S.

11:00-11:20 Title Not Available

Roland Potthast, Universitaet Goettingen, Germany

continued in next column

Friday, March 1

MS360

Numerical Methods for Biological Fluid Dynamics - Part I of II

9:45 a.m.-11:25 a.m.

Room: 206B

For Part 2 see MS392

Fluid flows are essential to life across a wide range of length and time scales, from the swimming of whales in a turbulent ocean environment to the motion of proteins inside of cells. The behaviors exhibited by biological fluids are rich and counterintuitive, given non-Newtonian effects generated by suspended structures that may be passively elastic or actively forcing the fluid out of equilibrium. Dealing with this complex and varied phenomenology requires original computational approaches. In this minisymposium we present a selection of promising recent developments in the field, including fluid-structure interaction and coarse-grained active matter models, and applications such as jellyfish swimming and intracellular motion.

Organizer: Thomas Fai

Brandeis University, U.S.

Organizer: Chris H. Rycroft

Harvard University, U.S.

9:45-10:05 Lubricated Immersed Boundary Method Simulations of Elastic Bodies in Near Contact

Thomas Fai, Brandeis University, U.S.

10:10-10:30 Multiphysics Computational Modeling of the Chick Embryonic Heart

Simone Rossi, Michael Bressan, and Boyce E. Griffith, University of North Carolina at Chapel Hill, U.S.

10:35-10:55 Fluid-structure Interactions in Active Matter

Saverio E. Spagnolie, University of Wisconsin, U.S.

11:00-11:20 Computational Mean-field Modeling of Confined Active Fluids

Maxime Theillard, University of California, Merced, U.S.; David Saintillan, University of California, San Diego, U.S.

Friday, March 1

MS361

Structure-exploiting Techniques for Approximation, Inference and Control of Complex Systems - Part I of II

9:45 a.m.-11:25 a.m.

Room: 206C

For Part 2 see MS393

Data-driven modeling and control of complex systems have received a lot of attention over the recent years, largely due to the availability of large datasets. Nevertheless, large scale data is not typically informative about the entire physical properties because it is often gathered through nonlinear or partially observable processes. In such situations, physically-principled priors that include symmetries and constraints are crucial to infer predictive models, typically with filtering or smoothing techniques. This minisymposium explores the latest developments and challenges in data-efficient techniques for modeling, inference and control of complex physical systems. This includes, but is not limited to, scalable approximation formats and spaces, algorithms for filtering and smoothing, sampling and exploration strategies for gathering "smart data", incorporation of symmetry-based priors, bounds and invariant preservation, information theoretic arguments, etc.

Organizer: Alex Gorodetsky

University of Michigan, U.S.

Organizer: Lionel Mathelin

CNRS, France

9:45-10:05 Kernel Koopman Spectral Analysis for Nonlinear Dynamical Systems

Yoshinobu Kawahara, Osaka University, Japan

10:10-10:30 Verification of Uncertain Pomdps via Lyapunov Functions and Barrier Certificates

Mohamadreza Ahmadi, University of Texas, U.S.

10:35-10:55 Clustering of Time Series via Dynamic Mode Decomposition and the Matrix Pencil Method

Leonid Pogorelyuk and Clarence Rowley, Princeton University, U.S.

11:00-11:20 Robust Reconstruction of Flow Fields from Limited Measurements

Jared L. Callaham, Kazuki Maeda, and Steven Brunton, University of Washington, U.S.

continued in next column

Friday, March 1

MS362

Structured Matrix Methods - Part I of II

9:45 a.m.-11:25 a.m.

Room: 206D

For Part 2 see MS394

Structured matrix methods utilize structures such as parameterized forms and low-rank approximations to gain high efficiency in matrix representation and arithmetic. They have demonstrated significant advantages in improving the efficiency and reliability of some large-scale computations and engineering simulations. Rank structured methods extend the fundamental ideas of multipole and panel-clustering methods to general non-local solution operators. While there exist various more or less closely related methods, the unifying aim of these methods is to explore efficient structured low-rank approximations, especially those exhibiting hierarchical or nested forms. These help the methods to achieve nearly linear complexity. In this minisymposium, we aim to present and exchange recent new developments on structured methods for some challenging numerical problems such as high frequencies, ill conditioning, eigenvalue solution, and stability. Studies of structures, algorithm design, and accuracy control will be discussed. The minisymposium will include experts working on a broad range of structured methods.

Organizer: Sabine Le Borne

Hamburg University of Technology, Germany

Organizer: Jianlin Xia

Purdue University, U.S.

9:45-10:05 Hierarchical Matrices in Stable Radial Basis Function Interpolation

Sabine Le Borne, Hamburg University of Technology, Germany

10:10-10:30 On the Best Approximation of the Hierarchical Matrix Product

Jürgen Dölz, Technische Universität Darmstadt, Germany; Helmut Harbrecht, Universität Basel, Switzerland; Michael D. Multerer, Università della Svizzera italiana, Switzerland

continued on next page

10:35-10:55 Inversion of Rank-structured Matrices: Which Formulas Should You Use?

Per-Gunnar Martinsson, University of Colorado Boulder, U.S.

11:00-11:20 Accelerated Interpolative Decompositions by General Proxy Point Methods

Xin Xing and Edmond Chow, Georgia Institute of Technology, U.S.

Friday, March 1

MS363

Summation-by-Parts: A Framework for the Development and Analysis of Modern Numerical Methods - Part I of II

9:45 a.m.-11:25 a.m.

Room: 207

For Part 2 see MS395
Featured Minisymposium

As we move towards an era of exascale HPC, flexible and robust algorithms will be necessary to take advantage of increasingly complex and failure prone computational resources. In this mini-symposium, the focus is on methods having the summation-by-parts (SBP) property, which are advantageous as they lead to high-order methods that are provably conservative and stable. In particular, the interest is in recent advancements leading to improved efficiency, flexibility, and robustness. Topics that will be covered include: Entropy stable algorithms Linearly stable algorithms h/p-adaptation Mixed formulations Wave propagation and inversion problems Optimized differentiation matrices

Organizer: David C. Del Rey Fernandez
NASA Langley Research Center and National Institute of Aerospace, U.S.

Organizer: Andrew R. Winters
University of Cologne, Germany

Organizer: Jesse Chan
Rice University, U.S.

Organizer: Jason E. Hicken
Rensselaer Polytechnic Institute, U.S.

9:45-10:05 Towards Robust Methods for Practical Aero Dynamic Flows: a Review of Recent Activities at NASA LaRC on Methods with the Summation by-parts Property

David C. Del Rey Fernandez, NASA Langley Research Center and National Institute of Aerospace, U.S.; Mark H. Carpenter, NASA Langley Research Center, U.S.

10:10-10:30 Stable and Accurate Filtering Procedures

Jan Nordström, Linköping University, Sweden; Tomas Lundquist, University of Cape Town, South Africa

10:35-10:55 Flexural Wave Propagation in Ice Shelfs

Jonatan Werpers and Ken Mattsson, Uppsala University, Sweden

11:00-11:20 An Algorithmic Search for Optimally Efficient High-order Summation-by-parts Operators

Kevin J. Mattalo and Pieter Boom, University of Toronto Institute for Aerospace Studies, Canada; David W. Zingg, University of Toronto, Canada

continued in next column

Friday, March 1

MS364

Bringing Algorithms and New Applications to Novel Architectures - Part I of II

9:45 a.m.-11:25 a.m.

Room: 401A

For Part 2 see MS396

Computational science and engineering now reaches beyond HPC simulation to diverse applications in machine learning and data analysis. Common architectures still focus on dense matrix multiplication. The result: Engineering solutions that require too much power and too much time. But the success of one particular matrix multiplication engine, the GPGPU, has opened people to design issues around "alternative" architectures. Improved software tools, larger flexible architectures (e.g. FPGAs), and lower fabrication costs have lowered the bar to entry. This minisymposium focuses on novel architectures addressing the new application challenges. These are real architectures and real systems for real problems. They provide testing grounds for applications, software frameworks, and our preconceptions on building both. We consider a range from novel memory systems and FPGAs to neuromorphic and quantum systems.

Organizer: Jason Riedy

Georgia Institute of Technology, U.S.

Organizer: Jeffrey Young

Georgia Institute of Technology, U.S.

9:45-10:05 Novel Architectures for Applications in Data Science and Beyond

Jason Riedy, Jeffrey Young, and Tom Conte, Georgia Institute of Technology, U.S.

10:10-10:30 Scientific Computer Architecture Beyond Moore's Law

John Shalf and David Donofrio, Lawrence Berkeley National Laboratory, U.S.

10:35-10:55 A Software Framework for Spiking Neuromorphic Computing Systems

Catherine Schuman, Oak Ridge National Laboratory, U.S.

11:00-11:20 Benchmarking and Modeling Sparse Memory Accesses for Heterogeneous Systems

Patrick Lavin, Georgia Institute of Technology, U.S.

Friday, March 1

MS365

High-order Discontinuous Galerkin and Finite Element Methods for CFD - Part I of II

9:45 a.m.-11:25 a.m.

Room: 401B

For Part 2 see MS397

Given recent developments in computational architectures, high-order discontinuous Galerkin and finite element methods are increasingly important for a range of applications in computational fluid dynamics. Their low dissipation, compact stencils, and high ratio of computation to communication make them well-suited for solving flow problems on next-generation many-core architectures. However, despite these desirable features, high-order methods suffer from robustness issues in the presence of discontinuous features, sharp gradients, and underresolution, which can limit their utility in resolving challenging flow features such as shock waves and turbulence. Furthermore, the efficient implementation of such methods on modern, massively-parallel, memory-constrained architectures presents interesting and important new challenges, prompting the development of sophisticated solver technology and matrix-free operators. This minisymposium aims to bring together researchers in the broad field of high-order discontinuous Galerkin and finite element methods to disseminate their research and discuss recent developments in discretizations, solvers, and software that aim to overcome the challenges facing the field. Specific topics of interest include shock capturing and tracking methods, stable and robust discretizations, efficient solvers and time integration (implicit, explicit, and IMEX), matrix-free algorithms, and algorithms for advanced computer architectures.

Organizer: Will Pazner

University of California, Berkeley, U.S.

Organizer: Matthew J. Zahr

University of California, Berkeley and Lawrence Berkeley National Laboratory, U.S.

9:45-10:05 An Optimization-based Discontinuous Galerkin Approach for High-order Accurate Shock Tracking

Marzieh Mirhose, University of Notre Dame, U.S.; Matthew J. Zahr, University of California, Berkeley and Lawrence Berkeley National Laboratory, U.S.; Per-Olof Persson, University of California, Berkeley, U.S.

10:10-10:30 Efficient Methods for Solving Implicit Space-time Multi-physics Systems

Scott Murman and Laslo Diosady, NASA Ames Research Center, U.S.

10:35-10:55 An Adaptive Stochastic Discontinuous Galerkin Method for Reliable Uncertainty Propagation

Geoff Donoghue and Masayuki Yano, University of Toronto, Canada

11:00-11:20 Efficient Matrix-free Solvers for the Incompressible Navier-Stokes Equations

Niklas Fehn, Technische Universität München, Germany; Wolfgang A. Wall, Technische Universitaet Muenchen, Germany; Martin Kronbichler, Technische Universität München, Germany

continued in next column

Friday, March 1

MS366

Hybrid Parallelization for Modern Architectures - Part I of II

9:45 a.m.-11:00 a.m.

Room: 401C

For Part 2 see MS398

Compute clusters are evolving away from compositions of many nodes that contain a few processors with a small amount of shared memory towards compositions of fewer nodes that contain many processors and a large amount of shared memory. This transition in computational architecture necessitates the development of hybrid parallelization implementations of scientific and engineering software to efficiently and flexibly utilize distributed and shared resources (processors and memory). Speakers will discuss algorithmic choices and implementation trade-offs for established and innovative methods. The impacts on computational cost (memory, CPU hours, and wall time) for production class simulations and implications for scaling to leadership class simulations will also be discussed.

Organizer: Stephen L. Wood
NASA Langley Research Center, U.S.

Organizer: Eric Nielsen
NASA Langley Research Center, U.S.

9:45-10:05 Hybrid Parallelization of the Stabilized Finite Element Library in FUN3D

Stephen L. Wood and William Anderson, NASA Langley Research Center, U.S.

10:10-10:30 Exploiting Node-level Performance in Sparse Linear Algebra

Hartwig Anzt, University of Tennessee, U.S.

10:35-10:55 Optimizing Single Node Performance of Parallel Trajectory Reconstruction

Anthony Williams and Justin Green, NASA Langley Research Center, U.S.

Friday, March 1

MS367

State-of-the-Art Auto-tuning: New Approaches and Algorithmic Innovations Towards Exascale Computing - Part I of II

9:45 a.m.-11:25 a.m.

Room:402A

For Part 2 see MS399

Exascale computers are expected to be deployed by 2021. Their architectures will be complex, building upon many-cores and GPUs, and offering unprecedented levels of parallelism. It is expected that auto-tuning (AT) research and technology will continue building upon its proven success for delivering high performance on a variety of computer architectures, and will enable optimized, high performance implementations of specific computations for those challenging architectures. For this, preliminary results suggest that the use of machine learning techniques is promising. In concert with the hardware evolution and complexity, algorithmic innovations will need to be devised for numerical computations that are essential to many applications, such as eigenvalue solvers and linear equation solvers. For example, reducing communications together with high performance implementations will be essential. While exascale computers will enable the solution of problems with unprecedented levels of details, the accuracy of the computations cannot be taken for granted: for example, codes that perform floating point sums of distributed data can produce results that are difficult to be reproduced. Given potentially conflicting goals for achieving performance and reproducibility, this minisymposium will discuss technology trends for AT and their interplay with AT frameworks, numerical algorithms, and accuracy assurance towards extreme levels of computing.

Organizer: Takahiro Katagiri
Nagoya University, Japan

Organizer: Osni A. Marques
Lawrence Berkeley National Laboratory, U.S.

Organizer: Toshiyuki Imamura
RIKEN, Japan

9:45-10:05 Auto-tuning of Preconditioners with Deep Learning

Takahiro Katagiri and Kenya Yamada, Nagoya University, Japan

10:10-10:30 The Big Data Approach to Autotuning of GPU Kernels Using the BONSAI Toolset

Jakub Kurzak and Piotr Luszczek, University of Tennessee, Knoxville, U.S.; Mark Gates, University of Tennessee, U.S.; Yaohung Tsai and Ahmad Abdelfattah, University of Tennessee, Knoxville, U.S.; Jack Dongarra, University of Tennessee, U.S.

10:35-10:55 Talk Title Not Available

Hiroyuki Takizawa, Tohoku University, Japan

11:00-11:20 Model-driven Auto-tuning for Optimizing Tensor Contractions

Ponnuswamy Sadayappan, Jinsung Kim, Rui Li, and Aravind Sukumaran-Rajam, Ohio State University, U.S.; Sriram Krishnamoorthy, Pacific Northwest National Laboratory, U.S.

continued in next column

Friday, March 1

MS368

Meshfree Methods: Computational Advances and Applications - Part I of II

9:45 a.m.-11:25 a.m.

Room: 402B

For Part 2 see MS400

Meshfree (or meshless) methods offer geometric flexibility, non-uniform resolution, and advantageous trade-offs between accuracy and computational costs. These properties make meshfree methods particularly useful for large-scale science and engineering applications featuring complex geometries and multiscale phenomena. This minisymposium focuses on computational advancements and applications of various meshfree discretization methods, including radial basis function generated finite differences (RBF-FD), partition of unity, and generalized moving least squares (GMLS). Recent computational advancements in scalable implementations, adaptive refinement, approximation space enrichment, stabilization, non-linear problems, and conservation techniques will be discussed together with applications to partial differential equations (PDEs) on manifolds, biological and geophysical fluid flows, and non-linear elasticity.

Organizer: Grady B. Wright

Boise State University, U.S.

9:45-10:05 A High-order Meshfree Semi-Lagrangian Method for Advection on Manifolds: Mass-conservation

Grady B. Wright, Boise State University, U.S.

10:10-10:30 The Fast Radial Basis Functions Orthogonal Gradients (RBF-OGr) Method for Solving PDEs on Arbitrary Surfaces

Cecile M. Piret, Michigan Technological University, U.S.

10:35-10:55 Meshless Methods for Manifolds: Hydrodynamics of Curved Fluid Interfaces and Related Applications

Ben Gross and Christian Bueno, University of California, Santa Barbara, U.S.; Nathaniel Trask, Sandia National Laboratories, U.S.; *Paul Atzberger*, University of California, Santa Barbara, U.S.

11:00-11:20 RBF-FD for PDEs on Manifolds: Stable and Stagnation-free Formulations

Varun Shankar, Akil Narayan, and Robert M. Kirby, University of Utah, U.S.

continued in next column

Friday, March 1

MS369

Molecular and Electronic Structure Theory Meets Data Science - Part I of II

9:45 a.m.-11:25 a.m.

Room: 402C

For Part 2 see MS401

The prediction of molecular and electronic structures using first principle calculations are among the most challenging and computationally demanding science and engineering problems. This minisymposium aims at presenting and discussing new developments of using machine learning techniques to reduce the computational cost and characterize important features and properties of materials and molecular systems. We bring together experts on molecular and electronic structure theories, which include mathematicians, physicists and chemists working actively in the field.

Organizer: Chao Yang

Lawrence Berkeley National Laboratory, U.S.

Organizer: Lin Lin

University of California, Berkeley and Lawrence Berkeley National Laboratory, U.S.

9:45-10:05 Overview of Machine Learning Techniques

Chao Yang, Lawrence Berkeley National Laboratory, U.S.

10:10-10:30 Deep Learning for Multi-scale Molecular Modelling

Linfeng Zhang, Weinan E, and Roberto Car, Princeton University, U.S.

10:35-10:55 Bayesian Inference of Molecular and Atomic Properties: Representations and Regression Techniques

Yu-Hang Tang, Lawrence Berkeley National Laboratory, U.S.

11:00-11:20 Coordinate-wise Descent Methods for the Full Configuration Interaction Calculation

Yingzhou Li, Jianfeng Lu, and Zhe Wang, Duke University, U.S.

Friday, March 1

CP19

Machine Learning in CSE

9:45 a.m.-11:05 a.m.

Room: 301

Chair: Heng Xiao, Virginia Tech, U.S.

9:45-10:00 Modified Kernel Regression for Building Low-dimensional Models of Complex Dynamical Systems

Elizabeth Armstrong and James C. Sutherland, University of Utah, U.S.

10:05-10:20 Bayesian Geometric Convolutional Neural Networks for Engineering Applications

Nicholas Geneva and Nicholas Zabaras, University of Notre Dame, U.S.

10:25-10:40 Using Machine Learning in Optimization of Structural Design of Thin-walled Vessels under External Pressure

Lydia S. Novozhilova, Western Connecticut State University, U.S.; Vladimir Urazhdin, Solutions With Innovation, LLC, U.S.

10:45-11:00 Physics-informed Machine Learning Approach for Augmenting Turbulence Models

Jinlong Wu and Carlos Michelen, Virginia Tech, U.S.; Jianxun Wang, University of Notre Dame, U.S.; *Heng Xiao*, Virginia Tech, U.S.

Friday, March 1

CP20

Numerical PDEs III

9:45 a.m.-11:05 a.m.

Room: 201B

Chair: Ayaboe K. Edoh, ERC Inc. and Air Force Research Laboratory , U.S.

9:45-10:00 Probably Stable and Conservative Spectrally-tunable Artificial Dissipation Operators

Ayaboe K. Edoh, ERC Inc. and Air Force Research Laboratory , U.S.

10:05-10:20 Numerical Advection Through PDE-Based Flow Map Composition

Chinmay S. Kulkarni and Pierre F. Lermusiaux, Massachusetts Institute of Technology, U.S.

10:25-10:40 Comparison of Weno Schemes to Second Order Methods for Capturing Compressible Turbulent Spectra

John P. Wakefield, University of Michigan, U.S.; Emmanuel Motheau, Lawrence Berkeley National Laboratory, U.S.

10:45-11:00 A Second Order Fully-discrete Linear Energy Stable Numerical Scheme of a Binary Compressible Viscous Fluid Model

Xueping Zhao, University of South Carolina, U.S.

Intermission

11:25 a.m.-11:30 a.m.

Friday, March 1

MS370

Tensor Decomposition for High Performance Data Analytics - Part II of II

11:30 a.m.-1:10 p.m.

Room:Ballroom 100BC

For Part 1 see MS337

Tensors are higher order generalization of matrices that map naturally to complex relationships found in the real world. In this regard, tensor decomposition has gained popularity in data analytics due its ability to identify hidden properties in large quantities of data. The talks in this minisymposium will explore state-of-the-art research into finding efficient and scalable solutions to decomposing tensors for data analysis on the latest processors and accelerators.

Organizer: Jee Choi

IBM T.J. Watson Research Center, U.S.

11:30-11:50 Performance Portable Parallel CP-APR Tensor Decompositions

Keita Teranishi, Sandia National Laboratories, U.S.

11:55-12:15 Talk Title Not Available

Rich Vuduc, Georgia Institute of Technology, U.S.

12:20-12:40 A Multi-dimensional Morton-ordered Block Storage for Mode-oblivious Tensor Omputations

Filip Pawlowski, Huawei Technologies, France; Bora Ucar, LIP-ENS Lyon, France; Albert-Jan N. Yzelman, Huawei Technologies, France

12:45-1:05 Scaling and Deepening Tensor Decompositions and Applications

Muthu M. Baskaran, Reservoir Labs, U.S.

Friday, March 1

MS371

Software Productivity and Sustainability for CSE and Data Science - Part II of II

11:30 a.m.-1:10 p.m.

Room: Conference Theater

For Part 1 see MS338

Software is the key crosscutting technology that enables advances in mathematics, computer science, and domain-specific science and engineering to achieve robust simulations and analysis for predictive science, engineering, and other research fields. While software is becoming more complex due to multiphysics and multiscale modeling, the coupling of data analytics, and disruptive changes in computer hardware (due to increases in typical system scale and heterogeneity, including GPUs and additional alternative architectures), software itself has not traditionally received focused attention in the CSE community or been rewarded by that community. The presenters in this minisymposterium will address work that addresses growing technical and social challenges in software productivity, quality, and sustainability, and thereby helps software fulfill its critical role as a cornerstone of long-term CSE collaboration. Having a minisymposterium for these topics provides a natural gathering point during poster sessions for informal conversation.

Organizer: David E. Bernholdt

Oak Ridge National Laboratory, U.S.

Organizer: Daniel S. Katz

University of Illinois at Urbana-Champaign, U.S.

Organizer: Catherine Jones

Science and Technology Facilities Council, United Kingdom

11:30-11:50 Supporting and Sustaining Open Source Software Development: the Commons Perspective

C. Titus Brown, University of California, Davis, U.S.

11:55-12:15 Contemporary Peer Code Review in Research Software

Jeffrey C. Carver, University of Alabama, U.S.

12:20-12:40 Acknowledging Scientific Software to Ensure the Future and Legacy of Scientific Research

Daina Bouquin, Harvard-Smithsonian Center for Astrophysics, U.S.

12:45-1:05 Supporting Continuous Integration at Large-scale HPC Centers

Todd Gamblin and Rob Neely, Lawrence Livermore National Laboratory, U.S.; David Montoya, Los Alamos National Laboratory, U.S.

continued in next column

Friday, March 1

MS372

Recent Development of Numerical Methods for Optics and Plasmonics - Part II of II

11:30 a.m.-1:10 p.m.

Room: 102A

For Part 1 see MS339

The main purpose of this workshop is to bring together specialists in the fields of plasmonics and optics. Driven by spectacular advances in the design capabilities of materials at the nanoscale, there has been recent exponential growth in the fields of plasmonics and nano-optics. In particular, algorithm design, analysis, application and implementations of mathematical models in the area are now at the heart everyday technologies such as seismic imaging, underwater acoustics, and biomathematics; nevertheless some very important issues still remain open. This mini symposium is meant to be a platform to exchange ideas on these problems.

Organizer: YoungJoon Hong

San Diego State University, U.S.

Organizer: David P. Nicholls

University of Illinois, Chicago, U.S.

11:30-11:50 A Plane Wave Discontinuous Galerkin Method with a Dirichletto-Neumann Boundary Condition for the Scattering Problem in Acoustics

Peter B. Monk, University of Delaware, U.S.

11:55-12:15 Recent Advances in High-order Simulations for Electromagnetics

Mi Sun Min and Yu-Hsiang Lan, Argonne National Laboratory, U.S.

12:20-12:40 A Non-overlapping Domain Decomposition Method for Simulating Localized Surface Plasmon Resonances: High Accuracy Numerical Simulation and Analyticity of Solutions

Xin Tong, University of Illinois at Chicago, U.S.

12:45-1:05 Computational Tools for Periodic and Non-periodic Nanophotonic Devices Design

Agustin Fernandez Lado and Oscar P. Bruno, California Institute of Technology, U.S.

Friday, March 1

MS373

Neutrino Transport Methods in Astrophysics - Part II of II

11:30 a.m.-1:10 p.m.

Room: 102B

For Part 1 see MS340

Core-collapse supernovae (CCSNe) and compact binary mergers (CBMs) are cataclysmic astrophysical events responsible for heavy element synthesis and the emission of photon, neutrino, and gravitational wave (GW) signals. They have long been targets of instruments covering most of the electromagnetic (EM) spectrum, and more recently of neutrino and GW detectors. The recent detection of a merging binary neutron star pair in EM and GW channels marked a breakthrough for multi-messenger astronomy. Harvesting insights into the physical processes driving these events from the wealth of data relies heavily on sophisticated models requiring extreme-scale, high-fidelity computing. These models solve a coupled system of equations for self-gravity, magneto-hydrodynamics, and neutrino transport. In CCSNe and CBMs, lepton and four-momentum exchange between neutrinos and matter plays a major role in the dynamics. However, neutrino-matter interactions occur under non-equilibrium conditions and a kinetic description based on the Boltzmann transport equation is warranted. As a six-dimensional phase-space problem, our ability to model these astrophysical events with satisfactory realism relies on advances in multi-physics and multi-scale algorithms, novel discretization techniques, fast solvers, and sustainable software. The goal of this minisymposium is to bring together researchers working on topics of relevance to neutrino transport modeling to discuss recent work and exchange ideas.

Organizer: Eirik Endeve
Oak Ridge National Laboratory, U.S.

Organizer: Christian Cardall
Oak Ridge National Laboratory, U.S.

Organizer: Reuben Budiardja
University of Tennessee, U.S.

11:30-11:50 Radiation Hydrodynamics in GenASiS

Christian Cardall, Oak Ridge National Laboratory, U.S.

11:55-12:15 FORNAX: A Flexible Code for Multiphysics Astrophysical Simulations

Aaron Skinner, Lawrence Livermore National Laboratory, U.S.

12:20-12:40 Hybrid Defect Correction Methods for Time-dependent Radiation Transport Simulations

Michael Crockatt and Andrew J. Christlieb, Michigan State University, U.S.; Cory Hauck, Oak Ridge National Laboratory, U.S.

12:45-1:05 Explosion Mechanisms in Multi-dimensional Stellar Core Collapse with Magnetic Fields

Martin Obergaulinger, University of Valencia, Spain

continued in next column

Friday, March 1

MS374

Stochastic Modeling and Algorithms for Complex Physical Systems - Part II of II

11:30 a.m.-12:45 p.m.

Room: 102C

For Part 1 see MS341

This minisymposium focuses on the fundamental problem of how to approximate solutions of both forward and inverse stochastic systems. Predicting the behavior of complex physical processes relies on constructing and solving stochastic systems in high dimensional spaces, particularly in the case when the input data (coefficients, forcing terms, initial and boundary conditions, geometry) are affected by large amounts of uncertainties. We will highlight recent advances in theory and algorithms for stochastic systems, as well as their applications in uncertainty quantification, machine learning, computational fluid dynamics, statistical mechanics, etc. This mini-symposium will bring together researchers from across the applied and computational mathematics communities to discuss recent advances, identify future direction, and promote new collaborations on stochastic methods.

Organizer: Miroslav Stoyanov
Oak Ridge National Laboratory, U.S.

Organizer: Tao Zhou
Chinese Academy of Sciences, China

11:30-11:50 Splitting Up Method for Backward Doubly SDEs and Applications to Nonlinear Filtering Problems

Yanzhao Cao, Auburn University, U.S.; Feng Bao, Florida State University, U.S.; He Zhang, Auburn University, U.S.

11:55-12:15 A Multilevel Reduced-basis Method for Linear Parameterized PDEs

Miroslav Stoyanov, Oak Ridge National Laboratory, U.S.

12:20-12:40 Learning Deep Neural Network Surrogate Models for High Dimensional and Multi-fidelity Uncertainty Quantification Problems

Rohit Tripathy and Ilias Bilionis, Purdue University, U.S.

Friday, March 1

MS375

Low Mach Number AMR Combustion Simulations with PeleLM - Part II of II

11:30 a.m.-1:10 p.m.

Room: 102D

For Part 1 see MS342

PeleLM is a code for evolving chemically reacting low Mach number flows with block-structured adaptive mesh refinement (AMR). The code features an iterative time step that efficiently couples together advection, diffusion and chemical reactions, each potentially evolving with different time scales, all while evolving the numerically conservative discretization across multiple levels of mesh refinement on a manifold where the equation of state is always satified. PeleLM is being used to study the details of turbulence-chemistry interaction in both premixed and diffusion flame configurations, from small-scale laboratory experiments to large-scale pool fires and wildfires, while it is being simultaneously improved and extended to incorporate new physics, and to run efficiently on emerging computing hardware platforms. In this set of talks, we outline the current capabilities of PeleLM via several different combustion-specific applications, and discuss improvements to the code that are under development and supported by DOE's Exascale Computing Project. These improvements include arbitrary (CAD-based) 3D geometries, electric field effects on charged particles, and incorporation of real gas and closed-chamber pressurization effects. We also discuss further extensions to the PeleLM time stepping scheme to support high-order methods, and implementation strategies used in PeleLM for many-core and GPU-based large-scale computing platforms.

Organizer: Marcus Day

Lawrence Berkeley National Laboratory, U.S.

11:30-11:50 An Efficient GPU Implementation for Time-implicit Integration of Arrhenius Combustion Kinetics

Anne Felden and Marcus Day, Lawrence Berkeley National Laboratory, U.S.

continued in next column

11:55-12:15 Towards the Distributed Burning Regime in Turbulent Premixed Flames

Andrew Aspden, Newcastle University, United Kingdom

12:20-12:40 Direct Numerical Simulations of Buoyancy-driven Flows Using PeleLM

Nicholas Wimer, University of Colorado Boulder, U.S.; Marcus Day, Lawrence Berkeley National Laboratory, U.S.; Amanda Makowiecki, Jeffrey Glusman, John Daily, Gregory Rieker, and Peter Hamlington, University of Colorado Boulder, U.S.

12:45-1:05 Direct Numerical Simulations of Compression Ignition Jet Flames

Evatt R. Hawkes, University of New South Wales, Australia

Friday, March 1

MS376

Data-augmented Reduced-order Modeling: Operator Learning and Closure/error Modeling - Part II of II

11:30 a.m.-1:10 p.m.

Room: 111A

For Part 1 see MS343

As the availability of experimental and computational data has exploded in recent years, an important question has arisen in computational science and engineering: How can these data be exploited to improve the predictive capacity of dynamical system models? This minisymposium focuses on reduced-order models (ROMs) of dynamical systems, which are attractive due to their low simulation cost. In this setting, we explore three emerging classes of techniques that leverage the availability of experimental and computational data to construct and improve ROMs: (1) methods that learn low-dimensional dynamical-system operators from data (e.g., dynamic mode decomposition, operator inference), and (2) closure and error-modeling techniques (e.g., Mori–Zwanzig) for ROMs.

Organizer: Kevin T. Carlberg

Sandia National Laboratories, U.S.

Organizer: Boris Kramer

Massachusetts Institute of Technology, U.S.

11:30-11:50 A Hybrid Approach for Model Order Reduction of Barotropic Quasi-geostrophic Turbulence

Sk. Mashfiqur Rahman and Omer San, Oklahoma State University, U.S.; Adil Rasheed, SINTEF Digital, Norway

11:55-12:15 Data-informed Reduced-order Models with Memory Effects

Eric Parish, Sandia National Laboratories, U.S.

12:20-12:40 Machine-learning Error Models for Rom Closure Via the ROMES Method

Stefano Pagani and Andrea Manzoni, Politecnico di Milano, Italy; Kevin T. Carlberg, Sandia National Laboratories, U.S.

12:45-1:05 Data-driven Correction for Reduced Order Modeling of Nonlinear Systems

Traian Iliescu and Muhammad Mohebujjaman, Virginia Tech, U.S.; Leo Rebholz, Clemson University, U.S.; Birgul Koc and Changhong Mou, Virginia Tech, U.S.

Friday, March 1

MS377

Inverse Problems in Machine Learning - Part II of II

11:30 a.m.-1:10 p.m.

Room: 111B

For Part 1 see MS344

Machine learning has been a fast growing area of research since the beginning of the 21st century due to the advent of computing resources and availability of data. Many of the questions of interest in machine learning can be formulated as inverse problems and vice versa. Therefore, it is not surprising that the developments in these two seemingly separate fields are in fact closely related. In this symposium we bring together researchers working at the intersection of machine learning and inverse problems to discuss recent advances, challenges and strategies in each area, aiming to promote cross-pollination and exchange of ideas that lead to new directions of research in both communities.

Organizer: Bamdad Hosseini

California Institute of Technology, U.S.

11:30-11:50 Randomized Methods for Regression-type Problems

Michael Mahoney and N. Benjamin Erichson, University of California, Berkeley, U.S.

11:55-12:15 Pdenetwork - The PDE's that Govern Deep Neural Networks

Eldad Haber, University of British Columbia, Canada

12:20-12:40 Generalized Graph Based Probit in the Continuum Limit

Bamdad Hosseini and Franca Hoffmann, California Institute of Technology, U.S.

12:45-1:05 Structured Bayesian Gaussian Process Latent Variable Model: Applications to Data-driven Dimensionality Reduction and High-dimensional Inversion

Steven Atkinson, University of Notre Dame, U.S.

Friday, March 1

MS378

Optimal Experimental Design for Inverse Problems - Part II of II

11:30 a.m.-1:10 p.m.

Room: 111C

For Part 1 see MS345

Good experimental data is critical for parameter inference in complex physical systems and enables subsequent robust decision making. The collection of useful experimental data, however, is arduous and expensive. It is therefore imperative to design experiments in an optimal manner, subject to physical or budgetary constraints while incorporating prior information. Leveraging simulation models, the computational cost of model-based optimal experimental design (OED) can be prohibitively expensive in many applications of interest. This minisymposium explores advances in numerical methods for OED of Inverse Problems.

Organizer: Xun Huan

University of Michigan, U.S.

Organizer: Jayanth Jagalur Mohan

Massachusetts Institute of Technology, U.S.

11:30-11:50 Sensitivity Analysis for Inverse Problems and Sensor Placement

Bart G. Van Bloemen Waanders, Sandia National Laboratories, U.S.; Joseph L. Hart, North Carolina State University, U.S.

11:55-12:15 Goal Oriented Information-Theoretic Observation Selection Strategies for Linear Bayesian Inverse Problems

Fengyi Li, Jayanth Mohan, and Youssef M. Marzouk, Massachusetts Institute of Technology, U.S.

12:20-12:40 Optimal Experimental Design under Model Uncertainty for Linear Inverse Problems

Karina Koval and Georg Stadler, Courant Institute of Mathematical Sciences, New York University, U.S.; Alen Alexanderian, North Carolina State University, U.S.

12:45-1:05 Selecting Multiple Borehole Locations for Maximizing Bayesian Information Gain on Past Ice Sheet Surface Temperatures

Andrew Davis, US Army Cold Regions Research and Engineering Lab (CRREL), U.S.; Xun Huan, University of Michigan, U.S.; Patrick Heimbach, University of Texas at Austin, U.S.; Youssef M. Marzouk, Massachusetts Institute of Technology, U.S.

continued in next column

Friday, March 1

MS379

Derivative-free and Global Optimization - Part II of II

11:30 a.m.-1:10 p.m.

Room: 300A

For Part 1 see MS346

Optimization problems arising in various science applications, such as engineering design, environmental applications, physics simulations, etc. require the development of new numerical optimization methods that efficiently and effectively find the optimal solutions. The objective functions are often multi-modal, computationally expensive, and black-box (their analytic descriptions and derivatives are not available). Often, the presence of stochasticity further increases the difficulty of solving these problems. New numerical optimization methods that do not rely on gradient information of the objective function and that are able to globally explore as well as locally search the parameter domain are therefore needed. The speakers in this minisymposium will present their recent developments of optimization methods that employ, among others, adaptive sampling methods, approximations of the objective functions, and the exploitation of any available problem structure to solve these difficult problems.

Organizer: Juliane Mueller
Lawrence Berkeley National Laboratory, U.S.

Organizer: Matt Menickelly
Argonne National Laboratory, U.S.

11:30-11:50 Derivative-free Robust Optimization

Matt Menickelly and Stefan Wild, Argonne National Laboratory, U.S.

11:55-12:15 A Stochastic Levenberg-Marquardt Method using Random Models with Application to Data Assimilation

El houcine Bergou, INRA, Jouy-en-Josas, France; Youssef Diouane, Institut Supérieur de l'Aéronautique et de l'Espace, France; Vyacheslav Kungurtsev, Czech Technical University, Prague, Czech Republic; *Clément W. Royer*, University of Wisconsin, Madison, U.S.

12:20-12:40 Scaling Up and Randomizing Derivative-free Optimization for Machine Learning

Albert S. Berahas, Northwestern University, U.S.; Katya Scheinberg, Lehigh University, U.S.; Jorge Nocedal, Northwestern University, U.S.; Richard H. Byrd, University of Colorado, U.S.

12:45-1:05 Complexity of ASTRO-DF: Adaptive Sampling Trust-region Methods for Derivative-free Simulation Optimization

Sara Shashaani, University of Michigan, U.S.; Raghu Pasupathy and Daniel Carvajal, Purdue University, U.S.

continued in next column

Friday, March 1

MS380

Advances in Global Sensitivity Analysis with Applications to Complex Systems - Part II of II

11:30 a.m.-1:10 p.m.

Room: 300B

For Part 1 see MS347

The traditional framework of global sensitivity analysis (GSA) considers scalar-valued quantities of interest (QoIs) that are functions of statistically independent inputs. The sensitivity of the QoIs to the inputs are determined using a number of classical methods such as Sobol' indices, Morris screening, or derivative based global sensitivity measures. To extend the applicability of GSA methods to broader classes of problems, such as problems with correlated inputs or problems with vectorial and functional outputs with high-dimensional parameters, new methods are needed. This minisymposium will highlight some of the recent advances in GSA and present recent generalizations and alternatives. Areas of particular interest are problems with dependent inputs, problems with vectorial or functional outputs, and applications of GSA in complex physical systems.

Organizer: Alen Alexanderian
North Carolina State University, U.S.

Organizer: Helen Cleaves
North Carolina State University, U.S.

Organizer: Pierre Gremaud
North Carolina State University, U.S.

11:30-11:50 Global Sensitivity Analysis for PDE-constrained Optimization

Joseph L. Hart, North Carolina State University, U.S.; Bart G. Van Bloemen Waanders, Sandia National Laboratories, U.S.

11:55-12:15 Application of Global Sensitivity Analysis and Monte Carlo Filtering for Constrained Models

Sergei S. Kucherenko, Dimitris Giamalakis, and Nilay Shah, Imperial College London, United Kingdom

continued on next page

12:20-12:40 Sensitivity Indices for Computer Codes Taking Values in General Metric Spaces

Ricardo Fraiman, Universidad de la República, Uruguay; Fabrice Gamboa, University of Toulouse, France; *Thierry Klein* and Agnès Lagnoux, Universite de Toulouse, France; Leonardo Moreno, Universidad de la República, Uruguay

12:45-1:05 Alternative to Sobol' Indices for the Multivariate Response Models

Matieyendou Lamboni, University of Guyane, France

Friday, March 1

MS381

Quantum Computing for Computational Science Applications

11:30 a.m.-1:10 p.m.

Room: 300C

Quantum computing is a paradigm of computation where data is encoded on the state of particles governed by the law of quantum physics. Prominent algorithms such as Shor's algorithm for integer factorization and Grover's search algorithm show that this computational viewpoint on vector spaces can be very powerful. From an algorithmic perspective, using a quantum computer amounts to finding a list of elementary operations admissible by the quantum memory which realize a desired unitary operation. The field of quantum computing recently witnessed a surge of interest from both academic and industrial actors. As a result, research on quantum algorithms shifted from a theoretical understanding of their complexity to a concrete study of their applications in computational science and engineering. This mini-symposium will focus on recent advances in the field of practical quantum computing. The speakers, coming from both academia and industry, will explain how scientific applications can benefit from quantum computing and will address the main issues in providing efficient quantum simulation algorithms.

Organizer: Marc Baboulin
University of Paris-Sud, France

Organizer: Neil J. Ross
Dalhousie University, Canada

Organizer: Benoît Valiron
CentraleSupelec, France

11:30-11:50 Accelerating Scientific Applications with Quantum Computing

Travis Humble and *Alexander McCaskey*, Oak Ridge National Laboratory, U.S.

11:55-12:15 Optimizing Quantum Circuit Synthesis Using Numerical Methods

Timothée Goubault de Brugière and Marc Baboulin, University of Paris-Sud, France; Benoît Valiron, CentraleSupelec, France; Cyril Allouche, ATOS, France

12:20-12:40 Seeking Mid-term Quantum Computing Supremacy with Numerical Simulation

Cyril Allouche, ATOS, France

12:45-1:05 Modern Methods in Quantum Compiling

Neil J. Ross, Dalhousie University, Canada

continued in next column

Friday, March 1

MS382

Multiscale and Domain Decomposition Approaches for PDEs with Rough Coefficients - Part II of II

11:30 a.m.-1:10 p.m.

Room: 302B

For Part 1 see MS350

Classical Lagrangian finite element methods are challenged by multi-scale problems due to the range of scales present in the problem. Over the last years, many discretization methods have been proposed to enable the accurate, efficient, and robust solution of these complex models. Examples include the Multi-scale Finite Element Method (and its variants, such as the Generalized Multi-scale Finite Element Method -- GMsFEM), the mixed multi-scale finite element, the Heterogeneous Multi-scale Method (HMM), multiscale methods based on a orthogonal decomposition of the space (such as the LOD), adaptive multi-scale methods, ... Such discretization methods often rely on approximation subspaces that incorporate specialized knowledge of the partial differential equation. Similar ideas appear in the context of preconditioners (or linear solvers) based on domain decomposition to build enhanced coarse spaces. The objective of this minisymposium is to gather presentations where approximation subspaces incorporating specialized knowledge of the multi-scale partial differential equation are used either for the discretization method and/or for the linear solver.

Organizer: Ulrich Hetmaniuk

University of Washington, U.S.

Organizer: Frederic Legoll

Ecole Nationale des Ponts et Chaussées, France

11:30-11:50 Sparse Compression of Expected Solution Operators

Daniel Peterseim, Universität Augsburg, Germany

11:55-12:15 Constraint Energy Minimizing Generalized Multiscale Finite Element for Flows in Heterogeneous Media

Tony Cheung, Texas A&M University, U.S.; *Wing Tat Leung* and Mary F. Wheeler, University of Texas at Austin, U.S.; Eric Chung, The Chinese University of Hong Kong, Hong Kong; Yalchin Efendiev, Texas A&M University, U.S.

12:20-12:40 On the Convergence Rates of GMsFEMs for Heterogeneous Elliptic Problems Without Oversampling Techniques

Guanglian Li, Imperial College London, United Kingdom

12:45-1:05 On Some New Variants of the Multiscale Finite Element Method

Frederic Legoll, Ecole Nationale des Ponts et Chaussées, France

continued in next column

Friday, March 1

MS383

Recent Advances in Numerical Methods for Multiphase Flow Problems - Part II of II

11:30 a.m.-12:45 p.m.

Room: 303A

For Part 1 see MS351

In nature and environment, rain, snow, fog, avalanches, mud slides, sediment transport, debris flows are all examples of multiphase flow where the behaviour of the phases are studied in different fields of natural science. Phase field method, or called diffusive interface approach, as a powerful numerical tool for simulating free interfacial motions, had been widely used in modeling and numerical approaches for multiphase flow problems. In this mini-symposium, the experts are expected to communicate the new idea, share their recent advances in the modeling, numerical Methods for, but not limited to, the phase field approach, as well as their applications in material science, fluid mechanics, etc.

Organizer: Xiaofeng Yang

University of South Carolina, U.S.

Organizer: Jia Zhao

Utah State University, U.S.

11:30-11:50 Numerical Analysis of a Variable Step Bdf2 Scheme for the Cahn-Hilliard Equation

Xiaoming Wang, Florida State University, U.S. and Fudan University, China; Wenbin Chen and Zhuying Zhang, Fudan University, China

11:55-12:15 Energy-stable Open Boundary Conditions and Associated Algorithm for Multiphase Flows of N Immiscible Incompressible Fluids

Zhiguo Yang and *Suchuan Dong*, Purdue University, U.S.

12:20-12:40 Super-convergence of the HDG Method for the Cahn-Hilliard Equation

Daozhi Han, Missouri University of Science and Technology, U.S.

Friday, March 1

MS384

Recent Developments in Model Order Reduction Methods - Part II of II

11:30 a.m.-1:10 p.m.

Room: 303B

For Part 1 see MS352

An increase in demands for simulations of complex systems has motivated active developments in model order reduction. We discuss several recent developments in model order reduction methods addressing challenges such as high dimensional parameter space, expensive offline phase. We also discuss how model reduction algorithms are developed in specific domains such as fluid control of quadratic systems, advection-dominated high-gradient structure systems, coupled flow-geomechanics problems, finite time Lyapunov exponents, and poroelastic media.

Organizer: Youngsoo Choi

Lawrence Livermore National Laboratory, U.S.

11:30-11:50 Reduced-order Modeling of Coupled Flow-geomechanics Problems

Louis J. Durlofsky, Zhaoyang Larry Jin, Timur T. Garipov, and Oleg Volkov, Stanford University, U.S.

11:55-12:15 Lagrangian Data-driven Reduced Order Modeling of Finite Time Lyapunov Exponents.

Changhong Mou, Virginia Tech, U.S.; Xuping Xie, Oak Ridge National Laboratory, U.S.; Peter Nolan, Shane D. Ross, and Traian Iliescu, Virginia Tech, U.S.

12:20-12:40 Multiscale Two-stage Solver for Poroelastic Media

Sergey Klevtsov, Stanford University, U.S.; Nicola Castelletto, Lawrence Livermore National Laboratory, U.S.; Hamdi Tchelepi, Stanford University, U.S.

12:45-1:05 Hierarchical Solvers for Parametric Problems

Simona Perotto, Politecnico di Milano, Italy

Friday, March 1

MS385

Mathematics of Energy Materials - Part II of II

11:30 a.m.-1:10 p.m.

Room: 201C

For Part 1 see MS353
Featured Minisymposium

Dramatically improved performance of energy storage devices, e.g., batteries and capacitors, is a prerequisite for successful energy transition. Porous materials are a key component in most of such devices. This minisymposium deals with innovative approaches to modeling, simulations, and design of energy storage devices, including development of new porous materials.

Organizer: Ilenia Battiato

Stanford University, U.S.

Organizer: Daniel M. Tartakovsky

Stanford University, U.S.

Organizer: John H. Cushman

Purdue University, U.S.

11:30-11:50 A New Type of Flow Battery Based on Oxidation State Instability

John H. Cushman, Purdue University, U.S.

11:55-12:15 Small-scale Acoustofluidics in Rechargeable Battery Technologies

James Friend, University of California, San Diego, U.S.

12:20-12:40 A Diffuse Interface Method for Electrochemical Systems

Katsuyo Thornton, University of Michigan, U.S.

12:45-1:05 Dimensional Considerations in Probing the Upper Limits of Energy Storage in Nanostructured Electrochemical Devices

Hidenori Yamada, Rajaram Narayan, and *Prabhakar R. Bandaru*, University of California, San Diego, U.S.

Friday, March 1

MS386

Next Generation FFT Algorithms in Theory and Practice: Parallel Implementations, Sparse FFTs, and Applications - Part II of II

11:30 a.m.-1:10 p.m.

Room: 202A

For Part 1 see MS354

The fast Fourier Transform (FFT) is an algorithm used in a wide variety of applications, yet does not make optimal use of many current hardware platforms. Hardware utilization performance on its own does not however imply optimal problem solving. The purpose of this minisymposium is to enable exchange of information between people working on alternative FFT algorithms such as sparse and non uniform FFTs, to those working on FFT implementations, in particular for parallel hardware.

Organizer: Daisuke Takahashi

University of Tsukuba, Japan

Organizer: Mark Iwen

Michigan State University, U.S.

Organizer: Samar A. Aseeri

King Abdullah University of Science & Technology (KAUST), Saudi Arabia

Organizer: Benson K. Muite

University of Tartu, Estonia

11:30-11:50 FFT Applications and Benchmarks

Samar A. Aseeri, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

11:55-12:15 High-dimensional Sparse FFT

Bosu Choi, University of Texas at Austin, U.S.; Andrew J. Christlieb, Michigan State University, U.S.; Yang Wang, Hong Kong University of Science and Technology, Hong Kong

12:20-12:40 Rank-1 Lattice Based High-dimensional Approximation and FFT

Toni Volkmer, Chemnitz University of Technology, Germany

12:45-1:05 A Periodic Treecode Method for Electrostatics In Molecular Dynamics Simulations

Henry A. Boateng, Bates College, U.S.

Friday, March 1

MS387

High-order Discretizations and Quadrature for Integral Equation Methods - Part II of II

11:30 a.m.-1:10 p.m.

Room: 202B

For Part 1 see MS355

Integral equation methods are a powerful set of techniques for simulating physical systems accurately and efficiently. They have been employed with great success in the design of devices in many fields such as audio equipment, microfluidic devices, photonic devices, and biomedical equipment. Recent years have seen tremendous improvement in the precision and complexity of devices which can be manufactured necessitating the development of new computationally-efficient high-precision numerical tools for complex geometries. In integral equation methods the solutions are typically represented as singular integrals, requiring efficient discretization and quadrature methods. Existing methods include skeletonization based approaches, generalized Gaussian quadrature, and quadrature by expansion. When coupled with fast algorithms the resulting discretizations yield solvers scaling linearly in the number of degrees of freedom. Two critical avenues of research are improving performance in complex geometries and in three dimensions, as well as developing automatically adaptive discretizations which resolve complexity in geometry and input data. The talks in this mini-symposium are focused on the construction of efficient Nystrom methods and associated high-order quadrature methods for surfaces in three dimensions; the design of automatically adaptive discretizations in moving geometries; and the detailed analysis of singularities of solutions near corners and edges.

Organizer: Manas N. Rachh
Simons Foundation and Flatiron Institute, U.S.

Organizer: Jeremy Hoskins
Yale University, U.S.

continued in next column

Organizer: Kirill Serkh
Yale University, U.S.

11:30-11:50 A Fast Boundary Integral Method for Generating High-order Surface Meshes

Felipe Vico, Universidad Politecnica de Valencia, Spain

11:55-12:15 Analytical Expressions for the Solutions of Elliptic PDEs Near Corners, Edges, and Conical Points

Kirill Serkh, Yale University, U.S.

12:20-12:40 Multiple Junctions and Transmission Problems in Elliptic PDEs

Jeremy Hoskins, Yale University, U.S.

12:45-1:05 Optimized Quadratures for the Sommerfeld Representations of Acoustic and Electromagnetic Fields

Zydrunas Gimbutas, National Institute of Standards and Technology, U.S.

Friday, March 1

MS388

Modeling Resource Utilization and Contention in HPC System-application Interactions - Part II of II

11:30 a.m.-1:10 p.m.

Room: 202C

For Part 1 see MS356

Application performance, system throughput and stability, and procurement specifications can be improved by better understanding of applications' resource requirements and the performance impacts of resource contention within a site's workload. Production data can potentially provide the most accurate insights, however, quantification of resource utilization and contention impact may be difficult empirically because relationships between such data and performance are largely inferential and controlled conditions for experimentation are difficult to obtain in production. Conversely, theoretical models developed in the design phase of large-scale systems have missed large-scale contention scenarios and may not capture the range of behavior of time-dependent production workloads. Advances in instrumentation and large data analytical and visual techniques can improve our ability to develop data-driven models with the fidelity to capture events of interest. Efficient techniques for assessing near-term system behavior with respect to models can enable runtime responses to improve operations and be used for model feedback to improve confidence. We bring together experts to discuss current techniques and experiences in developing and validating models and characterizations of resource utilization and contention; exploring data to determine and affirm model dependencies; and extracting insights from models to dynamically effect improved performance and operations.

Organizer: James Brandt
Sandia National Laboratories, U.S.

continued on next page

Organizer: Ann Gentile
Sandia National Laboratories, U.S.

Organizer: Zbigniew Kalbarczyk
University of Illinois at Urbana-Champaign, U.S.

11:30-11:50 Characterization of Contention and Failure in Large-scale Systems
Zbigniew Kalbarczyk, University of Illinois at Urbana-Champaign, U.S.

11:55-12:15 Rethinking HPC Performance and the Role of Learning-based Analytics
Emre Ates and Ayse K. Coskun, Boston University, U.S.

12:20-12:40 Performance Modeling Experience with MILC on Blue Waters
Greg Bauer, National Center for Supercomputing Applications (NCSA), U.S.; Celso Mendes and William Kramer, University of Illinois, U.S.

12:45-1:05 A Survey of Analysis Techniques Relevant to HPC Resource Modeling
David Thompson and *Thomas Hastings Greer*, Kitware, Inc., U.S.

Friday, March 1

MS389

Divide and Conquer Strategies for Large-scale Eigenvalue Problems - Part II of II

11:30 a.m.-1:10 p.m.

Room: 203

For Part 1 see MS357

This minisymposium presents recent advances in large-scale eigenvalue computations based on divide-and-conquer strategies. These new methodologies aim to handle problems that are often found challenging for traditional algorithms, namely, extracting a large number of eigenvalues from a large and sparse matrix, or extracting eigenvalues that are deep inside the spectrum. The presented techniques include spectrum slicing methods, polynomial and rational filtering, algorithms based on contour integration, etc, along with Krylov subspace methods or subspace iterations. We bring together numerical analysts who study these techniques with different methodologies and also physicists who initiate an effort to assemble the different ideas to solve problems in real applications.

Organizer: Ruipeng Li
Lawrence Livermore National Laboratory, U.S.

Organizer: Yuanzhe Xi
Emory University, U.S.

Organizer: Edoardo A. Di Napoli
Jülich Supercomputing Centre, Germany

11:30-11:50 The Chase Library for Large Hermitian Eigenvalue Problems
Edoardo A. Di Napoli, Jülich Supercomputing Centre, Germany

11:55-12:15 Stopping Criteria for Computing Low Rank Approximations with Iterative Methods
Andreas Stathopoulos and Eloy Romero, College of William & Mary, U.S.

12:20-12:40 Computing Planetary Interior Normal Modes with a Highly Parallel Polynomial Filtering Eigensolver
Jia Shi, Rice University, U.S.; Ruipeng Li, Lawrence Livermore National Laboratory, U.S.; Yuanzhe Xi, Emory University, U.S.; Yousef Saad, University of Minnesota, U.S.; Maarten de Hoop, Rice University, U.S.

12:45-1:05 Inexact Methods for Symmetric Stochastic Eigenvalue Problems
Kookjin Lee, Sandia National Laboratories, U.S.; Bedrich Sousedik, University of Maryland, Baltimore County, U.S.

continued in next column

Friday, March 1

MS390

Advances in Multi-method Time Discretizations of Evolutionary PDEs - Part II of II

11:30 a.m.-1:10 p.m.

Room: 205

For Part 1 see MS358

Simulations of systems modeled by large time dependent partial differential equations faces challenges posed by the interaction of multiple physical processes and the existence of multiple dynamical scales. Time discretization is an essential ingredient of the solution process that determines the overall accuracy and stability of the solution, the computational efficiency, and the data dependencies in a parallel environment. This minisymposium highlights recent developments in time stepping methods for large evolutionary PDEs. Topics include, but are not restricted to, partitioned methods for multiphysics systems, multirate schemes, approaches based on a minimal amount of implicitness to achieve stability without sacrificing computational efficiency, methods that preserve special properties of the underlying system (such as monotonicity), and approaches that improve the scalable parallel performance in the context of large applications.

Organizer: Emil M. Constantinescu

Argonne National Laboratory, U.S.

11:30-11:50 High-order Operator-Splitting Methods for Cardiac Tissue Simulations

Raymond J. Spiteri, University of Saskatchewan, Canada

11:55-12:15 Stable Time Integration for Coupled Ocean-atmosphere Models

Hong Zhang, Argonne National Laboratory, U.S.; Paula Egging, University of Nebraska, Lincoln, U.S.; Emil M. Constantinescu and Robert Jacob, Argonne National Laboratory, U.S.

12:20-12:40 Application of Multirate Methods to Multiphysics Problems

John Loffeld, Carol S. Woodward, and Slaven Peles, Lawrence Livermore National Laboratory, U.S.

12:45-1:05 Implicit Multirate Generalized Additive Runge-Kutta Methods

Steven Roberts, Arash Sarshar, and Adrian Sandu, Virginia Tech, U.S.

Friday, March 1

MS391

Computational Methods for Data Assimilation and Uncertainty Quantification - Part II of II

11:30 a.m.-1:10 p.m.

Room: 206A

For Part 1 see MS359

Model-based computer simulations about large-scale physical phenomena such as the atmosphere provide invaluable, yet uncertain, prediction of the evolution of the physical system. The model forecasts are corrected based on other available sources of information including measurements of reality, i.e. observations. Observational data are fused with model-based forecast, through the process of data assimilation, to estimate the trajectory or state of dynamical systems. This method has been extensively used in meteorology, hydrology, oceanography among other applications, to make accurate predictions about the state of these systems. The quality of the output of a data assimilation procedure is critically dependent on the data acquisition strategy, and thus an optimal experimental design is required for sensor placement. The focus of this minisymposium is on recent advances in computational methods for data assimilation, uncertainty quantification, sensor placement, and data acquisition.

Organizer: Emil M. Constantinescu

Argonne National Laboratory, U.S.

Organizer: Vishwas Rao

Argonne National Laboratory, U.S.

11:30-11:50 Assimilating Data in Models with Stochastic Parameters

Emil M. Constantinescu, Argonne National Laboratory, U.S.

11:55-12:15 Prior and Posterior Inflation for Ensemble Filters: Theoretical Formulation and Application to Community Atmosphere Model

Mohamad El-Gharamti, Kevin Raeder, and Jeffrey Anderson, National Center for Atmospheric Research, U.S.; Xuguang Wang, University of Oklahoma, U.S.

12:20-12:40 Adaptive Covariance Tuning in Hybrid Ensemble-variational Data Assimilation

Dacian N. Daescu, Portland State University, U.S.

12:45-1:05 Variational Data Assimilation with Model Errors

Vishwas Rao, Emil M. Constantinescu, and Julie Bessac, Argonne National Laboratory, U.S.; Adrian Sandu, Virginia Tech, U.S.

continued in next column

Friday, March 1

MS392

Numerical Methods for Biological Fluid Dynamics - Part II of II

11:30 a.m.-1:10 p.m.

Room: 206B

For Part 1 see MS360

Fluid flows are essential to life across a wide range of length and time scales, from the swimming of whales in a turbulent ocean environment to the motion of proteins inside of cells. The behaviors exhibited by biological fluids are rich and counterintuitive, given non-Newtonian effects generated by suspended structures that may be passively elastic or actively forcing the fluid out of equilibrium. Dealing with this complex and varied phenomenology requires original computational approaches. In this minisymposium we present a selection of promising recent developments in the field, including fluid-structure interaction and coarse-grained active matter models, and applications such as jellyfish swimming and intracellular motion.

Organizer: Thomas Fai

Brandeis University, U.S.

Organizer: Chris H. Rycroft

Harvard University, U.S.

11:30-11:50 Flagellar Swimming in Complex Fluids

Robert D. Guy, University of California, Davis, U.S.

11:55-12:15 Dynamics of Elastic Filaments in Viscous Fluids

Mehdi Jabbarzadeh and Henry Fu, University of Utah, U.S.

12:20-12:40 Simulating the Swimming Motion of \textit{C. elegans} and Amoeboids in Viscoelastic Fluids via the Immersed Boundary Method

Jeremy Binagia, Christopher Guido, and Eric S. Shaqfeh, Stanford University, U.S.

12:45-1:05 A Regularized Stokeslet Approach to the Motion and Behavior of Microscopic Swimmers

Alexander Hoover, University of Akron, U.S.; Ricardo Cortez, Tulane University, U.S.

Friday, March 1

MS393

Structure-exploiting Techniques for Approximation, Inference and Control of Complex Systems - Part II of II

11:30 a.m.-1:10 p.m.

Room: 206C

For Part 1 see MS361

Data-driven modeling and control of complex systems have received a lot of attention over the recent years, largely due to the availability of large datasets. Nevertheless, large scale data is not typically informative about the entire physical properties because it is often gathered through nonlinear or partially observable processes. In such situations, physically-principled priors that include symmetries and constraints are crucial to infer predictive models, typically with filtering or smoothing techniques. This minisymposium explores the latest developments and challenges in data-efficient techniques for modeling, inference and control of complex physical systems. This includes, but is not limited to, scalable approximation formats and spaces, algorithms for filtering and smoothing, sampling and exploration strategies for gathering "smart data", incorporation of symmetry-based priors, bounds and invariant preservation, information theoretic arguments, etc.

Organizer: Alex Gorodetsky

University of Michigan, U.S.

Organizer: Lionel Mathelin

CNRS, France

11:30-11:50 A Bayesian Framework for Robust Decisions in the Presence of Unobserved Heterogeneity

Chi Feng and Youssef M. Marzouk, Massachusetts Institute of Technology, U.S.

11:55-12:15 Representing Model Inadequacy in Interacting Systems with Constrained Stochastic Operators

Rebecca E. Morrison, University of Colorado Boulder, U.S.; Youssef M. Marzouk, Massachusetts Institute of Technology, U.S.

12:20-12:40 Model-Order Reduction for Data Assimilation from Partial Observations

Cedric Herzet and Patrick Heas, Inria Rennes Bretagne Atlantique, France; Ronan Fablet, IMT - Atlantique, France

12:45-1:05 Improving Stability of Numerical Methods for Recovering Governing Equations from Noisy Data

Paul Diaz and Alireza Doostan, University of Colorado Boulder, U.S.

continued in next column

Friday, March 1

MS394

Structured Matrix Methods - Part II of II

11:30 a.m.-1:10 p.m.

Room: 206D

For Part 1 see MS362

Structured matrix methods utilize structures such as parameterized forms and low-rank approximations to gain high efficiency in matrix representation and arithmetic. They have demonstrated significant advantages in improving the efficiency and reliability of some large-scale computations and engineering simulations. Rank structured methods extend the fundamental ideas of multipole and panel-clustering methods to general non-local solution operators. While there exist various more or less closely related methods, the unifying aim of these methods is to explore efficient structured low-rank approximations, especially those exhibiting hierarchical or nested forms. These help the methods to achieve nearly linear complexity. In this minisymposium, we aim to present and exchange recent new developments on structured methods for some challenging numerical problems such as high frequencies, ill conditioning, eigenvalue solution, and stability. Studies of structures, algorithm design, and accuracy control will be discussed. The minisymposium will include experts working on a broad range of structured methods.

Organizer: Sabine Le Borne
Hamburg University of Technology, Germany

Organizer: Jianlin Xia
Purdue University, U.S.

11:30-11:50 On the Accuracy and Stability of Hierarchical Structured Approximate Factorizations
Jianlin Xia, Purdue University, U.S.; Yuanzhe Xi, Emory University, U.S.; Zixing Xin, Purdue University, U.S.

11:55-12:15 Identifying 2D SSS Representations
Shivkumar Chandrasekaran, University of California, Santa Barbara, U.S.

12:20-12:40 Robust and Efficient QR Decomposition of HODLR Matrices with Applications
Daniel Kressner and Ana Susnjara, École Polytechnique Fédérale de Lausanne, Switzerland

12:45-1:05 Generalizations to the Superfast Divide-and-Conquer Eigenvalue Algorithm
James Vogel, Purdue University, U.S.

continued in next column

Friday, March 1

MS395

Summation-by-Parts: A Framework for the Development and Analysis of Modern Numerical Methods - Part II of II

11:30 a.m.-1:10 p.m.

Room: 207

For Part 1 see MS363
Featured Minisymposium

As we move towards an era of exascale HPC, flexible and robust algorithms will be necessary to take advantage of increasingly complex and failure prone computational resources. In this mini-symposium, the focus is on methods having the summation-by-parts (SBP) property, which are advantageous as they lead to high-order methods that are provably conservative and stable. In particular, the interest is in recent advancements leading to improved efficiency, flexibility, and robustness. Topics that will be covered include:

• Entropy stable algorithms
• Linearly stable algorithms h/p-adaptation
• Mixed formulations
• Wave propagation and inversion problems
• Optimized differentiation matrices

Organizer: David C. Del Rey Fernandez
NASA Langley Research Center and National Institute of Aerospace, U.S.

Organizer: Andrew R. Winters
University of Cologne, Germany

Organizer: Jesse Chan
Rice University, U.S.

Organizer: Jason E. Hicken
Rensselaer Polytechnic Institute, U.S.

continued on next page

11:30-11:50 Entropy Stable Discontinuous Galerkin Schemes on Moving Meshes for Hyperbolic Conservation Laws

Gero Schnuecke, Andrew R. Winters, and Gregor Gassner, University of Cologne, Germany

11:55-12:15 Recent Developments in the Convergence Theory of Numerical Solutions for Compressible Flows

Magnus Svard, University of Bergen, Norway

12:20-12:40 Summation-by-Parts Methods for Inverse Problems in Exploration Seismology

Martin Almquist, Eric M. Dunham, and Joseph Jennings, Stanford University, U.S.

12:45-1:05 Optimization of Free Parameters in Cubature Rules and Accuracy Equations to Develop Optimal Multi-Dimensional Summation-by-Parts Operators

Andre Marchildon and David Zingg, University of Toronto Institute for Aerospace Studies, Canada

Friday, March 1

MS396

Bringing Algorithms and New Applications to Novel Architectures - Part II of II

11:30 a.m.-1:10 p.m.

Room: 401A

For Part 1 see MS364

Computational science and engineering now reaches beyond HPC simulation to diverse applications in machine learning and data analysis. Common architectures still focus on dense matrix multiplication. The result: Engineering solutions that require too much power and too much time. But the success of one particular matrix multiplication engine, the GPGPU, has opened people to design issues around "alternative" architectures. Improved software tools, larger flexible architectures (e.g. FPGAs), and lower fabrication costs have lowered the bar to entry. This minisymposium focuses on novel architectures addressing the new application challenges. These are real architectures and real systems for real problems. They provide testing grounds for applications, software frameworks, and our preconceptions on building both. We consider a range from novel memory systems and FPGAs to neuromorphic and quantum systems.

Organizer: Jason Riedy
Georgia Institute of Technology, U.S.

Organizer: Jeffrey Young
Georgia Institute of Technology, U.S.

11:30-11:50 Heterogeneous Computing in ORNL's Experimental Computing Laboratory (excl)

Jeffrey S. Vetter, Oak Ridge National Laboratory, U.S.

11:55-12:15 Early Application Successes on D-Wave Quantum Computers

Steve P. Reinhardt, D-Wave Systems, Inc., U.S.

12:20-12:40 Unstructured Computations on Emerging Architectures

Mohammed A. Al Farhan and David E. Keyes, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

12:45-1:05 A Variable Neighbourhood Descent Heuristic for Conformational Search using a Quantum Annealer

Brad Woods, Dominic Marchand, Gili Rosenberg, and Moslem Noori, 1QBit, Canada

continued in next column

Friday, March 1

MS397

High-order Discontinuous Galerkin and Finite Element Methods for CFD - Part II of II

11:30 a.m.-1:10 p.m.

Room: 401B

For Part 1 see MS365

Given recent developments in computational architectures, high-order discontinuous Galerkin and finite element methods are increasingly important for a range of applications in computational fluid dynamics. Their low dissipation, compact stencils, and high ratio of computation to communication make them well-suited for solving flow problems on next-generation many-core architectures. However, despite these desirable features, high-order methods suffer from robustness issues in the presence of discontinuous features, sharp gradients, and underresolution, which can limit their utility in resolving challenging flow features such as shock waves and turbulence. Furthermore, the efficient implementation of such methods on modern, massively-parallel, memory-constrained architectures presents interesting and important new challenges, prompting the development of sophisticated solver technology and matrix-free operators. This minisymposium aims to bring together researchers in the broad field of high-order discontinuous Galerkin and finite element methods to disseminate their research and discuss recent developments in discretizations, solvers, and software that aim to overcome the challenges facing the field. Specific topics of interest include shock capturing and tracking methods, stable and robust discretizations, efficient solvers and time integration (implicit, explicit, and IMEX), matrix-free algorithms, and algorithms for advanced computer architectures.

Organizer: Will Pazner
University of California, Berkeley, U.S.

Organizer: Matthew J. Zahr
University of California, Berkeley and Lawrence Berkeley National Laboratory, U.S.

continued in next column

11:30-11:50 Entropy Stable Gauss Collocation Discontinuous Galerkin Methods

Jesse Chan, Rice University, U.S.

11:55-12:15 A Moving Discontinuous Galerkin Finite Element Method with Interface Condition Enforcement for Viscous Flows

Andrew Kercher, Andrew Corrigan, and David Kessler, Naval Research Laboratory, U.S.

12:20-12:40 A High-order Fully Implicit Incompressible Navier-Stokes DG Solver

Marian Piatkowski, Heidelberg University, Germany

12:45-1:05 Efficient High-order Discontinuous Galerkin Methods on Modern Architectures

Yohann Dudouit, Lawrence Livermore National Laboratory, U.S.; Johann Dahm, IBM Research, U.S.; Steve Rennich, NVIDIA, U.S.; Tzanio Kolev, Veselin Dobrev, and Jean-Sylvain Camier, Lawrence Livermore National Laboratory, U.S.

Friday, March 1

MS398

Hybrid Parallelization for Modern Architectures - Part II of II

11:30 a.m.-1:10 p.m.

Room: 401C

For Part 1 see MS366

Compute clusters are evolving away from compositions of many nodes that contain a few processors with a small amount of shared memory towards compositions of fewer nodes that contain many processors and a large amount of shared memory. This transition in computational architecture necessitates the development of hybrid parallelization implementations of scientific and engineering software to efficiently and flexibly utilize distributed and shared resources (processors and memory). Speakers will discuss algorithmic choices and implementation trade-offs for established and innovative methods. The impacts on computational cost (memory, CPU hours, and wall time) for production class simulations and implications for scaling to leadership class simulations will also be discussed.

Organizer: Stephen L. Wood
NASA Langley Research Center, U.S.

Organizer: Eric Nielsen
NASA Langley Research Center, U.S.

11:30-11:50 Unstructured-grid CFD Algorithms for a Many-core Landscape

Eric Nielsen and Aaron Walden, NASA Langley Research Center, U.S.; M Zubair, Old Dominion University, U.S.; Justin Luitjens, NVIDIA, U.S.

11:55-12:15 Accelerating Dynamic Load Balancing with GPUs

Cameron Smith, Gerrett Diamond, Lucas Davis, and Mark S. Shephard, Rensselaer Polytechnic Institute, U.S.

12:20-12:40 On Directives-based Refactoring of Legacy CFD Codes for Accelerator-based Systems

Christopher P. Stone and *Bracy H. Elton*, HPCMP PETTT, Engility Corp., U.S.

12:45-1:05 Design of a Portable, Multithreaded, and Generic C/C++ Library for SIMD Processing

Eduardo Ponce and Gregory D. Peterson, University of Tennessee, U.S.

Friday, March 1

MS399

State-of-the-Art Auto-tuning: New Approaches and Algorithmic Innovations Towards Exascale Computing - Part II of II

11:30 a.m.-1:10 p.m.

Room: 402A

For Part 1 see MS367

Exascale computers are expected to be deployed by 2021. Their architectures will be complex, building upon many-cores and GPUs, and offering unprecedented levels of parallelism. It is expected that auto-tuning (AT) research and technology will continue building upon its proven success for delivering high performance on a variety of computer architectures, and will enable optimized, high performance implementations of specific computations for those challenging architectures. For this, preliminary results suggest that the use of machine learning techniques is promising. In concert with the hardware evolution and complexity, algorithmic innovations will need to be devised for numerical computations that are essential to many applications, such as eigenvalue solvers and linear equation solvers. For example, reducing communications together with high performance implementations will be essential. While exascale computers will enable the solution of problems with unprecedented levels of details, the accuracy of the computations cannot be taken for granted: for example, codes that perform floating point sums of distributed data can produce results that are difficult to be reproduced. Given potentially conflicting goals for achieving performance and reproducibility, this minisymposium will discuss technology trends for AT and their interplay with AT frameworks, numerical algorithms, and accuracy assurance towards extreme levels of computing.

Organizer: Takahiro Katagiri
Nagoya University, Japan

Organizer: Osni A. Marques
Lawrence Berkeley National Laboratory, U.S.

continued in next column

Organizer: Toshiyuki Imamura
RIKEN, Japan

11:30-11:50 High Performance Eigensolver Exploiting an Online Tuning Mechanism

Toshiyuki Imamura, RIKEN, Japan

11:55-12:15 Test Matrices for Numerical Computations of Linear Systems and Eigenvalue Problems

Katsuhisa Ozaki, Shibaura Institute of Technology, Japan; Takeshi Ogita, Tokyo Woman's Christian University, Japan

12:20-12:40 A Contour Integral Eigensolver with Divide-and-Conquer and Locking Techniques for Organic Material Simulations

Yuhsiang M. Tsai, National Taiwan University, Taiwan; Takeo Hoshi, Tottori University, Japan; Weichung Wang, National Taiwan University, Taiwan

12:45-1:05 FFTX and SpectralPack: A First Look

Franz Franchetti, Carnegie Mellon University, U.S.

Friday, March 1

MS400

Meshfree Methods: Computational Advances and Applications - Part II of II

11:30 a.m.-1:10 p.m.

Room: 402B

For Part 1 see MS368

Meshfree (or meshless) methods offer geometric flexibility, non-uniform resolution, and advantageous trade-offs between accuracy and computational costs. These properties make meshfree methods particularly useful for large-scale science and engineering applications featuring complex geometries and multiscale phenomena. This minisymposium focuses on computational advancements and applications of various meshfree discretization methods, including radial basis function generated finite differences (RBF-FD), partition of unity, and generalized moving least squares (GMLS). Recent computational advancements in scalable implementations, adaptive refinement, approximation space enrichment, stabilization, non-linear problems, and conservation techniques will be discussed together with applications to partial differential equations (PDEs) on manifolds, biological and geophysical fluid flows, and non-linear elasticity.

Organizer: Grady B. Wright
Boise State University, U.S.

11:30-11:50 Well-balanced Meshless Methods for the Shallow-water Equations

Jörn Behrens, University of Hamburg, Germany; Alex Bihlo and Rüdiger Brecht, Memorial University of Newfoundland, Canada; *Scott Maclachlan*, Memorial University, Newfoundland, Canada

11:55-12:15 Meshfree Simulation of Human Respiratory Muscles

Nicola Cacciani, Karolinska Institutet, Sweden; Elisabeth Larsson and *Igor Tominec*, Uppsala University, Sweden; Pierre-Frederic Villard, University of Lorraine, France

continued on next page

Friday, March 1

MS400

Meshfree Methods: Computational Advances and Applications - Part II of II

continued

12:20-12:40 Stable Computations with Flat Anisotropic Gaussians Using Hermite Polynomials

Anna Yurova and Katharina Kormann, Max Planck Institute for Plasma Physics, Germany; Caroline Lasser, Technische Universität München, Germany

12:45-1:05 An Adaptive Least Squares Radial Basis Function Partition of Unity Method for PDEs

Elisabeth Larsson, Uppsala University, Sweden; Danilo Stocchino and Stefano De Marchi, University of Padova, Italy

Friday, March 1

MS401

Molecular and Electronic Structure Theory Meets Data Science - Part II of II

11:30 a.m.-1:10 p.m.

Room: 402C

For Part 1 see MS369

The prediction of molecular and electronic structures using first principle calculations are among the most challenging and computationally demanding science and engineering problems. This minisymposium aims at presenting and discussing new developments of using machine learning techniques to reduce the computational cost and characterize important features and properties of materials and molecular systems. We bring together experts on molecular and electronic structure theories, which include mathematicians, physicists and chemists working actively in the field.

Organizer: Chao Yang

Lawrence Berkeley National Laboratory, U.S.

Organizer: Lin Lin

University of California, Berkeley and Lawrence Berkeley National Laboratory, U.S.

11:30-11:50 Accelerating Molecular Dynamics with Multi-scale Neural Networks

Leonardo Zepeda-Nunez, Lawrence Berkeley National Laboratory, U.S.; Lin Lin, University of California, Berkeley and Lawrence Berkeley National Laboratory, U.S.; Lexing Ying and Yuwei Fan, Stanford University, U.S.

11:55-12:15 Advances in Machine Learned Potentials for Molecular Dynamics Simulation

Kipton Barros, Nicholas Lubbers, and Justin S. Smith, Los Alamos National Laboratory, U.S.

12:20-12:40 Toward the Systematic Generation of Hypothetical Atomic Structures: Neural Networks and Geometric Motifs

Tess Smidt, Lawrence Berkeley National Laboratory, U.S.

12:45-1:05 Predicting Molecular Energies Using Machine-learned Electron Densities

Leslie Vogt-Maranto, New York University, U.S.

Friday, March 1

CP21

Computational Mechanics

11:30 a.m.-12:50 p.m.

Room: 301

Chair: Andrew Reid, National Institute of Standards and Technology, U.S.

11:30-11:45 An Elastic Linear Deformation Model with Discrete Fractures at Core Scale

Martin A. Diaz-Viera, Manuel Coronado, Edscott Wilson-Garcia, Pedro Aguilar-Gastelum, and Mario Vadillo-Saenz, Instituto Mexicano del Petróleo, México

11:50-12:05 Numerical Methods for Thermally Stressed Shallow Shell Equations

Hangjie Ji, University of California, Los Angeles, U.S.; *Longfei Li*, University of Louisiana at Lafayette

12:10-12:25 Plasticity in the Oof Finite-Element Solver for Materials Science

Andrew Reid, National Institute of Standards and Technology, U.S.

12:30-12:45 A Projection Method for Simulating Incompressible Poroelastic Media

Chris H. Rycroft and Nicholas Derr, Harvard University, U.S.; Christoph Weber, Max Planck Institute for Physics of Complex Systems, Germany; L Mahadevan, Harvard University, U.S.

Notes

CSE19 Abstracts

SIAM Conference on
Computational Science and Engineering

February 25-March 1, 2019
Spokane Convention Center
Spokane, Washington, U.S.

IP1

Modelling 100 Percent Renewable Electricity

In the past few decades, power grids across the world have become dependent on markets that aim to efficiently match supply with demand at all times via a variety of pricing and auction mechanisms. These markets are based on models that capture interactions between producers, transmission and consumers. The recent explosion in the use of renewable supply such as wind, solar and hydro has led to increased volatility in this system at a number of different time scales. Nevertheless, various countries have announced target dates for a 100% renewable electricity system. Such targets require long-term investment planning, medium-term storage management (including batteries and pumped storage), as well as a short-term analysis of demand-response, involuntary load curtailment and transmission congestion. We develop models that aim to ensure enough generation capacity for the long term under various constraints related to environmental concerns, and consider the recovery of costs for this enhanced infrastructure. We demonstrate how risk can impose significant costs on the system that are not modeled in the context of socially optimal power system markets and highlight the use of contracts to reduce or recover these costs. We discuss how models can be used to compare several '100% renewable' investment scenarios, and consider the performance in terms of investment and fuel costs, emissions, electricity prices, and lost-load.

Michael C. Ferris
University of Wisconsin
Department of Computer Science
ferris@cs.wisc.edu

IP2

Deep Learning for Inverse Problems - Some Recent Approaches

In this talk we discuss the idea of data-driven regularisers for inverse imaging problems. We are in particular interested in the combination of model-based and purely data-driven image processing approaches. In this context we will make a journey from shallow learning for computing optimal parameters for variational regularisation models by bilevel optimization to the investigation of different approaches that use deep neural networks for solving inverse imaging problems. Alongside all approaches that are being discussed, their numerical solution and available solution guarantees will be stated.

Carola-Bibiane Schönlieb
DAMTP, University of Cambridge
cbs31@cam.ac.uk

IP3

Reduced Order Methods for PDEs: State of the Art and Perspectives with Applications in Industry, Medicine and Environmental Sciences

We provide the state of the art of Reduced Order Methods (ROM) for parametric Partial Differential Equations (PDEs), and we focus on some perspectives in their current trends and developments, with a special interest in parametric problems arising in Computational Fluid Dynamics (CFD). Efficient parametrisations (random inputs, geometry, physics) are very important to be able to properly address an offline-online decoupling of the computational procedures and to allow competitive performances. Current ROM developments in CFD include: a better use of stable high fidelity methods, considering also spectral element method and finite volume discretisations; efficient sampling techniques to reduce the number of the basis functions, retained as snapshots, as well as the dimension of online systems; the improvements of the certification of accuracy based on residual based error bounds and of the stability factors, as well as the the guarantee of the stability of the approximation. For nonlinear systems, also the investigation on bifurcations of parametric solutions are crucial and they may be obtained by a reduced eigenvalue analysis of the linearised operator. All the previous aspects are crucial to analyse in real time complex parametric industrial, environmental and biomedical flow problems. We will focus on few benchmarks, including fluid-structure interaction problems, flow control and shape optimization.

Gianluigi Rozza
SISSA University, Italy
Trieste, Italy
grozza@sissa.it

IP4

Fluid-Structure Interaction in Medicine and Biology: Methods, Models, and Applications

Fluid-structure interaction (FSI) is ubiquitous in nature and occurs at molecular to environmental scales, from the writhing of DNA in nucleoplasm, to the beating of cilia and flagella and the projection of lamellipodia and bleb-like protrusions by motile cells, to the flow of blood in the heart, to swimming fish and flying birds and insects, to the dispersal of seeds and pollen in the wind. This talk will describe numerical methods and computational infrastructure for FSI, focusing on extensions of the immersed boundary (IB) method for fluid-structure interaction and applications of these methods to various models in medicine and biology. Different approaches are needed for FSI involving rigid and elastic structures, but both can be addressed within the framework of the IB method. I will discuss IB methods for FSI with prescribed structural kinematics and methods for FSI involving flexible bodies that use nonlinear structural dynamics formulations. I also will describe new extensions of these methods that aim to achieve higher-order accuracy for applications involving realistic biological and physiological geometries. I will survey applications of these IB methods in biology and medicine, including flagellar mechanics, aquatic locomotion and neuro-mechanical feedback, and esophageal transport. I will also detail ongoing work to develop IB models of the heart and its valves and applications to cardiovascular medical devices.

Boyce E. Griffith
Departments of Mathematics and Biomedical Engineering
University of North Carolina at Chapel Hill
boyceg@email.unc.edu

IP5

Communication Avoiding: The Past Decade and the New Challenges

In this talk I will review 10 years of research in designing robust algorithms for linear algebra that reduce or even minimize in some cases data movement between processors, hence enhancing scalability while preserving accuracy. I will focus on two major operations in linear algebra. The first one, ubiquitous in scientific computing, involves solving large sparse linear systems of equations arising from

the discretization of PDEs with strongly heterogeneous coefficients. In particular, I will discuss here enlarged Krylov subspace solvers and algebraic preconditioners that are robust, in the sense that they allow to bound the condition number of the preconditioned matrix. The second operation, ubiquitous in the data analysis applications of scientific computing, involves the compression of large volumes of data while preserving information. I will present a unified view of deterministic and randomized algorithms for computing the low rank approximation of a matrix that allow to attain a given precision with low communication cost. I will discuss the efficiency of the proposed methods in the context of several different real life applications. I will conclude the talk with tensors in high dimensions presenting some recent results while emphasizing many remaining open questions.

Laura Grigori
INRIA
France
Laura.Grigori@inria.fr

IP6

Stochastic Gradient Descent, in Theory and Practice

Stochastic Gradient Descent (SGD) is a first-order stochastic optimization algorithm which is the algorithm of choice behind powerful deep learning architectures which are becoming increasingly omnipresent in society. Still, there remains is a wide gap between the setting where SGD theoretical guarantees and the setting where SGD is most effective and useful in practice. We discuss recent theoretical results which make progress towards closing this gap. First, we present the first theoretical guarantees for "adaptive learning rate" SGD algorithms such as AdaGrad which are used widely in practice, as they make SGD less sensitive to choice of hyperparameters such as the step-size, but which are challenging to analyze theoretically as they are nonlinear stochastic dynamical systems. Second, we provide new guarantees for the convergence of SGD to global minimizers of certain non-convex optimization problems via a novel analysis for products of independent random matrices. We conclude by discussing several open problems.

Rachel Ward
Department of Mathematics
University of Texas at Austin
rward@math.utexas.edu

IP7

Physical, Numerical, and Computational Challenges in Modelling Oceans for Climate

The worlds oceans play a critical role in determining the evolution of Earths climate. While changes in surface climate and weather extremes are usually presented in the context of the atmosphere and weather, the long time scales in the climate systems belong to the ice caps, land, and oceans. In the ocean, short spatial- and fast timescale processes, such as wave breaking and turbulence, ultimately control the planetary scale circulation which evolves and equilibrates over millennia. These same processes make the oceans an important part of the climate system at long time-scales, but computationally challenging to simulate. Until recently most ocean climate models were non-eddy resolving; the dominant modes of variability, mesoscale eddies - analogous to weather systems, were not resolved but parameterized with lower order diffusive closures. Models from this prior era used numerical methods and algorithms that emphasized efficiency over accuracy. Contemporary computing resources allow the use of eddy-permitting ocean models for climate, along with short-duration eddy-resolving calculations. A 2-D turbulent eddy cascade introduces new constraints on ocean models. Modern computing architectures present new challenges for simulating the long time-scales and near adiabatic nature of the ocean interior. We review innovations in numerical and computational methods used in ocean modelling, and examine some of the unique challenges posed by the physics of the ocean.

Alistair Adcroft
Princeton University
aadcroft@princeton.edu

IP8

Role of Tensors in Machine Learning

Tensors are higher order extensions of matrices that can incorporate multiple modalities and encode higher-order relationships in data. After an introduction to tensor methods, I will present ways in which tensor methods can be used in deep learning, as well as in probabilistic modeling. I will show that tensor contractions, which are extensions of matrix products, provide high rates of compression in a variety of neural network models. I will also demonstrate the use of tensors for document categorization at scale through probabilistic topic models. These are available in a python library called Tensorly that provides a high-level API for tensor methods and deep tensorized architectures.

Anima Anandkumar
Engineering and Applied Sciences
CalTech
anima@caltech.edu

SP1

SIAG CSE Best Paper Prize Lecture: Propagating Uncertainty from Data to Prediction with a Model of the Antarctic Ice Sheet

Given a (large-scale) model containing uncertain parameters, (possibly) noisy observational data, and a prediction quantity of interest, how do we construct efficient and scalable algorithms to (1) infer the model parameters from the data (the deterministic inverse problem), (2) quantify the uncertainty in the inferred parameters (the Bayesian inference problem), and (3) propagate the resulting uncertain parameters through the model to issue predictions with quantified uncertainties (the forward uncertainty propagation problem)? We present efficient and scalable algorithms for this end-to-end, data-to-prediction process under the Gaussian approximation and in the context of modeling the flow of the Antarctic ice sheet and its effect on loss of grounded ice to the ocean. We show that the work required for executing this data-to-prediction processmeasured in number of forward (and adjoint) ice sheet model solvesis independent of the state dimension, parameter dimension, data dimension, and the number of processor cores. Despite their large size, the observational data typically provide only sparse information on model parameters. This property can be exploited to construct a low rank approximation of the linearized parameter-to-observable map via randomized SVD methods and adjoint-based actions of Hessians of the data misfit functional.

Tobin Isaac

Georgia Tech
tisaac@cc.gatech.edu

Noemi Petra
University of California, Merced
npetra@ucmerced.edu

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

Omar Ghattas
The University of Texas at Austin
omar@ices.utexas.edu

SP2

SIAG/CSE Early Career Prize: Data-Driven Discovery and Control of Complex Systems: Uncovering Interpretable and Generalizable Nonlinear Models

Accurate and efficient reduced-order models are essential to understand, predict, estimate, and control complex, multiscale, and nonlinear dynamical systems. These models should ideally be generalizable, interpretable, and based on limited training data. This work develops a general framework to discover the governing equations underlying a dynamical system simply from data measurements, leveraging advances in sparsity-promoting techniques and machine learning. The resulting models are parsimonious, balancing model complexity with descriptive ability while avoiding overfitting. This perspective, combining dynamical systems with machine learning and sparse sensing, is explored with the overarching goal of real-time closed-loop feedback control of unsteady fluid systems. We will discuss how to enforce known constraints, such as energy conserving quadratic nonlinearities in incompressible fluids, to bake in known physics. Next, we will demonstrate that higher-order nonlinearities can approximate the effect of truncated modes, resulting in more accurate models of lower order than Galerkin projection. Finally, we will discuss the use of intrinsic measurement coordinates to build nonlinear models, circumventing the well-known issue of continuous mode deformation associated with methods based on the proper orthogonal decomposition. This approach is demonstrated on several relevant systems with low-dimensional dynamics.

Steven Brunton
University of Washington
sbrunton@uw.edu

SP3

SIAM/ACM Prize in Computational Science and Engineering: The Singular Value Decomposition: Anatomy of an Algorithm, Optimizing for Performance

The computation of the singular value decomposition, or SVD, has a long history with many improvements over the years, both in its implementations and algorithmically. Here, we look at the evolution of SVD algorithms for dense matrices, discussing the motivation and performance impacts of changes. We will look at the changes over time by testing various historical and current implementations on a common, modern multicore machine and a distributed computing platform. We show that algorithmic and implementation improvements have increased the speed of the SVD by several orders of magnitude, while using up to 40 times less energy.

Jack J. Dongarra
University of Tennessee, Oak Ridge National Laboratory, USA
dongarra@icl.utk.edu

SP4

James H. Wilkinson Prize for Numerical Software Prize Lecture: Solving the Two Language Problem in Scientific Computing and Machine Learning with Julia

Bridging cultures that have often been distant, Julia combines expertise from the diverse fields of computer science and computational science to create a new approach to numerical computing. Julia is designed to be easy and fast and questions notions generally held to be laws of nature by practitioners of numerical computing: 1. High-level dynamic programs have to be slow. 2. One must prototype in one language and then rewrite in another language for speed or deployment. 3. There are parts of a system appropriate for the programmer, and other parts that are best left untouched as they have been built by the experts. We introduce the Julia programming language and its designa dance between specialization and abstraction. Specialization allows for custom treatment. Multiple dispatch, a technique from computer science, picks the right algorithm for the right circumstance. Abstraction, which is what good computation is really about, recognizes what remains the same after differences are stripped away. Abstractions in mathematics are captured as code through another technique from computer science, generic programming. Julia has been downloaded over 3,000,000 times. The Julia community has contributed almost 2,500 Julia packages, while also building connectivity to C, C++, Fortran, MPI, Python, R, Java and others. Julia shows that one can achieve machine performance without sacrificing human convenience.

Jeffrey Bezanson, Stefan Karpinski
Julia Computing, Inc.
jeff@juliacomputing.com, stefan@juliacomputing.com

Viral Shah
Julia Computing, Inc.
Julia Computing, Inc.
viral@juliacomputing.com

CP1

Parameter Identification for a Viscoplastic Model with Damage and Effect of Conditions on Results using Bayesian Approaches

The state of materials and accordingly the properties of structures are changing over the period of use, which may influence the reliability and quality of the structure during its life-time. Therefore identification of the model parameters and states of the system is a topic which has attracted attention in the content of structural health monitoring.
In this work the focus is on identification of material parameters and states of a viscoplastic damaging material. It is proposed to use Bayesian inverse methods for this. To do so, two steps are considered, solving the forward and inverse problem. Therefore, first the propagation of the a priori parametric uncertainty through the model including damage describing the behaviour of a steel struc-

ture is studied. A non-intrusive Stochastic Finite Element Method (SFEM) based on polynomial chaos is applied.
From the forward model, material parameters and interval unobservable state variables can be identified using measurement data such as displacement via Bayesian approaches. In this study, two methods are applied. The first is a Polynomial chaos based update method which is a modification of Kalman filter and the second one is a transitional Markov chain Monte Carlo method. At the end, the results of both methods are compared.
We study the effect of load conditions and sensor placements, which provide us the observation, on the identification procedure and how much information they provide to the identification process.

Ehsan Adeli
Institute for Scientific Computing
Brunswick University of Technology
e.adeli@tu-braunschweig.de

CP1

Surrogate Model Design for Non-smooth and Computationally Expensive Radiation Detection Models

Surrogate models are increasingly required for applications in which first-principles simulation models are prohibitively expensive to employ for design, control, or uncertainty analysis. They can also be used to approximate models whose discontinuous derivatives preclude the use of gradient-based optimization or data assimilation algorithms. We consider two problems; the first is to infer the two-dimensional location and intensity of an urban radiation source using a ray tracing model based on Boltzmann transport theory. The resulting likelihood exhibits discontinuous derivatives due to the presence of buildings. Whereas this code is relatively efficient, the second problem considered is an extension to three-dimensional MCNP simulations. These less efficient simulations preclude Bayesian inference, which typically requires thousands to millions of simulations. To address these issues, we discuss the construction of surrogate models for optimization, Bayesian inference, and uncertainty propagation. Specifically, we consider surrogates based on Legendre polynomials, multivariate adaptive regression splines (MARS), radial basis functions, Gaussian processes, and neural networks. We detail strategies for computing training and test points and discuss the merits and deficits of each method.

Jared A. Cook
North Carolina State University
jacook8@ncsu.edu

CP1

Removal of Multiplicative Poisson Noise using Variance Based Joint Sparsity Recovery

There has been much effort put into the use of multiple measurement vectors (MMV) and ℓ_2 or ℓ_1 regularization to recover signals and images. Specifically, recent research has suggested that the effects of bad data on a signal or image may be lessened by using the variance based joint sparsity recovery method (VBJS), which weights the original jointly sparse (JS) signal recovery method, making signal recovery more accurate and cost efficient. This method, however, only accounts for additive noise supplement to the deblurring model for reconstruction of the true signal. Realistically, there is automatically inconsistency in the deblurring model, as well as other factors, that will also create multiplicative noise in the model for reconstruction. In this talk, this multiplicative noise in the deblurring model will be addressed, specifically multiplicative Poisson error distribution noise. Understanding the effects of multiplicative noise on the deblurring model as well as how to reduce its effects on the reconstruction of the true signal are essential for using the VBJS recovery method in data fusion problems, where measurements may be obtained through many different techniques and therefore must be combined for an accurate reconstruction. This proposal is considered in 1-dimension, but through future work may also be applied to 2- and 3-dimensional scenes.

Carley R. Walker
The University of Southern Mississippi
College of Arts and Sciences
Carley.Walker@usm.edu

CP1

Sampling Input Parameters of Stochastic Models using Consistent Pull-back Measures

Using a computational model to make accurate predictions often requires generating samples of uncertain input parameters. Prior knowledge, expert opinions, and other sources of information can often be used to describe a plausible initial distribution on these input parameters. Using the model to define the push-forward of this initial distribution provides a prediction distribution on model outputs. When data are available on some of the model outputs, we seek to use discrepancies in the predicted to observed distributions in outputs to update the initial distribution. We focus on a recently developed method that uses the ratio of observed versus predicted data distributions on model quantities to generate an updated distribution that is consistent with the observed data, I.e., we generate a pull-back measure that is regularized by the initial distribution in directions not informed by the data. Whereas previous research on this method has focused on applications where the parameter-to-outputs map is deterministic, we present a framework for applying this method to stochastic maps. Numerical results with applications to model inadequacy are reported and implications discussed.

Tian Yu Yen
University of Colorado Denver
tian.yen@ucdenver.edu

CP1

Parameter Inference and Uncertainty Quantification in Simulating Blood Flow in Coronary Arteries

Non-invasive measurement of fractional flow reserve (FFR), the key index in the diagnosis of coronary artery disease (CAD), is a promising alternative to traditional way of performing invasive coronary angiography (ICA). Computed FFR is one of the most feasible ways: by discretizing reconstructed coronary geometry from computed tomography angiography (CTA) and imposing proper boundary conditions, a full-order computational fluid dynamics solver could numerically reveal the pressure drop over the lesion non-invasively. Nonetheless, the computational cost to perform such simulations for a three-dimensional (3D) full-order geometry is unavailable in most clinical settings, greatly limiting the potential beneficiaries. One promising alternative is to perform one-dimensional (1D) reduced-order model based on the centerline extracted from the 3D

geometry, where 1D simulation could be performed on a personal computer within a couple of minutes. The motivation of this work is to (1)validate the accuracy of the 1D model versus 3D model, (2)study the sensitivity of physiological parameters to computed FFR and (3)calibrate the boundary condition by a probabilistic-based multi-fidelity framework.

Minglang Yin, Alireza Yazdani, George Karniadakis
Brown University
minglang_yin@brown.edu, alireza_yazdani@brown.edu, george_karniadakis@rown.edu

CP2

Equilibrium Statistics and Entropy Computation for Vortex Filaments on a Cubic Lattice: A Modified Pivot Algorithm

We present an extension of the results obtained by Chorin and others in the early 90s on the equilibrium statistics of the vortex filaments constrained to the cubic lattice. We present the pivot algorithm for generation of self-avoiding walks and its modification that allows to extend the computational results to a much wider range of temperatures, both positive and negative. We also discuss a way to reliably estimate the entropy of such filaments using the hypothetical scanning method of Meirovitch and another, computationally efficient algorithm if average energies are known.

Pavel Belik
Mathematics Department
Augsburg College
belik@augsburg.edu

Aleksandr Lukanen, Robert Laskowski
Augsburg University
lukanena@augsburg.edu, laskowsr@augsburg.edu

Douglas Dokken, Kurt Scholz, Mikhail M. Shvartsman
University of St. Thomas
dpdokken@stthomas.edu, k9scholz@stthomas.edu, mmshvartsman@stthomas.edu

CP2

Future Proofing a Massive Parallel Framework for Real-time Flood Prediction

Due to the rising numbers of natural disasters such as floods, numerical simulation to accurately predict their impact on human infrastructure, ensuring the safety of large populations becomes increasingly important. Having the simulation results available in time to be able to undertake viable countermeasures poses a difficult challenge and requires the efficient usage of todays top tier parallel hardware. To this end, a CFD Framework has been developed at our chair which has been successfully deployed on up to 140,000 cores on two of Germanys top-tier supercomputing systems. While the core concepts of the framework, a dedicated distributed data structure based on hierarchically refined non-overlapping block-structured Cartesian grids and a custom-tailored parallel multigrid-like solver have been proven, a complete overhaul has become necessary to address its shortcomings. In this presentation, we address lessons learned from this continued process with emphasis on the core structure of the framework for easy extension of new functionality. Furthermore, the possibility to calculate different physics on grids in the domain or use varying discretisations, allowing for free-surface flows, fluid-structure interaction and techniques like 2D-3D coupling is presented. Finally, a decentral domain management using individual sub-domains without global knowledge is shown, ensuring the scalability of the framework even towards exascale.

Christoph M. Ertl, Bobby Minola Ginting
Technical University of Munich
Chair for Computation in Engineering
christoph.ertl@tum.de, bobbyminola.ginting@tum.de

Ralf-Peter Mundani
TUM, Faculty of Civil, Geo and Environmental Engineering
Chair for Computation in Engineering
mundani@tum.de

CP2

Visual Exploration of Flood Simulations in Cave-based Environments

According to the US President's Council of Advisors on Science and Technology "high-performance computing must now assume a broader meaning, encompassing not only flops, but also the ability, for example, to efficiently manipulate vast and rapidly increasing quantities of both numerical and non-numerical data" [T. Kalil, 2015]. Real-time or in situ approaches provide users the possibility to interact with the application during runtime, nevertheless they often suffer from hardware and/or algorithmic limitations hindering fast and efficient visual exploration in order to experience phenomena that would not be possible or accessible in reality. Within our research, floods and impact of floods on infrastructure should be predicted, reaching from the river down to the scale of built infrastructure such as railway, subway, tunnels, waste water channels, buildings, and building infrastructure. For better visual comprehension of the computed results an immersive virtual reality facility such as a Cave Automatic Virtual Environment (CAVE) should be coupled to the running application. While standing in the CAVE, users have the possibility moving around and zooming into the data, moving forward/backward in time, and manipulating simulation parameters to evaluate various scenarios. Key feature of this approach is a data exploration technique that allows for online/offline visualisation of huge data sets (up to hundreds of billions of unknowns), here extended to a CAVE environment.

Ralf-Peter Mundani
TUM, Faculty of Civil, Geo and Environmental Engineering
Chair for Computation in Engineering
mundani@tum.de

Mario Silaci
Technical University of Munich
mario.silaci@tum.de

Christoph M. Ertl
Technical University of Munich
Chair for Computation in Engineering
christoph.ertl@tum.de

CP2

Exploring the Predictibility of Random Forests for the Sedov-Von Neumann-Taylor Blast Wave Solu-

tion

Blast waves correspond to the gas dynamics motion due to the instantaneous deposition of energy in an infinitesimal volume. They occur in a variety of natural and manmade applications such as supernovae, volcanic eruptions, asteroid impacts, and explosions. Semi-analytical solutions for the propagation of a blast wave, were derived independently by L.I. Sedov, J. von-Neumann, and Sir G.I. Taylor; hence termed the SNT solution. For a set of initial energy and gas adiabatic index, the SNT solution computes the density, pressure, and velocity as a function of the distance from the origin at a given time instant. Our research goals are to evaluate several Machine Learning (ML) models in order to predict the values of the initial energy and the time of the energy release for spherical configurations. Hence, we investigated the predictive capabilities of several ML models applied to the SNT problem, including Polynomial Regressors, and Random Forest with Decision Trees estimators. We found that the Polynomial Regressors produced a high bias error indicating poor predicitbility and the Random Forest was more successful in capturing the nonlinearity in the data indicating better predictive capabilities. Further, if the number of training samples is increased in the region where the sensitivity to the adiabatic index is high, the accuracy of the model will improve significantly. These results will be discussed in the presentation along with more recent analysis.

Mokbel Karam
PhD student, Chemical Engineering Department, University of Utah
mokbel.karam@chemeng.utah.edu

Fady M. Najjar
Lawrence Livermore National Laboratory
najjar2@llnl.gov

Ming Jiang
Ph.D., Computer Scientist, Lawrence Livermore National Labor
jiang4@llnl.gov

James C. Sutherland
Department of Chemical Engineering
The University of Utah
james.sutherland@chemeng.utah.edu

Tony Saad
Department of Chemical Engineering
University of Utah
tony.saad@chemeng.utah.edu

CP2

Reference Mapping: Application in Eulerian Hydrocodes

In the simulation of coupled high-rate solid and fluid mechanical processes, the Eulerian frame for hydrodynamics has an advantage over the Lagrangian in that large deformations, fracture, fragmentation and collision can be modeled without the disastrous degradations caused by mesh tangling. A major disadvantage is the diffusion of material properties as material flows through the mesh. Related are the ambiguities in material history caused by mixing materials with different properties or states, such as yield strength, damage characterization, etc. in a computational cell. Finally, there is the disrepect of advection schemes for material subtleties such as maintaining deviators on a yield surface. The reference map technique promises to remedy all of these deficiencies. It was first described in 2009 by Kamrin and Nave [Ken Kamrin and Jean-Christophe Nave. An Eulerian approach to the simulation of deformable solids: Application to finite-strain elasticity. http://arxiv.org/abs/0901.3799v2, 2009.]. Vitali and Benson [Efrem Vitali and David J. Benson. Modeling localized failure with arbitrary-lagrangian-eulerian methods. Computational Mechanics, 49:197-212, 2012.] independently described it in 2012 to treat the diffusion of rapidly varying material properties. We will discuss the details of one of the first implementation of this technique in a typical Eulerian code. The effectiveness of the algorithm will be tested on engineering examples.

Serge B. Ndanou
LOS ALAMOS NATIONAL LABORATORY
sndanou@lanl.gov

Gary Dilts
LANL
gad@lanl.gov

Thomas Masser
Continuum Dynamics Group, CCS-2
CCS Division, Los Alamos National Laboratory
tmasser@lanl.gov

CP3

Characterising Failure in Hypersonic Vehicles using Unified Bayesian Networks

We present a prototype system for probabilistic assessment of thermo-structural failure in hypersonic vehicles, based on a unified Bayesian network. A probabilistic characterization of failure in such systems presents significant challenges due to the complex dependence among a large number of uncertain parameters, the availability of an ensemble of models each describing a particular aspect of the underlying physics, and the presence of multiple mission critical components like thermal protection system (TPS) and control surfaces. We start by considering an idealized model of the Space Shuttle orbiter, where parameters are assigned to the vehicle geometry, the flight trajectory and TPS material properties. We use a simplified method to compute the heat flux and pressure load histories on the TPS surface and a two-dimensional plane strain characterization of its structural response. Additionally, failure is assessed based on thermal stress and maximum temperature in operation. Using a Bayesian network with discrete random variables, we perform inference using brute-force approaches and message passing. We also investigate how the results are affected by observations from pressure and temperature sensors.

Sean R. Breckling
Air Force Institute of Technology (AFIT)
sbreckli@nd.edu

Daniele E. Schiavazzi
University of Notre Dame
dschiavazzi@nd.edu

CP3

Combining Physical Models and Gaussian Processes to Control Risk-sensitive Systems with Lim-

ited Data

In the context of Real Time Optimization, there are problems which suffer from two conditions: 1) there is no precise model known for the system under consideration (model-system mismatch), leading to inaccurate predictions and convergence to suboptimal solutions and 2) the system is risk-sensitive hence exploration is inconvenient or dangerous. Simple examples of these problems are chemical reactors and aircraft controllers. In these cases, the implementation of data-driven optimization approaches is appealing to handle model-system mismatches but may represent an additional risk due to their exploratory nature. To address this challenge, we propose to combine the (approximate) physical known model with a data-driven model which we have chosen to be a Gaussian Process. The physical model allows us to make decisions minimizing risks, while the data-driven model refines the model-system mismatch as more data becomes available. The choice of a Gaussian Process allows us to hedge the risk given our ignorance in the model-system mismatch versus the gains in a performance index. Based on this philosophy we propose two algorithms where risk aversion and risk seeking can be predefined depending on the application.

Ehecatl A. Del Rio Chanona
Imperial College London
a.del-rio-chanona@imperial.ac.uk

CP3

An Adaptive Sampling Approach for Surrogate Modeling of Expensive Computer Experiments

In the design of complex engineering systems like aircraft/rotorcraft/spacecraft, computer experiments offer a cheaper alternative to physical experiments due to high-fidelity(HF) models. However, such models are still not cheap enough for application to *Global Optimization*(GO) and *Uncertainty Quantification*(UQ) to find the best possible design alternative. In such cases, *surrogate models* of HF models become necessary. The construction of surrogate models requires an offline database of the system response generated by running the expensive model several times. In general, the training sample size and distribution for a given problem is unknown apriori and can be over/under predicted, which leads to wastage of resources and poor decision-making. An adaptive model building approach eliminates this problem by sequentially sampling points based on information gained in the previous step. However, an approach that works for highly non-stationary response is still lacking in the literature. Here, we use Gaussian Process(GP) models as surrogate model. We employ a novel *process-convolution* approach to generate parameterized non-stationary GPs that offer control on the process smoothness. We show that our approach outperforms existing methods, particularly for responses that have localized non-smoothness. This leads to better performance in terms of GO, UQ and mean-squared-prediction-errors for a given budget of HF function calls.

Ashwin Renganathan
Georgia Institute of Technology
ashwinsr@gatech.edu

CP3

Long-term Risk Analysis of a Geologic CO2 Storage Project During the Post-injection Period

Geologic storage (GS) of carbon dioxide is an important technology for reducing greenhouse gas emissions. Geologic modeling and risk assessment are two integral components that run in tandem over the entire course of any GS project. The scope of the components may change as a project progresses from the pre-injection phase, through the injection period, to the post-injection phase. To date, there have been few industrial-scale GS projects that provide guidance for site operators on how to efficiently manage a storage complex during the post-injection period. A post-injection monitoring plan is required to produce metrics which meet the regulatory requirements to demonstrate that the storage complex (which might have legacy wells with integrity problems) has evolved to a state with an acceptably low level of risk to humans, resources, and the environment and can be closed. We propose a systems-level framework that can be used for the development of the post-injection monitoring plan and can be tuned for a specific GS project. We utilize detailed physics-based simulations to characterize the fluid and pressure behavior within storage reservoirs. We focus on the risk analysis of incidents that may occur at a GS site and simulations of conditions under which an operator can demonstrate that the risk of these events is sufficiently low to allow for site closure. This framework adapts to site-specific conditions, advances in risk assessment, and new regulatory requirements.

Veronika S. Vasylkivska
Department of Mathematics
Oregon State University
Veronika.Vasylkivska@netl.doe.gov

Gregory Lackey
Oak Ridge Institute for Science and Education
National Energy Technology Laboratory
gregory.lackey@netl.doe.gov

Seth King
AECOM Corporation
National Energy Technology Laboratory
seth.king@netl.doe.gov

Andrew Wentworth
Oak Ridge Institute for Science and Education
National Energy Technology Laboratory
andrew.wentworth@netl.doe.gov

Nicolas Huerta
National Energy Technology Laboratory
nicolas.huerta@netl.doe.gov

Christopher Creason, Jennifer DiGiulio
AECOM Corporation
National Energy Technology Laboratory
christopher.creason@netl.doe.gov,
jennifer.digiulio@netl.doe.gov

Ya-Mei Yang
Oak Ridge Institute for Science and Education
National Energy Technology Laboratory
ya-mei.yang@netl.doe.gov

Robert Dilmore
National Energy Technology Laboratory

robert.dilmore@netl.doe.gov

CP4

Immersed Interface Method for Modeling Fluid-structure Interaction of Flexible Faceted Surfaces

Fluid-structure interaction phenomena are abundantly found in biology and medicine wherein flexible bodies with complex geometries are immersed in a fluid domain. An accurate general-purpose immersed interface method is presented for arbitrary interfacial geometries described by general faceted surfaces that are amenable to nodal C-0 finite element discretization. The method that sharply resolves stress discontinuities at fluid-solid interfaces despite low order representation of the geometry. Smooth and accurate fluid traction forces are calculated and used in a strongly coupled FSI approach where the solid part is modeled by a volumetric mesh with an additional set of surface degrees of freedom to represent the position of the fluid-solid interface. The surface discretization is mechanically coupled to the volumetric solid mesh via stiff spring penalty forces. This fluid-structure interaction model is extensively tested against different numerical and experimental benchmark problems in the literature. The approach is then used to accurately simulate the FSI of deformable blood clots in the cardiovascular system.

Ebrahim (amin) M. Kolahdouz
University of North Carolina, Chapel Hill
ebrahimk@email.unc.edu

Amneet P.S. Bhalla
San Diego State University
asbhalla@sdsu.edu

Kenneth I. Aycock, Brent A. Craven
United States Food and Drug Administration
kenneth.aycock@fda.hhs.gov, brent.craven@fda.hhs.gov

Boyce E. Griffith
University of North Carolina, Chapel Hill
boyceg@email.unc.edu

CP4

Reference Map Technique: A Fully Eulerian Method for Fluid-structure Interactions

Conventional computational methods often create a dilemma for fluid-structure interaction problems. Typically, solids are simulated using a Lagrangian approach with a grid that moves with the material, whereas fluids are simulated using an Eulerian approach with a fixed spatial grid, requiring some type of interfacial coupling between the two different perspectives. Here, a fully Eulerian method for simulating structures immersed in a fluid will be presented. By introducing a reference map variable to model finite-deformation constitutive relations in the structures on the same grid as the fluid, the interfacial coupling problem is highly simplified. A key challenge in this technique is to extrapolate field values to new grid points as the solid moves across the fixed grid. We develop and test a least-squares regression-based algorithm that is more robust and suitable for parallel computation than the existing approaches. The reference map technique is particularly well suited for simulating soft, highly-deformable materials and many-body contact problems. We demonstrate the method by 3D simulations using a large-scale, three-dimensional, parallel implementation of it.

Yuexia Lin, Nicholas Derr
Harvard University
y_lin@g.harvard.edu, derr@g.harvard.edu

Ken Kamrin
MIT
kkamrin@mit.edu

Chris H. Rycroft
Harvard University
chr@seas.harvard.edu

CP4

Volumetric Stabilization of the Immersed Boundary Method

The immersed boundary (IB) method is a formulation for fluid-structure interaction between an incompressible Newtonian fluid and an incompressible solid. The hyperelastic extension generalizes this model to include a wide variety of nonlinearly elastic solid materials. When the equations are discretized, incompressibility of the solid is relaxed to near incompressibility, which allows small changes in the structure's volume. The behavior of the solid is encoded in a strain energy functional that describes how it changes in volume and shape. To correct for the changes in structural volume under large deformations, we use a deviatoric Cauchy stress and a pressure stabilization in the solid region. The performance of the proposed stabilization is explored via standard quasi-static structural mechanics benchmarks.

Ben Vadala-Roth
University of North Carolina at Chapel Hill
benlvr@live.unc.edu

Simone Rossi
Department of Mathematics
University of North Carolina at Chapel Hill
simone.rossi@unc.edu

Boyce E. Griffith
Departments of Mathematics and Biomedical Engineering
University of North Carolina at Chapel Hill
boyceg@email.unc.edu

CP4

An Inverse Problem Formulation of the Immersed Boundary Method

We formulate the immersed-boundary method (IBM) as an inverse problem. A control variable is introduced on the boundary of a larger domain that encompasses the target domain. The value of the control variable is determined such that the mismatch in the desired boundary condition is minimized along the immersed target-domain boundary. As with most inverse problems, this inverse IBM formulation is ill-posed. For example, in the case of the Laplace equation, we prove that the solution is unique but it fails to depend continuously on the data. For the linear advection equation, even solution uniqueness fails to hold. The ill-posedness is rectified by two complimentary strategies. The first strategy is to ensure that the enclosing domain tends to true domain as the mesh is refined. The second strategy is to include a novel (parameter-free) regularization that is based on penalizing the difference between

the control and the state on the boundary. The proposed inverse IBM is applied to the diffusion, advection, and advection-diffusion equations using a high-order discontinuous Galerkin discretization. The numerical experiments demonstrate that the regularized scheme achieves optimal rates of convergence and that the reduced Hessian of the optimization problem has a bounded condition number as the mesh is refined.

Jianfeng Yan
Rensselaer Polytechnic Institute
yanj4@rpi.edu

Jason E. Hicken
Rensselaer Polytechnic Institute
Assistant Professor
hickej2@rpi.edu

CP5

Basin-preserving Lossy Compression for Optimization of a Molecular Dynamics System

The growing disparity between speed of computation and I/O in the next generation of supercomputers heightens the need for the use of data compression within their software. A highlighting example, found in an application under development for the Exascale Computing Project (ECP), finds and catalogs the exponentially many optimal states of a molecular dynamics system. To reduce time required for writing to storage, the optimizers must be compressed in a lossy fashion whilst guaranteeing the compressed optimizer remains in the same basin of attraction under an optimization algorithm. We develop a framework to determine the acceptable level of compression for an optimizer by applying variants of the Kantorovich Theorem, using binary digit rounding as a compression technique. Choosing the Lennard-Jones potential function as a model problem, we exhibit a method for determining the local Lipschitz constant of the Hessian with low computational cost thus allowing the use of our technique in real-time computation. We show numerical results validating our scheme's scalability to systems of the size under consideration for ECP.

Michael A. Retzlaff
University of Maryland, Baltimore County
mretzla1@umbc.edu

Zichao (Wendy) Di
Argonne National Laboratory
wendydi@mcs.anl.gov

CP5

Study Type B Aortic Dissection using a Deconvolution-based Nonlinear Filter

Progressive false lumen aneurysmal degeneration in type B aortic dissection (TBAD) is a complex process with a multi-factorial etiology. Patient-specific computational fluid dynamics simulations provide spatial and temporal hemodynamic factors and registration methods quantify the morphological changes of the false lumen (FL). By considering both simultaneously, correlations can be established to potentially help to understand the intertwining between hemodynamics and false lumen progression. A deconvolution-based nonlinear filter is applied with Navier-Stokes equations with a specific calibration method for the boundary conditions. [H. Xu et al., Coupled morphologicalhemodynamic computational analysis of type b aortic dissection: A longitudinal study. Ann Biomed Eng., pages 113, 2018.] Moreover, the backflow stabilization effect of the filter at the Neumann boundary and global sensitivity of the filter parameters using polynomial chaos expansion are studied. The model is able to capture the flow properties accurately with a large reduction in computational time. Correlations between hemodynamics and the deformation of the FL are investigated. The time-averaged wall shear stress is found to correlate positively with FL dilation ($r^2 = 0.44$). A mild negative correlation is found between the oscillatory shear index and deformation ($r^2 = 0.29$). High relative residence time is suspected to be correlated with thrombose absorption.

Huijuan Xu
Georgia Institute of Technology
huijuanxu.xu@gmail.com

Alessandro Veneziani
MathCS, Emory University, Atlanta, GA
ale@mathcs.emory.edu

Davide Baroli
University of Luxembourg Faculty of Science, Technology and Communication
davide.baroli@uni.lu

CP6

A Single Fluid Flow Poroelastic Model for a Fractured Porous Medium with Transverse Isotropy

In this work, a poro-elastic model, with a full coupling of fluid flow to rock mechanics for an isothermal fractured porous medium with transverse isotropy, and saturated by one monocomponent fluid phase (water or oil) is proposed. The mathematical model includes one mass balance equation and one force equilibrium equation, respectively. The fractures are modeled following a discrete fracture approach, considering fractures as a simplified geometric entities with one lower dimension than the domain. The model is numerically resolved by using the finite element method, which is implemented in a PDE solver that is included in a CFD software. The model is applied to simulate a monophasic fluid flow in a homogeneous and fractured deformable porous medium with transverse isotropy. Finally, the model is validated by comparing numerical results, of pore pressure and rock displacements, against theoretical data from homogeneous and fractured systems.

Pedro Aguilar-Gastelum
INSTITUTO MEXICANO DEL PETROLEO
faguilar@imp.mx

Martín Díaz-Viera
Instituto Mexicano del Petróleo
mdiazv@imp.mx

Manuel Coronado
INSTITUTO MEXICANO DEL PETROLEO
mcoronad@imp.mx

Edscott Wilson-García, Mario Vadillo-Sáenz
Instituto Mexicano del Petróleo
edscott@imp.mx, mevadillo@imp.mx

CP6

Efficient Simulations for Contamination of Ground-

water Aquifers under Uncertainties

In many countries, groundwater is the strategic reserve, which is used as drinking water and as an irrigation resource. Therefore, accurate modeling of the pollution of the soil and groundwater aquifer is highly important. As a model, we consider a density-driven groundwater flow problem with uncertain porosity and permeability. This problem may arise in geothermal reservoir simulation, natural saline-disposal basins, modeling of contaminant plumes and subsurface flow. This strongly non-linear problem describes how salt or polluted water streams down building 'fingers'. The solving process requires a very fine unstructured mesh and, therefore, high computational resources. Consequently, we run the parallel multigrid solver UG4 on Shaheen II supercomputer. The parallelization is done in both - the physical space and the stochastic space. The novelty of this work is the estimation of risks that the pollution will achieve a specific critical concentration. We solve Elder's problem in 2D and 3D domains, where unknown porosity and permeability are modeled by random fields. For approximations in the stochastic space, we use the generalized polynomial chaos expansion. We compute different quantities of interest such as the mean, variance and exceedance probabilities of the concentration. As a reference solution, we use the solution, obtained from the quasi Monte Carlo method.

Alexander Litvinenko
Stochastic Numerics and Bayesian Computational groups, KAUST
alitvinen@gmail.com

Dmitry Logashenko
Extreme Computing Research Center
King Abdullah University of Science and Technology
dmitry.logashenko@kaust.edu.sa

David E. Keyes
King Abdullah University of Science and Technology (KAUST)
david.keyes@kaust.edu.sa

Gabriel Wittum
Extreme Computing Research Center
King Abdullah University of Science and Technology
gabriel.wittum@kaust.edu.sa

Raul F. Tempone
Mathematics, Computational Sciences & Engineering
King Abdullah University of Science and Technology
raul.tempone@kaust.edu.sa

CP6

Large-scale Stochastic Particle Systems with Hydrodynamic Interactions

Soft matter and complex fluids are characterized by the presence of dispersed microscopic particles within a viscous fluid media governed by the Stokes equation. The microscopic structure deforms and flows upon an applied force. The objective is to map the transient particles structure to the material properties. For this purpose, the particle trajectories need to be accurately and efficiently computed. This requires statistical mechanics for the particles and continuum mechanics for the media. The particles dynamics is correlated through the fluid which mediates their interaction, so called many body hydrodynamic interaction (MBHI). Accelerated Stokesian Dynamics (ASD) allows to include accurately the media influence and to compute particle trajectories without solving explicitly the Stokes equation. Large-scale particle simulations with HI are required to study systems such as colloidal glasses and gels. For this purpose, we propose scalable and parallel algorithms of ASD in a distributed memory architecture. We present the parallelization of the sparse near-field two-body lubrication interactions, the full many-body dense far-field interactions and the near-field far-field coupling. We treat in details the parallelization of shear flow, correlated Brownian motion and inter-particle force. The parallelization of the different phases of the algorithm are analyzed and we show scaling up to O(10, 000) processors for O(1, 000, 000) particles.

Gaddiel Ouaknin
Stanford University
gaddyouaknin18@gmail.com

Yu Su
Cornell University
ys576@cornell.edu

Roseanna Zia
Stanford University
rzia@stanford.edu

CP6

Regularized Single and Double Layer Integrals in 3D Stokes Flow

We will present a numerical method for computing the single layer (Stokeslet) and double layer (stresslet) integrals in Stokes flow. The method applies to smooth, closed surfaces in three dimensions, and achieves high accuracy both on and near the surface. Evaluating the integrals accurately for points near the boundary, e.g., when two interfaces are close together, is the most difficult case, and few works are available addressing this problem. Our method is based on regularizing the kernels and adding corrections which are derived analytically. For the case of evaluating the integrals on the surface, as needed when solving integral equations, we design high order regularizations for both kernels that do not require corrections. This approach is direct in that it does not require grid refinement or special quadrature near the singularity, and therefore does not increase the computational complexity of the overall algorithm. Our numerical results demonstrate the uniform convergence rates for several surfaces in both the singular and near singular cases, as well as the importance of corrections when two surfaces are close to each other.

Svetlana Tlupova
Farmingdale State College
tlupovs@farmingdale.edu

J. Thomas Beale
Duke University
beale@math.duke.edu

CP6

Chaotic Dynamics and Self-organization in a Depth-averaged Model of Turbidity Currents

We demonstrate a forward stratigraphic model based on depth-averaged governing equations for the flow of submarine turbidity currents over an erodible bed. We obtain a system of nonlinear hyperbolic PDEs, with an additional so-called Exner equation for modeling the flow-

bed sediment exchange and their bedload transport. The Exner equation plays a key role since a (slow time scale) change in the gradient of the bed influences the (fast time scale) momentum of the flow. The transport equations, along with closure models for sediment transport, TKE balance, and water entrainment, are solved using a first-order finite-volume method with a HLLC approximate Riemann solver and integrated using an explicit Euler scheme. We show the emergence of self-organized patterns in the deposits, including the creation of bedforms, channel formation, and avulsions, consistent with observations of modern systems and lab experiments. These occur even with uniform boundary conditions and symmetric initial conditions. The initial disturbances that trigger these mechanisms are ostensibly sourced by floating-point roundoff errors. We analyze an ensemble of simulations with slightly different initial conditions and show statistics on shapes of geomorphic elements and grain size distributions. Our objective is to assess whether and under what conditions such a numerical model can be predictive and quantify the uncertainty in the results arising due to the irreducible chaos in the dynamical system.

Ramanathan Vishnampet
Department of Mechanical Science and Engineering
University of Illinois at Urbana-Champaign
ramanathan.vishnampet@exxonmobil.com

Mrugesh Shringarpure, Huafei Sun, Amit Kumar, David Hoyal
ExxonMobil Upstream Research Company
mrugesh.s.shringarpure@exxonmobil.com,
huafei.sun@exxonmobil.com,
amit.x.kumar@exxonmobil.com,
david.c.hoyal@exxonmobil.com

CP7

aSP-AMG: Adaptive Smoothing and Prolongation Algebraic MultiGrid

The numerical simulation of modern engineering problems can easily incorporate millions or even billions of degrees of freedom. In several applications, these simulations require the solution of sparse linear systems of equations, and Algebraic MultiGrid (AMG) methods are often standard choices as iterative solvers or preconditioners. Despite carrying the name algebraic, most of these methods still rely on additional information other than the global assembled sparse matrix, for instance, the knowledge of the operator near kernel. This fact somewhat limits their applicability as black-box solvers. In this presentation, we introduce a novel adaptive AMG approach featuring the adaptive Factored Sparse Approximate Inverse (aFSAI) method as a flexible smoother, two strategies for uncovering the near-null space of the system matrix and, lastly, a new approach for computing the prolongation operator dynamically. We assess the performance of the proposed AMG through the solution of real-world test cases arising from engineering applications. Comparisons are made with the aFSAI and BoomerAMG preconditioners, showing that our new method proves to be superior to the first method and with performance comparable to the second one while being particularly attractive in the solution of linear elasticity problems.

Victor A. Paludetto Magri
University of Padova
victor.magri@dicea.unipd.it

Andrea Franceschini
Stanford University
af1990@stanford.edu

Carlo Janna
University of Padova
janna@dmsa.unipd.it

CP7

Multigrid Methods for Optimal Control Problems Constrained by Partial Differential Equations with Random Data

In this work we construct multigrid preconditioners to be used in the solution process of pathwise optimal control problems constrained by elliptic partial differential equations with random coefficients. We combine a sparse-grid collocation approach to discretize in the stochastic space with multigrid techniques in the physical space. Numerical results show that the proposed multigrid preconditioners lead to significant computational savings, with the number of preconditioned conjugate gradient iterations decreasing as the resolution increases.

Ana Maria Soane
U.S. Naval Academy
Mathematics Department
soane@usna.edu

CP7

Improved Convergence of Highly Parallel Additive Multigrid with Scaled Correction

Additive multigrid provides for a highly parallel algorithm, but convergence is typically observed to be poor in comparison to multiplicative multigrid. In order to improve convergence, a minimization problem is formulated such that residual difference between one additive and multiplicative multigrid cycles is minimized. Solving this minimization problem provides for an analytical expression for a scalar multiplier to the correction in additive multigrid. Such a scaled correction showed a significant improvement in the convergence of additive multigrid. Results of comparisons of additive multigrid with and without scaled correction, and multiplicative multigrids for the Poisson and advection-diffusion equations on uniform and directionally stretched structured grids are presented. Strong and weak parallel scaling results demonstrate the effectiveness of the approaches on highly parallel frameworks.

Syam Vangara, Siva Nadarajah
McGill University
syam.vangara@mail.mcgill.ca, siva.nadarajah@mcgill.ca

CP8

Effects of Induced Magnetic Field and Slip Velocity on Transient Rotating Magnetohydrodynamic Electroosmotic Micropumps in a Slowly Varying Periodically Micro-channel

In this work, I have studied the effects of magnetic fields on rotating electro-osmotic flow in a non-uniform microchannel. Electro-osmotically driven fluid flow takes place in a micro-channel with flexible walls in a rotating system with induced magnetic fields. I assume that the entire system is rotating about the height of the channel. The non-linear PoissonBoltzmann equation is analytically using the integral transform methods. Analysis of the results was conducted using Mathcad package. The effects of physical

parameters on the flow characteristics was analyzed.

Mohammed Abdulhameed
Federal Polytechnic Bauchi
amohd@fptb.edu.ng

CP8

Fast Algorithms for the Self-assembly of Charged Dielectric Spheres

We introduce fast algorithms for investigating the self-assembly of charged dielectric spheres. Specifically, we present efficient and higher order accurate methods for solving the electrostatic Poisson equation in the presence of multiple dielectric spherical interfaces. Involved techniques including the boundary element method, method of moments, and image method, etc. We further use our methods to explore the role of dielectric effect in the self-assembly of colloidal particles. And we will show that the dielectric effect can play an important role in those systems and lead to interesting phenomena observed in experiments such as like-charge attraction and varies self-assembled structures.

Zecheng Gan
Department of Mathematics, University of Michigan
zecheng@umich.edu

Weihua Geng
Southern Methodist University
wgeng@smu.edu

Robert Krasny
Department of Mathematics
University of Michigan
krasny@umich.edu

CP8

Modeling Diffusion Signal Decay of Magnetic Resonance Images

Magnetic resonance imaging (MRI) is one of the most useful diagnostic imaging technologies for identifying tissue damage with high contrast and resolution. Diffusion-weighted MRI characterizes the movement of water molecules in tissue by displaying the reduction in signal intensity following the application of diffusion-sensitizing gradients. Since the mobility of water is obstructed (hindered and restricted) by tissue components and compartments, the decay of the diffusion signal is characteristic of cellular and sub-cellular structures. Interpretation of this signal requires mathematical models whose parameters reflect key features (e.g., porosity, tortuosity, surface to volume ratio, permeability) of tissue composition. Current practice uses the apparent diffusion coefficient (ADC, mm2/sec) derived from a simple exponential model of hindered diffusion to describe signal decay. New models based on simulations, stochastic processes, tissue structure, or physical and physiological constraints are needed. In this study, we describe a generalization of the diffusion-weighted signal decay that captures tissue complexity in terms decay functions for different tissue features. We explore a subset of these models to illustrate their characteristics both theoretically and experimentally by considering applications in known heterogeneous environments such as Sephadex gels, and in complex biological tissue such as healthy gray and white brain matter, and brain tumors.

Richard Magin
Department of Bioengineering
University of Illinois at Chicago
rmagin@uic.edu

CP8

Simulating Nonlinear Faraday Waves on a Cylinder

In 1831, Faraday observed the formation of standing waves on the surface of a vibrating fluid body. Subsequent experiments have revealed the existence of a rich tapestry of patterned states that can be accessed by varying the frequency and amplitude of the vibration and have spurred vast research in hydrodynamics and pattern formation. Recently, there has been some work towards numerical simulations in various geometries. These methods however possess low orders of accuracy, making them unsuitable for nonlinear regimes. We present a technique for fast and accurate simulations of nonlinear Faraday waves in a cylinder. Beginning from the viscous potential flow model of Dias et al, we generalize the Transformed Field Expansion to this geometry for finding the highly non-local Dirichlet-to-Neumann map for the Laplace equation. A spectral method relying on Zernike polynomials is developed to compute the bulk potential. The free surface evolution equations are solved in time using fourth-order Runge-Kutta and refined to a high order by Picard iterations. The results are in perfect agreement with the instability thresholds and surface patterns predicted for the linearized problem. The nonlinear simulations reproduce several qualitative features observed experimentally. In addition, by enabling one to switch between various nonlinear regimes, the technique allows a precise determination of the mechanisms triggering various experimental observations.

Saad Qadeer
University of North Carolina Chapel Hill
saadq@email.unc.edu

CP9

Multigrid Method for Incompressible Flow Simulation on Arbitrarily Complex Domain using Cartesian Mesh

A novel efficient and reliable determination procedure of signed distance function (SDF) based on adaptive mesh refinement (AMR) strategy is proposed together with a robust way to apply multigrid (MG) on irregular domain. Through the AMR procedure, the number of the SDF computation process can be optimized by effectively detecting the interface cells. This new AMR based algorithm is proved to be an order of magnitude faster than a naive SDF computation on uniform mesh, and the intelligent refinement criterion, namely edge/face-based criterion, ensures that any small and sharp feature of the object can be effectively detected. The validity of computed SDF is demonstrated by volume convergence test. As the second contribution of this study, a novel geometric MG algorithm that can be applied on arbitrary complex domain is introduced. By analyzing the analytically defined irregular domain problem, the optimal performance of MG, i.e. convergence with fixed number of MG cycles regardless of mesh resolution, is clearly observed on irregular domain with our algorithm. In order to simulate flow around a complex shaped body, these two novel algorithms, i.e. AMR based SDF computation and irregular domain MG methods, are combined with a stable solution method for incompressible flow based on semi-Lagrangian method.

Gwangsoo Go, Hyung Taek Ahn
University of Ulsan

cfd.ggs@gmail.com, htahn@ulsan.ac.kr

CP9

Multigrid Preconditioning for Space-time Distributed Optimal Control of Parabolic Equations

In this work we contrast two multigrid preconditioning strategies for space-time distributed control of parabolic equations. The focus is on the reduced problem resulted from eliminating the state and adjoint variables from the KKT system. Using a classical backward Euler in time and continuous piecewise element discretization in space, our first strategy is based on coarsening both in space and time; we show that the resulting two-grid preconditioner is optimal with respect to both space and time discretization parameters. The second strategy uses only space coarsening, and again, it is shown to be of optimal quality with respect to the mesh size. In this work we discuss the comparative efficiency of these two techniques.

Mona Hajghassem, Andrei Draganescu
Department of Mathematics and Statistics
University of Maryland, Baltimore County
mhajghassem@ubalt.edu, draga@umbc.edu

CP9

Effect of Matrix Reordering on the Performance of Asynchronous Preconditioners for Fluid Flow Problems

In order to exploit massively parallel node architectures for the time-implicit solution of problems in computational fluid dynamics, we explore the use of asynchronous preconditioners. Inspired by chaotic relaxation of the 1960s, the idea of asynchronous iterations can be used to develop fine-grain parallel preconditioners for the non-symmetric ill-conditioned matrices arising in time-implicit schemes for compressible flow problems. In prior work, we developed point-block versions of asynchronous preconditioners which were more suitable for systems of PDEs than the existing scalar asynchronous preconditioners. However, while the asynchronous point-block ILU(0) preconditioner works reasonably well for inviscid flows, it was found that the rate of convergence for viscous problems was generally insufficient. It has been found that the ordering of the matrix can play an important role in the convergence of asynchronous iterations. Therefore, we present a study of the effect of different matrix reorderings on the performance of asynchronous ILU preconditioners when applied to compressible flow problems. We carry out parallel scaling tests for these reordered preconditioners and identify ways to improve their convergence and scalability.

Aditya Kashi
Department of Mechanical Engineering, McGill University
aditya.kashi@mail.mcgill.ca

Siva Nadarajah
McGill University
siva.nadarajah@mcgill.ca

CP9

Matrices, Moments, Quadrature and PDEs

Krylov subspace spectral (KSS) methods are high-order accurate, explicit time-stepping methods with stability characteristic of implicit methods. This "best-of-both-worlds' compromise is achieved by computing each Fourier coefficient of the solution using an individualized approximation, based on techniques from "matrices, moments and quadrature' due to Golub and Meurant for computing bilinear forms involving matrix functions. The result is superior scalability to that of other time-stepping approaches, which motivates continued development of KSS methods for high-resolution simulation of physical phenomena. This talk will present an overview of their derivation and essential properties, including new theoretical results, and also highlight ongoing projects aimed at enhancing their performance and applicability.

James V. Lambers
University of Southern Mississippi
Department of Mathematics
James.Lambers@usm.edu

CP9

P-multigrid Block Reduction Preconditioning for a Fully-implicit High-order Reconstructed Discontinuous Galerkin Flow Solver

We introduce a multigrid block reduction preconditioner for solving linear systems arising from a fully-implicit high-order Discontinuous Galerkin discretization of the all-speed Navier-Stokes equations with phase change. The equations are discretized in a conservative form with a reconstructed Discontinuous Galerkin (rDG) method and integrated with fully-implicit time discretization schemes. To robustly converge the numerically stiff systems, we use the Newton-Krylov framework with a primitive-variable formulation (pressure, velocity, and temperature), which is better conditioned than the conservative-variable form at low-Mach numbers. In the limit of large acoustic CFL number, there is a tight coupling between the low and high order degrees of freedom, leading to a 3x3 block system per equation for a 4th -order rDG scheme and a 6x6 block system per equation for a 5th-order rDG scheme. To robustly solve the large ill-conditioned systems, we utilize a nested Schur complement approach to reduce the NxN high-order blocks to N 1x1 blocks. The nested Schur complement preconditioner is a polynomial multigrid (p-multigrid) reduction technique, where the reduced systems can be effectively solved with an algebraic multigrid method. Numerical results for the p-multigrid Schur complement preconditioned Newton-GMRES solver in the limit of large CFL numbers for low-Mach lid-driven cavity flow and a compressible internally heated convection problem.

Brian Weston
University of California, Davis
Lawrence Livermore National Laboratory
weston8@llnl.gov

Robert Nourgaliev
LLNL
nourgaliev1@llnl.gov

CP10

Multigrid Preconditioners for Optimization-based Domain Decomposition of Elliptic Equations

We present a multigrid preconditioning technique in support of the optimization-based non-overlapping domain decomposition method for elliptic partial differential equations developed by Gunzburger, Heinkenschloss, and Lee (2000). The proposed method is rooted in our earlier work on multigrid preconditioners for boundary control of elliptic equations, and results, as expected, in a computational

cost that is decreasing with increasing resolution relative to the cost of solving the elliptic equations on each subdomain. Moreover, we show that the quality of the preconditioner decreases only mildly with increasing number of subdomains.

Andrei Draganescu
Department of Mathematics and Statistics
University of Maryland, Baltimore County
draga@umbc.edu

Michael A. Retzlaff
University of Maryland, Baltimore County
mretzla1@umbc.edu

CP10

Duality and the Level-set Method for Convex Optimization

Level-set methods for convex optimization are predicated on the idea of solving a sequence of subproblems by flipping the role of the objective and constraints. This idea emerges time and again across a range of algorithms for convex problems. We take a geometric viewpoint of this approach and make connections to duality theory, in order to show that strong duality is a necessary condition for the level-set approach to succeed. We discuss examples demonstrating the failure of the level-set method, and give a geometric explanation for this phenomenon.

Ron Estrin
ICME
Stanford University
restrin@stanford.edu

Michael P. Friedlander
Dept of Computer Science
University of British Columbia
mpf@cs.ubc.ca

CP10

Null Space Gradient Flows for Constrained Optimization with Applications to Shape Optimization

A dynamical system is proposed for solving generic optimization problems set on Hilbert spaces with a finite number of differentiable equality and inequality constraints. Our ultimate goal is non parametric shape optimization, for which more classical optimization methods are often difficult to use because of the infinite dimensionality or the need for tuning algorithm parameters. We rely on a variant of classical gradient flows for equality constrained problems: the search direction is a combination of a null space step, decreasing the objective function tangentially to the constraint, and a Gauss-Newton direction decreasing smoothly the violation of the constraints. Inequality constraints are specifically addressed by solving a dual quadratic programming subproblem of size the number of active or violated constraints, which allows to detect the subset of these to which the optimization trajectory should remain tangent. All stationary points of our dynamical system satisfy the Karush Kuhn and Tucker optimality conditions. We then detail the implementation of our method for non parametric optimal shape design based on Hadamard's method of boundary variations. We illustrate our method for the design optimization of a bridge structure subject to nine load constraints within the framework of a level set based mesh evolution method.

Florian Feppon
MIT
florian.feppon@polytechnique.edu

CP10

Optimization of Programmable Elastic Materials to Produce Targeted Deformed States

Programmable materials are those whose geometry and composition are tailored for a specific action or enhanced performance. A subclass of such objects are those whose elastic properties can be chosen to ensure that the deformable object undergoes a desired specific shape change under an input stimulus. As an example, consider a material made up of components of different thermal expansivity. How should the distribution of these components be chosen so that the shape of the object under a given temperature rise matches a pre-selected target shape? Or, similarly, how should the configuration of a multi-ply composite plate be selected so that it buckles into a desired shape on cool-down? In this talk I will present a general framework for solving such deformed-shape optimization problems. This comprises a finite element implementation of a PDE-constrained optimization problem. The canonical setting of linear thermoelasticity will be discussed first before demonstrating the more general applicability of the method in geometrically nonlinear models such as composite plates. Practical computational issues and a number of potential applications will be discussed, including morphing composites for aerospace applications.

Gareth W. Jones
Harvard University
gareth.jones-10@manchester.ac.uk

CP11

Sparse Low-rank Separated Representation Models for Regression

We consider the problem of learning a multivariate function from a set of scattered observations using a Sparse low-rank Separated Representation (SSR) model structure. The SSR model is promising for high-dimensional learning problems; however, existing training algorithms based on alternating least-squares (ALS) are known to have convergence issues, particularly when the rank of the model is greater than one. In the present work, the SSR model structure is supplemented with ℓ_1 sparsity constraints to ensure the well-posedness of the approximation problem. We propose two fast training algorithms to estimate the model parameters: (i) a cyclic coordinate descent algorithm and (ii) a block coordinate descent algorithm. While the first algorithm is not provably convergent due to the nonconvexity of the optimization problem, the block coordinate descent algorithm guarantees convergence to a Nash equilibrium point. The computational cost of proposed algorithms is shown to scale linearly with respect to all of the parameters in contrast to methods based on ALS. Numerical studies indicate that the proposed SSR model structure is very promising for a range of real-world regression datasets.

Christophe Audouze
University of Toronto
Institute for Aerospace Studies
c.audouze@utoronto.ca

Prasanth B. Nair
University of Toronto
pbn@utias.utoronto.ca

CP11

High Performance Qr Factorization of Ill-conditioned Matrices Based on the Cholesky Qr Algorithm

We aim for computing the QR factorization of an ill-conditioned matrix whose condition number is larger than 10^8 (and up to 10^{16}) by the Cholesky QR type algorithm only using double precision arithmetic. Cholesky QR is suitable for high performance computing but has a numerical instability. Recently, it is shown that this instability can be relieved by repeating Cholesky QR twice (CholeskyQR2), however we still have a problem that CholeskyQR2 breaks down for ill-conditioned matrices; the Cholesky factorization of the Gram matrix is numerically failed. To avoid this problem, we proposed an idea of reducing the condition number by right-multiplying an upper triangular matrix, which based on the observation that Cholesky QR belongs to triangular orthogonalization. We obtain such a good matrix by the Cholesky factorization of the Gram matrix with positive diagonal shift. We employ this process as a preconditioning step before CholeskyQR2; we call resulting algorithm shifted CholeskyQR3. In this poster, we first give an overview of shifted CholeskyQR3 and related theoretical results including roundoff error analysis. Then, we show that shifted CholeskyQR3 outperforms other algorithms for computing the QR factorization on recent computer systems. We also demonstrate that the shifted CholeskyQR3 algorithm efficiently computes the QR factorization in the oblique inner product defined by a sparse real symmetric positive definite matrix.

Takeshi Fukaya
Hokkaido University
fukaya@iic.hokudai.ac.jp

Ramaseshan Kannan
Arup, Manchester
ramaseshan.kannan@arup.com

Yuji Nakatsukasa
National Institute of Informatics
nakatsukasa@nii.ac.jp

Yusaku Yamamoto
The University of Electro-Communications, Japan
yusaku.yamamoto@uec.ac.jp

Yuka Yanagisawa
Waseda University
yuuka@aoni.waseda.jp

CP11

Comparing Randomized and Deterministic Approaches for Computing Low-rank Approximations

Computing a low-rank approximation of large, sparse matrices is an important tasks in scientific computing and data science. The goal is to represent a matrix A as $M_1 M_2^T \approx A$, where M_1 and M_2 are low-rank matrices. This representation should capture important characteristics of the original matrix. In this work, we compare two approaches to obtain a low-rank approximation: a probabilistic algorithm, proposed by Halko et al., that samples the original matrix with random test vectors, and a deterministic one, proposed by Grigori et al., that achieves the goal by selecting columns and rows using a tournament-like process on a hierarchical decomposition of the matrix. We efficiently implemented the algorithms and investigated their performance on random test matrices, as well as on some matrices from the MatrixMarket and the SuiteSparse matrix collection. We illustrate how the runtimes of the two different approaches vary with the approximation quality and with the problem size. In particular, we look for a quantitative understanding of the break-even points given a required precision or a predefined approximation rank, and illustrate that the problem type has a strong influence on whether the deterministic or the probabilistic approach is more efficient for the problem at hand.

Carlos Pachajoa, Viktoria Mayer
University of Vienna
carlos.pachajoa@univie.ac.at, viktoria.mayer@univie.ac.at

Wilfried N. Gansterer
Department of Computer Science
University of Vienna
wilfried.gansterer@univie.ac.at

CP11

Communication-optimal Qr Decomposition Using Task-based Parallelism

Current-day computer architectures tend towards heterogeneous combinations of processors with complex memory hierarchies. As the cost of floating point operations exponentially decreases, data movement within these architectures dominates overall compute time. Applications run on these machines are typically programmed using explicitly defined parallelism (MPI, OpenMP, etc.), usually resulting in limited portability. This changing nature of computer architectures requires a shift away from current programming models to one that focuses on the structure of program data if optimal performance is to be achieved. One such programming model is Legion. Legion provides several advantages over traditional practices, including explicit declaration of data properties, partitioning of data, and the assignment of privileges (such as read/write). As a result, parallelism is implicitly extracted, and data movement operations are no longer the responsibility of the programmer. This allows for easy performance tuning and porting to different architectures. The QR decomposition is commonly used in the numerical solution of multi-physics problems, which has motivated the development of communication-optimal algorithms. In this presentation, we describe the process of implementing a communication-optimal tall, skinny QR decomposition using Legion's programming language, Regent, and we present comparisons of performance and scaling results with that of an existing Fortran/MPI application.

Heather Pacella
Stanford University
hpacella@stanford.edu

Gianluca Iaccarino
Stanford University
Mechanical Engineering
jops@stanford.edu

CP11

A Partial Singular Value Decomposition Algorithm

and its Implementation on Distributed-memory Systems.

Current dense singular value (SVD) implementations suffer from the lack of concurrency during the reduction step toward a condensed bidiagonal matrix form and typically achieve a small fraction of theoretical peak performance. Indeed, whether the reduction to the condensed form employs a single or multi-stage approach as in the state-of-the-art numerical libraries, the memory-bound nature of some of the computational kernels is still an inherent problem. Moreover, the standard SVD implementations do not support more economical partial spectrum requirements. We introduce a new partial SVD decomposition algorithm based on a light-weight version of QR-based Diagonally Weighted Halley (QDWH) algorithm. The QDWH-based algorithm exposes more parallelism and runs closer to the theoretical peak performance of the system, thanks to more compute-bound matrix operations. The QDWHpartial-SVD algorithm selectively calculates the most significant singular values and their corresponding singular vectors. We develop a high-performance implementation of the QDWHpartial-SVD on distributed-memory manycore systems and demonstrate its numerical robustness. We perform a benchmarking campaign of QDWHpartial-SVD against ScaLAPACK PDGESVD and KSVD across various matrix sizes using up to 36864 MPI processes. QDWHpartial-SVD achieves gains up to 4X, 2X against PDGESVD and KSVD, respectively.

David E. Keyes
King Abdullah University of Science and Technology (KAUST)
david.keyes@kaust.edu.sa

Hatem Ltaief
King Abdullah University of Science & Technology (KAUST)
hatem.Ltaief@kaust.edu.sa

Yuji Nakatsukasa
University of Oxford
yuji.nakatsukasa@maths.ox.ac.uk

Dalal Sukkari
KAUST
dalal.sukkari@kaust.edu.sa

CP12

Efficient Multilevel Methods for the Optimal Control of Large-scale Stochastic Partial Differential Equations.

We investigate the design and analysis of multilevel preconditioners for optimal control problems constrained by elliptic equations with stochastic coefficients. Assuming a generalized polynomial chaos expansion for the stochastic components, our approach uses a stochastic Galerkin finite element discretization for the PDE, thus leading to a discrete optimization problem. The key aspect is solving the potentially very large linear systems arising when solving the system representing the first-order optimality conditions. We show that the multilevel preconditioning technique from the optimal control of deterministic elliptic PDEs has a natural extension to the stochastic case, and exhibits a similar optimal behavior with respect to the mesh size, namely the quality of the preconditioner increases with decreasing mesh-size at the optimal rate. Moreover, under certain assumptions, we show that the quality is robust also with respect the two additional parameters that influence the dimension of the problem radically: polynomial degree and stochastic dimension.

Sumaya Alzuhairy
University of Maryland Baltimore County
alsum1@umbc.edu

Andrei Draganescu
Department of Mathematics and Statistics
University of Maryland, Baltimore County
draga@umbc.edu

Bedrich Sousedik
University of Maryland Baltimore County
sousedik@umbc.edu

CP12

Advances in Greedy Optimization Algorithm in Locating Mfs Nodes and Collocation Points

The method of fundamental solutions and other meshless numerical techniques involve development of approximation functions that are linear combinations of basis functions defined over the problem domain, Ω, and boundary, Γ. Selected with respect to the governing partial differential equation (PDE), basis functions are being approximated with boundary conditions given on Γ. Fundamental basis functions, typically potential problems that solve the PDE over $\Omega \cup \Gamma$, involve nodal points where singularities are defined as source points. Typical methods uniformly distribute source points on the exterior of the problem boundary when using them in the approximation. This paper examines locations of corresponding nodes and optimization of their placement, reducing both the computational error and number of nodes involved when basis functions have singularities at source points. Benefits include reducing matrix solver requirements in solving dense systems that may tax traditional computation methods and finds better placement of such nodes in the approximation to improve computational accuracy for the boundary value problem. This novel algorithm provides an optimization capability with 3D case studies in steady state heat transport to show that by including the nodal position coordinates as additional variables being optimized, the resulting approximation function improves in computational accuracy with the degrees of freedom to be optimized.

Aidan Doyle, Anton Nelson, Matthew Yuan
United States Military Academy
aidan.doyle@usma.edu, anton.nelson@usma.edu, matthew.yuan@usma.edu

Noah DeMoes, Bryce Wilkins
Massachusetts Institute of Technology
noahdemoes@gmail.com, bdwilkins95@gmail.com

Kameron Grubaugh
Independent Researcher
grubaugh.kameron@aol.com

Theodore Hromadka
United States Military Academy
tedhromadka@yahoo.com

CP12

Model Reduction for Hamilton-Jacobi-Bellman Equations Resulting from Intraday Trading of Elec-

tricity

Due to the growth of renewable energy, the future perspective of energy markets is seen in short-term trading markets. In this talk, we consider the intraday trading of electricity and derive a Hamilton-Jacobi-Bellman (HJB) equation for this setting. As the intraday products are traded hourly, our aim is to find an optimal trading strategy within every hour using the most recent information on the market. To obtain such an optimal trading strategy, we solve the nonlinear HJB equation for multiple parameters. As no closed-form solution exists for this particular problem, we require a fine discretization such that the full complexity of the problem can be represented. This can result in long computation times, which is why we use the reduced basis method (RBM) to derive a reduced model. Due to the nonlinearity of the problem, the RBM cannot be applied directly. Instead, we use the piecewise constant policy approximation [Reisinger,Forsyth; *Piecewise Constant Policy Approximations to Hamilton-Jacobi-Bellman Equations*, 2016], where a compact control set is discretized into J control points. In each time-step, we solve J elliptic equations and perform a maximum operation at the end. We derive error estimates for which we provide numerical investigations regarding the approximation quality, efficiency and reliability. This is joint work with Christoph Reisinger, University of Oxford as well as N. von Luckner and R. Kiesel, University Duisburg-Essen.

Silke Glas
University of Ulm, Germany
silke.glas@uni-ulm.de

Karsten Urban
Institute of Numerical Mathematics, University of Ulm
karsten.urban@uni-ulm.de

CP12

Robust Residual-based and Residual-free Greedy Algorithms for Reduced Basis Methods

The Reduced Basis Method (RBM) is a rigorous model reduction approach for solving parametrized partial differential equations. RBM relies on residual-based error indicators or a posteriori error bounds to guide construction of the reduced solution subspace. Unfortunately, it is well-known that the standard algorithm for residual norm computation suffers from premature stagnation at the level of the square root of machine precision. In this paper, we develop two alternatives to the standard offline phase of reduced basis algorithms. First, we design a robust strategy for computation of residual error indicators that allows RBM algorithms to enrich the solution subspace with accuracy beyond root machine precision. Secondly, we propose a new error indicator based on the Lebesgue function in interpolation theory. This error indicator does not require computation of residual norms, and instead only requires the ability to compute the RBM solution. This residual-free indicator is rigorous in that it bounds the error committed by the RBM approximation, but up to an uncomputable multiplicative constant. Because of this, the residual-free indicator is effective in choosing snapshots during the offline RBM phase, but cannot currently be used to certify error that the approximation commits. However, it circumvents the need for a posteriori analysis of numerical methods, and therefore can be effective on problems where such a rigorous estimate is hard to derive.

Jiahua Jiang
Virginia Tech
jiahua@vt.edu

Yanlai Chen
Department of Mathematics
University of Massachusetts Dartmouth
yanlai.chen@umassd.edu

Akil Narayan
University of Utah
akil@sci.utah.edu

CP12

Scalable Incremental Nonconvex Optimization Approach for Phase Retrieval from Minimal Measurements

We aim to find a solution $x \in R^n/C^n$ to a system of quadratic equations of the form $b_i = |\langle a_1, x\rangle|^2, i = 1, 2, \ldots, m$, e.g., the well-known phase retrieval problem. It has been proved that the number $m = 2n - 1$ of generic random measurement vectors $a_i \in R^n$ is sufficient and necessary for uniquely determining the n-length real vector x up to a global sign. We revert to the convex relaxation semidefinite programming (SDP) approach and propose to indirectly minimize the convex objective by successive and incremental nonconvex optimization, termed as IncrePR, to overcome the excessive computation cost of typical SDP solvers. IncrePR avoids sensitive dependence of initialization of nonconvex approaches and achieves global convergence. For real Gaussian model, IncrePR achieves perfect recovery from $m = 2n - 1$ noiseless measurement and the recovery is stable from noisy measurement. When applying IncrePR for structured (non-Gaussian) measurements, such as transmission matrix and oversampling Fourier measurement, it can also locate a reconstruction close to true reconstruction with few measurements. Extensive numerical tests show that IncrePR outperforms other state-of-the-art methods in the sharpest phase transition of perfect recovery for Gaussian model and the best reconstruction quality for other non-Gaussian models.

Ji Li
Beijing Computational Science Research Center
keelee@csrc.ac.cn

Jian-Feng Cai
Department of Mathematics
Hong Kong University of Science and Technology
jfcai@ust.hk

Hongkai Zhao
University of California, Irvine
Department of Mathematics
zhao@math.uci.edu

CP13

Recovery-based Discontinuous Galerkin Method for the Cahn-Hilliard Equation

Handling high-order derivatives, say, the diffusion terms in the Navier-Stokes equations, with discontinuous Galerkin (DG) method is a nontrivial task, because the numerical solutions are represented by discontinuous piecewise polynomials of degree p. In 2005, a novel recovery-based discontinuous Galerkin method (RDG) was introduced by van Leer and Nomura for diffusion, where a polynomial of degree $2p + 1$ is reconstructed on the two adjacent elements. It achieved a surprising order of accuracy of $3p + 1$ for

odd p and $3p + 2$ for even p in terms of cell-average error for the heat equation. Here we illustrate how to apply the idea of recovery to solving partial differential equations with high-order derivatives. We developed a RDG method for the Cahn-Hilliard equation. To enable analysis of RDG schemes designed for non-linear problems, we suggested a new way to analyze the accuracy of RDG schemes via Taylor expansion. The new form of analysis explains the accuracy of the RDG scheme developed for the Cahn-Hilliard equation in one space dimension. Numerical experiments show that the new RDG scheme has property of superconvergence. Furthermore, it's demonstrated that the new RDG scheme is more accurate than the established local discontinuous Galerkin (LDG) approach.

Fanchen He
University of MIchigan
wzfche@umich.edu

Philip Johnson, Eric Johnsen
University of Michigan
phedjohn@umich.edu, ejohnsen@umich.edu

CP13

Homogenization Wavelet Reconstruction for Partial Equations with Elliptic Terms

In this talk an algorithm for the reconstruction of solutions of partial differential equations with elliptic terms using a wavelet basis conditioned on homogenization of piecewise discontinuous coefficients. A forward transform is developed that produces the correct homogenized value at a coarsest scale with details on all intermediate scales on a nested dyadic mesh. The details from the forward transform can be used to reconstruct the elliptic coefficient at all intermediate scales. This is done using the inverse transform. In addition, the same details along with the perturbation series from the definition of the homogenization process can be used to reconstruct a solution of the partial differential equation at all intermediate scales. Computational examples including a simple porous media problem and the Cahn-Hilliard equation will be presented.

Joseph V. Koebbe
Utah State University
Department of Mathematics
koebbe@usu.edu

Abibat Lasisi
Virginia Military Institute
abibat.lasisi@aggiemail.usu.edu

CP13

The Multiscale Perturbation Method for Elliptic Equations

In the formulation of multiscale methods for second order elliptic equations that are based on domain decomposition procedures, typically the computational domain is decomposed into subdomains, and for each subdomain a set of multiscale basis functions is numerically constructed. Consider the application of such a method to solve a multiphase flow problem where through an operator splitting algorithm, the velocity-pressure and transport equations are solved sequentially. From one time step to the next the multiscale basis functions should be recomputed, because of the coupling of the underlying PDEs. Instead of recomputing all multiscale basis function every time step of a numerical solution, we propose the Multiscale Perturbation Method (MPM). In MPM an approximate solution of velocity and pressure for a new time is obtained by combining regular perturbation theory with multiscale basis functions computed in a earlier time. A novel parallel preconditioner is then used to obtain an improvement on the result obtained from the naive approach. An efficient parallel algorithm is implemented in multi-core machines and the method also fits well in GPU clusters. Numerical experiments, where the perturbation theory results are compared with direct fine grid solutions, are presented and discussed.

Abdallah Alsadig, Het Y. Mankad
The University of Texas at Dallas
alsadig.ali@utdallas.edu, hym130030@utdallas.edu

Felipe Pereira
Department of Mathematical Sciences
University of Texas at Dallas, Dallas, Texas
felipepereira62@gmail.com

Fabricio Sousa
Institute of Mathematics and Computational Sciences
University of Sao Paulo
f.s.sousa@gmail.com

CP13

High-order Stochastic Semi-Lagrangian Method to Solve Transport Equations

A local, explicit high-order accurate semi-Lagrangian spectral element method is developed for the solution of stochastic transport equations which model reactive turbulent flows. The semi-Lagrangian method solves for the multi-variate probability density function. The high-dimensional probability distribution function (PDF) that depends on the number of species is obtained using a Monte-Carlo approach, reducing the computational cost as compared to solving the Fokker-Planck type of equations for the PDF. The method is local, highly parallel and explicit and it is therefore consistent with discontinuous Galerkin Navier-Stokes solvers that provide the velocity field for the transport of species. We achieve this by seeding Lagrangian particles at the Gauss quadrature collocation nodes within each element. The particle is advected by integrating the stochastic differential equation in time. A new interpolant is constructed from the advected nodes with a least squares method constrained by boundary conditions and mass conversation. The interpolant maps the function back to the Gauss quadrature nodes. With the stable explicit time restrictions particles cannot leave the element's bounds. The method is local and hence can be easily parallelized. The method is shown to have spectral convergence. We present the results for one and two dimensional advection and diffusion equations.

Hareshram Natarajan, Gustaaf Jacobs
San Diego State University
hareshram.n@gmail.com, gjacobs@sdsu.edu

CP13

Fast Finite Difference Schemes for the Time-fractional Diffusion Equation with a Weak Singularity at the Initial Time

In this paper, the numerical solution is considered for the fractional diffusion equation when the solution has a weak singularity near the initial time. In order to avoid the huge storage and computation cost causing by the

nonlocal information of the fractional derivative, an efficient and accurate fast difference scheme is constructed with the help of $L1$ formula on the graded meshes proposed in [Stynes et al,Error analysis of a finite difference method on graded meshes for a time-fractional diffusion equation] and the approximation of the kernel by sum-of-exponentials in [Zhang et al, Commun. Comput. Phys. 21 (2017), 650-678]. The stability and convergence of the fast difference scheme are proved. The convergence order is $O(N^{-\min\{r\alpha,2-\alpha\}} + h^2 + \epsilon)$ in the maximum norm. In addition, a sharp estimate is given for the $L1$ formula on the graded meshes. Some numerical examples are provided. Numerical results verify the convergence rate and illustrate that the fast difference scheme is more advantageous over the direct difference scheme in computational complexity when N is large.

Jinye Shen, Zhizhong Sun
Southeast University
goldlead0811@sina.com, zzsun@seu.edu.cn

Rui Du
southeast university
rdu@seu.edu.cn

CP14

Padé Time Stepping Method of Rational Form for PDEs

One of the difficulties of the numerical solution of time-dependent ordinary and partial differential equations is the severe restrictions on time-step sizes for stability in explicit methods. Otherwise, challenging, generally nonlinear systems of equations in implicit schemes would be imposed to solve such problem. A class of explicit methods based on use of Padé approximation is introduced. They are inexpensive per time-step and they possess stability restrictions similar to the one offered by implicit schemes. Due to the rationality form of PTS, some numerical error occurs and then some a kind of control is imposed. We find that the Padé time stepping (PTS) showed a preferable behaviour when reaction diffusion equations are considered. We also notice that the PTS schemes have less computational time than the compared ones. Finally, numerical runs are conducted to obtain the optimal local error control threshold.

Said Algarni
Assistant Professor - Department of Mathematics
King Fahd University of Petroleum and Minerals
garnis@kfupm.edu.sa

CP14

Efficient Multistage Time Integrators for Incompressible Flows using Projection Methods

We propose a rational framework to analyze the time accuracy of pressure projection methods for incompressible flows. Pressure projection methods are used to enforce mass conservation in computing the Navier-Stokes equations. They consist of replacing the mass continuity equation by a global Poisson equation for the pressure whose solution requires the use of global linear solvers. With the advent of heterogeneous computing architectures, the need to increase arithmetic intensity (ratio of computation to data transfers) is more pressing than ever. One way to accomplish this is to increase time accuracy. For fluid flow simulation, this entails embedding projection methods within high-order integrators such as Runge-Kutta schemes. However, these necessitate solving the Poisson equation for the pressure at each intermediate stage of the Runge-Kutta integrator which increases the data transfer overhead due to global communication required by the linear solver. We propose a time-accuracy analysis framework that enables us to rationally avoid solving the Poisson equation at intermediate stages while maintaining formal order of accuracy. For example, we will show that only one pressure Poisson equation is needed for second and third order Runge-Kutta schemes. This results in savings of up to 30% in the global cost of a single timestep. We will also show how our framework enables the option of selecting the locations where to solve the Poisson equation.

Mokbel Karam
PhD student, Chemical Engineering Department,
University of Utah
mokbel.karam@chemeng.utah.edu

Mike Hansen
Postdoctoral researcher, Sandia National Laboratories
mahanse@sandia.gov

James C. Sutherland
Department of Chemical Engineering
The University of Utah
james.sutherland@chemeng.utah.edu

Tony Saad
Department of Chemical Engineering
University of Utah
tony.saad@chemeng.utah.edu

CP14

Automatic Construction of Scalable Time-stepping Methods for Stiff PDEs

Krylov Subspace Spectral (KSS) Methods have been demonstrated to be highly scalable time-stepping methods for stiff nonlinear PDEs. However, ensuring this scalability requires analytic computation of frequency-dependent quadrature nodes from the coefficients of the spatial differential operator. This thesis describes how this process can be automated for various classes of differential operators to facilitate public-domain software implementation.

Vivian A. Montiforte
The University of Southern Mississippi
vivian.mclain@usm.edu

James V. Lambers
University of Southern Mississippi
Department of Mathematics
James.Lambers@usm.edu

CP15

A Comparison of Some Recently Developed Langevin Integrators on a Course-grained Polymer Melt

Current computational ability does not allow for fully atomistic simulations to reach timescales necessary for the modeling of a wide range of phenomena. Coarse-graining of these systems has been one way in which there has been success in this direction. Consequently, there is significant interest in understanding how current Langevin integrators perform on non-trivial coarse-grained molecular systems, and in particular how large of a timestep one can take with-

out introducing too much discretization error into phase space averages. To investigate this, we examined the recently developed BAOAB and G-JF Langevin integrators, in addition to the well-established BBK integrator, on a coarse-grained polymer melt and computed key statistical properties for each. Our findings indicate that G-JF should be the preferred method.

Joshua Finkelstein
Temple University
joshua.finkelstein@temple.edu

CP15

Reducing Biogeochemical Models using Computational Singular Perturbation

The state of the ocean is constantly changing; it is necessary to produce efficient, simplified models to affordably investigate the role of the ocean in the global carbon cycle and understand how the climate and Earth will evolve. Modeling elementary chemical reactions in ocean fluid dynamics simulations requires significant computing resources, which can be diminished with model reduction techniques. This study uses computational singular perturbation (CSP) methods from combustion research to simplify detailed biogeochemical reaction mechanisms that model ocean carbonate chemistry, reducing the computational cost. The CSP method is applied to three different biogeochemical systems: a 7-species carbonate system, a 17-constituent Biogeochemical Flux Model, and a 56-constituent Biogeochemical Flux Model. This technique analyzes the dynamics of the biogeochemical systems and eliminates modes causing fast timescales. The efficient, simplified models produce similar results to the full models.

Emily F. Klee, Kyle E. Niemeyer
Oregon State University
kleeem@oregonstate.edu, Kyle.Niemeyer@oregonstate.edu

CP15

Hyperelastic Immersed Boundary Finite Element Model of the Human Heart

In this talk we describe the construction and implementation of a computer model of the human heart, which uses hyperelastic solids coupled to the background fluid by the immersed boundary finite element method. In this approach, solid displacements and forces are approximated in the reference configuration on a finite element mesh, while fluid velocity and pressure are approximated on a locally refined, fixed, Cartesian grid. The model contains all four valves, the aorta and pulmonary artery, and realistic fiber reinforced constitutive models for the elastic solid stress. In addition to discussing the numerical approach, implementation, and model construction, we elaborate on some strategies for prescribing boundary conditions to obtain physiological pressures and flows during a cardiac cycle.

Charles Puelz
Courant Institute of Mathematical Sciences
New York University
puelz@cims.nyu.edu

Margaret Anne Smith, Simone Rossi
Department of Mathematics
University of North Carolina at Chapel Hill
masmath@email.unc.edu, simone.rossi@unc.edu

Boyce E. Griffith
Departments of Mathematics and Biomedical Engineering
University of North Carolina at Chapel Hill
boyceg@email.unc.edu

CP15

Two-dimensional Local Fourier Image Reconstruction via Domain Decomposition Fourier Continuation Method

Magnetic resonance imaging (MRI) is a commonly used medical technique. The MRI image is obtained in the spatial domain from the given Fourier coefficients in the frequency domain. However, the Gibbs phenomenon occurs if the function is piecewise smooth. In this talk, I propose an efficient and accurate local reconstruction method with the lower frequency Fourier data that yields sharp image profile near the local edge. The proposed method utilizes only the small number of image data in the local area. Thus the method is efficient. Furthermore the method is accurate because it minimizes the global effects on the reconstruction near the weak edges shown in many other global methods for which all the image data is used for the reconstruction. To utilize the Fourier method locally based on the local non-periodic data, the proposed method is based on the Fourier continuation method. The proposed method first divides the MRI image in the spatial domain into many subdomains and applies the Fourier continuation method on this local area. Then the proposed method reconstructs the local image based on L_2 minimization regularized by the L_1 using sparsity to sharpen the image near edges. The numerical results show that the solution is free of Gibbs oscillations. The numerical results also suggest that the proposed method should be utilized in dimension-by-dimension manner instead of in a global manner for both the quality of the reconstruction and computational efficiency.

Ruonan Shi
University at buffalo
ruonansh@buffalo.edu

Jae Hun Jung
Department of Mathematics
State University of New York at Buffalo
jaehun@buffalo.edu

Ferdinand Schweser
Neurology and Biomedical Engineering Department
State University of New York at Buffalo
schweser@buffalo.edu

CP15

Computational Modeling for Blood Flow Through Stenotic Artery

In this paper, the blood flow, together with magnetic particles in a stenosed artery is studied using magnetohydrodynamic approach. The blood flow without particles interaction is obtained as a particular case. Approximate solutions to the blood velocity, particles velocity and the shear stress are obtained in series forms, in the assumption that the Womersley frequency parameter has small values. It is found that the parameters of flow through the stenotic arteries are significantly disturbed compared to flow through the arteries without stenosis. In the area

where the diameter of stenosed artery is close in size to the diameter of the circular artery, the axial velocity of blood slightly increases, but, increases considerably and reaches its maximum value in the stenosis throat.

Abdulsalam G. Ya'u
Abubakar Tafawa Balewa University Bauchi, P.M.B. 0248
asgital@gmail.com

CP16

A Highly-accurate Finite Element Method for the Solution of the Dirichlet Problem of the Generalized Helmholtz Equation in Non-convex Polygons

In this study, a highly-accurate, conforming finite element method is developed and justified for the solution of the Dirichlet problem of the generalized Helmholtz equation on domains with re-entrant corners. The $k-th$ order Lagrange elements are used for the discretization of the variational form of the problem on exponentially compressed polar meshes employed in the neighbourhood of the corners whose interior angle is $\alpha\pi$, $\alpha \neq 1/2$, and on the triangular and curved mesh formed in the remainder of the polygon. The exponentially compressed polar meshes are constructed such that they are transformed to square meshes using the Log-Polar transformation, simplifying the realization of the method significantly. For the error bound between the exact and the approximate solution obtained by the proposed method, an accuracy of $O(h^k)$, h mesh size and $k \geq 1$ an integer, is obtained in the H^1-norm. Numerical experiments are conducted to support the theoretical analysis made. The proposed method can be applied for dealing with the corner singularities of general nonlinear parabolic partial differential equations with semi-implicit time discretization.

Emine Celiker
University of Lincoln
eceliker@dundee.ac.uk

Ping Lin
University of Dundee
plin@dundee.ac.uk

CP16

An Deformation Based Adaptive Moving Grid Method

The grid generation is very crucial for the accuracy of the numerical solution of partial differential equations, especially for problems with very rapid variations or sharp layers, such as shock waves, wing leading and trailing edges, regions of separation, and boundary layers. The adaptive grid generation is an iterative approach to accommodate these complex structures. In this talk, we introduce a deformation based adaptive grid generation method, in which a differentiable and invertible transformation from computational domain to physical domain is constructed such that the cell volume of the new grid is equal to a prescribed monitor function. A vector field is obtained by solving the div-curl system and can be used to move the grids to the desired locations. By computing the inverse of Jacobian, any deformed grids can also be transformed back to the uniform grid. Several numerical results in two dimensions are presented. Some applications in image registration are discussed.

Yinlin Dong
University of Central Arkansas
ydong5@uca.edu

CP16

A Communication Algorithm for the Patch-based Overlapping Grids Applications

The overlapping grid technique can be used to solve partial differential equations defined on complex computational domains. However, large-scale realistic applications using overlapping grid technique under distributed memory systems are not easy. The grid points do not meet point by point and interpolation is needed. Applications with millions of grid points may consist of many blocks. A proper method that specifies the connectivity among so many blocks may be a challenge in large-scale simulations. Furthermore, a domain-decomposition method is always used in parallel computing. A proper communication schedule among a large number of subdomains may be another challenge. In this talk, we describe a Communication Algorithm for Overlapping Grids Applications with patch-based data structure. Splitting overlapping grids into patches in order to lead to a more computationally balanced workload. The communication algorithm includes a grid mapping method and a communication schedule. Grid mapping method searches donor interpolating grid points based on patch. A tree with patch boxes as its nodes is designed in order to reduce searching cost. A ghost patch assignment strategy is used to facilitate messages pack/unpack routine. A coarse grained communication schedule is implemented in order to reduce communication latency. According to our test results, applications based on this communication algorithm can be run efficiently on thousands of cores.

Hong Guo
Institute of Applied Physics and Computational Mathematics
guo_hong@iapcm.ac.cn

CP16

Generation of Structured Meshes for Shape Optimization

The generation of stable and accurate meshes is a fundamental task in the area of shape optimization. As it usually comes as an iterative process, re-meshing is a key bottleneck in the quest for efficient optimization procedures. Hence, current approaches aim to avoid re-meshing by rather moving the grid points. Although these approaches show good results, re-meshing can't be avoided as the meshes become degenerated after some steps of the iteration. Therefore we present a mixed re-meshing/recycling approach based on the idea of composite finite elements introduced by Hackbusch and Sauter. This approach leads to very regular meshes, where we can recycle most of the grid and the entries of the governing PDE, which lead to a speed up in both the mesh generation and the computing of the matrix. Also, the regular structure contributes to means of efficient implementation on modern computer architectures. We show results for the setting of optimizing the failure probabilities of mechanical components under load (e.g. tensile load), governed by the linear elasticity PDE.

Camilla Hahn
Universität Kassel
chahn@uni-wuppertal.de

Matthias Bolten
Bergische Universität Wuppertal

Fakultät für Mathematik und Naturwissenschaften
bolten@math.uni-wuppertal.de

CP16

A Posteriori Error Estimates and Adaptive Mesh Refinement for the Stokes-Brinkman Problem

The Stokes-Brinkman equation models flow in heterogeneous porous media by combining the Stokes and Darcy models of flow into a single equation. With suitable parameters, the equation can model either flow without detailed knowledge of the interface between the two regions. Thus, the Stokes-Brinkman equation provides an alternative to coupled Darcy-Stokes models. After a brief review of the Stokes-Brinkman equation and its discretization using Taylor-Hood finite elements, we present a residual-based a posteriori error estimate and its use in driving an adaptive mesh refinement process. We compare several strategies for the mesh refinement, and demonstrate its effectiveness by numerical experiments in both 2D and 3D. Finally, we outline possible preconditioning strategies used in iterative solvers of the underlying linear systems.

Kevin Williamson
University of Maryland, Baltimore County
kwillia1@umbc.edu

Pavel Burda
Czech Technical University in Prague
pavel.burda@fs.cvut.cz

Bedrich Sousedik
University of Maryland, Baltimore County
bedrich.sousedik@gmail.com

CP17

A Visual Programming and Debugging Approach for Parallel Numerical Simulation Development

The scientific computing community is sufferring from a lack of good development tool that can handle well the unique problems of coding for high performance computing. It is much more diffcult for domain experts to parallelize inherited serial codes written in FORTRAN which are very common in CSE research field. An interactive visual programming approach is proposed for rapid development of large-scale numerical simulations. A domain specific graphical user interface is designed to facilitate the algorithm design and parallel program construction based on current serial codes. A debugging approach is developed to locate parallel errors automatically. Real applications demonstrated that is is helpful on developing complex numerical applications and enabling increased software productivity greatly.

Li Liao, Aiqing Zhang
Institute of Applied Physics and Computational Mathematics
liliao@iapcm.ac.cn, zhang_aiqing@iapcm.ac.cn

CP17

Superlinear Convergence of the Methods of Quasi-Newton Method Based on Assumptions About the Initial Point

Broyden's method is a quasi-Newton method which is used to solve a system of nonlinear equations. Almost all convergence theory in the literature assumes existence of a root and bounds on the nonlinear function and its derivative in some neighborhood of the root. All these conditions cannot be checked in practice. The motivation of this talk is to derive a convergence theory for Broydenwhere all assumptions can be verified, and the existence of a root and its superlinear rate of convergence are consequences of the theory. The method of BFGS is a quasi-Newton method for unconstrained minimization. Also, all known convergence theory assume existence of a solution and bounds of the function in a neighborhood of the minimizer. The other part of this talk would be on convergence theory of BFGS method where all assumptions are verifiable and existence of a minimizer and also discussing the superlinear convergence of the iteration.

Sarah Nataj
grad Student
ebrahimr@myumanitoba.ca

Shaun Lui
Professor, University of Manitoba
shaun.lui@umanitoba.ca

CP17

Recomputations and Communication Costs of Fast Matrix Multiplication with 2x2 Base Case

Communication costs, between processors and across the memory hierarchy, often dominate the runtime of algorithms. Can we trade these costs for recomputations? Most communication cost lower bounds assume no recomputation, hence do not address this fundamental question. Recently, Bilardi and De Stefani (2017), and Bilardi, Scquizzato, and Silvestri (2018) showed that recomputations cannot reduce communication costs for Strassen's fast matrix multiplication and for FFT. We extend the former and show that recomputations cannot reduce communication costs for a few other fast matrix multiplication algorithms.

Roy Nissim
The Hebrew University of Jerusalem
roynis1@gmail.com

Oded Schwartz
Hebrew University of Jerusalem
odedsc@cs.huji.ac.il

CP17

ArboX: A New Parallel Bounding Volume Hierarchy Implementation Using MPI+kokkos

Search on a set of geometric objects is an important component used in many scientific applications. These applications typically involve a large numbers of search queries through point clouds containing thousands or millions of points, done multiple times within the application run. The performance of search algorithms is thus crucial for the overall performance. In this talk, we present a novel implementation of the distributed bounding volume hierarchy (BVH) construction and search algorithms (k-nearest neighbors and radius) using MPI for inter-node and Kokkos for intra-node shared parallelism. We demonstrate its performance on both CPU and GPU, and compare our results with well known Boost Geometry R-tree and nanoflann libraries.

Andrey Prokopenko, Damien Lebrun-Grandie, Bruno Turcksin

Oak Ridge National Laboratory
prokopenkoav@ornl.gov, lebrungrandt@ornl.gov,
turcksinbr@ornl.gov

Stuart Slattery
Computational Sciences and Engineering Division
Oak Ridge National Laboratory
slatterysr@ornl.gov

CP17

Approximating the Generalized Singular Value Expansion

The singular value expansion (SVE) of a compact operator is invaluable for analyzing Tikhonov regularization. For semi-norm regularization, where a regularization operator such as the gradient is also involved, the generalized singular value expansion (GSVE) can be used in place of the SVE. In this talk, we discuss the approximation of the GSVE of a pair of operators and present sufficient conditions for the convergence of a sequence of discretizations.

Matthew J. Roberts
Graduate Student Michigan Technological University
majrober@mtu.edu

CP18

Hyperbolic Method for Incompressible Navier-Stokes Equations on Unstructured Grids

Recently, a new, efficient, and accurate method for reformulating the diffusion equation into a set of subsequent first order hyperbolic system is proposed by Nishikawa[JCP, 2007]. This method is based on the reformulation of elliptic nature of the viscous fluxes into a set of augmented variables that makes the entire system to be hyperbolic. In this paper, a hyperbolic incompressible Navier-Stokes (Hyper-INS) system is presented and the solution is obtained by cell-centered finite volume method using unstructured meshes. A very promising result, the uniform order of accuracy not only for the solution variables but also for their gradients are obtainable, will be presented. Efficiency and accuracy of the method will be presented with rigorous analysis of representative test cases, including Kovasznay flow, driven cavity problem, flow past a circular cylinder in 2D, and flow past a sphere in 3D.

Hyung T. Ahn
University of Ulsan
htahn@ulsan.ac.kr

CP18

On the Calculation of Exact Cumulative Distribution Statistics for Burgers Equation

Analytical solutions of nonlinear PDEs with uncertain initial and/or boundary data are invaluable benchmarks in assessing uncertainty quantification techniques. A mathematical procedure is presented for the calculation of "exact' (modulo root finding) cumulative distribution statistics for a viscosity-free variant of Burgers nonlinear partial differential equation (PDE) in a single periodic space dimension and forward time dimension subject to sinusoidal initial data with uncertain (random variable) phase shift. Specifically, for probability event outcomes, $\omega \in \Omega$

$$\partial_t u_{\mathbf{X}} + \partial_x u_{\mathbf{X}}^2/2 = 0$$

$$u_{\mathbf{X}}(x,0,\omega) = A\sin(2\pi(x+\mathbf{X}(\omega))) \quad (1)$$

where $u_{\mathbf{X}}(x,t,\omega) : [0,1] \times \mathbf{R}^+ \times \mathbf{\Omega} \to \mathbf{R}$ and $A > 0$. Our keen interest in this problem comes from the observation that the resulting solution is only piecewise nonsmooth in both physical and random variable dimensions. This can significantly complicate the task of computing statistics for derived output quantities of interest. Thus, this problem provides a severe benchmark test case for approximate uncertainty quantification techniques. In the present work, we prove that this cumulative distribution function can be computed exactly, i.e. the calculation requires, at most, root finding of a strictly convex function containing an isolated root in a bracketed interval.

Timothy J. Barth
NASA Ames Research Center
Timothy.J.Barth@nasa.gov

Jonas Sukys
Swiss Federal Institute of Aquatic Science and Technology
jonas.sukys@eawag.ch

CP18

t Going to Infinity Limit of Solutions for Linear Advection Problems

Consider the linear advection equation,

$$\phi_t + \vec{v}.\nabla\phi = 0. \quad (2)$$

which moves a scalar quantity $\phi(\vec{x},t)$ by a given velocity field $\vec{v}(\vec{x},t)$. This equation (1) is augmented with initial condition $\phi(\vec{x},0) = \Phi(\vec{x})$ and boundary conditions. There are several areas where we need to solve advection with a large number of time steps and that motivates us to look for a method which can produce nice solution when t is very large or one can ask for the method which can produce nice solution when $t \to \infty$. we propose a scheme known as the jet scheme with nonlinear interpolant which can produce a nice solution at least for linear advection with constant coefficient. We fix a final time t_f and solve the problem by a numerical scheme for a fixed resolution h to get the solution $\phi^h(t_f)$. Then we take limit first as $t_f \to \infty$ to get $\phi^h(\infty)$ and then limit as $h \to 0$ to get $\phi(\infty)$, which is our solution at $t \to \infty$. On the other hand we can first take the limit as $h \to 0$ to the solution $\phi^h(t_f)$ to get $\phi(t_f)$, the solution at time t_f and then take the limit as $t_f \to \infty$ to get $\phi_1(\infty)$. Now we ask the question is $\phi(\infty) = \phi_1(\infty)$ for this numerical scheme? The method we propose gives an affirmative answer to this question.

Abhijit Biswas
Temple University
N/A
tug14809@temple.edu

CP18

Discontinuous Galerkin Methods for the Nonlinear Shallow Water Equations with Horizontal Temperature Gradients

Shallow water equations with horizontal temperature gradients, also known as the Ripa model, are used to model flows when the temperature fluctuations play an important role. These equations admit steady state solutions where the fluxes and source terms balance each other. We will present well-balanced discontinuous Galerkin methods for the Ripa model, which preserve the lake-at-rest and

isobaric steady states as well as the more general moving water equilibrium. The key idea is the recovery of well-balanced states, source term approximations, and appropriate approximations of the numerical fluxes. We will also discuss the positivity preserving property of our methods. Numerical examples will also be presented to verify the well-balanced property, high order accuracy, and good resolution for both smooth and discontinuous solutions.

Jolene Britton
Department of Mathematics
University of California Riverside
jhout001@ucr.edu

Yulong Xing
The Ohio State University
xing.205@osu.edu

CP18

Ssp and Positivity for Implicit Methods

The state-of-the-art theory for strong stability preservation (SSP) and forward invariance of sets (positivity) relies on the same property of the Forward Euler (FE) method with a positive step size. However, this theory does not apply in many cases of practical importance (e.g. FEM without lumping, spectral methods). In this talk we shall present results for diagonally implicit methods (e.g. diagonally implicit Runge-Kutta methods) that use assumtions only on the property of the Backward Euler (BE) method, instead of the FE condition. We quantify the SSP/positivity step sizes in terms of a measure of the BE-step on the given problem and a new parameter of the scheme. In addition, the new arguments propose an efficient implementation procedure as well. To illustrate the new theoretical findings, we shall present results of computational experiments for CFD problems arising in automotive industry.

Zoltan Horvath
Széchenyi István University
horvathz@sze.hu

CP19

Modified Kernel Regression for Building Low-dimensional Models of Complex Dynamical Systems

Many dynamical systems with a high-dimensional state-space have low-dimensional attractive manifolds that occupy a small fraction of the original state space. Combustion is an example of such a system. One challenge in taking advantage of these low-dimensional manifolds is identifying suitable bases for their parameterization. We show that Principal Component Analysis (PCA) can be used to build low-dimensional kinetic models for complex chemically reacting systems. Producing robust models, however, can be challenging due to unconstrained behavior of source term regression models in regions off the training manifolds. Even the smallest of deviations from the training manifold can result in unphysical and/or unstable solutions due to the undefined behavior in areas devoid of training data. We have addressed this problem by developing a kernel regression technique that modifies source terms at the edges of training data to attract the system back to the original training manifold. We show results for this approach applied to several systems.

Elizabeth Armstrong
University of Utah
elizabeth.armstrong@chemeng.utah.edu

James C. Sutherland
Department of Chemical Engineering
The University of Utah
james.sutherland@chemeng.utah.edu

CP19

Bayesian Geometric Convolutional Neural Networks for Engineering Applications

Deep convolutional neural networks (CNNs) are a corner stone for the machine learning field providing the foundation for accurate image recognition, prediction and generation. However, traditional CNN implementations have the fundamental limitation of being constrained to structured, euclidean domains. Thus standard CNNs are not practical for non-uniform, unstructured domains/meshes ubiquitous to many engineering applications including computational fluid dynamics and solid mechanics. Recently the extension of CNNs to unstructured domains, such as graphs, meshes and point clouds, has become a topic of interest. In this work we take these generalizations of CNNs, referred to as geometric deep learning, and extend them to dynamical systems of engineering interest for modeling and uncertainty quantification. We will discuss the incorporation of prior physical knowledge provided by the unstructured domain and the implementation of such deep learning architectures in the Bayesian paradigm. Applications will have a focus on high-dimensional non-linear fluid simulation such as flow through porous media.

Nicholas Geneva, Nicholas Zabaras
University of Notre Dame
ngeneva@nd.edu, nzabaras@gmail.com

CP19

Physics-informed Machine Learning Approach for Augmenting Turbulence Models

Reynolds-averaged Navier-Stokes (RANS) models have large model-form uncertainties in engineering applications that involve turbulent flows. Recently, data-driven methods have been proposed as a promising alternative by using existing databases of experiments or high-fidelity simulations. In this talk, we present a comprehensive framework for augmenting turbulence models with physics-informed machine learning. The learned model satisfies two key requirements in turbulence modeling: Galilean invariance and coordinate rotational invariance. This framework consists of three components: (1) extracting invariant inputs and outputs to ensure important invariances of the machine-learning-assisted model, (2) constructing a functional mapping between inputs and outputs via machine learning techniques, (3) assessing the prediction confidence a priori by evaluating the distance between different flows in the mean flow features space, and (4) propagating the predicted Reynolds stress field to mean velocity field by using physics-informed stabilization. With the physics-informed machine learning augmentation, significant improvements over the baseline RANS simulation are observed for the Reynolds stress and the mean velocity fields.

Jinlong Wu
Dept. of Aerospace and Ocean Engineering, Virginia Tech
jinlong@vt.edu

Carlos Michelen

Dept. of Aerospace and Ocean Engineering
Virginia Tech
unknown

Jianxun Wang
University of Notre Dame
jwang33@nd.edu

Heng Xiao
Dept. of Aerospace and Ocean Engineering, Virginia Tech
hengxiao@vt.edu

CP20

Probably Stable and Conservative Spectrally-tunable Artificial Dissipation Operators

The incorporation of artificial dissipation (AD) is often associated with scheme stabilization; however, a concomitant role involves the mitigation of numerical error (e.g., dispersion, aliasing, etc.) particularly for time-accurate calculations. As such, spectrally-tunable AD methods are attractive as they can be tailored to characteristics of the base numerical method as well as the physical problem. In the case of scale-discriminant (i.e., spectrally sharp) attenuation, one is further challenged by sudden changes in topological features such as shocks/sharp gradients or grid discontinuities. Accommodating such situations requires a stencil adaptation strategy that provides resolution while also avoiding the generation of false over/undershoots. In the current work, these needs are addressed by a domain decomposition technique. Previous efforts in the development of explicit and implicit filtering operators on bounded domains are leveraged in order to build provably stable and conservative artificial dissipation discretizations. The applicability of such schemes for fluid flow simulations will thus explored.

Ayaboe K. Edoh
ERC, Inc. (AFRL)
ayaboe.edoh.ctr@us.af.mil

CP20

Numerical Advection Through PDE-Based Flow Map Composition

We propose a novel numerical methodology to compute the advective transport of a passive tracer under the influence of a dynamic velocity field. This method yields numerical solutions almost devoid of any numerical errors, while allowing for direct parallelization in the temporal direction. The tracer advection is computed by implicitly solving the characteristic evolution through a modified transport partial differential equation and domain decomposition in the temporal direction followed by composition with the known initial condition. We prove that there exists an optimal value of the numerical time-step which yields the minimum total error in the computations. Finally, applications of this methodology are illustrated using three examples, namely an analytical swirl flow, a two-dimensional flow exiting a strait, and a realistic flow in the Bismarck Sea. These examples highlight the methods theoretical aspects as well as its applicability in realistic simulations.

Chinmay S. Kulkarni
Massachusetts Institute of Technology
chinmayk@mit.edu

Pierre F. Lermusiaux
MIT
pierrel@mit.edu

CP20

Comparison of Weno Schemes to Second Order Methods for Capturing Compressible Turbulent Spectra

High-order methods are often advocated as superior because the rate of convergence of the discretization error decreases faster compared to low-order methods. In the context of compressible simulations, WENO schemes are interesting because of their ability to capture shocks and discontinuities, while keeping an overall high order rate of convergence. However, very few studies exist about the performance of these numerical methods in a realistic context where shock waves and a broadband turbulence flow evolve in the same domain. The proposed lecture will compare different recently developed WENO formulations (WENO-M, WENO-Z, MDCD, TENO) implemented in a fourth-order conservative finite volume method, to a finite volume second order Godunov-type method with PPM interpolation. In particular, we are interested in energy dissipation of turbulent spectra prevalent in WENO schemes, both considering dissipation of total energy over time as well as density and vorticity throughout the power spectrum. Though WENO schemes are formally higher order, it turns out that the second order method provide a solution that is nearly as accurate but with less computational cost.

John P. Wakefield
University of Michigan
jwake@umich.edu

Emmanuel Motheau
Lawrence Berkeley National Lab
Center for Computational Sciences and Engineering
emotheau@lbl.gov

CP20

A Second Order Fully-discrete Linear Energy Stable Numerical Scheme of a Binary Compressible Viscous Fluid Model

A thermodynamically consistent hydrodynamic phase field model of binary compressible fluid flow mixtures derived using the generalized Onsager Principle, which warrants not only the variational structure, but also the mass, linear momentum conservation, and the energy dissipation law in the isothermal case. We present a linear, second order fully discrete numerical scheme to solve this mathematical model on a staggered grid. The fully discrete scheme respects a discrete energy dissipation law. We present the scheme in 2D space for simplicity. Results apply to a 3D case as well. We prove the unique solvability of the linear scheme rigorously. We present several numerical examples, including phase separation due to the spinodal decomposition of two polymeric fluids and calculation of the equilibrium states of a gas-liquid mixture, to show the convergence property, stability, and efficiency of the new scheme.

Xueping Zhao
University of South Carolina

zhaoxueping0802@gmail.com

CP21

An Elastic Linear Deformation Model with Discrete Fractures at Core Scale

In this work, an elastic linear deformation model for a linear, homogeneous and isotropic elastic solid with fractures at core scale under laboratory conditions is presented. The model is derived on base of stationary momentum balance equation and considering Hooke's law as a constitutive relationship of stress-strain. Whereas, the fractures are modeling following a discrete fracture approach considering fractures as simplified geometric entities with one dimension lower than the domain. For the numerical solution is applied a finite element method and its computational implementation is carried out in an open source platform. From the methodological point of view, each stage of model development (conceptual, mathematical, numerical and computational) is described. Finally, the resulting model is numerically validated in a case study for a triaxial test cell.

Martin A. Diaz-Viera, Manuel Coronado, Edscott Wilson-Garcia, Pedro Aguilar-Gastelum, Mario Vadillo-Saenz
INSTITUTO MEXICANO DEL PETROLEO
mdiazv@imp.mx, mcoronad@imp.mx, edscott@imp.mx, faguilar@imp.mx, mevadillo@imp.mx

CP21

Numerical Methods for Thermally Stressed Shallow Shell Equations

We develop efficient and accurate numerical methods to solve a class of shallow shell problems of the von Karman type. The governing equations form a fourth-order coupled system of nonlinear biharmonic equations for the transverse deflection and Airys stress function. A second-order finite difference discretization with three iterative methods (Picard, Newton and Trust-Region Dogleg) are proposed for the numerical solution of the nonlinear PDE system. Three simple boundary conditions and two application-motivated mixed boundary conditions are considered. Along with the nonlinearity of the system, boundary singularities that appear when mixed boundary conditions are specified are the main numerical challenges. Two approaches that use either a transition function or local corrections are developed to deal with these boundary singularities. All the proposed numerical methods are validated using carefully designed numerical tests, where expected orders of accuracy and rates of convergence are observed. A rough run-time performance comparison is also conducted to illustrate the efficiency of our methods. As an application of the methods, a snap-through thermal buckling problem is considered. The critical thermal loads of shell buckling with various boundary conditions are numerically calculated, and snap-through bifurcation curves are also obtained using our numerical methods together with a pseudo-arclength continuation method.

Hangjie Ji
UCLA Department of Mathematics Mathematical Sciences 7354
hangjie@math.ucla.edu

Longfei Li
University of Louisiana at Lafayette
longfei.li@louisiana.edu

CP21

Plasticity in the Oof Finite-Element Solver for Materials Science

The Object-Oriented Finite Element code (OOF) developed at NIST is a long-standing project to deliver high-quality mathematics and computational capabilities to an audience of materials scientists. The code features tools to easily mesh complex 3D microstructures, and a scheme for the addition of custom constitutive rules, allowing users to conduct sophisticated structure-property explorations. In collaboration with domain experts from the computational mechanics community, the development team has added a crystal plasticity modeling capability to the this tool, allowing it to model systems which acquire a permanent deformation in response to loads. Being history-dependent, this expanded the OOF project beyond its initial scope of what were essentially divergence equations. Incorporating the insights and algorithms of the computational mechanics experts into this general-purpose tool posed several challenges to the OOF team, but the result is a tool which can generate microstructural meshes from images, and allow users to plug in various crystal-plastic constitutive rules to conduct high-value virtual experiments on yielding systems.

Andrew Reid
Center for Theoretical and Computational Materials Science
National Institute of Standards and Technology
andrew.reid@nist.gov

CP21

A Projection Method for Simulating Incompressible Poroelastic Media

A variety of materials such as contractile gels and living tissues are well-modeled as incompressible poroelastic media, whereby they are treated as a mixture of fluid and solid phases, subject to a constraint that the total mass flux is divergence-free. In this talk, we show that the projection method of Chorin (1968) for incompressible fluid flow can be extended to simulate incompressible poroelastic media. The method involves solving a Poisson problem for pressure, which is sufficient to project the velocities of both fluid and solid phases to maintain the incompressibility constraint. We apply the method to an active gel model, and examine the instabilities arising from active contractile stresses.

Chris H. Rycroft, Nicholas Derr
Harvard University
chr@seas.harvard.edu, derr@g.harvard.edu

Christoph Weber
Max Planck Institute for the Physics of Complex Systems
weber@pks.mpg.de

L Mahadevan
Harvard University
lm@seas.harvard.edu

MS1

Neural Ordinary Differential Equations

We introduce a new family of deep neural network mod-

els. Instead of specifying a discrete sequence of hidden layers, we parameterize the derivative of the hidden state using a neural network. The output of the network is computed using a blackbox differential equation solver. These continuous-depth models have constant memory cost, adapt their evaluation strategy to each input, and can explicitly trade numerical precision for speed. We demonstrate these properties in continuous-depth residual networks and continuous-time latent variable models. We also construct continuous normalizing flows, a generative model that can train by maximum likelihood, without partitioning or ordering the data dimensions. For training, we show how to scalably backpropagate through any ODE solver, without access to its internal operations. This allows end-to-end training of ODEs within larger models.

Ricky T. Q. Chen, David Duvenaud, Yulia Rubanova, Jesse Bettencourt
University of Toronto
rtqichen@cs.toronto.edu, duvenaud@cs.toronto.edu, rubanova@cs.toronto.edu, jessebett@cs.toronto.edu

MS1

Simultaneous Parallel-in-layer Optimization for Training of Deep Residual Networks

The forward propagation through residual neural networks (ResNet) has been shown to be equivalent to solving an initial value problem whose nonlinear dynamics are parameterized by the weights. Training a ResNet is therefore related to minimizing a regularized loss function with respect to those weights. In order to develop faster, more efficient, and more effective training algorithms, we leverage recent advances in parallel-in-time integration and optimization methods. We replace the forward and backward propagation through the ResNet by an iterative multigrid-in-time technique that allows for computational concurrency across the neural network layers. Distributing the time domain of the nonlinear dynamics across multiple processors allows for a parallel-in-layer optimization technique where speedup over serial-in-layer optimization methods can be achieved through the greater concurrency. Additionally, the multigrid iterations enable a simultaneous optimization framework where weight updates are based on inexact gradient information. Weight updates can thus be employed at an early stage of the propagation process which potentially further reduces the overall training runtime.

Stefanie Guenther
RWTH Aachen University
stefanie.guenther@scicomp.uni-kl.de

Lars Ruthotto
Department of Mathematics and Computer Science
Emory University
lruthotto@emory.edu

Jacob B. Schroder
Lawrence Livermore National Laboratory
schroder2@llnl.gov

Nicolas R. Gauger
TU Kaiserslautern
nicolas.gauger@scicomp.uni-kl.de

MS1

Adversarial Regularizers in Inverse Problem

Inverse Problems in medical imaging and computer vision are traditionally solved using purely model-based methods. Among those variational regularization models are one of the most popular approaches. We discuss a new framework for applying data-driven approaches to inverse problems, using a neural network as a regularization functional. The network learns to discriminate between the distribution of ground truth images and the distribution of unregularized reconstructions. Once trained, the network is applied to the inverse problem by solving the corresponding variational problem. Unlike other data-based approaches for inverse problems, the algorithm can be applied even if only unsupervised training data is available.

Sebastian Lunz
University of Cambridge
sl767@cam.ac.uk

Carola-Bibiane Schönlieb
DAMTP, University of Cambridge
cbs31@cam.ac.uk

Ozan ktem
KTH - Royal Institute of Technology, Sweden
ozan@kth.se

MS1

Simple and Effective PDE Based Approaches to Deep Learning

We will draw from some of our new results in optimization, related to partial differential equations, to improve the performance of algorithms ranging from data dependent activation in deep learning, training quantized neural networks, optimizing neural networks, new approaches to gradient descent and solving the phase lift problem.

Bao Wang
UCLA
wangbaonj@gmail.com

Stanley J. Osher
University of California
Department of Mathematics
sjo@math.ucla.edu

Li Zhen
Hong KHong University of Science and Technology
zli12@mails.tsinghua.edu.cn

Wei Zhu
Department of Mathematics
University of California, Los Angeles
weizhu731@math.ucla.edu

Zuoqiang Shi
Tsinghua University
zqshi@math.tsinghua.edu.cn

Adam M. Oberman
Department of Mathematics and Statistics
adam.oberman@mcgill.ca

Pratik A. Chaudhari, Penghang Yin
UCLA
pratikac@ucla.edu, yph@g.ucla.edu

Xiyang Luo
University of California, Los Angeles

xylmath@gmail.com

Minh Pham, Alex Lin
UCLA
minhrose@math.ucla.edu, atlin@math.ucla.edu

MS2

Development and Integration Workflows for Large Complex Distributed CSE Software Efforts

Larger research-driven Computational Science and Engineering (CSE) projects pose a variety unique challenges. Many of these projects involve multiple institutions, teams of developers, and software repositories. In addition, the push to exascale requires having to run performantly on a variety of bleeding-edge, multi-core and GPU accelerator architectures. This talk will describe the software engineering development and integration challenges and solution approaches required to effectively drive these projects that balance speed of development and integration with the stability of the software needed to keep developers, domain experts and analysts productive as they push the scientific frontier. This will include experiences from some larger CSE projects including the Consortium for the Advanced Simulation of Light-water reactors (CASL) DOE Energy Innovation Hub and the Exascale Computing Projects (ECP) Advanced Technology Deployment Mitigation (ATDM) project involving the Sandia National Labs Trilinos development repository and customer application codes.

Roscoe A. Bartlett
Sandia National Laboratories
rabartl@sandia.gov

MS2

Runtime Data Analysis for CSE Applications

CSE simulations in high-performance computers are still costly. They often involve the selection of many computational parameters and options. The set-up of such parameters is usually a trial-and-error process even for experienced users. In this talk, we will show how to extend in-situ visualization techniques with in-transit data analysis to provide information to help control the simulations at runtime. Often, by only observing a region of interest, an experienced analyst can infer that something is not going well, deciding to stop it or change parameters. However, to do that, visual information should be complemented with information regarding the evolution of quantities of interest. We use for the simulations the libMesh library, which provides a platform for parallel, adaptive, multiphysics finite element computations. We discuss the integration of libMesh with in-situ visualization and in-transit data analysis tools. We present a parallel performance analysis showing that the overhead for both in-situ visualization and in-transit data analysis is negligible. Our tools enable monitoring the quantities of interest at runtime and steer the simulation based on the solver convergence or other data and visual information. The data analysis tool registers the provenance of the simulation data for reproducibility, including registering the runtime changes on simulation parameters.

Alvaro Coutinho
Dept. of Civil Engineering
Federal University of Rio de Janeiro
alvaro@nacad.ufrj.br

Marta Mattoso
Fedral University of Rio de Janeiro
marta@cos.ufrj.br

José Camata
NACAD - High Performance Computing Center
COPPE/Federal University of Rio de Janeiro
camata@nacad.ufrj.br

Vitor Souza, Linda Gsenhues, Renan Souza
Federal University of Rio de Janeiro
vitor.silva.sousa@gmail.com, linda@nacad.ufrj.br, renanfrancisc@gmail.com

MS2

How to Assure Quality of Software Preservation Early in a Project Life Cycle and Ongoing Efforts

With the increasing importance of software for research, preservation of software has become a challenge that many researchers face. Version control systems such as GitHub and containerizations such as Docker can deliver part of the solution but preservation of software requires tools for assuring reusability of the computational methods, reproducibility of results as well as advanced preservation features, e.g., keyword assignment, fixity tools and indicators for the quality of the preservation. Researchers typically reach out to digital librarians for support for the preservation process at the end of the lifecycle of projects. The point of time creates not only a tight schedule and is more labor intensive but also risks the loss of important intermediate data. The project PresQT (Preservation Quality Tool) funded by IMLS tackles these challenges by engaging stakeholders in a collaborative effort to enhance reproducibility and more open sharing of software and data through open source development of services. PresQT is repository and technology agnostic and engages partners for testing against diverse existing solutions like ReproServer, Fedora, HUBzero and JupyterHub, as well as Software Preservation Networks EaaSI emulation service, and dashboards from the National Data Service. Advisors and testing collaborators from Confederation of Open Access Repositories, CERN, RDA, Science Gateways Community Institute, and West Big Data Innovation Hub inform priorities.

Sandra Gesing, Natalie K. Meyers, Rick Johnson, Zheng Wang
University of Notre Dame
sandra.gesing@nd.edu, natalie.meyers@nd.edu, rick.johnson@nd.edu, zheng.wang@nd.edu

MS2

Toward Automatic Generation of Scientific Software Artifacts

Software qualities, like maintainability, reproducibility and verifiability, often suffer for Scientific Computing Software (SCS) because of inadequate documentation, traceability, change-enabling design, and testing. Software developers would like to spend more time on documentation, and other software activities, but time and resource pressures frequently make this an unrealistic goal. Ideally, developers should be able to create traceable documentation, design, code, build scripts and tests, without the drudgery of writing and maintaining them. A potential solution is to generate the documentation, code and tests automatically by using Domain Specific Languages (DSLs) over a base of scientific, computing and documentation knowledge. This is the

approach that is proposed for a new scientific software development framework called Drasil. By having one source of knowledge, along with rules for transforming and consistency checking, the qualities of completeness, consistency and traceability can be achieved by construction. Moreover, these qualities can be maintained as requirements are modified, design decisions changed, and documentation standards are varied. New projects will reuse portions of the existing knowledge and expand the base when there is new information. Improving the documentation level qualities should mean an improvement in the indirect qualities of maintainability, reproducibility and verifiability.

Spencer Smith
McMaster University, Canada
smiths@mcmaster.ca

MS3

Compressed Anomaly Detection with Multiple Mixed Observations

We consider a collection of independent random variables that are identically distributed, except for a small subset which follows a different, anomalous distribution. We study the problem of detecting which random variables in the collection are governed by the anomalous distribution. Recent work proposes to solve this problem by conducting hypothesis tests based on mixed observations (e.g. linear combinations) of the random variables. Recognizing the connection between taking mixed observations and compressed sensing, we view the problem as recovering the "support" (index set) of the anomalous random variables from multiple measurement vectors (MMVs). Many algorithms have been developed for recovering jointly sparse signals and their support from MMVs. We establish the theoretical and empirical effectiveness of these algorithms at detecting anomalies. We also extend the LASSO algorithm to an MMV version for our purpose. Further, we perform experiments on synthetic data, consisting of samples from the random variables, to explore the trade-off between the number of mixed observations per sample and the number of samples required to detect anomalies.

Natalie Durgin
Spiceworks
njdurgin@gmail.com

Rachel Grotheer
Department of Mathematics
Goucher College
rachel.grotheer@goucher.edu

Chenxi Huang
Center for Outcomes Research & Evaluation (CORE)
Yale University
chenxi.huang@yale.edu

Shuang Li
Colorado School of Mines
shuangli@mymail.mines.edu

Anna Ma
Claremont Graduate University
annama@ucsd.edu

Deanna Needell
University of California Los Angeles
deanna@math.ucla.edu

Jing Qin
Department of Mathematics
University of California, Los Angeles
jxq@ucla.edu

MS3

Path-based Spectral Clustering: Guarantees, Robustness to Outliers, and Fast Algorithms

This talk will discuss new performance guarantees for robust path-based spectral clustering and an efficient approximation algorithm for the longest leg path distance (LLPD) metric, which is based on a sequence of multiscale adjacency graphs. LLPD-based clustering is informative for highly elongated and irregularly shaped clusters, and we prove finite-sample guarantees on its performance when random samples are drawn from multiple intrinsically low-dimensional clusters in high-dimensional space, in the presence of a large number of high-dimensional outliers. More specifically, we derive a condition under which the Laplacian eigengap statistic correctly determines the number of clusters for a large class of data sets, and prove guarantees on the number of points mislabeled by our method. Our methods are quite general and provide performance guarantees for spectral clustering with any ultrametric. We also propose a fast algorithm for implementing path-based spectral clustering which has complexity quasilinear in the number of data points.

Anna Little
Jacksonville University
littl119@msu.edu

James Murphy
Tufts University
jmurphy@math.jhu.edu

Mauro Maggioni
Johns Hopkins University
mauro@math.jhu.edu

MS3

Stochastic Methods for High Dimensional Data with Multiple Measurement Vectors

Nowadays, data is exploding at a faster rate than computer architectures can handle. For that reason, mathematical techniques to analyze large-scale data must be developed. The use of sparsity and stochastic designs has gained attention due to its ability to capture small amounts of intrinsic information embedded in high dimensional signals using few measurements. On the recovery side, stochastic methods can accurately estimate sparse, large-scale signals from a small number of measurements in the underdetermined setting, as well as solve large-scale systems in the highly overdetermined setting. In this talk, we present three approaches that address large-scale problems in the presence of multiple measurement vector (MMV) for various settings. The work presented is a product of the Women in Data Science and Mathematics (WiSDM) workshop at Institute for Computational and Experimental Research in Mathematics (ICERM).

Natalie Durgin
Spiceworks
njdurgin@gmail.com

Rachel Grotheer
Department of Mathematics

Goucher College
rachel.grotheer@goucher.edu

Chenxi Huang
Center for Outcomes Research & Evaluation (CORE)
Yale University
chenxi.huang@yale.edu

Shuang Li
Colorado School of Mines
shuangli@mymail.mines.edu

Anna Ma
Claremont Graduate University
annama@ucsd.edu

Deanna Needell
University of California Los Angeles
deanna@math.ucla.edu

Jing Qin
Department of Mathematics
University of California, Los Angeles
jxq@ucla.edu

MS3

Heuristic Framework for Multi-scale Testing of the Multi-manifold Hypothesis

Global linear models often overestimate the number of parameters required to analyze or efficiently represent datasets, for example when a data set in sampled from a manifold of lower dimension than the ambient space. The manifold hypothesis consists in asking whether data lies on or near a d-dimensional manifold or is sampled from a distribution supported on a manifold. In this talk, we outline a heuristic framework for a hypothesis test suitable for computation and empirical data analysis. We consider two datasets made of multiple manifolds and test our manifold hypothesis on a set of spline-interpolated manifolds constructed based variance-based intrinsic dimensions computed from the data.

Karamatou Yacoubou Djima
Amherst College
kyacouboudjima@amherst.edu

Ness Linda
Dimacs
Rutgers University
nesslinda@gmail.com

Patricia Medina
Worcester Polytechnic Institute
patriciamg90@gmail.com

Melanie Weber
Princeton University
mw25@math.princeton.edu

MS4

Rank Approximation of a Tensor with Application in Image Processing

The goal of this talk is to find a low-rank approximation for a given third order tensor. In particular, I give an iterative technique for approximating the rank of a tensor by formulating it as an sparse optimization problem via l1-regularization. The performance of the algorithm is tested on randomly generated data as well as real data such as videos and images.

Ramin Goudarzi
University of Alabama at Birmingham
ramin@uab.edu

MS4

Optimal Closures in Hydrodynamic Models

Turbulent flows at high Reynolds numbers continue to challenge both scientists studying their fundamental properties and engineers interested in diverse technical applications involving fluid mechanics. In particular, accurate and efficient numerical simulation of turbulent flows will remain an open problem in computational science for the foreseeable future. As a result, one must rely on solving various simplifications and obtain approximate solutions of the flow problem. One such approach which has gained widespread popularity in engineering practice relies on the so-called Large-Eddy Simulation (LES). In this talk, I will introduce a computational framework for determining optimal closures of the eddy-viscosity type for LES of a broad class of PDE models, such as the Navier-Stokes equation. The proposed framework is thoroughly tested on a model problem involving the LES of the 1D Kuramoto-Sivashinsky equation, where the optimal closure relations are obtained as generalizations of the standard Smagorinsky model. Since this leads to a PDE optimization problem with a nonstandard structure, the solution is obtained computationally with a flexible and efficient gradient approach, which will be discussed along with some results.

Pritpal Matharu
McMaster University
mathap1@mcmaster.ca

Bartosz Protas
Dept of Mathematics and Statistics
McMaster University
bprotas@mcmaster.ca

MS4

A Two-Way Coupled Fluid-Kinetic Solver for the Vlasov-Poisson System that Conserves Mass, Momentum, and Energy

High-order semi-Lagrangian time-stepping schemes such as those developed in Rossmanith and Seal [J. Comput. Phys., 230 (2011), pp. 6203-6232] are generally able to produce accurate and efficient numerical solutions of kinetic plasma systems such as the Vlasov-Poisson equations. One shortcoming of these schemes is that they do not exactly conserve momentum and energy. On the other hand, numerical methods that directly discretize the fluid equations (i.e., the equations of mass, momentum, and energy conservation) can be easily constructed to achieve exact conservation of these quantities; however, fluid models, unlike kinetic models, are generally unable to capture solutions far from thermodynamic equilibrium. In this work, we develop a two-way coupled fluid-kinetic solver that achieves benefits from both approaches. The resulting solver is consistent with all flow regimes, near to or far from thermodynamic equilibrium, and is guaranteed to locally conserve mass and total energy and globally conserve momentum. The resulting numerical scheme is validated on sev-

eral standard plasma physics problems.

Lopamudra Palodhi
Iowa State University
lpalodhi@iastate.edu

James A. Rossmanith
Iowa State University
Department of Mathematics
rossmani@iastate.edu

MS4

Application of Vector Autoregression to the Complex Dynamics of Multivariate Unigram Frequency Time-Series

There are three key features we observe on the unigram time-series data from the Google unigram corpus. (1) The time-series of each word are autoregressive, (2) each time-series affects other time-series intertemporally, and (3) the peaks and valleys of the word time-series are reflective of the volatile behaviors of economic systems. In this study, we explore the dynamics of the unigram frequency evolution of eight languages from the Google books corpus. Similar to the behaviors of stock market trends and the changes happening in molecular evolution, the frequency of words are susceptible to cultural and political changes. It is our hypothesis that natural language systems collectively deviates from neutral evolution because of at least one word undergoing volatile lexical changes in frequency. We developed a vector autoregression inspired model with an external forcing term to simulate rapid changes in word frequency while accounting for the frequency change relationships among words. Our results establish a generative model that can simulate word frequency evolution and have shown that the volatile changes in frequency and the collective behaviors among words are important.

Alex John Quijano
University of California Merced
aquijano4@ucmerced.edu

Rick Dale
University of California Los Angeles
rdale@ucla.edu

Suzanne Sindi
University of California, Merced
ssindi@ucmerced.edu

MS4

Viscous Surface Wave Problem with Generalized Surface Energies

The viscous surface wave problem consists of the study of a three-dimensional incompressible fluid in a horizontally periodic domain with finite depth whose boundary is the graph of a function. The fluid is subject to gravity and generalized forces arising from a surface energy. The surface energy incorporates both bending and surface tension effects. It can be proven that for initial conditions sufficiently close to equilibrium the problem is globally well-posed and solutions decay to equilibrium exponentially fast, in an appropriate norm. In this talk I will discuss how the proof is centered around a nonlinear energy method that is coupled to careful estimates of the fully nonlinear surface energy.

Antoine Remond-Tiedrez, Ian Tice
Carnegie Mellon University
aremondt@andrew.cmu.edu, iantice@andrew.cmu.edu

MS4

Compressive Hermite Interpolation: Sparse, High-Dimensional Approximation from Gradient-Augmented Measurements

Many physical problems involve approximating high-dimensional functions with a limited number of sampling points. It is seen that the high-dimensional function interpolation problem has various applications such as uncertainty quantification. For example, in order to compute a quantity of interest for the parametric PDE, high-dimensional function approximation is often required. In this talk, we present the work of interpolating high-dimensional functions using the weighted l1 minimization technique when both function values and their derivative values are sampled. For both deterministic high-dimensional function and parametric PDE examples, with additional derivative samples, we see a numerical improvement over the case with sampling function values only. Theoretically, we show that with exactly the same complexity as in the case of function samples only, the gradient enhanced method gives a better error bound in the Sobolev-type norm as opposed to the L2 norm.

Yi Sui, Ben Adcock
Simon Fraser University
ysui@sfu.ca, ben_adcock@sfu.ca

MS5

Permutation-invariant Interatomic Potentials Based on Ab-initio Data

For many applications ab initio computations are too expensive, so that interatomic potentials, often cheap to compute but less accurate, are in use. Originally derived from empirical models, these potentials have over the past few years mainly been developed from machine-learning methods. In this talk, I will present an interatomic potential combining a data-driven approach, using ab-initio data, with physically motivated functional forms - the many-body expansion. I will notably report our initial experiences with a practical implementation.

Geneviève Dusson
University of Warwick
g.dusson@warwick.ac.uk

MS5

Recent Developments of Quantum Embedding Theories

In the community of quantum physics and chemistry, quantum embedding theories refer to a class of numerical schemes trying to couple together different levels of approximation to the many body Schrodinger equation. These can be viewed as multiscale methods, though there is only one underlying physical scale. In this talk I will provide a brief overview of quantum embedding theories, and discuss some recent mathematical understanding of the projection based embedding theory and its numerical performance (joint work with Leonardo Zepeda).

Lin Lin
University of California, Berkeley
Lawrence Berkeley National Laboratory

linlin@math.berkeley.edu

MS5

Optical Phenomena on Layered Heterostructures

Plasmonic crystals, a class of particularly interesting metamaterials, consist of stacked metallic layers arranged periodically with subwavelength distance, and embedded in a dielectric host. These structures have made it possible to observe aberrant behavior like no refraction, referred to as epsilon-near-zero (ENZ), and negative refraction. This level of control of the path and dispersion of light is of fundamental interest and can lead to exciting applications. In particular, plasmonic metamaterials offer significant flexibility in tuning permittivity or permeability values. In this talk we present analytical and computational results for the simulation of optical phenomena on layered structures. In particular, a homogenization result for layered heterostructures is presented and a corresponding heterogeneous multiscale method is derived. We conclude by hilighting physical implications of our numerical investigation.

Matthias Maier
University of Texas A&M
maier@math.tamu.edu

Mitchell Luskin
School of Mathematics
University of Minnesota
luskin@umn.edu

MS5

Electronic Structure of Incommensurate Materials using Momentum Space

To make the investigation of electronic structure of incommensurate heterostructures computationally tractable, effective alternatives to Bloch theory must be developed. In previous work, we developed and analyzed a real space scheme that exploits spatial ergodicity and nearsightedness. Here we present an analogous scheme formulated in momentum space, which we prove has significant computational advantages in specific incommensurate systems of physical interest, e.g., bilayers of a specified class of materials with small rotation angles. We use our theoretical analysis to obtain estimates for improved rates of convergence with respect to total CPU time for our momentum space method that are confirmed in computational experiments.

Daniel Massatt
University of Minnesota
massa067@umn.edu

MS6

Characterizing Parallel Scalability of Asynchronous Krylov Solvers

As large-scale supercomputers approach exascale performance, algorithms must be designed with less communication overhead in order to scale efficiently in time and energy usage. Among fundamental algorithms to consider are Krylov subspace methods, a class of linear solvers and eigensolvers which are used as black boxes inside numerous scientific computing applications. Researchers have proposed asynchronous Krylov variants (conjugate gradient [Ghysels and Vanroose, 2014], GMRES [Ghysels et al., 2013], BiCGstab [Cools and Vanroose, 2017], etc.) which overlap communication and computation to achieve lower communication overhead at the expense of worse convergence behavior. While researchers have demonstrated promising theoretical and experimental results, there are still open question about how these methods scale in practice, especially under the influence of system noise, node failures, and rounding error. In this talk, we focus on the impact of system noise on the convergence behavior of CG and GMRES. Using existing LogGPS simulation tools [Hoefler et al., 2010], we model the parallel performance of asynchronous Krylov solvers under the impact of realistic HPC system noise.

Kevin DeWeese
University of Washington
Computer Science Department
deweeskg@uw.edu

MS6

Why Worry about Rounding Errors in Hermitian Krylov Subspace Methods?

The conjugate gradient and Lanczos algorithms are widely used methods for solving Hermitian positive definite linear systems, computing eigenvalues of Hermitian matrices, and, more generally, for computing the product of a function of a Hermitian matrix with a given vector. It is well-known that the methods may behave differently in finite precision arithmetic than they would in exact arithmetic, which means that whenever different implementations are considered – such as implementations that make better use of parallelism – one must be careful about the effects on actual performance and accuracy. I will review the assumptions required for the type of error analysis initiated by Paige and others in the 1980's and 1990's and discuss various implementations and whether or not they satisfy these assumptions. For those that do not, I will discuss whether other types of analysis might be more appropriate for explaining their behavior. I believe there are still many things that we do not understand about the behavior of Hermitian Krylov subspace methods in finite precision arithmetic.

Anne Greenbaum
University of Washington
greenbau@uw.edu

MS6

Spectral Clustering of Graph Vertex Subsets via Krylov Subspace Model Reduction

Clustering via graph-Laplacian spectral imbedding is ubiquitous in machine learning. It provides a low-dimensional parametrization of the data manifold which simplifies the subsequent clustering. However, it becomes less efficient for large data sets due to two factors. First, the partial eigendecomposition of the graph-Laplacian typically requires a large Krylov subspace. Second, after the spectral imbedding is done, clustering is performed with conventional algorithms, which lose robustness or become too expensive computationally for large data sets. We propose two novel algorithms for spectral clustering of a subset of the graph vertices (target subset) based on model order reduction. They rely on realizations of a reduced order model (ROM) that accurately approximate the diffusion transfer function of the original graph for inputs and outputs restricted to the target subset. The ROM realization is computed via repeated application of Lanczos algorithm and can be ultimately transformed to a form resembling

a graph-Laplacian. While our focus is limited to the target subset, our algorithms produce its clusterization that is consistent with the overall structure of the graph. Moreover, working with a small target subset reduces greatly the required dimension of Krylov subspace and allows to exploit the conventional clustering algorithms in the regime when they are most robust and efficient, as verified by the numerical experiments with both synthetic and real data.

Alexander V. Mamonov
University of Houston
mamonov@math.uh.edu

Vladimir Druskin
Worcester Polytechnic Institute
vdruskin@gmail.com

Mikhail Zaslavsky
Schlumberger-Doll Research
mzaslavsky@slb.com

MS6

Computing Matrix Functions with the Finite Precision Lanczos Method: New Convergence Bounds and Limitations

We first present a positive result: even though roundoff error destroys orthogonality in the Lanczos process, the finite precision Lanczos method can still effectively approximate almost any matrix function. Our result applies when the working precision exceeds $\log(C)$, where C is a bound on the target function's maximum value. In this case, the the number of iterations required by the Lanczos method can be bounded by the degree of the best uniform approximating polynomial to the function, exactly as can be done in an exact arithmetic analysis. This result generalizes work by Druskin, Greenbaum, Knizhnerman and others, who provide similar bounds for specific matrix functions, including the matrix exponential and inverse. While the above result can be quite strong, for certain functions, including the matrix inverse, the exact arithmetic Lanczos and conjugate gradients methods can converge much faster than predicted by the best uniform approximating polynomial. This is true when eigenvalues are sparse, clustered, or have a few outliers. It is natural to ask how limited precision affects these scenarios. In the second part of the talk, we provide theoretical evidence that improvements are substantially tempered by roundoff error. Specifically, we describe a class of matrices where it is possible to prove (and observe experimentally) that the finite precision conjugate gradient method converges exponentially slower than it does in exact arithmetic.

Christopher Musco
Massachusetts Institute of Technology
cmusco@cs.princeton.edu

MS7

A Novel Compression Algorithm for the Construction of Hierarchically Semi-separable Matrices for Kernel Matrices

We present a novel compression algorithm for the construction of hierarchically semi-separable matrices for kernel matrices. In this context, we show that the new compression algorithm is preferred instead of the classic HSS construction algorithm that uses a partially matrix-free, adaptive randomized projection scheme to determine the maximum off-diagonal block rank. Numerical experiments consider popular datasets from the machine learning field for the task of binary classification.

Gustavo Chavez
Lawrence Berkeley National Laboratory
gichavez@lbl.gov

Rebrova Elizaveta
University of California, Los Angeles
erebrova@umich.edu

Pieter Ghysels
Lawrence Berkeley National Laboratory
Computational Research Division
pghysels@lbl.gov

Sherry Li
Lawrence Berkeley National Laboratory
xsli@lbl.gov

MS7

Fast Approximations of Hessians

Second order optimization methods in machine learning, and inverse problems generally, have shown to be scalable and robust algorithms in practice but suffer from the difficulties of storing and operating on the Hessians that are formally dense. For many Hessians that appear in practical applications, the matrices involved can be well approximated by matrices that have hierarchically low rank structures. Hierarchical matrix representations hold promise to overcome the high complexity of dense representations and provide effective data structures and matrix operations that have only log-linear complexity. In this presentation, we describe algorithms for constructing and updating hierarchical matrix approximations of Hessians and for constructing their inverses. Data parallel versions of these algorithms, appropriate for GPU execution, will also be presented.

George M. Turkiyyah
American University of Beirut
gt02@aub.edu.lb

Wajih Halim Boukaram
KAUST
wajihhalim.boukaram@kaust.edu.sa

David E. Keyes
King Abdullah University of Science and Technology (KAUST)
david.keyes@kaust.edu.sa

MS7

Using Deep Learning for the Solution of Partial Differential Equations

Deep learning is a state-of-the-art technique that is widely used in various fields like computer vision and natural language processing. However, only recently has it been utilized to solve engineering problems. In engineering, many problems involve a low-dimensional input but low-dimensional output (the quantity of interest). The classical example is solving the solution of partial differential equations (PDEs). We present an implementation and experimental results using a novel recurrent neural network (RNN) with customized cells that solve systems described by PDEs. In such systems, given low-dimensional inputs,

we use an RNN with customized cells to directly solve for the low-dimensional output. Benchmarks will be presented for the following three problems:

1. Solve for the highest temperature in a two-dimensional time-dependent heat transfer problem. The error rate achieved, with a smaller number of parameters than long short-term memory (LSTM), is lower than with LSTM and gated recurrent unit (GRU).
2. Find the location where the maximum speed occurs in Burgers equation. Again, our self-designed cell achieves higher precision than LSTM and GRU.
3. We also show that our RNN-based method has the potential to solve the solution on the full grid, although the solution precision can have further improvment.

Ziyi Yang
Stanford University
zy99@stanford.edu

MS7

Kronecker Factorization for Second Order Optimization in Deep Learning

As the size of deep neural networks and the data they consume grow rapidly, it is becoming increasingly important to parallelize their training. Computing on massively parallel computers forces the optimization algorithms for deep learning to operate in a regime they were not originally designed for. Stochastic gradient descent relies on the mini-batches being small in order to generalize well to data it was not trained on. The increase in parallel processes leads to an increase in the batch size and deprives the SGD of the stochasticity necessary for the generalization. In this regime, techniques such as batch normalization also behave very differently, and initialization and normalization of the weights and regularization of the loss become more important than before. We show that when enough care is taken for the initialization, normalization, and regularization, second order optimization methods start to show their inherent advantage over first order methods.

Rio Yokota
Tokyo Institute of Technology
rioyokota@gsic.titech.ac.jp

MS8

Computer Model Calibration and Uncertainty Quantification for Binary Black Hole Formation

Abstract not available.

Derek Bingham
Department of Statistics and Actuarial Science
Simon Fraser University
dbingham@stat.sfu.ca

MS8

Tensorised Approximations for Intractable Probability Distributions

Sampling intractable multivariate probability distributions, e.g., those arising in Bayesian inference, is a fundamental task in statistical learning. We present a method for efficiently representing and sampling multivariate probability distributions based on tensor factorisation. Our approach exploits the intrinsic sparsity and low-rank structures of the target distribution through the discovery of function-form tensor factorisations of the target. In this talk, we will discuss our ongoing research on the structure-preserving tensor factorisation method (such as monotonicity-preserving) which is inevitable for handling probability distributions, as well as preconditioning strategies to accelerate the factorisation. A range of examples in statistical learning, including statistical inverse problems, nonlinear regression, and dynamical systems, are used to demonstrate various aspects of our approach.

Tiangang Cui
Monash University
tiangang.cui@monash.edu

Sergey Dolgov
University of Bath
Department of Mathematical Sciences
s.dolgov@bath.ac.uk

Robert Scheichl
University of Heidelberg
r.scheichl@uni-heidelberg.de

MS8

Derivative Informed MCMC Methods for Subsurface Models with Faults

In this work, we consider Bayesian inversion of fault transmissibility parameters in subsurface flow models. We first construct a low-rank-based approximation of the Hessian at the MAP point that captures the data-informed subspace (in parameter space), and then employ the approximate Hessian as a preconditioner for the preconditioned Crank-Nicolson MCMC method of Stuart et al. to explore the posterior probability distribution. This ensures both parameter-dimension independence as well as exploitation of curvature information in the data-informed directions.

Jeonghun J. Lee
Institute for Computational Engineering and Sciences
The University of Texas at Austin
Jeonghun_Lee@baylor.edu

Omar Ghattas, Tan Bui
The University of Texas at Austin
omar@ices.utexas.edu, tanbui@ices.utexas.edu

Umberto Villa
University of Texas at Austin
uvilla@wustl.edu

MS8

Scalable Bayesian Inference for Inverse Problems

In this talk, we will discuss several algorithms for proper uncertainty quantification in computationally intensive inverse problems. More specifically, we present a class of scalable computational techniques based on a state-of-the-art Markov chain Monte Carlo method, namely, Hamiltonian Monte Carlo. The key idea is to explore and exploit the structure and regularity in parameter space for the underlying probabilistic model to construct an effective approximation of its geometric properties. To this end, we build a surrogate function to approximate the target distribution using properly chosen random bases. We will evaluate our proposed methods based on an elliptic PDE inverse problem, which involves inference of the diffusion coefficient in an elliptic PDE commonly used to model isothermal steady

flow in porous media.

Cheng Zhang
fred hutchinson cancer research center
zc.rabbit@gmail.com

Babak Shahbaba
University of California, Irvine
babaks@uci.edu

Hongkai Zhao
University of California, Irvine
Department of Mathematics
zhao@math.uci.edu

MS9

The Development and use of in Situ Visualization and Analysis Approaches for the U.S. Exascale Program

Changes in the hardware architecture of exascale supercomputers will make current post-processing approaches to analysis and visualization difficult, resulting in disruptive changes to the scientific workflow. A major concern is that exascale system concurrency is expected to grow by five or six orders of magnitude, yet system memory and input/output (I/O) bandwidth/persistent capacity are only expected to grow by one and two orders of magnitude, respectively. A primary goal of our exascale in situ visualization and analysis effort is to overcome bandwidth and storage bottlenecks at the exascale. This will enable analysts and scientists, through the integration of analysis and visualization algorithms with simulations, to create powerful data extracts that can be efficiently processed, stored, and explored. At a high level, the process of in situ data analysis consists of three main objectives: (1) automatic data selection, (2) automatic data reduction that produces data extracts, and (3) storage, processing and presentation of these data extracts. In this talk, I will describe this process in more detail, highlighting data reduction algorithms including topological analysis and statistical sampling and their application in exascale simulation workflows.

James Ahrens
Los Alamos National Laboratory
ahrens@lanl.gov

MS9

GPU Accelerated Feature Extraction on Unstructured Grids using Contour Trees

Turbulent flows exist in many fields of science and occur in a wide range of engineering applications. In the past broad knowledge has been established regarding the statistical properties of turbulence at a range of Reynolds numbers, However, it remains to elucidate various aspects of the small-scale dynamics i.e. the behaviours and interactions of many individual vortical structures. To tackle this problem, we propose a framework that uses state-of-the-art contour tree construction algorithms to identify, classify, track, and visualise vortices in turbulent flow fields produced by large-scale simulations using high-fidelity massively-parallel computational fluid dynamics solvers such as PyFR. The proposed framework is designed from the ground up to exploit multi-GPU systems. In order to enable contour tree construction on grids of linear hexahedral elements a new method to identify and classify saddle points on such meshes has been developed. Since disk capacity and I/O have become a bottleneck for large-scale simulations, the proposed framework will be applied *in-situ*, while relevant data is still in device memory.

Marius K. Koch
Imperial College
marius.koch16@imperial.ac.uk

Paul Kelly
Imperial College London
p.kelly@imperial.ac.uk

Peter E. Vincent
Department of Aeronautics
Imperial College London
p.vincent@imperial.ac.uk

MS9

In-transit Data Analytics for Large-scale CFD

The term 'in situ' is broad and encompasses many strategies for computing data products from the simulation as it runs. Typically, in situ means that data products are generated in the same address space as the simulation by operating directly on simulation data structures while the calculation is paused. 'In transit' techniques are a variation on in situ; these methods permit the simulation to quickly send its raw data to additional compute resources, usually another set of analysis nodes on the HPC system which perform reductions to create the desired data products. In transit has the benefit that the simulation writes its data and then continues its computations, while the analytics code executes at the same time on additional HPC nodes. This design is robust and has the benefit that the simulation is free to continue its computations after sending its data to the analysis nodes. Images or movies have been the most common data products of these processes, enabling post-hoc visualization. Intelligent Light advocates the use of 3D 'extracts' as data products so that post-hoc analysis can be more flexible and numerical as well as visual. Over the past few years, we have been developing in transit reduced order surrogate models via POD and dynamic mode decomposition (DMD) in the extract process, enabling greater compression and post-hoc analysis based on the eigen basis functions of the underlying dynamical system, via signatures of complex unsteady CFD computations.

Steve Legensky
Intelligent Light
sml@ilight.com

MS9

In Situ Analysis and Visualization Overview

Given the time and the energy costs for data movement and the limited storage available on current generation high-performance computing (HPC) systems, the notion of performing analysis and visualization while data remains in memory has become an alternative to the traditional ad-hoc post-processing of simulation results. The concept of in situ analysis and visualization has a lengthy history covering decades of research and development work. At present, there exist production-ready in situ infrastructures, such as the open source ParaView Catalyst, that integrate seamlessly with state-of-the-art simulators. This infrastructure can be configured through editions to be extremely lightweight with small library sizes and memory footprints to ensure the most efficient use of HPC resources. ParaView Catalyst has been used to generate computa-

tional fluid dynamics in situ results on a 1 million MPI rank run on Argonnes Mira BG/Q.

Patrick O'Leary
Kitware, Inc., USA
patrick.oleary@kitware.com

MS10

Energy Norm Regularized Sparse Simultaneous Reconstruction of Solutions to Parameterized PDEs

We present and analyze a novel sparse polynomial approximation method for the solution of PDEs with stochastic and parametric inputs. Our approach, denoted *simultaneous compressed sensing*, treats the parameterized problem as a problem of sparse signal reconstruction. Using a reformulation of the standard basis pursuit denoising problem, involving a new mixed ℓ_1 and energy norm-based method of regularization, we are able to achieve global reconstruction of solutions to parameterized elliptic PDEs over both physical and parametric domains. In addition, we are able to show that, with minimal sample complexity, error estimates comparable to quasi-optimal best s-term approximations [Tran, Webster, Zhang, 2017] are achievable, while requiring only *a priori* bounds on truncation errors in energy norms. Finally, we present extensive numerical experiments on several high-dimensional parameterized elliptic PDE models to demonstrate the superior recovery properties of the proposed approach.

Nick Dexter
University of Tennessee
dexter.nicholas@gmail.com

Hoang A. Tran
Oak Ridge National Laboratory
Computer Science and Mathematics Division
tranha@ornl.gov

Clayton G. Webster
University of Tennessee and Oak Ridge National Laboratory
webstercg@ornl.gov

MS10

A Multi-fidelity Spectral Technique for Uncertainty Quantification

The use of model reduction has become widespread as a means to reduce the computational cost for uncertainty quantification of PDE systems. In this work we presenta model reduction technique that exploits the low-rank structure of the solution of interest, when exists, for fast propagation of high-dimensional uncertainties via spectral techniques. To construct this low-rank approximation, the proposed method utilizes models with lower fidelities (hence cheaper to simulate) than the intended high-fidelity model. Given realizations to the lower fidelity models, a stochastic reduced basis is identified and utilized to generate a spectral representation of the solution using a small set of high-fidelity realizations. In addition to the construction of this bi-fidelity approximation, we present convergence analysis and numerical results showcasing its efficacy on high-dimensional problems.

Alireza Doostan
Department of Aerospace Engineering Sciences
University of Colorado, Boulder
Alireza.Doostan@Colorado.EDU

Felix Newberry
University of Colorado Boulder
felix.newberry@colorado.edu

MS10

Adaptive Multi-index Collication for Balancing Statistical and Deterministic Errors

The computational cost of many high-fidelity models prohibit the use of Monte Carlo sampling and even many surrogate based methods for quantifying uncertainty. In this talk we present an adaptive multi-index approach that leverage a convergent hierarchy of models to build approximations of the parameter-output map of high-fidelity predictions. Our work is an extension of multi-index stochastic collocation (MISC) which, unlike traditional multi-level methods, can be applied to models with multiple hyperparameters controlling the model discretization, such as the number of finite elements and the number of time steps. The adaptive algorithm we present automatically balances stochastic approximation and physical discretization errors, whilst minimizing cost, and identifies the parameter combinations that significantly influence the model uncertainty. We demonstrate the efficacy of our algorithm on a state-of-the art aerospace model.

John D. Jakeman
Sandia National Labs
jdjakem@sandia.gov

Michael S. Eldred
Sandia National Laboratories
Optimization and Uncertainty Quantification Dept.
mseldre@sandia.gov

Gianluca Geraci
Sandia National Laboratories, NM
ggeraci@sandia.gov

Alex Gorodetsky
University of Michigan
goroda@umich.edu

MS10

Sparse Grid Approximation of Elliptic PDEs with Lognormal Diffusion Coefficient

This talk is concerned with efficient sparse grid approximations of u, solution of an elliptic PDE whose diffusion coefficient is modeled as a lognormal random field. In the first part, we build upon previous works available in the literature to establish a convergence result (in L^2 norm in probability) for the approximation of u by sparse collocation with Gauss-Hermite points. More specifically, we first link the error to the size of the multi-index set defining the sparse collocation and then derive a bound on the number of points in the associated sparse grid. The result of the analysis is an algebraic convergence rate of the approximation error with respect to both the size of the multi-index set and the number of points in the sparse grid; interestingly, the analysis gives also an explicit a-priori estimate of the optimal multi-index set. We validate the results by numerical tests in which we consider a family of random fields parameterized by a coefficient that sets the spatial smoothness of the field. As expected, the convergence rate for very rough fields turns out to be quite slow, even for

optimized grids (be it the above-mentioned a-priori grids or the classical a-posteriori-adaptive grids). Thus, in the second part of the talk, we propose a remedy based on using the solution of the PDE on a smoothed version of the random field as control variate for a Monte Carlo sampling of u.

Lorenzo Tamellini
Istituto di Matematica Applicata e Tecnologie Informatiche
"E. Magenes" (CNR-IMATI)
tamellini@imati.cnr.it

MS11

Towards Scalable Scientific Machine Learning: Motivation and Approaches

There is a growing amount of data generated from scientific instruments and computational simulations. The volume of data has driven an interest in scientific machine learning as an automated way to determine relationships that could otherwise only be discovered through painstaking data analysis. This talk will present several applications that motivate some of the mathematical and algorithmic challenges in scientific machine learning. In addition, focusing on deep neural networks, we will explore the scalability of recent approaches and their applicability to problems of interest to the scientific enterprise.

Eric C. Cyr
Computational Mathematics Department
Sandia National Laboratotories
eccyr@sandia.gov

MS11

Differential Equations with Unknown Constitutive Relations as Recurrent Neural Networks

We describe a data-driven approach to differential equations of known structure but with unknown constituent functions, exploiting formal connections between discrete dynamical systems and LSTM-style recurrent neural networks. Our method learns the unknown functions from discretized time-series data, using a loss function that takes into account the persistent consequences of errors. The method supports training in missing-data scenarios, and the resulting learned dynamics are discretization-independent. The method is implemented as an extension of Tensorflow's recurrent neural network architecture.

Tobias Hagge, Panos Stinis
Pacific Northwest National Laboratory
tobias.hagge@pnnl.gov, panagiotis.stinis@pnnl.gov

Enoch Yeung
University of California Santa Barbara
eyeung@ucsb.edu

Alexander Tartakovsky
Pacific Northwest National Laboratory
alexandre.tartakovsky@pnnl.gov

MS11

Energy Minimizing Deep Neural Networks for Solving High-dimensional Stochastic Partial Differential Equations

Stochastic partial differential equations (SPDEs) have a ubiquitous presence in engineering and computational science. Stochasticity in partial differential equation (PDE) arises from unknown/random boundary/initial conditions or field parameters (e.g., the permeability of the ground in flow through porous media, the thermal conductivity in heat transfer) and, thus, it is inherently high-dimensional. In this regime, traditional uncertainty propagation techniques fail to learn the high-dimensional response surfaces because of the curse of dimensionality. The viable alternative is Monte Carlo, but this method needs a large number of realizations for solution statistics convergence hence incurring excessive computational burden. In this work, we focus on elliptic SPDEs. To overcome the curse of dimensionality, we represent the solution of the SPDE as a linear combination of the basis functions estimated from spatial/time inputs, whereas the coefficients are estimated using a deep neural network (DNN) with stochastic inputs. However, contrary to the plethora of existing methodologies, we do not rely on observations of input-output examples from an external PDE solver to train the network. Our DNN is trained by minimizing a physics-informed loss function which is based on the energy functional of the PDE. We demonstrate our approach by solving steady-state heat equation with 100-dimensional random conductivity.

Sharmila Karumuri, Ilias Bilionis
Purdue University
skarumur@purdue.edu, ebilionis@gmail.com

Jitesh Panchal
School of Mechanical Engineering
Purdue University, West Lafayette, Indiana, 47907
panchal@purdue.edu

MS11

Data-driven Approximation of High-dimensional Flow Maps and Probability Distributions using Deep Neural Networks

Efficient approximation of the evolution of probability density functions (PDFs) on the state of high-dimensional nonlinear dynamical systems remains a central challenge in uncertain optimal control. Existing methods rely on trajectory sampling to evaluate performance metrics, but meaningful evaluations require many trajectories which greatly increases the dimension of the optimal control problem. The problem rapidly becomes computationally intractable, and thus these methods are limited to only a few uncertain variables. We propose a new approach to PDF and flow map approximation based on the partial differential equations (PDEs) governing the evolution of the PDF and flow map. In particular, using small sets of sample trajectories in addition the governing PDEs, we train deep neural networks to approximate the time dependent PDF and flow map of the system. We further demonstrate that these methods scale well to high-dimensional systems, laying the foundation for new possibilities in uncertainty quantification and uncertain optimal control.

Tenavi Nakamura-Zimmerer
University of California, Santa Cruz
tenakamu@ucsc.edu

Daniele Venturi
Department of Applied Mathematics and Statistics
University of California Santa Cruz
venturi@ucsc.edu

Qi Gong

University of California, Santa Cruz
qigong@soe.ucsc.edu

MS12

Abstraction Layer for Standardizing APIs of Task-based Engines

Task-based programming models have become ubiquitous in scientific applications and numerical libraries (e.g., DPLASMA, MUMPS, HiCMA, Chameleon, etc.). There exist a myriad of dynamic runtime systems to support the task-based programming model on shared and distributed memory systems, including hardware accelerators: OpenMP, StarPU, Quark, OmpSs, PaRSEC, ParalleX, etc. The lack of API standardization makes it difficult for task-based developers to switch between various runtimes. This hinders user productivity and also discourages developers from assessing the performance of specific features of a given runtime system. AL4SAN is a standalone software library that provides a collection of APIs to unify the expression of tasks and their data dependencies from some of the aforementioned task-based runtimes. AL4SAN APIs permit to code applications once for deployment on various underlying systems, including shared and distributed-memory manycore systems. AL4SAN has front-end APIs that are exposed to the end-users, and back-end APIs that support different task-based runtimes. AL4SAN is integrated to two linear algebra software packages for performance assessment: Chameleon and HiCMA, supporting compute-bound and memory-bound workloads, respectively. The overhead of the AL4SAN APIs is negligible for typical scientific workloads. This integration has been tested on shared and distributed-memory manycore systems, as well as systems equipped with GPU hardware accelerators.

Rabab Omairy
KAUST
rabab.omairy@kaust.edu.sa

Hatem Ltaief
King Abdullah University of Science & Technology (KAUST)
hatem.Ltaief@kaust.edu.sa

David E. Keyes
King Abdullah University of Science and Technology (KAUST)
david.keyes@kaust.edu.sa

MS12

Automatic Task-based Parallelization of Python Codes

Dealing with distributed computing issues when developing applications is a remarkable obstacle for many scientists. This talk introduces AutoParallel, a module based on sequential programming to ease the development of distributed Python applications. This talk also discusses several linear algebra applications regarding both, performance and code complexity. Python has risen to the top of the list of the programming languages for scientific applications thanks to the simplicity of its syntax, while still achieving a good performance even being an interpreted language. AutoParallel is a Python module built on top of PyCOMPSs and PLUTO to automatically find an appropriate task-based parallelization of affine loop nests and execute them in parallel in distributed computing infrastructures. It only contains a small annotation in the form of Python decorators so that anyone with little programming skills can scale up an application to hundreds of cores. Moreover, the parallelization can also include the building of data blocks to increase tasks granularity in order to achieve a good execution performance.

Cristian Ramon-Cortes
Barcelona Supercomputing Center (BSC)
cristian.ramoncortes@bsc.es

Ramon Amela, Jorge Ejarque
Barcelona Supercomputing Center
ramon.amela@bsc.es, jorge.ejarque@bsc.es

Philippe Clauss
INRIA
ICube Lab. - Université de Strasbourg
philippe.clauss@inria.fr

Rosa M. Badia
Barcelona Supercomputing Center / CSIC
rosa.m.badia@bsc.es

MS12

Task-based Programming for Scientific Computing: Runtime Supports

Abstract not available.

Thomas Sterling
Indiana University
tron@indiana.edu

MS12

Task-based Programming, from Sequential-looking Source Code to Exascale: What a Runtime System Can Help You With

Task-based programming has gained more and more attention in the past decade. Thanks to adapted programming interfaces, it is possible to write source code which looks sequential and can indeed be run and debugged in sequential order, but which also safely expresses parallelism through a task graph. A runtime system can then be used to automatically execute it on a wide range of systems, including accelerators, disks for out-of-core, and even over the network, using state-of-the art optimization heuristics. We will present the broad range of such features that runtime can provide seamlessly, how the programmer can safely provide some hints to improve performance even further, and discuss how well this can potentially scale to very large clusters.

Samuel Thibault
University of Bordeaux
samuel.thibault@labri.fr

MS14

On Conservation of Dissipative Relations in Model Order and Complexity-reduced Systems

The development of discretization methods for differential equations, which take into account fundamental underlying physical principles, has been an active field of research for quite a few years. For model-order and complexity-reduction structure preservation is a rather open field of research. In this contribution a variational framework for the construction of online-efficient reduced models will

be presented, which guarantees the preservation of desired dissipative relations, or respective discrete counterparts. It takes into account all approximation steps, which are space- and time-discretization, projection-based model-order-reduction, and complexity-reduction of the nonlinearities to avoid the so called lifting bottleneck. The methodology is widely applicable and general in its nature as will be demonstrated.

Bjoern Liljegren-Sailer
Faculty IV- Mathematics
Trier University
bjoern.sailer@uni-trier.de

Nicole Marheineke
Trier University
marheineke@uni-trier.de

MS14

Structure Preserving Reduced Order Models for Hamiltonian Problems with a Non-linear Poisson Structure

We develop structure preserving reduced basis methods for a large class of problems by resorting to their semi-discrete formulation as Hamiltonian dynamical systems. The Poisson manifold structure that characterizes the phase space of Hamiltonian problems encodes the physical properties, symmetries and conservation laws of the dynamics. Failing to preserve such structure in performing a model order reduction may result in the violation of the conservation laws and in the onset of spurious artifacts and instabilities. We design structure preserving reduced basis methods for the general case of nonlinear state-dependent degenerate Poisson structures based on a two-step approach. First, via a local approximation of the Poisson tensor we split the Hamiltonian dynamics into an 'almost' symplectic component and the trivial evolution of the Casimir invariants resulting from the degeneracy of the Poisson structure. Second, we apply canonically symplectic reduced basis techniques only to the nontrivial component of the dynamics. With this approach the global Poisson structure and the conservation properties of the phase flow are retained by the reduced model up to errors in the Poisson tensor approximation. The proposed reduction scheme is combined with a discrete empirical interpolation method (DEIM) to deal with nonlinear Hamiltonian functionals and ensure a computationally competitive reduced model.

Cecilia Pagliantini
Computational Mathematics and Simulation Science (MCSS)
Ecole Polytechnique Fédérale de Lausanne (EPFL)
cecilia.pagliantini@epfl.ch

Jan S. Hesthaven
EPFL
jan.hesthaven@epfl.ch

MS14

Novel Structure Preserving Model Reduction Schemes

In this talk, we discuss a novel model reduction framework for generalized linear parametric systems, with transfer function given by

$$\mathbf{H}(s,\mathbf{p}) = \mathcal{C}(s,\mathbf{p})\mathcal{K}(s,\mathbf{p})^{-1}\mathcal{B}(s,\mathbf{p}).$$

Such systems arise, *e.g.*, in the context of second-order linear systems and delay systems, and they may also have parameter dependencies. Our main goal is to construct reduced-order models for this class of systems. To this aim, we first characterize reachability and observability for such structured parametric systems. Thus, we can construct the numerical minimal realization by projection. Further, we extend the connection between interpolation-based and Loewner approaches, known for standard linear systems to the considered class of generalized linear parametric systems, which allows us to construct reduced-order models. Additionally, we extend balancing-type model reduction schemes for this class of systems. Furthermore, special attention is given to computational aspects of the approaches, and we discuss their applicability to large-scale problems. We illustrate the efficiency of the proposed approach with several numerical benchmark examples.

Igor Pontes Duff
Max Planck Institute for Dynamics of Complex Technical Syste
pontes@mpi-magdeburg.mpg.de

Pawan Goyal
Max Planck Institute Magdeburg for
Dynamics of Complex Technical Systems
goyalp@mpi-magdeburg.mpg.de

Peter Benner
Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany
benner@mpi-magdeburg.mpg.de

MS14

Perspective on Structure Preserving Model Order Reduction

Naive application of the conventional model order reduction (MOR) techniques can destroy the dynamical structure (for instance, Hamiltonian structure, stability, passivity, input-output behavior, switch threshold, delay, etc.). Scalable novel and prodigious computational tools are thus needed to mitigate the above challenge, while alleviating numerical burden and increasing simulation speed at quantifiably minimal error accrual. Until a decade ago, MOR techniques were rarely designed with the structure preservation in mind. Structure preserving MOR originated in the electronics industry and has now found a place in several other industrial settings. Problems of practical interest, such as power grid systems and gas networks, constitute ever more large scale and complex dynamical systems, and demand coupling different physical domains (structurally separated simulation environments) across many scales. This calls for the preservation of network structure of coupled systems, energy exchange across subsystems, preservation of power maps, conservation of dissipative relations, causality, feedback structure, synchronization property to name a few. We introduce the notion of structure preservation or mimetics in discretization and MOR, discuss the recent advances in structure preserving MOR and demonstrate its benefits via industrial case studies, encompassing the problems from the field of electronic circuit simulation, structural and fluid dynamics, systems and control, etc.

Wil Schilders
Eindhoven University of Technology

w.h.a.schilders@TUE.nl

MS15

Redemption of the High Frequencies: Scalable Algorithms for 3D Helmholtz

There is much truth to the conventional wisdom that computational wave propagation is harder when the frequency is higher. Ten years ago, it was unclear that scalable sequential algorithms could even exist for the Helmholtz equation. Today, linear complexity is not only available in many scenarios of interest, but it is becoming clear that parallelism can take us much further. I will show recent results that indicate that genuinely sublinear parallel runtimes are possible in the 3D case, both with respect to the total number of unknowns, and the number of right-hand sides. The ideas that enable such scalings do not seem to be available in the low-frequency regime, hence the title of the talk. Joint work with Matthias Taus (MIT), Leonardo Zepeda-Nunez (Berkeley), Russell Hewett (Total), Adrien Scheuer (UCL).

Laurent Demanet
Department of Mathematics, MIT
laurent@math.mit.edu

Matthias Taus
University of Texas
matthias.taus@gmail.com

Leonardo Zepeda-Núñez
Massachusetts Institute of Technology
lzepeda@math.mit.edu

Russell Hewett
Total
russell.hewett@total.com

Adrien Scheuer
Université catholique de Louvain
adrien.scheuer@uclouvain.be

MS15

Fast Huygens Sweeping (FHS) Methods for Highly Oscillatory Phenomena

In the first part of the talk, we discuss fast Huygens sweeping methods for Schrondinger equations in the semi-classical regime by incorporating short-time Wentzel-Kramers-Brillouin-Jeffreys (WKBJ) propagators into Huygens principle. Even though the WKBJ solution is valid only for a short time period due to the occurrence of caustics, Huygens principle allows us to construct the global-in-time semi-classical solution. Then, we will also present a more recent algorithm for simulating the multi-color optical self-focusing phenomena in nematic liquid crystals. This is a joint work with Wingfai Kwan, Jianliang Qian, Susana Serna and Xiaoping Wang.

Shingyu Leung
Hong Kong University of Science and Technology
masyleung@ust.hk

MS15

Hadamard-Babich Ansatz for Point-source Maxwell's Equations in Inhomogeneous Media

Starting from Hadamard's method, we extend Babich's ansatz to the frequency-domain point-source (FD Maxwell's equations in an inhomogeneous medium in high-frequency regime. First, we develop a novel asy totic series, dubbed Hadamard's ansatz, to form the damental solution of the Cauchy problem for the ti domain point-source (TDPS) Maxwell's equations in region close to the source. Governing equations for the knowns in Hadamard's ansatz are then derived. In o to derive the initial data for the unknowns in the ans we further propose a condition for matching Hadama ansatz with the homogeneous-medium fundamental s tion at the source. Directly taking the Fourier tr form of Hadamard's ansatz in time, we obtain a ansatz, dubbed the Hadamard-Babich ansatz, for FDPS Maxwell's equations. Next, we elucidate the r tion between the Hadamard-Babich ansatz and a rece proposed Babich-like ansatz for solving the same FI Maxwell's equations. Finally, incorporating the first terms of the Hadamard-Babich ansatz into a planar-ba Huygens sweeping algorithm, we solve the FDPS Maxw equations at high frequencies in the region where caus occur. Numerical experiments demonstrate the accu of our method.

Jianliang Qian
Department of Mathematics
Michigan State University
qian@math.msu.edu

MS15

Parallel-in-time Multiscale Algorithm for W Equations

We propose a new general framework for building r tiscale time-domain decomposition methods. The fra work is inspired by the parareal schemes introduced by ons, Maday and Turinici for parallel-in-time computa of evolutionary problems. In this talk, we shall desc our new framework which rely on computed data to bu on-the-fly, effective propagators. We will demonstrate algorithms that we developed for solving second order v equations.

Yen-Hsi Richard Tsai
Department of Mathematics
The University of Texas at Austin
ytsai@math.utexas.edu

MS16

Inverse Problems: Integrating Data with Mod under Uncertainty

This tutorial aims to introduce methods that facili the integration of data with complex physics based m els to quantify and reduce uncertainties in model pre tions. In particular, we will discuss mathematical and c putational aspects of inverse problems governed by tial differential equations (PDEs). Topics discussed clude ill-posedness and regularization, first and second der adjoints, inexact matrix-free Newton-CG for nonli least squares optimization, and connections to linear G sian Bayesian inverse problems. We will also discuss demonstrate implementation aspects by using hIPPY an extensible software framework for large-scale inv problems developed at The University of Texas at Au and the University of California, Merced. This framew implements state-of-the-art scalable algorithms for inv problems governed by PDEs, and provides a unique pability for algorithmic developments for large-scale

terministic and Bayesian) inversion. We will show how to implement various features that commonly arise in typical inverse problems, such as inversion for the initial condition in a one-dimensional heat equation, and for the coefficient field in an elliptic PDE.

Noemi Petra
University of California, Merced
npetra@ucmerced.edu

MS16

Global Sensitivity Measures: Theory to Practice

Stochastic or uncertain inputs to a model are a commonly observed in computational science and engineering. A strategy for investigating the impact of uncertain inputs on the output is therefore of remarkable interest. Global sensitivity measures (GSMs) provide an effective means for quantifying the relative contribution of such model inputs to the variability in its output. Estimating the measures could however be computationally challenging depending upon the computational effort associated with the simulation. This tutorial will introduce the commonly used GSMs, highlight connections between them, and demonstrate efficient techniques for estimating them. Additionally, the application of GSMs for dimension reduction will be demonstrated using simple illustrative examples as well as relatively complex research problems.

Manav Vohra
Vanderbilt University
manav.vohra@vanderbilt.edu

MS17

Relaxation Algorithms for Nonsmooth Phase Retrieval

Classic inverse problems are formulated using smooth penalties and regularizations. However, nonsmooth and nonconvex penalties/regularizers have proved to be extremely useful in underdetermined and noisy settings. Problems with these features also arise naturally when modeling complex physical phenomena, such as in phase-retrieval and related applications. We propose a 'relax-and-split' technique for solving a broad range of nonsmooth, nonconvex problems. The technique can be implemented on a range of problems in a simple and scalable way. In particular, we typically need only solve least squares problems, as well as implement custom separable operators. We discuss the problem class, reformulation and algorithms, and give examples of promising numerical results.

Aleksandr Aravkin, Peng Zheng
University of Washington
saravkin@uw.edu, zhengp@uw.edu

MS17

Joint Ptycho-tomography Reconstruction Through Alternating Direction Method of Multipliers

We present the extension of ptychograhpy for three-dimensional object reconstruction in a tomography setting. In this talk, we describe a general inversion framework for recovering an object from ptychographic measurements. We use the alternating direction method of multipliers (ADMM) algorithm to efficiently solve the nonlinear optimization problem. We demonstrate the performance of our algorithm via simulation study.

Selin Aslan, Doga Gürsoy
Argonne National Laboratory
ssariaydin@anl.gov, dgursoy@aps.anl.gov

Sven Leyffer
Argonne National Laboratory
leyffer@mcs.anl.gov

MS17

Hyperspectral Ptychography and Beyond

Ptychographic phase retrieval is a popular imaging technique to study materials at high resolution. By changing energy of the x-ray photons across resonant absorption energies it is possible to differentiate between different chemical or magnetic components. Here we describe the phase retrieval problem across energies with known and learned spectral contrast.

Huibin Chang
School of Mathematical Sciences, Tianjin Normal University
huibinchang@lbl.gov

Pablo Enfedaque, Hari krishnan
Lawrence Berkeley National Laboratory
pablo.enfedaque@gmail.com, hkrishnan@lbl.gov

Stefano Marchesini
Lawrence Berkeley National Lab
smarchesini@lbl.gov

MS17

New Methods for Partially Coherent Ptychography

As a multi-measurement scanning coherent diffraction imaging (CDI) method, X-ray ptychography is aimed to realize high spatial-resolution X-ray imaging beyond the resolution limit set by X-ray optics. A fundamental issue is that the high spatial resolution requires that the sample be illuminated by a high dose scale which is approximately inverse fourth power of the resolution, thus making higher fluxes very desirable to push the resolving power of CDI techniques. At the same time, the high fluxes available from synchrotron sources enable theoretical probe scan speeds that exceed what can be achieved due to mechanical limitations in scanning optics. Thus, the only way to achieve high scan speeds with an effectively stationary probe is to reduce the flux by avoiding detector integration while the probe is in motion, dramatically reducing the total flux acquired. In this talk, we introduce two new algorithms that are tailored to compensate for incoherent blur in ptychography acquisition systems. The first method compensates for use of a hard X-ray broadband source (DMM) in experiments performed at Argonne APS. The method works by sampling the broadband radiation source to a set of probe modes with discrete wavelengths, then the recorded diffraction intensity is the incoherent summation of diffraction intensities for each wavelength. The second method utilizes Algorithmic Differentiation to compensate for the incoherent blur caused by "fly-scan" mode of operation.

Oliver Cossairt, Yudong Yao
Northwestern University

ollie@eecs.northwestern.edu, yaoyud1990@gmail.com

MS18

Conservation Laws in Diffusive Chemical Reaction Networks for Semiconductor Defects

Modern semiconductor devices rely upon a variety of dopants to achieve the best possible performance. Traditional simulations treat these dopants as static background doping or stationary ionizing traps, neglecting both the slow movement of defects over time and chemical reactions between defects. However, technologies such as Cadmium Telluride exhibit metastability and accurate simulations require careful tracking of defect profiles. These defect profiles vary over many orders of magnitude, ranging from trace defects to the density of the lattice itself and accurate simulation requires careful numerical techniques to control conservation of mass and charge. In this talk we present a model case for fabrication of CdTe Solar Cells with active defects. We implement a chemical reaction network using mass-action kinetics and demonstrate a traditional projection method for ensuring conservation of mass and charge. Improvements are made by extending the conserved quantities to the complete set of invariants and by altering the projection direction to better correct the errors introduced by the numerical scheme.
This work is part of a unified II-VI semiconductor simulation in joint work with Arizona State (D. Vasileska, A. Shaik, C. Ringhofer) and First Solar (I. Sankin, D. Krasikov) and supported by the U.S. Department of Energys Office of Energy Efficiency and Renewable Energy (EERE) under Solar Energy Technologies Office (SETO) Agreement Number DE-EE0007536.

Daniel Brinkman
Department of Mathematics and Statistics
San Jose State University
daniel.brinkman@sjsu.edu

Christian Ringhofer
Arizona State University
ringhofer@asu.edu

MS18

High Order Sparse Grid Discontinuous Galerkin Methods for Nonlinear PDEs

This talk will present some recent progress on the development of high order sparse grid discontinuous Galerkin (DG) methods. The method resorts to a novel collocation approach, and can efficiently compute nonlinear PDE problems. The main advantage of the scheme lies in the reduced degrees of freedom for high dimensional problems, which is applicable to kinetic simulations.

Yingda Cheng
Department of Mathematics
Michigan State University
ycheng@math.msu.edu

MS18

Highly Accurate Quadrature-based Scharfetter-Gummel Schemes for Charge Transport in Degenerate Semiconductors

We introduce a family of highly accurate two-point flux approximations. They describe drift-diffusion charge transport in degenerate semiconductors and can be used in Voronoi finite volume discretizations. Unlike, the classical Scharfetter-Gummel scheme they also work well for nonlinear diffusion (non-Boltzmann statistics). This is particularly relevant in organic semiconductors or for semiconductors operating at very low temperatures. In this case, the numerical flux is given implicitly as the solution of a nonlinear integral equation. We solve it via quadrature and Newton's method. A convergence study reveals that the solution of the approximate integral equation converges exponentially in terms of the number of quadrature points. With very few integration nodes our method is already more accurate than a state-of-the-art reference flux, especially in the challenging physical scenario of highly nonlinear diffusion. Finally, we show that thermodynamic consistency is practically guaranteed.

Patricio Farrell
Weierstrass Institute for Applied Analysis and Stochastics
Weierstrass Institute
patricio.farrell@tuhh.de

MS18

Modeling Negative Capacitance and Inductive Loop in Perovskite Solar Cells

Negative capacitance and inductive loops in impedance spectroscopy at low or intermediate frequencies in perovskite solar cells have been described in several recent reports. The origin of these observations has, however, remained unknown so far and is currently under debate. In this contribution, we shed light on the physical mechanisms behind these two features. For this purpose, we employ a 1D model that includes electronic as well as ionic transport. The model is able to simulate current-voltage curves, transient responses and impedance spectra and naturally produces inductive loops and negative capacitance. Further, we discuss the negative capacitance occurrence in other devices such as organic light- emitting diode (OLEDs) and investigate its origin by including a thermal model.

Evelyne Knapp
ZHAW Zurich University of Applied Sciences
Organic Electronics & Photovoltaics
evelyne.knapp@zhaw.ch

Martin Neukom
Fluxim AG
martin.neukom@fluxim.com

Beat Ruhstaller
ZHAW Zurich University of Applied Sciences
Fluxim AG
beat.ruhstaller@zhaw.ch

MS19

Adaptive Product-convolution Approximation for Hessians, Interface Schur Complements, and other Locally Translation-invariant Operators

Hessians in PDE constrained inverse problems and interface Schur complements often exhibit local translation-invariance. In this talk we present an adaptive product-convolution approximation scheme for operators that are locally translation-invariant, even if these operators are high-rank. Unlike existing methods that do not exploit translation invariance, our scheme performs well regardless of the Peclet number when applied to the Hessian in an advection-diffusion inverse problem, and regardless of the mesh size when applied to an interface Schur comple-

ment for the Poisson operator.

Nick Alger
The University of Texas at Austin
Center for Computational Geosciences and Optimization
nalger225@gmail.com

Vishwas Rao
Argonne National Laboratory
vhebbur@anl.gov

Aaron Myers
University of Texas at Austin - ICES
aaron@ices.utexas.edu

Tan Bui-Thanh, Omar Ghattas
The University of Texas at Austin
tanbui@ices.utexas.edu, omar@ices.utexas.edu

MS19

Reduced Order Modeling for Time-dependent Optimal Control Problems with Variable Initial Values

A new ROM Hessian approximation for large-scale optimal control problems with initial value control is presented. Such problems arise in parameter identification of initial data, or in multiple shooting formulations of more general unsteady optimal control problems. No fixed ROM well-approximates the application of the Hessian to all possible initial data. The new approach selects a basic ROM and augments it by one vector to produce a substantially better approximation. This is joint work with Dr. Doerte Jando.

Matthias Heinkenschloss
Department of Computational and Applied Mathematics
Rice University
heinken@rice.edu

MS19

Accelerating Reconstruction Algorithms for Fluorescence Optical Tomography

Fluorescence optical tomography is a non-destructive imaging technique that aims at recovering the distribution of a fluorophore concentration inside an object by boundary measurements of fluorescent light for multiple illuminations. Due to the huge amount of measurement data and the huge number of degrees of freedom that are necessary to represent the fluorophore concentration with good resolution, fluorescence optical tomography amounts to a large-scale inverse problem. The goal of this talk is to present strategies that effectively reduce the computational costs for solving this inverse problem while maintaining a sufficiently high resolution.

Matthias Schlottbom
University of Twente
m.schlottbom@utwente.nl

Herbert Egger
Numerical Analysis and Scientific Computing
Darmstadt University of Technology
egger@mathematik.tu-darmstadt.de

MS19

Automatic Determination of Regularization Parameters in Krylov Methods for Large Scale Inverse Problems

In many inverse problems the regularization parameter is set ad hoc or as an outer loop to fit the discrepancy principle. For large scale problems, this practice is very expensive, since very similar large scale simulations are repeated for slightly different regularization parameters and very little information from previous simulations is reused. Leading to a long time to solution. In this talk we review several techniques to automatically set the regularization parameter, as part of the solver. Techniques such as Generalized Arnoldi Tikhonov or Regula Falsi based automatic regularization adapt the regularization parameter during the convergence of Krylov methods based on the current guess. Leading to a speed up. Similarly, it is possible to add the discrepancy principle to the system of equations and make the regularization parameter an unknown, just like the unknowns of the problem. The system is then solved simultaneously for the solution and the correct regularization parameter. This system is solved by a Krylov method with the Krylov method as an outer iteration and an inner iteration to determine the best regularization parameter for the current Krylov subspace. The methods are illustrated with examples from inverse scattering and model calibration.

Wim I. Vanroose
Universiteit Antwerpen
Belgium
wim.vanroose@ua.ac.be

Nick Schenkels
U. Antwerpen
Belgium
nick.schenkels@uantwerpen.be

MS20

Hermite Spectral Method for Boltzmann Equation

We propose a Hermite spectral method for the Boltzmann equation with three-dimensional velocity space. For the inverse-power-law model, we introduce an approximate quadratic collision model for the spatially homogeneous case, and then generalize it to the spatially inhomogeneous case. Different methods are proposed for low Mach number and high Mach number problems. The numerical method shows high efficiency especially for computing the first few moments of the distribution function.

Zhenning Cai
National University of Singapore
matcz@nus.edu.sg

Zhicheng Hu
Nanjing University of Aeronautics and Astronautics
huzhicheng@nuaa.edu.cn

Yanli Wang
Peking University
wang_yanli@pku.edu.cn

MS20

Numerical Study of the Maximum Entropy Moment System

Maximum entropy principle has attracted a lot of attention since it was proposed. Its good properties, including globally hyperbolic, preserving positivity of the distribu-

tion and infinite characteristic speed, made the researchers believe that the maximum entropy equations (MEE) would be one of the good choices for the moment method, even though its numerical simulation is not efficient. However, due to the complexity of the maximum entropy equations, till now there are few studies on the moment equations. In this talk, I will investigate the properties of MEE, including the closure, characteristic speeds, asymptotic behavior of the closure and numerical behavior of some classical problems.

Yuwei Fan
Stanford University
ywfan1989@gmail.com

MS20

Filtered Hyperbolic Moment Equations for Improved Convergence

Moment models are a successful way to solve the Boltzmann equation in reasonable runtime with relatively few unknowns while allowing for non-equilibrium effects of the gas. However, the convergence of moment models with respect to increasing number of moments is typically slow. In this talk, we present a method to improve the convergence of moment methods by introducing filtered hyperbolic moment equations (FHME) that results in virtually no additional computational overhead while reducing the error significantly. The filter approach is based on a careful study of averaging two moment solutions and the reformulation of the averaging using an artificial collision method that naturally gives rise to the filter. We study the properties of the filter and show numerical results for a shock tube that demonstrate the superior quality of the new filtered moment method.

Julian Koellermeier
RWTH Aachen UniversityFreie Universität Berlin
julian.koellermeier@gmx.de

MS20

Residual-based Model-adaptivity with Moment Equations

Gas flow in thermal non-equilibrium can be described by kinetic theory and Boltzmann equation at the expense of strongly increased computational costs. A hierarchical simulation approach based on moment equations should provide a single numerical framework in which a coarse representation can be used to compute gas flows as accurate and efficient as in computational fluid dynamics (CFD), but a subsequent refinement allows to successively improve the result to the complete Boltzmann result. The hierarchical nature of the method can be used to provide efficient model error estimates to increase the predictivity of a CFD computation. This talk will give a proof-of-concept approach for such a framework for the case of moment equations for the steady linearized Boltzmann equation, an implicit discontinuous Galerkin formulation and a residual-based model-error estimator.

Manuel Torrilhon
RWTH Aachen University
mt@mathcces.rwth-aachen.de

MS21

A Multi-dimensional Discontinuous Galerkin Method for Solid Dynamics

The modal Lagrangian discontinuous Galerkin (DG) method evolves piecewise polynomial expansions forward in time for specific volume, velocity, and specific total energy in Lagrangian hydrodynamics. Taylor polynomial expansions for deformation and velocity are limited to enforce boundedness at quadrature points. In this work, a specific kinetic energy polynomial expansion is derived and used to calculate a compatible specific internal energy distribution in the element that guarantees the conservation of the total energy. A merit of the compatible formulation is that the specific internal energy can be limited to enforce positivity and boundedness independent of the velocity limiting. Stress is a nodal quantity that is calculated at quadrature points and at the element corners using the constitutive model and the bounded specific internal energy and specific volume. Discontinuities in the velocity and stress at the element surface are resolved by solving a multi-directional approximate Riemann problem at the nodes. The new DG method is formulated for both fluids and solids in 2D and 3D for polytopal cells. Test problems are presented to demonstrate robustness and convergence order.

Donald E. Burton, Nathaniel Morgan
LANL
burton@lanl.gov, nmorgan@lanl.gov

Xiaodong Liu
Los Alamos National Laboratory
Computational Physics Division
xliu@lanl.gov

Michael Berry, Evan Lieberman
LANL
mrberry@lanl.gov, elieberman3@lanl.gov

MS21

Matrix-free Advection-based Remap Algorithms for Lagrangian/ALE Methods

The Arbitrary-Lagrangian-Eulerian framework is utilized for many different applications, such as multi-material hydrodynamics. Due to mesh movement and deformation during the Lagrangian phase, it is necessary to remap the solution onto a more well-conditioned mesh. This remap problem bears similarity to the scalar transport equation and is solved with similar strategies. A desired property of the numerical solution is monotonicity/positivity preservation, which is achieved by interpolating between a monotone and a more accurate, non-monotone approximation. I will describe a novelty approach of enforcing monotonicity in high order Discontinuous Galerkin methods and how it can be implemented in a matrix-free fashion. As it does not rely on the matrix-entries of the discrete transport/remap operator, our method can be used for significantly higher order approximations than the state of the art method. In my talk, I will present the mathematical foundation and numerical results for our scheme as well as a comparison to the state of the art.

Hennes Hajduk
TU Dortmund
Faculty of Mathematics
hennes.hajduk@math.tu-dortmund.de

MS21

An Interpolation-free Adaptive ALE Approach

with Multi-step Time Schemes

Mesh adaptation techniques are combined with the multi-step BDF time integrators to accurately solve the unsteady Euler equations over three-dimensional unstructured meshes. The Arbitrary Lagrangian-Eulerian (ALE) formulation and the finite-volume discretization are used to recover the solution on the adapted grid without any explicit interpolation [Re, Guardone, Dobrzynski. *J. Comput. Phys.*, 340, pp. 26–54, (2017)]. Thanks to a peculiar series of fictitious expansions and collapses, grid connectivity changes are described as continuous deformations of the finite volumes that compose the domain. The solution on the adapted grid is computed within the ALE framework taking into account the additional numerical fluxes generated by the volume swept by the cell interfaces during mesh adaptation. These contributions are also included in the definition of the interface velocities, so that the constraint known as GCL is fulfilled. The absence of explicit interpolation makes easier the use of high-order BDF schemes, since the solution at previous time steps is easily recovered through the history of the indexes of the degree of freedom. Particular care is only required to take into account the contribution of the removed nodes: if a p-steps BDF scheme is exploited, the solution on the removed nodes remains not null for p steps next to its deletion. The proposed scheme is assessed by numerical simulations of moving-body problems experiencing large displacements.

Barbara Re
University of Zurich
Institute of Mathematics
barbara.re@math.uzh.ch

Alberto Guardone
Politecnico di Milano
Department of Aerospace Science and Technology
alberto.guardone@polimi.it

Cecile Dobrzynski
IMB, IPB, Université de Bordeaux, France
Bacchus, INRIA Bordeaux Sud-Ouest, France
cecile.dobrzynski@math.u-bordeaux1.fr

MS21

A High Order Multidimensional Residual Distribution Scheme for Lagrangian Hydrodynamics

We present the high-order multidimensional Staggered Grid Residual Distribution (SGH RD) scheme for Lagrangian hydrodynamics. The SGH RD scheme is based on the staggered finite element discretizations. However, the advantage of the residual formulation over more classical FEM approaches consists in the natural mass matrix diagonalization which allows one to avoid the solution of the linear system with the global sparse mass matrix while retaining the desired order of accuracy. This is achieved by using Bernstein polynomials as finite element shape functions and coupling the space discretization with the deferred correction type timestepping method. Moreover, it can be shown that for the Lagrangian formulation written in non-conservative form, our residual distribution scheme ensures the exact conservation of the total energy. We shall also discuss stabilization techniques for the SGH RD schemes allowing to reduce the dissipation of the numerical solution. Thanks to the generic formulation of the staggered grid residual distribution scheme, it can be directly applied to both single- and multimaterial and multiphase models. Finally, we shall demonstrate computational results obtained with the proposed residual distribution scheme for several challenging test problems.

Svetlana Tokareva
Los Alamos National Laboratory
Applied Mathematics and Plasma Physics Group (T-5)
tokareva@lanl.gov

Remi Abgrall
University of Zurich
remi.abgrall@math.uzh.ch

Konstantin Lipnikov
Los Alamos National Laboratory
lipnikov@lanl.gov

Nathaniel Morgan
LANL
nmorgan@lanl.gov

MS22

Kokkos Libraries and Applications

With clear goals of portability, performance, and usability, Kokkos and its accompanying suite of tools established itself in the scientific computing with the use of modern C++. Kokkos is portable to many advanced many-core architectures that cover almost all supercomputing and computational science installations. Supported architectures include multicore CPUs, NVIDIA GPUs, Intel Xeon Phi processors, and potentially AMD Fusion. Kokkos maximizes the amount of user code (both application and library) that can be compiled without modification and run on these architectures. It also minimizes the amount of architecture-specific knowledge that a user is required to have. At the same time, Kokkos allows architecture-specific tuning to easily co-exist with the architecture-agnostic portion. Kokkos complies with modern C++ standards. Its testing and deployment on a broad range of compilers help it attain a good instruction mix on modern hardware. Code that uses Kokkos performs as well as architecture-specific code. Kokkos produces thread-scalable implementations that are lock free whenever possible. Finally, Kokkos is a small, straight-forward application programmer interface that neither constrains end users, nor compromises portability or performance.

David S. Hollman
Sandia National Labs
dshollm@sandia.gov

MS22

Benchmarking Modern C++ Abstraction Penalty

To understand a few fundamental techniques of C++, this presentation will present fundamental examples, among others, issues related to type safety, resource management, compile-time computation, error-handling, concurrency, performance, object-oriented programming, and generic programming. The presentation relies on and introduces a few features from the recent ISO C++ standard, C++11 (and maybe later). The emphasis is on simplification of the discussion of C++ fundamentals and modern expression style.

Andrew Lumsdaine
University of Washington

al75@uw.edu

MS22

Interfacing Dense Linear Algebra Libraries in C++

Established software libraries for dense linear algebra have enjoyed a consistent interface since their inception decades ago. However, the recent advances of the ISO C++ standard offer new opportunities to provide modern interfaces to the portable performance to BLAS and LAPACK and their customized implementations on a variety of hardware platforms. With features such as templates, structured binding, and range-based for-loops it is possible to succinctly express the common functionality and data access patterns to greatly expand applicability of the legacy API's while maintaining portability and performance of our new interface. In addition, the presented C++ BLAS and LAPACK interfaces are already made available with the widely used distribution of open source software from Netlib and new libraries are built on top of it. The primary example being SLATE – a ScaLAPACK replacement for modern hardware with heterogeneous hardware featuring accelerator components.

Piotr Luszczek
University of Tennessee
luszczek@icl.utk.edu

MS22

Modern C++ in Computation Science

In an effort to develop and implement robust algorithms and enabling technologies using modern object-oriented software design, we have built a common infrastructure that links together a number of software packages. As a disparate collection of modules, we now rely on some C++11 features: for example, the required dependence on Kokkos or Tpetra. At the same time it can be removed through the include guards to increase portability with support from non-compliant compilers.

Daniel Sunderland
Sandia National Laboratories
dsunder@sandia.gov

MS23

Hybrid High-order Methods on Unfitted Meshes

We design and analyze a Hybrid High-Order (HHO) method on unfitted meshes and we study its application to the approximation of elliptic and Stokes problems. The curved interface can cut through the mesh cells in a very general fashion. The unfitted HHO method uses cell and face unknowns and is formulated by means of local reconstruction and stabilization operators, together with a Nitsche-type technique to account for the unfitted interface. Cell unknowns can be eliminated locally leading to a global problem coupling only the face unknowns by means of a compact stencil. We prove stability estimates and optimal error estimates in the H^1-norm. Robustness with respect to cuts is achieved by a local cell-agglomeration procedure taking full advantage of the fact that HHO methods support polyhedral meshes. This is a joint with Prof. Alexandre Ern and Prof. Erik Burman.

Guillaume Delay
Institut de Mathématiques de Toulouse
guillaume.delay@enpc.fr

MS23

A Higher Degree Immersed Finite Element Method Based on a Cauchy Extension for Elliptic Interface Problems

We present a p-th degree immersed finite element (IFE) method for solving the elliptic interface problems with meshes independent of the coefficient discontinuity in the involved partial differential equations. In this method, an IFE function is the extension of a p-th degree polynomial from one subelement to the whole interface element by solving a local Cauchy problem on interface elements in which the jump conditions across the interface are employed as the boundary conditions. This approach completely addresses the existence issue of higher degree IFE functions circumventing further discussion based on different interface element configurations. And the analysis is performed to prove that the proposed IFE method converges optimally in both the L2 and H1 norms.

Ruchi Guo
Virginia Polytechnic Institute and State University
The Department of Mathematics
ruchi91@vt.edu

Tao Lin
Virginia Tech
tlin@vt.edu

MS23

A Posteriori Error Estimation for Cut Finite Element Method

Abstract not available.

Cuiyu He
Purdue University
c.he@ucl.ac.uk

MS23

A Finite Element Method For PDEs in Time-dependent Domains

In the talk we discuss a recently introduced finite element method for the numerical solution of partial differential equations on evolving domains. The approach uses a completely Eulerian description of the domain motion. The physical domain is embedded in a triangulated computational domain and can overlap the time-independent background mesh in an arbitrary way. The numerical method is based on finite difference discretizations of time derivatives and a standard geometrically unfitted finite element method with an additional stabilization term in the spatial domain. The performance and analysis of the method rely on the fundamental extension result in Sobolev spaces. Theoretical findings include a complete stability and error analysis, which accounts for discretization errors resulting from finite difference and finite element approximations as well as for geometric errors coming from a possible approximate recovery of the physical domain. We show numerical examples that illustrate the theory and demonstrate the practical efficiency of the method.

Maxim A. Olshanskii
Department of Mathematics
University of Houston

molshan@math.uh.edu

MS24

Hierarchical Fidelity to Meet the Evolving Needs of Basic and Applied Combustion Simulations

Contemporary combustion research for advanced internal combustion engines requires a multi-faceted approach to understand the design space as well as develop fundamental understanding of turbulence-chemistry interaction in relevant regimes. The Pele project involves a multi-scale treatment anchored in two algorithms to treat the problem: a compressible formulation for gas exchange and fuel injection processes and a low Mach formulation to treat the lower speed flows found during, for example, the closed portion of an internal combustion engine cycle efficiently. Multi-physics capability and complex geometry, represented by cartesian cut cells, along with adaptivity, are necessary. The two algorithms are being built out into an application based on AMReX that will ultimately allow the researcher to use a hierarchy of fidelities in concert to attack the problem from parameter estimation to fundamental research for understanding and model development.

Ray W. Grout
National Renewable Energy Laboratory
ray.grout@nrel.gov

MS24

MFIX-Exa: Using AMReX to Modernize a Legacy Code for Fossil Energy Applications

The combination of computational fluid dynamics and discrete element method (CFD-DEM) offers an accurate way to model gas-solids reactors. However, CFD-DEM simulations are computationally expensive, limiting their application to small systems containing, at most, tens of millions of particles. To apply CFD-DEM to accelerate fossil energy technology development, the modeling of pilot-scale reactors containing hundreds of billions of particles is needed. To address the computational and data management challenges, a new code is being developed by migrating the core hydrodynamic models from the widely used, open-source code MFIX into the AMReX framework. Parallel performance of the code has been improved by implementing logical tiling, enabling OpenMP threading over tiles, and employing a dual grid approach for fluid and particle load balancing. To handle complex reactor geometries, AMReX embedded boundary (EB) data structures and iterators were incorporated and an algorithm to address particle collisions with EB walls was implemented.

Jordan Musser
National Energy Technology Laboratory (NETL)
jordan.musser@netl.doe.gov

Madhava Syamlal
Fluent
MADHAVA.SYAMLAL@netl.doe.gov

Ann S. Almgren
Lawrence Berkeley National Laboratory
asalmgren@lbl.gov

John B. Bell
Lawrence Berkeley National Lab
Center for Computational Sciences and Engineering
jbbell@lbl.gov

MS24

Overview of the AMReX Framework - Capabilities and Scientific Applications

AMReX is a DOE Exascale Computing Project (ECP) co-design center that provides a mature, publicly-available software framework for block-structured adaptive mesh refinement (AMR) applications for current and next-generation architectures. A primary mission of the co-design center is to support user applications. I will describe current capabilities including particles, embedded boundaries, linear solvers, multicore support, visualization, and profiling tools, and discuss new programming models in the code. AMReX is being used by researchers to study a wide range of problems; in this minisymposium, speakers will highlight scientific achievements in solid mechanics, multiphase flow, combustion, astrophysics, accelerator physics, and fluctuating hydrodynamics. AMReX also includes extensive documentation and tutorial codes that allow users to develop their own block-structured mesh applications.

Andy J. Nonaka
Lawrence Berkeley National Laboratory
ajnonaka@lbl.gov

MS24

Modeling of Interface-driven Microstructure Evolution in Metals using the Multiphase Field Method

Abstract not available.

Brandon S. Runnels
University of Colorado, Colorado Springs
brunnels@uccs.edu

MS25

Data-informed Subspace Identification using a Consistent (Data-oriented) Bayesian Method

Bayesian inversion framework is the most popular approach that can incorporate most, if not all, uncertainties in inverse solutions. Bayesian prior can be understood as a regularization that penalizes parameters uninformed by the data, and hence making the inverse problem well-posed. Ideally it should not interfere with well-informed parameters. Unfortunately, well-informed parameters are not known a priori and none of the existing prior elicitation approaches could avoid polluting data-informed directions in the parameter space. We tackle the challenge of establishing a new statistical inversion paradigm, a data-consistent approach, that automatically prevents any prior from regularizing the well-informed parameter directions.To accomplish this seemingly impossible task, we start from scratch using the disintegration theory, a more general conditioning theory than Bayesian one, to construct posterior that respects the data in the aforementioned sense. Rigorous and numerical results will be presented to support the data-consistent statistical inversion theory.

Tan Bui-Thanh
The University of Texas at Austin
tanbui@ices.utexas.edu

MS25

Scalable Approximations for the Consistent Bayes

Method

We present a scalable approach to the solution of linear and nonlinear stochastic inverse problems using the recently developed consistent Bayesian method. In particular, the use of kernel density estimation in the consistent Bayesian method results in poor scaling with the dimension of the data. Our approach avoids the need for kernel density estimation and is therefore scalable with the data dimension. In addition, we adapt a randomized dimension reduction technique developed for the classical Bayesian maximum a posteriori (MAP) point method to develop a consistent Bayesian MAP point method which scales with the data, state, and parameter dimensions. We demonstrate our scalable approach on a linear inverse problem govern by an advection-diffusion equation and a nonlinear inverse problem govern by a hyperelasticity equation.

Brad Marvin
University of Texas at Austin
brad_marvin@utexas.edu

Timothy Wildey
Sandia National Laboratories
tmwilde@sandia.gov

Tan Bui-Thanh
The University of Texas at Austin
tanbui@ices.utexas.edu

MS25

Accelerating Prediction under Uncertainty with Dimension Reduction

Making predictions under uncertainty often involves solving stochastic inverse problems that have high-dimensional spaces of input parameters. Furthermore, the solution spaces of underlying forward problems are also often high-dimensional. This high-dimensionality may cause the problem of making a prediction under uncertainty extremely computationally expensive and possibly infeasible with traditional methods. We combine methods to effectively reduce the dimension of input parameter spaces and to reduce the dimension of the PDE solution space to accelerate predictions under uncertainty that are suffering the curse of dimensionality. These methods include active subspaces, reduced basis approximations, and model order reduction. Additionally, adaptivity is incorporated to further increase accuracy without greatly increasing the computational costs.

Steven A. Mattis
TUM
mattis@ma.tum.de

Barbara Wohlmuth
Technical University of Munich
barbara.wohlmuth@ma.tum.de

MS25

Dimensionality Reduction for Subsurface Flow Models

The progress of modern computers and sensing devices has led to subsurface flow models that have a large number of input parameters and these models must often be calibrated to match a large number of observations. These high dimensional input and output spaces can make it challenging to use these models effectively, rendering dimensionality reduction desirable in many contexts. Here we consider the application of dimensionality reduction techniques based on randomized matrix algorithms and the physical properties of the system to subsurface flow problems in traditional and fractured aquifers.

Daniel O'Malley
Los Alamos National Laboratory
omalled@lanl.gov

MS26

Parallel Multigrid in Time for Hyperbolic Problems

The focus of this talk is on the application of the multigrid reduction in time (MGRIT) method to hyperbolic partial differential equations. We first review an MGRIT approach developed in [De Sterck, Falgout, Howse, MacLachlan, Schroder, "Parallel-in-Time Multigrid with Adaptive Spatial Coarsening for the Linear Advection and Inviscid Burgers Equations', LLNL-JRNL-737050] that coarsens in both time and space. In the case of explicit time-stepping, spatial coarsening is needed to ensure stability on all levels, but it is also useful for implicit time-stepping by producing cheaper multigrid cycles. Unfortunately, uniform spatial coarsening results in extremely slow convergence when the wave speed is near zero, even if only locally. An adaptive spatial coarsening strategy helps to address this issue for the variable coefficient linear advection equation and the inviscid Burgers equation using first-order explicit or implicit time-stepping methods. Numerical results show that this offers significant improvements over classical coarsening. We then argue that spatial coarsening is not a good idea in general for hyperbolic problems, even for explicit methods, and discuss alternative approaches based on approximating the ideal Petrov-Galerkin coarse-grid operator. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Robert D. Falgout
Lawrence Livermore National Laboratory
rfalgout@llnl.gov

MS26

Convergence Analysis for Parallel-in-Time Solution of Hyperbolic Systems

Parallel-in-time algorithms have been successfully employed for reducing solution times of a variety of problems, especially of diffusive problems such as parabolic-type partial differential equations. A major drawback of parallel-in-time approaches, however, is that most methods show instabilities for hyperbolic problems. This talk focuses on understanding this behavior for time-multigrid methods. Three analysis tools are considered that differ, in particular, in the treatment of the time dimension: (1) space-time local Fourier analysis, using a Fourier ansatz in space and time, (2) semi-algebraic mode analysis, coupling standard local Fourier analysis approaches in space with algebraic computation in time, and (3) a two-level reduction analysis, reducing the analysis to the coarse grid. We discuss insights into convergence from the different tools for a one-dimensional linear advection problem and for two-dimensional incompressible linear elasticity.

Stephanie Friedhoff
University of Wuppertal

friedhoff@uni-wuppertal.de

Hans De Sterck
Monash University, Australia
Applied Mathematics
hdesterck@uwaterloo.ca

Alexander Howse
University of Waterloo
ajmhowse@gmail.com

Scott Maclachlan
Department of Mathematics and Statistics
Memorial University of Newfoundland
smaclachlan@mun.ca

MS26

Space-time Reduction Methods for Linear PDEs

Parallel time integration of partial differential equations (PDEs) has become an increasingly important topic over the past two decades. Many solution approaches have been developed, particularly for a range of diffusion-dominated problems. An example of one such algorithm is that of multigrid reduction-in-time (MGRIT). MGRIT is primarily a time-only coarsening algorithm, and thus is not directly suitable for explicit discretizations with stability restrictions, as in the case of hyperbolic problems, for example. A possible solution is to combine MGRIT with spatial coarsening, but it has been found that MGRIT combined with naive spatial coarsening does not yield a good solver for model hyperbolic PDEs. In an attempt to remedy this issue, we develop space-time algorithms that use a combination of both spatial and temporal multigrid reduction. We will discuss progress to date, including results for model hyperbolic and parabolic PDEs.

Oliver A. Krzysik, Hans De Sterck
Monash University
oliver.krzysik@monash.edu, hans.desterck@monash.edu

Scott Maclachlan
Department of Mathematics and Statistics
Memorial University of Newfoundland
smaclachlan@mun.ca

Alexander Howse
University of Waterloo
ajmhowse@gmail.com

Stephanie Friedhoff
University of Wuppertal
friedhoff@uni-wuppertal.de

MS26

A Relaxation Approach to Parallel-in-Time

The sequential nature of time-integration often creates bottlenecks for high-fidelity numerical simulations. One approach for accelerating the time-integration aspect of a simulation, is to compute a (faster) low-fidelity approximation to the solution, and in parallel, correct the low-fidelity approximation to generate a highly accurate solution. This approach has been studied in a variety of settings, such as parallel deferred correction methods and waveform relaxation methods. Here, we present a preliminary investigation into the benefits of propagating inaccurate solutions within an outer iterative process while correcting, in parallel, inner iterative processes.

Benjamin W. Ong
Michigan Technological University
ongbw@mtu.edu

Felix Kwok
Hong Kong Baptist University
felix_kwok@hkbu.edu.hk

MS27

Parallel Algorithms for Nonlinear Time-space Fractional PDEs

Abstract not available.

Toheeb A. Biala
Middle Tennessee State University
tab7v@mtmail.mtsu.edu

MS27

Source Determination for a Two-parameter Anomalous Diffusion with Local Time Datum

Anomalous diffusion is a well-documented phenomenon in many complex systems. We present an inverse source problem for a two-parameter fractional diffusion equation subject to nonlocal nonself-adjoint boundary conditions and two local time distinct datum. This problem models several physical processes, among them the microwave heating. The two local time conditions spare us from measuring the fractional integral initial conditions commonly associated with fractional derivatives. On the other hand, they lead to delicate 2 2 linear systems for the Fourier coefficients of the source term and of the fractional integral of the solution at t = 0. The asymptotic behavior and estimates of the generalized Mittag-Leffler function are used to establish the solvability of these linear systems. Analytical and numerical examples will presented.

Khaled Furati
Department of Mathematics & Statistics
King Fahd University of Petroleum & Minerals
kmfurati@kfupm.edu.sa

MS27

Sparse Approximation of Non Local Operator Equations via Compressed Sensing

In the field of signal processing, the shift to digitizing analog signals raised a lot of questions regarding the efficient acquisition, compression and storage of large datasets. The size of data and the limited resolution of sensing devices motivated the development of sampling algorithms which adjust to the intrinsic structure of signals. Compressed sensing is one of such techniques which exploits the sparsity of signals in certain bases to guarantee the accurate recovery from the small number of random measurements via the solution of an optimization problem. The method has proven its efficiency in several applications including image processing and approximation theory. Recently it has been considered in the context of solving differential equations by noticing the analogy between the sampling of continuous signals and the finite-dimensional approximation of boundary-value problems in Petrov-Galerkin formulation. In the current talk, we elaborate on this idea and extend it to a class of nonlocal operators for which the Galerkin projection schemes result in the fully dense matrix equa-

tions. In this setting, the compressive sensing approach can lead to significant computational savings by sampling in the appropriate test space with a rate proportional to the sparsity of the solution in the trial space. We discuss both theoretical and practical aspects of the proposed algorithm and provide its numerical justification on a set of examples.

Viktor Reshniak
Oak Ridge National Lab
reshniakv@ornl.gov

MS28

AMReX and Applications

AMReX is a software framework that supports the development of block-structured AMR algorithms for solving systems of partial differential equations (PDEs) with complex boundary conditions on exascale architectures. AMR reduces the computational cost and memory footprint compared to a uniform mesh while preserving the essentially local descriptions of different physical processes in complex multiphysics algorithms. AMReX supports a number of different time-stepping strategies and spatial discretizations, and incorporates multilevel data containers and iterators for mesh-based fields, particle data and irregular embedded boundary (cut cell) representations of complex geometries. Current AMReX applications include accelerator design, astrophysics, combustion, cosmology, microfluidics, materials science, multiphase flow and others. In this talk I will give an overview of the software components provided by AMReX and how those components are used to support different applications that use the framework.

Ann S. Almgren
Lawrence Berkeley National Laboratory
asalmgren@lbl.gov

MS28

Communication and Synchronization Reduction in Sparse Factorization and Triangular Solution

Abstract not available.

Xiaoye S. Li
Lawrence Berkeley National Laboratory
xsli@lbl.gov

MS28

Solving Nonlinear Eigenvalue Problems with Pad Approximate Linearization

We present a strategy for finding nontrivial solutions (λ, x) for a class of nonlinear eigenvalue problems of the form $F(\lambda)x = 0$. We use rational functions to approximate the nonlinear terms of the problem, together with Padé approximants. This leads to a linearization of the original problem, through a generalized eigenvalue of dimension larger than the original problem. We show the impact of the degree of the Padé approximants in the linearization process and convergence, and alternatives for solving the resulting linearized problem. As study cases, we focus on problems related to the modeling of waveguide-loaded accelerator cavities through a finite element discretization of Maxwell's equation, where some of the matrices involved exhibit low-rank properties and favor the Padé approximation strategy.

Osni A. Marques
Lawrence Berkeley National Laboratory
Berkeley, CA
oamarques@lbl.gov

Zhaojun Bai
Departments of Computer Science and Mathematics
University of California, Davis
bai@cs.ucdavis.edu

MS28

Unstructured Mesh Support for Particle-in-C Simulations

Particle-in-cell (PIC) method are becoming more c monly applied to resolve fine scale physics in problem which PDE based methods alone are not sufficient. In methods the motion of the particles is dictated by the m based field, which in general is then altered due to the ticle motion. In applications with complex geometries highly varying and/or anisotropic fields, there is a desir employ unstructured meshes with millions of element which the data for both the mesh and the particles need be distributed. The presentation will provide an overv of a set of tools, PUMIpic, being developed to support calculations of distributed unstructured meshes. The components of PUMIpic is a distributed mesh with l overlaps to support effective particle push steps, effec continuum equation construction in a manner consis with the particle driven mesh distribution, and procedu that dynamically migrate particles between parts to m tain load balance. The current state of the developmer a distributed mesh version of the XGC gyrokinetic fu plasma code using PUMIpic components will be provi

Mark S. Shephard
Rensselaer Polytechnic Institute
Scientific Computation Research Center
shephard@rpi.edu

Eisung Yoon
Ulsan National Institute of Science and Technology
esyoon@unist.ac.kr

Agnieszka Truszkowska, Gopan Perumpilly
Rensselaer Polytechnic Institute
trusza@rpi.edu, gopan.p@gmail.com

E. Seegyoung Seol
Rensselaer Polytechnic Institute
Scientific Computation Research Center
seols@rpi.edu

Onkar Sahni
Rensselaer Polytechnic Institute
sahni@rpi.edu

Cameron Smith
Scientific Computation Research Center
Rensselaer Polytechnic Institute
smithc11@rpi.edu

Gerrett Diamond, William R. Tobin
Rensselaer Polytechnic Institute

diamog@rpi.edu, tobinw2@rpi.edu

MS29

Regionally Implicit Discontinuous Galerkin Methods for Solving the Relativistic Vlasov-Maxwell System

In the relativistic limit, the Vlasov-Maxwell system introduces numerical difficulties as explained in (Suzuki, 2010). We develop an efficient solver for the relativistic Vlasov-Maxwell (RVM) system in order to model laser-plasma interactions; and in particular, the acceleration of electrons or ions to relativistic energies. In doing so we expand on the so-called Locally Implicit Discontinuous Galerkin method (LIDG) developed in (Qiu, 2005) by defining the Regionally Implicit Discontinuous Galerkin Methods. We create small implicit systems (regions) consisting of a cell and its neighbors in order to form a space-time prediction for each cell, which is then corrected by time-averaging fluxes. We show that these high order methods allow a much larger CFL number when compared to the high order LIDG and RKDG methods, and thus offer a improved efficiency. Here we demonstrate the efficiency of the RIDG method for nonlinear problems (compared to RKDG). We then discuss application of the RIDG method to the relativistic Vlasov-Maxwell system.

Pierson Guthrey
Michigan State University
Computational Mathematics Science and Engineering
piersonguthrey@gmail.com

James A. Rossmanith
Iowa State University
Department of Mathematics
rossmani@iastate.edu

MS29

Comparison of Boltzmann and Fokker-Planck Collision Operators for Moderately Coupled Plasmas

Abstract not available.

Jeffrey Haack
Los Alamos National Laboratory
haack@lanl.gov

MS29

Verification Test Cases of Grid-based Direct Kinetic Modeling Framework for Plasma Flows

Verification and validation are critical procedures when developing a new numerical model. We will present a few test cases for grid-based kinetic models, which we call the direct kinetic (DK) method. A conservative, positivity-preserving numerical scheme is employed, which is necessary for robust calculations of problems particularly when including ionization. The test cases include standard (linear and nonlinear) Landau damping problems, plasma sheaths with electron emission, injection and expansion of neutral atom flow in a two-dimensional configuration, and a comparison between particle-in-cell and DK methods modeling in a Hall thruster discharge plasma. For a purely collisionless case, a small negative value of the velocity distribution function using a high-order scheme, e.g., WENO, can be tolerated. However, it was found that the simulation sometimes can be unstable due to the small negativity when ionization term is included. Through the verification studies, we have corrected the kinetic theory of space charge limited sheaths, accounting for the non-Maxwellian distributions for all constituents, which illustrates that a numerical model can help plasma theory as well. Other recent test cases will also be discussed.

Ken Hara
Texas A&M University
khara@tamu.edu

Kyle Hanquist
University of Michigan
hanquist@umich.edu

MS29

Solving the Continuum Kinetic Equations for Plasmas using the Discontinuous Galerkin Method

The discontinuous Galerkin method has been used to solve equations in plasma physics in fluid and kinetic regimes. Some challenges in modeling plasmas include the need to resolve disparate spatial and temporal scales with physically-relevant dispersive and diffusive phenomena. A direct discretization of the Vlasov equation using the discontinuous Galerkin method is used to study applications in kinetic regimes. Typically, kinetic simulations use particle-in-cell codes for each of the ion and electron species, but particle codes are subject to noise. Direct discretization of the Vlasov equation can provide smooth solutions of the distribution function in phase space. Continuum kinetic simulations of plasmas require a careful construction of initial conditions, boundary conditions, and collision operators to maintain the conservative properties of the scheme. Kinetic simulations of two-component plasmas, the two components being ions and electrons, will be presented using the Vlasov-Maxwell and Boltzmann-Maxwell equations.

Bhuvana Srinivasan, Petr Cagas
Aerospace and Ocean Engineering
Virginia Tech
srinbhu@vt.edu, pcagas@vt.edu

Ammar Hakim
Princeton Plasma Physics Laboratory
ahakim@pppl.gov

MS30

Construction of Massive Protein Similarity Networks using Distributed Sparse Matrices

Sequence clustering methods, in general, rely on a similarity graph as input. However, all-pairs sequence alignment approaches are prohibitively expensive and unnecessary as not all pairs would result in statistically significant alignments. Seed-and-extend provides an efficient alternative approach. However, the current sequential methods take on the order of days to align large datasets. This talk focuses on using distributed sparse matrices to efficiently scale the construction of similarity networks for datasets over 250 million sequences.

Saliya P. Ekanayake
Lawrence Berkeley Lab
esaliya@lbl.gov

MS30

Parallel Algorithms for Degree Constrained Sub-

graphs with Applications

Given a weighted graph, G, the objective of the Degree Constrained Subgraph problem is to find an optimal (or approximate) weighted subgraph of G, subject to particular degree constraints on the vertices of the subgraph. In this talk, we will concentrate on two types of constraints; namely lower bound and upper bound on degree. We will discuss how these problems relate to better known combinatorial algorithms, and will show a list of approximation algorithms to solve these problems. Our focus on designing such algorithms is to obtain scalable parallel algorithms. To motivate our study on this problem, we will show the applicability of the degree constrained subgraph problem in areas such as anonymization of large data sets and sparsification of dense graphs for machine learning purposes.

S M Ferdous
Purdue University
sferdou@purdue.edu

Alex Pothen
Purdue University
Department of Computer Science
apothen@purdue.edu

MS30

A Scalable Algorithm for Data Anonymization

We explore the problem of sharing data that pertains to individuals with anonymity guarantees, where each user requires a desired level of privacy. We discuss the first shared-memory as well as distributed memory parallel algorithms for the adaptive anonymity problem that achieves this goal, and produces high quality anonymized datasets. The new algorithm is based on an optimization procedure that iteratively computes weights on the edges of a dissimilarity matrix, and at each iteration computes a minimum weighted b-Edge Cover in the graph. We describe how a 2-approximation algorithm for computing the b-Edge Cover can be used to solve the adaptive anonymity problem in parallel. We are able to solve adaptive anonymity problems with hundreds of thousands of instances and hundreds of features on a supercomputer in under five minutes. Our algorithm scales up to $8K$ cores on a distributed memory supercomputer, while also providing good speedups on shared memory multiprocessors. On smaller problems where an algorithm based on Belief Propagation is feasible, our algorithm is two orders of magnitude faster.

Arif Khan
Pacific Northwest National Lab
ariful.khan@pnnl.gov

MS30

A Scheduling Problem Motivated by Cybersecurity and Adaptive Machine Learning

We describe a scheduling problem motivated by cybersecurity. We briefly describe PLADD (Probabilistic Learning Attacker, Dynamic Defender), an abstract game introducted by Jones et. al. that approximates the competition for a resource between the resource owner (defender) and an attacker. PLADD models some elements of moving-target defenses in cyber systems. One can also use the same abstractions to study machine-learning-based classifiers under concept drift. In this case, the attacker is concept drift and the defender is the machine-learning model which must adapt to the drift. The goal of these models is to develop optimal defender/update strategies. One way to study defense strategies is to sample attack scenarios from the probability distributions representing time to attacker success. We can then compute the optimal sequence of defender actions for that finite attack scenario set using stochastic programming. In this case, the stochastic program has a scheduling interpretation and an integer-programming formulation. We will discuss progress towards formal complexity analysis, approximation algorithms, and report on computational experiments using cybersecurity data.

Ojas Parekh, Cynthia Phillips
Sandia National Laboratories
odparek@sandia.gov, caphill@sandia.gov

Vladlena Powers, Nouri Sakr, Cliff Stein
Columbia University
vp2342@columbia.edu, n.sakr@columbia.edu, cliff@ieor.columbia.edu

MS31

Hybrid-mixed Finite Element Solvers for Compressible Atmospheric Equations

There is a current explosion of interest in new numerical methods for operational weather forecasting and climate simulation. A driving force behind this is the need to be able to efficiently simulate atmospheric dynamics on massively parallel computer systems. This goal is a central aspect in the design of the next generation dynamical core, as part of the "Gung-Ho" project at the UK Meteorological Office (Met Office). In this talk, we present mixed finite element spatial discretizations for three-dimensional compressible atmospheric equations. We focus on a solution procedure known as "hybridization," which leads to a sparse problem on mesh interfaces. Once the interface unknowns are determined, the original fields can then be recovered in a local manner. We present some results using the Firedrake finite element library, as well as discuss the implementation of hybridization in the new Met Office dynamical core.

Thomas H. Gibson
Department of Mathematics
Imperial College London
t.gibson15@imperial.ac.uk

MS31

The E3SM Non-hydrostatic Dynamical Core at Cloud Resolving Resolution - DCMIP 2016 Super Cell Thunderstorm

We present results from a simulation of an idealized Super Cell thunderstorm using the HOMME non-hydrostatic dynamical core with local grid refinement. We seek to demonstrate the performance of the variable resolution method and investigate this solution relative to a uniformly fine resolution model. We will consider details of performance, accuracy, and interactions with physical parameterizations that arise as a consequence of the irregular mesh.

Jorge E. Guerra
Department of Land, Air and Water Resources
University of California, Davis

jeguerra@ucdavis.edu

MS31

High-order Finite-volume Weno Discretizations that Effectively use Accelerated HPC for Atmospheric Dynamics

Many real-world atmospheric dynamics applications have a throughput constraint that is significantly more stringent than that of other scientific fields. To run these simulations fast enough, especially for high-resolution climate, the model must be scaled out to many compute nodes. As this happens, the costs of transferring data between compute nodes begins to dominate, even to the point that most of the runtime is spent moving data rather than doing computations. This talk approaches this scaling challenge from an algorithmic point of view, considering the role of higher-order accuracy and more expensive (yet significantly more accurate) limiting mechanisms. The goal is to give a general discussion of the challenges, ideas, and approaches including the form of the numerical operator itself, limiting mechanisms, and temporal discretizations, all in the context of modern, accelerated HPC constraints.

Matthew R. Norman
Oak Ridge National Laboratory
Oak Ridge, TN
normanmr@ornl.gov

MS31

Evaluation of Additive Runge-Kutta Methods for the E3SM Non-hydrostatic Dycore

HOMME-NH is the non-hydrostatic dynamical core (dycore) for the Exascale Earth System Model (E3SM). As a dycore, it solves a system comprised of the Euler equations coupled with a modified ideal gas law for the equation of state. This system is numerically stiff from vertically shallow computational grid cells combined with fast-moving acoustic waves that have little effect on the overall solution. As such, the class of Implicit-Explicit, Additive Runge-Kutta (ARK) methods is chosen to advance the discretized system in time. The Suite of Nonlinear and Differential/Algebraic Equation Solvers (SUNDIALS) package is used to quickly incorporate ARK methods into HOMME-NH. The evaluation of these methods includes various metrics that exist in the dycore domain: self-convergence error, global energy conservation error, maximum stable timestep size, and time to obtain an acceptable surface pressure solution. The impact on these metrics from hyperviscosity, a common tool used to stabilize the spectral element spatial discretization, is also investigated. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Christopher J. Vogl
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
vogl2@llnl.gov

Daniel R. Reynolds
Southern Methodist University
Mathematics
reynolds@smu.edu

Paul Ullrich
University of California, Davis
paullrich@ucdavis.edu

Andrew J. Steyer
Sandia National Laboratories
asteyer@sandia.gov

Carol S. Woodward, David J. Gardner
Lawrence Livermore National Laboratory
woodward6@llnl.gov, gardner48@llnl.gov

MS32

Error Estimation for Multiplicative Schwarz Domain Decomposition Method

Domain decomposition methods are widely used for the numerical solution of partial differential equations on parallel computers. This work forms an error analysis for the overlapping multiplicative Schwarz domain decomposition method by employing tools from variational analysis, adjoint operators and computable residuals. The error is decomposed into a component arising due to finite iteration between the subdomains and a component due to discretization of each spatial subdomain. Moreover, the error estimate accounts for errors arising from each subdomain and hence guides in the choice of discretization employed to reduce the error in the solution.

Jehanzeb H. Chaudhry
University of New Mexico
jehanzeb@unm.edu

Simon Tavener, Don Estep
Colorado State University
tavener@math.colostate.edu, estep@stat.colostate.edu

MS32

Discretization Error Control for Constrained Aerodynamic Shape Optimization

In modern aircraft design, numerical optimization based on CFD simulations has been heavily used to improve vehicle performance. In order to achieve reliable and practical designs, discretization error has to be carefully controlled during the optimization. In this work, we present a method to control the discretization error in constrained aerodynamic shape optimization problems using meshes adapted via adjoint-based error estimates. The optimization constraints may involve outputs that are not directly targeted for optimization, and hence also not for error estimation and mesh adaptation. However, discretization error in these outputs often indirectly affects the evaluation of the objective function. The proposed method includes this effect into objective error estimates by introducing a coupled adjoint solution, so that the mesh can be adapted to predict both the objective outputs and constraint outputs with appropriate accuracy. A multi-fidelity optimization framework is developed by taking the advantage of variable fidelity offered by adaptive meshes. The multi-fidelity optimization process prevents over-optimization on coarse meshes and over-refinement on undesired designs. We demonstrate the accuracy and efficiency of the proposed method on several airfoil shape optimization problems.

Guodong Chen
Aerospace Engineering department
University of Michigan, Ann Arbor
cgderic@umich.edu

Krzysztof Fidkowski
University of Michigan

kfid@umich.edu

MS32

An A Posteriori Error Estimation for Numerical Solution of Richards Equation

We first discuss several time integration techniques of system of ordinary differential equations characterized by strong nonlinear coupling of the unknown variables. This system is a result of spatial discretization (using such as finite element or finite difference) of the Richards Equation, which is a governing mathematical principle for modeling water infiltration through a subsurface. The nature of Richards Equation is further complicated by the fact that the rate of change of the quantity of interest represented by a time derivative is also nonlinear. We formulate a general framework of the numerical time integration as a discontinuous Galerkin method in temporal variable. The actual implementation of a particular scheme is realized by imposing certain finite element space in time variable to the variational equation and appropriate "variational crime" in the form of numerical quadrature for calculating the integration in the formulation. Once this is in place, we derive an adjoint-based error estimator for the approximate solution from the scheme. The adjoint problem is obtained from appropriate linearization of the nonlinear system. Several numerical examples are presented to illustrate performance of the scheme and the error estimator.

Victor E. Ginting
Department of Mathematics
University of Wyoming
vginting@uwyo.edu

MS32

Variational Formulation and Error Estimation for Explicit Time Integrators

Space-time variational formulations of Partial Differential Equations are widely used to represent the error of the solution in some quantity of interest. Such error representation allows performing goal-oriented adaptivity for space-time problems. Many authors employ implicit methods in time to perform goal-oriented adaptivity, like Backward Euler or Crank-Nicholson, as it is known that these methods present variational structure. Recently, some multi-step Implicit-Explicit (IMEX) schemes have been expressed as Galerkin methods in time by using specific quadrature rules for time integration. However, a general variational formulation of explicit methods in time remains elusive. In this work, we derive a space-time variational formulation of any stage explicit Runge-Kutta method employing a discontinuous-in-time Petrov-Galerkin method. We build trial and test functions for linear PDEs and we use exact integration for the time integrals. Then, we derive an error representation and an explicit-in-time goal-oriented adaptive algorithm that enables dynamic meshes in space. Finally, we provide numerical results in 1D of the proposed explicit algorithm.

Judit Muñoz-Matute
University of the Basque Country (UPV/EHU), Leioa, Spain
judit.munoz@ehu.eus

David Pardo
Department of Applied Mathematics, Statistics, and Operational Research, University of the Basque Country
dzubiaur@gmail.com

Victor M. Calo
Curtin University, Perth, Australia
vmcalo@gmail.com

Elisabete Alberdi
University of the Basque Country (UPV/EHU), Leioa, Spain
elisabete.alberdi@ehu.es

MS33

Community Tree Persistence: A Topological Method for Community Structure Analysis in Dynamic Networks

The detection and evolutionary analysis of community structures in complex networks have garnered considerable attention in the recent past, as the structures and their dynamics provide valuable information on the innate partitions and intrinsic connectivity of the network nodes. Here, we propose a method that extends persistence-based topological data analysis, which is typically used to characterize shapes, to unweighted networks. We first introduce the concept of community trees that summarizes the topological structures in a network from a persistence perspective. A community tree is a tree structure established based on clique communities from the clique percolation method (CPM). To study the stability of community trees, total star number (TSN), a quantity that presents an upper bound on the change of community trees, is derived. Furthermore, we develop efficient algorithms to construct and update the community trees by maintaining a series of clique graphs in the form of spanning forests, in which each spanning tree is built on an underlying Euler Tour tree. With the information revealed by community trees and the corresponding persistence diagrams, our proposed approach is able to detect clique communities and keep track of major structural changes during their evolution for a given TSN threshold. The results demonstrate its effectiveness in extracting useful structural insights for time-varying networks across multiple domains.

Ashis G. Banerjee
University of Washington
ashisb@uw.edu

Wei Guo
Industrial & Systems Engineering
University of Washington
weig@uw.edu

Ruqian Chen
Mathematics
University of Washington
ruqian@uw.edu

Yen-Chi Chen
Statistics
University of Washington
yenchic@uw.edu

MS33

Topological Data Analysis of Vascular Disease

Cardiovascular disease is the leading cause of death worldwide. The current diagnosis methodology mainly relies on on the measurement of percent stenosis and fractional flow

reserve (FFR). In this talk, we present a new additional functional measure for more accurate diagnosis of the disease by using topological data analysis (TDA). We compute the homology group of the topological space composed of vascular variables and compute the persistence. We show several ways of using persistence by TDA for the diagnosis with barcode and persistence diagram.

Jae-Hun Jung
SUNY at Buffalo
jaehun@buffalo.edu

Megan Johnson
The University at Buffalo, State University of New York
meganjoh@buffalo.edu

John Nicponski
SUNY at Buffalo
johnnicp@buffalo.edu

MS33

R Package TDA for Statistical Inference on Topological Data Analysis

This presentation gives an introduction to the R package TDA, which provides some tools for Topological Data Analysis. The salient topological features of data can be quantified with persistent homology. The R package TDA provide an R interface for the efficient algorithms of the C++ libraries GUDHI, Dionysus, and PHAT. Specifically, The R package TDA includes functions for computing the persistent homology of the Rips complex, alpha complex, and alpha shape complex, and a function for the persistent homology of sublevel sets (or superlevel sets) of arbitrary functions evaluated over a grid of points or on data points. The R package TDA also provides a function for computing the confidence band that determines the significance of the features in the resulting persistence diagrams.

Jisu Kim
Carnegie Mellon University
jisu.kim@inria.fr

MS33

Stitch-fix for Mapper

Mapper is a powerful construction from topological data analysis that is designed for the analysis and visualization of multivariate data. We investigate a method for stitching a pair of mappers together, and study a topological notion of information gain as well as a general measure of correlation during such a process. We are inspired by the ideas of stepwise regression for model selection and of scatterplot matrices for visualization, and introduce a topological analogue of the scatterplot matrix for the mapper, and study the degree of topological correlation between filter functions.

Nathaniel Saul
Washington State University
nathaniel.saul@wsu.edu

Bala Krishnamoorthy
Department of Mathematics
Washington State University-Vancouver
bkrishna@math.wsu.edu

Bei Wang
University of Utah
beiwang@sci.utah.edu

MS34

Risk-Averse Control of Nonlocal Pdes

In this talk, we introduce and analyze a new class of optimal control problems constrained by uncertain nonlocal PDEs. We utilize risk measures to formulate the resulting optimization problem. We develop a functional analytic framework, study the existence of solution and rigorously derive the first-order optimality conditions. Additionally, we employ a sample-based approximation for the uncertain variables and the finite element method to discretize in space. We prove the rate of convergence for the optimal risk neutral controls and conclude with several illustrative examples.

Harbir Antil
George Mason University
Fairfax, VA
hantil@gmu.edu

Drew P. Kouri
Optimization and Uncertainty Quantification
Sandia National Laboratories
dpkouri@sandia.gov

Johannes Pfefferer
Technische Universität München
Zentrum Mathematik
pfefferer@ma.tum.de

MS34

Chance Constrained Optimization in Banach Spaces with Application to PDE Systems

This work studies chance constrained optimization (CCOPT) problems on Banach-spaces. Typical examples are optimization problems governed by random operator equations or state-constrained partial differential equations with random data. Since such CCOPT problems are generally nonsmooth, nonconvex and intractable, methods of nonsmooth and variational analysis on Banach spaces are used to address these issues. As a result, two differentiable families of functions are introduced as smoothing approximations to the probability function of chance constraints. This leads to two smooth parametric optimization problems IA_τ and OA_τ, where the feasible sets of IA_τ are always subsets (inner approximation) and the feasible sets of OA_τ always supersets (outer approximation) of the feasible set of CCOPT, respectively. Furthermore, the following results are shown

- the feasible sets of the inner and outer approximations converge to the feasible set of CCOPT w.r.t. the approximation parameter τ
- estimations to the generalized subdifferential of the probability function are given though the derivatives of the smoothing approximation functions
- any limit-point of a sequence of optimal solutions of IA_τ (resp. OA_τ) is a stationary point of CCOPT

Finally, a numerical example on a chance constrained PDE optimization problems demonstrations the viability of the proposed approach.

Abebe Geletu
Technische Universitat Ilmenau
abebe.geletu@tu-ilmenau.de

Armin Hoffmann
TU-Ilmenau
Institute of Mathematics, 98693 Ilmenau, Germany
armin.hoffmann@tu-ilmenau.de

Pu Li
TU-Ilmenau
Institute of Automation and Systems Engineering
pu.li@tu-ilmenau.de

MS34

A Primal-dual Algorithm for Risk Minimization with Application to PDE-constrained Optimization

Many science and engineering applications necessitate the optimization of systems described by PDEs with uncertain inputs including noisy problem data and unknown boundary or initial conditions. One can formulate such problems as risk-averse optimization problems in Banach space, which upon discretization, become enormous risk-averse stochastic programs. For many popular risk models including the coherent risk measures, the resulting risk-averse objective function is non-smooth. This lack of differentiability complicates the numerical approximation of the objective function as well as the numerical solution of the optimization problem. To address these challenges, I present a general primal-dual algorithm for solving large-scale non-smooth risk-averse optimization problems. This algorithm is motivated by epigraphical regularization of risk measures. As a result, the algorithm solves a sequence of smooth optimization problems using Newton-type methods. I prove convergence of the algorithm even when the subproblem solves are performed inexactly and conclude with numerical examples demonstrating the efficiency of this method.

Drew P. Kouri
Optimization and Uncertainty Quantification
Sandia National Laboratories
dpkouri@sandia.gov

MS34

How to Supplement Risk Regulations to Avoid Industrial Catastrophes

Risk regulations/constraints in various engineering areas are based on Probability of Exceedance (POE). E.g., in Material Design, it is the probability that the load will exceed the material strength; in Nuclear Safety, it is probability that release of radioactive materials exceeds some threshold; and in Finance, the probability that liabilities will exceed assets. Although POE is quite popular, it has a significant conceptual drawback: it does not take into account magnitude of outcomes exceeding the threshold. Therefore, large industrial catastrophes, such as Fukushima nuclear accident, do not lead to a violation of safety requirements. Recently developed risk function, called Buffered Probability of Exceedance (bPOE) takes into account the magnitude of tail outcomes of the distribution. The presentation explains how to design safety regulations/constraints with bPOE to control large low probability outcomes.

Stan Uryasev, Giorgi Pertaia
University of Florida
uryasev@ufl.edu, gpertaia@ufl.edu

MS35

Deep Inversion - Latent Space Analysis of Autoencoders for Inverse Problems

Data-driven deep learning methods have revolutionized many application fields in computer vision and speech recognition. Recently, learned variational methods, e.g. learned proximal point methods and learned primal-dual optimization algorithms, have enabled deep convolutional neural networks (CNN) to tackle more complicated inversion tasks. However, in the context of basic CNN inversion methods, one fundamental aspect of inverse problems theory is still largely missing: regularization theory addressing ill-posedness and uncertainty quantification. In this talk, we present a latent space analysis of Autoencoding networks (VAEs, GANs) which offers new mathematical tools and insights for addressing this limitation. A synthetic deconvolution problem and photoacoustic tomography inversion illustrate our theoretical findings.

Yoeri Boink
University of Twente
y.e.boink@utwente.nl

Stephan A. Van Gils
Department of Applied Mathematics
University of Twente
s.a.vangils@math.utwente.nl

Srirang Manohar
Biomedical Photonic Imaging
University of Twente
s.manohar@utwente.nl

Christoph Brune
Department of Applied Mathematics
University of Twente
c.brune@utwente.nl

MS35

A Mean-field Optimal Control Formulation of Deep Learning

Recent work linking deep neural networks and dynamical systems opened up new avenues to analyze deep learning. In particular, it is observed that new insights can be obtained by recasting deep learning as an optimal control problem on difference or differential equations. However, the mathematical aspects of such a formulation have not been systematically explored. This work introduces the mathematical formulation of the population risk minimization problem in deep learning as a mean-field optimal control problem. Mirroring the development of classical optimal control, we state and prove optimality conditions of both the Hamilton-Jacobi-Bellman type and the Pontryagin type. These mean-field results reflect the probabilistic nature of the learning problem. In addition, by appealing to the mean-field Pontryagin's maximum principle, we establish some quantitative relationships between population and empirical learning problems. This serves to establish a mathematical foundation for investigating the algorithmic and theoretical connections between optimal control and deep learning.

Jiequn Han
Princeton University

jiequnh@princeton.edu

MS35

Structured Models for Convolutional Neural Networks

Abstract not available.

Xiaojin Tan
University of British Columbia
xtan@eoas.ubc.ca

Eldad Haber
Department of Mathematics
The University of British Columbia
haber@eoas.ubc.ca

MS35

Extensions and Algorithms for Neural Net Training

Neural Net (NN) is considered the most powerful tool in modern machine learning and data processing tasks. In this work, we incorporate a robust "trimming' mechanism into NN training. While fitting the model, trimming variables automatically pick out the outliers in the data. We use trimming-convolutional-neural-network (TCNN) as an example to classify the CFAIR10 data set and illustrate the trimming mechanism.

Peng Zheng
University of Washington
Applied Mathematics Department
zhengp@uw.edu

Aleksandr Aravkin
University of Washington
saravkin@uw.edu

MS36

Software Engineering Guidelines for Scientists - A Practical Handout for the Developing Researcher

In research facilities, scientists often develop software. Most of them do not have any specific education in software development. Usually they had programming courses at university or they self-taught some programming skills. Therefore their knowledge about software engineering and adjacent topics is quite limited. To support scientists, we created a set of software engineering guidelines. These guidelines give advice in different fields of software development (e.g., requirements management, design and implementation, change management). To make it easy to start with them, we developed a simple classification scheme taking aspects into account like expected software size or software lifetime. This scheme is useful to filter the guidelines and to fit them to the right context. Besides providing written guidelines and explanations, we created check list in different formats (e.g., Markdown, Word) to offer scientist a light-weight and easy-to-use tool. In this talk, we provide an overview about the concept of the guidelines and report about experiences introducing them at the German aerospace center (DLR) a large research facility in Germany. At DLR around 2000 to 3000 persons develop software in part or full time.

Carina S. Haupt
DLR, Germany
carina.haupt@dlr.de

MS36

Role of Requirements in Scientific Software

When starting a new software project or enhancing preexisting software through the addition of new features, a common and necessary first step is to consider what functionality needs to be implemented and what constraints are imposed upon any possible implementation of the functionality. While this step likely occurs at varying levels of detail and formality as part of a normal workflow, this process and the results of the process could be improved by adopting ideas from requirements engineering or software specification. For the purpose of establishing a general understanding of definitions as well as possible benefits of generating requirements, we will discuss requirements and the different phases of requirement engineering as related to developing and maintaining CSE software. We will compare and contrast the creation of requirements at different levels of formality and relate requirements engineering to other elements of software development such as documentation and verification.

Jared O'Neal
Argonne National Laboratory
joneal@anl.gov

MS36

Repository Analysis of Open-source and Scientific Software Development Projects

Software development processes in scientific computing have not been studied extensively, with most software engineering research focusing on industry or other open-source projects. In particular, understanding individual developers' productivity, as well as overall project health, remains a challenging task. We consider a number of metrics that can help characterize typical and anomalous project development patterns by analyzing each project's repository data collected over the lifetime of the project. Our goal is to produce insights into what is working well for a particular project, as well as to enable detection of potential problems before they can negatively impact the success of a project. To illustrate these analyses, we present our studies of several collaborative HPC software development efforts and other open-source projects.

Boyana Norris
University of Oregon
norris@cs.uoregon.edu

Kanika Sood
University Of Oregon
1991
kanikas@cs.uoregon.edu

Anshu Dubey
Argonne National Laboratory
adubey@anl.gov

Lois Curfman McInnes
Mathematics and Computer Science Division
Argonne National Laboratory

curfman@mcs.anl.gov

MS36

Symbiosis between Software and Scientific Career

Abstract not available.

Matthew J. Turk
NCSA, UIUC
matthewturk@gmail.com or mjturk@illinois.edu

MS37

On theService Capacity Region of Content Access from Coded Storage

Abstract not available.

Sarah Anderson
University of St. Thomas
ande1298@stthomas.edu

MS37

Fast and Efficient Distributed Matrix-vector Multiplication using Rateless Fountain Codes

Large-scale machine learning and data mining applications require computer systems to perform massive computations that need to be parallelized across multiple nodes, for example, massive matrix-vector and matrix-matrix multiplication. The presence of straggling nodes – computing nodes that unpredictably slowdown or fail – is a major bottleneck in such distributed computations. We propose a rateless fountain coding strategy to alleviate the problem of stragglers in distributed matrix-vector multiplication. Our algorithm generates linear combinations of the m rows of the matrix and assigns them to different worker nodes, which then perform row-vector products with the encoded rows. The original matrix-vector product can be decoded as soon as slightly more than m row-vector products are collectively completed by the nodes. This strategy enables fast nodes to steal work from slow nodes, without requiring the knowledge of node speeds. Compared to recently proposed fixed-rate erasure coding strategies which ignore partial work done by straggling nodes, rateless coding achieves significantly lower overall delay, as well as small computational overhead.

Gauri Joshi
Carnegie Mellon University
gaurij@andrew.cmu.edu

MS37

Data Analysis and Automated Sleep Apnea Diagnosis

Obstructive sleep apnea (OSA) affects a growing number of patients with potentially fatal results. Diagnosis involves an expensive and time-consuming battery of tests, including a lengthy survey, scans of the patient airway, and a sleep study where multiple measurements are taken over several hours. In other words, it is a multifactorial disorder, requiring many different types of data, including DNA sequences, multiple time series, metabolites, and geometric analysis of airways. Diagnosis is currently done by hand by trained technicians. In this talk, we present preliminary methods for automating diagnosis that draw from statistics, machine learning, computational topology, and geometry to analyze survey data, time series measurements, and airway obstruction. The results from these indep dent analyses provide meaningful results for clinical stuc and can be used as guidance for clinical practitioners, b in determining the relevant measurements for diagnosis well as for the diagnosis itself.

Kathryn Leonard
Occidental College
kleonard.ci@gmail.com

MS37

Automated Obstructive Sleep Apnea Diagnosis through High-dimensional Multi-Source Data Analysis

Abstract not available.

Xu Wang
Washington and Lee University
xwang@wlu.ca

MS39

Practical Computation of the Noncommutative Kubo Formula for the Conductivity in Incommensurate Multilayers

Weak van der Waals interactions between 2D materials l ers do not impose limitations on integrating highly d parate materials such as graphene, phosphorene and ma others. This is both a blessing, allowing the realization many more stable assemblies, and a curse from a moc ing perspective due to the loss of periodicity. In this ta we discuss how abstract C*-algebras formulations of tig binding electronic structure models, introduced by Be sard, can be implemented directly on computers lead to new algorithms for computing Kubo formulae asso ated with macroscopic observables such as the conduct ity. This allows to perform calculations outside the sc of usual methods involving finite-volume periodic appro mants.

Paul Cazeaux
Department of Mathematics
University of Kansas
pcazeaux@ku.edu

MS39

Exploiting Sparsity in the Chebyshev Coefficients for Faster Conductivity Calculations

The Fermi Operator Expansion is a well-established te nique for replacing a sum over the eigenstates of a spa Hamiltonian operator $H \in \mathbb{R}^{n\times n}$ with a trace $\mathrm{Tr}\big(p(H$ of a polynomial in H such that the computational c reduces from $\mathcal{O}\big(n^3\big)$ for computing the eigenstates $\mathcal{O}\big(n\,\deg(p)\big)$ for evaluating the matrix polynomial. T conductivity of a system of electrons is given by a s over pairs of eigenstates, hence the Fermi Operator l pansion takes the form of a bivariate matrix polynom $p(H,H)$ and a naive analysis suggests that the costs sc as $\mathcal{O}\big(n\,\deg(p)^2\big)$ with $\deg(p)$ the one-dimensional poly mial degree. We prove that after a suitable modificat by one-dimensional polynomials, the bivariate polynom $p(E,E')$ has at most $\mathcal{O}\big(\deg(p)^{3/2}\big)$ significant coefficie and present numerical evidence that the number of s nificant coefficients is in fact only $\mathcal{O}\big(\deg(p)^{6/5}\big)$. Th observations allow to reduce the runtime of conductiv calculations by at least two orders of magnitude and s

nificantly increase the range of computationally tractable physical parameters. This is joint work with Stephen Carr, Mitchell Luskin, Daniel Massatt and Christoph Ortner.

Simon Etter
University of Warwick
s.etter@warwick.ac.uk

MS40

The s-Step Conjugate Gradient Method in Finite Precision

While synchronization-reducing variants of Krylov subspace methods can improve per-iteration performance, they do come with a cost; it is well-known that finite precision computation can cause delayed convergence and decreased attainable accuracy in standard Krylov subspace methods, and even slight modifications to the standard recurrences can amplify these effects. This talk will focus on s-step variants of the conjugate gradient and Lanczos methods and present theoretical and experimental results regarding their finite precision behavior. Based on the analysis, we discuss potential remedies for improving accuracy while still maintaining desirable synchronization-reducing properties.

Erin C. Carson
New York University
carson@karlin.mff.cuni.cz

MS40

Numerical Analysis of the Maximal Attainable Accuracy in Communication-hiding Pipelined Conjugate Gradients

On massively parallel hardware the performance of Krylov subspace methods (KSMs) is typically limited by communication latency rather than floating point performance. With HPC hardware advancing towards the exascale regime the need for scalable alternatives to traditional KSMs becomes ever more apparent. One such alternative is the class of pipelined KSMs, which overlap global reduction phases with local arithmetic operations, thus 'hiding' communication latency behind computations. To obtain the overlap the traditional KSM algorithm is reformulated by introducing auxiliary vector quantities which are computed recursively. Although pipelined KSMs are equivalent to traditional KSMs in exact arithmetic, local rounding errors induced by the multi-term recurrence relations in finite precision may in practice affect convergence significantly. In this talk we present a numerical stability study that characterizes the possibly detrimental propagation of local rounding errors in various pipelined versions of the Conjugate Gradient method. Expressions for the gaps between the true and recursively computed variables that are used to update the search directions are derived. It is shown how these results can be used to analyze and correct the effect of local rounding error propagation on the maximal attainable accuracy of pipelined CG methods. The theoretical analysis is supplemented by various numerical and performance experiments.

Siegfried Cools, Jeffrey Cornelis
University of Antwerp
siegfried.cools@uantwerp.be, jeffrey.cornelis@uantwerp.be

Emmanuel Agullo
INRIA
emmanuel.agullo@inria.fr

Emrullah Fatih Yetkin
Istanbul Technical University, Informatics Institute
fatih.yetkin@khas.edu.tr

Luc Giraud
Inria
luc.giraud@inria.fr

Wim Vanroose
Antwerp University
wim.vanroose@uantwerpen.be

MS40

On the Convergence Rate of Different Variants of the Conjugate Gradient Algorithm

Several variants of the conjugate gradient algorithm have been proposed for computational purposes. Although equivalent in exact arithmetic, these variants behave differently in finite precision arithmetic, both in the maximum attainable accuracy and in the rate of convergence. We focus on the latter and demonstrate numerically how the variants converge differently on a wide range of test problems. For small problems where the eigenvalues of the coefficient matrix A are evenly distributed, these variants converge in similar rates. However in some extreme eigenvalue cases and in particular for large problems, some implementations require far more iterations than others to reach a certain level of accuracy. The delayed convergence could result in a longer running time even if the variant reduces the computational time per iteration. We, therefore, wish to understand the reasons behind the delayed convergence for the studied variants and to propose conditions that could guarantee a reasonable convergence rate, which would aid in the future design of new algorithms.

Hexuan Liu
Applied Mathematics
University of Washington
hl65@uw.edu

MS40

Exploiting Multiprecision in Krylov Subspace Methods

Variable floating-point arithmetic precision (beyond IEEE single and double) is becoming increasingly available to users. We show how this can be exploited in Krylov subspace methods for solving linear systems of equations. Specifically, we show how to use inexact matrix-vector products and inexact reorthogonalization in the Arnoldi algorithm without affecting the convergence rate of MGS-GMRES. We discuss extensions to the method of conjugate gradients. This is joint work with Serge Gratton and Ehouarn Simon (INPT-ENSEEIHT) and Philippe Toint (Namur).

David Titley-Peloquin
McGill University
david.titley-peloquin@mcgill.ca

MS41

Scalable Hierarchical Matrix Methods in Machine Learning

I will give an overview of different ways to exploit advances in H-matrix methods in problems in machine learning. In particular, I will discuss H-matrix methods in terms of ker-

nel methods, covariance approximation, and Hessian operator approximation. First, I will review results in scalable approximation of dense matrices using H-matrix ideas. Then, I will briefly review kernel regression, kernel clustering, and Hessian approximation for simple feed-forward neural networks. I will conclude with numerical examples.

George Biros
The Institute for Computational Engineering and Sciences
The University of Texas at Austin
gbiros@acm.org

MS41

Randomized Sparse PCA using the Variable Projection Method

Many fields from genomics to environmental science must deal with large collections of correlated (redundant) data in the form of an array. Dimensionality reduction is commonly used to summarize the data and reveal any underlying coherent structure. Principal component analysis (PCA) is the most popular linear method for this task, however, it can be difficult to interpret the principal components (PC) in a high-dimensional data setting. Sparse PCA (SPCA) tries to overcome this issue and provides a more parsimonious summary of the data. More concretely, SPCA aims to form the PCs as linear weighted combinations of only a few of the original variables. In this talk we revisit the seminal work by Zou, Hastie and Tibshirani who formulated SPCA as a regularized regression-type problem. While their algorithm is popular, the shortcoming is that their approach does not scale to big data. To ease the computational demands, we use the variable projection method which was developed by Golub and Pereyra in the 70's. Following this powerful paradigm to minimize two sets of variables, we recast the original problem as a value-function optimization problem. Next, we show how randomized methods for linear algebra can be leveraged to extend the approach to the large-scale data setting. The proposed algorithms are applied to both synthetic and real world data, showing exceptional computational efficiency and diagnostic performance.

N. Benjamin Erichson
University of California, Berkeley
erichson@berkeley.edu

MS41

Randomized Algorithms to Compute Low-rank Approximation for Data Analysis and Machine Learning Applications

In this talk we discuss the potential of the randomized algorithm to compute low-rank approximation, focusing specifically on data analysis and machine learning applications.

Ichitaro Yamazaki
University of Tennessee, Knoxville
iyamazak@icl.utk.edu

Yuechao Lu
Osaka University
yc-lu@ist.osaka-u.ac.jp

Michael Mahoney
UC Berkeley
mmahoney@stat.berkeley.edu

Jack J. Dongarra
University of Tennessee, Oak Ridge National Laboratory, USA
dongarra@icl.utk.edu

MS41

Online Optimization for Continuous Hyperparameters with Applications to Traffic Flow Prediction

Hyperparameter tuning is a critical step to obtain good predictive performance of many machine learning models. The computation cost of traditional grid search grows exponentially with the dimension of hyperparameters. In applications where the data distribution is not stationary, one may need to perform hyperparameter tuning and update the model periodically. In this work, we first describe hyperparameter tuning as a bi-level optimization problem, then present a principle way to adjust the continuous hyperparameters of a model based on online projected hypergradient descent. The theoretical properties of the algorithm are studied using local regret analysis. The performance of the algorithm is demonstrated on both synthetic data and real traffic flow data.

Hongyuan Zhan
Pennsylvania State University
hongyuan.zhan@gmail.com

MS42

EnKF-Based Estimators for Aerodynamics

Ensemble Kalman Filters (EnKF) are efficient estimators for high-dimensional systems arising in computational fluid dynamics. We demonstrate their efficacy in relatively simple aerodynamic flows over two-dimensional airfoils at low Reynolds numbers, computed using a twin-experiment strategy with the incompressible Navier-Stokes equations. Using conventional techniques such as covariance inflation and state augmentation, the EnKF is able to accurately track the flow state by assimilating a few surface pressure measurements, even in the presence of uncertain inputs including time-varying gusts. We consider strategies for estimating bias errors associated with coarse resolution in the estimator. In particular, we represent the truncation error with a low-rank spatial basis and a colored-noise process to represent the dynamics; estimator accuracy is increased considerably compared to bias-blind strategies. We also adapt the EnKF to reduced-order estimators based on regularized vortex elements, and show how such models can be made more accurate, and less fragile, by estimating the leading-edge suction parameter online rather than using a fixed value.

Tim Colonius, Andre da Silva
Division of Engineering and Applied Science
California Institute of Technology
colonius@caltech.edu, andrefcs@caltech.edu

Darwin Darakananda, Jeff D. Eldredge
University of California, Los Angeles
Mechanical & Aerospace Engineering
darwindarak@ucla.edu, eldredge@seas.ucla.edu

MS42

Ensemble Kalman Inversion: A Derivative-free Technique for Machine Learning Tasks

The standard probabilistic perspective on machine learn-

ing gives rise to empirical risk-minimization tasks that are frequently solved by stochastic gradient descent (SGD) and variants thereof. We present a formulation of these tasks as classical inverse or filtering problems and, furthermore, we propose an efficient, gradient-free algorithm for finding a solution to these problems using ensemble Kalman inversion (EKI). Applications of our approach include offline and online supervised learning with deep neural networks, as well as graph-based semi-supervised learning. The essence of the EKI procedure is an ensemble based approximate gradient descent in which derivatives are replaced by differences from within the ensemble. We suggest several modifications to the basic method, derived from empirically successful heuristics developed in the context of SGD. Numerical results demonstrate wide applicability and robustness of the proposed algorithm.

Nikola Kovachki
California Institute of Technology
nkovachki@caltech.edu

Andrew Stuart
Computing + Mathematical Sciences
California Institute of Technology
astuart@caltech.edu

MS42

Interdisciplinary Data Assimilation and Bayesian Learning

Extracting all of the pertinent information from observations is critical for many interdisciplinary scientific research and societal applications. This is especially true for coupled physical-acoustical and physical-biogeochemical ocean dynamical systems. A theory for the nonlinear non-Gaussian estimation of the state, parameters, parameterizations, constitutive relations, and model formulations of such systems is developed. Bayesian and deep machine learning methods are combined and results are showcased in a range of multi-physics applications, including medical ones. This is work in collaborations with our MSEAS group at MIT.

Pierre F. Lermusiaux
MIT
pierrel@mit.edu

MS42

Greedy Sensor Placement for Controlling High-dimensional Dynamics

Optimal sensor and actuator placement is one of the foremost challenges in the estimation and control of high-dimensional complex systems. Often it is impractical to monitor or actuate every state, or perform brute-force combinatorial searches among all possible placements, and existing convex optimization and gradient descent methods do not scale to high-dimensional gridded domains in climate and fluid dynamics. We propose an efficient placement method based on the pivoted matrix QR factorization of a low-rank representation, for which the runtime and number of sensors scale with the intrinsic rank of the system. We show that this method generalizes to different objective functions suitable for different dimensionality reduction methods and downstream tasks. The benefits of this approach are demonstrated on a variety of datasets in manufacturing, climate, and fluid dynamics. Furthermore, we generalize this method to actuator placement for optimal control.

Krithika Manohar
University of Washington
Dept of Applied Mathematics
kmanohar@uw.edu

J. Nathan Kutz
University of Washington, Seattle
Dept of Applied Mathematics
kutz@uw.edu

Steven Brunton
University of Washington
sbrunton@uw.edu

MS43

In-situ Data Visualization Approaches in the SU2 Analysis and Design Framework

The size and number of datasets analyzed by post-processing and visualization tools are rapidly outpacing our ability to extract useful information from complex simulations. At the same time, the disk-read data transfer rate is only doubling every 36 months and is destined to be the bottleneck for traditional post-processing architectures. To eliminate this bottleneck, a sub-zone load-on-demand (SZL) visualization architecture has been developed which only loads the data needed to create the desired plot. For datasets on the order of a billion cells this can reduce the amount of data loaded by 2 orders of magnitude. The SU2 open-source framework has been modified to write SZL files including the indexing required to allow selective loading of subzones. The performance of SU2 and Tecplot while writing/loading these SZL files will be presented for large unsteady CFD simulations carried out with both the SU2 DG-FEM and finite-volume solvers. The performance of the implementation on two separate test cases will be presented. The same I/O bottleneck also affects the writing of data from CFD codes. Recently, Tecplot has extended the SZL technique to write only those subzones needed to create the desired set of plots. The new technique is a compromise between in-itu visualization techniques and traditional full-volume data file export from the CFD code. During recent tests with slices and isosurfaces, file sizes are reduced to less than 0.8% of the full volume dataset.

Scott Imlay
Tecplot, Inc.
s.imlay@tecplot.com

Juan J. Alonso
Department of Aeronautics and Astronautics
Stanford University
jjalonso@stanford.edu

Thomas D. Economon
Bosch
thomas.economon@us.bosch.com

MS43

Siac Filtering for Visualization and Feature Detection

As high-order continuous and discontinuous Galerkin methods continue their proliferation into various applied engineering disciplines, it is important that the visualization techniques and corresponding data analysis tools that

act on the results produced by these methods faithfully represent the underlying data. To state this in another way: the interpretation of data generated by simulation needs to be consistent with the numerical schemes that underpin the specific solver technology. As the verifiable visualization literature has demonstrated: visual artifacts produced by the introduction of either explicit or implicit data transformations, such as data resampling, can sometimes distort or even obfuscate key scientific features in the data. In this talk, we focus on the handling of elemental continuity, which is often only C0 continuous or piecewise discontinuous, when visualizing or extracting features from primary or derived fields from high-order data. We propose using the line-SIAC filter as a way of handling elemental continuity issues in an accuracy-conserving manner with the added benefit of casting the data in a smooth context even if the representation is element discontinuous. This presentation based upon work done in collaboration with Bob Haimes (MIT), Julia Docampo-Sanchez (MIT) and Jennifer Ryan (U. East Anglia).

Mike Kirby, Ashok Jallepalli
University of Utah
School of Computing
kirby@cs.utah.edu, ash.nani@gmail.com

MS43

Advanced in Situ Analysis of High Order Dg Solutions with Application to Scale-resolving Simulations of Complex Turbulent Flows

With the rise of petascale and soon exascale simulations, the amount of data generated by high-fidelity simulations will become impossible to analyse a posteriori. In the meantime, high-order (HO) methods have recently gained considerable attention in both the research community and the industry. In this context, Cenaero has developed several capabilities to perform in-situ visualization and analysis of data generated by its massively parallel HO discontinuous Galerkin CFD code Argo. The first capability is related to in-situ visualization of HO solutions. The adopted strategy is based on an interface to the Catalyst library from Kitware. This interface is a direct extension of a ParaView plugin developed specifically for the visualization of HO solutions under the Gmsh format. The second capability concerns the injection of realistic turbulence inflow conditions for turbomachinery applications through co-simulation of the main domain (e.g. blade profile) with traditional configurations in fluid dynamics such as homogeneous isotropic turbulence and channel flows. In order to pilot the co-simulation, in-situ analysis of the main characteristics of the turbulence must be implemented, which includes advanced statistics such as spatial correlations, integral length scales and energy spectra. These new capabilities will be demonstrated at scale with the simulation of a low pressure turbine blade cascade.

Michel Rasquin
University of Colorado Boulder
michel.rasquin@cenaero.be

Koen Hillewaert
Cenaero
Cenaero
koen.hillewaert@cenaero.be

MS43

Online Compression of High-order CFD Solutions using Machine Learning

Unsteady computational fluid dynamics (CFD) simulations have the potential to generate petabytes of solution data; well beyond what can be handled by current generation I/O and storage sub-systems. As such, it is typically not practical to write out a sequence of finely spaced solution snapshots. Instead, the user must instrument the simulation in advance. However, this requires the user to have some a priori knowledge about the dynamics and evolution of the system—something which negates many of the advantages inherent to simulation. In this talk I will show how, within the context of high-order methods, a combination of convolutional autoencoders, principal component analysis, and interpolation can be employed to perform online lossy compression of unsteady CFD simulations. Such compression has the potential to enable routine snapshotting of even the largest simulations, thus facilitating far richer far richer a posteriori analysis. Numerical experiments showcasing a five-fold reduction in archival storage requirements will be presented and the resulting flow fields will be compared and contrasted against the reference data.

Freddie Witherden
Department of Aeronautics and Astronautics
Stanford University
freddie.witherden08@imperial.ac.uk

Jeremy Morton
Department of Aeronautics & Astronautics
Stanford University
jmorton2@stanford.edu

Liqian Peng, Kevin T. Carlberg
Sandia National Laboratories
pengliqian@gmail.com, ktcarlb@sandia.gov

Mykel Kochenderfer
Stanford University
mykel@stanford.edu

Antony Jameson
Professor, Department of Aeronautics & Astronautics
Stanford University
jameson@baboon.stanford.edu

MS44

A Multi-Level Optimization Based Monte-Carlo Sampler for Large Scale Inverse Problems

The Markov Chain Monte Carlo (MCMC) method is one of the pillars of Bayesian inverse problems. However, this approach typically faces several challenges in large-scale inverse problems: classical MCMC algorithms rely on constructing a sequential Markov chain, which makes it hard to fully parallelise; it is often challenging to derive efficient transition kernels; and simulating the Markov chain can be computationally costly, as the posterior density evaluation involves expensive forward model solves. We present an integrated approach based on the multilevel Monte Carlo method and the randomise-then-optimise (RTO) method to address these challenges. The use of RTO allows us to derive efficient and parallelisable MCMC or importance sampling estimators for solving inverse problem. With the help of the multilevel Monte Carlo, we can further accelerate RTO and reduce the variance of resulting estimators. We will demonstrate the efficacy of our approach on inverse

problems governed by PDE and ODE.

Chuntao Chen
Monash University
chuntao.chen@monash.edu

MS44

Large-scale Uncertainty Propagation via Overlapping Domain Decomposition

For large-scale decomposable systems (large number of coupled subsystems), high-fidelity simulation and uncertainty quantification (UQ) are typically only achievable at the subsystem level. This is due to long full-system simulation times and challenges in integrating subsystem models characterized by different physics/scales. Even full-system surrogate models, which are typically employed to facilitate UQ, rely on intractable full-system simulations. We propose a novel UQ framework for decomposable systems that (1) performs UQ at the subsystem level, and (2) enforces compatibility in the uncertainties at inter-subsystem interfaces. To enforce inter-subsystem compatibility, we adopt domain decomposition (DD) methods as they enable full-system analysis while relying on subsystem-level simulations. In the DD context for UQ, the inter-subsystem coupling corresponds to random variables, represented with polynomial chaos expansions (PCEs). We define inter-subsystem compatibility as equivalence of the PCEs of interface unknowns, enforced using fixed-point iterative methods (e.g., multiplicative Schwarz) from overlapping DD. The method comprises of (1) an inner loop performing subsystem UQ (e.g., non-intrusive spectral projection), and (2) an outer loop enforcing inter-subsystem compatibility. We apply the framework to nonlinear stochastic heat equations, with boundary condition and parametric uncertainties, while examining convergence and scalability of the proposed solver.

Kevin T. Carlberg
Sandia National Laboratories
ktcarlb@sandia.gov

Sofia Guzzetti
Department of Mathematics and Computer Science
Emory University
sofia.guzzetti@emory.edu

Mohammad Khalil, Khachik Sargsyan
Sandia National Laboratories
mkhalil@sandia.gov, ksargsy@sandia.gov

MS44

Stability-preserving Model Order Reduction for Random Dynamical Systems

We consider linear dynamical systems of ordinary differential equations or differential-algebraic equations including random parameters. The unknown random processes can be expanded into a series with orthogonal basis polynomials due to the concept of the polynomial chaos. The stochastic Galerkin method yields a large coupled linear dynamical system, whose solution represents an approximation of the unknown coefficient functions in the expansion. We apply projection-based model order reduction to this large dynamical system. However, a reduced-order model may be unstable, even though the original systems as well as the stochastic Galerkin system are asymptotically stable. Transformations to dissipative systems guarantee a preservation of stability in the reduced-order models. Lyapunov equations yield the required transformation matrices. We investigate and compare two approaches: the original systems are transformed followed by a Galerkin-projection or, vice versa, the Galerkin-projected system is transformed. Parameter-dependent Lyapunov equations or a single high-dimensional Lyapunov equation have to be solved, respectively. We present results of numerical computations for test examples.

Roland Pulch
University of Greifswald
pulchr@uni-greifswald.de

MS44

Adaptive Construction of Quadrature Rules for Bayesian Inference and Prediction

Bayesian inverse problems remain a challenging topic due to the large number of samples that are necessary to resolve the posterior. In this talk, a new approach is proposed to tackle this problem that consists of two ingredients: an adaptive quadrature rule and interpolation of the posterior. The novel quadrature rule is constructed using samples from a distribution, has positive weights, and can be constructed such that a nested sequence of rules is obtained. By enforcing that the rule has a high degree, convergence for sufficiently smooth functions is guaranteed. Moreover the rule is numerically stable and correctly predicts that positive integrands have a positive valued integral. By interpolation of the posterior an approximate distribution is obtained, which is used to efficiently draw samples and obtain a new (and more accurate) quadrature rule. This iterative process converges to the true posterior. We discuss the conditions that are necessary for interpolation of the posterior and present some numerical examples.

Laurent van den Bos, Benjamin Sanderse
Centrum Wiskunde & Informatica (CWI)
Amsterdam, the Netherlands
l.m.m.van.den.bos@cwi.nl, b.sanderse@cwi.nl

MS45

DeepHyper: Scalable Asynchronous Hyperparameter Search for Deep Learning

Hyperparameters employed by deep learning methods play a substantial role in the performance and reliability of these methods in practice. Hyperparameter search methods typically do not target scalability within a highly parallel machine, portability across different machines, experimental comparison between different methods, and tighter integration with workflow systems. We present DeepHyper, a Python asynchronous hyperparameter hyperparameter search package that provides a common interface for the implementation and study of scalable hyperparameter search methods. We evaluate the efficacy of the asynchronous hyperparameter search relative to approaches such as random search, genetic algorithms, Bayesian optimization, and hyperband on DL benchmarks on CPU- and GPU-based HPC systems.

Prasanna Balaprakash
Argonne National Laboratory
pbalapra@mcs.anl.gov

MS45

Data-Driven Multi-scale Decompositions for Fore-

casting and Model Discovery

Many physical systems of practical interest exhibit simultaneous dynamics on highly disparate time scales. Deconstructing such systems according to this property can offer valuable insight: behavior on two very different time scales can often be modeled separately, with only a limited coupling between them. In this work I will present a data-driven approach to decomposing time series data with multiscale properties which could serve as a valuable precursor to a variety of scale-separated analysis tasks. This method makes use of Dynamic Mode Decomposition (DMD), which identifies spatial and temporal coherencies to approximate sample data with a linear superposition of complex exponentials. DMD is applied on a sliding window over the input data, and the resultant frequency spectra are then clustered to identify dominant time scale components. Scale-separated reconstructions of the original signal are produced simply by summing over each cluster of modes separately. Results are presented for a simple polynomial toy model and for a system of three-body planetary motion. To conclude I will briefly discuss how this method might be used in model discovery and forecasting for multiscale systems.

Daniel Dylewsky
Applied Mathematics
University of Washington
dylewsky@uw.edu

J. Nathan Kutz
University of Washington, Seattle
Dept of Applied Mathematics
kutz@uw.edu

Molei Tao
School of Mathematics, Georgia Institute of Technology
mtao@gatech.edu

MS45

Data-driven Modeling and Analysis of Non-stationary Fluid Flows

Dynamic mode decomposition (DMD) is an algorithm that decomposes a time series of data into spatiotemporal coherent modes associated with dynamic growth/decay and oscillations in time. To date, DMD has been successful applied to analyze the dynamics of various stationary fluid flows. Here we extend the DMD framework for the analysis of non-stationary flows by leveraging the Gabor transform from time-frequency analysis. The resulting DMD spectrogram is demonstrated using various nonlinear systems of fluid flows that undergo intermittent transitions among distinct dynamical regimes, including droplet formation from a nozzle and oscillations of spherical bubbles. The results indicate that the DMD spectrogram successfully captures the non-stationary dynamics of the flows.

Supported by the eScience Institute and the Air Force Office of Scientific Research (AFOSR FA9550-16-1-0650 and FA9550-18-1-0200).

Kazuki Maeda, Steven Brunton
University of Washington
maeda@uw.edu, sbrunton@uw.edu

MS45

PDE-inspired Deep Neural Networks for Machine Learning and Inverse Problems

This talk bridges the gap between partial differential equations (PDEs) and neural networks and shows two possible uses. First, it presents a new mathematical paradigm that simplifies designing, training, and analyzing deep neural networks. It shows that training deep neural networks can be cast as a dynamic optimal control problem similar to path-planning and optimal mass transport. Second, it shows how neural networks can be used to improve inverse problems such as PDE parameter estimation.

Lars Ruthotto
Department of Mathematics and Computer Science
Emory University
lruthotto@emory.edu

MS46

Dynamic Algorithms with PaRSEC

Next generation of HPC platforms presents a variety of challenges, including a growing need for asynchrony, increased hardware diversity, performance non-uniformity, and decreased hardware reliability. These threaten the maintainability of increasingly complex scientific code bases and questions the continued viability of message passing as a model for direct, application-level interactions. This talk presents an extension to the PaRSEC task-based runtime, providing users with a Domain Specific Language allowing the dynamic description of dataflow application. This approach singularly improves the performance portability and maintainability of codes. Qualitative data about the new programming model and some of it's applications as well as quantitative performance data will be presented.

George Bosilca
University of Tennessee - Knoxville
bosilca@icl.utk.edu

Thomas Herault
Univ. of Tennessee - Knoxville
herault@eecs.utk.edu

Aurélien Bouteiller
University of Tennessee, Knoxville, US
bouteill@icl.utk.edu

Damien Genet
University of Tennessee, Knoxville
dgenet@icl.utk.edu

Jack J. Dongarra
University of Tennessee, Oak Ridge National Laboratory, USA
dongarra@icl.utk.edu

MS46

SPETABARU: A Task-based Runtime System with Speculative Execution Capability

While task-based programming models allow expressing the parallelism of algorithms finely, the traditional data accesses used in the sequential task-flow model (STF) can restrict the parallelism and hide useful information. In this presentation, we describe how more precise data accesses can be used to get better performance, and how uncertain modifications of the data by the tasks open the possibility for speculative execution. We detail different speculative execution models when this uncertainty exists. We also

introduce our speculative runtime system, SPETABARU, and provide examples with the parallelization of the Monte Carlo and replica exchange Monte Carlo simulations.

Berenger Bramas
INRIA
berenger.bramas@inria.fr

MS46

Performance and Implementation of UPC++

UPC++ is a C++ library that supports Partitioned Global Address Space (PGAS) programming. This presentation will describe a new implementation that differs substantially from the prior version of UPC++, and will discuss the reasons for our design decisions. We will also present new features, and show microbenchmark performance numbers demonstrating that one-sided remote memory access (RMA) in UPC++ is compet- itive with MPI-3 RMA. We will also showcase the benefits of UPC++ with irregular applications through a pair of application motifs, including a distributed hash table and a sparse solver component. UPC++s combination of low-overhead one-sided communication, remote procedure calls, and asynchronous API is the key to delivering high performance with irregular applications and enabling improved programmer productivity.

Steven Hofmeyr
Lawrence Berkeley National Laboratory
shofmeyr@lbl.gov

MS46

How to Adjust an OpenMP Runtime to Better Reach Hardware Performance on HPC Applications: libKOMPs Experiences?

Tasks are a good support for composition. During our development of HPC applications, we have experimented to manage parallelism using OpenMP tasks. Since version 4.0, the OpenMP standard proposes a model with dependent tasks that seems very attractive: it enables the description of dependencies between tasks generated by different parts of the application without breaking maintainability constraints of software development. In this talk we present how we have modified an existing OpenMP runtime (LLVM runtime) to better map the task parallelism to the real multi-cores NUMA architectures. This presentation will focus on three points: a better description of the parallelism by extending the 'task' construct to capture parallelism between tasks making concurrent updates of data; an extension of the task construct to describe scheduling affinity between tasks and data and its effective implementation in a scheduling algorithm; and finally n optimisation of tasks management by identifing important hot spots in the implentation of the LLVM runtime. We will illustrate our proposals by illustrating performances on classical dense matrix factorisation algorithms, and on applications coming from plasma physics. All experimental results have been performed using our modified OpenMP runtime libkomp `http://gitlab.inria.fr/openmp/libkomp`.

Thierry Gautier, Philippe Virouleau
Inria
thierry.gautier@inrialpes.fr, philippe.virouleau@inria.fr

MS47

Data Assimilation Subject to Practical Imperfections

Abstract not available.

Henry Huang
Pacific Northwest National Laboratory, USA
Zhenyu.Huang@pnnl.gov

MS47

An Improved Particle Filter and its Application for Tracking in a Wireless Sensor Network

We developed a novel particle filter algorithm, called drift homotopy likelihood bridging particle filter (DHLB-PF). The DHLB-PF is designed to surmount the degeneracy problem of generic particle filter framework by employing a Markov chain Monte Carlo (MCMC) procedure after the resampling step. DHLB-PF considers a sequence of pertinent stationary distributions which is designed based on dynamics of moving targets. These intermediate stationary distributions facilitate the MCMC step as well as explore the state space with higher degree of freedom. We evaluated the DHLB-PF algorithm using a target tracking problem in a wireless sensor network where no fusion center is required for data processing. The observations are collected only from the informative sensors, which are sensing only nearby moving targets. The detection of those informative sensors, which are typically a small portion of the sensor network, is performed by employing a sparsity-aware matrix decomposition technique. Simulation results demonstrated that the DHLB-PF outperforms current popular tracking algorithms.

Kai Kang
National Institutes of Health
kai.kang@nih.gov

MS47

Ensemble Kernel-based Learning for Analytic Continuation

In this talk, we introduce a novel data-driven learning method for the analytic continuation problem using ensemble kernel-based regression. The many-body quantum physics analytic continuation problem forms a standard inverse problem. Motivated by the development of machine learning methods, we apply the classic ensemble methods with kernel regression to this problem. The results show that our approach perform very well over the traditional regularization method.

Xuping Xie
Oak Ridge National Lab
xiex@ornl.gov

MS47

An Adaptive Multi-fidelity PC Based Iterative Ensemble Kalman Method for Inverse Problem

The polynomial chaos (PC) expansion has been widely used as a surrogate model in the Bayesian inference to speed up the Markov chain Monte Carlo (MCMC) calculations. However, the use of a PC surrogate introduces the modeling error, that may severely distort the estimate of the posterior distribution. This error can be corrected by increasing the order of the PC expansion, but the cost for building the surrogate may increase dramatically. In this talk, we seek to address this challenge by proposing an

adaptive procedure to construct a multi-fidelity PC surrogate. This new strategy combines (a large number of) low-fidelity surrogate model evaluations and (a small number of) high-fidelity model evaluations, yielding an adaptive multi-fidelity approach. Here the low-fidelity surrogate is chosen as the prior-based PC surrogate, while the high-fidelity model refers to the true forward model. The key idea is to construct and refine the multi-fidelity approach over a sequence of samples adaptively determined form data, so that the approximation can eventually concentrate to the posterior distribution. We illustrate the performance of the proposed strategy through two nonlinear inverse problems. It will be shown that the proposed adaptive multi-fidelity approach can improve significantly the accuracy, yet without a dramatic increase in the computational complexity. Applications of the adaptive idea in the ensemble Kalman methods will also be discussed.

Tao Zhou
Institute of Computational Mathematics
Chinese Academy of Sciences
tzhou@lsec.cc.ac.cn

MS48

Diversifying the Workforce at Lawrence Livermore National Laboratory

This information session will help attendees learn how LLNL builds STEM pipelines in computational science and engineering. Our efforts reflect striving for a diverse and broader population by understanding bias, creating strategic collaborations and building inclusive working environments for a more productive and engaged organization.

Tony Baylis
Lawrence Livermore National Laboratory
baylis3@llnl.gov

MS48

Creating and Sustaining a Research Network for Women in Mathematics of Materials

In this talk, I will talk about the process of creating and sustaining a research network for women in math of materials. In particular, I will discuss the organization of a research workshop and the ways to keep the network connected.

Malena I. Espanol
Department of Mathematics
The University of Akron
mespanol@uakron.edu

MS48

Effective Workplace Communication

Effective communication skills are vital in the professional workplace. In this session, I will begin by presenting five rules for professional communication. In a workplace, conflicts will inevitably occur. I have developed a four-step method for dealing with the inner turmoil that can arise from conflict. Combining these two techniques will help session attendees handle all workplace interactions with confidence.

Rebecca J. Hartman-Baker
Lawrence Berkeley National Laboratory
rjhartmanbaker@lbl.gov

MS48

Sustainable Research Pathways: Building Connections Across Communities to Diversify the National Laboratory Workforce

The Sustainable Research Pathways (SRP) program is a partnership between the Computing Sciences Division of the Lawrence Berkeley National Laboratory, a Department of Energy National Laboratory, and Sustainable Horizons Institute, a nonprofit organization. SRP aims to create research opportunities for students and faculty from underrepresented, low-income, and first-generation communities that lead to long-term, fruitful relationships and research collaborations with DOE Laboratory researchers. To initiate and realize the full potential of these relationships the program organizes an annual matching workshop followed by summer internships at the laboratory packed with research and educational activities focused on computational science and high-performance computing. Visiting faculty and students are recruited from a variety of institutions including minority serving, womens, liberal arts, community colleges and other educational institutions. We present data on recruitment, the matching workshop, and research experiences, illustrating how the program has successfully created opportunities that changed the professional trajectory of many participants, infused a new dimension of diversity awareness among Laboratory staff, brought people together that would probably never have met otherwise, started new productive collaborations, and provided vibrant research experiences for faculty who otherwise have scarce opportunities for research.

Mary Ann E. Leung
Sustainable Horizons Institute
mleung@shinstitute.org

MS49

Model Reduction of Synchronized Lur'e Networks

In this talk, we investigate a model order reduction scheme that reduces the complexity of uncertain dynamical networks consisting of diffusively interconnected nonlinear Lure subsystems. We aim to reduce the dimension of each subsystem and meanwhile preserve the synchronization property of the overall network. Using the upper bound of the Laplacian spectral radius, we first characterize the robust synchronization of the Lure network by a linear matrix equation (LMI), whose solutions can be treated as generalized Gramians of each subsystem, and thus the balanced truncation can be performed on the linear component of each Lure subsystem. As a result, the dimension of the each subsystem is reduced, and the dynamics of the network is simplified. It is verified that, with the same communication topology, the resulting reduced network system is still robustly synchronized, and the a priori bound on the approximation error is guaranteed to compare the behaviors of the full-order and reduced-order Lure subsystem.

Xiaodong Cheng, Jacquelien M.A. Scherpen
ENgineering and TEchnology institute Groningen
(ENTEG)
University of Groningen

x.cheng@rug.nl, j.m.a.scherpen@rug.nl

MS49

Operator Splitting and Model Reduction

Modeling and simulation for large scale nonlinear systems can be very costly. Thus, the need for model reduction arises. POD is a commonly used method for approximating large-scale nonlinear systems with a reduced order model. However, POD may not be able to capture what has not been observed i.e. the results of POD might be input-dependent. In order to mitigate this limitation we propose an approach that combines operator splitting and model reduction for solving nonlinear systems. In other words, we consider the linear and nonlinear terms in the dynamical system separately. First, we reduce the linear terms using some optimal model reduction algorithm. Then, we approximate the nonlinear terms using POD. Once we have reduced the model, we apply operator splitting i.e. in each step we numerically integrate the linear and nonlinear parts separately. The implemented operator splitting first evolves the linear terms, and then uses the result to evolve the nonlinear terms.

Klajdi Sinani
Virginia Tech
klajdi@vt.edu

Serkan Gugercin
Virginia Tech
Department of Mathematics
gugercin@vt.edu

Christopher A. Beattie
Virginia Polytechnic Institute and State University
beattie@vt.edu

Jeff Borggaard
Virginia Tech
Department of Mathematics
jborggaard@vt.edu

MS49

Randomized Residual-based Error Estimators for Parametrized Equations

We propose a randomized a posteriori error estimator for reduced order approximations of parametrized (partial) differential equations. The error estimator has several important properties: the effectivity is close to unity with prescribed lower and upper bounds at specified high probability; the estimator does not require the calculation of stability (coercivity, or inf-sup) constants; the online cost to evaluate the a posteriori error estimator is commensurate with the cost to find the reduced order approximation; the probabilistic bounds extend to many queries with only modest increase in cost. To build this estimator, we first estimate the norm of the error with a Monte-Carlo estimator using Gaussian random vectors whose covariance is chosen according to the desired error measure, e.g. user-defined norms or quantity of interest. Then, we introduce a dual problem with random right-hand side the solution of which allows us to rewrite the error estimator in terms of the residual of the original equation. In order to have a fast-to-evaluate estimator, model order reduction methods can be used to approximate the random dual solutions. Here, we propose a greedy algorithm that is guided by a scalar quantity of interest depending on the error estimator. Numerical experiments on a multi-parametric Helmholtz problem demonstrate that this strategy yields rather low-dimensional reduced dual spaces.

Kathrin Smetana
University of Twente
Faculty of Electrical Engg, Mathematics and Computer Science
k.smetana@utwente.nl

Olivier Zahm
Univ. Grenoble Alpes, Inria
olivier.zahm@inria.fr

Anthony T. Patera
Massachusetts Institute of Technology
Department of Mechanical Engineering
patera@mit.edu

MS49

From Data to Structured Reduced-order Models

Projection-based model reduction schemes require access to the internal state-space representation of the (complex) physical system. In practical applications, the model might be available implicitly via a simulation code and even depend on look-up tables. Thus, instead of internal dynamics, only measurements of the system or simulation results are available. The possible unmanageably large (computational) cost that is required to obtain this data may be avoided if a reliable surrogate model can be used instead. One popular method to obtain such a model solely from data is dynamic mode decomposition (DMD).

Although the system under investigation is considered a black-box, it should reflect certain physical properties like preservation of energy, stability, or passivity, and of course, it is desirable to enforce such structural properties in the surrogate model. To encode these properties in the non-intrusive surrogate model, we introduce a modification of DMD that produces a discrete-time port-Hamiltonian realization. To this end, we split the problem into several different tasks, including a least-squares solution of the dissipation inequality and the solution of a skew-symmetric Procrustes problem. Our theoretical findings are illustrated with numerical examples.

Benjamin Unger
TU Berlin
unger@math.tu-berlin.de

MS50

High-order Galerkin Method for Helmholtz and Laplace Problems on Multiple Open Arcs

We present a spectral numerical scheme for solving Helmholtz and Laplace problems with Dirichlet boundary conditions on a finite collection of open arcs in $\mathbf{R}^2$. An indirect boundary integral method is employed, giving rise to a first kind formulation whose variational form is discretized using weighted Chebyshev polynomials. Well-posedness of both continuous and discrete problems is established as well as spectral convergence rates under the existence of analytic maps to describe the arcs. In order to reduce computation times, a simple matrix compression technique based on sparse kernel approximations is developed. Numerical results provided validate our claims.

Carlos Jerez-Hanckes

Pontificia Universidad Catolica de Chile
cjerez@ing.puc.cl

MS50

A Stochastic Greens Function - Integral Equation Method for High Frequency, Chaotic Wave Systems

Even though we are seeking the highest fidelity, the computer representation is not exactly the same compared to the real world. These uncertainties may arise from imprecise knowledge of the system, small differences in manufacturing, or numerical errors in the simulation. For integrable, regular wave systems, these small differences can be considered as local perturbations of the entire system. The numerical solution is still a good approximation to the exact solution of the physical problem. However, the situation is very different in non-integrable, wave-chaotic dynamics. The wave solutions can be extremely sensitive to details and initial conditions. We propose a rigorous and versatile mathematical model in prediction and exploration of wave dynamics in chaotic environments. A novel stochastic Greens function (SGF) method is proposed, which quantitatively describes the universal statistical property of chaotic media through random matrix theory (RMT). Based on the SGF, we develop a stochastic integral equation formulation that rigorously resolves the transmitting correlation, propagation correlation, and receiving correlation in the dense multipath, chaotic environment. Moreover, it leads to nonstandard statistical models, which utilize the macroscopic properties of environments instead of detailed environmental specifics. The theoretical investigation is evaluated and validated through representative experiments.

Zhen Peng
University of New Mexico
Department of Electrical and Computer Engineering
pengz@unm.edu

MS50

Frozen Gaussian Approximation for 3-D Elastic Wave Propagation and Seismic Tomography

Three-dimensional (3-D) elastic wave propagation and seismic tomography is computationally challenging in large scales and high-frequency regime. We propose the frozen Gaussian approximation (FGA) to compute the 3-D elastic wave equation and use it as the forward modeling tool for seismic tomography with high-frequency data. The derivation requires to do asymptotic expansion in the week sense so that one is able to perform integration by parts. In particular, we obtain the diabatic coupling terms for SH- and SV-waves, with the form closely connecting to the concept of Berry phase that is intensively studied in quantum mechanics and topology (Chern number). The accuracy and parallelizability of the FGA algorithm is illustrated by comparing to the spectral element method for 3-D elastic wave equation. With a parallel FGA solver built as an computational engine, we explore various applications in 3-D seismic tomography, including seismic traveltime tomography, full waveform inversion, and optimal transport theory-based seismic tomography, respectively. We notice the nonconvergence of Wasserstein distance-based seismic tomography for high-frequency seismic data under certain circumstances, and explain the reason by a simple example of high-frequency wave. Global minimization for seismic tomography is investigated based on particle swarm algorithm. We also apply the FGA algorithm to train a neural network to learn the depth of Moho surface and the velocities below and above it.

James Hateley
Department of Mathematics
University of California, Santa Barbara
jhateley@ucsb.edu

Lihui Chai
School of Mathematics
Sun Yat-sen University, China
chailihui04@gmail.com

Ping Tong
Nanyang Technological University, Singapore
tongping@ntu.edu.sg

Xu Yang
Department of Mathematics
University of California, Santa Barbara
xuyang@math.ucsb.edu

MS50

Learning Dominant Wave Directions For High-frequency Helmholtz Equations

We present a ray-based finite element method (ray-FEM) for the high-frequency Helmholtz equation in smooth media, whose basis is learned adaptively from the medium and source. The method requires a fixed number of grid points per wavelength to represent the wave field; moreover, it achieves an asymptotic convergence rate of $O(k^{-\frac{1}{2}})$, where k is the wave number. The local basis is motivated by the geometric optics ansatz and is composed of polynomials modulated by plane waves propagating in a few dominant ray directions. The ray directions are learned by processing a low-frequency wave field that probes the medium with the same source. Once the local ray directions are extracted, they are incorporated into the local basis to solve the high-frequency Helmholtz equation. This process can be continued to further improve the approximations for both local ray directions and high-frequency wave fields iteratively. By using a hybrid approach, combining the ray-FEM with a high-order asymptotic expansion for the near field close to the source, we can properly capture the singularity of a point source solution.

Jun Fang
UCI
junf2@uci.edu

Jianliang Qian
Department of Mathematics
Michigan State University
qian@math.msu.edu

Hongkai Zhao
University of California, Irvine
Department of Mathematics
zhao@math.uci.edu

Leonardo Zepeda
Lawrence Berkeley National Laboratory
lzepeda@lbl.gov

MS51

Introduction to Neural Networks and Deep Learn-

ing

In this tutorial, we provide participants with a quick introduction to neural networks and deep learning. Specifically, will explore the key ideas of (i) activation functions, (ii) stochastic gradient descent, (iii) back propagation, and (iv) regularization techniques. We will then explore some of the basic ideas around key architectures such as convolutional neural networks and recurrent neural networks. Everything will be coded in Python, which is relatively easy to follow even if you don't know the language, and code snippets will be provided. Participants will get the most out of this if they have the Anaconda software suite installed on their laptops (see anaconda.com to download) and follow along interactively, but the material should still be accessible for those who just want to observe passively.

Jeff Humphreys
Brigham Young University
jeffh@math.byu.edu

MS51

Mathematical Optimization: Modern Algorithms for Applications in Data Science

We take a closer look at optimization algorithms which have become important in the context of new applications in data science. The main goal of this tutorial consists in the mathematical derivation and convergence analysis of these algorithms. Only minor background knowledge of optimization is required. In particular, we consider first order methods, but also second order methods specialized to the particular application. Among those are methods like coordinate descent or conditional gradient methods. For second order methods we consider inexact Newton methods with trust region techniques. The training of neutral networks and the LASSO problem serve as one of the main applications for this tutorial.

Ekkehard W. Sachs
University of Trier
sachs@uni-trier.de

MS52

High Performance Denoising on X-ray Ptychography

The success of ptychographic imaging experiments strongly depends on achieving high signal-to-noise ratio. This is particularly important in nanoscale imaging experiments when diffraction signals are very weak and the experiments are accompanied by significant parasitic scattering, outliers or correlated noise sources. It is also critical when rare events such as cosmic rays, or bad frames caused by electronic glitches or shutter timing malfunction take place. During this talk we are going to introduce a novel iterative algorithm named Advanced Denoising for X-ray Ptychography (ADP). The algorithm exploits the direct forward model for different ptychographic noise sources to achieve a thorough characterization and removal of both structured and random noise. We will present a high performance implementation of the algorithm, which, exploiting GPU acceleration and MPI multiprocessing, achieves comparable execution times with respect to similar HPC solutions, while retrieving much sharper features, cleaner background, and higher contrast.

Pablo Enfedaque
Lawrence Berkeley National Laboratory
pablo.enfedaque@gmail.com

Stefano Marchesini
Lawrence Berkeley National Lab
smarchesini@lbl.gov

Huibin Chang
School of Mathematical Sciences, Tianjin Normal University
huibinchang@lbl.gov

MS52

On the use of Automatic Differentiation as a Generic Phase Retrieval Framework

The recent rapid development in Coherent Diffraction Imaging (CDI) methods has enabled nanometer-resolution imaging in a variety of experimental modalities. Image reconstruction with such CDI methods involves solving the phase retrieval problem, where we attempt to reconstruct an object from only the amplitude of its Fourier transform. This can be framed as a nonlinear optimization problem which we can solve using a gradient-based minimization method. Typically, such approaches use closed-form gradient expressions. For complex imaging schemes, deriving this gradient can be difficult and laborious. This restricts our ability to rapidly prototype experimental and algorithmic frameworks. In this work, we use the *reverse-mode automatic differentiation* method to implement a generic gradient-based phase retrieval framework. With this approach, we only need to specify the physics-based forward propagation model for a specific CDI experiment; the gradients are exactly calculated automatically through a sequential application of the chain rule in a reverse pass through the forward model. Our gradient calculation uses the popular Tensorflow deep learning library, which includes state-of-the-art accelerated gradient descent algorithms. We demonstrate the generic nature of this phase retrieval method through numerical experiments in transmission (far-field and near-field) ptychography, multi-angle Bragg ptychography, tomography, and ptychographic tomography geometries.

Saugat Kandel, Ming Du
Northwestern University
saugat.kandel@u.northwestern.edu,
mingdu2015@u.northwestern.edu

Siddharth Maddali
Materials Science Divison
Argonne National Laboratory
b216692@anl.gov

Marc Allain
Institute Fresnel
marc.allain@univ-amu.fr

Chris Jacobsen
Advanced Photon Source
Argonne National Laboratory
cjacobsen@aps.anl.gov

Stephan O. Hruszkewycz
Materials Science Division
Argonne National Laboratory
shrus@anl.gov

Youssef Nashed
Mathematics and Computer Science Division
Argonne National Laboratory

ynashed@anl.gov

MS52

Ptychography: An Algorithmic Playground

To collect data for ptychographic phase retrieval involves five steps. 1) A small area of a sample is illuminated by a beam of coherent radiation the probe. 2) The beam scattered by the sample propagates to form a diffraction pattern, which is recorded by a CCD camera. 3) The sample is moved laterally, ensuring that there is some overlap between the new and the previously illuminated regions. 4) Another diffraction pattern is recorded. 5) Steps 3 and 4 repeat to cover an area of interest. This straightforward experimental procedure belies the richness of the resulting inversion problem, which involves re-phasing the diffraction data to form a quantitative phase/amplitude image of the sample. The inversion has been tackled by a myriad of optimisation schemes, but in this talk I will focus on the original approach, the Ptychographic Iterative Engine (PIE), a form of stochastic gradient descent. I will show how PIE can be accelerated by adapting the algorithms used to train neural networks, and I will demonstrate how this simple algorithm can be extended to correct experimental errors, recover corrupt or missing data, and image a sample in 3D.

Andrew M. Maiden
University of Sheffield
a.maiden@sheffield.ac.uk

MS52

Sparse Dictionary Learning Methods for Coherent X-ray Diffractive Imaging

Coherent X-ray Diffractive Imaging (CXDI) is a popular technique at highly coherent, brilliant x-ray sources for characterizing the behavior and structure of materials at the nanoscale. New x-ray sources coming on line in the next few years will produce orders of magnitude more data than is produced currently. In the last decade or so, great strides have been achieved in the literature on sparse representations and processing of large data sets. We will explore the use of these results in CXDI, an area to which their application has thus far been limited. We will develop a class of image retrieval algorithms based on novel non-convex and nonlinear numerical optimization techniques utilizing analysis and synthesis sparse dictionary learning methods. These algorithms will also incorporate all known physical constraints on the experimental geometry and sample, yielding a unified image retrieval algorithm of great utility to the coherent imaging community.

Ashish Tripathi, Brendt Wohlberg, Richard Sandberg
Los Alamos National Laboratory
atripathi@lanl.gov, brendt@lanl.gov, sandberg@lanl.gov

MS53

Multi-scaled Simulations and Modelling of Novel Electronic Devices

Recently interest in electronic, magnetic and optical materials based on inorganic and organic nano-materials has increased significantly. Such new materials and novel device architectures have a potential to bring the technology into the more than Moore and beyond Moore era. The most time efficient and cost-effective approach to find the best materials and devices for specific application is to perform modelling and simulations.
To show the importance of the device simulations, in this talk, I will present an exciting research carried out in the field of molecular electronics and quantum transport of the next generation transistors for the semiconductor industry. The main aim is to establish a link between the material properties, device architecture and performance using hierarchical, multi-scaled simulation modelling. For example, in order to evaluate the performance of molecular based flash cells we developed a simulation flow that links the density function theory (DFT) with a three-dimensional (3D) numerical device simulator. The main advantage of that framework is that once the charge for the molecule is obtained from the DFT method and it is transferred to the 3D numerical TCAD simulator, a quantum corrected drift-diffusion transport formalism is applied. It includes quantum corrections by means of the density-gradient approach. In this way the computation flow has capabilities to evaluate not only the material capabilities but also the device performance.

Vihar Georgiev
School of Engineering
University of Glasgow, UK
vihar.georgiev@glasgow.ac.uk

MS53

Algorithms for Large-scale Ab Initio Quantum Transport Simulations

Accurately predicting the performance of current and future electronic components (logic switches, memory cells, thermoelectric generators, light-emitting devices) requires to go beyond classical simulation approaches such as the drift-diffusion equations and to adopt quantum mechanical concepts. Also, structural variations at the atomic scale can only be captured by modelling schemes that take each individual atom into account. Finally, the possibility to stack several materials with very different properties together makes the parameterisation of empirical electronic structure models extremely difficult, thus favouring the usage of first-principles methods. In this context, quantum transport device simulators based on density-functional theory (DFT) and the Non-Equilibrium Green's Function (NEGF) formalism have recently seen tremendous developments. The latter have been made possible through algorithmic innovations that have contributed to a drastic reduction of the simulation times and therefore an increase of the device dimensions that can be handled. In this talk, a few algorithm examples will be presented, both to compute the open boundary conditions that connect the simulation domain to its environment and to solve the resulting Schrödinger equation. It will also be shown how the unique features of graphics processing units (GPUs) can be leveraged to minimise the computational burden of large-scale ab initio quantum transport simulations.

Mathieu Luisier
Integrated Systems Laboratory
ETH Zurich
mluisier@iis.ee.ethz.ch

Fabian Ducry, Sascha Brück, M. Hossein Bani-Hashemian
ETH Zurich
ducryf@iis.ee.ethz.ch, bruecks@iis.ee.ethz.ch,
hossein.banihashemian@mat.ethz.ch

Mauro Calderara
Integrated Systems Laboratory
ETH Zurich

mcalderara@iis.ee.ethz.ch

MS53

GPU vs CPU - Comparative Performance Evaluation of Parallel Linear Solvers for TCAD

Will typical technology computer aided design (TCAD) applications in electronics benefit from GPUs? This talk sheds some light on the domain-specific challenges in TCAD and presents results on using GPUs for a range of representative TCAD examples. Performance results will be presented for two common machine configurations: Single workstations on the one hand, and small clusters, as they are typical for in-house simulations, on the other hand. Our results allow for better decisions on when additional development time for GPU acceleration may be well spent.

Karl Rupp
TU Wien
me@karlrupp.net

MS53

Recent Advances in High Performance Process TCAD

Process technology computer-aided design (TCAD) deals with simulating semiconductor device fabrication steps, such as etching and deposition, to enable computer-based device designs. The simulation backends are based on a variety of numerical methods, e.g., particle transport, surface advection, diffusion, and stress calculation, underlining the inter-disciplinary nature of this topic. The rapidly evolving device concepts in electronics more and more utilize the third dimension - in contrast to previous planar technologies - to sustain the demand for higher integration densities: More compute performance or storage capacity is required whilst simultaneously limiting power consumption and if possible reducing device sizes. Future device designs will continue on this three-dimensional (3D) design path which will further increase the already dire need for fast and accurate 3D simulation capabilities to fully capture 3D effects arising during the individual fabrication steps. In this talk, recent advances in high performance process TCAD will be presented. The focus will be on accelerated flux calculation approaches and re-distancing algorithms: both important for a wide range of processing steps but also potentially relevant to other application areas. *The financial support by the Austrian Federal Ministry for Digital and Economic Affairs and the National Foundation for Research, Technology and Development is gratefully acknowledged.*

Georgios Diamantopoulos, Paul Manstetten, Lukas Gnam, Vito Simonka, Luiz Felipe Aguinsky, Michael Quell, Alexander Toifl
Christian Doppler Laboratory for High Performance TCAD
Institute for Microelectronics, TU Wien
diamantopoulos@iue.tuwien.ac.at, manstetten@iue.tuwien.ac.at, gnam@iue.tuwien.ac.at, simonka@iue.tuwien.ac.at, aguinsky@iue.tuwien.ac.at, quell@iue.tuwien.ac.at, toifl@iue.tuwien.ac.at

Andreas Hössinger
Silvaco Europe Ltd.
andreas.hoessinger@silvaco.com

Josef Weinbub
Christian Doppler Laboratory for High Performance TCAD
Institute for Microelectronics, TU Wien
weinbub@iue.tuwien.ac.at

MS54

LAP A Fast Solver for Coupled Imaging Problems

Coupled nonlinear inverse problems arise in numerous imaging applications, and solving them is often difficult due to ill-posedness and high computational cost. In this work, we introduce LAP, a linearize and project method for coupled nonlinear inverse problems with two (or more) sets of coupled variables. LAP is implemented within a Gauss–Newton framework. At each iteration of the Gauss–Newton optimization, LAP linearizes the residual around the current iterate, eliminates one block of variables via a projection, and solves the resulting reduced dimensional problem for the Gauss–Newton step. The method is best suited for problems where the subproblem associated with one set of variables is comparatively well-posed or easy to solve. LAP supports iterative, direct, and hybrid regularization and supports element-wise bound constraints on all the blocks of variables. This offers various options for incorporating prior knowledge of a desired solution. We demonstrate the advantages of these characteristics with several numerical experiments. We test LAP for two and three dimensional problems in super-resolution and MRI motion correction, two separable nonlinear least-squares problems that are linear in one block of variables and nonlinear in the other. We also use LAP for image registration subject to local rigidity constraints, a problem that is nonlinear in all sets of variables. These two classes of problems demonstrate the utility and flexibility of the LAP method.

James L. Herring
Department of Mathematics
University of Houston
herring@math.uh.edu

MS54

Title Not Available

Abstract not available.

James G. Nagy
Emory University
Department of Math and Computer Science
nagy@mathcs.emory.edu

MS54

Efficient Marginalization-based MCMC Approaches for Hierarchical Bayesian Inverse Problems

Hierarchical models in Bayesian inverse problems are characterized by an assumed prior probability distribution for the unknown state and measurement error precision, and hyper-priors for the prior parameters. Combining these probability models using Bayes' law often yields a posterior distribution that cannot be sampled from directly, even for a linear model with Gaussian measurement error and Gaussian prior. Gibbs sampling can be used to sample from the posterior, but problems arise when the dimension of the state is large. This is because the requisite Gaussian sample required for each iteration can be prohibitively expensive to compute and because the statistical

efficiency of the Markov chain degrades as the dimension of the state increases. The latter problem can be mitigated using marginalization-based techniques, but these can be computationally prohibitive as well. In this paper, we combine the low-rank techniques for approximating the prior-preconditioned misfit part of the Hessian along with a one-block MCMC sampler to obtain an approximate one-block sampler. We consider two variants of this approach: delayed acceptance and pseudo-marginalization. We provide a detailed analysis of the acceptance rates and computational costs associated with our proposed algorithms and compare their performances on two numerical test cases—image deblurring and inverse heat equation.

Arvind Saibaba
North Carolina State University
asaibab@ncsu.edu

Johnathan M. Bardsley
University of Montana
bardsleyj@mso.umt.edu

Andrew Brown
Clemson
ab7@g.clemson.edu

Alen Alexanderian
NC State University
alexanderian@ncsu.edu

MS54

Title Not Available

Abstract not available.

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

MS55

Unified Gas Kinetic Particle Method for Multiscale Photon Transport

We consider the solution of linear kinetic models to propose a unified gas-kinetic particle (UGKP) method for the simulation of multiscale photon transport. The multiscale nature of the particle method mainly comes from the recovery of the time evolution flux function in the unified gas-kinetic scheme (UGKS) through a coupled dynamic process of particle transport and collision. In the diffusive limit, the UGKP method could recover the solution of the diffusion equation with the cell size and time step being much larger than the photons mean free path and the mean collision time. In the free transport limit, it presents an exact particle tracking process as the original Monte Carlo method. Several numerical examples covering all transport regimes from the optically thin to optically thick are computed to validate the accuracy and efficiency of the current scheme.

Weiming Li
Beijing Computational Science Research Center
Hong Kong University of Science & Technology
liweiming@csrc.ac.cn

Chang Liu
The Hong Kong University of Science & Technology
cliuaa@connect.ust.hk

Yajun Zhu
Northwestern Polytechnical University
zhuyajun@mail.nwpu.edu.cn

Jiwei Zhang
Beijing Computational Science Research Center
jwzhang@csrc.ac.cn

Kun Xu
University of Science & Technology
Hong Kong
makxu@ust.hk

MS55

Nonlinear Moment Model for Radiative Transfer Equation

I will introduce some moment model as numerical approximations to radiative transfer equation. The Mn model based on the maximum entropy principle is equipped with the perfect theory, while it is not applicable to practical simulation due to the lack of explicit closure. As a consequence, the models often used can be collected in the flavour of Pn model and Sn model. We are devoted to develop some models with explicit closure in the spirit of maximum entropy principle. The models are expected to depict the specific intensity both in isotropic mode and in delta function.

Ruo Li
School of Mathematical Science
Peking University
rli@math.pku.edu.cn

MS55

Moment Methods for Polydisperse Sprays with Evaporation

Polydisperse multiphase flows are difficult to model. Lagrangian methods can become overwhelmingly expensive as particle count increases and traditional Eulerian methods can be inaccurate. This talk presents a new moment-based model for polydisperse flows that are typical of sprays. The new model builds on an existing moment technique that was developed for the treatment of solid particles of constant size and adds an extension for evaporation. This model is derived from the maximum-entropy moment framework and is described by fifteen first-order hyperbolic partial differential equations that describe statistics of the evolution of particles that display an evolving range of velocities and diameters at every point in space. The model is currently derived assuming a Stokes law for drag, but extensions to more accurate drag laws are possible. In this talk, the derivation of the model is shown. A numerical implementation for its solutions is demonstrated. Numerical results are shown for several simple problems for which exact solutions to the kinetic equation are possible. It is demonstrated that the simple model produces very accurate results for these cases.

James McDonald, Mathieu Giroux
University of Ottawa
jmcdona6@uottawa.ca, mgiro027@uottawa.ca

Lucian Ivan
Canadian Nuclear Laboratories

lucian.ivan@cnl.ca

MS55

Asymptotic-preserving Lax-Wendroff Discontinuous Galerkin Schemes for Quasi-exponential Moment-closure Approximations of Kinetic Models

In many applications, the dynamics of gas and plasma can be accurately modeled using kinetic Boltzmann equations. These equations are integro-differential systems posed in a high-dimensional phase space, which is typically comprised of the spatial coordinates and the velocity coordinates. If the system is sufficiently collisional the kinetic equations may be replaced by a fluid approximation that is posed in physical space (i.e., a lower dimensional space than the full phase space). The precise form of the fluid approximation depends on the choice of the moment-closure. In general, finding a suitable robust moment-closure is still an open scientific problem. In this work we consider a specific moment-closure based on a nonextensible entropy formulation. In particular, the true distribution is replaced by a Maxwellian distribution multiplied by a quasi-exponential function. We develop a high-order, locally-implicit, discontinuous Galerkin scheme to numerically solve resulting fluid equations. The numerical update is broken into two parts: (1) an update for the background Maxwellian distribution, and (2) an update for the non-Maxwellian corrections. We also develop limiters that guarantee that the inversion problem between moments of the distribution function and the parameters in the quasi-exponential function is well-posed.

James A. Rossmanith
Iowa State University
Department of Mathematics
rossmani@iastate.edu

MS56

Multi-dimensional High-Order FEM Methods for Electron Kinetics in ALE Hydrodynamics

In Plasma physics, ALE hydrodynamics belongs to the most efficient simulation tools while modeling ions and electrons as interacting fluids. We support a different approach, where the kinetic model of electrons is used instead of the electron fluid. Our high-order space and time electron transport method (7D) then provides an efficient and physically relevant alternative to the multi-fluid approach. We use an arbitrary order mixed finite elements or discontinuous Galerkin (DG) spatial discretizations along with diagonally implicit Runge-Kutta (DIRK) time discretization with an arbitrary time step. We also address the aspects of the electron transport coupling to the hydrodynamic temperature in the case of local and nonlocal transport regimes with respect to the iteration count.

Milan Holec
Lawrence Livermore National Laboratory
holec1@llnl.gov

Ben S. Southworth
Department of Applied Mathematics
University of Colorado, Boulder
ben.southworth@colorado.edu

MS56

Using Virtual Element Method for Data Remap on Polytopal Meshes

A remap of physical quantities between two meshes is an important step of arbitrary Lagrangian-Eulerian (ALE) simulations. This step becomes more challenging for high-order discontinuous Galerkin schemes that require high-order conservative remap schemes. We propose and analyze a new framework to the development of remap schemes on polytopal meshes based on theory of the virtual element method. As the first step in applying this framework, we derive conservative L^2 -stable remap schemes that are second-order accurate in space and time. The convergence of these schemes is illustrated with a few numerical experiments on structured and polytopal meshes.

Konstantin Lipnikov
Los Alamos National Laboratory
lipnikov@lanl.gov

Nathaniel Morgan
LANL
nmorgan@lanl.gov

MS56

Exploration of a High-order Lagrangian Discontinuous Galerkin Method on Quadratic Triangular Meshes

We present a Lagrangian discontinuous Galerkin (DG) hydrodynamic method using subcell mesh stabilization (SMS)for compressible flows on quadratic triangular meshes in two-dimensional (2D) Cartesian coordinates. The physical evolution equations are discretized using a modal DG method with Taylor series polynomials. The Riemann velocity at the vertices of a curvilinear cell, and the corresponding surface forces, are calculated by solving a multidirectional approximate Riemann problem. Quadratic triangular cells have many deformational degrees of freedom, and with these cells, they can deform in unphysical ways. Likewise, the Riemann solution at an edge vertex differs from the one at the corner of a cell. With SMS, each quadratic triangular cell is decomposed into four quadrilateral subcells, that move in a Lagrangian manner. The edge vertex is surrounded by six subcells so that enough information can be obtained. The difference between these two density fields detected by SMS is used to correct the stress (pressure) input to the Riemann solver. This SMS scheme enables stable mesh motion and accurate solutions in a context of the Lagrangian high-order DG method. We also present effective limiting strategies that ensure monotonicity of the primitive variables with the high-order DG method. A suite of test problems are calculated to demonstrate the designed order of accuracy of this method and the robustness for the strong shock problems.

Xiaodong Liu
Los Alamos National Laboratory
Computational Physics Division
xliu@lanl.gov

Nathaniel Morgan, Donald E. Burton
LANL
nmorgan@lanl.gov, burton@lanl.gov

MS56

Invariant Domains Preserving ALE Approximation of Hyperbolic Systems with Continuous Finite El-

ements

The objective of the present paper to propose a second-order continuous finite element technique for solving hyperbolic systems in the Arbitrary Lagrangian Eulerian framework. This is done by revisiting the some recent results from [J.-L. Guermond, B. Popov, L. Saavedra, and Y. Yang. Invariant domains preserving arbitrary Lagrangian Eulerian approximation of hyperbolic systems with continuous finite elements. SIAM J. Sci. Comput., 39(2):A385A414, 2017.] and adapting the recently proposed convex limiting technique from [J.-L. Guermond, , M. Nazarov, B. Popov, and I. Tomas. Second-order invariant domain preserving approximation of the euler equations using convex limiting. SIAM, J. Sci. Comput., 2018, In press.] to the ALE setting. The main properties of the method presented in the paper is that, provided the user-defined ALE velocity is reasonable, the approximate solution produced by the algorithm is formally second-order accurate in space, is conservative, satisfies the so-called discrete geometric conservation law, preserves convex invariant domains of the system. The second-order accuracy is shown to hold in the maximum norm on numerical examples with smooth solutions.

Jean-Luc Guermond
Department of Mathematics
Texas A&M, USA
guermond@math.tamu.edu

Bojan Popov
Department of Mathematics
Texas A&M University
popov@math.tamu.edu

Laura Saavedra
Polytechnic University of Madrid
laura.saavedra@upm.es

MS57

Automated Fortran-C++ Bindings for Scientific Applications

This presentation will cover a new tool that may be used to integrate legacy but still evolving code bases for large scale scientific calculations that have been maintained over years and have been verified for a number of predictive features. We will show how it is possible, in an automated session, to expose established and large libraries written in C++ as native Fortran modules in a same way as modern Fortran code can be packaged.

Seth Johnson
Oak Ridge National Laboratory
johnsonsr@ornl.gov

MS57

The Simulation Development Environment (SDE): A C++ Framework for Reusable Computational Chemistry

Three major high performance quantum chemistry computational packages, NWChem, GAMESS and MPQC have been developed by different research efforts following different design patterns. The goal is to achieve unified environment that encompasses the functionality of these packages. Overcoming the challenges caused by the different communication patterns and software design of each of these packages is among the goals. Creating connections between these incompatible packages is the major motivation of the proposed work. This interoperability is achieved by bringing the benefits of modern C++. It easies the Software Engineering through a component framework with the ultimate goal of presenting a strategy and process used for interfacing two widely used and important computational chemistry methodologies: Quantum Mechanics and Molecular Mechanics.

Ryan M. Richard
Ames Laboratory
Iowa State University
ryanmrichard1@gmail.com

MS57

Data Flow Graph Programming for High-performance Scientific Computing in C++

Our open-source TiledArray parallel tensor framework is for composable massively parallel block-sparse tensor framework that may be used for Clustered Low-Rank Tensor Format. We extensively use C++ to ease the burden of generating tensor contractions in a form of high-performance implementations. In addition, we use data-flow scheduling based on representation of contractions as a graph of interdependent tasks. With modern C++ syntax we show how to conveniently and compactly represent the algebra of atomic potential interactions.

Edward F. Valeev
Virginia Tech
efv@vt.edu

MS58

A Conforming Enriched Finite Element Method for Interface Problems

A new conforming enriched finite element method is presented for interface problems with interface-unfitted meshes. For elliptic interface problems, the conforming enriched finite element space is constructed based on the P1-conforming finite element space. Approximation capability of the conforming enriched finite element space is analyzed. The standard conforming Galerkin method is considered without any penalty stabilization term. Our method does not limit the diffusion coefficient of the elliptic interface problem to a piecewise constant. Finite element errors in H1-norm and L2-norm are proved to be optimal. For Stokes/Stokes-elliptic interface problems, the conforming enriched finite element pair is constructed based on the MINI element pair. A ghost penalty term is used in the standard discretization form as a stabilization term. An inf-sup stability result is derived, which is uniform with respect to the mesh size. Finite element errors are proved to be optimal. Our method breaks the limit of the immersed finite element method which can only deal with the constant jump coefficient and the case of identical governing equations on either side of the interface. Numerical experiments are carried out to validate theoretical results.

Jinru Chen
Nanjing Normal University
05129@njnu.edu.cn

MS58

An Unfitted Interface Penalty Finite Element Method for Elliptic Interface Problems

An unfitted interface penalty finite element method

(UIPFEM) is proposed for the elliptic interface problems, in which the Nitsche's method together with the ideas of merging elements and harmonic weighting fluxes are used. Both the convergence rate of the UIPFE solution and the condition number of the algebraic system are optimal and independent of the interface position. Furthermore, the error estimates do not depend on the ratio of the discontinuous coefficients. Numerical examples are also given to confirm the theoretical results.

Peiqi Huang
Nanjing University
pqhuang@njfu.edu.cn

MS58

Unconditionally Stable Cutfem for Dynamic Interfaces in a Fluid Structure Interaction Problem

Abstract: Interface problems arise in several applications including heart models, cochlea models, aquatic animal locomotion, blood cell motion, front-tracking in porous media flows and material science, to name a few. One of the difficulties in these problems is that solutions are normally not smooth across interfaces, and therefore standard numerical methods will lose accuracy near the interface unless the meshes align to it. However, it is advantageous to have meshes that do not align with the interface, especially for time dependent problems where the interface moves with time. Remeshing at every time step can be prohibitively costly, can destroy the structure of the mesh, or can deteriorate the well-conditioning of the stiffness matrix, and affect the stability of the problem. For a simple moving interface fluid-membrane interaction, we present a formal second-order finite element discretization in space and first-order in time where the finite element triangulation does not fit the interface and it is unconditionally stable in time independently of mesh parameters and fluid viscosity and membrane stiffness. This is a joint work with Kyle Dunn and Roger Lui from WPI.

Sarkis Marcus
Worcester Polytechnic Institute (WPI)
msarkis@wpi.edu

MS58

Numerical Study of Lateral Phase Separation in Two-component Bio-membranes with TraceFEM

Lipid-driven separation of immiscible fluidic phases is likely a factor in the formation of rafts in cell membranes and bacteria. Also lipid vesicles have been recently designed to form phase-separated lipid domains in order to enhance functionality for drug delivery. We consider two models, Cahn-Hilliard and Allen-Cahn, of phase separation posed on a regular surface representing interplay between two-dimensional phases that form a cell membrane or a vesicle. Trace finite element method is applied to resolve the problem numerically on different shapes of vesicles. Long-term behavior of phases is studied and results of numerical analysis are presented.

Vladimir Yushutin
University of Houston
yushutin@math.uh.edu

MS59

Castro: A Compressible Astrophysical Hydrodynamics Code

Castro is an adaptive mesh refinement code that solves the compressible hydrodynamics equations using an Eulerian grid with simultaneous refinement in space and time. As part of the AMReX astrophysics suite, Castro follows a hybrid parallelism approach based on MPI+tiled OpenMP and it is available on github.com/AMReX-Astro/Castro. In Castro, the fluxes through the cell interfaces are constructed using a dimensionally-unsplit second order Godunov method. Castro supports the use of a general equation of state and nuclear reaction networks, gray and multigroup radiation under a flux limited diffusion approximation, and self gravity. The applications using Castro include simulations of different progenitor models of Type Ia supernovae, and studies of the dynamics in core-collapse supernovae. In this presentation I will expand on Castros capabilities, ongoing development in the code, and astrophysics applications.

Maria G. Barrios Sazo
Stony Brook University
maria.barriossazo@stonybrook.edu

Ann S. Almgren
Lawrence Berkeley National Laboratory
asalmgren@lbl.gov

Max Katz
Nvidia Corporation
mkatz@nvidia.com

Donald E. Willcox
Lawrence Berkeley National Laboratory
dewillcox@lbl.gov

Weiqun Zhang
Lawrence Berkeley National Laboratory
Center for Computational Sciences and Engineering
weiqunzhang@lbl.gov

Michael Zingale
Department of Astronomy and Astrophysics
UCSC
michael.zingale@stonybrook.edu

MS59

Solving Low Mach Number Stellar Hydrodynamics using MAESTROeX

We present our low Mach number hydrodynamics code, MAESTROeX, which is used for modeling low-speed astrophysical flows over long times that are too computationally expensive to solve using traditional compressible codes. We use the code to model astrophysical phenomena that occur in the low Mach number regime, including but not limited to the convective phase of Type Ia supernovae, convection in stars, and Type I X-ray bursts. MAESTROeX is suitable for two- and three-dimensional flows and uses the AMReX framework to perform simulations on Cartesian grids with adaptive mesh refinement. In the modeling of a three-dimensional full-star, a one-dimensional radial base state is used to map the base state to and from the Cartesian grid. Here, we present the results obtained from MAESTROeX for various astrophysical problems.

Duoming Fan
Lawrence Berkeley National Laboratory

dfan@lbl.gov

MS59

Fluctuating Hydrodynamics Simulations of Reactive Electrolyte Solutions

In the last decade, there has been significant progress in the computational fluctuating hydrodynamics (FHD) approach, which provides accurate description of fluids at small scales. We have developed an FHD formulation and numerical methodology for stochastic simulation of reactive liquid mixtures containing electrolytes. The numerical framework is based on a structured-grid finite-volume approach with cell-averaged densities and pressure, and face-averaged (staggered) velocities. This talk will elaborate on how our numerical code is constructed from the AMReX framework and how discrete fluctuation-dissipation balance is achieved. In addition, numerical challenges such as GMRes solver, random number generators, and large-scale parallelization will be discussed.

Changho Kim
UC Merced
ckim103@ucmerced.edu

Andrew Nonaka
Lawrence Berkeley National Laboratory
ajnonaka@lbl.gov

John B. Bell
Lawrence Berkeley National Lab
Center for Computational Sciences and Engineering
jbbell@lbl.gov

Alejandro Garcia
San Jose State University
algarcia@algarcia.org

Aleksandar Donev
Courant Institute of Mathematical Sciences
New York University
donev@courant.nyu.edu

MS59

WarpX: Exascale Modeling of Advanced Particle Accelerators with AMReX

Turning the current experimental plasma accelerator state-of-the-art from a promising technology into mainstream scientific tools depends critically on high-performance, high-fidelity modeling of complex processes that develop over a wide range of space and time scales. As part of the U.S. Department of Energys Exascale Computing Project, a team composed of LBNL, SLAC and LLNL researchers is developing a new plasma accelerator simulation tool: WarpX. We will present the code structure and how it articulates around its main components: the new Particle-In-Cell Scalable Application Resource (PICSAR) and the adaptive mesh refinement library AMReX, which are combined with redesigned elements of the Warp code, in the new WarpX software. The status, examples of convergence, scaling and applications will be presented.

Maxence Thevenet
Lawrence Berkeley National Laboratory
mthevenet@lbl.gov

Rémi Lehe
Lawrence Berkeley National Laboratory
ATAP
rlehe@lbl.gov

Andrew Myers
Lawrence Berkeley National Lab
atmyers@lbl.gov

Weiqun Zhang
Lawrence Berkeley National Laboratory
Center for Computational Sciences and Engineering
weiqunzhang@lbl.gov

Mark Hogan
SLAC National Accelerator Laboratory
mjhogan@stanford.edu

Cho-Kuen Ng
Stanford Linear Accelerator Center
cho@slac.stanford.edu

Olga Shapoval
LBNL
OShapoval@lbl.gov

Jaehong Park
Lawrence Berkeley National Laboratory
ATAP
jaehongpark@lbl.gov

David Grote
Lawrence Livermore National Laboratory
dpgrote@lbl.gov

John B. Bell
CCSE
Lawrence Berkeley Laboratory
jbbell@lbl.gov

Lixin Ge
Stanford Linear Accelerator Center
lge@slac.stanford.edu

Ann S. Almgren
Lawrence Berkeley National Laboratory
asalmgren@lbl.gov

Jean-Luc Vay
Lawrence Berkeley National Laboratory
Berkeley, CA
JLVay@lbl.gov

MS60

Lifting Nonlinear Systems: More Structure, More Opportunities for ROM?

We present a structure-exploiting nonlinear model reduction method for systems with general nonlinearities. First, the nonlinear model is lifted to a model with more structure via variable transformations and the introduction of auxiliary variables. The lifted model is equivalent to the original model: it uses a change of variables, but introduces no approximations. When discretized, the lifted model yields a polynomial system of either ordinary differential equations or differential algebraic equations, depending on the problem and lifting transformation. Proper orthogonal decomposition (POD) is applied to the lifted models,

yielding a reduced-order model for which all reduced-order operators can be pre-computed. Thus, a key benefit of the approach is that there is no need for additional approximations of nonlinear terms, in contrast with existing nonlinear model reduction methods requiring sparse sampling or hyper-reduction. Application of the lifting and POD model reduction are presented on a cumbustion test problem with strong nonlinearities.

Boris Kramer, Karen E. Willcox
Massachusetts Institute of Technology
bokramer@mit.edu, kwillcox@ices.utexas.edu

MS60

Scalable Online Nonlinear Goal-oriented Inference with Physics-Informed Maps

In the goal-oriented inverse problem, one uses observed data to infer often high-dimensional unknown model parameters that are then used to calculate some low-dimensional quantity of interest (QoI). In time-constrained settings where the QoI is needed to make a decision soon after observations are obtained, it is infeasible to solve the full inverse problem. We present an algorithm to rapidly map observations to QoI outputs for the deterministic nonlinear goal-oriented inverse problem. We apply tensor decompositions to approximate goal-oriented inverse problems to obtain small (multi-)linear physics-informed observation-to-QoI maps that can be calculated and stored in an offline preparatory phase, then quickly applied to online observations to get QoI estimates. We apply our algorithm to a goal-oriented tomography problem, where we achieve small QoI errors in much less time than solving the full inverse problem. Compared to black-box regression methods such as k-nearest-neighbors, we achieve greater accuracy in the QoI estimate for small training sets, and are more robust in the case where the training set does not well represent the test set.

Harriet Li
MIT
kameeko@mit.edu

Karen E. Willcox
Massachusetts Institute of Technology
kwillcox@ices.utexas.edu

MS60

Exploiting Low-dimensional Structure for Multifidelity Simulations

Reduced order modeling for simulations of parameterized systems frequently make use of low-rank structure in the solution manifold to construct efficient and accurate surrogates. We show in this presentation that this low-rank structure can persist across scales of fidelity. Given an inexpensive low-fidelity model and an expensive high-fidelity model, we use a greedy procedure to identify low-rank structure in the low-fidelity model. We then "lift" this detected low-rank structure to the high-fidelity model, which constructs a surrogate for the high-fidelity model. The resulting surrogate is as efficient to query as the low-fidelity model, but can achieve accuracy comparable to the high-fidelity model. We discuss theoretical guarantees for this problem and demonstrate through many examples that such low-dimensional structure across fidelity levels exists in practice and can be effectively exploited.

Akil Narayan
University of Utah
akil@sci.utah.edu

MS60

Exploiting Low-dimensional Structure to Efficiently Perform Stochastic Inference for Prediction

Uncertainty quantification is challenging for large-scale multiphysics applications where the number of uncertain parameters may be large, the number of high-fidelity model evaluations may be limited, and the available data may be corrupted by significant noise. The recently developed consistent Bayesian approach solves a specific stochastic inverse problem based on the measure-theoretic principles. This approach produces a pullback density on the parameters that is consistent in the sense that the push-forward of this density matches the given distribution on the observable data. While the consistent Bayesian approach is theoretically sound and conceptually simple, it does require approximating the push-forward of an initial probability density through the computational model. While this is certainly nontrivial, we can leverage advanced approaches for forward propagation of uncertainty to reduce the online computational burden. In this presentation, we discuss the utilization of dimension reduction strategies within this inversion framework with particular emphasis on the inference-for-prediction problem. Numerical results will be presented to demonstrate the effectiveness of combining dimension reduction and stochastic inference-for-prediction.

Timothy Wildey
Sandia National Laboratories
tmwilde@sandia.gov

Troy Butler
University of Colorado Denver
troy.butler@ucdenver.edu

John D. Jakeman
Sandia National Labs
jdjakem@sandia.gov

MS61

Multigrid Based Strategies for the Solution of Nonlinear Space-time Problems

We present a parallel and efficient multilevel solution strategy for solving non-linear time-dependent problems. We consider in particular the monodomain model, a non-linear reaction-diffusion equation arising from a problem in electrophysiology: the electrical activation in the human heart. We propose a semi-geometric multigrid solver, for which the coarse level approximation spaces are created using arbitrary hierarchies of non-nested meshes. Interpolation and restriction in the multilevel context are then realized by means of a discrete L2-projection between the non-matching meshes. This approach allows for creating the coarser levels of a multigrid hierarchy, even if only a single "fine" mesh is available. Hence, multigrid hierarchies can be created for arbitrary geometries in any dimension. We discuss how this approach can be applied to the monodomain equation discretised with space-time finite elements. While we use continuous finite elements in space, for stability reasons we adopt discontinuous elements in time. We discuss shortly the properties of this time discretization scheme. We investigate how different block smoothers, coarsening strategies and ordering of the space-time variables affect the overall convergence and robustness

of the solver. Finally, we investigate numerically the scalability and the convergence of our multilevel solver.

Pietro Benedusi
Institute of Computational Science
USI Lugano
pietro.benedusi@usi.ch

Rolf Krause, Patrick Zulian
Università della Svizzera italiana, Switzerland
rolf.krause@usi.ch, patrick.zulian@usi.ch

MS61

Adaptive Parareal: Introducing Time-grid Adaptivity to Improve the Stability of the Parareal Algorithm

When applied to an hyperbolic equation in a regime that falls outside of its stability region, the typical behaviour of the Parareal algorithm shows a dramatic increase in the error introduced by the prescribed update for early iterations. As the iterations proceed, and the sequential application of the fine solver catches up with later instants in the time domain, the error is eventually reduced, but this makes time-parallelisation extremely ineffective in these situations. We propose in this talk a procedure which aims at identifying this negative behaviour in early stages, and counteracting it by opportunely redistributing the computational power so to improve the quality of the update. The procedure here introduced is akin to the ones employed in adaptive stepsize methods, and necessitates a proper error estimate as well as an adequate refining procedure in order to be effective and still preserve its parallel efficiency. We discuss the various alternatives investigated so far and their effect on improving the stability of the classical Parareal algorithm when applied to various model problems.

Federico Danieli
University of Oxford
federico.danieli@trinity.ox.ac.uk

MS61

Parallel-in-Time Simulation of the Schrdinger Equation

We discuss the applicability of parallel-in-time integration methods to the Schrdinger equation. Modern supercomputers consist of millions of cores. Hence, to use these machines efficiently for numerical simulations, implementations need to be designed to have minimal serial parts. Many current parallel solvers for time-dependent partial differential equations compute a sequence of time-steps. Here, each core computes the solution for a portion of the spatial domain, which can be done in parallel for each time-step. Each time-step, however, is computed sequentially, often limit the scalability of the implementation. Parallel-in-time integrators avoid this limit by allowing to compute multiple time-steps in parallel. While this approach works particularly well for diffusive problems, parallel-in-time integration of oscillatory problems, like the Schrdinger equation, is more difficult. The Schrdinger equation describes the time evolution of quantum mechanical systems. To understand atomic and molecular phenomena it is often useful to inspect these time evolutions. Hence, efficient ways of solving the equation are needed. Being essentially the heat equation in imaginary time, it is similarly easy to implement. It is, however, more challenging to obtain parallel speedup in time due to its oscillatory nature. Solving the Schrdinger equation hopefully provides insight on how to solve other oscillatory problems more efficiently as well.

Hannah Rittich
University of Wuppertal
h.rittich@fz-juelich.de

MS61

pySDC: A Parallel Framework for Spectral Deferred Corrections

The parallel full approximation scheme in space and time (PFASST) allows to integrate multiple time-steps simultaneously. Based on iterative spectral deferred corrections (SDC) methods, PFASST uses a space-time hierarchy with various coarsening strategies to maximize parallel efficiency. In numerous studies, this approach has been used on up to 448K cores and coupled to space-parallel solvers which use finite differences, spectral methods or even particles for discretization in space. However, since the integration of SDC or PFASST into an existing application code is not straightforward and the potential gain is typically uncertain, we have developed the Python prototyping framework pySDC. It allows to rapidly test new ideas and to implement first toy problems more easily. It also provides interfaces to the finite element library FEniCS and the PETSc toolkit. In this talk, we describe the structure of the code and show different examples to highlight various aspects of the implementation, the capabilities and its usage. Moreover, this code sparked the development of a C++-based PFASST implementation, making use of the DUNE framework. The dune-PFASST module specifically targets finite element simulations and is used for reaction-diffusion examples. We will thus also discuss the transition from pySDC to dune-PFASST and show first numerical results.

Robert Speck
Juelich Supercomputing Centre
Forschungszentrum Juelich GmbH
r.speck@fz-juelich.de

MS62

Numerical Methods for Multidimensional Fractional Reaction-diffusion Systems in Developmental Biology

Reaction-diffusion (RD) models have been developed to capture the behavior of many natural systems. These models have been generalized to include a fractional differential operator, which facilitates accurate models of super diffusion. Natural systems can have many components that react and diffuse across multiple spatial dimensions. Modern research of complex, nonlinear systems is made possible by efficient numerical methods. In this talk, we discuss an efficient scheme for numerically solving such problems. For space discretization, we apply a Fourier spectral method that is suited for multidimensional systems. Efficient approximation of time-stepping is accomplished with an exponential time difference approach. Focus is given to aspects of fractional diffusion that are relevant to developmental biology.

Ashlin Harris
MIddle Tennessee State University
aph3b@mtmail.mtsu

Abdul M. Khaliq
Middle Tennessee State University
Department of Mathematical Sciences

akhaliq@mtsu.edu

MS62

Anomalous Diffusion, Dilation, and Erosion in Image Processing

In this talk, anomalous sub- and super-diffusion arising in image processing is considered and is modeled by a diffusion equation with fractional time derivative. It might serve as a building block for the construction of various filters. The resulting partial differential equation is discretised in space with centered differences and in time with the explicit or implicit Euler method, respectively. A numerical investigation is performed to illustrate new and interesting results. Additionally, the time derivative of the partial differential equation describing dilation and erosion is replaced by a fractional time derivative and then solved numerically. Interesting new questions arise from the presented numerical results.

Andreas Kleefeld
Forschungszentrum Jülich GmbH
Jülich Supercomputing Centre
a.kleefeld@fz-juelich.de

MS62

Modeling and Finite Difference Computing of Fractional Convection Equation

Abstract not available.

Changpin Li
Department of Mathematics
Shanghai University, Shanghai, China
lcp@shu.edu.cn

MS62

Inverse Source Problems for Fractional Differential Equations with Nonlocal Boundary Conditions

The purpose of this talk is twofold, some inverse source problems for time-fractional diffusion equations with nonlocal boundary conditions will be considered. The fractional derivative in time is defined in Riemann-Liouville or Caputo's sense (or composite of these two usually known as Hilfer derivative). Due to nonlocal boundary conditions the spectral problem is non-self-adjoint and we need a bi-orthogonal system of functions to apply the Fourier's method. The results about existence, uniqueness and stability of the solutions of the inverse problems will be discussed. Secondly, some inverse problems with multi-term fractional derivatives in time for space-time fractional differential equations will be discussed. Some future perspectives and problems will be shared with the audience.

Salman A. Malik
COMSATS University Islamabad
salman_amin@comsats.edu.pk

MS63

Sparse Representations of Data for HPC

Sparse representation of data is powerful tool for large datasets, with particular datasets being accurately represented by only a sparse set of function approximation coefficients. In the case of high performance computing (HPC), complex simulations often have a bottleneck in storing solution states, especially when these solutions are in high dimensions. This talk will introduce scalable sparse functional representations of data and demonstrate the unique problems that occur for HPC.

Rick Archibald
ORNL
archibaldrk@ornl.gov

MS63

UQTk, a FASTMath C++/Python Toolkit for Uncertainty Quantification: Overview and Applications

The FASTMath UQ Toolkit (UQTk) is a collection of libraries, tools and apps for the quantification of uncertainty in numerical model predictions. UQTk offers intrusive and non-intrusive methods for forward uncertainty propagation, tools for sensitivity analysis, sparse surrogate construction, Bayesian inference via various flavors of MCMC, model error assessment, as well as many other capabilities. The core libraries are implemented in C++ but a Python interface is available for easy prototyping and incorporation in UQ workflows. The talk will give an overview of UQTk capabilities and illustrate its application to complex scientific workflows typical of the SciDAC program.

Bert Debussschere, Khachik Sargsyan, Cosmin Safta
Sandia National Laboratories
bjdebus@sandia.gov, ksargsy@sandia.gov, csafta@sandia.gov

MS63

Nonlinear Optimization using the Toolkit for Advanced Optimization

In this talk, I will discuss the latest developments with the Newton, quasi-Newton, and nonlinear conjugate gradient methods for solving unconstrained and bound-constrained optimization problems in the Toolkit for Advanced Optimization. I will focus on the essential algorithm characteristics and make use of extensive computational studies to compare and contrast the performance.

Adam Denchfield
University of Illinois at Chicago
adenchfield@anl.gov

Alp Dener
Rensselaer Polytechnic Institute
Department of Mechanical, Aerospace, and Nuclear Engineering
adener@anl.gov

Todd Munson
Argonne National Laboratory
Mathematics and Computer Science Division
tmunson@mcs.anl.gov

MS63

Uncertainty Quantification in Computations of Large Scale Physical Systems

Uncertainty quantification (UQ) in computations of large scale physical systems is rendered feasible, at least for smooth observables, with state of the art, robust and efficient UQ methods and software. This talk will present an overview of these methods, focusing on forward UQ methods addressing challenges associated with high-

dimensionality and high computational costs. Illustrative demonstrations in computations of physical systems will be presented, highlighting these capabilities and associated software tools.

Habib N. Najm
Sandia National Laboratories
Livermore, CA, USA
hnnajm@sandia.gov

MS64

Simulation of Interpenetrating Plasmas in 2D using the Grid Based Continuum Code LOKI

Counter-streaming interpenetration is an important problem in plasma physics, with application to laboratory scale experiments performed on NIF, NOVA, etc, as well as astrophysical relevance. Physical instabilities often dominate the dynamic evolution, and in this work we seek a detailed understanding of these instabilities in multiple dimensions using continuum grid based simulation. Our 2D+2V simulation code, LOKI, uses fully sixth-order accurate conservative finite differences to reduce computational cost, and employs MPI parallelism to efficiently use large computational resources. Using LOKI, simulations of collisionless interpenetrating plasmas are diagnosed to extract instability growth rates. The results are compared to linear theory, and good agreement is observed. Finally, we address the effects of particle collisions using a simplified collision operator.

Jeffrey W. Banks
Rensselaer Polytechnic Institute
banksj3@rpi.edu

Richard Berger
AX DIVISION
Lawrence Livermore National Laboratory
berger5@llnl.gov

Stephan Brunner
Centre de Recherches en Physique des Plasmas
Ecole Polytechnique Fédérale de Lausanne
stephan.brunner@epfl.ch

Thomas Chapman
Lawrence Livermore National Laboratory
chapman29@llnl.gov

Andris Dimits
dimits1@llnl.gov
dimits1@llnl.gov

Debojyoti Ghosh
Lawrence Livermore National Laboratory
Center for Applied Scientific Computing
ghosh5@llnl.gov

MS64

Implicit, Conservative Multiscale Algorithms for Hybrid Kinetic (Particle) Ion/Fluid Electron Models

The quasi-neutral hybrid model with kinetic ions and fluid electrons is a promising approach to model multi-scale problems in space and laboratory plasmas. However, current explicit schemes suffer from a number of key algorithmic issues related to the stable propagation of whistler waves, and finite-grid instabilities for cold ion beams [P. W. Rambo, *J. Comput. Phys.*, **118**, 152-158 (1995).] due to non-conservation of discrete momentum or energy. Fully implicit methods have been recently explored [B. Sturdevant, S. E. Parker, Y. Chen and B. B. Hause, *J. Comput. Phys.*, **316**, 519 (2016).] to step over fast timescales, but these schemes are not conservative. Here, we present a novel particle-based non-linear hybrid algorithm that features discrete conservation of mass, momentum and energy [A. Stanier, L. Chacon, G. Chen, arxiv e-print:1803.07158]. The scheme combines a cell-centered discretization of the field equations with implicit-midpoint time advance. A fluid moment-based preconditioner is used to accelerate convergence when stepping over fast normal modes. The scheme is also generalized to allow for adaptive integration of the ion orbits with orbit averaging and conservative smoothing to reduce finite-particle noise. We demonstrate the unique conservation and stability properties of the scheme.

Luis Chacon, Adam Stanier, Guangye Chen
Los Alamos National Laboratory
chacon@lanl.gov, stanier@lanl.gov, gchen@lanl.gov

MS64

A Hybrid Discontinuous Galerkin Method for 2D-2V Vlasov-Poisson Problems with Geometry

Abstract not available.

David C. Seal
United States Naval Academy
seal@usna.edu

MS65

Partitioning Strategies for Exascale Computing

Partitioning is a crucial technique for load-balancing parallel computing. Current graph partitioners do not scale well in that the quality deteriorates as the number of processors increases. Geometric methods are more scalable but partition quality is often worse. Spectral partitioning based on the graph Laplacian can be slow, but has the advantage one can leverage efficient parallel software for linear algebra. We cover the state of the art in parallel partitioning and discuss future directions for exascale computing.

Erik G. Boman
Center for Computing Research
Sandia National Labs
egboman@sandia.gov

Karen D. Devine, Siva Rajamanickam, Michael W. Wolf
Sandia National Laboratories
kddevin@sandia.gov, srajama@sandia.gov, mmwolf@sandia.gov

MS65

Graph Algorithms in Scalable Implementations of Computational Chemistry Methods

Abstract not available.

Sriram Krishnamoorthy
Pacific Northwest National Laboratory

sriram@pnnl.gov

MS65

Accelerating Biological Network Clustering with Sparse Matrix Kernels on Large-scale System

Biological networks represent the relationships between biological entities and analyses on them reveal crucial information about the complex biological systems in which these entities are contained. Common examples of biological networks include protein-protein interaction networks, sequence similarity networks, gene regulatory networks and gene co-expression networks. With networks of size of tens of billions edges, utilization of fast and scalable algorithms that run on large-scale systems becomes imperative. In this work, we investigate the feasibility of utilizing GPU accelerators to reduce the parallel execution time of a large-scale biological network clustering algorithm named HipMCL, which stands for High-performance Markov Cluster Algorithm. The clustering of biological networks of interest makes use of common building blocks for sparse matrix computations, for which the HPC literature offers a wide variety of optimizations. Adopting certain optimization techniques and porting compute-intensive routines to accelerators, our goal is to improve the scale of the networks for which the analysis can be performed within hours.

Oguz Selvitopi
Lawrence Berkeley Lab
roselvitopi@lbl.gov

Ariful Azad, Aydin Buluc
Lawrence Berkeley National Laboratory
azad@iu.edu, abuluc@lbl.gov

MS65

Parallelism in Deep Learning through Linear Algebra

We present a formulation of the forward and backward propagations of a neural network as the solutions of two single lower and upper block triangular systems of equations. For standard feedforward neural networks (FNNs), the triangular systems are reduced to block bidiagonal. For recurrent neural networks (RNNs), they are reduced to two-level block bidiagonal. For a general computation graph (directed acyclic graph), they can have more complex triangular sparsity patterns. We focus on the former two types and discuss the direct and iterative algorithms that can be used for the parallelism of the solutions to the systems. By choosing some popular activation functions such as ReLU and leaky ReLU, we show that the systems can be expressed as low-rank updates to systems with unit diagonals, which facilitate better parallelism. We also show that for FNNs and RNNs with k layers and τ time steps the propagations can be performed in parallel in $O(\log k)$ and $O(\log k \log \tau)$ steps, respectively.

Yu-Hong Yeung
Purdue University
yyeung@purdue.edu

Alex Pothen
Purdue University
Department of Computer Science
apothen@purdue.edu

Maxim Naumov
NVIDIA
maxim.a.naumov@gmail.com

Mahantesh Halappanavar
Pacific Northwest National Laboratory
mahantesh.halappanavar@pnnl.gov

MS66

Development of Terrestrial Dynamical Core for Earth System Models

Developing a predictive understanding of the terrestrial water cycle at local to global scale is essential for accurate assessment of water resources. Higher spatial resolution in the land component of the Energy Exascale Earth System Model (E3SM) alone is insufficient to meet the U.S. Department of Energys (DOEs) 10-year vision. Next-generation terrestrial models need to move beyond one-dimensional systems by including scale appropriate three-dimensional physics formulations. To this end, we are developing a dynamical core which is tailored to solving thermal mass balance. The method of choice should converge linearly in the velocity in the presence of discontinuous coefficients and non-orthgonal grids. We choose to implement a mixed finite element based on the Brezzi-Douglas-Marini (BDM1) mixed finite element spaces analyzed by Wheeler and Yotov [M. F. Wheeler and I. Yotov, A multipoint flux mixed finite element method, SIAM J. Numer. Anal. 44:5 (2006) 2082-2106.]. In their method, the drawback of the resulting saddle point problem is overcome by using a vertex-centered quadrature rule, decoupling the velocity systems around the vertices of the mesh. This allows for the velocity system to be eliminated locally, resulting in a cell-centered stencil for pressure only. In this talk we describe the challenges in efficiently implementing such a method as well as present strategies for scaling the solution from the watershed to continental scales.

Nathan Collier
Oak Ridge National Laboratory
nathaniel.collier@gmail.com

MS66

Modeling Antarctic Ice Sheet Dynamics using Adaptive Mesh Refinement

The response of the Antarctic Ice Sheet (AIS) remains the largest uncertainty in projections of sea level rise. The AIS (particularly in West Antarctica) is believed to be vulnerable to collapse driven by warm-water incursion under ice shelves, which causes a loss of buttressing, subsequent grounding-line retreat, and large (up to 4m) contributions to sea level rise. Very fine (finer than 1km) spatial resolution is needed to resolve ice dynamics around shear margins and grounding lines (the point at which grounded ice begins to float). The BISICLES ice sheet model uses adaptive mesh refinement (AMR) to enable sufficiently-resolved modeling of full-continent Antarctic ice sheet response to climate forcing. This talk will discuss recent progress and challenges modeling the sometimes-dramatic response of the ice sheet to climate forcing using AMR.

Daniel Martin
Lawrence Berkeley National Laboratory
dfmartin@lbl.gov

MS66

Diagnosing Physically-simulated Mixing in MPAS-O and E3SM via Online Lagrangian Particle Track-

ing with LIGHT

Ocean models are typically formulated in the Eulerian frame of reference, using discretizations and numerical algorithms developed over many decades for increasing fidelity of physical simulation. A recent ocean model, the Model for Prediction Across Scales Ocean, leverages advanced mimetic discretization and an unstructured mesh, providing a multiscale capability to resolve global ocean flows down to coastal scales in high performance computing environments. An online diagnostic, Lagrangian Insitu Global High-performance particle Tracking (LIGHT), provides the capability to understand mixing using the Lagrangian frame of reference via particle diagnostics at native model spatio-temporal resolution, which is essential as ocean modeling moves toward Exascale. Mixing simulated across the range of scales simulated by MPAS-O varies and depends upon resolution of mesoscale eddies and is largely Peclet number invariant as mixing is dominated by stirring resulting from geostrophic turbulence. LIGHT is used in idealized cases to compute in-situ estimates of mixing occurring due to the Eulerian flow and new results from MPAS-O simulations in the Energy Exascale Earth System Model (E3SM) are used to understand transport for biogeochemistry applications within the global ocean.

Phillip J. Wolfram, Todd Ringler, Luke Van Roekel
Los Alamos National Laboratory
pwolfram@lanl.gov, ringler@lanl.gov, lvanroekel@lanl.gov

Mark R. Petersen
Los Alamos National Laboratory
Computational and Computer Science Division
mpetersen@lanl.gov

Mathew Maltrud
LANL
maltrud@lanl.gov

Riley Brady
University of Colorado Boulder
rileyx@gmail.com

Jon Wolfe, Adrian Turner, Philip Jones, Xylar Asay-Davis
Los Alamos National Laboratory
jonbob@lanl.gov, akt@lanl.gov, pwjones@lanl.gov, xylar@lanl.gov

MS67

An Analysis of the Approximations and Errors in a Boundary Integral Solution of the Poisson-Boltzmann Equation

In biomolecular applications, the implicit solvent model simulates the mean field potential using continuum electrostatics on an infinite domain (the solvent, usually water with salt) with a cavity (the dissolved biomolecule). This results in a coupled system with the Poisson equation inside the cavity, the Poisson-Boltzmann equation in the outer region, and appropriate interface conditions on the molecular surface. We formulate the resulting system in terms of boundary integral equations that run on the molecular surface only, and solve them with a boundary element method. We can distinguish several sources of errors in the resulting algorithm, for example, from the molecular surface discretization, the numerical integration, the linear solver tolerance, and acceleration strategies. These affect not only the computed property of interest (solvation energy), but also the condition number of the stiffness matrix, which may impact the time to solution. In this talk, we will review the sources of error in our boundary integral solver, and discuss their effect on the numerical solution and on the behavior of the matrix condition number.

Christopher D. Cooper
Department of Mechanical Engineering
Universidad Tecnica Federico Santa Maria
christopher.cooper@usm.cl

Paulo Fernandez
Universidad Técnica Federico Santa María, Chile
paulo.fernandezm@alumnos.usm.cl

MS67

On a Goal-oriented Formulation for the PGD Method

The subject of the talk will deal with a mathematical formulation for constructing reduced-order models suited for the approximation of quantities of interest. The main idea is to formulate a minimization problem that includes an equality or inequality constraint on the error in the goal functional so that the resulting model is capable of delivering estimations of the quantity of interest at a higher rate of convergence. The formulation will be applied to, and tested on, the so-called Proper Generalized Decomposition (PGD) method. Such a paradigm represents a departure from classical goal-oriented approaches in which a reduced model is first derived by minimizing the energy of the problem, or residual functionals, and then adapted via a greedy approach by controlling the error with respect to quantities of interest using dual-based error estimates. Numerical examples will be presented in order to demonstrate the efficiency of the proposed approach. In particular, we will consider the case of a delaminated composite material simulated using the Proper Generalized Decomposition approach.

Serge Prudhomme
École Polytechnique de Montréal
serge.prudhomme@polymtl.ca

Kenan Kergrene
Ecole Polytechnique de Montréal
kenan.kergrene@polymtl.ca

Marc Laforest
École Polytechnique de Montréal
marc.laforest@polymtl.ca

Ludovic Chamoin
LMT, ENS Cachan, CNRS, Université Paris-Saclay
chamoin@lmt.ens-cachan.fr

MS67

Reduced Basis Methods with Least Squares Finite Elements

Reduced Basis methods for parametrized Partial Differential Equations (PDEs) are techniques that compute accurate approximations to high-fidelity or full-order numerical solutions at a much reduced computational cost. Their measure of accuracy relies on an assumption that the full-order solutions can be computed to within arbitrary accuracy of the analytical solution. Solutions that are not sufficiently smooth may result in full-order solutions with large errors, and overly optimistic accuracy measures for the re-

duced order model. Least Squares Finite Element Methods (LSFEMs) are an alternative to standard Galerkin variational methods, that replace a PDE with a system of first-order PDEs and minimize the residual over a suitable Hilbert space. One attractive feature of LSFEMs is that it provides a cheap error estimate with respect to the true solution using the residual of the transformed system. In this work, we use the LSFEM method in the framework of Reduced Basis methods, in order to develop direct error estimates between reduced-order solutions and analytical solutions, using the residual error estimator. Numerical results are shown for linear elliptic PDEs.

Peter Sentz
University of Illinois at Urbana-Champaign
sentz2@illinois.edu

Jehanzeb H. Chaudhry
University of New Mexico
jehanzeb@unm.edu

Luke Olson
University of Illinois Urbana-Champaign
lukeo@illinois.edu

MS67

Estimating and Controlling Communication Errors in Multiphysics Problems

Multiphysics problems by their very nature involve the communication or transfer of information between distinct physical processes, domains or time scales. This communication necessarily results in the transfer of numerical (and modeling) error between different physical processes, domains or time scales, which must be estimated and controlled as part of a successful computation. We develop adjoint-based a posteriori error techniques to estimate the error in a quantity of interest and implement refinement strategies to reduce these three very different types of communication error. When both stochastic and deterministic processes are present, additional complications arise. We consider a very general system of equations $g(\lambda, u) = 0$, $\lambda \in D \subset \mathbb{R}^d$, $u \in \mathbb{R}^n$, where λ is a single realization of a random variable Λ and the solution $U = G(\Lambda)$ is also a random variable. We assume that the solution $G(\lambda)$ must be obtained via an iterative (fixed point) procedure and address the convergence of *distributions* of solutions U or of a quantity of interest $q(U)$.

Simon Tavener, Don Estep
Colorado State University
tavener@math.colostate.edu, estep@stat.colostate.edu

MS68

Overcomplete Sparse Tensor Decomposition

To enhance interpretability of multimodal datasets, one may wish to represent the data in a sparse tensor format - with a sparse core multidimensional tensor and several dictionary matrices corresponding to each mode of the data. Common compressed sensing techniques may be used to obtain such a representation. However, when the data is best expressed as a sum of such decompositions, the separation and recovery of the individual component tensors and their sparse representations is more challenging, with standard alternating schemes having high computational cost and failing to exploit the special structure of the representation. We propose a system of optimization problems to recover the sparse core tensor coefficients for each component simultaneously, rather than using an alternating scheme. Moreover, we introduce an iterative algorithm, based on soft-thresholding shrinkage techniques, to solve the proposed optimization problem working on the separate components in parallel. The proposed model and algorithm are demonstrated to effectively separate and recover sparse tensor components on both synthetic and real-world datasets.

Ashley Prater-Bennette
Air Force Research Laboratory
ashley.prater-bennette@us.af.mil

Kenneth Carr
Northeastern University
carr.ken@husky.neu.edu

MS68

Complexity Bounds of First-order Methods for Large-scale Optimization

We consider a class of large-scale optimization with certain structure. In particular, the structure of the problem may involve the smoothness of certain loss functions, saddle-point formulation, and linear equality constraints. In order to understand the efficiency of first-order methods, it is important to study the impact of the specific structure on the complexity of the problem. We will design worst-case instances that yield lower complexity bounds of any first-order method.

Yuyuan Ouyang
Department of Mathematical Sciences
Clemson University
yuyuano@clemson.edu

MS68

A Distributed ℓ^1 Minimization Algorithm

The ℓ^1 minimization models have been widely used in many applications such as image processing and machine learning. There have been extensive research on the numerical algorithms of solving ℓ^1 minimization models. However, when the data has a huge size, a single computer would not be able to store and process all the data at one time. We would discuss in this talk a distributed algorithm that divides the huge minimization problem into several intermediate sub-problems such that each sub-problem could be solved separately.

Guohui Song
Clarkson University
gsong@clarkson.edu

MS68

Stability and Generalization for SGD in Pairwise Learning

Pairwise learning refers to a learning task which involves a loss function depending on pairs of instances. Most notable examples of pairwise learning include bipartite ranking, metric learning, AUC maximization and minimum error entropy (MEE) principle. We establish the stability results for SGD algorithms for pairwise learning in both convex, strongly-convex and non-convex settings. As a consequence, we derive their generalization error bounds. Finally, we describe our stability results by illustrating some

specific examples of pairwise learning such as AUC maximization, metric learning and MEE. The motivation comes from a previous recent work and the results we obtain complement it in the setting of pointwise learning.

Yiming Ying
SUNY Albany
yying@albany.edu

MS69

Adaptive Local Reduced Basis Approach for Stochastic Optimization with Nonlinear PDEs

We present an inexact trust-region (TR) algorithm from with a local, reduced basis (RB) approximation to efficiently solve risk-averse optimization problems with linear and nonlinear PDE constraints. The main contribution of this work is a numerical framework for systematically constructing surrogate models for the TR subproblem and objective function using local sample-based approximations. Under standard assumptions, the inexact TR algorithm is guaranteed to converge from any initial guess, provided that errors in the evaluation of the objective function and its gradient using our RB approach are adequately bounded. We provide conditions for which our RB approximations satisfy the necessary bounds and demonstrate the performance of our proposed approach through numerical examples. These examples demonstrate that we can efficiently solve risk-averse PDE-constrained optimization problems with significant computational savings when compared to Monte Carlo methods.

Wilkins Aquino
Professor
Duke University
wa20@duke.edu

Zilong Zou
Duke University
zilong.zoou@duke.edu

Drew P. Kouri
Optimization and Uncertainty Quantification
Sandia National Laboratories
dpkouri@sandia.gov

MS69

Measuring Uncertainty and Performance of Machine Learning Algorithms with Buffered Probability of Exceedance

In the machine learning community, the relevance of constrained optimization is on the rise due, in part, to the growing desire to incorporate secondary data-dependent performance requirements, such as constraints on fairness or false alarm rates, while optimizing a primary objective. Commonsense formulations of performance requirements, however, can often be non-convex or discontinuous, e.g. expected 0-1 loss. We discuss the use of Buffered Probability of Exceedance, a new idea from risk management, for formulating performance requirements and the tractable optimization problems they often generate. In particular, we formulate new binary classification problems with precise control over false alarm rates.

Matthew Norton
Navel Postgraduate School
mnorton@nps.edu

MS69

Low-Rank Solvers for Optimal Control Problems Constrained by PDEs with Uncertain Inputs

Optimization problems constrained by deterministic unsteady PDEs are computationally challenging because one would need to solve a system of PDEs coupled globally in time and space, and time-stepping methods quickly reach their limitations due to the enormous demand for storage. Yet more challenging are problems constrained by unsteady PDEs involving parametric or uncertain inputs. A viable solution approach to optimization problems with stochastic constraints employs the spectral stochastic Galerkin finite element method (SGFEM). However, the SGFEM often leads to *curse of dimensionality*, in the sense that it results in prohibitively high-dimensional linear systems with tensor product structure. Moreover, a typical model for an optimal control problem with stochastic inputs (OCPS) will usually be used for the quantification of the statistics of the system response - a task that could in turn result in additional enormous computational expense. In this talk, we consider two prototypical model OCPS and discretize them with SGFEM. We exploit the underlying mathematical structure of the discretized systems to derive and analyze low-rank iterative solvers and robust block-diagonal preconditioners for solving the resulting stochastic Galerkin systems. The developed solvers are efficient in the reduction of temporal and storage requirements of the high-dimensional linear systems. Finally, we illustrate the effectiveness of our solvers with numerical experiments.

Akwum Onwunta
Max Planck Institute, Magdeburg, Germany
onwunta@cs.umd.edu

MS69

Scalable Quasi-Newton Methods for the Optimization of Engineering Systems under Uncertainty using High-performance Computing

The optimization of many engineering systems reaches extreme size and unprecedented computational complexity due to the need to consider high fidelity physics and to account for uncertainties. We discuss various paradigms of optimization under uncertainty governed by differential equations and present scalable computational methods for solving such optimization problems. We delve in quasi-Newtwon methods for constrained optimization in function spaces and present our recent work in infinite-dimensional quasi-Newton algorithms, corresponding decomposition methods, parallelization strategies, and solver developments. Finally, we present on the effectiveness of our methodology in the optimization of electricity dispatch operations in large-scale power grid systems and in high-resolution structural topology optimization using the U.S. Department of Energy's high-performance computing infrastructure.

Cosmin G. Petra
Lawrence Livermore National
Center for Advanced Scientific Computing
petra1@llnl.gov

MS70

Statistical Learning Approaches for Automatic Parameter Tuning to Increase Parallel Performance

in Large CFD Simulations

The first part of the talk focuses on auto-tuning parameters to boost parallel performance Parallelizable simulations often feature tunable parameters that affect their parallel performance. These parameter spaces are often continuous rather than discrete and only offer efficient results within a small subset of the space. Manual investigation of these spaces is impractical, and exhaustive searches often impossible. Instead, machine learning techniques can be applied to the problem to automatically tune these parameters, generating efficient performance results on a variety of simulations without the need to empirically search the solution space to ensure the efficiency of each case. A case of auto-tuning cache block variables within a fluid flow PDE simulation will be examined towards this effect. The second part of the talk, we demonstrate the use of simple ML techniques to study the uncertainty in numerical weather prediction models due to the interaction of multiple physical processes. The first problem addressed herein is the estimation of systematic model errors in output quantities of interest at future times, and the use of this information to improve the model forecasts. The second problem considered is the identification of those specific physical processes that contribute most to the forecast uncertainty in the quantity of interest under specified meteorological conditions.

Rachel Cooper
Virginia Tech.

Azam Moosavi
Tesla
azmosavi@vt.edu

Adrian Sandu
Virginia Polytechnic Institute
and State University
sandu@cs.vt.edu

Danesh Tafti
Department of Mechanical Engineering
Virginia Polytechnic Institute & State University
dtafti@vt.edu

Vishwas Rao
Argonne National Laboratory
vhebbur@anl.gov

MS70

A Learning-based Approach for Data Assimilation

Data assimilation is used to solve many forecasting problems. Specifically, in meteorology, DA is used to obtain analysis of the state of the atmosphere that can be used as initial condition in weather forecast problem. The most common approach for solving DA problems is the Kalman Filter (KF) family of method. The ensemble Kalman Filter (EnKF) is one of the common approaches that can handle the non-linear systems efficiently. EnKF is a Monte-Carlo algorithm that solves the data assimilation problem by sampling the probability distributions involved in Bayes' theorem. EnKF works by operating on an ensemble of the system state which are statistical representative of the system state. If the sample size is too small it fails to reflect the true state well which is called under-sampling error In typical weather forecasting applications, the model state space has dimension $10^9 - 10^{12}$, while the ensemble size typically ranges between 30 - 100 members. Consequently all family of EnKF are fundamentally prone to sampling errors when the ensemble size is small. Under-sampling error leads to the issues of malfunctioning the DA algorithm and manifest themselves as inbreeding and long-range spurious correlations and have been shown to cause filter divergence.

Azam Moosavi
Tesla
azmosavi@vt.edu

Ahmed Attia
Argonne National Laboratory
attia@anl.gov

Adrian Sandu
Virginia Polytechnic Institute
and State University
sandu@cs.vt.edu

MS70

Inferring Black Box Functions under a Limited Computational Budget

Numerous engineering problems are characterized by objective functions, where the evaluation of the objective via experiment and/or simulation comes at the expense of vast resources. In some problems, running a simple instance of the objective can take days or weeks. The goal of these problems is to identify regions of design space which satisfy a set of criteria defined at the outset (which includes optimization cases). In many engineering applications, it is of key interest to understand which design variables drive changes in the objectives to compute the relative order of importance. This question is answered via a combination of data-driven modeling and global sensitivity analysis where so-called sensitivity indices are computed. In this work, we develop an optimal acquisition strategy for obtaining the most relevant regions of the design space if one aims to obtain information about a function of the expensive objective like the order of sensitivities. This framework guides the designer towards evaluating the objective function to acquire information about the sensitivities sequentially. The framework can be extended more generally to acquire information about any arbitrary quantity of interest to the engineer or scientist. We verify and validate the proposed methodology by applying it on synthetic test problems. We then demonstrate our approach on a real-world industry engineering problem of optimizing a compressor for oil applications.

Piyush Pandita
School of Mechanical Engineering, Purdue University,
West Lafayette, Indiana, 47906
piyush100078@gmail.com

Ilias Bilionis
Purdue University
ebilionis@gmail.com

Jesper Kristensen
Cornell University
jespertoftkristensen@hotmail.com

Jitesh Panchal
School of Mechanical Engineering
Purdue University, West Lafayette, Indiana, 47907

panchal@purdue.edu

MS70

Bayesian Model Selection, Calibration and Uncertainty Quantification for Thermodynamic Properties

Models of the thermodynamic properties of materials form the basis for technological applications including the calculation of phase diagrams and simulation of microstructure evolution during processing both of which play an important role in the design of materials for improved performance. Currently, the weighting of datasets, removal of outliers and the selection of model forms rely on expert judgements and do not provide uncertainty intervals. In this work we present a Bayesian framework for the selection, calibration and uncertainty quantification of thermodynamic property models. The framework is enabled by recent advances in numerical sampling methods. In addition, we present intuitive modifications that automatically weight datasets, improve robustness of outlier treatments, and ensure consistency of thermodynamically related models. We demonstrate the power of the approach through the construction of models for the specific heat, enthalpy, entropy and Gibbs free energy of Hafnium metal for the alpha, beta and liquid phases at temperatures ranging between 0 and 4900K.

Noah Paulson, Elise Jennings, Marius Stan
Argonne National Laboratory
npaulson@anl.gov, ejennings@anl.gov, mstan@anl.gov

MS71

Machine-learning Error Models for Approximate Solutions to Parameterized Systems of Nonlinear Equations

This work proposes a machine-learning framework for constructing statistical models of errors incurred by approximate solutions to parameterized systems of nonlinear equations. These approximate solutions may arise from early termination of an iterative method, a lower-fidelity model, or a projection-based reduced-order model. The proposed statistical model comprises the sum of a deterministic regression-function model and a stochastic noise model. The method constructs the regression-function model by applying regression techniques from machine learning (e.g., support vector regression, artificial neural networks) to map features (i.e., error indicators such as sampled elements of the residual) to a prediction of the approximate-solution error. The method constructs the noise model as a mean-zero Gaussian random variable. This work considers a wide range of feature-engineering methods and regression techniques that aim to ensure that (1) the features are cheaply computable, (2) the noise model exhibits low variance (i.e., low epistemic uncertainty introduced), and (3) the regression model generalizes to independent test data. Numerical experiments performed on several computational-mechanics problems and types of approximate solutions demonstrate the ability of the method to generate statistical models of the error that satisfy these criteria and significantly outperform more commonly adopted approaches for error modeling.

Brian A. Freno, Kevin T. Carlberg
Sandia National Laboratories
bafreno@sandia.gov, ktcarlb@sandia.gov

MS71

No Equations, No Variables, No Parameters, No Space, No Time: Data and the Modeling of Complex Systems

Abstract not available.

Yannis Kevrekidis
Department of Chemical and Biomolecular Engineering
Johns Hopkins University
ykevrek1@jhu.edu

MS71

Learning Parameters and Constitutive Relationships with Physics Informed Deep Neural Networks

Abstract not available.

Alexander Tartakovsky
Pacific Northwest National Laboratory
alexandre.tartakovsky@pnnl.gov

MS71

Projection-based Model Reduction: Formulations for Physics-based Machine Learning

We consider parametric surrogate models for applications in science and engineering where the goal is to predict high-dimensional output quantities of interest, such as pressure, temperature and strain fields. Our methodology combines a low-dimensional parametrization of these quantities of interest using the proper orthogonal decomposition (POD) with machine learning methods to learn the map between the input parameters and the POD expansion coefficients. The use of particular solutions in the POD expansion provides a way to embed physical constraints, such as boundary conditions. Case studies demonstrate the importance of embedding physical constraints within learned models, and also highlight the important point that the amount of model training data available in an engineering setting is often much less than it is in other machine learning applications, making it essential to incorporate knowledge from physical models.

Renee Swischuk
Texas A&M University
swischuk@mit.edu

Laura Mainini
Politecnico di Torino
lmainini@mit.edu

Benjamin Peherstorfer
Courant Institute of Mathematical Sciences
New York University
pehersto@cims.nyu.edu

Karen E. Willcox
Massachusetts Institute of Technology
kwillcox@ices.utexas.edu

MS72

Recent Advances in Primal-dual Coordinate Meth-

ods for ERM

We revisit the empirical risk minimization (ERM) problem with a convex-concave reformulation. With this formulation, recently a line of work proposed to solve the ERM with variants of the primal-dual coordinate methods. This kind of methods match the state-of-the-art overall convergence rate and benefit in some particular settings. Specifically, we consider the problem of ERM when features and observations are of very large-scale, but the optimal model only interacts with a small portion of features (aka the optimal primal variable is sparse). The method we propose is called Doubly Greedy Primal-Dual Coordinate Descent algorithm, which is able to exploit sparsity in both primal and dual variables. It enjoys a low cost per iteration and our theoretical analysis shows that it converges linearly with a good iteration complexity, provided that the set of primal variables is sparse. We then extend this algorithm further to leverage active sets. The resulting new algorithm is even faster in practice and experiments on large-scale Multi-class data sets show that our algorithm achieves up to 30 times speedup on several state-of-the-art optimization methods.

Qi Lei
University of Texas, Austin
Institute for Computational Engineering and Sciences(ICES)
cecilialeiqi@gmail.com

Ian En-Hsu Yen
Canegie Mellon University
eyan@cs.cmu.edu

Chao-yuan Wu
the University of Texas at Austin
cywu@cs.utexas.edu

Inderjit S. Dhillon
University of Texas at Austin
inderjit@cs.utexas.edu

Pradeep Ravikumar
Carnegie Mellon
pradeepr@cs.cmu.edu

MS72

Efficient Coordinate-wise Leading Eigenvector Computation

We develop and analyze efficient "coordinate-wise" methods for finding the leading eigenvector, where each step involves only a vector-vector product. We establish global convergence with overall runtime guarantees that are at least as good as Lanczos's method and dominate it for slowly decaying spectrum. Our methods are based on combining a shift-and-invert approach with coordinate-wise algorithms for linear regression.

Jialei Wang
University of Chicago
jialei@uchicago.edu

Weiran Wang
Amazon Alexa
weiranwang@ttic.edu

Dan Garber
Technion
dangar@technion.ac.il

Nathan Srebro
Toyota Technical Institute - Chicago
nati@ttic.edu

MS72

Coordinate-Wise Descent Methods for Leading Eigenvalue Problems

In quantum physics and quantum chemistry, an important class of problems is to find the leading eigenvalue and eigenvector of a matrix. However, the challenge is that the dimension of the matrix grows exponentially with the system size and thus can be extremely huge. In the fast-developing machine learning field, leading eigenvalue problems for huge sparse matrix is also an important task. However, traditional eigensolvers are powerless for huge sparse matrix. Coordinate-wise descent methods (CDMs) are a class of optimization methods which minimize the objective function along coordinate directions. In our work, we reformulate the leading eigenvalue problem as a non-convex optimization problem and apply CDMs. We review some existing CDMs and propose some new algorithms. The convergence of algorithms is analyzed and compared. Finally, the numerical examples on quantum many-body problems also show the efficiency of CDMs over traditional methods.

Yingzhou Li
Duke University
yingzhouli0417@gmail.com

Jianfeng Lu
Mathematics Department
Duke University
jianfeng@math.duke.edu

Zhe Wang
Duke University
zhe.wang3@duke.edu

MS72

Primal-dual Block Coordinate Update Methods for Multi-block Structured Affinely Constrained Problems

The alternating direction method of multipliers (ADMM) has been popularly used in many areas including imaging, statistics, and machine learning. However, its direct extension to multiple block problems is not guaranteed to converge under merely convexity assumption. This talk will give two variants of multi-block ADMM. One is based on mixing Jacobi and Gauss-Seidel updates, and the other one applies the randomization technique. For the second variant, an asynchronous parallel version will also be presented to handle extremely large-scale problems. Convergence and also numerical results will be shown for all three methods.

Yangyang Xu
Rensselaer Polytechnic Institute
xuy21@rpi.edu

MS73

Random Sampling and Efficient Algorithms for

Multiscale PDEs

Multiscale PDEs problems are computationally challenging because it usually requires a fine mesh to resolve small scales. In this talk, I will describe an efficient framework for multiscale PDE problems that uses random sampling to capture low-rank local solution spaces arising in a domain decomposition framework. In contrast to existing techniques, our method does not rely on detailed analytical understanding of specific multiscale PDEs, in particular, their asymptotic limits. [Chen, K., Li, Q., Lu, J., and Wright, S. J. (2018). Random Sampling and Efficient Algorithms for Multiscale PDEs. arXiv:1807.08848.]

Ke Chen
University of Wisconsin-Madison
kchen222@wisc.edu

Qin Li
University of Wisconsin
qinli@math.wisc.edu

Jianfeng Lu
Mathematics Department
Duke University
jianfeng@math.duke.edu

Stephen Wright
University of Wisconsin
Dept. of Computer Sciences
swright@cs.wisc.edu

MS73

Stochastic Modeling of Neuronal Transport in Various Cellular Geometries

We present a mathematical framework to analyze the transport processes inside a neuron. Our model captures spatial dynamics and interactions of a motor and cargo particles through a system of coupled stochastic differential equations. We study the transport on a parallel arrangement of microtubules inside axon (axonal transport), as well as various tangled networks of microtubules inside soma (somatic transport). In all cases, we derive the effective velocity and diffusion coefficient at the macroscopic scale.

Abhishek Choudhary
Rensselaer Polytechnic Institute
abhi.achoudhary@gmail.com

Peter R. Kramer
Rensselaer Polytechnic Institute
Department of Mathematical Sciences
kramep@rpi.edu

MS73

An Adjoint Method using Fully Implicit Runge-Kutta Schemes for Optimization of Flow Problems

We introduce the fully discrete adjoint equations and corresponding adjoint method for unsteady PDE-constrained optimization problems. Specifically, we consider conservation laws on deforming domains that are temporally discretized by high-order fully implicit Runge-Kutta (IRK) schemes. Through a change of variables, the linear systems arising in the primal and dual problem are transformed, leading to computationally cheaper systems that compare competitively with those derived from diagonally implicit Runge-Kutta (DIRK) schemes. Quantities of interest that take the form of space-time integrals are discretized in a solver-consistent manner. Our fully discrete, IRK adjoint method is used to compute exact gradients of quantities of interest with respect to the optimization parameters. These quantities of interest and their gradients are used for gradient-based PDE-constrained optimization. Our implementation of this IRK adjoint method is tested by computing the energetically optimal trajectory of a 2D airfoil in flow governed by the compressible Navier-Stokes equations. We also analyze the parallel performance of our IRK adjoint method and the DIRK adjoint method, showing that our implementation is computationally comparable.

Michael Franco
University of California, Berkeley
michaelfranco@berkeley.edu

Per-Olof Persson
Dept. of Mathematics
University of California, Berkeley
persson@berkeley.edu

Will Pazner
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
pazner1@llnl.gov

Matthew J. Zahr
University of Notre Dame
mzahr@nd.edu

MS73

Bifurcations and Dynamics in a Growth-First Discrete-Time Host-Parasitoid Model

We investigate the dynamics of a specific discrete-time host-parasitoid model originally proposed by May et al. (1981). While the model includes standard functional forms for parasitism and density dependent growth in the host species, the timing of these events in the life cycle leads to unexpected dynamics and rich mathematical behavior for certain parameters. We discuss the bifurcations in the model, including a period-doubling bifurcation, bistability, and a strange attractor.

Kelsey Marcinko, Mark Kot
University of Washington
kmar517@uw.edu, mark_kot@comcast.net

MS73

Recovery of Material Map by Full waveform inversion of Ultrasonic Phased Array Data

Typically, imaging algorithms within the ultrasonic nondestructive testing community assume that waves travel at a constant wave speed within the material under inspection. This is a poor assumption when the medium under inspection is of a heterogeneous nature and it can contribute to the misplacement of defects. Prior knowledge of the materials internal structure would allow corrective measures to be taken, but this is seldom available. The work presented here endeavours to reconstruct the spatially varying wave speed maps of heterogeneous media from ultrasonic phased array measurements. This is achieved via a semi-analytical approach to full waveform inversion and the reversible-jump Markov Chain Monte Carlo (rj-MCMC)

method.

Russell Niven
University of Strathclyde
russell.niven@strath.ac.uk

Tony Mulholland
university of Strathclyde
anthony.mulholland@strath.ac.uk

MS73

A Micro-Macro Acceleration Method for Simulations of Stiff Stochastic Differential Equations

Stochastic differential equations (SDEs) are commonly used to model physical systems from traffic flow to plasma dynamics and polymer networks. Many practical systems are however high dimensional, rendering grid-based techniques too expensive so that we must resort to Monte Carlo methods. At the same time, multiple time scales can be present between the fast microscopic paths and some slow macroscopic state variables. Simulating such systems is very computationally demanding. Explicit methods need many steps as they have a small stability domain, while implicit methods suffer from the high dimensionality. In this talk, I will present a new micro-macro acceleration method that tries to overcome stiffness and is fully explicit. The micro-macro acceleration method performs a short Monte Carlo simulation of the individual fast paths, before extrapolating some macroscopic state variables of interest over a larger time step. After extrapolation, we construct a new probability distribution that is consistent with the extrapolated macroscopic state variables, while minimizing the perturbation of the distribution available at the end of the Monte Carlo simulation. The final step is called matching and is the hardest step of the algorithm. During the talk I will explain micro-macro acceleration with a focus on matching. If time permits, I will show its performance on a numerical example.

Hannes Vandecasteele
University of Leuven
hannes.vandecasteele@cs.kuleuven.be

Przemyslaw Zielinski, Giovanni Samaey
Department of Computer Science, K. U. Leuven
przemyslaw.zielinski@cs.kuleuven.be, giovanni.samaey@cs.kuleuven.be

MS74

Transmission Eigenvalues for Materials with a Conductive Boundary Condition

In this talk we will investigate the inverse scattering problem associated with an inhomogeneous media with a conductive boundary. We consider the corresponding interior transmission eigenvalue problem. This is a new class of eigenvalue problem that is not elliptic, not self-adjoint, and non-linear, which gives the possibility of complex eigenvalues. We investigate the convergence of the eigenvalues as the conductivity parameter tends to zero as well as prove existence and discreteness for the case of an absorbing media.

Isaac Harris
Purdue University
Department of Mathematics
iharris1107@gmail.com

MS74

Determination of Electromagnetic Bloch Modes in a Medium with Frequency-dependent Coefficients

Photonic crystals, used to manipulate light, often involve periodic arrays of elements. In order to design the elements of such a periodic lattice, it is useful to determine the Bloch variety for the structure. This represents the relationship between a wave vector and the corresponding frequencies. In particular gaps where no propagation is possible for a given frequency are important. If the material is not frequency dependent, the problem reduces to finding resonant frequencies for Maxwell's equations with quasi-periodic boundary conditions given from the wave vector. However, if the constituent parts of the photonic crystal have different frequency dependent electromagnetic properties, this direct approach is no longer applicable. Instead we propose to compute the wave vector as a function of frequency. For a give frequency (and hence given electromagnetic parameters) this results in a quadratic eigenvalue problem for the magnitude of the wave vector. We show that this approach can be formulated as a mixed variational problem and, after linearization, becomes the problem of finding the generalized eigenvalues of a compact operator. Numerical results show that our approach can indeed compute eigenvalues for frequency dependent structures.

Shixu Meng
IMA
University of Minnesota
shixumen@umich.edu

MS74

Method of Fundamental Solution for Computing Elastic Interior Transmission Eigenvalues

Elastic interior transmission eigenvalues arise in the study of inverse scattering problems as a precursor for qualitative reconstruction methods. As such, their accurate and fast calculation is a desired but challenging task due to the non-elliptic, non-self-adjoint and non-linear structure of the underlying eigenproblem. We present an improved solver for the efficient calculation of approximate interior transmission eigenvalues based on the method of fundamental solutions. It distinguishes from most of the currently utilized techniques as being mesh- and integration free. A variety of numerical examples proves our simple approach to be more than competitive for not too exotic scatterers. Theoretical results will then complete our approximation analysis.

Lukas Pieronek
Forschungszentrum Jülich GmbH
Jülich Supercomputing Centre
l.pieronek@fz-juelich.de

MS75

Large Scale Unfitted Finite Elements for Fluid and Solid Mechanics Problems

We will present the aggregated finite element method for the efficient and accurate simulation of fluid and solid mechanics problems. This method solves all the conditioning issues of embedded, unfitted, or extended finite element methods, its implementation is quite simple, and it does not involve any additional artificial dissipation for co-

ercive problems or a very limited one for incompressible fluid problems. Numerical analysis, including stability and a priori convergence results, will also be included. We will also focus on the implementation of the method at hand on distributed memory platforms, taking into account all the different steps in the simulation pipeline, from the geometry to the linear system solve. With regard to the linear system, we will provide scalability results of our strategy using black-box algebraic multigrid solvers in the PETSc library without any kind of customization.

Santiago Badia
Universitat Politecnica de Catalunya
sbadia@cimne.upc.edu

Alberto F Martin, Francesc Verdugo
CIMNE
amartin@cimne.upc.edu, fverdugo@cimne.upc.edu

MS75

Unfitted FEM/DEM Methods for Modeling Variably Saturated Granular Media

We will present several methods for modeling two-phase incompressible air/water flows interacting with moving granular solids. The methods include conforming Arbitrary Lagrangian Eulerian formulations with deforming, boundary-fitted meshes and immersed and embedded Finite Element Methods on fixed methods with dynamic, implicit solid boundary representations. The incompressible two-phase Navier-Stokes equations use a conservative level set method to model air/water interface. In the current research, we focus on the coupling between different phases in order to avoid instability and get high accuracy and robustness during solid-solid contacts. While we enforce the no-slip boundary condition at the solid using a weak (Nitsche-type) method, we test several options for approximating the boundary integrals at implicit boundaries, including various smoothed Dirac penalization and the Shifted Boundary Method (SBM). We will present results for modeling such immersed solids in two-phase flows relevant to sediment transport modeling.

Chris Kees
U.S. Army Engineer Research and Development Center
Coastal and Hydraulics Laboratory
christopher.e.kees@usace.army.mil

Yong Yang, Manuel Quezada De Luna
Texas A&M University
yytamu@math.tamu.edu, mquezada@math.tamu.edu

MS75

The Shifted Interface Method for Multiphase Flows Computations

In modern engineering, numerical simulations with complex geometries remain challenging. Numerical techniques such as embedded and immersed boundary methods are gaining popularity since they avoid the burden of conformal mesh generation. However, the proper imposition of the boundary conditions is cumbersome since these methods do not preserve the optimal rate of convergence of the numerical schemes. To circumvent this issue, we propose to use the Shifted Interface Method. The key feature of the proposed approach is to shift the location of the interface to a surrogate interface, for which the interface conditions are appropriately modified and weakly enforced. In this work, the closest point projection on one side of the interface is used to define the surrogate interface, and a Taylor expansion enables to modify in a proper way the interface conditions. This method is shown to be favorable in terms of accuracy, robustness, and computational cost as demonstrated for the imposition of jump interface conditions in the case of moving interfaces problems such as multiphase flows with surface tension, using a front tracking or a front capturing algorithm.

Mehdi Khalloufi
Duke Univeristy
mehdi.khalloufi@duke.edu

Kangan Li, Guglielmo Scovazzi
Duke University
kangan.li@duke.edu, guglielmo.scovazzi@duke.edu

MS75

Coupling Multi-material Flows and Structures using Embedded Methods on ALE Grids

We present a computational framework for solid-gas-liquid interactions. Two different strategies are utilized to handle the interfaces between materials. On the solid/fluids interface, the computational grid for fluids confirms with the Lagrangian solid mesh, hence our method employs an arbitrary Lagrangian-Eulerian (ALE) procedure to update the fluids' solutions. In order to capture the material interface between different fluids, however, an embedded boundary is used for capturing the large deformation and topological changes, for example, due to the Richtmyer-Meshkov instabilities. The embedded fluid material interfaces are captured implicitly by a projection level set method; and the transmission condition between different fluids are enforced by a ghost-fluid-like procedure equipped with two-material Riemann solvers, which enable us to handle very large density ratios as well as fluids in extreme conditions such as close to zero temperature.

Xianyi Zeng
University of Texas El Paso
Mathematical Sciences Dept.
xzeng@utep.edu

MS76

Robust Structure-exploiting Multilevel Solvers

Modern architectures continue to challenge the scalability and performance of multilevel solvers for the linear systems arising in physics-based simulations. Although algebraic multigrid methods are robust for solving a variety of problems, the increasing importance of data locality and cost of data movement in modern architectures motivates the need to carefully exploit structure in the problem. This need is renewing interest in meshing strategies that retain or enhance structure in the model's discrete representation, including single- and multi-block logically structured curvilinear meshes, overset meshes and block-structured adaptive mesh refinement. Robust single-block logically structured variational multigrid methods, such as Black Box Multigrid (BoxMG), maintain structure throughout the multigrid hierarchy. This avoids indirection and increased coarse-grid communication costs typical in parallel algebraic multigrid. Redistribution of coarse-grid problems, guided by a predictive performance model can be used to balance the cost of communication and computation on coarser grids to achieve excellent weak scaling. We demonstrate this guided redistribution algorithm on single-block meshes, and comment on its use in multilevel solvers

for multi-block and overset meshes.

J. David Moulton
Los Alamos National Laboratory
moulton@lanl.gov

Andrew Reisner
University of Illinois at Urbana-Champaign
areisne2@illinois.edu

Luke Olson
University of Illinois Urbana-Champaign
lukeo@illinois.edu

MS76

Black Box Multigrid for An Exact Penalty Formulation of Resistive MHD

Magnetohydrodynamics (MHD) provides a continuum level description of conducting fluids in the presence of electromagnetic fields. MHD is a useful tool for modelling astrophysical phenomenon, as well as controlled thermonuclear fusion experiments. In general, the governing equations of MHD are highly-coupled, nonlinear and non self-adjoint. For this reason, the linear algebraic system resulting from a discretization process is difficult to solve with generic iterative methods. One approach to rectify this issue is to construct a physics based preconditioner that utilizes a priori knowledge of the problem as an inner iteration for a Newton-Krylov solver. In this work we explore a variant of black box multigrid as a preconditioner for sub-blocks in the MHD sytem. Black box multigrid sits somewhere in between purely geometric and purely algebraic multigrid methods. The intergrid transfer operators are motivated by approximating the system Schur complement on coarser meshes. Despite its algebraic theoretical basis, black box multigrid does expect a regular discretization (logical (i, j, k) structure in 3D, for example).

Ari Rappaport
Sandia National Laboratories
Computational Mathematics
aerappa@sandia.gov

John Shadid
Sandia National Laboratories
Albuquerque, NM
jnshadi@sandia.gov

Ray S. Tuminaro
Sandia National Laboratories
Computational Mathematics and Algorithms
rstumin@sandia.gov

Jehanzeb H. Chaudhry
University of New Mexico
jehanzeb@unm.edu

Luc Berger-Vergiat
Sandia National Labs
lberge@sandia.gov

MS76

HyTeG: A High Performance Multigrid Framework on Hybrid Meshes

In this talk we will present HyTeG, the successor of the HHG framework. We utilize the well established idea of combining fully unstructured triangular or tetrahedral meshes with fast matrix-free multigrid solvers. Excellent performance is achieved by applying structural refinement inside each primitive which enables the use of stencil codes. Furthermore, this hierarchy of meshes also provides a natural basis for multigrid methods. By employing a fully distributed data structure, the foundation for perfect scalability and advanced load balancing techniques is created.

Dominik Thoennes, Nils Kohl
University of Erlangen-Nuremberg
dominik.thoennes@fau.de, nils.kohl@fau.de

Daniel Drzisga
Technical University of Munich
drzisga@ma.tum.de

Dominik Bartuschat
University of Erlangen-Nuernberg
dominik.bartuschat@fau.de

Ulrich J. Ruede
University of Erlangen-Nuremberg
Department of Computer Science (Simulation)
ulrich.ruede@fau.de

MS76

An Algebraic Multigrid Method Tailored to Semi-structured Grids

Hierarchical hybrid grids (HHG) have been proposed in conjunction with multigrid solvers to promote the efficient utilization of high performance modern computer architectures[Gmeiner, Kostler, Sturmer, and Rude, *Parallel multigrid on hierarchical hybrid grids: a performance study on current high performance computing clusters*, Concurrency and Comp.: Practice and Experience, 26 (2014), pp. 217–240.]. While HHG meshes provide some flexibility for unstructured applications, most of the multigrid calculations can be accomplished using structured grid ideas and kernels. In this paper, we generalilze the HHG idea in two ways to facilitate its adoption within the unstructured finite element community. First, we introduce an HHG framework that significantly reduces the overall coding effort required to adapt mature finite element applications to an HHG solver. This is accomplished by applying a dis-assembly process to a standard finite element matrix. Second, we modify the HHG solver so that it can address significantly more complex meshes. In particular, conventional HHG meshes are constructed by uniformly refining an unstructured mesh. Instead, we consider extensions such that the underlying finest mesh can contain a few completely unstructured regions. These extensions allow finite element practitioners to use unstructured meshes in specific areas (e.g., near jagged interfaces) where they are most natural.

Ray S. Tuminaro
Sandia National Laboratories
Computational Mathematics and Algorithms
rstumin@sandia.gov

Matthias Mayr
Bundeswehr University Munich
matthias.mayr@unibw.de

Luc Berger-Vergiat

Sandia National Labs
lberge@sandia.gov

Peter Ohm
Tufts University
peter.ohm@tufts.edu

MS77

Coupling Multirate and Parallel-in-time Approaches

On the one hand, parareal is a modern algorithm, which enables the usage of multigrid method in time domain. In fact, it can be thought of as a specific version of a multiple shooting algorithm. On the other hand, coupled systems do often exhibit widely separated, inherent time scales in the subsystems. Now, multirate approaches aim at employing adapted time steps for the subsystems in order to gain efficiency of the overall computation. In this talk, we approaches are discussed and we explore the multirate setting within the application of the parareal method. The simplest version of such an approach is the usage of an outer cosimulation-like procedure.

Andreas Bartel
University of Wuppertal
bartel@math.uni-wuppertal.de

MS77

Multirate Schemes for Differential-Algebraic Equations

In refined network analysis, after semidiscretizing the lumped part described by a PDE model using the method-of-lines approach, one is faced with the problem to simulate a large system of coupled differential-algebraic equations (DAEs). One approach to increase efficiency is to exploit the different time scales of the subsystem by applying multirate schemes. This talk will discuss multirate schemes for index-1 DAE arising in refined network analysis including convergence and stability analysis.

Michael Günther
Bergische Universität Wuppertal
Fachbereich C
guenther@math.uni-wuppertal.de

MS77

Multirate Parareal with Application in Electrical Engineering

Simulation of various physical phenomena often requires coupling of systems described by different formulations which exhibit multirate behaviour. This is, for example, the case for electrical machinery modelled by partial differential equations and connected to its surrounding circuitry given by a network model. Often the circuit model is quickly switching, whereas the machinery possesses a rather slowly-varying dynamics. In that case a classical single-rate treatment is not efficient. In the current contribution we propose a specific time-parallel approach for this application, based on the Parareal algorithm. It allows to speed-up the computational time by parallelising the time-domain simulation using a coarse and a fine propagator. Parareal has been recently applied to the simulation of an electric machine by the authors in [S. Schops, I. Niyonzima, and M. Clemens. Parallel-in-time Simulation of Eddy Cur- rent Problems using Parareal, IEEE Trans. Magn., 54 (3), 2018.]. A novel variant, suitable for systems excited with a (highly-oscillatory) pulse-width-modulated signal, was introduced in [M. J. Gander, et al. A new Parareal algorithm for problems with discontinuous sources, 2018. Submitted, https://arxiv.org/abs/1803.05503.]. This talk is about efficient choices of coarse propagators in the presence of the multirate phenomena.

Iryna Kulchytska-Ruchka
Graduate School of Computational Engineering
TU Darmstadt
kulchytska@gsc.tu-darmstadt.de

Idoia Garcia
Technische Universitat Darmstadt
cortes@gsc.tu-darmstadt.de

Sebastian Sch
Theorie Elektromagnetischer Felder (TEMF) and GSCE
Technische Universit
schoeps@gsc.tu-darmstadt.de

MS77

Multi-rate Time Integration on Overset Meshes

Overset meshes are an effective tool for the computational fluid dynamic simulation of problems with complex geometries or multiscale spatio-temporal features. When the maximum allowable timestep on one or more meshes is significantly smaller than on the remaining meshes, standard explicit time integrators impose inefficiencies for time-accurate calculations by requiring that all meshes advance with the smallest timestep. With the targeted use of multi-rate time integrators, separate meshes can be time-marched at independent rates to avoid wasteful computation while maintaining accuracy and stability. This work applies time-explicit multi-rate integrators to the simulation of the compressible Navier-Stokes equations discretized on overset meshes using summation-by-parts (SBP) operators and simultaneous approximation term (SAT) boundary conditions. We introduce a novel class of multi-rate Adams-Bashforth (MRAB) schemes that offer significant stability improvements and computational efficiencies for SBP-SAT methods. We present numerical results that confirm the numerical efficacy of MRAB integrators, outline a number of outstanding implementation challenges, and demonstrate a reduction in computational cost enabled by MRAB. We also investigate the use of our method in the setting of a large-scale distributed-memory parallel implementation where we discuss concerns involving load balancing and communication efficiency.

Cory Mikida
University of Illinois at Urbana-Champaign
cmikida2@illinois.edu

Andreas Kloeckner
University of Illinois
andreask@illinois.edu

Dan Bodony
University of Illinois at Urbana-Champaign
bodony@illinois.edu

MS78

Topology Optimization of Heat Exchanging Com-

ponents and Networks

In this presentation, we discuss topology optimization as a design tool for compact water-cooled heat sinks and latent heat storage devices. Simulating the heat transport in these applications by itself typically already incurs high computational costs, since it involves solving a coupled system of flow and thermal partial differential equations. With current manufacturing flexibility, finding the design that best enables the full potential of such a device is a tremendous challenge. Therefore, both considered applications benefit from the usage of adjoint-based numerical optimization methods to find innovative configurations. Nevertheless, the validity of the underlying flow and heat transport models is often challenged by the optimization procedure. Care is therefore needed when translating the design problem into a numerical optimization problem. To illustrate this, we demonstrate how ineffective designs with disconnected channels are obtained by numerically optimizing the lay-out of a pressure-driven heat sink and offer a solution. Furthermore, it is shown how numerical optimization methods can be applied to efficiently optimize fin arrays in heat sinks. By replacing direct numerical simulation of Navier-Stokes equations for the fin arrays by macro-scale models, computational costs are strongly reduced, while the optimization trends are maintained. Finally, it is shown how adjoint-based methods can be adapted to the topological optimization of district heating networks.

Maarten Blommaert, Martine Baelmans, Geert Buckinx, Wouter Dekeyser, Paul Lacko, Bart Peremans, Yang Zhou
KU Leuven
maarten.blommaert@kuleuven.be, tine.baelmans@kuleuven.be, geert.buckinx@kuleuven.be, wouter.dekeyser@kuleuven.be, paul.lacko@kuleuven.be, bart.peremans@kuleuven.be, yang.zhou@kuleuven.be

MS78

Coupled Discrete Adjoint Computations for CFD-related Multiphysics Problems

Obtaining gradients from discrete adjoint solutions have been proven to be a reliable and precise method for CFD-related optimization problems. To keep up with the recent developments, discrete adjoints do not only have to incorporate more and more complicated turbulence models but also have to cover multiphysics effects and couplings.
For example, these could be the change of the fluid flow domain due to tractions exterted on its boundaries in case they represent an interface to an elastic body or the heating and cooling effects at it, if it can conduct heat as well.
Let us abstractly denote a computer program driving such a simulation by a fluid and a solid zone iterator, $\mathcal{G}_F$ and $\mathcal{G}_S$, mapping the current vectors of conservatives $\left(U_F^i, U_S^i\right) =: U^i$ to U^{i+1} in order to obtain a steady-state solution U^*.
The fixed point iteration we set up to obtain the discrete adjoint solutions λ_F and λ_S with respect to an objective function J then takes the form

$$\left(\lambda_F^{i+1}, \lambda_S^{i+1}\right) = \frac{\partial}{\partial U} J(U^*) + \frac{\partial}{\partial U}\mathcal{G}_F^\top(U^*)\cdot\lambda_F^i + \frac{\partial}{\partial U}\mathcal{G}_S^\top(U^*)\cdot\lambda_S^i,$$

where all terms are evaluated by algorithmic differentiation in reverse mode, taking into account all cross terms like $\frac{\partial}{\partial U_S}\mathcal{G}_F \cdot \lambda_F^i$ in order to derive accurate gradients.

Ole Burghardt
TU Kaiserslautern
Germany
ole.burghardt@scicomp.uni-kl.de

Tim Albring, Nicolas R. Gauger
TU Kaiserslautern
tim.albring@scicomp.uni-kl.de, nicolas.gauger@scicomp.uni-kl.de

MS78

A Modified Search Direction Method with Weakly Imposed Karush-Kuhn-Tucker Conditions using Singular-value Decomposition for Large Constrained Optimization Problems

In this work, a modified search direction method with weakly imposed Karush-Kuhn-Tucker conditions (KKT) is presented, in which we apply the singular-value decomposition method for both the objective and constraint sensitivity in feasible domain. We propose an alternative search direction in stead of the steepest descent direction although there is no constraint active, such that the KKT gradient condition is approached iteratively during the optimization process. Various 2D inequality constraint examples demonstrate that the proposed method finds local minima directly and avoids traveling along the constraints. The results are shown for shape optimization problems with large numbers of design variables. In order to robustly deal with complex geometries, the Vertex Morphing method is used.

Long Chen
TU Munich
Germany
long.chen@tum.de

MS78

Numerical Optimization for Aeroacoustics

A hybrid noise prediction framework is developed in which a permeable surface Ffowcs Williams and Hawkings (FW-H) Equation solver is implemented and coupled with an unsteady Reynolds-averaged Navier-Stokes (URANS) solver.This framework is applied to a number of 2-D and 3-D noise m inimization cases via shape optimization. The lift and noise design objectives were shown to be competing in all cases studied noise minimization always leads to a marked loss of lift. Lift-constrained noise minimization were performed for all 2-D cases and shown to be able to successfully constrain the mean lift at its baseline level while still reducing noise. A number of unconventional optimal designs were obtained, including an airfoil design with wavy surfaces to reduce wake interaction noise. In the 3-D case, the baseline and optimized designs were also analyzed using a turbulence-resolving delayed detached-eddy simulation (DDES). The results indicate that the optimal configuration determined by the URANS-based optimization also performs well when analyzed with DDES and even exceeds the prediction of the original URANS simulation.

Beckett Zhou, Tim Albring, Nicolas R. Gauger
TU Kaiserslautern
yuxiang.zhou@rhrk.uni-kl.de, tim.albring@scicomp.uni-kl.de, nicolas.gauger@scicomp.uni-kl.de

MS79

A New Spectral Method for Elliptic Problems in

General Domains

An efficient spectral solver is developed for elliptic partial differential equations in complex domains. By a fictitious domain approach, an extended problem is built by embedding the original domain into an enlarged and regular one. Two types of Petrov-Galerkin formulations with special trial and test functions are presented and analyzed. A fast algorithm is developed for solving the corresponding linear system. We present ample numerical results to show the efficiency and accuracy of this approach.

Yiqi Gu
Purdue University
yiqigu1989@gmail.com

Jie Shen
Dept. of Mathematics
Purdue University
shen7@purdue.edu

MS79

Exact Smooth Piecewise Polynomial Sequences on Alfeld Splits

We develop exact polynomial sequences on Alfeld splits in any spatial dimension and any polynomial degree. An Alfeld split of a tetrahedron is obtained by connecting the vertices of an n-simplex with its barycenter. We show that, on these triangulations, the kernel of the exterior derivative has enhanced smoothness. Byproducts of this theory include characterizations of discrete divergence-free subspaces for the Stokes problem, commutative projections, and simple formulas for the dimensions of smooth polynomial spaces.

Michael J. Neilan
University of Pittsburgh
Department of Mathematics
neilan@pitt.edu

Johnny Guzman
Brown University
johnny_guzman@brown.edu

Guosheng Fu
Division of Applied Mathematics
Brown University
guosheng_fu@brown.edu

MS79

Primal-dual Weak Galerkin Finite Element Methods

Weak Galerkin (WG) finite element method is a numerical technique for PDEs where the differential operators in the weak form are reconstructed by using a framework that mimics the theory of distributions for piecewise polynomials. The usual regularity of the approximating functions is compensated by carefully-designed stabilizers. The fundamental difference between WG and other existing finite element methods is the use of weak derivatives and weak continuities in the design of numerical schemes based on conventional weak forms for the underlying PDE problems. Due to its great structural flexibility, WG methods are well suited to a wide class of PDEs by providing the needed stability and accuracy in approximations. The speaker will present a recent development of WG, called "Primal-Dual Weak Galerkin (PD-WG)", for problems for which the usual numerical methods are difficult to apply. The essential idea of PD-WG is to interpret the numerical solutions as a constrained minimization of some functionals with constraints that mimic the weak formulation of the PDEs by using weak derivatives. The Euler-Lagrange equation offers a symmetric scheme involving both the primal variable and the dual variable (Lagrange multiplier). PD-WG method is applicable to several challenging problems for which existing methods may have difficulty in applying; these problems include the second order elliptic equations in nondivergence form, Fokker-Planck equation, and elliptic Cauchy problems.

Junping Wang
National Science Foundation
jwang@nsf.gov

MS79

A Charge-conservative Mixed Finite Element Method for Inductionless MHD Equations

A charge-conservative mixed finite element method is proposed for inductionless and incompressible magnetohydrodynamic (MHD) equations in three dimensions. The divergence-free property of the electric current density, or the conservation of charges, plays an important role in simulating liquid metals like the lithium-lead experimental blanket of a TOKAMAK. The discrete current density by our mixed method is divergence-free exactly and globally in the domain. The second objective is to design a robust and quasi-optimal solver for the discrete problem. We propose a block preconditioner for the linearized algebraic system of discrete MHD problem. By extensive numerical examples for both stationary MHD problem and time-dependent MHD problem, we demonstrate the robustness of the preconditioner to Reynolds number and the quasi-optimality to meshes.

Weiying Zheng
Academy of Mathematics and Systems Science
Chinese Academy of Sciences
zwy@lsec.cc.ac.cn

MS80

Applications of Tensor Decompositions at Amazon

Abstract not available.

Anima Anandkumar
Engineering and Applied Sciences
CalTech
anima@caltech.edu

MS80

Detecting Malware using Tensor Decompositions

Abstract not available.

Charles Nicholas
University of Maryland-Baltimore County
nicholas@umbc.edu

MS80

Tensor Methods and Recommender Systems

Abstract not available.

Ivan Oseledets

Skolkovo Institute of Science and Technology
ivan.oseledets@gmail.com

MS80

Tensor Factorizations with Applications to Pattern and Topic Detection

Abstract not available.

Evangelos Papalexakis
University of California, Riverside
epapalex@cs.ucr.edu

MS81

Hybrid a Posteriori Error Estimators for Finite Element Approximations

Abstract not available.

Difeng Cai
Purdue University
cai92@purdue.edu

MS81

A Posteriori Error Estimates for the Virtual Element Method

A posteriori error analysis for the virtual element method (VEM) applied to a model Poisson problem is presented. The local error indicator is constructed based on a recovered equilibrated flux in the $H(\mathrm{div})$-conforming VEM space. Even though the solution itself is non-computable, the post-processed flux can be computed relying only on quantities available, namely the degrees of freedom and element-wise polynomial projections. An h-refinement strategy for polytopal meshes and hanging node automation are also discussed. This is a joint work with Long Chen.

Shuhao Cao
University of California, Irvine
shuhao.cao@uci.edu

MS81

Adaptive Least-squares Finite Element Methods for Linear Transport Equations

Abstract not available.

Shun Zhang
City University of Hong Kong
Department of Mathematics
shun.zhang@cityu.edu.hk

MS81

Title Not Available

Abstract not available.

Xu Zhang
Mississippi State University
xuzhang@math.msstate.edu

MS82

Session Details TBA

Input your abstract, including TeX commands, here. The abstract should be no longer than 1500 characters, including spaces. Only input the abstract text. Don't include title or author information here.

Mary Ann E. Leung
Sustainable Horizons Institute
mleung@shinstitute.org

MS83

A Unsteady Discrete Adjoint Solver for RANS and Hybrid-RANS/LES

Many fluid flows in industry are inherently unsteady and turbulent. RANS models provide results for mean quantities for a wide range of steady and unsteady flows. Often the modeling assumptions of the RANS models are no longer valid. In situations like that LES performs much better. Unfortunately it also increases the costs by a factor of 10 to 100 compared to the RANS computations. Hence there is a big interest in hybrid RANS/LES methods, to perform LES only where it is needed while using RANS in regions where it is reliable and efficient. Amongst others, here the so-called Detached Eddy Simulations (DES) proofed to have great potential in resolving the flow even behind full car and aircraft configurations. Going one step further, the natural question to ask is whether it is possible to extend these methods into the field of design optimization. In that regard it is possible to identify several challenges that have to be faced in order to be feasible for industrial applications. Topics like memory management, speed, consistency and the fundamental problem of the chaotic behavior of dynamical systems needs to be addressed and understood. Within the scope of this talk we want to investigate some of the problems mentioned above using the recent combination of the unsteady adjoint solver and the DDES features introduced in the open-source CFD solver SU2.

Nicolas R. Gauger, Max Sagebaum, Tim Albring
TU Kaiserslautern
nicolas.gauger@scicomp.uni-kl.de,
max.sagebaum@scicomp.uni-kl.de,
tim.albring@scicomp.uni-kl.de

MS83

Adjoint-based Mesh Optimisation and Optimal Initial Conditions in a Spectral Element Framework

The continuous adjoint framework implemented in the high-order spectral-element code Nek5000 is discussed for two applications: i) error estimators to drive adaptive mesh refinement and ii) computation of optimal initial conditions. In both cases, the base flow (the forward solution) is changing while integrating backwards in time. We use a multi-level checkpoint-and-revolve algorithm, which makes the current implementation to scale even to very large problems. Adaptive mesh refinement (AMR) relies on tools for automatic mesh refinement and error estimators which allow for optimal error control. We highlight the capability of Nek5000 to use h-refinement for mesh adaptation, coupled to goal-oriented estimator that takes into account both the local properties of the solution and the global dependence of the error in a functional of interest. We demonstrate their relation to more straightforward spectral error estimators, and discuss aspects of implementation. Optimal initial conditions are a useful tool to analyse flows undergoing transition to turbulence. These so-called "optimals" are flow structures that exhibit the most growth in time. The computation of these structures relies on the

minimisation of the energy of the initial condition, while still reaching turbulence. Adjoint methods are used to calculate the gradients, and the update is a variant of (rotated) conjugate gradients. We show the efficiency on the example of localised optimals in curved pipe flow.

Philipp Schlatter
KTH Royal Institute of Technology
pschlatt@mech.kth.se

Nicolas Offermans
KTH, Sweden
nof@mech.kth.se

Oana Marin, Elia Merzari
Argonne National Laboratory
oanam@anl.gov, emerzari@anl.gov

MS83

Large-scale PDE Constrained Optimization from the Perspective of a Library (PETSc)

From the perspective of production codes, optimization and control have long been regarded as add-ons to an existing large-scale solver which preserves its structure and parallelism. To facilitate the computation of adjoints, external libraries based on Automatic Differentiation (AD) have been included such that the forward solver is not altered in any significant way. The PETSc library, which was essentially a numerical algebra library, has newly added features for time integration, mesh handling, and adjoint capabilities, aiming to provide an integrated package for PDE-constrained optimization. The advantage PETSc holds over other existing libraries is the embedded TAO optimization solvers. This allows a user to port a solver already finely tuned for efficient computation of the spatial components into PETSc and from hereon solve a PDE-constrained optimization time-dependent (or independent) problem using all the tools available in PETSc (AD, optimization routines, etc.).

Barry F. Smith, Oana Marin
Argonne National Laboratory
bsmith@mcs.anl.gov, oanam@anl.gov

Emil M. Constantinescu
Argonne National Laboratory
Mathematics and Computer Science Division
emconsta@mcs.anl.gov

MS83

Obtaining Discrete Adjoint Sensitivities using Algorithmic Differentiation

In the context of CFD, most researchers in the past have applied algorithmic differentiation (AD) and adjoint methods selectively to certain numerical kernels or specifically tailored to a specific problem. In contrast our work pursues the opposite approach, to initially cover as much as possible of the used CFD framework by AD. The advantage of such an approach is two-fold: First, a initial differentiated version of the problem can be obtained very rapidly and without much analytical insight into the underlying problem. Starting from there, possible optimizations can then be identified, applied, and evaluated. Second, a full AD implementation gives the flexibility to pursue a wide range of optimization tasks, without needing to adapt the underlying CFD framework to each of these applications. To facilitate this a discrete adjoint framework, implemented in OpenFOAM, was developed. OpenFOAM is particularly suited for this approach, due to its split architecture into a general CFD framework, (which was fully covered by AD), and individual solvers and utilities based on the framework (which were covered by AD as needed). The application of AD to a full CFD iteration loop poses challenges, especially in terms of memory consumption. The talk will present methods to overcome those challenges and show case studies that demonstrate the feasibility of the approach.

Markus Towara
STCE, RWTH Aachen University
towara@stce.rwth-aachen.de

MS84

Free Boundary Solvers on Octree Grids and in Parallel Environments

Abstract not available.

Frederic Gibou
UCSB
fgibou@engineering.ucsb.edu

MS84

A Multi-scale and Multi-physics Tissue Simulation Schema to Assess Emergent Rules of Complex Systems

Abstract not available.

Pouria Mistani
UCSB
pouria@ucsb.edu

MS84

Pressure Based Compressible Solver for Two-phase Flows with Arbitrary Equation of State and Entropic Effects

Abstract not available.

Sebastien Tanguy
Universite de Toulouse
sebastien.tanguy@imft.fr

MS84

Parallel Octree Simulations of Incompressible Flows Coupled with Soluble Surfactant in a Level-set Framework

Abstract not available.

Fernando Temprano-Coleto
UCSB
ftempranocoleto@ucsb.edu

MS85

Algebraic Multigrid for Hypersonic Simulations

With an ever growing number of space launch and vehicles re-entering the atmosphere every year, hypersonic flight is an important research area. While the simulation of transonic and supersonic flow has been reasonably well studied, it is not the case for hypersonic flow and the application of multigrid to solve the linear systems of equations resulting

from the discretization of such problems is still an open area of research. In this talk I will present current efforts at Sandia to use MueLu, our multigrid framework, as a preconditioner to GMRES in SPARC, Sandia's hypersonic simulation code. Past and current preconditioning strategies will be demonstrated on examples and future work will be outlined.

Luc Berger-Vergiat
Sandia National Labs
lberge@sandia.gov

Ray S. Tuminaro
Sandia National Laboratories
Computational Mathematics and Algorithms
rstumin@sandia.gov

Jonathan J. Hu
Sandia National Laboratories
Livermore, CA 94551
jhu@sandia.gov

Chris Siefert
Sandia National Laboratories
csiefer@sandia.gov

Michael Brazell
University of Wyoming
mbrazell@uwyo.edu

Dimitri Mavriplis
Department of Mechanical Engineering
University of Wyoming
mavripl@uwyo.edu

MS85

Parallel High-order Space-time AIR for Hyperbolic PDEs

Parallel time integration of partial differential equations (PDEs) is becoming more prominent as core counts of high-performance machines rapidly increase and strong scaling limits in space are consequently reached. Unfortunately, there has been relatively little success in the parallel time integration of hyperbolic PDEs, with the most promising results pertaining to their parabolic counterparts. A promising algebraic multigrid method, known as Approximate Ideal Restriction (AIR) multigrid, was recently developed for upwind-type discretizations. Given that temporal discretizations are naturally upwind, we consider the possibility of using AIR as a solver for the space-time linear systems that arise from the discretization of hyperbolic PDEs. In particular, we consider the popular high-order weighted essentially non-oscillatory spatial discretization. In this nonlinear case, AIR is used as a key component of a Newton-multigrid solver. We will discuss progress to date, including the solution of Burgers' equation and the associated difficulties that arise when its solution is non-smooth.

Hans De Sterck
Monash University, Australia
Applied Mathematics
hdesterck@uwaterloo.ca

Oliver A. Krzysik
Monash University
oliver.krzysik@monash.edu

Scott Maclachlan
Department of Mathematics and Statistics
Memorial University of Newfoundland
smaclachlan@mun.ca

Ben S. Southworth
Department of Applied Mathematics
University of Colorado, Boulder
ben.southworth@colorado.edu

MS85

A Convergence Theory for Algebraic Multigrid (AMG) Applied to Nonsymmetric Systems

Algebraic multigrid (AMG) is one of the fastest numerical methods for solving large sparse linear systems, often arising in the study of graph Laplacians, Markov chains, and the discretization of partial differential equations (PDEs). For symmetric positive definite (SPD) matrices, convergence of AMG in the A-norm is well motivated, and AMG has proven to be an effective solver for many applications. Recently, several AMG algorithms have been developed that are effective on highly nonsymmetric linear systems. Although motivation was provided in each case, a general convergence theory for AMG applied to nonsymmetric linear systems is lacking. Effective algorithms are based largely on heuristics and/or incomplete theory. A handful of works have delved into convergence of nonsymmetric AMG, but there has yet to be a thorough study on conditions for convergence and, in particular, the practical implications for solver development. In this talk, a first step toward such a theory is presented. Why theory for SPD systems breaks down in the nonsymmetric setting is discussed and a general framework for convergence of nonsymmetric AMG is developed. The classical multigrid weak and strong approximation properties are generalized to a "fractional approximation property," and sufficient conditions are developed for two-grid and multigrid convergence in the $(A^*A)^{1/2}$-norm.

Thomas Manteuffel
University of Colorado
tmanteuf@colorado.edu

Ben S. Southworth
Department of Applied Mathematics
University of Colorado, Boulder
ben.southworth@colorado.edu

MS85

Nonsymmetric Algebraic Multigrid Based on Approximate Ideal Restriction (AIR)

A new algebraic multigrid (AMG) method has been developed recently based on an approximate ideal restriction (AIR), which has demonstrated potential as a fast, parallel solver for hyperbolic-type PDEs. In the hyperbolic limit of no diffusion, AIR is effectively a reduction-base algorithm. However, as diffusive components are introduced, classical AMG ideas based on accurately interpolating smooth modes must be incorporated into the algorithm as well. In this talk, we discuss the convergence properties of AIR, motivating choices in its development, and explaining the strong convergence attained on difficult scalar problems. In particular, new theory sheds light on how to attain scalable convergence and how to tune AIR for different types of problems.

Ben S. Southworth

Department of Applied Mathematics
University of Colorado, Boulder
ben.southworth@colorado.edu

Thomas Manteuffel, John Ruge
University of Colorado
tmanteuf@colorado.edu, jruge@colorado.edu

MS86

Software Library for Accelerating Chemical Kinetics on Hybrid Architectures: The SLACKHA Project

Combustion simulations with finite-rate chemistry rely on accurate and efficient methods for solving stiff ordinary differential equations (ODEs). Typical reacting-flow solvers decouple the ODEs involving chemical kinetics at each spatial location by operator splitting, allowing each to be solved concurrently. Efficient ODE solvers must take into account both numerical efficiency as well as the available parallelism of the underlying computational hardware being used to perform the simulations, especially on many-core coprocessors. This talk will summarize work on a collaborative effort developing techniques and software to reduce the computational expense of chemical kinetics on modern processing architectures. First, we will review past work studying explicit, exponential, and implicit integrators implemented for various many-core devices and using single instruction, multiple thread (SIMT) and single instruction, multiple data (SIMD) paradigms. We will next present and discuss results using SIMD and SIMT to evaluate chemical kinetic source terms and sparse Jacobian matrices, implemented in the open-source software tool pyJac. Then, we will discuss the efficacy of applying stiffness metrics to chemical kinetics ODEs to predict computational cost of integration. Finally, we will identify remaining questions and directions for future research, and describe related efforts applying the tools developed here to applications beyond combustion, such as ocean biogeochemistry.

Kyle E. Niemeyer, Andrew Alferman, Emily F. Klee
Oregon State University
Kyle.Niemeyer@oregonstate.edu,
alfermaa@oregonstate.edu, kleeem@oregonstate.edu

Nicholas Curtis
University of Connecticut
Department of Mechanical Engineering
nicholas.curtis@uconn.edu

Chih-Jen Sung
University of Connecticut
cjsung@engr.uconn.edu

MS86

Towards a Performance Portable Multi-fluid Plasma Capability

Supporting scalable and performant finite element simulations across next generation architectures can add significant code complexity. This presentation will discuss the design of general finite element assembly tools applied to magnetohydrodynamics, multi-fluid plasma and coupled hybrid particle-in-cell simulations. Performance portability is achieved via the Kokkos programming model. The assembly library uses a directed acyclic graph for composable physics kernels in a multiphysics setting. Embedded automatic differentiation, applied via templates and operator overloading, is used for generating machine precision sensitivities for implicit and IMEX solvers. Performance results will be shown for Intel Phi, NVIDIA GPU and Intel Haswell architectures.

Roger Pawlowski
Multiphysics Simulation Technologies Dept.
Sandia National Laboratories
rppawlo@sandia.gov

Eric C. Cyr
Computational Mathematics Department
Sandia National Laboratotories
eccyr@sandia.gov

Edward G. Phillips, Matthew Bettencourt
Sandia National Laboratories
egphill@sandia.gov, mbetten@sandia.gov

Eric Phipps
Sandia National Laboratories
Optimization and Uncertainty Quantification Department
etphipp@sandia.gov

Christian Trott, Paul Lin
Sandia National Laboratories
crtrott@sandia.gov, ptlin@sandia.gov

John Shadid
Sandia National Laboratories
Albuquerque, NM
jnshadi@sandia.gov

MS86

Solving Evolutionary Differential Equations on Heterogeneous Architectures

Modern hybrid heterogeneous architectures, consisting of a blend of manycore CPUs and accelerators (such as GPUs), have promised to circumvent traditional scalability and speedup barriers encountered by their homogenous CPU-based counterparts. However, they fall short in delivering that promise when used in many applications such as those involving evolutionary differential equations. In this talk, we will give an overview of our efforts to improve heterogeneous performance in the context of low-Mach turbulent combustion. Our abstractions are based on an embedded domain specific language within a task-based runtime framework. They enable rapid deployment of physical models at extreme scale without distracting application developers with fine tuning their code for different hardware backends. Furthermore, our runtime system explores concurrency at different granularities all the way from global tasks to fine grained data vectorization. Finally, we will show the role that an algorithm plays in increasing arithmetic intensity and its potential to deliver better performance on heterogeneous computers. The overall message is that algorithm design is no longer dependent on the physics, but also on the computing hardware.

Tony Saad
Department of Chemical Engineering
University of Utah
tony.saad@chemeng.utah.edu

James C. Sutherland
Department of Chemical Engineering
The University of Utah

james.sutherland@chemeng.utah.edu

MS87

The Commutation Error for Large Eddy Simulation Reduced Order Models (LES-ROMs)

We propose reduced order models (ROMs) for an efficient and relatively accurate numerical simulation of nonlinear systems. We use the ROM projection and the ROM differential filters to construct a novel data-driven correction ROM (DDC-ROM). We show that the ROM spatial filtering and differentiation do not commute for the diffusion operator. Furthermore, we show that the resulting commutation error has an important effect on the ROM, especially for low viscosity values. As a mathematical model for our numerical study, we use the one-dimensional Burgers equations with smooth and non-smooth initial conditions.

Birgul Koc
Virginia Tech
birgul@vt.edu

Muhammad Mohebujjaman
Department of Mathematics
Virginia Tech
jaman@vt.edu

Changhong Mou
Virginia Tech
cmou@math.vt.edu

Traian Iliescu
Department of Mathematics
Virginia Tech
iliescu@vt.edu

MS87

Model Reduction and Basis Adaptation for Uncertainty Quantification of Time-dependent PDEs

The last decade has witnessed a surge in the study of the subject of Uncertainty Quantification (UQ) for steady and time-dependent problems. This is an important question because we need to know how the lack of knowledge of the data of a problem (initial/boundary conditions, parameters etc.) propagates to the solution. As it was quickly understood, for many problems of practical interest the popular UQ approaches of Polynomial Chaos Expansions (PCE) or Monte Carlo (MC) simulations become inefficient. This is particularly true for time-dependent problems where the proliferation of activity across scales with time can compound the computational difficulties. We present two novel ways of addressing this issue. First, we have shown that the MZ formalism which has been used before to construct reduced models in a spatial sense can also be used to construct reduced models in the variables needed to describe the randomness (uncertainty). Second, we have extended the use of basis adaptation, a probabilistic dimension reduction method, from steady-state problems to time-dependent ones. Illustrative example will be provided.

Panos Stinis, Ramakrishna Tipireddy
Pacific Northwest National Laboratory
panos.stinis@pnnl.gov, ramakrishna.tipireddy@pnnl.gov

Alexandre M. Tartakovsky
PNNL
alexandre.tartakovsky@pnnl.gov

Jing Li
Pacific Northwest National Lab
jing.li@pnnl.gov

MS87

Ensemble Kalman Filter with Localization

Ensemble Kalman Filter with Localization Xin Tong National University of Singapore, Singapore Ensemble Kalman filter (EnKF) is an important data assimilation method for high dimensional geophysical systems. Efficient implementation of EnKF in practice often involves the localization technique, which updates each component using only information within a local radius. This paper rigorously analyzes the local EnKF (LEnKF) for linear systems, and shows that the filter error can be dominated by the ensemble covariance, as long as 1) the sample size exceeds the logarithmic of state dimension and a constant that depends only on the local radius; 2) the forecast covariance matrix admits a stable localized structure. In particular, this indicates that with small system and observation noises, the filter error will be accurate in long time even if the initialization is not. The analysis also reveals an intrinsic inconsistency caused by the localization technique, and a stable localized structure is necessary to control this inconsistency. While this structure is usually taken for granted for the operation of LEnKF, it can also be rigorously proved for linear systems with sparse local observations and weak local interactions. These theoretical results are also validated by numerical implementation of LEnKF on a simple stochastic turbulence in two dynamical regimes.

Xin T. Tong
National University of Singapore
Department of Mathematics
mattxin@nus.edu.sg

MS87

Accounting for Model Error in ENKF

In this talk, the use of discrete-time stochastic parameterization will be investigated by numerical experiments to account for model error due to unresolved scales in ensemble Kalman filters. The parameterization produces an improved non-Markovian forecast model, which generates high quality forecast ensembles and improves filter performance. Results are compared with the methods of dealing with model error through covariance inflation and localization (IL), using as an example the two-layer Lorenz-96 system.

Xuemin Tu
University of Kansas
xuemin@ku.edu

Alexandre Chorin
Department of Mathematics
University of California at Berkeley
chorin@math.berkeley.edu

Fei Lu
Johns Hopkins University

feilu@math.jhu.edu

MS88

Can Autotuning Compilers become Mainstream?

We describe research on mainstreaming autotuning compiler technology, whereby the compiler automatically explores a search space of alternative implementations of a computation to find the best implementation for a target architecture. Autotuning has demonstrated success in achieving performance portability as it enables the compiler to tailor optimization and code generation to a specific architectural context, starting from the same high-level program specification. Still, mainstream adoption requires availability in widely-used compilers and demonstrated impact on production application codes while under development. This talk will highlight an example of the impact of autotuning compiler technology, recent work on a brick data layout and associated code generator for stencil computations that uses fine-grained data blocking as a tunable abstraction for performance portability across CPUs and GPUs. It also will describe research on migrating autotuning technology into Clang/LLVM to support autotuning of OpenMP and complex loop transformation sequences.

Mary Hall
School of Computing
University of Utah
mhall@cs.utah.edu

Tuowen Zhao
University of Utah
ztuowen@gmail.com

MS88

Transformation-based Code Optimization for Finite Element Methods

To enable performance-driven, user-guided code transformation on an intermediate representation allowing high-fidelity composition, we present a code transformation tool loosely based on the polyhedral framework that enables novel cross-procedural transformations.
Transformation-based code generation is a promising approach to allow near-mathematical notation to be turned into high-performance, parallel code on CPUs and GPUs. This capability has been demonstrated in the literature mostly at the level of a single kernel.
As the complexity and fidelity of models in modern computational science applications increases, composition of an application from smaller pieces becomes increasingly common, in keeping with maintaining separation of concerns. However, many meaningful performance optimizations require a perspective that transcends these component boundaries, for example changes in data layout, cross-loop vectorization, and other complex loop fusion operations. We present an intermediate representation focused on human-in-the-loop code optimization along with transformation vocabulary operating on it that allows the systematic application of transformations to composite programs. We demonstrate the effectiveness of our framework in the context of the Firedrake finite element framework targeting CPU and GPU architectures, n-body calculations, and the discontinuous Galerkin method.

Kaushik Kulkarni
University of Illinois at Urbana-Champaign
kgk2@illinois.edu

Andreas Kloeckner
University of Illinois
andreask@illinois.edu

MS88

Designing for High Performance

Applications are becoming more varied. Architectures are becoming more complex. Yet software implemented by the expert still achieves higher performancethan many automatically generated code despite two decades of research in automatic empirical optimization systems. In this talk, I will discuss how the expert's implementation of matrix matrix multiplication can be captured using analytical models. In addition, I will also show that the same model yields portable performance across a large class of problems that spans multiple application domains.

Tze Meng
Carnegie Mellon University
lowt@andrew.cmu.edu

MS88

A DSL for Program Optimization

The complexity and diversity of todays computer architectures are requiring more attention from the software developers in order to harness all the computing power available. Furthermore, each architecture typically requires a different sequence of optimizations to attain a high fraction of its nominal peak speed. This complicates performance portability and code maintainability. A big challenge is to figure out how to manage different optimized versions of the same code tailored to different architectures and keep them up to date as new algorithmic features are added. We have developed the ICE system, which decouples the performance expert role from the application expert role (separation of concerns). It allows the use of architecture-specific optimizations while keeping the code maintainable in the long term. ICE orchestrates the use of multiple optimization tools to applications baseline version and perform an empirical search to find the best sequence of optimizations and their parameters. The optimizations and the empirical search are specified using a domain-specific language (DSL) in an external file. While other approaches attack the problem of facilitating optimizing applications, ICE contains a unique set of features to simplify the generation of optimized code.

Thiago Teixeira
University of Illinois at Urbana-Champaign
tteixei2@illinois.edu

MS89

Reproducibility and Performance of the Feltor Code on Parallel Architectures

Feltor is both a numerical library and a scientific software package built on top it. Its main target are two- and three-dimensional drift- and gyro-fluid simulations with discontinuous Galerkin methods as the main numerical discretization technique. Feltor allows developing platform independent code that runs on a variety of parallel computer architectures ranging from laptop CPUs to hybrid CPU+GPU distributed memory systems. We investigate reproducibility since we observe that numerical simulations of a recently developed gyro-fluid model produces non-deterministic results in parallel computations. We show

how we can restore bitwise reproducibility algorithmically and programmatically. Furthermore, we explore important performance tuning considerations and discuss latencies and bandwidths of elementary subroutines necessary to implement the aforementioned algorithms and equations. We propose a parallel performance model that predicts the execution time of algorithms implemented in Feltor and test our model on a selection of parallel hardware architectures. We are able to predict the execution time of more complex algorithms with a relative error of less than 25 % for problem sizes between 10^{-1} and 10^3 MB.

Roman Iakymchuk
KTH Royal Institute of Technology
riakymch@kth.se

Matthias Wiesenberger
Technical University of Denmark
mattwi@fysik.dtu.dk

Stef Graillat
University Pierre and Marie Curie (UPMC)
LIP6
stef.graillat@lip6.fr

MS89

Iterative Refinement for Singular Value Decomposition

Let A be a real $m \times n$ matrix. We present a refinement algorithm for the singular value decomposition (SVD) of A. The SVD plays an important role in many areas of scientific computing, such as signal processing, statistical analysis, and so forth. Throughout this talk, we assume that $m \geq n$. If $m < n$, then considering the SVD of A^{T} yields equivalent results. Let $\sigma_i \in \mathbf{R}$, $i = 1, \ldots, n$ denote the singular values of A. Let us consider the (full size) SVD of A such that

$$A = U\Sigma V^{\mathrm{T}}, \quad U \in \mathbf{R}^{m\times m}, V \in \mathbf{R}^{n\times n}, \ \Sigma \in \mathbf{R}^{m\times n},$$

where both U and V are orthogonal and Σ is diagonal with $\Sigma_{ii} = \sigma_i$. For simplicity, we assume that

$$\sigma_1 \leq \sigma_2 \leq \cdots \leq \sigma_n.$$

Recently, the authors have proposed an iterative refinement algorithm for symmetric eigenvalue decomposition [OA2018: T. Ogita, K. Aishima: Iterative refinement for symmetric eigenvalue decomposition, Japan J. Indust. Appl. Math., published online, 2018]. With the same spirit of the paper [OA2018], the structure of the algorithm is straightforward, primarily comprising matrix multiplications. The use of higher-precision arithmetic in our proposed algorithm is restricted to matrix multiplication, which accounts for most of the computational cost. Detailed discussions will be presented in the talk.

Takeshi Ogita
Tokyo Woman's Christian University
ogita@lab.twcu.ac.jp

Kensuke Aishima
Hosei University
aishima@hosei.ac.jp

MS89

Thin QR Decomposition using LU Factors and its Refinement

This talk concerns fast algorithms of thin QR decomposition for an m-by-n full column rank matrix A such that $A = QR$, where an m-by-n matrix Q has orthogonal columns and R is an upper triangular matrix. Recently, CholeskyQR2 algorithm was proposed as the refinement of orthogonality of the computed result by CholeskyQR algorithm [Takeshi Fukaya, et al. : CholeskyQR2: a simple and communication-avoiding algorithm for computing a tall-skinny QR factorization on a large-scale parallel system, Proceedings of the 5th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems. IEEE Press, 2014.]. However, neither CholeskyQR nor CholeskyQR2 algorithm can be applied for ill-conditioned matrices. We propose a preconditioning method of the CholeskyQR algorithm for ill-conditioned matrices using LU factors of A. We focus on the Doolittle LU decomposition that produces the product $A = P^T LU$, where L is a unit triangular matrix, U is an upper triangular matrix, and P is a permutation matrix. Since a condition number of L tends to be of the order of n in many cases, the CholeskyQR algorithms can be applied for the matrix L. Then, $A = P^T Q'SU = QR$, where Q', S are QR factors computed by the CholeskyQR algorithm and $R = SU$. We will talk about a QR decomposition using LU factors and refinement of orthogonality of computed QR factors.

Takeshi Terao
Shibaura Institute of Technology
nb17105@shibaura-it.ac.jp

Katsuhisa Ozaki
Shibaura Institute of Technology / JST
ozaki@sic.shibaura-it.ac.jp

Takeshi Ogita
Tokyo Woman's Christian University
ogita@lab.twcu.ac.jp

MS89

Verified Algorithm for High Order Differentiation Based on Hyper-dual Numbers

A verified algorithm for higher order derivatives using several features of hyper-dual numbers is proposed. Since the calculations of derivatives are used in many applications, there are many previous studies. For example, the automatic differentiation (AD) algorithm is well known as an accurate and efficient algorithm for lower order derivatives, however, for arbitrarily higher order derivatives the AD algorithm quickly grow very complicated. Recently, hyper-dual numbers are proposed by Fike and Alonso. The numbers are an extension of dual numbers, which are based on the non-real term $\epsilon^2 = 0$ where $\epsilon = 0$. Similar to this features, hyper-dual numbers are based on the non-real term $\epsilon_i (1 \leqq i \leqq n)$ where $\epsilon_i \neq 0, \epsilon_i^2 = 0, \epsilon_i \neq \epsilon_j, \epsilon_i\epsilon_j = \epsilon_j\epsilon_i \neq 0 (i \neq j)$. These numbers enable us to get accurate results of second or higher derivatives. In additon to that, the matrix representation of hyper-dual numbers are very useful for calculation of high order derivatives. In the AD algorithm, every operation needs to be made for calculation, however, the matrix representation of hyper-dual numbers enable us to avoid it. In this talk, we propose verification algorithm using features of hyper-dual numbers. Numerical examples are presented for illustrating effectiveness of the proposed algorithm.

Naoya Yamanaka
Meisei University

naoya.yamanaka@meisei-u.ac.jp

MS90

Adaptive Parallel-in-time PDE-constrained Optimization

Adjoint gradient computation for optimization problems governed by parabolic PDEs requires one solve of the state equation and one backward-in-time solve of the adjoint equation, which makes the iterative optimization process extremely costly. With today's modern computers, the time-to-solution can be decreased through massive parallelization, which is traditionally done in the spatial dimensions. In addition, time-parallel methods have received increasing interest in recent years, mainly due to the saturation in processor speeds. Iterative multilevel schemes such as PFASST (Parallel Full Approximation Scheme in Space and Time) are currently state of the art and can achieve significant parallel efficiency. In this talk, we introduce approaches to exploit PFASST for the time-parallel solution of optimization problems with parabolic PDEs. As PFASST is based on spectral deferred correction methods for the time integration, their iterative nature provides additional flexibility, e.g., by re-using previously computed solutions in the optimization loop to reduce SDC iterations required for solving state and adjoint equations at the cost of additional storage. We investigate the influence of inexact storage of solutions for warm-starting SDC iterations, and focus especially on adapting the accuracy of state and adjoint solves to the optimization progress.

Sebastian Götschel
Zuse Institute Berlin
goetschel@zib.de

Michael Minion
Lawrence Berkeley National Lab
mlminion@lbl.gov

MS90

Parallel-in-time Integration of the Shallow-water Equations on the Rotating Sphere

The accurate modeling of atmospheric processes over long periods of time is expensive. The high-order integration of the PDEs requires a large number of time steps, even when the terms accounting for the propagation of fast waves are treated implicitly. Therefore, parallel-in-time integration is attractive to compute multiple time steps concurrently and reduce the time-to-solution. We propose a multi-level parallel-in-time integration method combining the Parallel Full Approximation Scheme in Space and Time (PFASST) with the Spherical Harmonics to solve the shallow-water equations on the sphere. The iterative algorithm interweaves parallel fine corrections and serial corrections performed on a coarsened problem to propagate the updated initial conditions between time steps. We design a consistent coarsening methodology using the spectral basis of the Spherical Harmonics, and we show that this step is key for the accuracy and efficiency of the approach. We test our method with problems characterized by the fast generation of small-scale features in the solution. We study the convergence of the PFASST iterations to evaluate the robustness of the scheme on wave-propagation problems. We assess the impact of the coarsening strategy on the accuracy and efficiency of the approach and on its ability to capture the high-frequency modes accumulating in the solution. We show that PFASST can resolve the main features of the solution significantly faster than with serial SDC.

Francois P. Hamon
Lawrence Berkeley National Laboratory
fhamon@stanford.edu

Martin Schreiber
University of Exeter
College of Engineering, Mathematics and Physical Sciences
schreiberx@gmail.com

Michael Minion
Lawrence Berkeley National Lab
mlminion@lbl.gov

MS90

Tracking Fast Ions in Fusion Reactors using Spectral Deferred Corrections

Simulating the trajectories of fast ions in the magnetic field generated by fusion reactors requires solving the second order Lorentz equations over very long time intervals. The talk will present a variant of SDC based on the widely used Boris algorithm. It will discuss the adoption of convergence accelerators and present results for problems with inhomogeneous magnetic fields, in particular a mirror trap and trajectories in a Solev'ev equilibrium.

Krasymyr Tretiak
School of Mechanical Engineering
University of Leeds
k.tretiak@leeds.ac.uk

James Buchanan
Culham Centre for Fusion Energy
james.buchanan@ukaea.uk

Debasmita Samaddar
CCFE
UK Atomic Energy Authority
Debasmita.Samaddar@ukaea.uk

Rob Akers
Culham Centre for Fusion Energy
rob.akers@ukaea.uk

Daniel Ruprecht
School of Mechanical Engineering
University of Leeds
d.ruprecht@leeds.ac.uk

MS90

Improved Coupling of Hydrodynamics and Nuclear Burning in Astrophysical Flows using SDC

Simulations in stellar astrophysics involve the coupling of hydrodynamics and burning under a wide variety of conditions. Approximations such as Strang splitting are usually made to make the algorithm simpler, but this can comprimise the accuracy of the simulation as the hydrodynamics and reactions can decouple in regions of vigorous burning. Furthermore, Strang splitting does not have a straightforward extension to higher-order integration. We explore spectral deferred correction time-integration methods that improve the coupling of hydrodynamics and reactions. We focus on second- and fourth-order accurate schemes, implemented in the open-source hydrodynamics code Castro.

We show applications to nuclear deflagrations and detonations.

Michael Zingale
Department of Astronomy and Astrophysics
UCSC
michael.zingale@stonybrook.edu

MS91

An Acceleration Framework for Parameter Estimation using Implicit Sampling and Adaptive Reduced-order Models

Parameter estimation for partial differential equations is a computationally demanding endeavor, particularly when adopting a Bayesian setting since it amounts to sampling from a posterior probability distribution involving repeated solution of the governing equations. For efficiency, we use implicit sampling to focus samples in high probability regions of the posterior distribution. However, locating regions of high probability amounts to solving a PDE-constrained optimization, which can be expensive due to repeated queries to the primal and adjoint PDE. To remedy this cost, we replace the PDE by efficient, projection-based reduced-order models, embedded in a globally convergent trust region framework. Once we have located the maximum a posteriori point, the random maps sampling procedure is re-cast as a one-dimensional PDE-constrained optimization problem, which is also solved using the reduced-order model trust region method. The proposed method based on implicit sampling and reduced-order models is shown to substantially reduce the cost of parameter estimation for a subsurface flow model problem.

Robert Baraldi
North Carolina State University
rjbarald@ncsu.edu

Matthew J. Zahr
Lawrence Berkeley National Laboratory
University of California, Berkeley
mjzahr@lbl.gov

MS91

Accelerating Limited-memory Quasi-Newton Convergence for Optimization of PDE-governed Systems

Quasi-Newton methods are key kernels popular in PDE-constrained optimization due to the difficulty of computing exact Hessian information. However, the limited-memory variants commonly used in large-scale applications exhibit linear convergence rates and may require many expensive function evaluations. In this work, we propose a Hessian initialization techniques for quasi-Newton approximations that is based on the restricted Broyden update. We use the full-memory update formula, but construct only the diagonal of the Hessian to minimize the memory footprint of the initialization. A dynamic re-scaling for the Hessian initialization is performed on every update in order to generate well-scaled step directions. We conduct an exhaustive parameter study for convex combination and re-scaling factors in the initialization and analyze its efficacy on a wide range of unconstrained and bound constrained CUTEst problems. Our results demonstrate that the proposed Hessian initialization significantly accelerates convergence and improves robustness.

Alp Dener
Rensselaer Polytechnic Institute
Department of Mechanical, Aerospace, and Nuclear Engineering
adener@anl.gov

Todd Munson
Argonne National Laboratory
Mathematics and Computer Science Division
tmunson@mcs.anl.gov

MS91

Fast Solvers for Mixed-integer PDE-constrained Optimization

Many important design and inverse problems involve both discrete or integer decision variables and PDE constraints. Applications include topology design and certain classes of inverse problems. We review recent new developments that combine clever rounding-based heuristics with trust-region-like search techniques. We show that these methods can be implemented efficiently, and we provide conditions under which they converge to a local minimum. We illustrate the effectiveness of these approaches on a electro-magnetic cloaking problem, and a 3D inverse problem.

Sven Leyffer
Argonne National Laboratory
leyffer@anl.gov

MS91

Scalable Optimization of Discrete-time PDE Systems with Inequality Constraints

A key challenge in solving optimal control problems modeled by time-dependent PDEs is the serial bottleneck of forward and backward (adjoint) time integration. To overcome this challenge, we introduce a multigrid-in-time technique for the solution of optimality systems arising in a matrix-free trust-region sequential quadratic programming (TR-SQP) method for equality-constrained optimization. Additionally, we propose and study extensions of the matrix-free TR-SQP method, including augmented Lagrangian formulations, to handle general inequality constraints. We present numerical results for fluid dynamics control problems with control and state inequality constraints, and examine the weak and strong scaling of the combined multigrid-in-time TR-SQP approach on thousands of compute cores.

Denis Ridzal
Sandia National Laboratories
dridzal@sandia.gov

Eric C. Cyr
Computational Mathematics Department
Sandia National Laboratotories
eccyr@sandia.gov

MS92

Generalized Newton Methods for the Numerical Simulation of Viscoplastic Fluids

In this talk we give a general perspective of the application of generalized Newton methods to the numerical simulation of the Bingham and Herschel-Bulkley flows. Our approach is based on regularizations of the bi-viscosity type, which allows us to develop fast convergent algorithms (mainly superlinear convergent algorithms), taking advantage of the

semi smoothness of the involved functions in the described models. Further, we discuss the application of Multigrid algorithms in this context, where the smoothing strategies are based on the generalized Newton algorithms. Finally, we discuss the construction of these type of algorithms, without regularization, keeping the convergence properties.

Sergio González-Andrade
Research Center on Mathematical Modelling
Escuela Politécnica Nacional del Ecuador
sergio.gonzalez@epn.edu.ec

MS92

A Multigrid Optimization Algorithm for the Numerical Solution of Classical Viscoplastic Models

We propose a multigrid optimization algorithm (MG/OPT) for the numerical solution of the variational inequalities of the second kind associated with the classical viscoplastic models such as Bingham, Casson and HerschelBulkley fluids in a pipe. This approach is enabled by the fact that the solution of the variational inequality is given by the minimizer of a nonsmooth energy functional, involving the p-Laplace operator. We propose a Huber regularization of the functional and a finite element discretization for the problem. Further, we analyze the regularity of the discretized energy functional and we prove that its Jacobian is slantly differentiable. This regularity property is useful to analyze the convergence of the MG/OPT algorithm. In fact, we demonstrate that the algorithm is globally convergent by using a mean value theorem for semismooth functions.

Sofía López-Ordóñez, Sergio González-Andrade
Research Center on Mathematical Modelling
Escuela Politécnica Nacional del Ecuador
sofia.lopezo@epn.edu.ec, sergio.gonzalez@epn.edu.ec

MS92

The Viscoplastic Navier-Stokes Equations: Fast Defect Correction, Mimetic Discretisation and Applications in Bioprinting

We discuss two building blocks of a robust, accurate and efficient numerical solver for the incompressible Navier-Stokes equations coupled with a viscoplastic constitutive law: (1) a defect correction scheme to resolve the nonlinearity and nonsmoothness and (2) a suitable discretisation. While in the limit of Stokes flow the governing equations may be cast as a nonsmooth, convex optimisation problem, this optimisation structure is lost when considering inertial flow. Instead of resorting to inefficient Stokes linearisations of the full Navier-Stokes equations, we consider Oseen problems with Newton-type linearisations of viscoplastic terms along with suitable preconditioners for the resulting linear systems.
Care has to be taken when algorithms for viscoplastic flow problems are discretised, as many discretisation schemes lead to finite-dimensional algorithms that fail to converge unless the mesh is increasingly refined. We point out how to obtain a robust defect correction method with convergence independent of the chosen grid spacing. Our local discontinuous Galerkin scheme also preserves the physical laws encoded in the system at the discrete level.
Numerical examples of reactive flows that arise in the production of plastic-like biomaterials shall complement this presentation.

Timm Treskatis
Mathematics Department
University of British Columbia
timmt@math.ubc.ca

MS92

Flow of a Yield-stress Fluid over a Cavity: Experimental and Numerical Investigation of an Oldroyd's Boundary Layer

The ability of viscoplastic fluids to self-select their flow geometry through the formation of unyielded dead zones has important consequences for flows in wavy channels, flows over obstacles, etc. Yet, the mechanisms controlling the formation and dimensions of these dead zones remain poorly understood. We present a detailed cross-comparison of experimental and numerical results concerning channel flows of a viscoplastic fluid over a rectangular cavity filled by the same material. In all the configurations studied, which correspond to moderate values of the Bingham number, a continuous dead zone forms inside the cavity. Both numerical and experimental data reveal that, unlike at high Bingham numbers, the shear-rate profiles above the dead zone display an asymmetric shape. Accordingly, two different flow zones can be distinguished: a Poiseuille-like zone, in which the velocity profile is similar to that over a rigid wall, and a boundary layer ensuring the transition with the dead zone below. It is shown that the effective boundary condition felt by the Poiseuille-like layer is essentially controlled by incoming flow characteristics, such that the thickness of this layer does not obey simple relations with cavity length. Interestingly, however, the thickness of the boundary layer appears to follow a generalized Oldroyd's scaling with cavity length. Joint work with G. Chambon, A. Marly, L.-H. Luu and P. Philippe.

Paul Vigneaux
UMPA - CNRS UMR 5669 ENS de LYON
paul.vigneaux@math.cnrs.fr

MS93

A Weighted L1-Minimization Approach For Sparse Wavelet Reconstruction of Signals and Images

Real-world images and signals are known to be sparse in a wavelet basis. In this work, we propose a convex optimization approach, based on weighted l1-regularization, for reconstructing a signal of interest by identifying the coefficients in a sparse wavelet approximation. The vector of wavelet coefficients that form this approximation is obtained by finding the minimizer to our proposed weighted optimization problem where the weights are chosen to be the uniform norms of the wavelet basis functions. We illustrate the effectiveness of this method by solving the image inpainting and image denoising problems. The constraints of the optimization problem are either a set of subsamples in the case of the inpainting problem or a set of noisy observations in the case of the denoising problem.

Joseph Daws
University of Tennessee, Knoxville
jdaws@vols.utk.edu

Clayton G. Webster
University of Tennessee and Oak Ridge National Laboratory

webstercg@ornl.gov

MS93

Domain Decomposition Least-Squares Petrov–Galerkin for Nonlinear Model Reduction

While reduced-order models have demonstrated success in many applications in computational science and engineering, challenges arise when applying model reduction either to *extreme-scale models* or to *decomposable systems*. To address these problems, we propose a domain-decomposition least-squares Petrov–Galerkin (DD-LSPG) model-reduction method. Rather than constructing a low-dimensional subspace for *the entire state space*, these subspaces are constructed for different subdomains/components characterizing the original model. During the offline stage, only *subdomain/component training simulations* are required. During the online stage, the approach constructs a LSPG model for each subdomain/component (including hyper-reduction in the case of nonlinearities), and enforces (weak) compatibility on the 'ports' connecting them. We propose several different strategies for defining the ingredients characterizing the methodology: (i) three different ways to construct basis functions on the interface/ports of subdomains, (ii) different methods for enforcing compatibility (iii) four different solvers that expose different levels of concurrency and intrusiveness. Numerical results performed on problems in heat transfer and fluid dynamics demonstrate that the proposed method outperforms typical model-reduction techniques in terms of both accuracy and (parallel) computational time.

Chi K. Hoang
Sandia National Laboratory
ckhoang@sandia.gov

Kevin T. Carlberg
Sandia National Laboratories
ktcarlb@sandia.gov

MS93

Structure-preserving Galerkin POD Reduced-order Modeling of Hamiltonian Systems

The proper orthogonal decomposition reduced-order model (POD-ROM) has been widely used as a computationally efficient surrogate model in large-scale numerical simulations of complex systems. However, when it is applied to a Hamiltonian system, a straightforward application of the POD method can destroy the Hamiltonian structure in the reduced-order model. In this talk, we develop a new reduced-order modeling approach for Hamiltonian systems, which modifies the Galerkin projection-based POD-ROM so that the appropriate Hamiltonian structure is preserved. We derive a rigorous a priori error estimate for the structure-preserving ROM and demonstrate its effectiveness in several numerical examples.

Zhu Wang
Department of Mathematics
University of South Carolina
wangzhu@math.sc.edu

MS94

Asynchronous Two-level Domain Decomposition Solvers

Increasingly large problem sizes and more heterogeneous supercomputing systems make load balancing and network layout very challenging tasks for linear iterative solvers. In particular, global communication patterns such as inner products become increasingly limiting at scale. In this talk, we explore the use of an asynchronous scalable two-level domain decomposition solver based on one-sided MPI primitives. We will discuss practical issues encountered in the development of a scalable solver and show experimental results obtained on a variety of state-of-the-art supercomputer systems.

Christian Glusa
Sandia National Laboratories
caglusa@sandia.gov

MS94

Asynchronous Subgradient-push

We consider a multi-agent framework for distributed optimization where each agent in the network has access to a local convex function and the collective goal is to achieve consensus on the parameters that minimize the sum of the agents' local functions. We propose an algorithm wherein each agent operates asynchronously and independently of the other agents in the network. When the local functions are strongly-convex with Lipschitz-continuous gradients, we show that a subsequence of the iterates at each agent converges to a neighbourhood of the global minimum, where the size of the neighbourhood depends on the degree of asynchrony in the multi-agent network. When the agents work at the same rate, convergence to the global minimizer is achieved. Numerical experiments demonstrate that Asynchronous Subgradient-Push can minimize the global objective faster than state-of-the-art synchronous first-order multi-agent methods, is more robust to failing or stalling agents, and scales better with the network size. This is joint work with Mahmoud Assran.

Mike Rabbat
Facebook Research
mikerabbat@fb.com

MS94

An Asynchronous, Decentralized Solution Framework for the Large Scale Unit Commitment Problem

The Unit Commitment (UC) problem in power networks deal with determining which generators to turn on during the day and how much power each produces. Owing to the binary nature of decisions, the UC problem is Mixed Integer in nature and as a result computationally challenging. In this talk we present an asynchronous decentralized solution framework for the UC problem for large scale power networks. We exploit the inherent asynchrony arising out of imbalance in regional subproblems to boost computational efficiency in a region based decomposition. A two phase algorithm is proposed that relies on the convex relaxation and privacy preserving valid inequalities in order to deliver algorithmic improvements. Our algorithm employs a novel interleaved binary mechanism that locally switches from the convex subproblem to its binary counterpart based on consistent local convergent behavior. We develop a high performance computing (HPC) oriented software framework that uses Message Passing Interface (MPI) to benchmark our asynchronous method against a state of

the art synchronous decentralized as well as the centralized method. The results, obtained on the large scale IEEE 3012 bus case, demonstrate that the asynchronous method improves computational efficiency by a significant amount and provides a competitive solution quality rivaling the benchmark methods.

Paritosh P. Ramanan
School of Computational Science and Engineering
Georgia Institute of Technology
paritoshpr@gatech.edu

MS94

Asynchronous Optimized Schwarz Methods for the Solution of PDEs on Bounded Domains

In this talk, an asynchronous version of the optimized Schwarz method is presented for the solution of differential equations on a parallel computational environment. Convergence is proved under very mild conditions on the size of the subdomains, when optimal as well as approximate (non-optimal) interface conditions are utilized for Poisson's equation (and others) on the plane and on bounded rectangular domains. Numerical results are presented on large three-dimensional problems illustrating the efficiency of the proposed asynchronous parallel implementation of the method.

Daniel B. Szyld
Temple University
Department of Mathematics
szyld@temple.edu

Federeic Magoules
CentraleSupelec
France
frederic.magoules@centralesupelec.fr

Jose Garay
Lousiana State University
jgaray@lsu.edu

Mireille El Haddad
Universite Saint-Joseph
Beirut, Lebanon
tuh47970@temple.edu

MS95

Chapel's Language-based Approach to Performance Portability

In mainstream computing, programming languages have long served as a way to achieve performance portability by raising the level of abstraction while relying on a compiler to specialize for specific platforms. Fortran is a canonical example, generally credited with obviating the need to write machine code for each new system at a time when hardware was evolving rapidly. Similarly, C, Java, and Swift have each provided performance portability benefits for their respective domains.Notably, the most commonly adopted programming notations in High Performance Computing tend not to be language-based, relying instead on pragmas and libraries. Worse, like machine code, they tend to embed architecture-specific assumptions into the program text, such as data transfer paradigms, assumptions of shared memory, or SIMD kernel parameters. This causes pain for programmers wanting to run across disparate HPC systems today, to say nothing of the architectural diversity we can expect in the exascale era.In this talk, I'll describe how the Chapel programming language has been designed for performance portability through a combination of language features, multiresolution design, and runtime library design. As a result, Chapel programmers are able to write clean, high-level code that can be optimized for specialized HPC hardware features while also porting to less capable architectures. In closing, I'll characterize recent Chapel performance results and use cases.

Bradford L. Chamberlain
Cray Inc.
bradc@cray.com

MS95

Performance Optimisation of Finite Element Assembly in Firedrake

Modern CPUs increasingly rely on SIMD instructions to achieve higher throughput and better energy efficiency. It is therefore important to vectorize sequences of computations in order to sufficiently utilize the hardware today and in the future. Typical finite element assembly kernels suffer from issues that often preclude efficient vectorization. These include complicated loop structure, poor data access patterns, and unfriendly loop trip counts. General purpose compilers often perform poorly in generating efficient, vectorized code. In this work, we present a generic and automated solution in Firedrake based on cross element vectorization. We adapt our form compiler, TSFC, to generate Loo.py kernels for local assembly operations, and systematically generate data gathering and scattering operations across the mesh in PyOP2. Firedrake drives loop transformations using Loo.py to generate efficient code vectorized across a group of elements to fully utilize SIMD lanes. This toolchain automates the error-prone process of data layout transformation, loop unrolling and loop interchange, while being transparent to the users. We present experimental results on multiple kernels and meshes on different platforms and compilers. We achieve speed ups consistent with the vector architecture available compared to baseline which vectorizes inside the local assembly kernels. The global assembly computations consistently reach tens of percent of hardware peak arithmetic performance.

Tianjiao Sun
Imperial College London
tianjiao.sun14@imperial.ac.uk

Lawrence Mitchell
Durham University
lawrence.mitchell@durham.ac.uk

David Ham
Department of Mathematics
Imperial College London
david.ham@imperial.ac.uk

Paul Kelly
Imperial College London
p.kelly@imperial.ac.uk

MS95

Kokkos: Insights from 4 Years of Application Adoption Work

Todays High Performance Computing (HPC) applications face a diversifying supercomputer landscape, which each

pose their own unique challenges. One of the most fundamental problems is that each of those architectures comes with its own preferred on-node programming model, be it OpenMP for traditional multi core CPUs, or CUDA, OpenACC or ROCm based models for GPUs depending on the vendor. Reimplementing an application for each of those models is for many teams out of the question due to the maintenance and development burden. Kokkos solves this problem by providing a C++ Performance Portability Programming model, which maps to vendor specific models underneath. For this reason, dozens of applications with hundreds of developers have moved to adopt Kokkos during the last 4 years. These applications come from numerous HPC institutions and cover a vast range of use cases such as molecular dynamics, graph analytics, mechanical engineering, computational fluid dynamics, climate modelling and electronic structure calculations codes. This talk will provide insight into experience with adopting Kokkos in production and research applications, and what real world performance portability looks like.

Christian Trott
Sandia National Laboratories
crtrott@sandia.gov

MS96

Tuning Machine Learning Models with Derivative-free Optimization

With the exponential growth rate of digital data, the challenge of managing, understanding, and capitalizing on this data continues to grow. Machine learning modeling algorithms are commonly used to find hidden value in big data. These algorithms are governed by hyperparameters with no clear defaults agreeable to a wide range of applications. Ideal settings for these hyperparameters significantly influence the resulting accuracy of the predictive models. In this talk we discuss the use of derivative-free optimization for automated hyperparameter tuning. We present our Local Search Optimization (LSO) framework which implements a parallel hybrid derivative-free optimization strategy for problems with functions that are nonsmooth, discontinuous, or computationally expensive to evaluate directly. We also show how our initial tuner has been extended to make use of multiobjective optimization techniques. In the multiobjective setting, we trade-off model error and model complexity, providing users with an approximate Pareto-optimal set of nondominated solutions. We present tuning results for multiple examples.

Steven Gardner
SAS Institute Inc.
steven.gardner@sas.com

Oleg Golovidov, Joshua Griffin, Patrick Koch, Scott Pope
SAS Institute, Inc.
oleg.golovidov@sas.com, Joshua.Griffin@sas.com, patrick.koch@sas.com, scott.pope@sas.com

MS96

Decomposition in Global Derivative-free Optimization

Derivative-Free Optimization (DFO) is an important class of optimization algorithms that address problems where derivatives are unavailable, unreliable, or only available at a significant cost. Current DFO solvers have the limitation that they perform well only when the number of degrees of freedom is about a dozen or less. We propose a novel decomposition scheme for model-based DFO algorithms that significantly extends the scope of model-based solvers to larger-scale problems. A practical implementation is developed and demonstrated with the global model-based solver Stable Noisy Optimization by Branch and Fit (SNOBFIT). The impact of the subproblem dimension and the subproblem function evaluation limit on the performance of the overall algorithm is studied computationally on a collection of over 500 test problems of varying dimensions and complexity. Significant improvement in the quality of solutions obtained under a 2500 function evaluation limit is observed for a large fraction of the test problems.

Kaiwen Ma
Carnegie Mellon University
Pittsburgh, PA
kaiwenm@andrew.cmu.edu

Nikolaos Sahinidis
Carnegie Mellon University
sahinidis@cmu.edu

Satyajith Amaran, Scott Bury
The Dow Chemical Company
samaran@dow.com, sjbury@dow.com

MS96

Improving Surrogate Global Optimization Accuracy with RBF Trust Regions

A new algorithm GLOS improves accuracy over existing Radial Basis Function RBF Algorithms for surrogate continuous global optimization by using cycling of global and local optimization search. One of the features of the algorithm is to analyze prior searches around multiple local minima to decide where to do a derivative-free Trust Region search with ORBIT in the next iteration. The comparisons are based on performance and data profiles from multiple trials of each algorithm on 43 test problems with dimensions varying between 2 to 40. Comparisons of RBF surrogates to Bayesian Optimization to RBF surrogates are also provided, including a deep learning application. The results show that GLOS has the best overall accuracy, especially on high dimensional problems with a small to medium number of local minima, including improved performance over previous surrogate algorithms that used a mixture of local and global optimization. A theorem proves almost sure convergence of the GLOS algorithm.

Christine Shoemaker, Limeng Liu, Jin Yi
National University of Singapore
shoemaker@nus.edu.sg, limeng_liu@u.nus.edu, iseyij@nus.edu.sg

MS96

Targeted Design of Nanoparticle Synthesis using Derivative-free Optimization

Production of nanoparticles with uniform properties (e.g., size, composition, phase, morphology) can benefit a number of applications, such as catalysis and energy production. Microfluidic flow synthesis methods have allowed for production of metallic nanoparticles in a controllable fashion by a tight control over key physical parameters such as residence time, chemical concentrations and temperature. However, determining these parameters typically requires parametric testing, which becomes expensive for even a moderate number of parameters. We propose techniques aimed at determining optimal parameters for synthesizing

nanoparticles close to a target size. Since each experiment yields an unspecified number of particles, we require our method to take empirical distributions as input. With Efficient Global Optimization as a reference blackbox optimization method, we study the suitability of different objective functions measuring the uniformity of sample size. We then demonstrate the impact of incorporating the full empirical distribution in the surrogate modeling process. Although motivated by nanoparticle synthesis in continuous flow reactors, our methodology can be extended to other material fabrication systems and blackbox experimental settings.

Aekaansh Verma
Department of Mechanical Engineering,
Stanford University
aekaansh@stanford.edu

Stefan Wild
Argonne National Laboratory
wild@mcs.anl.gov

MS97

Coupled Monte Carlo Radiation Transport and CFD for Nuclear Reactor Modeling

Modeling and simulation of nuclear reactors requires simultaneously solving equations describing neutron transport, heat transfer, and fluid flow. The ExaSMR project is aimed at producing extremely high-fidelity simulations of small modular reactors (SMRs) by coupling continuous-energy Monte Carlo radiation transport and computational fluid dynamics (CFD) solvers for execution on some of the largest available computing platforms. In this talk, we will discuss algorithmic challenges associated with this multi-physics problem. In particular, we will talk about difficulties in the coupling arising due to stochastic noise from the Monte Carlo solver. Numerical results from a series of test problem will be presented.

Steven Hamilton
ORNL
hamiltonsp@ornl.gov

Thomas Evans
Oak Ridge National Laboratory
evanstm@ornl.gov

MS97

Multiphase Flow Modeling with MFIX-Exa

Multiphase flows containing solid particles in a gaseous flow can be modeled with a CFD-DEM approach, combining computational fluid dynamics (CFD) with the discrete element method (DEM). The traditional approach in the MFiX code has decoupled the fluid and particle motion and used the SIMPLE method for the fluid flow. In MFIX-Exa, a next-generation multiphase modeling code, we replace the SIMPLE algorithm with a modern projection method and explore improved coupling of the fluid and particle motion. In particular, we include the time variation of the volume fraction occupied by the particles in each grid cell in the evolution of the fluid. In this talk we will describe the new algorithm and present several test cases to assess its efficacy in the context of industrial relevant flows.

Michele Rosso
LBNL
mrosso@lbl.gov

MS97

Novel Computational Methods for High-fidelity Modeling of Plasma-based Particle Accelerators

Turning the current experimental plasma-based particle accelerator state-of-the-art from a promising technology into mainstream scientific tools depends critically on high-performance, high-fidelity modeling of complex processes that develop over a wide range of space and time scales. As part of the U.S. Department of Energy's Exascale Computing Project, a team composed of LBNL, SLAC and LLNL researchers is developing a new plasma accelerator Particle-In-Cell simulation tool: WarpX. The new software will harness the power of future exascale supercomputers for the exploration of outstanding questions in the physics of acceleration and transport of particle beams in chains of plasma channels. This will benefit the ultimate goal of compact and affordable high-energy physics colliders, and many spinoff applications of plasma accelerators along the way. The success of the project depends critically on the combination of advanced numerical algorithms that include pseudo-spectral Maxwell analytical solvers (PSATD), adaptive mesh refinement and modeling in a so-called 'Lorentz boosted frame of reference'. This in turn demands the development of novel algorithms for the application of Perfectly Matched Layers for open boundary conditions with PSATD and for large time step solvers. We will present the novel algorithms, analyses of their performances and examples of applications.

Olga Shapoval
LBNL
OShapoval@lbl.gov

Jean-Luc Vay
Lawrence Berkeley National Laboratory
Berkeley, CA
JLVay@lbl.gov

Rémi Lehe
Lawrence Berkeley National Laboratory
ATAP
rlehe@lbl.gov

Maxence Thevenet
Lawrence Berkeley National Laboratory
mthevenet@lbl.gov

Henri Vincenti
CEA
henri.vincenti@cea.fr

MS97

ExaWind: Multi-scale Challenges of Massively Parallel Blade-resolved Wind Turbine Simulations

This talk will describe experiences from the ExaWind Exascale Computing Project, which is in pursuit of predictive wind-turbine and wind-plant simulations. Predictivity will require blade-resolved moving meshes, high-resolution grids to resolve the flow structures, hybrid-RANS/LES turbulence modeling, fluid-structure interaction, and coupling to meso-scale flows. The modeling and algorithmic pathways of ExaWind include unstructured-grid finite volume spatial discretization and pressure-projection methods for incompressible flow, which are embodied in the Nalu-Wind

computational fluid dynamics code. Nalu-Wind utilizes the Hypre and/or Trilinos linear-system solvers, is built on the Trilinos Sierra Toolkit, and incorporates the Kokkos abstraction layer for next-generation platforms. A significant challenge is the strong-scaling limit of pressure-Poisson solvers and limits on time-step size. Using blade resolved simulations of a megawatt-scale turbine, we will describe efforts to improve time to solution through, e.g., MPI+X parallelism and improvements to the linear-solver algorithms. GPU acceleration of the pressure-Poisson solve will also be described.

Michael Sprague
NREL
michael.a.sprague@nrel.gov

Shreyas Ananthan
National Renewable Energy Laboratory
shreyas.ananthan@nrel.gov

Stephen Thomas
Scientific Computing Division
National Center for Atmospheric Research
stephethomas@gmail.com

Michael Lawson
National Renewable Energy Laboratory
michael.lawson@nrel.gov

Katarzyna Swirydowicz
National Renewable Energy Lab
Golden CO
katarzyna.swirydowicz@nrel.gov

MS98

Wireless Sensing for the Energy Industry over the Internet of Things

Sensors are everywhere across different sectors of the oil and gas industry. Seismic acquisition in upstream, pipeline monitoring in midstream and predictive maintenance in downstream are few examples of applications that need more sensors to ensure an acceptable level of safety and satisfy the pressing need for accuracy. In some cases, sensor data should be quickly aggregated, coordinated and transmitted, sometimes from harsh environments where crew intervention should be minimized in favor of safety and cost. Our solution to address this need consists of three components: 1) a network of cheap, low-power and long-range wireless sensors to minimize hardware and maintenance costs, 2) an appropriate wireless networking solution that easily scales up and can handle heterogeneous and simultaneous data streams and 3) cloud-based data storage and computation that would enable efficient and scalable accessing, handling and processing of data. We field tested two Low-Power Wide Area Network technologies readily available in the market: LoRa and NB-IoT. The first two tests were designed to understand the performance of LoRa for up- and downstream applications: wireless seismic data quality control and wireless sensing in a chemical plant. The third trail was designed to demonstrate the use of NB-IoT to collect and process in real-time on the cloud ambient seismic noise for monitoring purposes. We discuss details of our networking design and show results of the field trials.

Hadi Jamali-Rad
Shell International
hadi.jamali-rad@shell.com

Xander Campman
Shell Global Solutions International B.V.
xander.campman@shell.com

MS98

MyShake - Dense Global Seismic Network using Smartphones

Sensors play an important role in today's world to sense the environment, but what's the best way to scale the density of sensors up to have finer spatial recordings? MyShake is a project exploring along this line by turning the daily smartphones in your pocket to a recording device to monitor the earthquakes. With this approach, a dense global seismic network has been set up in recently, and started to detect earthquakes and potentially issue earthquake early warning in the near future. The higher density of this low-cost sensor network comes with a price of noisier data compare with the high-quality sensors, for example, the phones are moving around, the noise level on the phones are high, there are so many human activities in the data and so on. Therefore, special treatment is needed to address these issues. In this talk, I will take you on the journey of addressing them using machine learning approaches, such as Artificial Neural Network, clustering algorithms and so on. At the same time, data collected and observations from this network will be reported here to illustrate the effectiveness of this low-cost smartphone network.

Qingkai Kong
University of California Berkeley
kongqk@berkeley.edu

Richard Allen
University of California Berkeley
Berkeley Seismology Lab
rallen@berkeley.edu

MS98

Fast Radon Transforms for Ambient Noise Interferometry

Ambient noise interferometry is a process that extracts approximations to controlled-source seismic experimental data from a random diffuse wavefield recorded by an array of n seismic sensors. The first step is typically to perform $O(n^2)$ pairwise cross-correlations between sensors, yielding virtual source response estimates, which mimic data from a controlled source at each receiver location. Next a Radon transform is applied to each virtual source response estimate. The output of each Radon transform is the dispersion image for that virtual source location, showing how much energy travels at a variety of phase velocities and frequencies. Dispersion images are a common tool for near-surface characterization, and are often the input to a surface wave inversion process. Motivated by low-cost high-density arrays for continuous monitoring, we propose a new $O(n)$ algorithm which calculates the same dispersion images without requiring the intermediate step of $O(n^2)$ cross-correlations. The new algorithm is conceptually simple, embarrassingly parallel, uses streaming data, does not require regular spacing between sensors, and can be extended to several methods of estimating virtual source responses including cross-correlation, cross-coherence and deconvolution. We demonstrate the algorithm's effectiveness and scalability on data from a field trial of a trenched distributed acoustic sensing array at the Richmond Field

Station in California.

Eileen R. Martin
Virginia Tech
eileenrmartin@vt.edu

Biondo Biondi
Geophysics Department
Stanford University
biondo@sep.stanford.edu

Jonathan Ajo-Franklin
Lawrence Berkeley National Lab
jbajo-franklin@lbl.gov

MS98

Compressive Least-squares Migration with on-the-fly Fourier Transforms

Least-squares seismic imaging is an inversion-based approach for accurately imaging the earth's subsurface. However, in the time-domain, the computational cost and memory requirements of this approach scale with the size and recording length of the seismic experiment, thus making this approach often prohibitively expensive in practice. To overcome these issues, we borrow ideas from compressive sensing and signal processing and introduce an algorithm for sparsity-promoting seismic imaging using on-the-fly Fourier transforms. By computing gradients and functions values for random subsets of source locations and frequencies, we considerably limit the number of wave equation solves, while on-the-fly Fourier transforms allow computing an arbitrary number of monochromatic frequency-domain wavefields with a time-domain modeling code and without having to solve large-scale Helmholtz equations. The memory requirements of this approach are independent of the number of time steps and solely depend on the number of frequencies, which determine the amount of crosstalk and subsampling artifacts in the image. We show the application of our approach to several large-scale open source data sets and compare the results to a conventional time-domain approach with optimal checkpointing.

Philipp A. Witte
PhD student at the University of British Columbia
pwitte.slim@gmail.com

Mathias Louboutin
PhD student at the University of British Columbia
Seismic Laboratory for Imaging and Modeling
mloubout@eos.ubc.ca

Fabio Luporini
Department of Computing
Imperial College London
f.luporini12@imperial.ac.uk

Gerard J Gorman
Department of Earth Science and Engineering
Imperial College London
g.gorman@imperial.ac.uk

Felix Herrmann
Georgia Institute of Technology
felix.herrmann@gatech.edu

MS99

Challenges in MPI for Supporting Distributed-memory Graph Analytics

Distributed-memory graph analytics have proven to be one of the most complex problems to scale to large supercomputing systems, primarily owing to their significant bias towards data movement compared with traditional computation. While several graph partitioning techniques exist to improve this situation, a number of domains such as genome assembly still suffer from such overheads due to the fact that the partitioning of the graph is often as expensive as analyzing it. This leads to a situation where executing such graph analytics on a large supercomputer would result in a tremendous number of data messages to be injected on to the network by each node in the supercomputer system. In this talk, I'll discuss the role of the MPI communication library in improving the performance of such distributed-memory graph analytics problems. I'll use genome assembly as a case study, and showcase how recent improvements in MPI implementations, such as lightweight communication techniques and one-sided communication, have addressed some of these problems. I'll also discuss what's missing in MPI in handling such problems and what the next steps are.

Pavan Balaji
Argonne National Laboratory
balaji@anl.gov

MS99

Scalable Graph Community Detection using the Louvain Method

Graph clustering, popularly known as community detection, is a fundamental graph operation used in several applications of relevance for network analysis and cybersecurity. Community detection is an operation of partitioning a network into a number of communities such that each community represents a tightly-knit group of nodes with relatively sparser connections to the rest of the nodes in the network. The need to compute clustering on large-scale networks necessitates the development of efficient parallel algorithms capable of exploiting modern parallel architectures. However, due to their irregular and inherently sequential nature, many of the current algorithms for community detection are challenging to parallelize. In this talk, we present the design of a distributed memory implementation of the Louvain algorithm for parallel community detection. Our approach begins with an arbitrarily partitioned distributed graph input, and employs several heuristics to accelerate the computation of the different steps of the Louvain algorithm. We evaluate our implementation and its different variants using real-world networks from various application domains (including internet, biology, social networks). Our MPI+OpenMP implementation yields about 1.8x-46.18x relative to a parallel baseline version (on 4K processes of NERSC Cori supercomputer) without compromising on output quality.

Sayan Ghosh
Washington State University
sghosh1@eecs.wsu.edu

Mahantesh Halappanavar, Antonino Tumeo
Pacific Northwest National Laboratory
mahantesh.halappanavar@pnnl.gov,
Antonino.Tumeo@pnnl.gov

Ananth Kalyanaraman
Boeing Centennial Chair, Associate Professor
School of EECS, Washington State University

ananth@wsu.edu

Assefaw Gebremedhin
Washington State University
assefaw.gebremedhin@wsu.edu

MS99

Distributed-memory Graph Algorithms for Bioinformatics

This talk will review distributed-memory graph-based methods for common bioinformatic analyses such as *de novo* genome assembly, small variant detection, and DNA sequence clustering. The effective use of distributed-memory systems for bioinformatics requires lightweight preprocessing routines for static load-balancing, parallel I/O, inter-node communication-reducing mechanisms, and fully exploiting shared-memory multicore resources. We will present MetaPartMin, a DNA sequence partitioning tool employing an implicit distributed graph representation. MetaPartMin includes novel memory-reducing optimizations to reduce aggregate main memory use, enabling the partitioning of massive metagenomic datasets on a modest number of compute nodes. All steps of MetaPartMin exploit *hybrid* multicore and distributed-memory parallelism. On 32 nodes of the Stampede2 supercomputer, MetaPartMin can partition a 1.25 terabase soil metagenome in 6 minutes.

Kamesh Madduri
Pennsylvania State University
madduri@cse.psu.edu

MS99

Kronecker Graphs as Benchmarks for Distributed Graph Analytics

Graph analysis is an important facet of data analysis and scalable graph analytics are desired for massive sparse graphs, where the number of edges, m, is potentially trillions. Development of scalable graph algorithms requires graph generators with challenging properties of real-world graphs (e.g. small-world, scale-free, heavy-tailed degree distribution) and with efficiently calculable ground-truth of desired output, so researchers can check algorithmic correctness. Typically an $O(m^p)$ analytic is computable for a (non-stochastic) Kronecker graph, with adjacency matrix of the form $C = A \otimes B$, in $O(m^{p/2})$ cost, if a Kronecker formula exists. This is a highly attractive property, particularly in cases where $p > 1$, so we promote Kronecker graphs as an important class of benchmarks for HPC graph algorithms and implementations. Here we demonstrate how tailoring assumptions on matrices A and B often yields a massive Kronecker graph C with a simple Kronecker formula for a specific graph analytic in terms of related calculations on A and B. We repeat this exercise for a diverse set of graph computations including: triangle counts and vertex/edge participation rates in undirected/directed/labelled graphs, k-truss decomposition, graph diameter and vertex eccentricity, and graph clustering metrics. We discuss subtle important differences in assumptions and limitations of the approaches, so researchers can utilize these test graphs most effectively.

Geoff Sanders
Lawrence Livermore National Laboratory
sanders29@llnl.gov

Trevor Steil
University of Minnesota
steil016@umn.edu

Timothy La Fond
Lawrence Livermore National Lab
lafond1@llnl.gov

Roger Pearce
Lawrence Livermore National Laboratory
rpearce@llnl.gov

Keita Iwabuchi
Lawrence Livermore National Laboratory, USA
iwabuchi1@llnl.gov

Benjamin Priest
Dartmouth College
bwpriester@gmail.com

MS100

Using AMR for Scaling Globally Connected Radiation Calculations

We present recent developments in advancing the Uintah code. Uintah realizes dynamic adaptive meshing and achieves its scalability using a task-based subsystem for work sharing. We demonstrate algorithms for further reducing the overhead in task scheduling, honoring latest developments in hybrid architectures. The benefits are especially visible for non-local physical processes. We close with demonstrations of complex multi-physics computations that involve radiative heat transfer and multiscale combustion.

Martin Berzins, Alan Humphrey
Scientific Computing and Imaging Institute
University of Utah
mb@sci.utah.edu, ahumphrey@sci.utah.edu

MS100

Non-local Parallel Algorithms for Data Assignment and Analysis

Many relevant tasks in numerical simulation and pre-/post-processing require non-local data access in terms of the geometrical neighborhood of computational elements or points. Such cases arise in large-CFL fluid dynamics, in particle tracking, and when transfering data between render and writer processes in visualization. When we use distributed parallelism, it is a priori unclear how to identify the peer processes without all-to-all communication. We address these tasks with new algorithms based on a distributed forest of octrees, exploiting its particular encoding of the parallel partition. Specifically, we propose how to find non-neighbor remote objects and how to communicate data of irregular distribution in an efficient and scalable way. We illustrate our methods using large-scale examples of particle tracking and in-situ visualization.

Carsten Burstedde
Universität Bonn
burstedde@ins.uni-bonn.de

MS100

ForestClaw: An Adaptive, Finite Volume, Cartesian Grid Solver Based on p4est

We describe the software package ForestClaw, a PDE solver

for time dependent hyperbolic and parabolic solvers based on updating a solution on an adaptive quadtree mesh. A basic design principle behind ForestClaw is to abstract algorithmic components and make extensive use of polymorphism. To maintain ease of use and flexibility, only limited use is made of inheritance and encapsulation. The core ForestClaw routines manage input parameters, adaptive (multi-rate) time stepping, re-meshing, transferring the solution between old and new meshes, parallel communication, file I/O and diagnostics. A secondary focus of the talk will be on the underlying capabilities provided by p4est and the p4est application program interface (API). Applications from natural hazards modeling will be presented.

Donna Calhoun
Boise State University
donnacalhoun@boisestate.edu

Carsten Burstedde
Universität Bonn
burstedde@ins.uni-bonn.de

MS100

Parallel Variational Transfer for Scalable Multiphysics Simulations

The scalable simulation of coupled multiphysics problems requires parallel and efficient transfer routines between possibly non-matching and arbitrarily distributed unstructured meshes. Here, we present in the context of fluid-structure simulations a scalable and parallel approach, which allows for the variational transfer of discrete fields between non-matching unstructured meshes in parallel. Our parallel transfer operator uses a discrete L^2-projection for stability reasons. We employ this parallel transfer for developing a novel Fluid-Structure-Interaction approach, which is based on ideas from immersed boundary methods. The resulting FSI methods allows for the coupled simulation of on-linearly elastic materials in turbulent fluids. Here, we employ finite elements for the solid part and finite differences for the fluid part. Numerical examples, including the simulation of artificial heart valves, illustrate the properties of our approach.

Marco Favino
Institute of Computational Science
Università della Svizzera italiana
marco.favino@usi.ch

Patrick Zulian, Rolf Krause
Università della Svizzera italiana, Switzerland
patrick.zulian@usi.ch, rolf.krause@usi.ch

Maria Nestola
Institute of Computational Science, USI, Lugano, Switzerland
nestom@usi.ch

MS101

3D Visual Tour of Earth's Interior Based on Global Adjoint Tomography

Increase in the amount and quality of seismic data together with advances in numerical solvers and HPC systems have provided new opportunities to increase the resolution of seismic images thus our understanding of Earths interior and dynamics. After the first-generation global adjoint tomography model, GLAD-M15, its successor GLAD-M25* has been constructed with an almost six-times larger dataset of 1428 earthquakes. The models are often inspected visually thus robust & efficient visualization tools are necessary to thoroughly investigate large model files. In collaboration with Oak Ridge National Lab., we have developed effective visualization tools to produce 3D contour plots and movies for various seismic parameters to better visualize plume- and slab-like features as well as anisotropy in the mantle. We use VisIt (wci.llnl.gov/simulation/computer-codes/visit/) for initial exploration of the models and extraction of seismological features which reads models in ADIOS format (www.olcf.ornl.gov/center-projects/adios/) through a developed data reader plugin. Blender (www.blender.org) is used for the setup of lighting, materials, camera paths and rendering of geometry. Our aim is to make visualization an integral part of our adjoint tomography workflow. Furthermore we now implement our models into an immersive VR system, that allows for the interaction of multiple users, for their full 3D inspection. The developed tools will also serve to enhance education in earth sciences.

Ebru Bozdag
Colorado School of Mines
bozdag@mines.edu

David Pugmire
Oak Ridge National Laboratory
pugmire@ornl.gov

Cagdas Demirkan, Sebnem Duzgun, Ergin Isleyen
Colorado School of Mines
Department of Mining Engineering
demirkan@mymail.mines.edu, duzgun@mines.edu, isleyen@mymail.mines.edu

Wenjie Lei
Princeton University
lei@princeton.edu

Ridvan Orsvuran
University of Cote d'Azur
ridvan.orsvuran@geoazur.unice.fr

Youyi Ruan
Nanajing University
youyir@nju.edu.cn

Jeroen Tromp
Princeton University
Department of Geosciences
jtromp@princeton.edu

MS101

Opening Your Data: Libraries and Repositories Help You Go Public

With the February 2013 release of a White House policy memorandum mandating public access to federally funded research, multiple federal agencies began creating their public access plans, delineating which repositories would house the articles published as a result of federally funded research. However, many still leave the deposit of research data to a repository of the researcher's choice. In addition to federal public access mandates, many publishers now require authors make their data sets and other supplementary files publicly available but don't always have the technical capacity to host these files. Because libraries began hosting repositories nearly 20 years ago, they are uniquely

positioned to play an important role in the research and publication cycle by shepherding these files that comprise the ingredients to reproducible science. The University at Buffalo Institutional Repository (UBIR) is one such repository that can host any file type and assign persistent URLs to items. University Libraries can also assist researchers with licensing their data for reuse by the public.

Karlen Chase
University at Buffalo
University Libraries
klchase@buffalo.edu

MS101

Using 3D Virtual Reality for Interactively Exploring Models and Big Data

The coupling of HPC with immersive 3D data visualization is transforming computational fluid dynamics research, by rendering a new dimension of physical meaning to numerical models by enabling the user to grasp the model as if it were a tangible physical construct. Here we show how increases in computational resources drive a new avenue in geodynamics research focusing on the Earths natural subduction zones. We present virtual voyages through the Earth's interior that provide a snapshot into the four-dimensional evolution of plate tectonics, captured by the modern state of slab structure. By conceptualizing modern subduction system on Earth in 3D space, this approach facilitates the paradigm shift from a two-dimensional (2D) to 3D framework for interpreting subduction. We then show results from high-resolution, finite-element models of non-linear viscous flow in the Earth using a complex model design representative of natural plate boundaries. The 3D models contain over 100 million mesh nodes, over 500 million unknowns, with local resolution of 2.5 km, and cost over 15,000 compute hours per time-step. The large viscosity variations pose a challenge for the solvers. The full multigrid is used to reduce the residual and expedite convergence. Flow regions with large strain-rates emerge, resulting in low viscosity regions to due to the non-linear weakening. The result is a framework for 3D non-linear flow subduction zones that is a departure from the corner flow paradigm.

Margarete Jadamec
University at Buffalo
mjadamec@buffalo.edu

Oliver Kreylos
Center for Image Processing and Integrated Computing
University of California, Davis
kreylos@cs.ucdavis.edu

Burak Yikilmaz
University of California, Davis
mbyikilmaz@ucdavis.edu

Karen Fischer
Brown University
karen_fischer@brown.edu

Benjamin Chang
University of Houston
chang.benj@gmail.com

MS101

Creating Visualizations for Science Communication

Science centers use high quality visualization displays to explain concepts to their visitors. The modern digital planetarium is an immersive, all purpose (not just astronomy) visualization facility. However, while visualizing scientific data and computational simulations for public audiences, additional considerations need to be made. For example when using color, the visualizer needs to be aware of the mental associations that public audiences bring with them regarding color. In addition public audiences often require higher fidelity renderings with more timesteps than are required for scientific purposes. Elements that a scientist can be trained to ignore (e.g. artifacts, selection masks) can be distracting when presented to the public. This presentation will discuss the considerations that need to be made when adapting a visualization made for and by scientists to one used to communicate with the public. Ultimately, one wants to strike a harmonious balance between telling a story and presenting the data as accurately as possible.

Maria Weber
University of Chicago
Adler Planetarium
mweber@adlerplanetarium.org

MS102

A Scalable Global Optimization Algorithm for Stochastic Nonlinear Programs

We present a global optimization algorithm for two-stage stochastic nonlinear programs (NLPs). The algorithm uses a tailored reduced-space spatial branch and bound (BB) strategy to exploit the nearly decomposable structure of the problem. At each node in the BB scheme, a lower bound is constructed by relaxing the so-called non-anticipativity constraints and an upper bound is constructed by fixing the first-stage variables to the current candidate solution. A key advantage of this approach is that both lower and upper bounds can be computed by solving individual scenario subproblems. Another key property of this approach is that we only need to perform branching on the first-stage variables to guarantee convergence (branching on the second-stage variables is performed implicitly during the computation of lower and upper bounds). We present a serial implementation of the algorithm in Julia , that we call SNGO. The implementation is interfaced to the structured modeling language Plasmo.jl , which facilitates benchmarking and model processing. Our implementation incorporates typical features that help accelerate the BB search such as LP-based lower bounding techniques, local search-based upper bounding techniques, and relaxation-based bounds tightening techniques. We perform numerical experiments and compare the computational results against SCIP and demonstrate that the proposed approach achieves significant speedups.

Yankai Cao
University of British Columbia
yankai.cao@gmail.com

Victor Zavala
University of Wisconsin Madison
zavalatejeda@wisc.edu

MS102

Bilevel Mixed-binary Programming for Identifying

Critical Contingency Events in AC Power Systems

We present a bilevel mixed-binary programming approach for identifying critical contingency events in alternating current (AC) power systems. The upper level problem represents the attacker that maximizes the infeasibility of system operations for a given budge, and the lower level problem represents the system operator that minimizes the total amount of load lost for a contingency event. The problem is challenging because the bilevel mixed-binary program is nonconvex and discontinuous, and also because the lower level problem, representing ACOPF, is a nonconvex nonlinear program. We present an optimization algorithm that can effectively find a global optimum for the problem. We also present numerical results on IEEE systems.

Kibaek Kim
Mathematics and Computer Science Division
Argonne National Laboratory
kimk@anl.gov

Brian C. Dandurand, Sven Leyffer
Argonne National Laboratory
brian.dandurand@rmit.edu.au, leyffer@anl.gov

MS102

Using Machine Learning to Understand and Solve Power System Optimization Problems

Power systems optimization involves solving similar optimization problems over and over and over again, with slightly varying input parameters. We consider the problem of learning the solution to a new instance by leveraging information available in previous solutions. While we found that directly learning the solution is a challenging task, we propose to instead learn sets of active constraints, which allow us to both better understand and more efficiently solve future instances of the problem. We propose a rigorous framework for learning the relevant sets of active constraints, which comes with statistical performance guarantees. Applying the algorithm to the optimal power flow problem with renewable energy, we establish that the set of active constraints is efficiently learnable for OPF problems, and discuss theoretical and practical implications for power systems operation.

Line A. Roald
University of Wisconsin-Madison
lineroald@gmail.com

Sidhant Misra
Los Alamos National Laboratory
sidhant@lanl.gov

Yee Sian Ng
Massachusetts Institute of Technology
ngyeesian@gmail.com

MS102

A Framework for Scalable Multiperiod Security Constrained AC-based Optimal Power Flow

We present a scalable framework using StructJuMP, a parallel algebraic modeling framework for block structured optimization models in Julia, and PIPS, a scalable interior point method solver written in C++. The two stage optimization structure of StructJuMP is conveyed to the Schur complement based parallelization of PIPS. Crucial to the scaling is a loose binding of the first stage with the second stage and high number of second stage scenarios. These requirements are met by multiperiod ACOPFs. We provide scaling results and bottlenecks of our current framework implemtation used in the ExaSGD ECP project.

Michel Schanen
Argonne National Laboratory
mschanen@anl.gov

Cosmin G. Petra
Lawrence Livermore National
Center for Advanced Scientific Computing
petra1@llnl.gov

Mihai Anitescu
Argonne National Laboratory
Mathematics and Computer Science Division
anitescu@mcs.anl.gov

Francois Gilbert
Argonne National Laboratory
francois.gilbert.fg@gmail.com

MS103

A New Class of High-order, Flexible, IMEX Multirate Integrators for Multiphysics Applications

Differential equations that couple multiple physics phenomena together usually require special treatment of the possible different rates at which the multiphysics evolve. Extensive research into methods that can handle multiple time scales efficiently and at high orders is essential to areas such as atmospheric modeling or chemical reaction modeling. In this talk, we focus on the implementation of multiple time stepping (MTS) algorithms derived from classes of explicit one-step exponential Runge-Kutta methods. These algorithms are applied to test problems that mainly consist of a stiff and linear fast component that is easily computed with a micro time step using either an explicit or implicit Runge-Kutta method and a slow (usually nonlinear) component that requires a larger step size. Unlike other multiple time stepping algorithms that depend on exponential integrators, the proposed class of methods avoid the computation of matrix functions. We illustrate convergence of the MTS algorithms up to fifth order and show their competitiveness in efficiency with other recent and well known multirate integrators.

Rujeko Chinomona, Vu Thai Luan
Department of Mathematics, SMU, Dallas, TX
rchinomona@smu.edu, vluan@smu.edu

Daniel R. Reynolds
Southern Methodist University
Mathematics
reynolds@smu.edu

MS103

Coupling MOR and Multirate Techniques in Multiphysics Problems

Multiphysical system comprise parts differing by its physical description and time constants. To increase efficiency in simulation these coupled systems, model order reduction and multirate techniques have been applied successfully in the past. In this talk we propose to couple both ideas by applying multirate integration to overall coupled system, whereas the big, usually slow part is replaced by a reduced

model. We discuss convergence and stability issues of this approach and verify the feasibility of this approach by inspecting an example from refined network analysis.

Michael Günther
Bergische Universität Wuppertal
Fachbereich C
guenther@math.uni-wuppertal.de

MS103

Design of High-order Multirate General Additive Runge Kutta Schemes

Multirate time integration schemes apply different step sizes to different components of a system based on the local dynamics of the components. Local selection of step sizes allows increased computational efficiency while maintaining the desired solution accuracy. The multirating idea is elegant and has been around for decades, however, difficulties faced in construction of high order multirate schemes has hampered their application. Seeking to overcome these challenges, our work focuses on the design of high-order multirate methods using the theoretical framework of *generalized additive Runge-Kutta (GARK)* methods, which provides the generic order conditions and the stability analyses. Of special interest is deriving methods that avoid unnecessary coupling between the components of the system, and allow straightforward transition to different step sizes between the steps. We present *Multirate GARK* schemes of up to order four that are explicit-explicit, implicit-explicit, and explicit-implicit in different components. We present numerical experiments illustrating the performance of these new schemes.

Arash Sarshar
Virginia Polytechnic Institute and State University
USA
sarshar@vt.edu

Steven Roberts
Virginia Tech
steven94@vt.edu

Adrian Sandu
Virginia Polytechnic Institute
and State University
sandu@cs.vt.edu

MS103

Efficient Multirate Methods from High Order

Increasingly, computational science requires large-scale simulations that consistently and accurately couple distinct physical processes. While the mathematical models for specific processes often have a well-known type (hyperbolic, parabolic, etc.) and are suitable for classical numerical integrators, the same cannot be said for the coupled models. These multiphysics models are often of mixed type, may have limited differentiability, and involve processes that evolve at dissimilar rates. As such, many multiphysics simulations require more flexible time integrators that may be tuned for these complex problems. In this talk, we discuss recent work on constructing novel 'multirate' time integration methods for such applications. While integrators that utilize different stepsizes for distinct processes are not new, traditional approaches suffer from nonrobust stability or low accuracy. The proposed class of 'relaxed multirate infinitesimal step methods' pushes the limits of multirate integration, allowing for fourth-order accuracy, robust theory, and temporal error estimators to enable stepsize adaptivity. This talk focuses on the theoretical construction of explicit-explicit MIS methods (both fast and slow components are integrated explicitly), and presents convergence and efficiency results for such methods applied to a set of standard multirate test problems.

Jean Sexton
Lawrence Berkeley National Laboratory
jmsexton@lbl.gov

Daniel R. Reynolds
Southern Methodist University
Mathematics
reynolds@smu.edu

MS104

Accelerating Electronic Structure Calculations using GPUs: From Hand-coded Loops to Relying on Third-party Libraries

Traditionally, the electronic structure community has been somewhat reluctant to use third party libraries (TPL), besides standard 1D FFT, BLAS, and LAPACK. One main reason has been performance since the best iterative solvers for nonlinear equations — such as the Kohn-Sham equations — are often hard to formulate in terms of linear solvers or standard eigensolvers without affecting performance. Things may change for codes that target new GPU-accelerated architectures, while still supporting a base of users running on standard CPU platforms. Performance portability in particular, can be an issue, and rewriting large portions of codes to target a new architecture is time consuming and can lead to code maintenance nightmares. After reviewing some of the challenges in the field, we will discuss our recent efforts at developing libraries for electronic structure computations (BML and PROGRESS) which aim at facilitating code developments on modern hardware. The goal of these libraries is to support various operations/algorithms, some specific to electronic structure, and enable performance portability with architecture specific implementations, or calls to other TPL optimized for the targeted platform through a unified interface. This work is part of the Exascale Computing Projects Co-design Center for Particle Applications (CoPA) and initially is focused on the LATTE and MGmol applications.

Jean-Luc Fattebert
Oak Ridge National Laboratory
fattebertj@ornl.gov

Susan Mniszewski
Los Alamos National Laboratory
smm@lanl.gov

MS104

Accelerating Ab Initio Quantum Monte Carlo Calculations

QMCPACK (https://qmcpack.org) is a high performance Petascale open source quantum monte carlo (QMC) package for computing the electronic structure of atoms, molecules, and solids. Implemented methods include real space diffusion Monte Carlo and auxiliary field QMC. Parallelization uses a fully hybrid (OpenMP,CUDA)/MPI approach to optimize memory usage and to take advantage of the growing number of cores per SMP node, graphical processing units (GPUs), or accelerators. I will present performance data for the Summit supercomputer and de-

scribe new algorithms to improve the efficiency of the matrix inverse updates that dominate for large problem sizes. While the current code runs at high efficiency, we are currently investigating use of Kokkos and OpenMP offload to meet the challenge of supporting multiple Exascale architectures, minimizing vendor-specific code, and increasing maintainability. I will present our current performance analyses and strategies for implementation.

Paul Kent
Oak Ridge National Laboratory
kentpr@ornl.gov

MS104

Large Scale Hybrid Density Functional Theory Calculations with a Systematic Basis Set using GPUs

Density functional theory is routinely used for electronic structure simulations due to its excellent balance between accuracy and computational efficiency. Nonetheless, its accuracy can be limited by key physical and numerical approximations, notably the choice of exchange correlation (XC) functional and basis set. For larger systems, the high computational cost means that one frequently has to choose between using a less complete and potentially inaccurate basis set, or using a traditional semilocal XC functional, which might be less well suited to the problem than a more expensive hybrid functional. We have recently introduced a highly optimized GPU implementation of the exact exchange operator in the BigDFT code, which uses a systematic wavelet basis set. The efficient GPU parallelization significantly reduces the cost of using hybrid functionals relative to traditional semilocal functionals. This allows the treatment of hundreds of atoms using hybrid functionals in a well converged basis set within a reasonable time-to-solution [Ratcliff et al., Affordable and accurate large-scale hybrid-functional calculations on GPU-accelerated supercomputers, J. Phys.: Condens. Matter, 30 (2018), 095901]. Such an approach therefore avoids the need to compromise on either XC functional or basis set quality. In this talk, we will present the approach used to calculate exact exchange on GPUs, giving examples of potential areas of use, including an application to uranium dioxide.

Laura Radcliff
Imperial College
laura.ratcliff08@imperial.ac.uk

Augustin Degomme
CEA Grenoble
augustin.degomme@cea.fr

José Flores-Livas
University of Basel
jose.flores@unibas.ch

Stefan Goedecker
Univ Basel, Switzerland
Stefan.Goedecker@unibas.ch

Luigi Genovese
Institut Nanosciences et Cryogénie, CEA
luigi.genovese@cea.fr

MS104

Experiences Porting Quantum Espresso's PWscf Solver to GPUs with CUDA Fortran

First-principle computer simulations of materials have become widespread in many fields of science and in industry. This talk focuses on efforts to port aspects of Quantum ESPRESSO (QE), namely the Plane-Wave Self-Consistent Field (PWscf) solver to GPUs using CUDA Fortran. The porting effort is presented, from profiling and identifying time-consuming procedures to performance analysis of the GPU-accelerated solver on several benchmark problems running on small workstations to large distributed GPU clusters. This talk provides an in-depth look at both CUDA Fortran programming and how to apply various tools in the context of several examples. The benefits of porting existing Fortran programs natively, over mixed source approaches with CUDA C, will be illustrated through examples where existing Fortran source can be minimally modified to run on GPUs. In particular, instances where higher level programming using CUF kernel directives provides a simple and effective means to port nested loop structures to GPU with little code modification will be discussed. Conversely, examples where straightforward application of these high-level methods can fall short and how to address these issues with reorganization of the existing source or the use of lower level CUDA kernels will be provided. The talk concludes with performance analysis of the GPU-accelerated code on benchmark cases, with discussion related to the performance implications of GPU node topology and on-node resources.

Josh Romero
NVIDIA
joshr@nvidia.com

MS105

Numerical Relativity and the Challenge of Turbulent Post-mergers

Compact object binary mergers with neutron stars are detectable both electromagnetically and via gravitational waves, but modeling the post-merger evolution is a difficult general relativistic radiation magnetohydrodynamic (MHD) problem, involving multiple length and time scales. This has driven numerical relativists to pursue strategies for approximately capturing unresolvable turbulent dynamo and transport effects, with consequences for remnant collapses, gamma ray bursts, and outflows. In this talk, I will describe post-merger simulations with neutrinos and MHD carried out by the SXS collaboration and different models for dealing with subgrid scale effects.

Matthew Duez
Department of Physics and Astronomy
Washington State University
m.duez@wsu.edu

MS105

Senr/NRPy+: Numerical Relativity in Singular Curvilinear Coordinate Systems

NRPy+ is a user-friendly, open-source Python library that takes as input human-readable mathematical expressions for partial differential equations written in Einstein notation, and outputs highly optimized SIMD C code finite-difference kernels. Building atop the powerful, open-source SymPy computer algebra system, NRPy+ is designed to lower the barrier to entry for new developers in the field of numerical relativity, as it does not rely on expensive Mathematica or Maple licenses to function. Further, to minimize

the learning curve for new users, NRPy+ code documentation takes the form of interactive Jupyter notebook tutorials hosted in the cloud by mybinder.org, which are linked to from the NRPy+ homepage: blackholesathome.net. I will present some use cases in which NRPy+ greatly optimizes workflows, focusing especially on the BlackHoles@Home project, which aims to unlock the consumer-grade desktop as a powerful tool for black hole binary astrophysics.

Zachariah B. Etienne
Department of Mathematics
West Virginia University
zachetie@gmail.com

MS105

SpECTRE: A Task-based Discontinuous Galerkin Code for Relativistic Astrophysics

We will describe SpECTRE (https://github.com/sxs-collaboration/spectre), a new open-source relativistic astrophysics code that combines a discontinuous Galerkin method with a task-based parallelism model. The goal of SpECTRE is to achieve more accurate solutions for challenging relativistic astrophysics problems such as core-collapse supernovae and binary neutron star mergers, while making efficient use of the largest supercomputers. We will discuss some preliminary results from modeling accretion disks around a black hole.

Lawrence Kidder
Cornell University
kidder@astro.cornell.edu

MS105

Modelling Relativity in Software: We Need More Mathematical Rigour in our Codes for Both Clarity and Performance

Most, if not all, current numerical relativity codes model general relativistic concepts at a very low level: Concepts such as "manifold", "vector space", or "basis" are not rigorously defined in the code and are only expressed via low-level mechanisms such as "array", "loop", or "index". This makes it tedious to express physics concepts (such as the Einstein equations) in a code, and implementations then fail to be sufficiently generic. At the same time, modern computing hardware is becoming more complex, and respective necessary code optimizations are difficult to express, further complicating the code and obscuring the physics. I argue that we need a much more mathematical and rigorous definition of physics and computing concepts in our software, and will give examples how this can improve both code clarity and code performance. I further discuss a possible path towards achieving this with current programming languages.

Erik Schnetter
Perimeter Institute
eschnetter@perimeterinstitute.ca

MS106

Scalable Integral Equation Methods for Problems with Moving Boundaries

In this talk we will present boundary integral equation (BIE) methods and related software for elliptic boundary value problems. We will discuss application of these methods to simulating particulate flows and for computing equilibrium in magnetically confined plasmas. The boundary integral formulation requires fewer unknowns since we do not need to discretize the entire volume and this results in significant savings in work. However, these methods require accurate discretization of complex three-dimensional surfaces. For moving boundaries, we need to reparameterize the surface meshes between time steps to ensure that the surfaces are accurately resolved at all times. We also require robust high-order quadratures for computing singular and near-singular boundary integrals on these surfaces. We will discuss efficient implementation of these quadratures and their acceleration using our parallel fast multipole method (FMM) library. Finally, we will present numerical results to demonstrate convergence and parallel scalability of our methods.

Dhairya Malhotra
Institute of Computational Engineering and Sciences
The University of Texas at Austin
malhotra@cims.nyu.edu

Antoine Cerfon
Courant Institute NYU
cerfon@cims.nyu.edu

Michael O'Neil
New York University
oneil@cims.nyu.edu

MS106

High Order Layer Potential Evaluation in three Dimensions using Quadrature by Expansion

The practical application of integral equation methods requires the evaluation of singular or weakly singular integrals on complex geometries. Quadrature by expansion (QBX), a method developed by Barnett, Epstein, Greengard, Klöckner, and O'Neil, stems from the observation that a layer potential induces a field which is locally smooth when restricted to either the interior or the exterior of the boundary. QBX can be used as a quadrature method in two modes—global QBX where a smooth expansion for the layer potential is formed at an off-surface location; and local QBX where a smooth expansion is formed only for a local part of the layer potential. In this talk, we present linear CPU time fast multipole accelerated local and global QBX algorithms for evaluating layer potentials with high order accuracy in three dimensions. When solving integral equations using iterative methods, the layer potentials need to be evaluated on the same geometry but with different layer potential densities. In this setup, many local quantities can be precomputed to further accelerate subsequent evaluation of the layer potentials. We compare the CPU-time and memory performance of both of these algorithms when used in the context of solving boundary value problems.

Manas N. Rachh
Applied Mathematics Department
Yale University
mrachh@flatironinstitute.org

MS106

Parallel Space-time Boundary Element Methods for the Heat Equation

The standard approach in space-time boundary element methods for discretizing boundary integral equations is using space-time tensor product spaces, originating from a separate decomposition of the boundary and the time in-

terval. This space-time discretization technique allows us to parallelize the computation of the global solution of the whole space-time system. Instead of using tensor product spaces one can also consider an arbitrary decomposition of the whole space-time boundary into boundary elements. This approach additionally allows adaptive refinement in space and time simultaneously. In this talk we consider the heat equation as a model problem and introduce a parallel solver for the space-time system. The space-time boundary mesh is decomposed into a given number of submeshes. Pairs of the submeshes represent blocks in the system matrix. Due to the structure of the matrix one has to design a suitable scheme for the distribution of the matrix blocks to the computational nodes in order to get an efficient method. In our case the distribution is based on a cyclic decomposition of complete graphs. We present numerical tests to evaluate the efficiency of the proposed parallelization approach. The presented parallel solver is based on joint work with G. Of from TU Graz, J. Zapletal and M. Merta from the Technical University of Ostrava.

Olaf Steinbach
TU Graz, Institute of Computational Mathematics
o.steinbach@tugraz.at

Stefan Dohr
TU Graz
dohr@math.tugraz.at

MS106

Fast Algorithms with Error Bounds for Quadrature by Expansion

The design of low complexity algorithms for evaluation of for layer potentials must effectively combine a singular/near singular quadrature rule with an acceleration scheme. To maintain accuracy and efficiency, both components as a unit must have well-understood error behavior. This talk considers the combination of Quadrature by Expansion (QBX) with the Fast Multipole Method (FMM). It is shown that, unless carefully controlled for, the FMM may allow for inaccurate near-field contributions to reduce the accuracy of computed QBX expansions. A set of revised geometric criteria is discussed that leads to robust FMM error estimates for QBX expansions in two and three dimensions. The algorithm that results from incorporation of these criteria recovers a level of accuracy that is the same as of the 'point' FMM.

Matt Wala
University of Illinois
Department of Scietific Computing
wala1@illinois.edu

Andreas Klöckner
University of Illinois
andreask@illinois.edu

MS107

The Shifted Boundary Method for Embedded Domain Computations: Application to Solid Mechanics

Embedded/immersed boundary methods circumvent the challenge of representing complex geometries through their ease in mesh generation. On the other hand, with such a decision arises the need to integrate over the cut elements. To counter this dilemma and maximize on the advantages of embedded methods, we propose a novel approach, named shifted boundary method. The proposed method obviates the need to integrate over the cut boundary elements by weakly imposing an equivalent boundary condition on its surrogate (formed of un-cut elements) counterpart. We start by presenting the method for the Poisson problem along with its associated Lagrangian on the surrogate domain and its corresponding boundary. We then present the method in light of elasticity problems whereby we study stability and convergence of the method with the support of a series of tests that display the robustness and accuracy of the method.

Nabil Atallah, Alex Main, Guglielmo Scovazzi
Duke University
nabil.atallah@duke.edu, geoffrey.main@duke.edu, guglielmo.scovazzi@duke.edu

MS107

Embedded Stabilized Methods for Free Surface Flow Problems

The generation of Finite Element meshes on complex geometries can be a demanding and costly task, not to mention that usually it requires the assistance or support of trained personnel. The use of embedded geometries into background (structured or unstructured) meshes is very appealing in these situations since they obviate the need for continual re-meshing in many applications involving rapid prototyping and design. Unfortunately, many finite element embedded boundary methods for incompressible flow are also difficult to implement due to the need to perform complex cell cutting operations at boundaries, and the consequences that these operations may have on the overall conditioning of the ensuing algebraic problems. We present a new, stable, and simple embedded boundary method, which we call shifted boundary method (SBM), that eliminates the need to perform cell cutting. Boundary conditions are imposed on a surrogate discrete boundary, lying on the interior of the true boundary interface. We then construct appropriate field extension operators, with the purpose of preserving accuracy when imposing the boundary conditions. We demonstrate the SBM applied to several free surface flow problems.

Oriol Colomes, Leo Nouveau, Guglielmo Scovazzi
Duke University
oriol.colomes@duke.edu, leo.nouveau@duke.edu, guglielmo.scovazzi@duke.edu

MS107

The Shifted Interface Method for Embedded Interface Computations

We propose a new finite element method for embedded interface computations, which falls in the category of surrogate /approximate interface algorithms. The key idea of this method is shifting the location where interface jump conditions are applied from the true to the surrogate interface. In order to preserve optimal convergence rates of the numerical solution, the jump conditions will be appropriately modified on the surrogate interface. From our view of point, this method is simple, robust, and efficient. In this presentation, we apply this new concept to the Poisson interface problems. We present numerical tests for the Poisson equations on progressively more complex geometries.

Kangan Li, Nabil Atallah, Guglielmo Scovazzi
Duke University

kangan.li@duke.edu, nabil.atallah@duke.edu, guglielmo.scovazzi@duke.edu

MS107

The Shifted Boundary Method: A New Framework for Embedded Domain Computations

Embedded boundary methods obviate the need for continual re-meshing in many applications involving rapid prototyping and design. Unfortunately, many finite element embedded boundary methods are also difficult to implement due to the need to perform complex cell cutting operations at boundaries, and the consequences that these operations may have on the overall conditioning and numerical stability of the ensuing algebraic problems. We present a new, stable, accurate, and simple embedded boundary method, which we call shifted boundary method (SBM), that eliminates the need to perform cell cutting. Boundary conditions are imposed on a surrogate discrete boundary, lying on the interior of the true boundary interface. We then construct appropriate field extension operators, with the purpose of preserving accuracy when imposing the boundary conditions. We demonstrate the SBM on large-scale incompressible flow problems, free-surface problems, multiphase flow problems, solid mechanics problems, and shallow water flow problems.

Guglielmo Scovazzi, Alex Main, Nabil Atallah, Ting Song, Kangan Li, Oriol Colomés, Leo Nouveau
Duke University
guglielmo.scovazzi@duke.edu, geoffrey.main@duke.edu, nabil.atallah@duke.edu, ting.song@duke.edu, kangan.li@duke.edu, oriol.colomes@duke.edu, leo.nouveau@duke.edu

MS108

Nngp: Deep Neural Networks Combined with Gaussian Processes

Abstract not available.

George E. Karniadakis
Brown University
Division of Applied Mathematics
george_karniadakis@brown.edu

MS108

Deep Learning of Dynamics and Signal-noise Decomposition

A critical challenge in the data-driven modeling of dynamical systems is producing methods robust to measurement error, particularly when data is limited. Many leading methods either rely on denoising prior to learning or on access to large volumes of data to average over the effect of noise. We propose a novel paradigm for data-driven modeling that simultaneously learns the dynamics and estimates the measurement noise at each observation. Our method explicitly accounts for measurement error in the map between observations, treating both the measurement error and the dynamics as unknowns to be identified, rather than assuming idealized noiseless trajectories. We model the unknown vector field using a neural network, imposing a Runge-Kutta integrator structure to isolate this vector field, even when the data has a non-uniform time-step, thus constraining and focusing the modeling effort. We demonstrate the ability of this framework to form predictive models on a variety of canonical test problems of increasing complexity and show that it is robust to substantial amounts of measurement error.

Samuel Rudy
Applied Mathematics
University of Washington
shrudy@uw.edu

J. Nathan Kutz
University of Washington, Seattle
Dept of Applied Mathematics
kutz@uw.edu

Steven Brunton
University of Washington
sbrunton@uw.edu

MS108

A Sequential Sampling Strategy for Extreme Event Statistics in Nonlinear Dynamical Systems

We develop a method for the evaluation of extreme event statistics associated with nonlinear dynamical systems from a small number of samples. From an initial dataset of design points, we formulate a sequential strategy that provides the next-best data point (set of parameters) that when evaluated results in improved estimates of the probability density function (pdf) for a scalar quantity of interest. The approach utilizes Gaussian process regression to perform Bayesian inference on the parameter-to-observation map describing the quantity of interest. We then approximate the desired pdf along with uncertainty bounds utilizing the posterior distribution of the inferred map. The next-best design point is sequentially determined through an optimization procedure that selects the point in parameter space that maximally reduces uncertainty between the estimated bounds of the pdf prediction. We apply the method to estimate the extreme event statistics for a very high-dimensional system with millions of degrees of freedom: an offshore platform subjected to three-dimensional irregular waves.

Themistoklis Sapsis
Massachusetts Institute of Techonology
sapsis@mit.edu

Mustafa A. Mohamad
Mechanical Engineering
MIT
mmus@nyu.edu

MS108

Numerical Aspects for Approximating Governing Equations using Data

Recently there has been a growing interest in discovering governing equations of certain physical problems using observational data. In this talk, we are going to discuss effective numerical algorithms for locally recovering unknown governing differential equations from measurement data. We employ a set of standard basis functions, e.g., polynomials, to approximate the governing equation with high accuracy. Upon recasting the problem into a function approximation problem, we discuss several important aspects for accurate approximation. Most notably, we discuss the importance of using a large number of short burst of trajectory data, rather than using data from a single long trajectory. Several options for the numerical algorithms to perform accurate approximation will be presented, along

with an error estimate of the final equation approximation. We will present numerical examples of both linear and nonlinear systems to demonstrate the properties and effectiveness of the equation recovery algorithms.

Dongbin Xiu, Kailiang Wu
Ohio State University
xiu.16@osu.edu, kailiangmath@gmail.com

MS109

Spatio-temporal Tensor Super-resolution using Chemical Priors

Fusing a low-spatial resolution hyperspectral image (HSI) and a high-spatial resolution multispectral image (MSI) of the same scene leads to a super-resolution image (SRI), which is information rich spatially and spectrally. In this talk, I shall discuss two approaches to achieve spatio-spectral super-resolution. The first approach uses a spatial graph Laplacian regularizer defined on the MSI and the other via spectral basis correlation. Both regularizations are driven from prior knowledge of the chemical constituents of the imaged domain. Unlike earlier work, these approaches don't assume prior knowledge about the spatial degradation process from SRI to HSI, nor a perfectly aligned HSI and MSI tensor pair. (Joint work with Dr. Tianming Wang).

Chandrajit Bajaj
Dept. of Computer Science, University of Texas at Austin
bajaj@ices.utexas.edu

Pradeep Ravikumar
Carnegie Mellon
pradeepr@cs.cmu.edu

MS109

Physics-constrained Surrogates for Reduced-order Modeling and Uncertainty Quantification

Modern computer simulations of physical processes frequently involve a large number of random, input parameters x. While the statistical estimation of output quantities $y(x)$ (i.e. the forward UQ or uncertainty propagation problem,) may directly be carried out using Monte Carlo, it is particularly impractical when output variances are high and only few simulations/samples N are feasible. In this paper, we present a novel strategy to construct accurate probabilistic surrogates $p(y|x)$ based on very limited samples N and very high-dimensional inputs i.e. $d = \dim(x) >> 1$ and discuss their applicability to both uncertainty propagation as well as Bayesian inverse problems. To cope with the large d, small N regime, instead of using off-the-shelf machine learning methods, we directly incorporate physical principles by leveraging reduced-order models based on both simplified constitutive behavior and coarser spatial resolution than the forward code $y(x)$. By making use of information-theoretic criteria, the surrogate complexity is automatically adapted to the available amount of data and a user-specified ratio of accuracy to computational cost. Our approach is validated on both forward and inverse problems in random heterogeneous media with several thousands of stochastic input parameters x.

Constantin Grigo
Technical University of Munich
constantin.grigo@tum.de

Phaedon-Stelios Koutsourelakis
Technische Universität München
Germany
p.s.koutsourelakis@tum.de

MS109

Neural Networks Learning Quantum Chemistry

Historically, computational chemistry has been unable to overcome the orthogonal requirements of speed and accuracy. In this talk, we will present a fully transferable deep learning potential that is applicable to complex and diverse molecular systems well beyond the training dataset. Recently we introduced ANAKIN-ME (Accurate NeurAl networK engINe for Molecular Energies) or ANI in short. ANI is a new method and sampling procedure for training neural network potentials that utilizes a special kind of symmetry functions to build single-atom atomic environment vectors (AEV) as a molecular representation. The ANI method, which focuses on the use of large and diverse data sets in training new potentials, has consistently proven to be universally applicable to systems containing the atomic species in the training set. Focusing on parametrization for organic molecules, we have developed a universal ANI potential which is highly accurate compared to reference QM calculations at speeds 107 faster. The ANI potential is shown to accurately represent the underlying physical chemistry of molecules through various test cases including: chemical reactions (both thermodynamics and kinetics), thermochemistry, structural optimization, and molecular dynamics simulations. The results presented in this talk will provide evidence of the universal applicability of ANI deep learning potential to various chemistry problems involving organic molecules.

Olexandr Isayev
University of North Carolin
olexandr@olexandrisayev.com

MS109

Enforcing Physical Constraints in Machine Learning with Application to Fluid Flows

Recently, several researchers have adopted machine-learning techniques to the modeling of complex physical systems. However, important physical constraints such as conservation laws and high-order statistics are usually not well preserved in existing machine-learning techniques, demanding the development of improved machine-learning tools for science and engineering. In this work, we present a physics-informed generative adversarial network (GAN) by enforcing constraints of both conservation laws and certain statistics from the training data. We show that the physics-informed GAN is more robust and better captures high-order statistics. In addition, the physics-informed GAN generates more realistic flow fields than standard GANs. This work has great potential as an alternative to the explicit modeling of closures for unresolved physics, which account for a major source of uncertainties when simulating complex systems, e.g., turbulence and Earths climate.

Jinlong Wu
Dept. of Aerospace and Ocean Engineering, Virginia Tech
jinlong@vt.edu

Yang Zeng
Virginia Tech
yangzeng@vt.edu

Karthik Kashinath, Adrian Albert, Mr Prabhat
Lawrence Berkeley National Laboratory
kkashinath@lbl.gov, aalbert@lbl.gov, prabhat@lbl.gov

Heng Xiao
Dept. of Aerospace and Ocean Engineering, Virginia Tech
hengxiao@vt.edu

MS110

Approximate Bayesian Model Inversion for PDEs with Unknown Parameters and Constitutive Relations

We present two approximate Bayesian inference methods for model inversion, with application to systems governed by PDEs with unknown spatially heterogeneous model parameters and unknown constitutive relations. The methods presented employ parameterized multivariate Gaussian densities for the prior and posterior, where the prior hyperparameters and the posterior parameters are estimated by maximizing the expectation lower bound (ELBO), a variational lower bound on the marginal likelihood of the data. In the first method, we employ Expectation Maximization with a Laplace approximation to the posterior on the E-step to maximize the ELBO. In the second method, we employ doubly stochastic variational inference (DSVI) together with stochastic gradient for maximizing the ELBO and Gaussian backpropagation for computing noisy gradient estimates. The proposed methods are applied to identifying diffusion coefficients in linear and nonlinear stationary diffusion equations.

David A. Barajas-Solano
Pacific Northwest National Laboratory
david.barajas-solano@pnnl.gov

Alexandre M. Tartakovsky
PNNL
alexandre.tartakovsky@pnnl.gov

MS110

Maximum Entropy Inference for Learning Model Parameters when Data is Unreported

Physical models employed in predictive simulations are often constructed though comparison to experimental data. Techniques such as Bayesian inference can be employed in such settings to estimate model parameters with uncertainty from limited, typically noisy, data. However, raw experimental data are usually unreported, with the information from the experiment often summarized using low-dimensional summary statistics (e.g. means, standard deviations) of fitting model parameters fit to the data. Such presentation may exclude details about correlation structure between the parameters, hindering the ability to perform analyses such as uncertainty propagation. In order to recover important information such as correlation structure we employ a maximum information entropy statistical inference procedure, employing an Approximate Bayesian Computation (ABC) approach to construct inferences on the unreported experimental data itself in the context of the data model employed by the original experimentalists. In this way, the data model coupled with the available summary statistics and additional experiment context constrain a space of hypothetical data which are consistent with all available information which can then be analyzed in a consensus fashion using Bayesian inference to recover correlations. This approach also enables combination of hypothetical data across different experiments and replacement of the original fitting models with models of choice.

Tiernan Casey
Sandia National Laboratories
tcasey@sandia.gov

MS110

Multifideity Informed Neural Network in Reduced Order Modeling

In this talk, we will discuss a neural network-based reduced basis method with multi-fidelity models to accurately approximate the reduced solutions. The combination of low/high-fidelity solutions helps improve the accuracy of the neural networks as the approximator, and this method demonstrates its ability to produce more accurate results with a limited number of high-fidelity simulations with an affordable computational cost. We also provide several numerical examples to illustrate the effectiveness of this method.

Chuan Lu
University of Iowa
chuan-lu@uiowa.edu

Xueyu Zhu
University of Iowa
Department of Mathematics
xueyu-zhu@uiowa.edu

MS110

Physics Informed Deep Neural Networks in the Data Limited Regime

We present a physics informed deep neural network (DNN) method for estimating parameters and unknown physics (constitutive relationships) in partial differential equation (PDE) models in the presence of partial data. In particular, we will concentrate on the problem of estimating the unknown space-dependent diffusion coefficient in a linear diffusion equation and an unknown constitutive relationship in a non-linear diffusion equation. For the parameter estimation problem, we assume that partial measurements of the coefficient and states are available and demonstrate that under these conditions, the proposed method is more accurate than state-of-the-art methods e.g. Maximum a posteriori estimation. Finally, we present a systematic study of the method and demonstrate that the proposed method remains accurate in the presence of measurement noise.

Carlos Ortiz-Marrero
Pacific Northwest National Laboratory
carlos.ortizmarrero@pnnl.gov

MS111

Large Scale Parallel Solution Methods for Electromagnetic Simulations

A key challenge in supporting massively parallel electromagnetic simulations at Sandia is developing appropriate high performance linear solver technologies that are scalable and portable on differing advanced architectures. In this talk, we discuss the development of block solvers in Trilinos that target linear systems arising from electromagnetic simulations within a Sandia simulation code. A critical component of the solve is deploying a previously devel-

oped algebraic multigrid method for Maxwell's equations within the templated solver stack in the Sandia Trilinos project. We describe challenges that arise from requirements imposed by the simulation, highlight several solver approaches based on Schur complement formulations, and finally present scalability results to justify the current solution method.

Jonathan J. Hu
Sandia National Laboratories
Livermore, CA 94551
jhu@sandia.gov

Christian Glusa, Paul Lin
Sandia National Laboratories
caglusa@sandia.gov, ptlin@sandia.gov

MS111

Robust AMG Solvers for Linear Elasticity Problems

In this talk, we will present our recent development on algebraic multigrid (AMG) methods for solving sparse linear systems originated from linear elasticity problems. The new algorithms can integrate the rigid-body mode vectors that span the near null space of the global matrix, which are usually available "for free' from the discretizations. The information contained in the target vectors are used to determine the C/F splittings, and also to compute the interpolation coefficients by solving small constrained minimization problems locally. The presented approaches are based on the framework of AMGe, where the local operators are computed algebraically by defining local extension mappings as done in the element-free AMGe work. So, this work is a natural extension to element-free AMGe to be able to handle an arbitrary number of test vectors. With a proper smoother, numerical results suggested that compared with standard AMG, these approaches can be more robust and efficient for linear elasticity problems.

Ruipeng Li
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
li50@llnl.gov

Robert D. Falgout
Lawrence Livermore National Laboratory
rfalgout@llnl.gov

Victor A. Paludetto Magri
University of Padova
victor.magri@dicea.unipd.it

MS111

Node-aware Communication in Parallel AMG

Algebraic multigrid is a robust and highly effective solver for sparse linear systems for a range of problems. In parallel, the scalability of AMG is challenged by large interprocess communication costs within sparse matrix operations, namely the sparse matrix-matrix (SpGEMM) and sparse matrix-vector (SpMV) multiplication. The cost of communication is particularly high in inter-node messages, which often incur additional penalties due to injection bandwidth limits and contention of the network. We present a method for node-aware communication for SpGEMMs and SpMVs used in AMG, aggregating messages on each node before injecting data into the network. In addition, we present several performance models to highlight the inherent costs of AMG cycling and setup, the limits of node-to-node communication, and potential contention in the network during computation.

Luke Olson
University of Illinois Urbana-Champaign
lukeo@illinois.edu

Amanda Bienz, William D. Gropp
University of Illinois at Urbana-Champaign
bienz2@illinois.edu, wgropp@illinois.edu

MS111

Scalable Structured Solvers on Emerging Architectures

Line and plane relaxation are important components for robust variational multigrid methods on structured grids when using standard coarsening with anisotropic problems. These methods constitute the majority of time spent in a multilevel solve–driving the performance of the solver. Tridiagonal systems used in line relaxation are solved using a memory efficient, multilevel parallel partitioning algorithm. In plane relaxation, 2D multilevel cycles are run on each plane. This results in a series of independent parallel cycles. To improve parallel efficiency, the uniform communication pattern of these cycles can be exploited to aggregate communication across planes on a process. A straightforward application of this approach, when used with coarse-grid redistribution at scale, is limited by the maximum number of communicators in MPI so coordination is needed among planes. In this talk, we will present an approach for automating this aggregation using lightweight user-level threads.

Andrew Reisner
University of Illinois at Urbana-Champaign
areisne2@illinois.edu

Luke Olson
University of Illinois Urbana-Champaign
lukeo@illinois.edu

David Moulton
Los Alamos National Laboratory
Applied Mathematics and Plasma Physics
moulton@lanl.gov

MS112

Certified Reduced Basis Methods for Fractional Laplace Equations via Extension

Fractional Laplace equations are becoming important tools for mathematical modeling and prediction. Recent years have shown much progress in developing accurate and robust algorithms to numerically solve such problems, yet most solvers for fractional problems are computationally expensive. Practitioners are often interested in choosing the fractional exponent of the mathematical model to match experimental and/or observational data; this requires the computational solution to the fractional equation for several values of the both exponent and other parameters that enter the model, which is a computationally expensive many-query problem. To address this difficulty, we present a model order reduction strategy for fractional Laplace problems utilizing the reduced basis method (RBM). Our RBM algorithm for this fractional partial differential equation (PDE) allows us to accomplish significant acceleration compared to a traditional PDE solver

while maintaining accuracy. Our numerical results demonstrate this accuracy and efficiency of our RBM algorithm on fractional Laplace problems in two spatial dimensions.

Harbir Antil
George Mason University
Fairfax, VA
hantil@gmu.edu

Yanlai Chen
Department of Mathematics
University of Massachusetts Dartmouth
yanlai.chen@umassd.edu

Akil Narayan
University of Utah
akil@sci.utah.edu

MS112

A Low-rank Multigrid Method for the Stochastic Steady-state Diffusion Equation

We study a multigrid method for solving large linear systems of equations with tensor product structure. Such systems are obtained from stochastic finite element discretization of stochastic partial differential equations such as the steady-state diffusion problem with random coefficients. When the variance in the problem is not too large, the solution can be well approximated by a low-rank object. In the proposed multigrid algorithm, the matrix iterates are truncated to low rank to reduce memory requirements and computational effort. The method is proved convergent with an analytic error bound. Numerical experiments show its effectiveness in solving the Galerkin systems compared to the original multigrid solver, especially when the number of degrees of freedom associated with the spatial discretization is large.

Howard C. Elman, Tengfei Su
University of Maryland, College Park
elman@cs.umd.edu, tengfesu@math.umd.edu

MS112

Low-rank Solution Methods for Stochastic Eigenvalue Problems

We study efficient solution methods for stochastic eigenvalue problems arising from discretization of self-adjoint partial differential equations with random data. With the stochastic Galerkin approach, the solutions are represented as generalized polynomial chaos expansions. A low-rank variant of the inverse subspace iteration algorithm is presented for computing one or several minimal eigenvalues and corresponding eigenvectors of parameter-dependent matrices. In the algorithm, the iterates are approximated by low-rank matrices, which leads to significant cost savings. The algorithm is tested on two benchmark problems, a stochastic diffusion problem with some poorly separated eigenvalues, and an operator derived from a discrete stochastic Stokes problem whose minimal eigenvalue is related to the inf-sup stability constant. Numerical experiments show that the low-rank algorithm produces accurate solutions compared to the Monte Carlo method, and it uses much less computational time than the original algorithm without low-rank approximation.

Tengfei Su, Howard C. Elman
University of Maryland, College Park
tengfesu@math.umd.edu, elman@cs.umd.edu

MS112

Rank-adaptive Tensor Recovery Based Model Reduction for PDEs with High-dimensional Random Inputs

This work proposes a systematic model reduction approach based on rank adaptive tensor recovery for partial differential equation (PDE) models with high-dimensional random parameters. Since the outputs of interest of these models are typically approximate solutions on given physical grids which are also high-dimensional, we use kernel principal component analysis to construct stochastic collocation approximations in reduced-dimension spaces of the outputs. To address the issue of high-dimensional random inputs, we develop an efficient rank adaptive tensor recovery approach to compute the collocation coefficients. In particular, efficient stable initialization strategies for non-convex optimization problems involved in tensor recovery are developed in this work. We present a general mathematical framework of our overall model reduction, analyze its stability and demonstrate its efficiency with numerical experiments.

Kejun Tang, Qifeng Liao
ShanghaiTech University
tangkj@shanghaitech.edu.cn,
liaoqf@shanghaitech.edu.cn

MS113

An Approximate Empirical Bayesian Method for Large-Scale Linear-Gaussian Inverse Problems

Abstract not available.

Jinglai Li
Shanghai Jiao Tong University
lijinglai@gmail.com

MS113

Bayesian Identification of Discontinuous Fields with an Ensemble-based Variable Separation Multiscale Method

This work presents a multiscale model reduction approach to discontinuous fields identification problems in the framework of Bayesian inference. We use an ensemble-based variable separation method to approximate multiscale basis functions used to construct a coarse model. We construct the variable-separation expressions for stochastic multiscale basis functions based on the random field, which is treated Gauss process as prior information. To this end, multiple local inhomogeneous Dirichlet boundary condition problems are required to be solved, and the ensemble-based method is used to obtain variable separation forms for the corresponding local functions. The local functions share the same interpolate rule for different physical basis functions in each coarse block. This is a rough approximation but improves the efficiency of computation greatly. We obtain the variable separation expression of multiscale basis functions, which can be used to the models with different boundary conditions and source terms, once the expression constructed. The proposed method is applied to discontinuous field identification problems where the hybrid of total variation and Gaussian densities are imposed as the penalty.

Lijian Jiang

Los Alamos National Lab
ljjiang@tongji.edu.cn

Guang Lin
Purdue University
guanglin@purdue.edu

Na Ou
Hunan University
oyoungla@hnu.edu.cn

MS113

Bayesian Inference for Structural Error Quantification

Conventional Bayesian inference typically assumes that the computational model replicates the true mechanism behind data generation. As a result, calibrated model parameters are often biased, leading to deficient predictive skills. Augmenting model outputs with statistical correction terms may remove the predictive bias, but it can violate physical laws, make the calibrated model ineffective for predicting non-observable quantities, and experience identifiability challenges in distinguishing between data noise and model error. This work will present a framework for representing, quantifying and propagating uncertainties due to model structural errors by embedding stochastic correction terms in the model. The embedded correction approach ensures physical constraints are satisfied, and renders calibrated model predictions meaningful and robust with respect to structural errors. The physical inputs and correction parameters are simultaneously inferred via surrogate-enabled Markov chain Monte Carlo. With a polynomial chaos characterization of the correction term, the approach allows efficient decomposition of uncertainty that includes contributions from data noise, parameter posterior uncertainty, and model error. The developed structural error quantification workflow is implemented in UQ Toolkit (www.sandia.gov/uqtoolkit). We will discuss the challenges associated with the increased dimensionality of the associated Bayesian problem, and steps that are taken to alleviate them.

Khachik Sargsyan, Xun Huan
Sandia National Laboratories
ksargsy@sandia.gov, xhuan@umich.edu

Habib N. Najm
Sandia National Laboratories
Livermore, CA, USA
hnnajm@sandia.gov

MS113

Title Not Available

Abstract not available.

Xiang Zhou
Department of Mathematics
City Univeristy of Hong Kong
xiang.zhou@cityu.edu.hk

MS114

Parallel High-resolution Compact Schemes for Thermal Convection Simulation

High-resolution parallel simulation is a vital tool for the study of the atmospheric boundary layer. As a first step in the development of a high-resolution parallel compact simulator, we consider a classical problem of flow over a heated plane. In our considerations, we use the implicit large-eddy-simulation (ILES) approach with compact non-oscillatory high-order approximation schemes. This method produces an accurate resolution of the boundary layer convection. The main focus of this presentation is on the scalability of the proposed algorithm. The results of the numerical calculations on a sequence of grids in multicore and multi-node environment are presented. These results confirm the high efficiency of the developed approach.

Ronald L. Gonzales, Yuiry Gryazin
Idaho State University
gonzrona@isu.edu, gryazin@isu.edu

MS114

Parallel High-resolution Compact Krylov-FFT-type Algorithms for Subsurface Scattering Problems

In this paper, we consider a parallel Krylov-FFT type high-resolution algorithm for the solution of a subsurface electromagnetic scattering problem. The 3D Helmholtz equation is discretized by high-order compact finite-difference schemes. The resulting systems of finite-difference equations are solved by different preconditioned Krylov subspace-based methods. The FFT-based preconditioner is used for efficient implementation of the developed iterative approach. The main focus of the paper is the efficient parallel implementation of the proposed algorithm in the shared (OpenMP) and distributed (MPI) memory environment. The complexity and scalability of the parallel preconditioned iterative method are analyzed on scattering problems with realistic ranges of parameters in soil and mine-like targets.

Yury Grazin, Ronald L. Gonzales, Yun Teck Lee
Idaho State University
gryazin@isu.edu, gonzrona@isu.edu, leeyunt@isu.edu

MS114

Data Quality Challenges with Missing Values and Mixed Types in Joint Sequence Analysis

This work is focused on developing foundational techniques to process time series from decision science applications. More specifically, we concentrate on developing resilient algorithms to compare similarity of time series (or sequences) of mixed types for the purpose of classifying and predicting behaviors of energy consumers. For example, to understand the transportation trends and their impact on consumption, we need to identify relationships between lifecycle patterns/phases and certain types of decisions or choices (such as the choice to purchase a home, own a car, invest in new technologies, etc.). Well-known algorithms including dimensionality reduction and clustering might be effective for these tasks, however, a number of challenges in the consumer data make it difficult to even perform the basic comparison operations. One core challenge is that the time series are composed of categorical values that do not have an obvious metric for distance. Moreover, to compare two consumers, we have to compare a number of different types of attributes. Additionally, the time series frequently contain missing values, unknown values or corrupted values, which presents yet another challenge. Our aim is to investigate how missing values in time series data affect state-of-art similarity measures and clustering techniques.

Alina Lazar
Youngstown University
alazar@ysu.edu

MS114

Scalable Direct Solver for Compact Stencil Calculations on Rectangular Grids

The recent development of parallel CPU technologies on modern desktop computers makes sequential numerical algorithms a thing of the past. It is obvious that the main performance improvement in the upcoming years will be made based on increasing the number of cores on the modern computers. It shifts the focus of the algorithmic research from the development of the sequential numerical methods to the parallel methodology. In this paper, we consider an efficient parallel implementation of the direct compact six order approximation numerical solver for the Helmholtz equation on the multicores computers and multi-node clusters. The developed high-order parallel direct algorithms are based on a combination of the separation of variables technique and the Fast Fourier Transform (FFT) type methods. The results of implementation of these methods in OpenMP, MPI and Hybrid programming environment are presented. We also consider the generalization of the presented algorithms to the solution of linear systems obtained from approximations on the compact 27-point stencils on the rectangular grids with similar stencil coefficients. In this paper, we focus on scalability of the developed numerical technique.

Yun Teck Lee
Idaho State University
leeyunt@isu.edu

MS115

Cyber Security in a High Performance Computing Environment

High performance computing (HPC) environments are unique in their intent to speed up computation, while minimizing administrative overhead. The intent in most HPC environments is to get things done as fast and efficiently as possible. Security is sometimes seen as an unnecessary administrative burden that just takes the resource that is usually so scarce in HPC systems, time and bandwidth. Nevertheless, it is needed. There was a time when HPC systems were isolated from the outside world, and all traffic into and out of those systems went through a gateway that could be used to prevent intrusions and breaches. Unfortunately, malware has become so sophisticated, that attacks can be hidden in packed code that is hard to detect, and may slip past this gateway. Systems are needed that can monitor activities on an HPC system, detect malicious activity and intervene before significant damage is done to the system. These systems exist in the form of Autonomic Security Management Systems. These systems can monitor the traffic in an HPC system using only a minimum number of nodes, and can analyze that traffic for patterns that may be indicative of malicious actions. This talk will describe these environments, their uses, benefits, and drawbacks to present a case for their use in all HPC environments.

David Dampier
Professor and Department Chair Information Systems and Cyber
University of San Antonio, TX
david.dampier@utsa.edu

MS115

Resilience Modeling and Analysis of Scheduling Scientific Applications

Resilience to faults is one of the most important challenges for future HPC systems as the computing systems are expected to scale up in component count and component reliability is expected to decrease. The inherent instability of modern computing systems, the envisioned diversity of their faults and the limitations of generic one fits all fault tolerance approaches mandates reconsideration of these approaches to alleviate the issue of resilience. Moreover, designing scheduling algorithms integrated with fault tolerance techniques with the goal of achieving resiliency does not assure their robust performance, due to the risks (such as undetected errors, increase in number of faults, overhead) that the fault tolerance mechanisms themselves exhibit, due to their flexibility and complexity, unforeseen and sometime harmful behavior. The challenge is to formulate an approach to assess the balancing of such risks, and to soundly quantify the resiliency of the fault-tolerant scheduling algorithms in the presence of environmental perturbations (unpredictable errors etc.) in order to evaluate whether or not a scheduling technique is more resilient than another. This talk will cover a simulation-based methodology and a high level formalism based on Stochastic Process Algebra (SPA) used for performing performance and resilience analysis of algorithms for scheduling scientific application on high performance computing systems.

Nitin Sukhija
Slippery Rock University of Pennsylvania
nitin.sukhija@sru.edu

MS115

Cybersecurity for Trustworthy Science: The NSF Cybersecurity Center of Excellence

Think about cybersecurity and scientific research at the same time and you probably think about regulated data (e.g. HIPAA) or other compliance (e.g. NIST 800-171). These are certainly critical for any researcher or information security professional to consider, but what about other science that falls outside that scope - open science such as physics, astronomy, chemistry, biology, ecology, climate science, computer science, etc.? What role does cybersecurity play in open science? A key goal of scientific research is to produce trustworthy results and to be reproducible. Given the prevalence of computing in supporting scientific research, cybersecurity has a strong role to play in support any research by ensuring that the computing it relies on has appropriate assurances for integrity, is available, and, for some circumstances, provides confidentiality while minimizing impact on scientific productivity. Trusted CI, the NSF Cybersecurity Center of Excellence, has been addressing this challenge since its inception five years ago. It has worked closely with dozens of NSF projects on their cybersecurity challenges and impacts hundreds of projects through its best practices, training, and the annual NSF cybersecurity summit.This talk will cover the role and importance of cybersecurity in support scientific research and the resources available from Trusted CI for researchers and information security professionals.

Von Welch
Director of Center for Applied Cybersecurity Research

Indiana University
vwelch@iu.edu

MS115

Minisymposterium: Leveraging Data Analytics Tools for Fraud Prevention

In todays fast paced business world, data typically resides on information technology (IT) systems. As a result, automation of business processes has become more common in order to increase the efficiency and effectiveness of managing large volumes of data. Due to the increase in these automated processes, finance professionals have sought to rely on automated controls to prevent and detect fraudulent activity. Finance professionals have started to leverage the use of Audit Command Language (ACL), Microsoft Excel, IDEA and other data analytics tools in the fight against fraud. Fraud professionals use techniques such as rule based analytics, anomaly detection analytics and predictive analytics to identify known, unknown and complex patterns in voluminous data. The use of these data analytics tools for fraud prevention and detection allows large organizations to effectively test and monitor their internal control environment at the transaction level. When a robust and effective fraud data analytics program is implemented, it can provide an organization with a high level of assurance that its internal controls are operating appropriately and reduce the instance of asset misappropriation. This talk will cover how fraud professionals leverage the use of data analytics tools for fraud prevention in organizations.

Blair William
Corporate Director, Financial Compliance and Process Reengin
UPMC Corporate Services
williamba@upmc.edu

MS116

Combining Structured and Unstructured Meshing for Efficient and Robust 3D Fluid-structure Interaction Applications

We present a new approach to the solution of fluid-structure interaction problems in three dimensions, based upon the combination of a block-structured background mesh and a Lagrangian or ALE (arbitrary Lagrangian-Eulerian) unstructured tetrahedral mesh that evolves as the solid structure deforms. Our algorithm, which combines elements of the fictitious domain method (FDM) and the immersed finite element method (IFEM), solves a modified flow problem on the block-structured mesh and then updates the solid structure based upon an interpolated velocity field. The talk will describe details of the modified flow problem, to account for the presence of the solid structure, which is the main novelty of the proposed methodology. Details of the numerical solution methods will also be provided, including discussion of computational efficiency through adaptive mesh refinement (AMR) and preconditioning. Results will be presented for a selection of three-dimensional test problems involving large structural deformations. These problems are selected to validate the results and to demonstrate the robustness of the methodology across different material properties, geometries and flow regimes.

Yongxing Wang, Peter K. Jimack, Mark A. Walkley
School of Computing
University of Leeds
scywa@leeds.ac.uk, p.k.jimack@leeds.ac.uk, m.a.walkley@leeds.ac.uk

MS116

Data Remapping between Distributed General Unstructured Meshes

Data remapping between distributed general unstructured meshes is an often encountered and non-trivial step in the numerical simulations of multi-physics applications. Such applications require fast and scalable remapping strategies satisfying accuracy, conservation and bounds preservation constraints as the remapping may be used as frequently as at each time step. Towards this end, we present Portage, a flexible, customizable and scalable framework for remapping fields between distributed general unstructured meshes and particle clouds. Portage uses mesh-mesh intersection to conservatively interpolate single/multi-material fields between distributed source and target meshes. Portage is a modular and extensible library of remap components. By design, these components can be used independently or composed into a custom remapper as required by a specific application. It currently supports both distributed and on-node parallelism to run effectively on advanced architectures.

Gary Dilts
LANL
gad@lanl.gov

Rachel Ertl, Charles R. Ferenbaugh, Rao V. Garimella, Angela Herring, Eugene Kikinzon, Chris Malone
Los Alamos National Laboratory
ertl@lanl.gov, cferenba@lanl.gov, rao@lanl.gov, angelah@lanl.gov, kikinzon@lanl.gov, cmalone@lanl.gov

Navamita Ray
Argonne National Laboratory
nray@lanl.gov

Daniel Shevitz, Jan Velechovsky
Los Alamos National Laboratory
shvitz@lanl.gov, jan@lanl.gov

MS116

A Vision for Research in Unstructured Mesh Generation

In this talk, I will outline several important open problems in the area of unstructured mesh generation. Open problems may come from any area of unstructured mesh generation including high-order mesh generation, parallel mesh generation, mesh untangling, mesh adaptation, and mesh warping, as well as their applications. The goal of the talk is to establish a research vision for unstructured mesh generation. Audience members are encouraged to take part in the discussion and add their own ideas to the list of open problems.

Suzanne M. Shontz
University of Kansas
shontz@ku.edu

MS116

Size Field Driven Advancing Layer and Hex Core Hybrid Volume Meshing

A multi-component hybrid meshing technique for viscous

CFD applications is presented. Each mesh component is locally sized by a global field function defined either from user-set sources (geometric shapes with specified spacing and decay) or external point cloud (e.g., flow solution) data. Starting with a set of closed surface triangulations representing the inner/outer boundaries of the model, with mesh spacing controlled by the size field, an advancing layer method deforms vertices along directions orthogonal to the surface, forming tetrahedra underneath the advancing points. Vertices are halted during layer progression due to poor cell quality or proximity/collision with adjacent fronts. Other topology operators are applied as needed to drive the front towards isotropy at the final layer. Tets are then combined into pyramids, prisms and hexahedra using adjacency data recorded during layer advancement. The final anisotropic front seeds a Cartesian voxel block graded to match the size and positions on the front, as well as the global size field. The voxel block is intersected with the front triangles, and is flood-filled to pare the voxel set to fit wholly within the fronts. Finally, regions between the viscous layers and voxel boundaries are filled with tets using a modified-Delaunay scheme, with local cells driven by the size field and efficiently evaluated using radial basis functions. Examples of surface/volume adaption to point cloud data will be presented.

John P. Steinbrenner
Pointwise
john.steinbrenner@gmail.com

MS117

Batched Blas Moving Forward in MAGMA

The MAGMA library provides a wide range of GPU accelerated routines for batch linear algebra operations. Such routines are carefully optimized for small matrix operations, and provide a solid foundation for many applications, including sparse direct solvers, tensor contractions, machine learning, and others. This talk presents the latest developments in the MAGMA library and its impact on some real applications. The talk also highlights a forward-looking C++ API that provides significantly more flexibility than the existing C interfaces from both vendors and library developers.

Ahmad Abdelfattah
University of Tennessee, Knoxville
ahmad@icl.utk.edu

Azzam Haidar
NVIDIA
azzamhaidar@nvidia.com

Stanimire Tomov
University of Tennessee, Knoxville
tomov@icl.utk.edu

Jack J. Dongarra
University of Tennessee, Oak Ridge National Laboratory, USA
dongarra@icl.utk.edu

MS117

CUDA Linear Algebra and High Performance Tensor Primitives

With tensor formats appearing across deep learning, general tensor operations required by quantum computing simulation, and various scientific computing applications, basic tensor operations have been supported via batched linear algebra operations and convolution. General tensor contractions, on the other hand, have been left unsupported. We propose two batched linear algebra primitives and discuss their implementation in CUDA/CUTLASS to provide speed-of-light kernels that will allow nearly all tensor contractions to be performed at peak performance on GPUs.

Cris Cecka
NVIDIA
Stanford University
ccecka@nvidia.com

MS117

Standardization of the Batched BLAS

In this talk we shall motivate and describe the proposal for the Batched Basic Linear Algebra Subprograms (BBLAS). The BBLAS independently perform a large number of a specific BLAS operation, such as matrix multiplication, on small matrices. As with the existing BLAS, we have given a standard specification for the BBLAS so that code which requires the BBLAS can be portable, but at the same time utilise efficient versions of the BBLAS produced by vendors and other developers.

Sven J. Hammarling
The University of Manchester
sven.hammarling@btinternet.com

MS117

Toward Fast Eigensolvers for Electronic Structure Calculations using Low-rank Approximations

Leveraging low-rank matrix approximations/computations on manycore architectures represents perhaps one of the most challenging open research problems for the computational linear algebra community. The challenge comes from the mismatch between the low arithmetic intensity of the individual low-rank numerical kernels and the throughput-oriented massively parallel accelerators, such as GPUs. To reconcile both algorithmic and hardware trends, the idea consists in refactoring the original algorithm into a series of batched kernel calls, not only to reduce the kernel launch overhead, but also to increase the hardware occupancy. This talk highlights these necessary algorithmic adaptations toward developing fast eigensolvers for low-rank matrices arising from electronic structure calculations and presents preliminary results.

Kadir Akbudak
Bilkent University
kadir.cs@gmail.com

Ali M. Charara
KAUST
ali.charara@kaust.edu.sa

David E. Keyes
King Abdullah University of Science and Technology (KAUST)
david.keyes@kaust.edu.sa

Hatem Ltaief, Aleksandr Mikhalev
King Abdullah University of Science & Technology (KAUST)
hatem.Ltaief@kaust.edu.sa,
aleksandr.mikhalev@kaust.edu.sa

Dalal Sukkari
KAUST
dalal.sukkari@kaust.edu.sa

MS118

Low Mach Number Fluctuating Hydrodynamics for Electrolytes

At small scales where thermal fluctuations play an important role in the dynamics, the continuum Navier-Stokes equations do not provide an adequate description of fluid flow. Landau and Lifshitz proposed a modified version of the Navier-Stokes equations, referred to as the fluctuating hydrodynamcis equations that incorporates stochastic flux terms designed to incorporate the effect of fluctuations. These stochastic fluxes are constructed so that the fluctuating hydrodynamics equations are consistent with equilibrium fluctuations from statistical mechanics. Here we discuss the development of fluctuating hydrodynamics for electrolytes. The resulting model does not require that the solution be a dilute mixture of solutes in a distinguished solvent and allows for an arbitrary number of ionic species. We analyze fluctuating hydrodynamics for strong electrolyte mixtures to compute concentration corrections for chemical potential, diffusivity, and conductivity and show that these corrections are in agreement with the limiting laws of Debye, Hückel, and Onsager. We then introduce a finite volume method for solving the fluctuating hydrodynamics equations for electrolytes and present numerical simulations of several electrokinetic flows.

John B. Bell
CCSE
Lawrence Berkeley Laboratory
jbbell@lbl.gov

MS118

Fluctuating Hydrodynamics of Flow through Porous Membranes

The flow of liquids and gases through porous membranes has long been a topic of interest; wide ranging applications include, for example, water purification, drug delivery, and molecular sieves. In mesoscopic regimes, transport though these membranes may be significantly impacted by random thermal fluctuations. In this talk we discuss the use of fluctuating hydrodynamics to model such flows. Several methods for adding particle effusion to an existing compressible flow solver will be reviewed, including master equation, tau leaping, and Langevin equation approaches. These are illustrated using simple examples, and validated with comparison to DSMC results.

Daniel R. Ladiges
Center for Computational Science and Engineering
Lawrence Berkeley National Laboratory
drladiges@lbl.gov

MS118

Hybrid Discrete-continuous Dynamics Method for Chemically-powered Nanomotors

In this talk, we will present a new hybrid fluctuating hydrodynamics/kinetic Monte Carlo method for modeling heterogeneous catalysis at the nanoscale. The system is decomposed into the fluid domain and the catalyst-fluid interface. A simplified fluctuating hydrodynamics framework is used to model the diffusion of the chemical species in the fluid domain. The catalytic activity at the reactive interface is described by a chemical master. The method is validated through a simple one-dimensional (1D) linear model. Our approach is then applied to model spatially asymmetric catalysis on the surface of self-propelled nanoswimmers. Numerical simulations that estimate the uncertainty in the swimming velocity resulting from inherent stochastic nature of the chemical reactions at the catalytic interface will be also presented.

Nikolaos Voulgarakis
Washington State University
n.voulgarakis@wsu.edu

MS118

Nano Hydrodynamics Near Walls

We propose a generalization of Equilibrium Density Functional Theory (DFT) for simple fluids in motion in the presence of solid walls. By using the Theory of Coarse-Graining we obtain the structure of the equations that govern the dynamics of the fluid. [D. Camargo *et al.*, J. Chem. Phys. 148, 064107, 2018] These equations describe the irreversible forces that arise in the fluid because of the interaction with the walls. Here, the presence of the walls is reflected not through boundary conditions but rather in terms of reversible and irreversible surface forces that are highly localised near the solid walls. In the limit of macroscopic flows, one recovers the conventional Navier slip and impenetrability boundary conditions. The theory relies on a Markovian assumption that may not be valid at nanoscales. Therefore, we consider molecular dynamics simulations in confined and unconfined fluids in order to asses and validate this local in time assumption. We also discuss the plateau problem that emerges in the calculation of the non-local transport coefficients through the Green-Kubo formulae.

Jaime Arturo de la Torre
Physics Department, UNED
Madrid, Spain
jatorre@fisfun.uned.es

Diego Duque-Zumajo
UNED
d.duque@fisfun.uned.es

Diego Camargo
U. Nacional Colombia
dacamarg@unal.edu.co

Pep Español
Departamento de Física Fundamental
Universidad Nacional de Educación a Distancia, Madrid
pep@fisfun.uned.es

MS119

Dynamic Coupling of Full and Reduced Models via Randomized Online Basis Updates

Traditional model reduction approaches typically fail for convection-dominated problems because the manifolds induced by the solutions of the full models contain high-dimensional features. We exploit that manifolds are low dimensional in a local sense, so that they can be well-approximated locally with low-dimensional (reduced) spaces. We iteratively learn and adapt reduced spaces from randomly sampled data of the full models to locally approximate the solution manifolds. Furthermore, we discuss

deterministic sampling strategies for basis adaptation that take structure of the reduced spaces into account to guide the sampling. Numerical experiments to predict pressure waves in combustion dynamics demonstrate that our approach achieves significant speedups in contrast to classical, static reduced models, which can be even more expensive to evaluate than the full models.

Benjamin Peherstorfer
Courant Institute of Mathematical Sciences
New York University
pehersto@cims.nyu.edu

MS119

The Shifted Pod: A Data Based Model Reduction for Situations with Multiple Transports - Formulations and Application Aspects

Mode-based model-reduction is effective when the system dynamics or flow field can be described by a small number of spatial modes. For transport dominated phenomena often a large number of modes is needed to reproduce the full solution. Here, an approach is discussed, which decomposes a flow field into several co-moving frames, where each frame can be approximated by a low-rank field. The method of decomposition is formulated as an optimization problem. This allows to include boundary condition in a simple manner. We apply the method to the description of a pulse detonation combustion chamber, where the description is obtained from experimental data by a data assimilation technique.

Julius Reiss
TU Berlin
julius.reiss@tnt.tu-berlin.de

Philipp Schulze
Institute of Mathematics
TU Berlin
pschulze@math.tu-berlin.de

Volker Mehrmann
Technische Universität Berlin
mehrmann@math.tu-berlin.de

MS119

Model Reduction for Compressible Fluids: Application to Uncertainty Quantification

In this work, we focus on reduced order modeling (ROM) techniques for hyperbolic conservation laws with application in uncertainty quantification (UQ) and in conjunction with the well-known Monte Carlo sampling method. Because we are interested in model order reduction (MOR) techniques for unsteady non-linear hyperbolic systems of conservation laws, which involve moving waves and discontinuities, we explore the parameter-time framework and in the same time we deal with nonlinearities using a POD-EIM-Greedy algorithm (M. Drohmann et al., *Reduced basis approximation for nonlinear parametrized evolution equations based on empirical operator interpolation*, SIAM SISC, 34 (2012), pp. A937–A969). We provide under some hypothesis an error indicator, which is also an error upper bound for the difference between the high fidelity solution and the reduced one.

Remi Abgrall
University of Zurich
remi.abgrall@math.uzh.ch

Roxana Crisovan
Institute of Mathematics
University of Zurich
roxana.crisovan@math.uzh.ch

Davide Torlo
Universität Zürich
Switzerland
davide.torlo@math.uzh.ch

Svetlana Tokareva
Los Alamos National Laboratory
Applied Mathematics and Plasma Physics Group (T-5)
tokareva@lanl.gov

MS119

Reduced Basis Methods for Wave and Schrdinger Equations

It is well-known (see Ohlberger and Rave) that the decay of the Kolmogorov N-width for transport equations is only $N^{-1/2}$, which means that any model reduction technique using linear spaces has fairly limited speedup. The same holds true for other kind of problems such as hyperbolic, HJB and Schrödinger equations, just to name a few. Moreover, there are severe stability issues when it comes to solving the reduced system. In this talk, we report on recent developments concerning reduced basis methods for such kind or problems.

Karsten Urban
Institute of Numerical Mathematics, University of Ulm
karsten.urban@uni-ulm.de

MS120

Hybrid Methods for Radiation Transport

Recently, hybrid angular discretization methods have been used to improve the efficiency of numerical solutions to radiation transport problems. Although effective at improving the time to solution, these methods typically require more memory than a standard discretization strategy. In this talk, we show that additional hybridization in the spatial variable can help to reduce the memory load while maintaining accuracy and important asymptotic limits.

Cory Hauck
Oak Ridge National Laboratory
hauckc@ornl.gov

MS120

High-order Convex Limiting for Conservation Equations

I will present a continuous finite element technique to approximate hyperbolic systems that is high-order in space and guaranteed to be invariant domain preserving. The theory is illustrated with tests using third and fourth order approximation in time and space.

Ignacio Thomas
Sandia National Laboratories
itomas@tamu.edu

MS120

Entropy Stable and Well-balanced Discontinuous

Galerkin Methods for the Shallow Water Equations

In this presentation, we present our recent work on discontinuous Galerkin methods for the shallow water equations which are entropy stable, well-balanced and positivity preserving. The entropy stable property is achieved with an arbitrary quadrature rule for the numerical integration. The proposed methods are also well-balanced for the steady state and can preserve the non-negativity of water height numerically. Some 1D and 2D numerical tests are provided to demonstrate their performance.

Yulong Xing
The Ohio State University
xing.205@osu.edu

MS120

Fourth Order Accuracy and Discrete Maximum Principle of the Finite Difference Implementation of C^0-Q^2 Finite Element Method for Elliptic Equations

To implement the finite element method for a variable coefficient problem, suitable quadrature must be used to approximate the integrals in the bilinear form. The most convenient choice of quadrature for the tensor product of quadratic polynomials is the 3 by 3 Gauss-Lobatto quadrature since these quadrature points form a uniform grid and can also represent the numerical solution. We can prove that this simple quadrature will not affect the superconvergence of finite element method at these quadrature points thus such a scheme is a fourth order accurate finite difference scheme for elliptic equations. Moreover, discrete maximum principle can be proven for a variable coefficient Poisson equation. This is the first high order accurate scheme that can be proven to satisfy the discrete maximum principle for a variable coefficient problem.

Xiangxiong Zhang
Purdue University
Department of Mathematics
zhan1966@purdue.edu

MS121

Parallel Graph Algorithms to Remove Degenerate Features from Ice Sheet Meshes, and Determine Biconnectivity

As an ice-sheet simulation progresses, certain structural features of the underlying ice-sheet mesh can develop. These features can be problematic for the solvers driving the simulation. Since these solvers update the mesh iteratively, it is necessary to check for the appearance of these features at each step of the simulation. To quickly check for their existence, we reformulate the problem on the original mesh to an equivalent one based on graph connectivity. This allows a parallel and efficient label propagation-based algorithm to detect and remove these problematic features. We also show that our algorithm can be generalized to label the different biconnected components of an arbitrary undirected graph. Our work includes serial and distributed-memory implementations of algorithms for both the specialized ice-sheet and general cases. Our method targets the Albany simulation framework within Trilinos, and we improve upon the per-iteration solution time of the original approach by orders-of-magnitude.

Ian Bogle
Rensselaer Polytechnic Institute
boglei@rpi.edu

Karen D. Devine
Sandia National Laboratories
kddevin@sandia.gov

Mauro Perego
CCR Sandia National Laboratories
mperego@sandia.gov

Siva Rajamanickam
Sandia National Laboratories
srajama@sandia.gov

George M. Slota
Rensselaer Polytechnic Institute
slotag@rpi.edu

MS121

Inferring Past Temperature Changes in Greenland using the Unscented Transform

Cosmogenic nuclide dating of moraines provides new constraints on the retreat history of the southwest Greenland ice sheet between 11.6 - 7.3 ka before present. This period overlaps with the Holocene thermal maximum (HTM), a period of higher than modern temperatures similar to those projected to occur between 2090 and 2100. We capitalize on this new chronology of ice sheet retreat in order to infer past temperature on three paleo-flowlines in southwest Greenland. To model ice sheet extent, we use a flowline ice sheet model along with a positive degree day (PDD) model. Inferring past temperature is an inverse problem. Let $\boldsymbol{l}$ be a vector of glacier length observations through time and $\boldsymbol{t}$ be a vector of unknown temperatures at a fixed reference point through time. Given a multivariate Gaussian prior of the form $\boldsymbol{t} \sim \mathcal{N}(\boldsymbol{t_0}, P_0)$, selected to enforce temporal smoothness, the objective is to estimate the full posterior distribution $P(\boldsymbol{t}|\mathbf{l}) \propto P(\boldsymbol{l}|\boldsymbol{t})P(\boldsymbol{t})$. As an alternative to standard Markov chain Monte Carlo (MCMC) sampling methods, we approximate the posterior distribution as a multivariate Gaussian using the unscented transform (UT). Although the UT does not directly sample from the non-Gaussian posterior distribution, it provides a good approximation with low computational cost. In particular, the UT is trivially parallelizable and requires only 175 forward model runs per flowline.

Jake Downs
University of Montana
iostream333@gmail.com

Jesse Johnson
Dept of Computer Science
University of Montana
johnson@cs.umt.edu

Jason Briner
University of Buffalo
jbriner@buffalo.edu

NIcolas Young
Columbia University
nicolasy@ldeo.columbia.edu

Joshua Cuzzone
California Institute of Technology

joshua.k.cuzzone@jpl.nasa.gov

MS121

H-Div Conforming Methods for Geodynamic Stokes Flow

The Marker-and-Cell (MAC) discretization is a surprisingly efficient discretization for Stokes equations, converging at second order with a small number of degrees of freedom and sparse stencil. But that second order convergence can be explained by a cancellation of errors, which makes it fragile. Can we maintain second order convergence for MAC discretizations on deformed or non-conforming grids, or for creeping flows with anisotropic effective viscosities like those occuring in ice sheets? We develop an approach for a version of this problem on tetrahedral meshes, combining weakly-symmetric discretizations from the study of linear elasticity with multipoint flux methods to eliminate auxiliary fields that get introduced along the way, leaving a scheme with the same variables as a MAC scheme and robust second order convergence. Our approach also leads to systematic handling of "coarse-fine" interfaces in adaptively refined meshes.

Tobin Isaac
Georgia Tech
tisaac@cc.gatech.edu

MS121

Performance Portability in Albany Land Ice using Kokkos

High resolution simulations of polar ice sheets play a crucial role in the ongoing effort to develop more accurate and reliable Earth-system models for probabilistic sea-level projections. These simulations often require a massive amount of memory and computation from large, heterogeneous supercomputing clusters to provide sufficient accuracy and resolution. To avoid architecture specific programming, programming models such as Kokkos are becoming increasingly important to obtain efficient implementations. By heavily utilizing various performance portable packages within Trilinos, including the Kokkos programming model, a significant amount of progress has been made towards developing a performance portable implementation of Albany Land Ice (formerly known as Albany/FELIX). This presentation focuses on the key performance developments and future performance goals towards obtaining a robust and scalable land ice solver on a variety of different architectures including traditional CPUs, NVIDIA GPUs and Intel Xeon Phis.

Jerry Watkins, Irina K. Tezaur
Sandia National Laboratories
jwatkin@sandia.gov, ikalash@sandia.gov

MS122

High-order, Dispersionless, Spatio-Temporally Parallel "Fast-Hybrid" Wave Equation Solver at O(1) Sampling Cost

We propose and demonstrate a frequency/time hybrid integral-equation method for the time dependent wave equation in two and three-dimensional spatial domains. Relying on Fourier Transformation in time, the method utilizes a fixed (time-independent) number of frequency-domain integral-equation solutions to evaluate, with superalgebraically-small errors, time domain solutions for arbitrarily long times. The approach relies on two main elements, namely, 1) A smooth time-windowing methodology that enables accurate band-limited representations for arbitrary long time signals, and 2) A novel Fourier transform approach which, without causing spurious periodicity effects, delivers dispersionless spectrally accurate solutions. The algorithm can handle dispersive media, complex physical structures, it enables parallelization in time in a straightforward manner, and it allows for time leaping—that is, solution sampling at any given time T at $\mathcal{O}(1)$-bounded sampling cost, for arbitrarily large values of T, and without requirement of evaluation of the solution at intermediate times. The proposed frequency/time hybridization strategy, which generalizes to any linear partial differential equation in the time domain for which frequency-domain solutions can be obtained (including e.g. the time-domain Maxwell equations), provides significant advantages over other available alternatives such as volumetric discretization and convolution-quadrature approaches.

Thomas Anderson
Caltech
tanderson@caltech.edu

Oscar P. Bruno
California Institute of Technology
obruno@caltech.edu

Mark Lyon
University of New Hampshire
mark.lyon@unh.edu

MS122

Higher-Order Implicit Methods for the Navier-Stokes Equations and Other Linear and Nonlinear Waves

Many challenging CFD problems require very fine meshes to reproduce flow structures such as boundary layers, eddies, vortices, and turbulence. The celebrated Beam and Warming method–an alternating direction implicit (ADI) algorithm–provides one of the most attractive ways to overcome the strict stability constraints of explicit solvers. However, the Beam and Warming method has never been shown to have greater than first-order temporal accuracy. In this talk, we present a new class of ADI solvers with orders of accuracy between two and six, even in the presence of general boundary conditions. In fact, we show for the first time high-order time convergence curves for any ADI-based Navier-Stokes solver. Furthermore, the solvers are quasi-unconditionally stable in practice–i.e., they do not suffer from CFL stability constraints for adequately resolved flows. Theoretical results place on a solid basis the observed quasi-unconditional stability. Applications of the method to other linear and nonlinear wave equations are also presented.

Max Cubillos
Wright-Patterson Air Force Labs.
max.cubillos-moraga.ctr@us.af.mil

MS122

Approximating Eigenvalue Clusters using FEAST and Application to Beam Propagation

When using coupled mode approximations to simulate propagation of light through optical fibers, one needs to compute certain propagation constants which solve an eigenvalue problem. We apply the FEAST algorithm,

a popular eigenvalue solver in numerical linear algebra, known for its multiple layers of parallelism. The algorithm targets a cluster of eigenvalues enclosed within a contour. It can be viewed as a subspace iteration applied to an operator-valued contour integral. When used to approximate a part of the spectrum of an unbounded differential operator, it is typical to discretize the operator. We introduce an abstract theoretical framework to study the influence of such discretization errors in the final spectral approximations.

Jay Gopalikrishnan
Portland State University
gjay@pdx.edu

MS122

High-order Upwind Schemes for the Dispersive Maxwell Equations

An efficient and high-order accurate time-domain (FDTD) scheme for solving Maxwell's equations in second-order form with a generalized dispersive material (GDM) model is described. The equations are discretized using finite-differences, and advanced in time with a single-stage, three-level, space-time scheme which remains stable up to the usual explicit CFL-one restriction, as shown using mode analysis. Composite overlapping grids are used to treat complex geometries with boundary-conforming grids, and a high-order upwind dissipation is added to ensure robust and stable approximations on overlapping grids. Properties of the GDM model, which can represent arbitrary dispersive behaviours, are analysed. Numerical results for some newly developed exact solutions confirm the accuracy and stability of the new scheme.

William D. Henshaw
Rensselaer Polytechnic Institute
henshw@rpi.edu

MS123

Fractional Generalizations of Gradient Elasticity and Higher Order Diffusion

This contribution concerns the fractional and fractal generalizations of gradient elasticity and higher-order diffusion. Both of these theories were introduced by the second author three decades ago to model deformation and transport problems in media with micro/nanostructures. A new Laplacian multiplied by a corresponding internal length term was added in the standard constitutive equations of Hookean elasticity and Fickean diffusion to interpret experimental data that could not be modeled by classical theories. Among these were the elimination of undesirable elastic singularities (dislocation lines, crack tips), as well as diffusion processes in nanopolycrystals. The resulting internal length gradient (ILG) framework is extended herein to account for fractional and fractal considerations (fractional/fractal ILG framework). Fundamental solutions for static fractional/fractal gradient elasticity and steady-state higher-order diffusion are derived and certain example problems are discussed.

Elias Aifantis, Konstantinos Parisis
Aristotle University of Thessaloniki
mom@mom.gen.auth.gr, konspari@gmail.com

MS123

Modeling Diffusion Signal Decay of Magnetic Resonance Images

Magnetic resonance imaging (MRI) is one of the most useful diagnostic imaging technologies for identifying tissue damage with high contrast and resolution. Diffusion-weighted MRI characterizes the movement of water molecules in tissue by displaying the reduction in signal intensity following the application of diffusion-sensitizing gradients. Since the mobility of water is obstructed(hindered and restricted)by tissue components and compartments, the decay of the diffusion signal is characteristic of cellular and sub-cellular structures. Interpretation of this signal requires mathematical models whose parameters reflect key features(e.g., porosity, tortuosity, surface to volume ratio, permeability)of tissue composition. Current practice uses the apparent diffusion coefficient (ADC, mm2/sec) derived from a simple exponential model of hindered diffusion to describe signal decay. New models based on simulations, stochastic processes, tissue structure, or physical and physiological constraints are needed. In this study, we describe a generalization of the diffusion-weighted signal decay that captures tissue complexity in terms decay functions for different tissue features. We explore a subset of these models to illustrate their characteristics both theoretically and experimentally by considering applications in known heterogeneous environments such as Sephadex gels, and in complex biological tissue such as healthy gray and white brain matter,and brain tumors.

Richard L. Magin
University of Illinois at Chicago
rmagin@uic.edu

MS123

A Physical Interpretation of Fractional Viscoelasticity Based on the Fractal Structure of Media

The use of fractional order calculus operations to describe the rate-dependent viscoelastic response yielded a model with self-consistent model parameters. Therefore, the model performs better at predictions across a broad range of experimental rates. This led us to discover physical relations that connect the fractional time derivative applied to viscoelasticity with thermal diffusion and fractal structure in fractal media. These new relations, which date back to the work of Bagley and Torvik and Mandelbrot, show connections between fractional viscoelasticity and the fractal structure of materials. The derivation is based on the generalized molecular theory of Rouse and Zimm where generalized Gaussian structures (GGSs) are used to replace the Rouse matrix with the generalized Gaussian Rouse matrix (GRM). To validate the theory, we collect viscoelastic measurements, infrared imaging, and microscopy images of fractal structure. Bayesian uncertainty quantification is used to identify self-consistent parameters based on fractional properties and fractal structure on two different types of elastomers.

Somayeh Mashayekhi
Florida State University
smashayekhi@fsu.edu

M. Yousuff Hussaini
Department of Mathematics
Florida State University
yhussain@fsu.edu

William Oates
Department of Mechanical Engineering, Florida State

Universi
woates@eng.fsu.edu

MS123

Fractional Viscoelasticity of Auxetic Foams

Auxetic foams are useful in applications where high energy absorption and fracture resistance are desired making them excellent candidates for packing material, sports safety equipment, and shoe soles, for example. Since these materials are relatively new, their viscoelastic behavior has not been extensively studied and modeled. A knowledge of the viscoelastic behavior over a broad range of deformation rates is critical when considering the mechanical properties of the material. Previous research has demonstrated success in applying fractional order linear viscoelastic constitutive model to characterize the behavior of Very High Bond (VHB) 4949. In the present study, a similar model is applied while considering changes due to the structure of the material. Predictions across several orders of magnitude in deformation rates are validated against data using a single set of model parameters. The influence of the Poisson ratio is also examined through mathematical modeling and the use of uncertainty methods. All results are statistically validated using Bayesian uncertainty methods to obtain posterior densities for the fractional order parameters.

Eugenia Stanisauskis
Florida State University
Materials Science and Engineering
estanisauskis@gmail.com

Paul Miles
North Carolina State University
prmiles@ncsu.edu

William Oates
Department of Mechanical Engineering, Florida State Universi
woates@eng.fsu.edu

MS124

Electromagnetics Simulations with Vector-valued Finite Elements in MOOSE

Vector-valued finite elements have become an essential part of modern electromagnetic (EM) numerical simulation. This talk will describe the development and initial validation of a MOOSE-based EM solver application, ELK (**E**lectromagnetics **L**ibrary for **K**inetics and fluids), that utilizes the recently added vector-valued finite element system. The primary motivation for the development of ELK was to study EM wave propagation in radio-frequency (rf) driven systems, such as those encountered in the study of antenna structures, waveguides, and plasma physics. In simulating these scenarios with traditional scalar-valued finite elements, boundary conditions on tangential field components can sometimes be difficult to define, and field discontinuities encountered with certain geometric or structural features (e.g., abrupt corners and inhomogeneous materials) can often lead to non-physical results. Vector-valued elements can largely mitigate and even eliminate these difficulties. ELK verification and validation efforts thus far have focused on standard wave propagation benchmarks for antennas and waveguides, and the impact of vector-valued elements on the implementation and results of these simulations will be described. Progress using ELK for plasma physics simulation, where inhomogeneous simulation domains are commonplace, will also be discussed.

Casey T. Icenhour
North Carolina State University
cticenho@ncsu.edu

Alex Lindsay, Richard Martineau
Idaho National Laboratory
alexander.lindsay@inl.gov, richard.martineau@inl.gov

Steven Shannon
North Carolina State University
scshanno@ncsu.edu

MS124

Porous Media Thermal Hydraulics Simulations of Pebble Bed Nuclear Reactors using Pronghorn

Pebble bed nuclear reactors (PBRs) generate thermal power for electricity production by sustaining a fission chain reaction in a core containing hundreds of thousands of spherical solid uranium and graphite fuel elements known as “pebbles.’ Fission heat is removed by coolant flowing between pebbles and transferred to a secondary system. A challenge in thermal/hydraulic (T/H) modeling is capturing the effects of the fuel-coolant structure without requiring extremely detailed meshes and large computing resources. One medium-fidelity modeling approach, known as a porous media model, involves averaging the fluid flow and heat transfer equations in space and providing closure relationships for interphase momentum and energy transfer. Such models can be used for scoping analyses and to inform higher-fidelity exploratory studies with relatively low computing requirements. This talk presents a new porous media simulation code, Pronghorn, that is built on the Multiphysics Object-Oriented Simulation Environment (MOOSE). Pronghorn is the medium-fidelity simulation tool in a multi-scale T/H system based on the MOOSE framework. The governing equations, finite element discretization, and stabilization schemes used are reviewed. Example applications to both gas- and liquid-cooled PBRs and a non-nuclear natural circulation facility are presented. The advantages of contributing new software to the MOOSE framework are highlighted through examples encountered during Pronghorn development.

April Novak
University of California, Berkeley
novak@berkeley.edu

John Peterson, Ling Zou
Idaho National Laboratory
jw.peterson@inl.gov, ling.zou@inl.gov

Rachel Slaybaugh
UC Berkeley
Dept. of Nuclear Engineering
slaybaugh@berkeley.edu

Richard Martineau
Idaho National Laboratory
richard.martineau@inl.gov

MS124

An Overview of the Multiphysics Simulation Capa-

bilities of MOOSE

The Multiphysics Object Oriented Simulation Environment (MOOSE) framework enables complex multiphysics analysis tools to be built rapidly by scientists, engineers, and domain specialists. The full potential of MOOSE architecture is demonstrated by the breadth and depth of the included physics modules. From fully-coupled thermo-mechanical phase-field to stochastic defect analysis, MOOSE can be extended for solving a wide variety of physics and almost any field. An overview of the framework and current physics modules is presented along with future directions for the platform.

Cody J. Permann
Advanced Modeling and Simulation
Idaho National Laboratory
cody.permann@inl.gov

MS124

Moose Multiphysics Simulations Employing Physics-based, Transient Material Interfaces

A modern approach to modeling physical processes with large changes in field quantities over small scales compared to the size of the domain is the extended finite element method (XFEM). This process enriches the polynomial approximation space of the traditional finite element method (FEM) such that non-smooth solutions may be modeled independently from the original mesh without extensive refinement. XFEM is employed in contemporary solid mechanics to simulate cracking, voids, and dislocations, but we aim to demonstrate its efficacy in modeling processes involving transient material interfaces driven by coupled physics solutions. Oxidation is an example of a specific application of interest for this work. We will explore the current XFEM implementation in the Moose multiphysics FEM framework as well as its performance in heat transfer problems of increasing complexity constructed using the Method of Manufactured Solutions (MMS) building up to a coupled oxidation simulation.

James B. Tompkins
Texas A&M University
tompkins@tamu.edu

MS125

Acceleration of Prediction of Chemical Shift Structures

Experimental chemical shifts from solution and solid state magic-angle-spinning nuclear magnetic resonance spectra provide atomic level data for each amino acid within a protein or complex. However, structure determination of large complexes and assemblies based on NMR data alone remains challenging due to the complexity of the calculations. This talk will describe key computational challenges while presenting a hardware accelerated strategy that will expedite the prediction of estimation of NMR chemical-shifts of large macromolecular complexes by manyfold. Results will demonstrate scalability and feasibility of our approach in systems of increasing complexity ranging from 2,000 to 11,000,000 atoms. A performance yet portable solution demonstrated on heterogeneous computing systems will be presented.

Sunita Chandrasekaran, Juan Perilla
University of Delaware
schandra@udel.edu, jperilla@udel.edu

MS125

Challenges for Analysis and Visualization of Atomic-detail Simulations of Minimal Cells

The combination of experimental structure information, molecular dynamics (MD) simulation, interactive and parallel analysis, and high-fidelity visualization techniques with HPC computing resources creates a powerful "computational microscope" that permits molecular scientists to view the structure and dynamics of biomolecular complexes in atomic detail within realistic cellular environments. Pre-exascale supercomputers create new opportunities for the study of the structure and function of large biomolecular complexes such as viruses, and photosynthetic organelles, and the cellular envelope of engineered minimal cells permitting all-atom molecular dynamics simulations of hundreds of millions to billions of atoms. Such simulations present many challenges for simulation preparation, analysis, and visualization, producing terabytes of data that is impractical to transfer to remote facilities. It is therefore necessary to perform visualization and analysis tasks in-situ as the data are generated, by running interactive analysis and visualization tasks in remote sessions, and through batch-mode post-hoc analyses co-located with direct access to high performance storage systems. This talk will describe some of the key challenges arising for state-of-the-art MD simulations and methodologies, and how scientific tools such as NAMD and VMD address them using state-of-the-art algorithms and heterogeneous computing hardware platforms.

John E. Stone
Beckman Institute, U. Illinois at Urbana-Champaign
johns@ks.uiuc.edu

MS125

In Situ Data Analytics for Next Generation Molecular Dynamics Workflows

Molecular dynamics (MD) simulations studying the classical time evolution of a molecular system at atomic resolution are widely recognized in the fields of chemistry, material sciences, molecular biology and drug design; these simulations are one of the most common simulations on supercomputers. Next-generation supercomputers will have dramatically higher performance than do current systems, generating more data that needs to be analyzed (i.e., in terms of number and length of molecular dynamics trajectories). The coordination of data generation and analysis cannot rely on manual, centralized approaches as it is predominately done today. In this talk I will discuss how the combination of machine learning and data analytics approaches, workflow management methods, and high performance computing techniques can transition the runtime analysis of larger and larger MD trajectories towards the exascale era.

Michela Taufer
University of Tennessee
taufer@utk.edu

MS125

Capabilities, Collaboration and Cancer: Co-design for Advanced Computing Solutions for Cancer

New computational opportunities have emerged within the

cancer research and clinical application areas as the size, number, variety and complexity of cancer datasets have grown in recent years. One example, the Joint Design of Advanced Computing Solutions for Cancer, a collaboration involving the US Department of Energy, the National Cancer Institute, and several national laboratories, is piloting new computational approaches in large-scale computing as insights are sought into cancer research challenges in the molecular, cellular, and population domains. Additional examples are found in development of data management infrastructures, and public-private efforts working to significantly accelerate drug discovery and development. The presentation will provide an overview of these efforts, including the critical role of co-design driven by scientific challenges plays in developing advances in computing.

Eric Stahlberg
SAIC-Frederick/National Cancer Institute
OpenFPGA Inc.
estahlberg@gmail.com

George Zaki
Frederick National Lab. for Cancer Research
george.zaki@nih.gov

MS126

Strategies for Handling Uncertainties in Model Parameters and Measurements for Cardiovascular Applications

We will demonstrate methods for uncertainty quantification of cardiovascular model predictions and their sensitivity respect to uncertain model parameters. In particular, we will explore how spectral expansion and Monte Carlo strategies may be applied for various cardiovascular models to guide model parameter prioritization/fixation in an optimization process. In our presentation we will demonstrate how UQ and SA may be used for clinically relevant cardiovascular models, e.g. model predictions of arterial pressure and flow wave dynamics, model-based estimation of the fractional flow reserve (FFR) and model-based estimation of the total arterial compliance. Moreover, we will explore the performance of several strategies for parameter estimation models accounting for uncertainties in model parameters and measurements.

Leif Hellevik
Dept of Structural Engineering
Norwegian University of Science and Technology
leif.r.hellevik@ntnu.no

Lucas Omar Mueller
Norwegian University of Science and Technology l
lucas omar muller ¡lucas.o.muller@ntnu.no¿

MS126

Sensitivity Analysis, Model Reduction and Parameters Estimation for Thrombosis Modeling for CFD

Biomaterial surfaces used in medical devices initiate blood coagulation that may lead to device malfunction or thromboembolism. A numerical strategy is proposed where a boundary condition relates the species wall flux to the surface reaction rate of factor XII activation. This contrasts with conventional in-vivo models in which coagulation is initialized at user-defined injury sites where tissue factor is present. The approach makes it possible to predict the coagulation initiated by the contact activation system but it relies on the coagulation scheme proposed by Chatterjee et al. (2010) which involves as many as 63 reaction rates and is thus hardly tractable for the analysis of complex geometries. A reduced order scheme for the thrombin generation triggered by the contact activation system is thus proposed, based on the outcome of the Morris sensitivity analysis performed to identify the most sensitive parameters involved in the original scheme. The best possible values of the 9 remaining reaction rates are obtained by Bayesian inference, using data from thrombin generation tests as a reference. The resulting optimized reduced coagulation scheme is then coupled to a platelet-based thrombus growth model (Taylor et al., 2016) and implemented in the in-house YALES2BIO solver (www.math.univ-montp2.fr/˜yales2bio/); the results obtained for the backward-facing step experiment of Taylor et al. (2014) compare favorably with the measurements.

Rodrigo Mendez, Simon Mendez
IMAG, Univ Montpellier, CNRS, Montpellier, France
rodrigo.mendez-rojano@umontpellier.fr,
simon.mendez@umontpellier.fr

Franck Nicoud
CNRS - UMR5149
University of Montpellier
franck.nicoud@umontpellier.fr

Didier Lucor
Sorbonne Univ, UPMC Univ Paris 06
CNRS, UMR 7190, Inst Jean Le Rond d'Alembert
didier.lucor@upmc.fr

MS126

Data Assimilation Methods for Defective Boundary Data Problems in Computational Hemodynamics

One challenging aspect of numerical solving blood flow problems in clinical settings is that boundary data critical for the reliability of the results - are almost invariably missing. "Defective data" problems need to be solved with a combination of literature data, patient-specific measures, and empirical laws. The rigorous background provided by Data Assimilation techniques needs to be adapted to clinical timelines. We will present some recent methods inspired by variational principles used for aortic diseases.

Alessandro Veneziani
MathCS, Emory University, Atlanta, GA
ale@mathcs.emory.edu

Rodrigo M. Romarowski, Simone Morganti
University of Pavia
rodrigo.romarowski@unipv.it, simone.morganti@unipv.it

Ferdinando Auricchio
University of Pavia, Italy
auricchi@unipv.it

Adrien Lefieux
Emory University
adrien.lefieux@emory.edu

MS126

Improving Convergence of Iterative Ensemble Kalman Filter by Preventing Early Stopping via

a Resampling Approach

Iterative ensemble Kalman filter is widely used in nonlinear inverse problems in different areas. However, its convergence still remains to be an open problem. In this paper, we conducted a comprehensive convergence analysis and proposed a resampling method to improve its convergence. First, the need of iterations for nonlinear inverse problems is established by formulating it as a constrained optimization. Secondly, the standard iteration approach is studied and the shrinking effect of Kalman updates on the ensemble covariances is revealed. Thirdly, we showed that the interaction between the nonlinearity of the forward model and the covariance shrinking effect causes early stopping of the iterative ensemble updates, i.e. the Kalman gain converges to zero before the innovation is minimized. Furthermore, the steady state behavior of the early stopping phenomenon and its relation to observation uncertainty are investigated. After that, we proposed an approach to prevent the early stopping of the iterative ensemble Kalman filter by perturbing the covariance shrinkage with hidden parameter ensemble resampling. The ensemble mean and covariance are kept unchanged during the resampling process, which guarantees the change of the Kalman gain at each iteration caused by resampling remains small to yield correct Kalman update directions. Synthetic model tests are presented to demonstrate the early stopping effect and merits of the proposed resampling scheme.

Jian-Xun Wang
University of Notre Dame
jwang33@nd.edu

Jiacheng Wu
UC Berkeley
wujiacheng@berkeley.edu

Shawn Shadden
University of California, Berkeley
shadden@berkeley.edu

MS127

The Dividends of Investing in Computational Software Design

Growth through accretion is among the biggest obstacles to long term viability of a software project. In contrast, an upfront investment in extensible framework design gives returns in a long run that are several times the original investment. Software remains maintainable and grows manageably. Well designed codes have the potential to be adopted by more than one research community, thus leading to reuse of software artifacts. In this presentations I will present general benefits of software design obtained by a few community codes, with a specific example from FLASH, a multiphysics software for solving PDEs in several scientific domains.

Anshu Dubey
Argonne National Laboratory
adubey@anl.gov

MS127

Firedrake: High Productivity High Performance Simulation

A critical challenge in the numerical solution of PDEs is that the mathematical route from equations to simulation comprises many steps, each of which demands a high degree of sophistication. Choice of equations and discretisation, implementation of local operators, global assembly and parallelisation, low-level code optimisation, solvers, preconditioners: each of these is a domain of expertise in its own right on which theses and books have been written. However the simulation developer necessarily has limited resources and can only be truly expert in some parts of this chain. Firedrake addresses this challenge by modelling the finite element simulation creation process as a series of mathematical representations, and software which maps between them, automatically transforming the simulation problem from higher to lower mathematical abstractions until high performance compiled code is generated and executed. This makes every stage of the simulation process composable: users can benefit from advanced algorithms at every layer, and can change their choices without reimplementing. By tightly integrating with PETSc's composable solver interface, sophisticated nested preconditioners which create and solve the appropriate auxiliary differential operators on the fly can easily be created. This composable infrastructure yields a step-change in productivity for users, who become able to create sophisticated optimised PDE solvers much more efficiently than is otherwise possible.

David Ham
Department of Mathematics
Imperial College London
david.ham@imperial.ac.uk

Lawrence Mitchell
Durham University
lawrence.mitchell@durham.ac.uk

Thomas H. Gibson
Department of Mathematics
Imperial College London
t.gibson15@imperial.ac.uk

Tianjiao Sun
Imperial College London
tianjiao.sun14@imperial.ac.uk

Miklós Homolya
Department of Computing
Imperial College London
miklos.homolya14@imperial.ac.uk

Rob C. Kirby
Baylor University
robert_kirby@baylor.edu

Paul Kelly
Imperial College London
p.kelly@imperial.ac.uk

MS127

Modularity of Lowlevel Forest-of-octree Libraries

The forest of octrees meshing paradigm has been established primarily using cubic elements. Recently, however, we have seen scalable adaptive mesh refinement codes that use triangles, tetrahedra, and prisms. While it seems that different shapes of elements require different algorithms for meshing, this is not necessarily the case. In this talk, we propose a modular approach that separates the shape-specific, per-element logic from the high-level mesh modification algorithms. This way, the latter can be coded independently of the elements' shape and other properties. We reduce redundancy and allow for future extensions of

even more shapes and element orderings. We present examples of our approach using the p4est and t8code software libraries and selected numerical applications.

Johannes Holke
University of Bonn
johannes.holke@dlr.de

Carsten Burstedde
Universität Bonn
burstedde@ins.uni-bonn.de

MS127

Exahype - An Exascale Engine for Solving Hyperbolic PDEs

The ExaHyPE consortium develops an engine to write solvers for hyperbolic differential equation system solvers. Two major demonstrators built upon this engine are a solver for earthquakes and one solver for gravitational waves. ExaHyPE's core is designed to perform on large computers through an innovative combination of dynamically adaptive meshes, aggressively optimised compute kernels and the ADER-DG method. While this triad seems to be brilliant fit to what upcoming exascale computers might look like, prototyping brings us back down to earth. We observe naive implementations to suffer from complex concurrency dependencies with strong synchronisation, phases with memory access bursts which are bandwidth-bound, a lack of time to squeeze MPI messages through the system while cores are computing, and so forth. Data movements become a limiting factor. In this talk, we propose a novel *fused implementation* of DG methods which relies on optimistic time stepping. We further combine it with a novel mesh colouring scheme that we call *enclave tasking*. We demonstrate the benefit of these techniques by means of some benchmarks, and we classify them in the context of communication-avoiding techniques.

Tobias Weinzierl
Durham University
tobias.weinzierl@durham.ac.uk

MS128

Bernstein-Bezier Techniques in High Order Finite Element Analysis

High order finite element methods have been analysed extensively for a wide variety of applications and are known to be capable of producing xponential rates of convergence, even for challenging problems with singularities, sharp boundary layers and high frequency oscillations. High order polynomial approximations are commonplace in many areas of scientific computing including computer graphics, computer aided-geometric design, and spectral methods for PDEs. Whilst the spectral method is routinely used with approximation orders in the 100s or even 1000s. Yet, despite theory giving the nod to the use of very high order finite element methods, the range of polynomial degree used in practical finite element computations is rarely larger than *eighth* order! Few commercial codes allow the use of high order finite elements. The rather modest polynomial degrees seen in high order finite element analysis are due to efficiency considerations rather than any theoretical barriers. Bernstein-Bézier polynomials have a number of interesting properties that have led to their being the industry standard for visualisation and CAGD. We explore the use of Bernstein polynomials as a basis for finite element approximation.

Mark Ainsworth
Division of Applied Mathematics
Brown University
Mark_Ainsworth@brown.edu

Shuai Jiang
Division of Applied Mathematics, Brown Univ
shuai_jiang@brown.edu

MS128

A Performance Portable Implementation of the Spectral Elements Atmosphere Dynamical Core of E3SM

This work converts the atmospheric dynamical core (HOMME) of the Energy Exascale Earth System Model (E3SM) from the current CPU-centric implementation, in Fortran 90, to a new performance-portable implementation, in C++ with the Kokkos performance-portability framework. HOMME simulates the dynamics and physical processes of the atmosphere. It is the most computationally demanding part of E3SM. Kokkos provides performance-portable multidimensional arrays and intraprocess parallel execution constructs. These form an abstraction layer over the hardware architecture of a compute node within a supercomputer. We will present results for the performance of our implementation on conventional CPU, Intel Xeon Phi, and Nvidia GPU; compare performance with the original Fortran on CPU and Xeon Phi; and discuss details of the implementation.

Luca Bertagna
Florida State University
lbertag@sandia.gov

Michael Deakin
Sandia National Laboratory
mfdeakin@cs.unc.edu

Oksana Guba
Sandia National Laboratories
NM, USA.
onguba@sandia.gov

Daniel Sunderland, Andrew Bradley, Irina K. Tezaur, Mark A. Taylor
Sandia National Laboratories
dsunder@sandia.gov, ambradl@sandia.gov, ikalash@sandia.gov, mataylo@sandia.gov

Andrew Salinger
CSRI
Sandia National Labs
agsalin@sandia.gov

MS128

Efficient Implementation of a High Order Control Volume Finite Element Scheme for Low-mach Flow

High-fidelity wind energy applications require the low-dissipation, low-dispersion numerical methods in order to accurately propagate wake structures over the relatively large time scales where these structures are persistent. Because of this, and the favorable ratio of computational work to datum transfer relative to low-order methods, high-order methods are attractive for wind modeling applications at

exascale. This talk will discuss work on re-implementing a high-order "control volume finite element" scheme into the Nalu codebase that is higher-performing. The high-order CVFEM method uses nodal finite element basis functions for approximation of primitive variables, but the weak form results from applying the divergence theorem on control volumes around high-order solution points, as in finite volume. In this talk, the CVFEM method is formulated to allow a faster, tensor-contraction based residual evaluation using a numerical quadrature set specialized for the scheme. Results of the improved scheme on a turbulent, temporally evolving shear layer and an atmospheric boundary layer are discussed and compared with a second-order, edge-based finite volume scheme.

Robert Knaus
Sandia National Laboratory
rcknaus@sandia.gov

MS128

A Discontinuous Galerkin Spectral Element Method Compressible Flow Solver

We present a discontinuous Galerkin spectral element method compressible flow solver for the three-dimensional Euler and Navier-Stokes equations. The truly high-order nature of the current solver is retained from the spectral element method framework, which ensures fast tensor-product form operator evaluation and rapid convergence rates. The discontinuous Galerkin formulation has the added benefits of allowing for discontinuous quantities, upwinding numerical fluxes, and small communication requirements. For improved robustness, various stabilization techniques have been implemented: de-aliasing is used to treat the advection terms, low-pass filtering removes spurious high wavenumber oscillations, and a non-linear filter-based artificial viscosity method is employed for shock-capturing. To aid in temporal stability, we use a fully-implicit, Jacobian-free Newton-Krylov time-marching scheme. We developed efficient and scalable parallel implementations in the current solver for applications on modern high-performance computers.

Li Lu, Paul Fischer
University of Illinois at Urbana-Champaign
lilu3@illinois.edu, fischerp@illinois.edu

MS129

Raptor: Parallel Algebraic Multigrid with Node-aware Communication

In this talk we present RAPtor, a parallel algebraic multigrid solver that uses node-aware communication to improve performance and scalability of both the setup and solve phases. We investigate multiple methods of node-aware communication and analyze the trade-offs among these communication strategies. Furthermore, we investigate a variety of AMG optimizations to be used in combination with node-aware communication, such as an efficient MIS-2 algorithm for aggregation and improved convergence from increasingly dense interpolation operators. Finally, we present direct comparisons between the Ruge-Stuben and smoothed aggregation solvers in RAPtor, as well as comparisons with other AMG codebases.

Amanda Bienz
University of Illinois at Urbana-Champaign
bienz2@illinois.edu

Luke Olson
University of Illinois Urbana-Champaign
lukeo@illinois.edu

William D. Gropp
University of Illinois at Urbana-Champaign
Dept of Computer Science
wgropp@illinois.edu

MS129

Library Interface Design and Performance Portability

Library developers seek to effectively serve the needs of present and future users while providing extension points for new developers to start contributing and quickly realize their impact. This involves providing powerful abstractions that are not too abstract to reason about, often decomposing familiar algorithms into constituent parts that can be recomposed to create new algorithms or variants with little or no software development effort, and providing diagnostics and debugging tools. When striving for performance, library developers must pay attention to interface granularity to enable vectorization, various forms of blocking for locality, and fusion. The flavor of a library is often defined by its choice of programming technology, from traditional fixed-ABI to template metaprogramming to just-in-time (JIT) compilation. We provide examples of interface tradeoffs by way of libCEED, a new performance portable algebraic discretization library that provides optimized finite element/spectral element infrastructure to a range of applications.

Jed Brown, Jeremy Thompson, Valeria Barra
University of Colorado Boulder
jed.brown@colorado.edu, jeremy.thompson@colorado.edu, valeria.barra@colorado.edu

MS129

BLIS: A Case Study in Performance and Portability

HPC and ML communities rely on dense linear algebra software as fundamental building blocks for higher-level applications. Over the years, these applications have come to expect BLAS and BLAS-like functionality that yields high levels of performance on modern CPU architectures. Implementing these key operations of interest, such as matrix multiplication, is now relatively well-understood. However, it is easy for a library project to obsessively prioritize performance above all else, sometimes to the detriment of adding or maintaining support for new hardware, compilers, and build environments. The BLAS-like Library Instantiation Software (BLIS) project was established with the simple goal of providing a portable framework for rapidly instantiating dense linear algebra libraries capable of yielding high performance with relatively little developer effort. Many lessons were learned before and during the development of BLIS, and most of these lessons directly or indirectly facilitate the software's ability to minimize architecture-specific code and parameters while simultaneously supporting competitive performance. In this talk, we walk through key features, characteristics, and design decisions that collectively form the foundation for BLIS's software architecture. We illustrate how these features enables portability to support new hardware, as well as extensions of the traditional BLAS functionality.

Devangi N. Parikh

University of Texas at Austin
devangi.parikh@utexas.edu

Field G. Van Zee
Department of Computer Sciences
The University of Texas at Austin
field@cs.utexas.edu

MS129

Performance Portability of Communication Optimal GMRES

Communication optimal GMRES is a version of GMRES in which we reduce the number of required synchronizations by replacing the standard Level-1 BLAS MGS GMRES algorithm with quadratic communication costs by alternative low synchronization orthogonalization strategies. The method is well-suited for modern parallel architectures, both in terms of fine-grain and course-grain parallelism. Applications that employ GMRES, such as flow simulations, often use a myriad of parallelization strategies, i.e., openMP, MPI and cuda, amongst others. The applications are commonly written to run in multiple computational environments. We implemented communication-optimal GMRES as a part of HYPRE library. We discuss the challenges associated with performance optimization on various computational platforms and scales. For application developers, the main issue is typically time-to-solution and we address this issue using Nalu wind turbine simulations.

Katarzyna Swirydowicz
National Renewable Energy Lab
Golden CO
katarzyna.swirydowicz@nrel.gov

Stephen Thomas
NREL
stephen.thomas@nrel.gov

Julien Langou
University of Colorado Denver
julien.langou@ucdenver.edu

Shreyas Ananthan
National Renewable Energy Laboratory
shreyas.ananthan@nrel.gov

Ulrike Meier Yang
LLNL
yang11@llnl.gov¿

MS130

Computationally Efficient Parameter Sensitivity Analysis for PDE-Constrained Optimization

Abstract not available.

Joseph L. Hart
North Carolina State University
jlhart3@ncsu.edu

Bart G. Van Bloemen Waanders
Sandia National Laboratories
bartv@sandia.gov

MS130

From Nonlinear Partial Differential Equations to Lowdimensional Polynomial Approximations

Abstract not available.

Elizabeth Qian, Karen E. Willcox
Massachusetts Institute of Technology
elizqian@mit.edu, kwillcox@ices.utexas.edu

MS130

Mass Conserving Hamiltonian Structure Preserving Reduced Order Model for Mimetic Ocean Models

Abstract not available.

Kenneth Sockwell
Florida State University
kcs12j@my.fsu.edu

Konstantin Pieper
Florida State University
Dept of Scientific Computing
kpieper@fsu.edu

Max Gunzburger
Florida State University
Department of Scientific Computing
mgunzburger@fsu.edu

MS130

Reduced Order Methods for Parametrized Optimal Flow Control: Applications to Coronary Artery Bypass Grafts Assimilated with Patient-Specific Data

Abstract not available.

Zakia Zainib
SISSA International School for Advanced Studies
zzainib@sissa.it

MS131

Numerical Homogenization for Waves using Hermite Methods

We discuss challenges and advantages in using Hermite methods together with various numerical homogenization techniques such as HMM and LOD. In particular we will consider homogenization of so-called micro-polar elasticity in Cosserat materials that possess orientation as well as rotation at the microlevel.

Daniel Appelo
Department of Applied Mathematics
University of Colorado Boulder
daniel.appelo@colorado.edu

Thomas M. Hagstrom
Southern Methodist University
Department of Mathematics
thagstrom@smu.edu

MS131

Sparse Data-driven Reduced Order Models and Inverse Scattering

Geophysical seismic exploration, as well as radar and sonar imaging, require the solution of large scale inverse prob-

lems for hyperbolic systems of equations. In this talk, I will show how model order reduction can be used to address some intrinsic difficulties of these problems. We consider ROMs that capture properties of the large problem that are essential for imaging and that can be realized via sparse graph-Laplacian networks. The ROMs are data-driven, i.e., they learn the underlying PDE problem from the transfer function. Here I will focus on one recent applications of this approach: A direct, nonlinear imaging algorithm in strongly heterogeneous media, where the ROM is used to manipulate the data in such a way that multiply scattered waves are separated from the single scattered ones. The algorithm, known as Data-to-Born map, transforms multi-scattering data to the single-scattering one. The latter can be effectively processed by any off-shelf linear algorithm.

Vladimir Druskin
Worcester Polytechnic Institute
vdruskin@gmail.com

Liliana Borcea
University of Michigan
borcea@umich.edu

Alexander V. Mamonov
University of Houston
mamonov@math.uh.edu

Mikhail Zaslavsky
Schlumberger-Doll Research
Cambridge, MA, USA
mzaslavsky@boston.oilfield.slb.com

MS131

A Backscattering Model Based on Corrector Theory of Homogenization for the Random Helmholtz Equation

This work concerns the analysis of wave propagation in random media. Our medium of interest is sea ice, which is a composite of a pure ice background and randomly located inclusions of brine and air. From a pulse emitted by a source above the sea ice layer, the main objective is to derive a model for the backscattered signal measured at the source/detector location. The problem is difficult in that, in the practical configuration we consider, the wave impinges on the layer with a non-normal incidence. Since the sea ice is seen by the pulse as an effective (homogenized) medium, the energy is specularly reflected and the backscattered signal vanishes in a first order approximation. What is measured at the detector consists therefore of corrections to leading order terms, and we focus here on the homogenization corrector. We describe the propagation by a random Helmholtz equation, and derive an expression of the corrector in this layered framework. We moreover obtain a transport model for quadratic quantities in the random wavefield in a high frequency limit.

Olivier Pinaud
Colorado State University
pinaud@math.colostate.edu

MS131

Long Time Homogenization of the Wave Equation in Locally Periodic Media

[1] A. Abdulle and T. Pouchon, ***Effective models for the multidimensional wave equation in heterogeneous media over long time and numerical homogenization*, Math. Models Methods Appl. Sci., 26 (2016), pp. 2651–2684.**

[2] ——, *Effective models for long time wave propagation in locally periodic media*, SIAM J. Numer. Anal., 56 (2018), pp. 2701–2730.

Timothée Pouchon
Ecole Polytechnique Federale de Lausanne
timothee.pouchon@gmail.com

MS132

Algorithmic Differentiation for Shape Optimization Problems in the High Level Finite Element Framework FEniCS

Shape optimization problems constrained by partial differential equations (PDEs) are ever-present in the area of scientific computing, for example structural mechanics, computational fluid dynamics and acoustics. The PDEs considered in such applications spans a broad spectrum, ranging from steady state, linear equations, to transient non-linear equations. To efficiently solve the shape optimization problem, one has to consider shape sensitivities and shape Hessians of the PDEs. A popular approach is the adjoint method. Deriving the adjoint equations by hand is known to be cumbersome and error-prone. We propose an approach where this derivation is computed by a high level algorithmic differentiation tool. The shape sensitivities are computed by discrete derivatives of the mesh node sensitivities. An advantage of our approach is that only the forward problem has to be postulated as input. Then the algorithm creates a computational tape that tracks the propagation of variables and computes corresponding gradients and Hessians. By using the operator overloading approach we inherit the parallelism and performance of the software used to solve the PDE. Our software is overloading the framework FEniCS, which has a high level Python interface and generates parallel, optimized C++ code. We illustrate the efficiency and robustness of the code by presenting results for several PDEs with a wide range of solution methods.

Jørgen Dokken
Simula Research Laboratory
dokken@simula.no

Simon W. Funke
Center for Biomedical Computing
Simula Research Laboratory
simon@simula.no

Sebastian Kenji Mitusch
Simula Resarch Laboratory
sebastkm@simula.no

Stephan Schmidt
University Würzburg
stephan.schmidt@mathematik.uni-wuerzburg.de

MS132

Latest Advances in Topology and Shape Optimiza-

tion

There is a rising interest in level set topology optimization as the design method for additive manufacturing due to its capability to obtain unintuitive designs. One area of potential exploitation offered by additive manufacturing is the range of length scale where feature sizes crossing nine orders of magnitudes may be manufacturable without a significant cost increase. We present two approaches to the numerical methods for level set topology optimization to large scale design problems. One is the multiscale formulation where the design problem is decomposed to two scales and coupled via homogenization. The decomposition formulation constitutes one macroscale and N number of microscale optimization for N number of material varying spatially. When N is one, it considers a single uniform material and when N is large, it can consider an almost continuous variation of material throughout the macroscopic structural domain. Such a physical decomposition means local material scale functionalities and constraints can be additionally applied offering a greater of level of freedom to engineers. This type of architected material can exhibit metamaterial properties achieving a superior and unintuitive design solution. The other approach is to push the limits of problem size by developing a sparse computational method using the VDB data structure for single scale topology optimization. We demonstrate that billions of voxel size problems is achievable with a medium size desktop.

H. Alicia Kim
University of Bath, United Kingdom
alicia@ucsd.edu

Sandilya Kambampati, Zongliang Du
University of California San Diego
akambampati@eng.ucsd.edu, zodu@eng.ucsd.edu

MS132

A Stress Constrained Level Set Topology Optimization Method with the Cut Finite Element Method

We introduce a linear elasticity topology optimization framework for mass minimization and pointwise stress constraints based on the cut finite element method [Burman, E., Claus, S., Hansbo, P., Larson, M. G. and Massing, A. CutFEM: discretizing geometry and partial differential equations. Int. J. Numer. Meth. Engng. 104:472501 (2015)]. Combining an explicit definition of the geometry using level sets with an accurate integration of the embedded geometry by means of the cut finite element method, we are able to obtain an accurate stress field without having to waste computational resources to accommodate the mesh to the geometry changes. The level sets are parametrized using radial basis functions that regularize the problem. The pointwise stress constraint is lumped together into a single global constraint and attached to the cost function weighted with a penalty parameter. We then optimize and discretize to avoid the derivation of the cumbersome discrete shape gradients, at the expense of introducing inconsistencies in the shape gradient that are reduced with the mesh refinement. The optimization problem is solved using a Newton method with a line search method to ensure global convergence. We restrict the geometry intersection points with the finite elements to avoid cases that could damage the convergence. To show the efficiency of the method, we solve the classic L-bracket example and compare it with previous approaches in the literature.

Miguel Salazar de Troya
Lawrence Livermore National Laboratory
University of Illinois at Urbana-Champaign
salazardetro1@llnl.gov

Daniel Tortorelli
Lawrence Livermore National Laboratory
Univerisity of Illinios at Urbana-Champaign
tortorelli2@llnl.gov

MS132

H-Div Based ZZ Error Indicator for Adaptive Mesh Refinement in Topology Optimization

An H-Div based Zienkiewicz Zhu (ZZ) error indicator is employed as a refinement criteria for topology optimization (TO) with adaptive mesh refinement (AMR). The computed finite element stresses are projected onto a Raviart-Thomas (RT) finite element space via the solution to a mass-matrix linear system to serve as an estimate of the true stress. This technique identifies and refines elements that introduce the most finite element error, providing a means to achieve an accurate solution with as few elements as possible. The ZZ error indicator is compared with other finite element error indicators as a mesh refinement criteria. Volume fraction based refinement is also considered. Compliance minimization subject to a volume fraction constraint is investigated along with mass minimization subject to a stress constraint.

Kenneth Swartz
University of Illinois at Urbana-Champaign
Lawrence Livermor
swartz10@llnl.gov

Youngsoo Choi, Jun Kudo
Lawrence Livermore National Laboratory
choi15@llnl.gov, kudo4@llnl.gov

Miguel Salazar de Troya
Lawrence Livermore National Laboratory
University of Illinois at Urbana-Champaign
salazardetro1@llnl.gov

Dan White
Lawrence Livermore National Laboratory
white37@llnl.gov

Dan Tortorelli
Lawrence Livermore National Laboratory
Univerisity of Illinios at Urbana-Champaign
tortorelli2@llnl.gov

MS133

Pulling 3D Cloud Structure Out of Passive Overhead Imagery: Observational and Computational

Challenges

Looking down at Earth from space or aircraft with passive optical imaging systems typically "flattens" the scene. Accordingly, spectral and/or polarimetric samples can be processed into geophysical properties, but they will be column-integrals. If information about the vertical structure of the atmosphere is desirable, then an active down-looking lidar system is required. In that case, artificial highly-collimated pulsed light sources (i.e., lasers) are used to convert time-of-flight into range, hence altitude. But overhead lidar systems generally sample only one vertical transect of the atmosphere/surface below the moving platform. Apart from that inherently limited spatial sampling in the horizontal plane, lidar systems are costly and high-maintenance. We will discuss two workarounds for gaining information on vertical structure in clouds. One approach uses a special spectroscopy focused on the abundant O_2 molecule where, in essence, modulation of radiant energy according to absorption strength is interpreted as the Laplace transform of the waveform for the reflected laser pulse, thus carrying the same information about vertical structure. The other is akin to stereoscopy, but uses many more directional samples of the cloud scene. The very high-dimensional multi-angle imagery is then used to determine a 3D grid of opacity values by using advanced inverse problem theory and an efficient 3D physics-based rendering model (a.k.a. the radiative transfer equation).

Anthony B. Davis
Jet Propulsion Laboratory
California Institute of Technology
anthony.b.davis@jpl.nasa.gov

Quentin Libois
CNRM, Météo France
U. de Toulouse, CNRS
quentin.libois@meteo.fr

Nicolas Ferlay
Laboartoire d'Optique Atmospherique, CNRS
U. Lille-1
nicolas.ferlay@univ-lille.fr

Aviad Levis, Yoav Schechner
Technion Israel Institute of Technology
aviad.levis@gmail.com, yoav@ee.technion.ac.il

MS133

Phase Transformation Models as Regularizers for Electron Microscope Imagery

Advances in Electron Microscopy image analysis traditionally have been achieved through the use of forward modeling. More recently, significant advances have been achieved by the introduction of inversion-based image processing algorithms. In particular, Markov Random Field (MRF) priors significantly improve the quality of reconstruction. Since the MRF has analogues in the Ising and Potts models, originally developed to describe magnetic transitions, the MRF has a more direct connection to materials images than to those of natural scenes: it simulates the (real) interfacial energy at boundaries between phases. Quantitative MRF descriptions are problematic, since the length scale of the MRF model is the atomic dimension, while in microscopy, it is typically on the order of thens of nanometers to microns. Thus, we developed the more modern Phase Field model as a regularizer, whose length scale matches that of the observation. The resulting regularized inversion method allows for simultaneous denoising and segmentation of polycrystalline samples. We applied this to simulated phantoms and real observations in SiC polycrystalline materials, and achieved performance exceeding that of the state of the art Non-local Means and BM3D denoising algorithms and successfully soft segmented the individual grains. The significant differences between regularizing with quantitative physics-based models and the more generic MRF regularizers are discussed.

Jeff Simmons
Air Force Research Laboratory
jeff.simmons.3@us.af.mil

Amir Koushyar Ziabari
Purdue University
aziabari@purdue.edu

Jeffery Rickman
Department of Materials Science and Engineering
Lehigh University
jmr6@lehigh.edu

Lawrence Drummy
Air Force Research Laboratory
lawrence.drummy.1@us.af.mil

Charles Bouman
School of ECE
Purdue University
bouman@ecn.purdue.edu

MS133

Challenges in Quantitative Image Analysis for Science and Engineering

Imaging is an increasingly common method for measuring physical objects in science and engineering applications. Unlike many consumer applications of imaging in which object detection and classification is the goal of image processing, technical domains often require quantification. Applications that can identify common objects like cats in photographs are commonplace. Deriving physically meaningful quantities describing these objects from imagery is significantly more challenging. Natural images often contain features that lack crisp boundaries, or are affected by noise and variability in imaging conditions. As such, analysis must be robust to these uncertainties, and furthermore, must allow such uncertainty to be communicated as part of the image analysis process. Additionally, many recent advances in computer vision and image processing have been driven both by algorithmic advances and the existence of large labeled data sets such as ImageNet. While algorithmic advances in quantitative analysis have continued in scientific and technical imaging, there is a significantly smaller number of data sets available for training and evaluating them. This talk will discuss some of the issues that arise in quantitative image analysis in scientific

contexts (e.g., biomedical) as well as industrial engineering contexts (e.g., measurement and tracking of structural defects).

Matthew Sottile
Noddle.io
matt@noddle.io

MS133

Camera-Based, Mobile Disease Surveillance using Convolutional Neural Networks

By combining state-of-the-art deep learning approaches in computer vision with mobile computing, we developed a highly advanced prototype disease surveillance tool that can identify sick individuals from a camera image. To be more specific, we applied a machine learning method called transfer learning to a unique dataset containing a few thousands images of healthy and sick labeled faces. To identify a sick face from a healthy one, we experimented with various pretrained Deep Convolutional Neural Networks (CNN) from Google including MobileNet for mobile and embedded vision applications. By applying MobileNet, we retrained the model and deployed it to our prototype mobile disease surveillance android application. This retrained CNN model gives high training, validation, and test accuracy, suggesting the high potential for using deep learning to distinguish between healthy and sick human face images. The prototype research needs further work to be deployed as an actual disease surveillance system, but this research proves its potential. The results described here show the suitability of applying deep learning methods for disease classification and the ability to run such a model on a mobile device for field use.

Altansuren Tumurbaatar
Washington State University
altaa.tumurbaatar@wsu.edu

Lauren Charles
Pacific Northwest National Laboratory
lauren.charles-smith@pnnl.gov

Saiteja Abbu
The State University of New York, Albany
abbusaitejayadav@gmail.com

Olatunde Madandola
Georgia Southern University
oladimejimadandola@gmail.com

Liangyu Tan
Iowa State University
tly1212@gmail.com

MS134

Graph Structure Inference: Methodology and Applications

Graphs are a fundamental, ubiquitous model for analysis and prediction of complex systems. However, most applications produce data that is not natively in the form of a graph. The many considerations in transforming these datasets to a graph structure are often overlooked, and have significant impact on downstream machine learning models and hypothesis testing. In this talk we present (1) the challenges in constructing graphs from data, (2) several applications illustrating these challenges, (3) rigorous, end-to-end evaluation methodologies over a set of potential graphs, (4) common best-practices for constructing and evaluating graphs from data.

Ivan Brugere
University of Illinios at Chicago
ibruge2@uic.edu

Tanya Y. Berger-Wolf
University of Illinois at Chicago
tanyabw@uic.edu

MS134

Constructing a Graph from user Implicit Feedback

eBay is an online marketplace and the United States 9th most visited website. eBay is working to model user behavior and predict future outcomes based on historical user actions and current market behavior. Modeling user behavior as a graph problem can help in discovering latent relationships, user segmentation, clustering and analytics. Recently, we worked on constructing a graph based on implicit feedback data. The user implicit feedback data is a matrix of users as rows and hundred of thousands of columns representing event types. An example of an event type can be the number of times a user has searched in a certain category over a period. Many eBay users both buy and sell on eBay Marketplace. Here, we concentrate on users who sell on eBay . With a user graph constructed on a months of data we evaluated the graph for its predictive ability in the future. We compared the actual buying behavior over the next month with the edge weighted mean of the neighbors buying behavior. A successful prediction is when a seller actually buys an item from a category that they were not buying before and matches with one of the top 5 predicted item categories. We found a success rate of approximately 5% with this setup. If the edge weighted mean of neighbors' buying behavior is computed iteratively, the number of predicted item category count increases.

Haozhe Zhang
Iowa State
haozhe@iastate.edu

Yan Qu
eBay Inc
yanqu@ebay.com

Shad Kirmani
Ebay
skirmani@ebay.com

MS134

The Core Resilience of Networks

The concept of k-cores is important for understanding the global structure of networks, as well as for identifying cen-

tral or important nodes within a network. It is often valuable to understand the resilience of the k-cores of a network to attacks and dropped edges (i.e., damaged communications links). In this talk, we provide a formal definition of a networks core resilience, and examine the problem of characterizing core resilience in terms of the networks structural features: in particular, which structural properties cause a network to have high or low core resilience? To measure this, we introduce two novel node properties, Core Strength and Core Influence, which measure the resilience of individual nodes core numbers and their influence on other nodes core numbers. Using these properties, we propose the Maximize Resilience of k-Core (MRKC) algorithm to add edges to improve the core resilience of a network.

Sucheta Soundarajan
Department of Electrical Engineering and Computer Science
Syracuse University
ssoundarajan@gmail.com

Ricky Laishram
Syracuse University
rlaishra@syr.edu

MS134

The Pursuit of Reproducibility through Well Planned and Executed Evidence Trails

Abstract not available.

Eric Stephan
Pacific Northwest National Laboratory
eric.stephan@pnl.gov

MS135

Computation of Koopman Modes in Chaotic Systems

Abstract not available.

Nisha Chandramoorthy
Massachusetts Institute of Technology
nishac92@gmail.com

MS135

Accuracy and Performance Tradeoffs for Robust Porous Media Discretization with Non-orthogonal Grids and Tensor Coefficients

Mixed Finite Element Methods (FEM) and multipoint flux Finite Volume Methods (FVM) can be used to accurately discretize Richards equation for porous media with heterogeneous coefficients and distorted meshes. Techniques such as inexact quadrature (lumping) can be used to make more solver-friendly and efficient formulations in exchange for some problem-dependent accuracy. Different multilevel solvers such as balancing domain decomposition by constraints and algebraic multigrid are suitable for these different formulations. This work investigates tradeoffs in accuracy versus cost metrics such as number of degrees of freedom, memory usage/bandwidth requirements, and solver efficiency using the best available scalable techniques. We consider a range of test problems with varying degrees of heterogeneity and compare accuracy metrics relevant to hydrologic applications. Results using the multipoint flux FVM and mixed FEM with exact and inexact quadrature are analyzed using the method of manufactured solutions and heterogeneous benchmark problems.

Reagan Cronin
University of Colorado Boulder
reagan.cronin@colorado.edu

Jed Brown
Univeristy of Colorado Boulder
jed@jedbrown.org

MS135

Multi-scale Modeling of Paracrine PDGF-driven Glioma Growth and Invasion

The most common primary brain tumor in adults, glioma claims thousands of lives each year. It is a notoriously aggressive and invasive disease, and despite efforts to improve survival rates, the standard of care has remained unchanged for more than a decade. Platelet–derived growth factor (PDGF) is often over-expressed in gliomas, where it can drive tumor growth via autocrine and paracrine stimulation of PDGF receptor (PDGFR)-expressing glioma cells, as well as paracrine recruitment of non-neoplastic oligodendroglial progenitor cells (OPCs), which also express PDGFRs. Constructing multi-scale mechanistic mathematical models allows us to better examine the effects of paracrine PDGF signaling on glioma growth. Model simulations show that increased PDGF signaling can increase growth rates and alter the distribution of glioma cells infiltrating adjacent normal brain tissue. While the use of PDGF inhibitors has remained largely unsuccessful at improving patient outcomes in glioblastoma, this may be due to inadequate targeting of these agents to the best candidates. By incorporating different treatment simulations in our model, we show that PDGF inhibition results in decreased OPC recruitment, which leads to slower growing, but more diffusely infiltrating tumors.

Susan Massey
Mayo Clinic Arizona
Precision Neurotherapeutics Innovation Program
massey.susan@mayo.edu

MS135

An Uncertainty Representation for Model Inadequacy in a Field-scale Contaminant Transport Model

The advection-diffusion equation (ADE) is an inadequate predictor of mean contaminant transport through heterogeneous porous media at the field scale because of missing microscale dynamics affecting the mesoscale transport. With the goal of prediction, a representation of uncertainty caused by this inadequacy is embedded in the ADE. The so-called inadequacy representation is modeled as a stochastic operator acting on the state variable. It is formulated to be extrapolative by constraining the deterministic structure of the operator to reflect prior information about the

physical process and the nature of the inadequacy. Recent work incorporating the dependence of the inadequacy on velocity statistics is presented.

Teresa Portone, Robert D. Moser
University of Texas at Austin
tnportone@gmail.com, rmoser@ices.utexas.edu

MS136

Scalable Gaussian Process Computations using Hierarchical Matrices

We present a kernel-independent method that applies hierarchical matrices to the problem of maximum likelihood estimation for Gaussian processes. The proposed approximation provides natural and scalable stochastic estimators for its gradient and Hessian, as well as the expected Fisher information matrix, that are computable in quasilinear $O(n \log^2 n)$ complexity for a large range of models. To accomplish this, we (i) choose a specific hierarchical approximation for covariance matrices that enables the computation of their exact derivatives and (ii) use a stabilized form of the Hutchinson stochastic trace estimator. Since both the observed and expected information matrices can be computed in quasilinear complexity, covariance matrices for MLEs can also be estimated efficiently. After discussing the associated mathematics, we demonstrate the scalability of the method, discuss details of its implementation, and validate that the resulting MLEs and confidence intervals based on the inverse Fisher information matrix faithfully approach those obtained by the exact likelihood.

Christopher Geoga
Rutgers University
cgeoga@anl.gov

Mihai Anitescu
Argonne National Laboratory
Mathematics and Computer Science Division
anitescu@mcs.anl.gov

Michael Stein
University of Chicago
stein@galton.uchicago.edu

MS136

Structured Modeling and Decomposition Methods in Pyomo

Large-scale optimization problems are often solved using decomposition-based solution approaches that exploit certain structure in the problem. For example, Schur complement decomposition can be used to solve the linear systems arising from dynamic optimization problems and progressive hedging can be used to solve two-stage stochastic programming problems. While these solution approaches and the structures they exploit are well-established, very few general implementations exist that can be easily applied to a variety of models. One of the main challenges is that most algebraic modeling languages (AMLs) do not capture the high-level structure exploited by these decomposition algorithms. In this talk we demonstrate how the Pyomo AML is able to capture this high-level structure and therefore enable general implementations of several decomposition methods. We also give an overview of the large-scale applications that have been solved using decomposition methods in Pyomo with a focus on dynamic optimization problems.

Bethany Nicholson, Carl Laird
Sandia National Laboratories
blnicho@sandia.gov, cdlaird@sandia.gov

John D. Siirola
Sandia National Laboratories
Exploratory Simulation Technologies
jdsiiro@sandia.gov

MS136

Chordal Decompositions in Semidefinite Programming: Trading Stability for Scalability

Chordal decompositions in Semidefinite Programs (SDP) have allowed Interior Point Methods (IPMs) to solve large-scale SDPs. However, we show that the SDP decomposition is primal degenerate when the SDP has a low-rank solution even when the original SDP is not. This renders the decomposition numerically ill-conditioned near a solution. Numerical experiments demonstrate that the SDP decomposition results in worse ill-conditioning of the IPM linear system as compared to for the original SDP formulation. We describe a decomposition approach based on the recently proposed LDL^T direction that can avoid the ill-conditioning. Numerical results confirm the stability of the proposed approach.

Arvind Raghunathan
Mitsubishi Electric research Laboratories
raghunathan@merl.com

Larry Biegler
Carnegie Mellon University
lb01@andrew.cmu.edu

MS136

Parallel Multigrid Reduction in Time (MGRIT) Applied to Power Grid Simulations

Since clock speeds are no longer increasing, time integration is becoming a sequential bottleneck. The multigrid reduction in time (MGRIT) algorithm is an approach for creating concurrency in the time dimension that can be exploited to overcome this bottleneck and is designed to build on existing codes and time integration techniques. In this talk we present an approach for multi-step backward difference formula (BDF) integration of fully implicit differential algebraic equations (DAE) on adaptive variable timestep grids. Results of both application of this multistep approach and a multistage approach to power grid applications within the LLNL transmission grid modeling code, GridDyn, with and without discontinuities, will be presented. An overview of MGRIT and the corresponding open source XBraid library will also be given. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. LLNL-ABS-757660.

Carol S. Woodward, Robert D. Falgout

Lawrence Livermore National Laboratory
woodward6@llnl.gov, rfalgout@llnl.gov

Matthieu Lecouvez
CEA
matthieu.lecouvez@cea.fr

Jacob B. Schroder
Department of Mathematics and Statistics
University of New Mexico
jbschroder@unm.edu

Philip Top
Lawrence Livermore National Laboratory
top1@llnl.gov

MS137

DUNE Collaborating via Interfaces

DUNE is a well established toolbox for solving PDEs. Over the years the project has grown significantly and has undergone several changes in the organization. DUNE is developed as open source, inviting new researchers to contribute. A central idea in DUNE is the design and definition of interfaces, which enable a successful collaboration. These are fine grained interfaces modeled according to the underlying mathematical concepts. In order to ensure efficiency DUNE make excessive use of modern C++ techniques and generic programming. This design is also reflected in the modular code structure. There is not the DUNE software, but DUNE consists of a set of code modules, and many addon modules developed around this core functionality. These addons allow new contributions to be develop independent from the existing code. A particular challenge are fine grained interfaces as they enforce a relatively strong integration of the different components, so that it is hard or often impossible to use numerical methods developed in DUNE without the core DUNE stack. We will critically discuss how the eco-system of DUNE works, which lessons can be learned for future software development in the scientific community.

Christian Engwer
INAM, University of Münster
christian.engwer@uni-muenster.de

MS137

An Innovative Method for Integration of Simulation/Data/Learning in the Exascale/Post-Moore Era

Supercomputing in the Exa-scale and the Post-Moore Era is inherently different from that in the Peta-scale Era and before. Although supercomputers have been the essential tool for computational science in recent 30 years, they are now used for other purposes, such as data analytics and machine learning. The architecture of the next generation supercomputing system is essentially heterogeneous for these multiple purposes (simulations + data + learning). In the present work, we propose a new innovative method for integration of computational and data science (Big Data & Extreme Computing, BDEC) for sustainable promotion of new scientific discovery by supercomputers in the Exa-Scale/Post-Moore Era with heterogeneous architecture. **h3-Open-BDEC** (h3: hierarchical, hybrid, heterogeneous) is an open source infrastructure for development and execution of optimized and reliable codes for BDEC on such supercomputers.In this presentation, we will overview the h3-Open-BDEC, and the target supercomputer system, which will start operation in April 2021.

Kengo Nakajima
The University of Tokyo
Information Technology Center
nakajima@cc.u-tokyo.ac.jp

Takeshi Iwashita
Hokkaido University
iwashita@iic.hokudai.ac.jp

Takashi Shimokawabe
Global Scientific Information and Computing Center
Tokyo Institute of Technology
shimokawabe@cc.u-tokyo.ac.jp

Takahiro Katagiri
Information Technology Center
Nagoya University
katagiri@cc.nagoya-u.ac.jp

Hisashi Yashiro
RIKEN AICS
h.yashiro@riken.jp

Hiroya Matsuba
RIKEN R-CCS
hiroya.matsuba@riken.jp

MS137

Software Development Kits: A Software Integration Strategy for CSE

Depending on a large number of CSE software packages is challenging for many reasons, including incompatibility among software versions, needing to use several different build systems, lack of portability to new platforms, and unknown future commitment to software support. The Software Development Kit (SDK) effort in the Exascale Computing Project (ECP) focuses on these and other issues impacting the ability and practicality of scientific application teams to use and depend on other CSE software packages. The multifaceted approach includes a common delivery of dozens of ECP software products via Spack, the creation of a continuous integration testing capability for SDK member packages, and extension and generalization of community policies determined by the xSDK effort, which focuses numerical software. This presentation will describe the status of the SDK effort nearly a year after the start of the project, including lessons learned and difficulties we have encountered.

James Willenbring
Sandia National Laboratories

jmwille@sandia.gov

MS137

xSDK: Foundations of a Numerical Software Ecosystem for High-performance CSE

With the development of increasingly complex architectures with the potential for higher performance, software complexity is increasing due to multiphysics and multiscale modeling, the coupling of simulations and data analytics, and the demand for greater reproducibility in the midst of disruptive architectural changes. Applications increasingly require the combined use of independent software packages, which have diverse sponsors, priorities, and processes for development and release. These challenges create the unique opportunity to fundamentally change how scientific software is designed, developed, and sustained—with explicit work toward scientific software ecosystems. This presentation will introduce the xSDK, or Extreme-scale Scientific Software Development Kit, where community-defined policies are increasing the quality and interoperability across numerical libraries as needed by the DOE Exascale Computing Project. The community policies developed through the xSDK project, their expected benefits, and lessons learned from adoption of these policies into existing packages will be presented. Lastly, plans for new developments, e.g. new member libraries, or plans for increasing interoperability among member projects, will be discussed.

Ulrike M. Yang
Lawrence Livermore National Laboratory
umyang@llnl.gov

MS138

Petaflop Scale Electronic Structure Calculations with RMG and GPU Accelerators

Electronic structure calculations based on Density Funtional Theory DFT are widely used to investigate the properties of materials. While less computationally demanding than many-body approaches, DFT calculations still typically exhibit $O(N^3)$ scaling with the number of electronic orbitals so efficient computational techniques are required. RMG is an electronic structure code that solves the Kohn-Sham equations of DFT using real-space grids to represent objects of interest such as charge densities, electronic orbitals, and potentials. The solution process uses multigrid methods and iterative diagonalization to obtain self-consistent solutions to the Kohn-Sham equations. RMG is well suited for petascale level calculations and can accelerate time to solution using GPU acclerators when they are available. RMG is cross platform and available for Linux, Windows and Mac systems as well as clusters and supercomputers. Applications range from semiconductors to biological systems and a localized orbital variant of the code has been used for quantum transport calculations on systems containg up to 10,000 atoms.

Emil Briggs
NC State University
elbriggs@ncsu.edu

Wenchang Lu, Jerzy Bernholc
North Carolina State University
wlu2@ncsu.edu, bernholc@ncsu.edu

MS138

Development of Plane-wave and O(N) Methods in Nwchemex for Emerging Exascale Machines

Abstract not available.

Eric J. Bylaska
Fundemental Sciences Laboratory
Pacific Northwest National Laboratory
eric.bylaska@pnl.gov

MS138

GPU-accelerated Real Space Electronic Structure Theory on HPC Resources

We present our experiences implementing GPU acceleration in the massively parallel, real space FHI-aims electronic structure code for computational materials science. We highlight how FHI-aims achieves efficient scaling on HPC resources using real-space operations and demonstrate how the vectorized implementation originally designed for CPUs was ported to GPUs with minimal code rewrite. This is made possible by a domain decomposition scheme which divides the integration grid into groups ("batches") of points, which can be distributed in an embarrassingly parallel fashion amongst tasks, reducing the problem to dense serial linear algebra on each task. To assess the performance of our GPU implementation, we performed benchmarks on three different architectures using a 103-material test set. We find that operations which are particularly reliant on dense serial linear algebra show dramatic speedups from GPU acceleration: notably, SCF iterations including force and stress calculations show speedups ranging from 4.5x to 6.6x. This translates to an expected overall speedup between 3x-4x for the entire calculation, as is highlighted in the example of diamond silicon. Calculations for a 375-atom Bi2Se3 bilayer were performed on OLCF's Titan to verify that the GPU implementation scales for large-scale calculations.

William P. Huhn, Bjoern Lange, Victor Yu
Duke University
william.huhn@duke.edu, bjoern.lange@dacs-labs.com,
wenzhe.yu@duke.edu

Mina Yoon
CNMS, Oak Ridge National Laboratory
myoon@ornl.gov

Volker Blum
Duke University
Department of Mechanical Engineering and Materials Science
volker.blum@duke.edu

MS138

Large Scale Hybrid Functional Plane Wave Pseudopotential Density Functional Theory Calculation

on GPU Cluster

In this talk, we present an effective GPU implementation of a large scale plane hybrid functional density functional theory method. This implementation includes recent methods of both adptive compressed exchange operator (ACE) and PC-DIIS. Our tests show that for a 1000 atom Silicon system, the most time consuming part, namely the $O(N_e^2)$ Poisson-like solver, can be finished in 9 seconds on SummitDev Supercomputer with 200 GPUs. This is 10× faster than the CPU code run in full parallel with 2000 CPU cores on Edison supercomputer in National Energy Research Supercomputing Center. The success of our implementation relies mainly on: 1) mixed precison calculation during the ACE operator construction 2) overlapping GPU computation with CPU communication. 3) moving all calculations to the GPU. We show that the ratio between fock operator calculation and eigen solver down to 2-3 for the GPU code, which is about 10 in the CPU code.

Weile Jia
University of California at Berkeley
wjia@lbl.gov

MS139

Numerical Relativity, Artificial Intelligence, Big Data Analytics and High Performance Computing: An Emergent Transdisciplinary Paradigm to Realize Multimessenger Astrophysics

The fusion of highly parallel numerical relativity simulations, high performance computing and innovative AI algorithms has emerged as a novel paradigm to address grand computational challenges in Multimessenger Astrophysics. In this talk, I will review this emergent trend, and will discus the critical role of High Performance Computing to enable real-time discovery at scale in the Multimessenger Astrophysics era.

Elliu Huerta
National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign
elihu@illinois.edu

MS139

On the Development and use of ADER-DG Methods in Computational Relativistic Astrophysics

We present a new class of high-order accurate numerical algorithms for solving the equations of general-relativistic ideal magnetohydrodynamics in curved space-times. In this paper, we assume the background space-time to be given and static, i.e. we make use of the Cowling approximation. The governing partial differential equations are solved via a new family of fully discrete and arbitrary high-order accurate path-conservative discontinuous Galerkin (DG) finite-element methods combined with adaptive mesh refinement and time accurate local time-stepping. In order to deal with shock waves and other discontinuities, the high-order DG schemes are supplemented with a novel a posteriori subcell finite-volume limiter, which makes the new algorithms as robust as classical second-order total-variation diminishing finite- volume methods at shocks and discontinuities, but also as accurate as unlimited high-order DG schemes in smooth regions of the flow. We show the advantages of this new approach by means of various classical two- and three-dimensional benchmark problems on fixed space-times. Finally, we present a strongly hyperbolic first-order formulation of the Einstein equations based which combines the advantages of a conformal and traceless formulation, with the suppression of constraint violations given by the damping terms, but being first order in time and space, it is particularly suited for a discontinuous Galerkin (DG) implementation.

Luciano Rezzolla
Max-Planck Institut für Gravitationsphysik
rezzolla@aei.mpg.de

MS139

Numerical Relativity in the Age of Gravitational Wave Observations

The advent of gravitational wave astronomy has created opportunities to probe strong-field gravity as we measure the merger of black holes. Numerical relativity provides the means to confront the measurements with theoretical prediction. In this talk, Ill discuss the role numerical relativity played in the observed black hole binaries by LIGO and Virgo and the future potential for unveiling strong-field gravity in both future ground and space based detectors with an emphasis on the computational challenges.

Deirdre Shoemaker
Center for Relativistic Astrophysics, School of Physics
Georgia Tech
deirdre@gatech.edu

MS139

Numerical Relativity at the Extremes

Numerical simulations of highly-spinning and high-mass ratio binaries are generally quite slow compared to simulations of more modest binaries. This is due to the extreme resolution required for accurate evolutions and the disparate timescales set by the resulting Courant stability condition and the much larger dynamical timescale of the binary's inspiral. In this talk, I will describe the numerical challenges encountered in evolving 100:1 mass ratio binaries 0.99 spinning binaries, as well as techniques for improving simulation efficiency.

Yosef Zlochower
Rochester Institute of Technology
Center for Computational Relativity and Gravitation
yrzsma@rit.edu

MS140

A Fast Direct Solver for Multilayered Quasi-periodic Scattering

Abstract not available.

Adrianna Gillman
Rice University
Department of Computational and applied mathematics

adrianna.gillman@rice.edu

MS140

Symbolic Computation for Layer Potential Evaluation with Quadrature by Expansion

In this talk, I discuss a number of issues arising in the evaluation of complex layer potential operators in which symbolic computation is instrumental for good performance. The rapid evaluation of layer potentials plays an important role in the solution of elliptic boundary value problems by way of integral equation methods, in turn a key building block for many key simulations in computational science. The method of Quadrature By Expansion (QBX), recent advances in fast algorithms for QBX, and the symbolic approaches detailed here open an opportunity to create numerical tools for layer potential evaluation that are competitive in performance, generality, and convenience to competing finite element toolsets. Opportunities to exploit symbolic computation in this setting include the optimization of compositions of layer potential operators, the synthesis of high-perfomrance expansion and translation operators for arbitrary and particularly complexk kernels, and the automated derivation of an integral equation operator from a representation formula. In each of these settings, symbolic computation replaces laborious and error-prone manual coding that was previously required. I discuss symbolic manipulation techniques that lead to low (and in some cases, optimal or near-optimal) complexity for many of these use cases. The talk concludes with performance results for the demonstrated techniques in a number of application settings.

Andreas Kloeckner
University of Illinois
andreask@illinois.edu

MS140

Adaptive BEM with Inexact PCG Solver Yields Almost Optimal Computational Costs

We consider PCG with optimal preconditioner in the frame of BEM with adaptive mesh-refinement. As model problem serves some weakly-singular integral equation $V\phi = f$. Given an initial mesh $\mathcal{T}_0$, parameters $0 < \theta \leq 1$ and $\lambda > 0$, counters $j = 0 = k$, as well as a discrete initial guess $\phi_{00} \approx \phi$ on $\mathcal{T}_0$, our adaptive strategy reads as follows:

(i) Update $(j,k) \mapsto (j,k+1)$.

(ii) Do one PCG step to obtain ϕ_{jk} from $\phi_{j(k-1)}$.

(iii) Compute indicators $\eta_j(T,\phi_{jk})$ for all elements $T \in \mathcal{T}_j$.

(iv) If $\lambda^{-1}\,\|\phi_{jk}-\phi_{j(k-1)}\|^2 > \sum_{T\in\mathcal{T}_j}\eta_j(T,\phi_{jk})^2$, continue with (i).

(v) Otherwise determine elements $\mathcal{M}_j \subseteq \mathcal{T}_j$ such that $\theta\sum_{T\in\mathcal{T}_j}\eta_j(T,\phi_{jk})^2 \leq \sum_{T\in\mathcal{M}_j}\eta_j(T,\phi_{jk})^2$.

(vi) Refine all $T \in \mathcal{M}_j$ to obtain $\mathcal{T}_{j+1}$.

(vii) Update $(j,k) \mapsto (j+1,0)$ and continue with (i).

We prove that this algorithm does not only lead to linear convergence with optimal algebraic rates, but also to almost optimal computational costs, if $\mathcal{H}^2$-matrices are employed. In particular, we provide an optimal additive Schwarz preconditioner which can be computed in linear complexity. The talk is based on recent work (arXiv:1806.00313, 2018).

Stefan Schimanko
TU Wien
stefan.schimanko@asc.tuwien.ac.at

MS140

An Integral Equation Method for the Cahn-Hilliard Equation in the Wetting Problem

We present an integral equation approach to solving the Cahn-Hilliard equation equipped with boundary conditions that model solid surfaces with certain Young's angles. Discretization of the system in time using convex splitting leads to a modified biharmonic equation at each time step. To solve it, the basic idea is to split the solution into a volume potential computed with free space kernels, plus the solution to a second kind integral equation (SKIE). The volume potential is evaluated with a box-based volume-FMM method. For non-box domains, source density is extended by solving a biharmonic Dirichlet problem. The near-singular boundary integrals are computed using quadrature by expansion (QBX) with FMM acceleration. Our method has linear complexity and can achieve high order convergence with adaptive refinement. We showcase applications in simulating wetting transition and boundary effects on spinodal decomposition.

Shidong Jiang
Department of Mathematical Sciences
New Jersey Institute of Technology
shidong.jiang@njit.edu

Andreas Kloeckner
University of Illinois
andreask@illinois.edu

Xiao-Ping Wang, Xiaoyu Wei
Hong Kong University of Science and Technology
mawang@ust.hk, xywei@illinois.edu

MS141

Transfer Learning and Generalization in Physical Problems

Applications of machine learning in physical problems has mostly been focused on reconstructing or predicting properties of a given system, in the presence of some training data. Physical laws, however, are applicable across different systems and regimes. The question we raise in this talk is whether one can construct data-driven models that can learn from different systems, and transfer this physical knowledge to make predictions in other systems that are governed by similar physics. This defines a new paradigm of transfer learning. Taking steps towards this end, we embed neural networks within physics-based PDE models and train the hybrid model in an integrated fashion. An immediate benefit of the formulation is that the hybrid model satisfies the same physical constraints as the PDE model. Another key aspect is that the loss function for the neural network is defined based on the output of the hybrid model, thus enforcing consistency between the learning, inference

and model construction. Such hybrid models are trained across different systems that are representative of the underlying physics and the final model is one that embeds a neural network that achieves consensus across the different problems.

Karthik Duraisamy
University of Michigan Ann Arbor
kdur@umich.edu

MS141

Physics-informed Deep Generative Models

We consider the application of deep generative models in propagating uncertainty through complex physical systems. Specifically, we put forth an implicit variational inference formulation that constrains the generative model output to satisfy given physical laws expressed by partial differential equations. Such physics-informed constraints provide a regularization mechanism for effectively training deep probabilistic models for modeling physical systems in which the cost of data acquisition is high and training data-sets are typically small. This provides a scalable framework for characterizing uncertainty in the outputs of physical systems due to randomness in their inputs or noise in their observations. We demonstrate the effectiveness of our approach through a canonical example in transport dynamics.

Paris Perdikaris
University of Pennsylvania
pgp@seas.upenn.edu

MS141

Hidden Physics Models: Machine Learning of Nonlinear Partial Differential Equations

A grand challenge with great opportunities is to develop a coherent framework that enables blending conservation laws, physical principles, and/or phenomenological behaviours expressed by differential equations with the vast data sets available in many fields of engineering, science, and technology. At the intersection of probabilistic machine learning, deep learning, and scientific computations, this work is pursuing the overall vision to establish promising new directions for harnessing the long-standing developments of classical methods in applied mathematics and mathematical physics to design learning machines with the ability to operate in complex domains without requiring large quantities of data. To materialize this vision, this work is exploring two complementary directions: (1) designing data-efficient learning machines capable of leveraging the underlying laws of physics, expressed by time dependent and non-linear differential equations, to extract patterns from high-dimensional data generated from experiments, and (2) designing novel numerical algorithms that can seamlessly blend equations and noisy multi-fidelity data, infer latent quantities of interest (e.g., the solution to a differential equation), and naturally quantify uncertainty in computations. The latter is aligned in spirit with the emerging field of probabilistic numerics.

Maziar Raissi
Brown University
maziar_raissi@brown.edu

MS141

Deep Learning Algorithm for the Data-driven Simulation of Noisy Dynamical System

Data-driven reconstruction of a dynamical system has been of great interest due to its direct relevance to numerous applications across disciplines. While there is a large volume of literature, most of the classical "model-free' approaches rely on linearization and an assumption of Gaussian noise. Here, we present a deep learning model, DE-LSTM, for the data-driven simulation of a noisy nonlinear dynamical system. In DE-LSTM, the probability density function of a stochastic process is approximated by a numerical discretization, and the underlying nonlinear dynamics is modeled by the Long Short-Term Memory (LSTM) network. It is shown that the function approximation problem can be solved by a multi-label classification problem by using the discretization. A penalized maximum log likelihood method is proposed to impose a smoothness condition in the predicted probability distribution. We outline a Monte Carlo method to compute the time evolution of the probability distribution. The behavior of DE-LSTM is thoroughly investigated by using the Ornstein-Uhlenbeck process and noisy observations of nonlinear dynamical systems; Mackey-Glass equation and forced Van der Pol oscillators. It is shown that DE-LSTM makes a good prediction of the probability distribution without assuming distributional properties of the noise. For a multiple-step forecast, it is found that the prediction uncertainty, denoted by the 95% confidence interval, dynamically adjusts following the system dynamics.

Kyongmin Yeo
IBM
kyongmin@gmail.com

MS142

Neural Style Transfer for Enriching Turbulent Flow

Numerical simulations of turbulent flows in domains and conditions relevant to engineered systems remains a persistent challenge due to the vast range of spatial and temporal scales, chaotic dynamics, complex nonlinear interactions exhibited by turbulence. Data-driven machine learning techniques have recently yielded advances across many computational science domains, with noteworthy successes in turbulent flow applications related to turbulence modeling, data compression, and inflow generation. In this talk, we demonstrate the use of neural style transfer to enrich turbulent flows. We propose an analogy from image content and style distinctions to filtered and fluctuating components of turbulent velocity fields. Using neural style transfer techniques, we enrich filtered velocity fields with various small scale styles representing different turbulence conditions. We examine the choice of loss function, and discuss how perceptual and temporal considerations improve the physical fidelity of the enriched turbulent fields. Finally, we provide an example of turbulent flow enrichment via neural style transfer on the generation of a complex wind turbine wake flowfield, and discuss impacts on simulation convergence time, data compression, and engineering

loads calculations.

Ryan King, Peter Graf, Michael Sprague
NREL
ryan.king@nrel.gov, peter.graf@nrel.gov, michael.a.sprague@nrel.gov

MS142

Learning Atomic Force Fields with Compositional Neural Networks

Machine learning is moving from the purely engineering, commercial domain to having an ever greater impact in the sciences. One of the specific tasks for which the ML approach holds much promise is learning the atomic potentials or "force fields" experienced by atoms in complex environments. In this talk we give an overview of some recent neural networks approaches to this problem. In particular, we focus on how the compositional structure of the system and the underlying symmetries can be reflected in the structure of the network and why this is important for learning a physically faithful model.

Risi Kondor
University of Chicago
risi@cs.uchicago.edu

MS142

An Information-theoretic Approach to Fusion of Multi-scale Data and Models

We employ an information-theoretic approach that allows for seamless integration of multi-resolution data into multi-scale simulations to upscale/downscale hydraulic conductivity of heterogeneous porous formations. Available data (at either the fine- or the coarse-scale) are used to inform models at the opposite scale by setting a probabilistic equivalence between the fine and the coarse scale, with closures (parameters and/or constitutive laws) that are learnt via minimization of observables error and mutual information across scales. Finally, we investigate means to accelerate scaling of dynamic processes.

Hannah Lu, Francesca Boso
Stanford University
hannahlu@stanford.edu, fboso@stanford.edu

Daniel M. Tartakovsky
University of California, San Diego
tartakovsky@stanford.edu

MS142

Classical and Quantum Machine Learning with Tensor Networks

In the last few decades, there has been great progress in algorithms for working with many-body quantum wavefunctions. Though most many-body wavefunctions are prohibitively large, interesting wavefunctions such as ground states of Hamiltonians have structures of entanglement which allow them to be compressed. The compression is carried out by representing the wavefunction as a tensor network, which are equivalent to certain interesting quantum circuits. It turns out that tensor networks are actually a very general tool for compressing large tensors. I will discuss an interesting framework for machine learning where large tensors naturally arise and can be successfully represented by tensor networks of the same type used in quantum physics. The approach has many interesting payoffs for machine learning, such as significantly better scaling for kernel-learning models and adaptive training algorithms. One of the most interesting direction for using tensor networks for machine learning is that the same models one uses classically can be straightforwardly evaluated on quantum hardware. Models can be pre-trained classically then extended to use quantum resources. I will discuss proof-of-principle numerical experiments, including robustness to qubit noise, and conclude by discussing the feasibility of evaluating models for realistic data sets on small, near-term quantum computers.

Miles Stoudenmire
Flatiron Institute
mstoudenmire@flatironinstitute.org

MS143

Data-driven Approach of Quantifying Uncertainty in Complex Systems with Arbitrary Randomness

The challenge of quantifying uncertainty propagation in real-world systems is rooted in the high-dimensionality of the stochastic input and the frequent lack of explicit knowledge of its probability distribution. Traditional approaches show limitations for such problems, especially when the size of the training data is limited. To address these difficulties, we have developed a general framework of constructing surrogate models on spaces of stochastic input with arbitrary probability measure irrespective of dependencies and the analytical form. This framework includes a data-driven basis construction for arbitrary probability measure and a sparsity enhancement rotation procedure. This work builds on the previous rotation method, ensuring that the developed basis construction approach maintains the orthogonality/near-orthogonality with respect to the density of the rotated random vector. The developed framework leads to accurate recovery, with only limited training data, of a sparse representation of target functions depending on high-dimensional inputs. The effectiveness of our method is demonstrated in challenging problems such as PDE systems and realistic biomolecular systems where the underlying density is implicitly represented by a collection of sampling data, as well as systems with explicitly given non-Gaussian probabilistic measures.

Jing Li
Pacific Northwest National Lab
jing.li@pnnl.gov

Huan Lei, Peiyuan Gao, Nathan Baker
Pacific Northwest National Laboratory
huan.lei@pnnl.gov, peiyuan.gao@pnnl.gov, nathan.baker@pnnl.gov

MS143

Reducing Parameter Space for Neural Network

Training

For neural networks (NNs) with rectified linear unit (ReLU) or binary activation functions, we show that their training can be accomplished in a reduced parameter space. Specifically, the weights in each neuron can be trained on the unit sphere, as opposed to the entire space, and the threshold can be trained in a bounded interval, as opposed to the real line. We show that the NNs in the reduced parameter space are mathematically equivalent to the standard NNs with parameters in the whole space. The reduced parameter space shall facilitate the optimization procedure for the network training, as the search space becomes (much) smaller. We demonstrate the improved training performance using numerical examples.

Tong Qin, Ling Zhou, Dongbin Xiu
Ohio State University
qin.428@osu.edu, zhou.2568@osu.edu, xiu.16@osu.edu

MS143

Physics-informed GANs for Stochastic Differential Equations

We developed a new class of physics-informed generative adversarial networks (PI-GANs) to solve in a unified manner forward, inverse and mixed stochastic problems based on a limited number of scattered measurements. We encoded into the architecture of GANs the governing physical laws in the form of stochastic differential equations (SDEs) using automatic differentiation. In particular, we applied Wasserstein GANs with gradient penalty (WGAN-GP) for its enhanced stability compared to vanilla GANs. We tested WGAN-GP in approximating Gaussian processes of different correlation lengths based on data realizations collected from simultaneous reads at sparsely placed sensors, and obtained good approximation of the generated stochastic processes to the target ones even for a mismatch between the input noise dimensionality and the effective dimensionality of the target stochastic processes. We also studied the overfitting issue for both the discriminator and generator, and we found that overfitting occurs also in the generator in addition to the discriminator as previously reported. Subsequently, we solved forward, inverse, and mixed SDE problems without changing the framework of PI-GANs, obtaining both the means and standard deviations of the stochastic solution and the diffusion coefficient in good agreement with benchmarks. PI-GANs could tackle very high dimensional problems given more sensor data with low-polynomial growth in computational cost.

Liu Yang
Division of Applied Mathematics
liu_yang@brown.edu

Dongkun Zhang
Division of Applied Mathematics
Brown University
dongkun_zhang@brown.edu

George Em Karniadakis
Brown University
george_karniadakis@brown.edu

MS143

Coupling of Non-intrusive and Conventional RB Models for Large-scale Structural Problems

In solid mechanics, linear structures exhibit (local) nonlinear behaviors when they are close to failure, for instance, a structure's elastic deformation turns into plastic by being deformed beyond recovery. To properly assess this kind of problems in real-life application, we need fast and multi-query evaluations of coupled linear and nonlinear structural systems, whose solutions are not straight forward and computationally expensive to obtain. We propose a linear-nonlinear domain decomposition where the two systems are coupled through the solutions on the linear-nonlinear interface. We then utilize different model order reduction techniques to address the linear and nonlinear problems individually. We adopt non-intrusive method, e.g. Gaussian Processes regression, to tackle the nonlinearity, while conventional reduced basis (RB) method is employed for solutions in the linear domain.

Zhenying Zhang
EPFL
zhenying.zhang@epfl.ch

Mengwu Guo
Ecole Polytechnique Fédérale de Lausanne
mengwu.guo@epfl.ch

Jan S. Hesthaven
EPFL
jan.hesthaven@epfl.ch

MS144

Multigrid Solver Integrated Tree-Based Adaptive Refinement (SITAR) in PETSc

We discuss advanced algebraic solver methods currently under development in PETSc (www.mcs.anl.gov/petsc) in the Solver Integrated Tree-based Adaptive Refinement (SITAR) infrastructure. Algebraic multigrid (AMG) is a popular scalable, fairly general, algebraic solver and PETSc provides three distinct AMG solvers (ML, hypre and the native GAMG). While AMG is and will continue to be used extensively in the foreseeable future, AMG does have limitations and can be difficulty to use at extreme-scale. These limitations come from the simple, by design, interface between the solver and application. Geometric multigrid (GMG) is mathematically more powerful in some ways than AMG, but GMG requires coarse grid mesh generation and access to the discretization for optimal nonlinear, matrix-free methods. SITAR provides infrastructure to develop GMG and AMG solvers with integrated discretizations, mesh management and solvers. PETSc provides high order finite element, finite volume and discontinuous Galerkin discretizations, structured, semi-structured and unstructured mesh adaptivity, and hybrid AMG and GMG solvers. This basic approach was taken by Rudi et. al. in, for example, the SC15 Gordon Bell prize paper (special category), and it provides a path for robust, fast and

scalable solvers at extreme scale for emerging architectures.

Mark Adams
Lawrence Berkeley Laboratory
mfadams@lbl.gov

Matthew G. Knepley
University at Buffalo
Department of Computer Science and Engineering
knepley@gmail.com

Tobin Isaac
Georgia Tech
tisaac@cc.gatech.edu

Johann Rudi
Argonne National Laboratory
The University of Texas at Austin
jrudi@anl.gov

MS144

Performance of Parallel Approximate Ideal Restriction Multigrid (pair) with Applications in Transport

Algebraic multigrid (AMG) methods have been widely used to solve symmetric positive definite linear systems. AMG is ideally both algorithmically scalable, where the total operations scale linearly with problem size, and scalable in parallel, where communication costs scale logarithmically with the number of processors [R. D. Falgout, Computing in Science & Engineering, vol. 8, no. 6, pp. 24-33]. Recently, a classical AMG method based on approximate ideal restriction (AIR) was developed for nonsymmetric matrices. AIR was shown to be effective for solving linear systems arising from application of upwind discontinuous Galerkin (DG) finite element to advection-diffusion problems, including the hyperbolic limit of pure advection [T. A. Manteuffel, J. W. Ruge, and B. S. Southworth, SIAM J. on Scientific Computing, submitted]. Here, parallel scaling results are presented for AIR applied to the upwind DG discretization of the steady state transport equation. Specific parallel performance characteristics of each step in the setup phase of AIR are presented, and it is shown that for typical situations, calculation of the restriction operator is the most costly step. Results from the profiling tool TAU are presented characterizing the relative costs of communications versus local computations for this step, which informs implementation improvements [S. Shende and A. D. Malony, International Journal of High Performance Computing Applications, vol. 20, no. 2, pp. 287-311].

Joshua Hanophy
Texas A&M
josh.hanophy@gmail.com

Ben S. Southworth
Department of Applied Mathematics
University of Colorado, Boulder
ben.southworth@colorado.edu

Ruipeng Li
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
li50@llnl.gov

MS144

Multilevel Convergence Theory for Multigrid Reduction in Time (MGRIT)

Parallel-in-time integration methods explore parallelism in the temporal domain and can offer a reduction in wall clock time compared to sequential time stepping algorithms. Multigrid reduction in time (MGRIT) is an $O(N)$ solver that employs a hierarchy of time grids and a parallel, iterative coarse-grid correction scheme based on multigrid reduction. In this talk, we will present a multilevel convergence theory for MGRIT that generalizes previous two-level results [Dobrev et al. (2018)]. Firstly, we will discuss necessary and sufficient conditions for convergence of two-level MGRIT and present tight upper and lower bounds that are asymptotically exact. Secondly, we will discuss error and residual propagation of a multilevel V-cycle algorithm with F- or FCF-relaxation and upper and lower bounds. Lastly, we compare MGRIT convergence for various A- and L-stable Runge-Kutta schemes for spatial eigenvalues in the complex plane.

Robert D. Falgout
Lawrence Livermore National Laboratory
rfalgout@llnl.gov

Andreas Hessenthaler
University of Stuttgart
hessenthaler@mechbau.uni-stuttgart.de

Oliver Röhrle
Institut für Mechanik (Bauwesen), Germany
roehrle@mechbau.uni-stuttgart.de

Jacob B. Schroder
Department of Mathematics and Statistics
University of New Mexico
jbschroder@unm.edu

Ben S. Southworth
Department of Applied Mathematics
University of Colorado, Boulder
ben.southworth@colorado.edu

MS144

Parallel Performance of Algebraic Multigrid with Domain Decomposition

Algebraic multigrid (AMG) is a widely used and highly scalable solver and preconditioner for large-scale linear systems resulting from the discretization of a wide class of elliptic PDEs. The need for constant communication between processors in order to perform matrix-vector multiplications during a multigrid cycle is often a significant bottleneck for AMG in parallel, however. This talk examines the design, implementation, and parallel performance of a novel algorithm designed specifically to limit communication, Algebraic Multigrid with Domain Decomposition (AMG-DD). The goal of AMG-DD is to provide a low-communication alternative to standard AMG V-cycles by enabling significantly more independent computational

work to be done between communication steps. Thus, AMG-DD is particularly well suited to computational environments where the cost of communication is high compared to the cost of computation. Parallel performance results for AMG-DD are shown for a variety of algorithm design choices and for a variety of elliptic PDE problems in 2 and 3 dimensions.

Wayne Mitchell
University of Colorado at Boulder
wayne.b.mitchell@gmail.com

MS145

A Riemannian View on Active Subspaces

Active subspaces provide an intuitive change of basis for studying differentiable functions with Euclidean domain of dimension greater than or equal to two. Namely, the active subspace represents directions in a function's domain that change the function more, on average. We wish to extend this work to understand the implications of differentiable functions defined on smooth manifolds M. In particular, when the smooth manifold is endowed with some smooth inner-product on its tangent space we seek geodesics as paths over this Riemannian manifold, (M, g), that change the differentiable function more, by an analogous notion of the average.

Zach Grey
University of Colorado, Boulder
zachary.grey@colorado.edu

Paul Constantine
University of Colorado Boulder
Department of Computer Science
pconstan@mines.edu

MS145

Reduced Models for Uncertainty Quantification in the Cardiovascular Network via Domain Decomposition

Characterizing the variability of the parameters in mathematical models of the cardiovascular system (CVS) via uncertainty quantification (UQ) requires the identification of an efficient method that guarantees enough accuracy for medical purposes. While the computational resources required by 3D models may not be easily accessible by users due privacy or time constraints, 1D models may be inaccurate in capturing blood flow anomalies in pathological conditions. This study aims at designing UQ solvers that promote parallelism and scalability, and at enhancing the accuracy of 1D models without incurring computational cost. The Domain Decomposition (DD) Uncertainty Quantification approach performs UQ at the subsystem level and propagates uncertainties through polynomial chaos coefficients via overlapping DD techniques, allowing for a reduction of the computational time. The transverse dynamics discarded by 1D models can be retrieved via educated reduced models such as the Transversally Enriched Pipe Element Method or Hierarchical Model reduction. The axial dynamics and the transverse components are solved via Finite Elements and Spectral Methods, respectively, to guarantee high accuracy at low computational cost. Preliminary results show that: The computational cost of UQ in the CVS can be drastically reduced by avoiding full-system simulations; The solver is scalable; Educated reduced models improve the accuracy of 1D solvers at approximately the same computational cost.

Sofia Guzzetti
Department of Mathematics and Computer Science
Emory University
sofia.guzzetti@emory.edu

Alonso M. Alvarez, Pablo J. Blanco
Laboratorio Nacional de Computacao Cientifica
Brasil
lalvarez@lncc.br, pablo.j.blanco@gmail.com

Kevin T. Carlberg
Sandia National Laboratories
ktcarlb@sandia.gov

Alessandro Veneziani
MathCS, Emory University, Atlanta, GA
ale@mathcs.emory.edu

MS145

A Stein Variational Newton Method for Optimal Experiment Design Problems

In this talk, we present an efficient, accurate, and scalable method based on Stein variational Newton sampling to solve Bayesian optimal experiment design (OED) problems. A critical challenge arises when we use the the double-loop Monte Carlo method to evaluate the expected information gain – the inner loop involves computation of a normalization constant as an integral of the likelihood function with respect to the prior distribution, which is typically intractable especially in high dimensions. To tackle this challenge, we propose an adaptive Stein variational Newton importance sampling method to generate samples that closely follow the posterior distribution in every step of the optimization and derive the formula to compute the densities of the generated samples to facilitate the computation of the normalized constant. The computation for the sample generation, density evaluation, and resulted optimization is scalable with respect to the parameter dimension, i.e., the total number of function evaluations depend only on the intrinsic dimension of the problem informed by the effective rank of the Hessian of the potential (log likelihood function). We use several numerical examples ranging from low-dimensional analytic models to high-dimensional PDE models, in particular Helmholtz equations for scattering problems, to demonstrate the efficiency, accuracy, and scalability of our proposed method.

Keyi Wu
University of Texas Austin
keyiwu@math.utexas.edu

Peng Chen
UT Austin
peng@ices.utexas.edu

Omar Ghattas

The University of Texas at Austin
omar@ices.utexas.edu

MS145

A Physics-informed Gaussian Process Regression Method for Data Model Convergence

In this talk, we introduce a new Gaussian process regression (GPR) method: physics informed Kriging (PhIK). In the standard data-driven Kriging, the unknown function of interest is treated as a Gaussian process with assumed stationary covariance with hyperparameters estimated from data. In PhIK, we compute the mean and covariance function from realizations of available stochastic models, e.g., from realizations of governing stochastic partial differential equations solutions. Such constructed Gaussian process is in general non-stationary, and does not assume a specific form of the covariance function. Our approach avoids the costly optimization step in data-driven GPR methods to identify the hyperparameters. More importantly, we prove that the physical constraints in the form of a deterministic linear operator are guaranteed in the resulting prediction. In addition, we provide an error estimate in preserving the physical constraints when errors are included in the stochastic model realizations. To reduce the computational cost of obtaining stochastic model realizations, we propose a multilevel Monte Carlo estimate of the mean and covariance functions. Further, we present an active learning algorithm that guides the selection of additional observation locations. The efficiency and accuracy of PhIK are demonstrated for reconstructing a partially known modified Branin function and learning a conservative tracer distribution from sparse concentration measurements.

Xiu Yang
Pacific Northwest National Laboratory
xiu.yang@pnnl.gov

Alexandre M. Tartakovsky
PNNL
alexandre.tartakovsky@pnnl.gov

Guzel Tartakovsky
Pacific Northwest National Laboratory
guzel.tartakovsky@pnnl.gov

MS146

An Adaptive Reduced Basis Anova Method for High-dimensional Bayesian Inverse Problems

In Bayesian inverse problems sampling the posterior distribution is often a challenging task when the underlying models are computationally intensive. To this end, surrogates or reduced models are often used to accelerate the computation. However, in many practical problems, the parameter of interest can be of high dimensionality, which renders standard model reduction techniques infeasible. In this work, we present an approach that employs the ANOVA decomposition method to reduce the model with respect to the unknown parameters, and the reduced basis method to reduce the model with respect to the physical parameters. Moreover, we provide an adaptive scheme within the MCMC iterations, to perform the ANOVA decomposition with respect to the posterior distribution. With numerical examples, we demonstrate that the proposed model reduction method can significantly reduce the computational cost of Bayesian inverse problems, without sacrificing much accuracy.

Qifeng Liao
ShanghaiTech University
liaoqf@shanghaitech.edu.cn

Jinglai Li
Shanghai Jiao Tong University
lijinglai@gmail.com

MS146

Fast Convolutional Encoder-decoder Networks with Quantified Uncertainty for Heterogeneous Elliptic Partial Differential Equations on Varied Domains

In this talk, we introduce a framework for constructing light-weight numerical solvers for partial differential equations using convolutional neural networks. A theoretical justification for the neural network approximation to partial differential equation solvers on varied domains is established based on the existence and properties of Greens functions. These solvers are able to effectively reduce the computational demands of traditional numerical methods into a single forward-pass of a convolutional network. The network architecture is also designed to predict pointwise Gaussian posterior distributions, with weights trained to minimize the associated negative log-likelihood of the observed solutions. This setup facilitates simultaneous training and uncertainty quantification for the networks solutions, allowing the solver to provide pointwise uncertainties for its predictions. The performance of the framework is demonstrated on three distinct classes of PDEs consisting of two linear elliptic problem setups and a nonlinear Poisson problem.

Nick Winovich, Karthik Ramani, Guang Lin
Purdue University
nwinovic@purdue.edu, ramani@purdue.edu, guanglin@purdue.edu

MS146

Uncertainty Quantification with Non-Gaussian Correlated Uncertainties, and Applications to Electronic and Photonic Circuits

Stochastic spectral methods have been used for the uncertainty analysis of vast engineering problems including semiconductor chip design influenced by process variations. However, existing stochastic spectral methods mainly assume that process variations are described by independent or Gaussian correlated random parameters. This talk presents our recent results to address the following challege: how can we quantify the impact of non-Gaussian correlated uncertainties (without using error-prone pre-processing techniques such as PCA or independent component analysis)? Our main results include: (1) numerical techniques to generate orthornormal basis functions of non-Gaussian correlated random

parameters; (2) numerical techniques and theoretical gaurantees of a projection-based stochastic collocation framework; (3) numerical techniuqes (e.g., tensor train) to address the scalability issue caused by high-dimensional random parameters. Some simulation results of electronic and photonic IC will be presented.

References:
1. C. Cui and Z. Zhang, "Stochastic collocation with non-Gaussian correlated process variations: Theory, algorithms and applications," submitted to IEEE Trans. Components, Packaging and Manufacturing Technology, (arXiv:1808.09720)
2. C. Cui and Z. Zhang, "Uncertainty quantification of electronic and photonic ICs with non-Gaussian correlated process variations," ACM Intl. Conf. Computer-Aided Design, 8 pages, San Diego, CA, Nov. 2018.

Chunfeng Cui, Zheng Zhang
University of California, Santa Barbara
chunfeng@ece.ucsb.edu, zhengzhang@ece.ucsb.edu

MS146

An Efficient Multiscale Finite Element Method for Multiscale Elliptic PDEs with Random Coefficients

We propose a generalized multiscale finite element method (GMsFEM) based on a clustering algorithm to study the elliptic PDEs with random coefficients in the multi-query setting. The new GMsFEM can be applied to multiscale SPDE starting with a relatively coarse grid, without requiring the coarsest grid to resolve the smallest-scale of the solution. The new method offers considerable savings in solving multiscale SPDEs. Numerical results are presented to demonstrate the accuracy and efficiency of the proposed method for several multiscale stochastic problems without scale separation.

Zhiwen Zhang
Department of Mathematics,
The University of Hong Kong,
zhangzw@maths.hku.hk

Eric Chung
The Chinese University of Hong Kong
Department of Mathematics
tschung@math.cuhk.edu.hk

Yalchin Efendiev
Dept of Mathematics
Texas A&M University
efendiev@math.tamu.edu

Wing Tat Leung
University of Texas at Austin
sidnet123@gmail.com

MS148

Machine Learning Enabled Suicide Prevention Research using ICU Patient Data

Recently, large amounts of Electronic Health Record (EHR) data such as MIMIC-III have been made available. They are a rich source for data-driven precision medicine research and public health policy studies. We surveyed available tools and provided best practices and strategies for EHR data preparation and data analysis. Despite some rules of thumb which have emerged through years of experience and observations, suicide prevention remains by-and-large an unsolvable problem. Deep learning based knowledge discovery models have been proven effective in learning from noisy data. However, they require large training data and strong computing power. With access to large datasets, deep learning models, and supercomputing power, scientists hope to advance research tackling this difficult problem. The EHR record for a typical hospital stay usually includes several parts: 1) structured data such as medications, prescriptions, hourly vitals, examinations, procedures, etc. 2) unstructured data which is mostly extensive notes taken by nurses, physicians, social workers, case managers, psychiatrists, etc. 3) imaging data such as X-Rays, CT, MRI, etc. 4) genomic data, and 5) demographic data and lifestyle surveys. Our overarching goal is to help a wide health care community working on EHR-enabled machine learning research.

Rafael Zamora, Shirley Wang, Cheng Ding, Shahzeb Khan, Ryan Kingery
LBL
rzamoraresendiz@lbl.gov, xiangewang@lbl.gov, chengding@lbl.gov, shahzeb@lbl.gov, rkingery@lbl.gov

Xinlian Liu
Hood College
xinlianliu@lbl.gov

Silvia N. Crivelli
Lawrence Berkeley National Laboratory
University of California, Davis
sncrivelli@lbl.gov

MS148

A Deep Learning Approach to Protein Structure Prediction

Proteins are responsible for most functions in our body. Determining proteins 3D structure is key to understanding how they work, why they cause diseases and how to design drugs to block or activate their functions. However, experimental methods are expensive and time consuming. To overcome this limitation, computational methods been applied to structural biology. Capsule Networks have great potential to tackle problems in structural biology because of their attention to hierarchical relationships. This work describes the implementation and application of a Capsule Network architecture to the classification and prediction of protein structures on GPU-based computational resources. The results show an accuracy and performance improvement compared to traditional convolutional networks, while improving interpretability through visualization of activation vectors.

Wilson Rivera
University of Puerto Rico at Mayaguez (UPRM)
wilson.riveragallego@upr.edu

Dan Rosa de Jesus, Julian Cuevas

LBNL & University of Puerto Rico at Mayaguez
dan.rosa@upr.edu, julian.cuevas1@upr.edu

Silvia N. Crivelli
Lawrence Berkeley National Laboratory
University of California, Davis
sncrivelli@lbl.gov

MS148

Computational Tools for Adaptive Resolution Molecular Dynamics Simulation: Application to Biological Molecules

Despite the increasing computational power and ongoing efforts to enhance the efficiency of molecular dynamic (MD) algorithms, all-atom (atomistic) MD simulations are often incapable to cover all the phase-space needed to understand several important chemical and biological processes, such as ligand-protein recognition, DNA-protein or protein-protein interactions, and signaling, evolving in much longer time (usually milliseconds or more). In the attempt to bridge the gap between time scales of feasible MD simulations and those of biologically relevant motions, we have developed tools to run the simulation with an adaptive resolution (AdResS) method [Praprotnik, et. al., J. Chem. Phys., vol. 126, p. 134902, 2007], in which mixed-resolution systems allow for a dynamical change in resolution of select groups of atoms - a low-resolution force field for the part of the system distant from an active site (for efficiency) and an atomistic force field for the active site and its direct environment (for accuracy), during the MD simulation. Our tools include Gaussian Processing algorithm to simplify a development of the low-resolution potentials and to produce thermodynamic force which acts on molecules to ensures the thermodynamics equilibrium of the system during the AdResS simulation. Then we have made use of an improved strategy to compute an appropriate density profile of the system. Here we present the application of new tools to biological molecules.

Masa Watanabe
St. Mary's College
masa.watanabe@stmary.edu

Branden Sanchez, Zachary Douglas
University of Saint Mary
branden.sanchez@my.stmary.edu,
zachary.douglas@my.stmary.edu

Yu-Hang Tang
Lawrence Berkeley National Laboratory
tang@lbl.gov

Bert de Jong
Lawrence Berkeley National Lab
wadejong@lbl.gov

Silvia N. Crivelli
Lawrence Berkeley National Laboratory
University of California, Davis
sncrivelli@lbl.gov

MS148

Learning From Protein Structures using Graph Convolutional Networks

As of 2018, the Protein Data Bank (PDB) contains a total of 140,824 biological structures determined through methods such as x-ray crystallography, NMR spectroscopy, and electron microscopy. In addition, the PDB provides functional annotations for each protein within the database allowing for various types of supervised learning tasks. With the ever-increasing availability of scientific data and hardware acceleration provided by GPU-based architectures, Deep Neural Networks (DNNs) are becoming a popular machine learning technique for various feature extraction, pattern recognition, and classification tasks conducted on scientific datasets. We propose a novel approach for protein structure classification which takes advantage of graph-based convolution and pooling operations to learn from descriptive graph representations of protein structures. By defining a proteins structure as a graph, explicit spatial and structural relationships between amino-acids within the protein can be expressed with greater fidelity than that of 3D volumetric CNN models. Models were trained on two classification tasks: 1) classifying tumor suppressor genes (TSGs) and proto-oncogenes (OGs) structures, and 2) classifying active and inactive kinase conformations. The experimental results demonstrate promising performance with a 0.95 AUC on the TSG/OG dataset and 0.985 AUC on the active/inactive Kinase dataset.

Rafael Zamora-Resendiz
Hood College
Lawrence Berkeley National Laboratory
rzamoraresendiz@lbl.gov

MS150

An Efficient Parallel Algorithm for Multiphysics Simulations on 3D Unstructured Meshes

General unstructured meshes are necessary for a wide range of applications with complex geometry; for example, fluid-structure interaction in hemodynamics, neutron transport in nuclear reactors, etc. However, parallel algorithms used for constructing an unstructured-mesh based simulation are challenging to design and develop on a modern supercomputer with a lot of processor cores since the related data structure is irregular and the underlying multiphysics problems become more difficult to solve. In this work, we propose an efficient parallel algorithm to address these issues. The proposed parallel algorithm consists of a hierarchal mesh partitioning method for partitioning the unstructured mesh into a large number of submeshes, an inexact Newton solver for solving the nonlinear system of equations and an efficient parallel preconditioner for calculating the Jacobian system. We numerically demonstrate that the proposed algorithm scales well up to 10,000 processor cores for applications on unstructured meshes with billion unknowns.

Fande Kong
Department of Modeling and Simulation

Idaho National Laboratory
fdkong.jd@gmail.com

Derek R. Gaston, John W. Peterson
Idaho National Laboratory
derek.gaston@inl.gov, jwpeterson@gmail.com

Cody J. Permann
Advanced Modeling and Simulation
Idaho National Laboratory
cody.permann@inl.gov

Andrew Slaughter, Alex Lindsay, Richard Martineau
Idaho National Laboratory
andrew.slaughter@inl.gov, alexander.lindsay@inl.gov, richard.martineau@inl.gov

MS150

Optimization-based PDE-constrained Discontinuity Tracking with High-order Curved Unstructured Meshes

We present a framework for generating curved moving meshes which align the element faces with discontinuities in an evolving solution. The solution comes from solving a conservation law discretized using the discontinuous Galerkin (DG) method and an arbitrary Lagrangian-Eulerian formulation. The meshes are evolved by solving an optimization problem, where a carefully chosen indicator penalizes misaligned faces. Our discontinuity indicator monotonically approaches a minimum as element faces approach the discontinuity surface, which allows for efficient gradient-based optimizers. We also include a mesh skewness measure to ensure the meshes are well-shaped. For problems with large deformations, we use local element topology changes such as edge flips on the curved elements to improve the mesh qualities. We demonstrate our methods on a number of problems with moving discontinuities, such as convection problems and flow problems with shocks.

Andrew Shi
University of California, Berkeley
andrewshi94@gmail.com

Per-Olof Persson
Dept. of Mathematics
University of California, Berkeley
persson@berkeley.edu

MS150

High-order Mesh Untangling Based on Angles

In this talk, we will discuss a heuristic for correcting tangled second- and third-order meshes. For each interior edge, our method minimizes an objective function based on the angles of the pair of triangles that share the edge. We will present several numerical examples in two dimensions with second- and third-order elements that demonstrate the capabilities of our method for untangling invalid meshes.

Mike Stees, Suzanne M. Shontz
University of Kansas
mstees@ku.edu, shontz@ku.edu

MS150

Simulation-driven Optimization of High-order Meshes by the Target-matrix Optimization Paradigm

We present a method for simulation-driven optimization of high-order curved meshes via node movement. Pointwise mesh quality metrics are defined by utilizing sub-zonal information. These metrics can measure shape, size or alignment of the region around a given point. Mesh optimization towards a chosen metric is performed by our high-order extension of the Target-Matrix Optimization Paradigm (TMOP). TMOP uses pre-defined target (or perfect) elements, which provide a way for the users to incorporate application-specific information into the optimization. The combination of targets and quality metrics is used to optimize the node positions, so that they are as close as possible to the shape/size/alignment of their targets. The construction of target-matrices is enhanced by utilizing discrete fields of interest. As these fields are defined only with respect to the initial mesh, their values on the intermediate meshes (produced during the optimization process) must be computed; we present two approaches for obtaining such values. The mesh optimization problem is posed as a global nonlinear system for all node positions. The benefits of the new methods are illustrated by applying them in high-order arbitrary Lagrangian-Eulerian simulations. Our implementation is freely available in the open-source library MFEM (http://mfem.org). LLNL-ABS-743026.

Veselin Dobrev
Lawrence Livermore National Laboratory
dobrev1@llnl.gov

Patrick Knupp
Dihedral LLC
zg37rd@gmail.com

Tzanio Kolev
Lawrence Livermore National Laboratory
kolev1@llnl.gov

Ketan Mittal
University of Illinois Urbana-Champaign
kmittal2@illinois.edu

Vladimir Z. Tomov
Lawrence Livermore National Laboratory
tomov2@llnl.gov

MS151

AcroTensor: Flexible Tensor Contractions on GPUs

As generalized linear algebra operations, tensor contractions lie at the heart of many computationally expensive fields including computer vision, machine learning, and finite element simulation. Work in these fields yields contraction equations with tensors of varying rank and dimension. Fortunately, large scale tensor contractions are suited

to GPU computation, however, to extract maximum performance requires tailoring code to each of these tensor equations and dimension combinations. In some cases the tensor dimensions are set at runtime and require dynamic sizing. In order to maximize performance under these constraints we developed AcroTensor, an open source GPU enabled C++ library utilizing Just In Time (JIT) compilation and optimization of CUDA Kernels for tensor contractions. This way "fixed" sized loops can be utilized and unrolled, and decisions can be made about the use of GPU resources like shared memory with the addition of this sizing information. We present our work on the AcroTensor library including the basic structure and usage of the library, the JIT optimizations employed, and the GPU performance tests we have completed. Special attention will be paid to the performance of AcroTensor in the Modular Finite Element Method (MFEM) library that now employs GPU enabled tensor contractions.

Aaron Fisher, Tzanio Kolev
Lawrence Livermore National Laboratory
fisher47@llnl.gov, kolev1@llnl.gov

Johann Dahm
IBM
Johann.Dahm@ibm.com

MS151

A Proposal for Next-generation BLAS

The Basic Linear Algebra Subroutines (BLAS) have enabled writing portable, high-performance linear algebra software for the past 40 years. However, with advances such as mixed precision and reproducible computations, and the introduction of additional data types such as half precision, quad precision, and double-double, the BLAS are no longer sufficient as basic building blocks. In response, we propose a Next-Generation BLAS API, with both an easy-to-use high-level interface in C++ and Fortran, and a low-level naming scheme at the library level. This new API is extensible, supporting a variety of data types, extended precision, mixed precision, and reproducible computations. It also improves error handling and propagation of nan and inf values.

Mark Gates
Computer Science Department
University of Tennessee
mgates3@icl.utk.edu

James W. Demmel
University of California
Division of Computer Science
demmel@berkeley.edu

Greg Henry
Intel Corporation
greg.henry@intel.com

Xiaoye S. Li
Lawrence Berkeley National Laboratory
xsli@lbl.gov

Jason Riedy
Georgia Institute of Technology
School of Computational Science and Engineering
jason.riedy@cc.gatech.edu

Peter Tang
Intel Corporation
2200 Mission College Blvd, Santa Clara, CA
peter.tang@intel.com

MS151

Adventures in Batched Linear Algebra in Intel® Math Kernel Library

Many numerical algorithms and machine learning applications require a large number of dense linear algebra operations that can be performed independently. As these operations are often performed on relatively small matrices, batching the operations allows exploiting the parallelism available on today's multi- and many-core systems. Batched BLAS and LAPACK subroutines can operate on either the canonical data layout or on a compact data layout that allows for cross-matrix SIMD vectorization. We consider the latest developments in the Intel® Math Kernel Library for improving performance of batched linear algebra operations on the latest Intel® architectures.

Sarah Knepper
Intel Corporation
sarah.knepper@intel.com

Peter A. Caday
Rice University
peter.caday@intel.com

Kazushige Goto, Louise Huot, Mesut Meterelliyoz, Arthur Mitrano, Shane Story
Intel Corporation
kazushige.goto@intel.com, louise.huot@intel.com, mesut.meterelliyoz@intel.com, arthur.araujo.mitrano@intel.com, shane.story@intel.com

MS151

Batched Linear Algebra in Kokkos Kernels

Many applications, such as PDE based simulations and machine learning, apply BLAS/LAPACK routines to large groups of small matrices. While existing batched BLAS APIs provide meaningful speedup for this problem type, a non-canonical data layout enabling cross-matrix vectorization may provide further significant speedup. In this talk, we present a new compact data layout that interleaves matrices in blocks according to the SIMD vector length. We combine this compact data layout with a new interface to BLAS/LAPACK routines that can be used within a hierarchical parallel application. Our layout provides significant speedups against OpenMP loops around optimized BLAS/LAPACK kernels. The compact batched BLAS/LAPACK implementations are available in the KokkosKernels library.

Siva Rajamanickam
Sandia National Laboratories

srajama@sandia.gov

MS152

Effects of Subscale Fluctuations on Predictions of Hot Spot Formation in Energetic Materials

A multiscale stochastic approach is used to understand thermal localization, known as hotspots, in a hyperbolic two-phase mixture model of energetic materials under dynamic compaction. A hybrid reduced-order heat model coupled with a one-dimensional granular compaction model is developed and implemented to characterize the effect of microstructure heterogeneity on thermal fluctuations. The variance in porosity is shown to correlate with the pore surface temperature at the shock front. Numerical and analytical techniques to solve the problem are explored. We first use a multilevel Monte Carlo (MLMC) method for finding the probability distribution of reaction initiation due to pressure and temperature fluctuations. Then we derive an equation for the joint probability density function of key problem variables contributing to the formation of hotspots. Using the results from our MLMC simulations, we employ machine learning methods to approximate closure terms in the PDF equation.

Joseph Bakarji
University of California, San Diego
jbakarji@eng.ucsd.edu

MS152

Hydrodynamic Fluctuations in Quasi-two Dimensional Diffusion

We study diffusion of colloids on a fluid-fluid interface using particle simulations and fluctuating hydrodynamics. Diffusion on a two-dimensional interface with three-dimensional hydrodynamics is known to be anomalous, with the collective diffusion coefficient diverging like the inverse of the wavenumber. This unusual collective effect arises because of the compressibility of the fluid flow in the plane of the interface, and leads to a nonlinear nonlocal convolution term in the diffusion equation for the ensemble-averaged concentration. We study the magnitude and dynamics of density and color density fluctuations using a novel Brownian dynamics algorithm, as well as fluctuating hydrodynamics theory and simulation. WWe also examine nonequilibrium fluctuations in systems with two-dimensional hydrodynamics, such as thin smectic films in vacuum. We find that nonequilibrium fluctuations are colossal and comparable in magnitude to the mean, and can be accurately modeled using numerical solvers for the nonlinear equations of fluctuating hydrodynamics.

Aleksandar Donev
Courant Institute of Mathematical Sciences
New York University
donev@courant.nyu.edu

R. Delgado-Buscalioni
Universidad Autónoma de Madrid
r.delgado-buscalioni@uam.es

MS152

Towards an Irreversible Thermodynamics for Active Systems

Active matter systems are composed of constituent elements capable of self-directed motion. They provide an opportunity to revisit the notions of statistical mechanics and condensed matter physics from a fresh nonequilibrium perspective. In particular they allow us to ask how stress, pressure, and transport phenomena are affected due to the non-equilibrium driving forces. One important question is whether the formalism of traditional thermodynamics can be extended to active matter systems. In this talk, I will discuss our recent efforts towards addressing this question by working with stochastic models of active particles that contain different types of driving forces, including convective and rotary. In particular, I will discuss derivations concerning the continuum balance equations for these models including balances of mass, linear momentum, angular momentum and energy. We show how the different microscopic active forces, which break time reversal symmetry, manifest at the collective level. We further discuss conditions under which bulk flows or currents can be supported.

Katherine Klymko
Lawrence Berkeley National Labs and Berkeley University
Berkeley, CA
kek2134@berkeley.edu

MS152

Fluctuating Hydrodynamics of Janus Particles Assembled via Long-range Attractive Hydrophobic Interactions

In macroscopic models, the well-known Helfrich Hamiltonian membrane model has been extensively used to capture macroscopic physical properties of a lipid bilayer membrane. Some phenomena such as membrane fusion and micelle formation are, however, challenging to be described using a macroscopic framework, and including all the molecular details has its challenges, from a numerical simulation perspective. Therefore, in order to include the salient molecular details in a coarse-grained manner, we study the dynamics of lipid bilayer membrane using Janus-type particle configurations to represent collections of lipids. These coarse-grained lipid molecules interact through an action field that measures water activity due to the presence of nearby hydrophobic surfaces, leading to specific boundary conditions on each Janus particle. We adopt two numerical frameworks to investigate the particle dynamics: (1) the finite element method (FEM); (2) the quadrature by expansion method (QBX). We also examine the numerical accuracy and qualitative comparisons for large system simulations.

Yuan-Nan Young
New Jersey Institute of Technology
yuannan.young@gmail.com

MS153

Transported Snapshot Model Order Reduction for

Parametric, Steady-state Fluid Flows Containing Parameter Dependent Shocks

A new model order reduction approach is proposed for parametric steady-state nonlinear fluid flows characterized by shocks and discontinuities whose spatial locations and orientations are strongly parameter dependent. In this method, solutions in the predictive regime are approximated using a linear superposition of parameter dependent basis. The sought after parametric reduced-basis are obtained by transporting the snapshots in a spatially and parametrically dependent transport field. Key to the proposed approach is the observation that the transport fields are typically smooth and continuous, despite the solution themselves not being so. As a result, the transport fields can be accurately expressed using a low-order polynomial expansion. Similar to traditional projection-based model order reduction approaches, the proposed method is formulated mathematically as a residual minimization problem for the generalized coordinates. The proposed approach is also integrated with well-known hyper-reduction strategies to obtain significant computational speed-ups. The method is successfully applied to the reduction of a parametric 1-D flow in a converging-diverging nozzle, a parametric 2-D supersonic flow over a forward facing step and a parametric 2-D jet diffusion flame in a combustor.

Nirmal Nair, Maciej Balajewicz
University of Illinois at Urbana-Champaign
njn2@illinois.edu, mbalajew@illinois.edu

MS153

Model Reduction of Multi-dimensional Hyperbolic Conservation Laws

We propose methods for model reduction of multi-dimensional nonlinear hyperbolic systems of conservation laws, by building on the previous works of the authors [Rim and Mandli, *Preprint* (2018)]. Solutions to parametrized hyperbolic problems do not yield a low-rank structure in the usual sense, since their energy is localized with respect to time or parameter, and consequently this has hindered the success of many popular methods. Our method exploits the low-rank structure in the transport map between the solutions, and makes use of the Approximate Discrete Radon transform (ADRT) as a dimensional splitting tool to transform a multi-dimensional wave into a family of single dimensional waves.

Donsub Rim
Columbia University
dr2965@columbia.edu

Kyle T. Mandli
Columbia University
Applied Physics and Applied Mathematics
kyle.mandli@columbia.edu

MS153

Interpolation of Solutions of Parametric Hyperbolic PDEs with Changing Shock Structure

Despite significant progress in parametric and stochastic PDEs, solutions with jumps, as e.g. in the hyperbolic case, still pose severe difficulties for their numerical approximation. In order to handle this problem, in recent years, several methods have emerged that use some variants of transforms or shifts to align jumps and achieve significant improvements in approximation error versus required degrees of freedom. Nonetheless, changes in the jump set topology, such as emerging or colliding jumps, cannot be handled by simple transforms alone. To this end, we interpolate the solution by transformed snapshots with "aligned' jumps, but somewhat similar to ENO schemes, only use those which share the same topological structure as the solution we wish to reconstruct. This entails that we need a good and efficient prediction of the topological structure of an unknown solution, or rather the set of active snapshots. Intuitively, at a structural change, the transform maps must develop some kind of singularity. We demonstrate a rigorous version of this intuition, which we use to locate the topology changes. In addition, we present numerical methods that deal with the arising practical difficulties such as learning the singular transforms form snapshots alone and their efficient approximation.

Gerrit Welper
Texas A&M University
gerrit.welper@ucf.edu

MS153

Sparsity Promoting Acceleration for Reduced Order Modeling of Multi-scale Transport Problems

New variants of sparse acceleration is proposed for projection-based reduced order models of flows involving turbulence and combustion. Separate empirical interpolation-based sparse sampling is applied on the conserved variables, flux, and source terms individually to improve both the accuracy and efficiency over the traditional acceleration methods. The benefits are more apparent when sharp gradients and instabilities are involved in the flow dynamics. Optimality of the proposed method is analyzed in detail and evaluated in the case of both static and adaptive basis choices. Numerical tests are conducted on rocket combustor models and robustness and efficiency of the resulting reduced order model is assessed.

Jiayang Xu
University of Michigan Ann Arbor
davidxu@umich.edu

Cheng Huang
University of Michigan Ann Arbor
Department of Aerospace Engineering
huangche@umich.edu

Karthik Duraisamy
University of Michigan Ann Arbor
kdur@umich.edu

MS154

Convex Limiting

A new second-order method for approximating the compressible Euler equations is introduced. The method preserves all the known invariant domains of the Euler system:

positivity of the density, positivity of the internal energy and the local minimum principle on the specific entropy. The technique combines a first-order, invariant domain preserving, guaranteed maximum speed method using a graph viscosity (GMS-GV1) with an invariant domain violating, but entropy consistent, high-order method. Invariant domain preserving auxiliary states, naturally produced by the GMS-GV1 method, are used to define local bounds for the high-order method which is then made invariant domain preserving via a *convex limiting* process. Numerical tests confirm the second-order accuracy of the new GMS-GV2 method in the maximum norm, where 2 stands for second-order. The proposed convex limiting is generic and can be applied to other approximation techniques and other hyperbolic systems.

Jean-Luc Guermond
Department of Mathematics
Texas A&M, USA
guermond@math.tamu.edu

Ignacio Thomas
Sandia National Laboratories
itomas@tamu.edu

Bojan Popov
Department of Mathematics
Texas A&M University
popov@math.tamu.edu

Murtazo Nazarov
Uppsala University
murtazo@csc.kth.se

MS154

Asymptotic Preserving Schemes for Euler Equations of Gas Dynamics and Shallow Water Equations with Coriolis Forces

We are interested in deriving robust asymptotic preserving (AP) numerical methods for Euler equations of gas dynamics and shallow water equations with Coriolis forces. It is well-known that both the Euler and shallow water equations become stiff in low Mach and low Froude number regimes, respectively. In these regimes, the applicability of explicit schemes is limited due to severe stability limitation on the mesh size: $\Delta x \sim \epsilon$ and $\Delta \sim \epsilon^2$, where ϵ is the Mach/Froude number. In order to design substantially more efficient numerical methods, one typically needs to design an AP scheme, which approximates the incompressible equation obtained in the limiting $\epsilon \to 0$ case. Such schemes can be designed using either a flux splitting implicit-explicit (IMEX) or fully implicit approach. In this talk, I will discuss both approaches and present several different AP schemes.

Alexander Kurganov
Southern University of Science and Technology
kurganov@math.tulane.edu

MS154

Second Order Unconditional Positivity-preserving Schemes for Poisson-Nernst-Plank Equations

The Poisson-Nernst-Planck (PNP) system is a widely accepted model for simulation of ionic channels. In this paper, we design, analyze, and numerically validate a second order unconditional positivity-preserving scheme for solving a reduced PNP system, which can well approximate the 3D ion channel problem. Positivity of numerical solutions is proven to hold true independent of the size of time steps and the choice of the Poisson solver. The scheme is easy to implement without resorting to any iteration method. Several numerical examples further confirm the positivity-preserving property and demonstrate the accuracy, efficiency, and robustness of the proposed scheme, as well as the fast approach to steady states.

Hailiang Liu
Iowa State University
USA
hliu@iastate.edu

Wumaier Maimaitiyiming
Iowa State University
wumaierm@iastate.edu

MS155

Large Ensemble Modelling of Greenland's Contribution to Sea Level Rise

The Greenland Ice Sheet holds around 7.2 m of sea-level equivalent. In recent decades rising atmosphere and ocean temperatures have led to an acceleration in mass loss, adding an average of about 0.5 mm per year to global mean sea-level between 1991 and 2015. Current ice margin recession in Greenland is led by the retreat of outlet glaciers, but only recently have high-resolution ice thickness maps became available that allow prognostic simulation of complex flow patterns. Here we run an ensemble of outlet glacier-resolving ice sheet models with model parameters systematically distributed using latin hypercube sampling to estimate Greenland's contribution to sea-level over the next millennium under different climate forcings. This ensemble approach allows us to robustly quantify mass loss uncertainty, which shows that impacts from outlet glaciers are important for the next few centuries, and that Greenland will very likely become ice-free within a millennium without significant reductions in greenhouse gas emissions. Furthermore, we use Sobol indices to compute the relative variance in mass loss estimates attributable to variance in model parameters. We find that over the next century, better estimates of parameters governing both climate-mass balance coupling and ice dynamics will substantially improve ice sheet forecasts. In the longer term, uncertainty in climate models and climate-mass balance coupling dominate the variability in simulated ice sheet response.

Andy Aschwanden, Mark Fahnestock, Martin Truffer
University of Alaska Fairbanks
aaschwanden@alaska.edu, mfahnestock@alaska.edu, mtruffer2@alaska.edu

Douglas Brinkerhoff
University of Montana
douglas1.brinkerhoff@umontana.edu

Regine Hock, Constantine Khroulev
University of Alaska Fairbanks
rehock@alaska.edu, ckhroulev@alaska.edu

Ruth Mottram
Danish Meteorological Institute
rum@dmi.dk

Abbas Khan
DTU Space
abbas@space.dtu.dk

MS155

Reconciling Observations of Driving Stress with Observations of Velocity

High quality, gridded observations with low errors now exists for the surface elevation and velocity of the large ice sheets in Greenland and Antarctica. With such measurements, we can return to an older piece of glaciological theory suggesting the driving stress reflects an averaging of the surface slope over several ice thicknesses. Performing this averaging, I find that there are significant differences in the observed velocity and driving stress directions. Specifically, for the regions of the ice sheet having a speed of greater than 50 meters per year, where reported errors are lowest, the difference in directions averages 11.6 degrees. Theory suggests the observed flow direction and driving stress differ due to the action of longitudinal and transverse stresses. Here, I attempt to understand the differences in four ways. First, I consider that the differences are due to changes in the ice surface elevation that followed or preceded the surface velocity observations. Second, the strain heating that arises from the stress fields corresponding to velocity observations is used to test the plausibility of the observations. Third, the tortured flow around basal obstructions is explored as an explanation. Finally, a map-plane, hybrid momentum balance model is used to assimilate the observations, positing that directional differences can not be resolved with the inverse model suggest an inconsistency with the observed physics.

Jesse Johnson
Dept of Computer Science
University of Montana
johnson@cs.umt.edu

MS155

Modeling Thermally Activated Sliding in Ice Sheet Flow

Thermo-mechanical feedbacks are a key driver of patterning in the velocity field of ice sheets, and must therefore be included in ice sheet models used to predict the evolution of fast flowing regions. In this context, our work is concerned with thermally activated sliding, and in particular with deriving from first principles the simplest ice flow model that captures this process. In the first part of this talk, we consider the paradigmatic case of sliding switching on where bed temperature (T_{bed}) reaches the melting point. Through a boundary layer analysis we show that this set up localizes the acceleration from no to finite sliding over a distance comparable with the ice thickness, and that such localization leads to refreezing. We thus conclude that a viable transition has to span a distance asymptotically larger than the ice thickness, which is possible only if regelation and premelting are accounted for via a temperature dependent sliding law. In the second part of this talk we derive and analyze an asymptotically reduced model for ice flow with temperature dependent sliding. Our results illustrate that (i) small deviations of T_{bed} from the melting point must be resolved, and (ii) that the subtemperate region is unstable with structure emerging at the ice thickness scale. We conclude by discussing how multiple scale expansions could be used in some cases to parameterize this small-scale structure in ice sheet models that resolve only the ice sheet scale.

Elisa Mantelli
Stanford University
emantell@stanford.edu

Christian Schoof
Dept. of Earth and Ocean Sciences
University of British Columbia
cschoof@eos.ubc.ca

MS155

The Case for Large Ice Sheet Model Ensembles from Theory and Practice

Predictions of future marine ice sheet evolution are subject to considerable uncertainty due to poorly understood processes and internal variability of the climate system. We argue here that large ensembles are necessary to quantify the amount and structure of uncertainty in ice sheet projections. We start with a stochastic perturbation analysis [Moon and Wettlaufer 2013] of a simple prognostic model of grounding line migration [Robel et al. 2018], to derive analytical approximations for the second and third moments of the ensemble distribution of ice sheet projections. The evolving ensemble distribution indicates that the marine ice sheet instability widens and skews the distribution of uncertainty in ice sheet predictions, at a rate which depends on bed topography, nonlinearity in grounding line flux, and the temporal persistence of climate variability. We test the intuition gained from this mathematical analysis by calculating large (∼500 member) ensembles of the future evolution of Thwaites Glacier in West Antarctica using the Ice Sheet System Model [Larour et al. 2012], with stochastic ocean forcing. We discuss how the degrees of freedom in the stochastic forcing and parameter uncertainty determines the necessary number of ensemble members for converged distributions of future ensemble projections. We conclude by discussing how modern uncertainty quantification methods may be used to reduce the computational burden of ensemble ice sheet modeling.

Alexander Robel
California Institute of Technology
robel@eas.gatech.edu

Helene Seroussi
Jet Propulsion Laboratory
helene.seroussi@jpl.nasa.gov

Gerard Roe

University of Washington
groe@uw.edu

MS156

Galerkin Difference Approximations for Biharmonic Equations

We discuss the Galerkin Difference Difference Spline method in providing high order accurate, energy stable schemes for biharmonic PDE systems related to beam and plate motion. The method uses a *Galerkin Difference* (GD) framework where the basis functions are constrained to be C^1 continuous. This smoothness is obtained by requiring the underlying polynomial interpolant not only match both the discrete data associated with element corners but also finite difference approximations to the derivative, a so-called *Difference Spline* (D-Spline). Computational results verifying high-order accuracy for the dynamic behavior of the 1D Euler–Bernoili Beam and the 2D Kirchoff–Love Plate are presented.

John Jacangelo
Rensselaer Polytechnic Institute
jacanj@rpi.edu

MS156

Fast, Higher-order Direct/Iterative Hybrid Solver for Scattering by Inhomogeneous Media with Application to High-frequency and Discontinuous Refractivity Problems

A fast high-order method for the solution of two-dimensional problems of scattering by penetrable inhomogeneous media will be presented, with application to high-frequency configurations containing a (possibly) discontinuous refractivity. The method relies on a combination of a differential volumetric formulation and a boundary integral formulation. Thus, in the proposed method the entire computational domain is partitioned into large numbers of volumetric spectral approximation patches, which are then grouped into sub-domains consisting of adequately-chosen groups of patches, and, finally, an overarching integral equation formulation on the overall domain boundary. The resulting algorithm can be quite effective: after a modestly-demanding precomputation stage (whose results for a given frequency can be repeatedly used for arbitrarily chosen incidence angles), the proposed algorithm can accurately evaluate scattering by very large objects, and with very high contrasts in the refractive index variations (including possibly refractive-index discontinuities), in single-core computing times of a few seconds.

Ambuj Pandey, OSCAR P Bruno
Caltech
ambuj@caltech.edu, obruno@caltech.edu

MS156

Mesh Refinement Interfaces with Hanging Nodes for the Elastic Wave Equation

Summation by parts (SBP) operators have traditionally been derived without ghost points. Boundary conditions can then be imposed by the projection method, or by adding penalty (SAT) terms to the differential equation. For wave equations in second order formulation, an alternative way of enforcing boundary conditions is to modify the SBP operator to use ghost points. They provide an extra degree of freedom that allows the boundary condition to be enforced strongly. There is a simple recipe for adding ghost points to the traditional SBP operators, such that energy stability can be maintained. The ghost point approach generalizes to discretizations on piecewise structured meshes with hanging nodes along the refinement interface. In this talk we present a novel technique that combines SBP operators with and without ghost points to enforce the interface conditions for the elastic wave equation on curvilinear meshes. The method is energy conserving and stable under essentially the same time step restriction as for the spatially periodic problem. The proposed tecnique is currently being implemented in the wave propagation code SW4, which models ground motion due to earthquakes, or other seismic activities.[1]

N. Anders Petersson
Lawrence Livermore National Laboratory
petersson1@llnl.gov

Siyang Wang
Chalmers University of Technology
University of Gothenburg
siyang.wang@chalmers.se

Bjorn Sjogreen
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
sjogreen2@llnl.gov

MS156

High Order Modeling of Visco-elastic Attenuation in Seismic Wave Computation

In earthquake simulations, a common model is to consider rocks and sediments underground as elastic materials, and propagate the seismic waves according to the elastic wave equation. However, some parts of the underground, especially softer sediments near the surface do not behave as perfectly elastic materials. For such materials, a visco-elastic model is more appropriate. The visco-elasticity introduces damping of the waves. A standard assumption is that the damping (Q-factors) does not depend on the frequency of the waves. This leads to a non-local visco-elastic model, which is usually approximated by polynomials in Fourier space. Inverting to the time domain gives a number of auxiliary variables that satisfy ordinary differential equations (ODEs) that are coupled to the elastic wave equation. In previous work, we developed a second order accurate finite difference discretization of the attenuation model, that was used together with a second order accurate discretization of the elastic wave equation. For this coupled discretization, we were able to prove fully discrete (both space and time) energy norm stability. In more recent work, we developed the fourth order accurate seismic simulation code SW4. The attenuation model in SW4

[1]This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344.

is a fourth order accurate discretization, that generalizes the previous second order accurate model. In this talk we describe the fourth order accurate attenuation model, and illustrate its performance by numerical examples.

Bjorn Sjogreen
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
sjogreen2@llnl.gov

N. Anders Petersson
Lawrence Livermore National Laboratory
petersson1@llnl.gov

MS157

Nonlocal Operators in Dynamics and Control: Theories and Applications

In this talk, a new stability criterion, called fractional Chebyshev collocation (FCC), is presented to study the stability of linear dynamical systems with nonlocal operators such as delays, fractional operators, and periodic coefficients. The FCC stability criterion can be applied to broad classes of linear dynamical systems in comparison to the current stability criteria that are particularly limited to a few classes of linear dynamical systems. It can be used for examining the stability of linear dynamical systems with several nonlocal operators. As a result of this fundamental study, some open problems of this field were solved. In addition, the FCC framework is developed for obtaining a better control performance for linear dynamical systems with a large degree of freedom. The FCC framework can be applied to a broad range of applications such as flexible robots, consensus control, and autonomous vehicles.

Arman Dabiri
Eastern Michigan University
armandabiri@email.arizona.edu

MS157

The Zolotarev Fractional Derivative: Modeling Anomalous Diffusion and Dissipative Wave Propagation

Anomalous diffusion characterized by a linear spreading rate with respect to time scales like pure advection. This linear scaling of solutions with respect to time coupled with skewness is observed in many applications, including hydrology, nuclear physics, and viscoelasticity. In this talk, we derive a governing equation for anomalous diffusion with linear scaling and arbitrary skewness by utilizing a new operator called the Zolotarev fractional derivative, which is valid for all orders between 0 and 2. The Zolotarev derivative is defined using the Fourier symbol related to the characteristic function of a stable random variable in the Zolotarev M parameterization. The Zolotarev derivative is also used to model wave propagation in dissipative media where the attenuation coefficient follows a power law with respect to frequency. We prove that this operator is continuous with respect to order and present generator, Caputo, and Riemann-Liouville forms. As an application, we propose a continuous power law wave equation that models attenuation coefficients that scale linearly with respect to frequency. The solution of diffusion equations utilizing the Zolotarev derivative with an impulsive initial condition and Green's functions of the continuous power law wave equation are shifted and scaled stable densities in the M parameterization. The Zolotarev derivative is discretized using a Grunwald-Letnikov formula derived from the theory of semi-fractional derivatives.

James F. Kelly
Dept. of Statistics & Probability, Michigan State Universi
kellyja8@stt.msu.edu

Mark Meerschaert
Dept. of Statistics and Probability
Michigan State University, USA
mcubed@stt.msu.edu

Robert J. McGough2
Michigan State University
mcgough@egr.msu.edu

MS157

Numerical Techniques to Approximate Fractional-order Nonlinear Viscoelasticity in Soft Elastomers

Dielectric elastomers are employed for a variety of adaptive structures. Many of these soft elastomers exhibit significant rate-dependencies in their response. Accurately quantifying this viscoelastic behavior is non-trivial and in many instances a non-linear modeling framework is required. Fractional-order operators have been applied to modeling viscoelastic behavior for several years, and recent research has shown fractional-order methods to be effective for nonlinear frameworks. This implementation can become computationally expensive to achieve an accurate approximation of the fractional-order derivative. In this presentation, we demonstrate the effectiveness of using quadrature techniques in approximating the Riemann–Liouville definition for fractional derivatives in the context of developing a non-linear viscoelastic model.

Paul Miles, Graham Pash, Ralph Smith
North Carolina State University
prmiles@ncsu.edu, gtpash@ncsu.edu, rsmith@ncsu.edu

William Oates
Department of Mechanical Engineering, Florida State Universi
woates@eng.fsu.edu

MS157

Fractional Modeling of Visco-elasto-plastic Materials

In this talk, we present multiple time-fractional (constitutive) models for visco-elastic and visco-elasto-plastic materials. We make use of fractional Scott-Blair elements instead of standard Hookean springs and Newtonian dashpots, which introduces power-law stress-strain behaviors in both visco-elasticity and visco-plasticity. The thermodynamic consistency of the models is verified using fractional free-energy potentials with the classical Clausius-Duhem

inequality. We develop proper numerical methods for long and accurate time-integration through the efficient computation of the history load and the use of fractional return-mapping algorithms. Finally, we analyze these models in the context of structural analysis by incorporating the resulting constitutive laws in a finite element framework.

Jorge L. Suzuki, Mohsen Zayernouri
Michigan State University
suzukijo@egr.msu.edu, zayern@egr.msu.edu

MS158

Implementation of a State-based Peridynamic Analysis in MOOSE Framework

This study presents an ordinary state-based peridynamics (PD) element for modeling structures under mechanical and thermal loads in the presence of viscoelastic and creep behavior. The behavior of the viscous material is modeled in terms of Prony series. The constitutive constants are the same as those necessary in the classical history-integral model, and they are also readily available from relaxation tests. The PD element enables the peridynamic interactions among the finite element nodes in a non-uniform discretization with a variable horizon. As part of the Multiphysics Object-Oriented Simulation Environment (MOOSE) framework, this element has been used to model the fracture and thermoelastic response of elastic materials. In the MOOSE implementation, the line elements are used to represent the interactions among the nodes. The ultimate goal of this work is to model the fracture behavior in metallic nuclear fuels, which experience significant nonlinear creep and swelling before fracture. Unlike the previous PD models, this PD element does not require any surface and volume corrections. Also, it invokes the engineering material constants directly without any calibration procedure. As in previous PD models, crack nucleation and growth can be based on the local stress or stretch based criteria. This model has been verified against benchmark solutions in elasticity and heat conduction, and its damage prediction capability against experimental observations.

Deepak K. Behera
University of Arizona
deepak7@email.arizona.edu

Hailong Chen
Idaho National Laboratory
University of Kentucky
hailong.chen@uky.edu

Benjamin Spencer
Idaho National Laboratory
benjamin.spencer@inl.gov

Erdogan Madenci
The University of Arizona
madenci@email.arizona.edu

MS158

Simulation of Molten Salt Reactors with Moltres

The Advanced Reactors and Fuel Cycles (ARFC) group models and simulates the design, safety, and performance of advanced nuclear reactors. For these simulations coupling between physics such as neutron transport, thermal-hydraulics phenomena, and fuel performance must be taken into account. Our group performs high fidelity simulation of Gen IV reactor designs through development of models and tools for representing unique materials, complex geometries, and physical phenomena. Current work introduces an extension of the MOOSE framework, Moltres, to appropriately model coupled thermal-hydraulics and neutronics of promising liquid-fueled Molten Salt Reactor designs. Initial simulations of the Molten Salt Reactor Experiment (MSRE) have been conducted on Blue Waters supercomputer with deterministic multiphysics. Steady state, transient, and fuel cycle analysis simulations have been run in 2D as well as 3D and compared against the Molten Salt Reactor Experiment. These simulations have occupied up to many hundreds of nodes simultaneously and have resulted in rich datasets for use in reactor design and analysis. This talk will describe how coupling between neutronics and thermal hydraulics have been established in the Moltres as well as the validation and verification efforts which have been completed.

Andrei Rykhlevskii
University of Illinois at Urbana-Champaign
andreir2@illinois.edu

Alexander Lindsay
Idaho National Laboratory
alexlindsay239@gmail.com

Kathryn Huff
University of Illinois at Urbana–Champaign
kdhuff@illinois.edu

MS158

Multiphysics Analysis of Pebble Bed Reactors in MOOSE

Pebble Bed Reactors (PBRs) are a candidate for the next generation nuclear power plant, because they exhibit several desirable features: passive safety, high thermal efficiency, low excess reactivity. However, their mode of operation poses unique challenges to modeling and simulation. Typical PBRs consist of roughly 400,000 graphite pebbles (D=6 cm) containing TRISO nuclear fuel particles that are 1 mm in diameter. The pebbles move through the cylindrical vessel measuring roughly 3 m in diameter and 10 m in height within several hundred days. Many design challenges of PBRs directly relate to the multi-physics interaction of the pebble flow with temperature, coolant flow, and reactor power: dust production and transport, earthquakes. We present an approach of coupling the discrete element method simulating pebble motion with the FEM-based MOOSE framework under the umbrella of the Magpie (Mesoscale Atomistic Glue Program for Integrated Execution) code. The MOOSE herd already contains capability for reactor-physical (MAMMOTH) and fluid-dynamic (Pronghorn) simulation of PBR cores. Pebble interaction is simulated using LAMMPS with the granular particle package. We demonstrate the capability on a PBR design inspired by Kairos power, where pebbles are suspended in molten fluoride salt floating from the bottom to the top of

the vessel due to the pebbles buoyancy. By coupled simulations, we explore the behavior of the system in response to earthquakes.

Sebastian Schunert
Idaho National Lab
sebastian.schunert@inl.gov

Daniel Schwen
Idaho National Laboratory
daniel.schwen@inl.gov

MS158

Parallel Phasefield Simulation with MOOSE and Marmot

Materials in nuclear reactors are subject to extreme driving forces, such as radiation damage, thermal and mechanical loading, and chemical changes. These driving forces influence the evolution microstructural features in the materials. These features include solid precipitates, gas bubbles, the materials grain size and shape distributions, and phase boundaries. Material microstructure to a large extent determines material properties and its time evolution needs to be understood to predict changes in the material properties. In nuclear reactors these properties determine the operational performance and safety margins. The variety of driving forces experienced by materials in the nuclear fuel cycle pose a challenging Multiphysics problem, which we approach using INLs finite element Multiphysics Object-Oriented Simulation Environment (MOOSE). We will demonstrate microstructural evolution using tightly coupled phase field, heat conduction, and mechanics solves for strain periodic representative volume models suitable for simulating effective material properties in complex composite materials.

Daniel Schwen, Larry K. Aagesen
Idaho National Laboratory
daniel.schwen@inl.gov, larry.aagesen@inl.gov

MS159

An Analysis of Parameter Sensitivity in Continuous Data Assimilation of 2D Incompressible Navier-Stokes Equations

One of the challenges of the accurate simulation of turbulent flows is that initial data is often incomplete. Data assimilation circumvents this issue by continually incorporating the observed data into the model. Recently a new approach to data assimilation known as the Azouani-Olson-Titi (AOT) algorithm introduced a feedback control term to the 2D incompressible Navier-Stokes equations in order to incorporate sparse measurements. It was proven that the solution to the AOT algorithm for continuous data assimilation converges exponentially to the true solution of 2D incompressible Navier-Stokes equations with respect to the given initial data. In this talk we examine the parameter sensitivity of viscosity. Specifically, we analyze how perturbations of viscosity affect the convergence of the AOT algorithm to the 2D incompressible Navier-Stokes equations. Analytical and computational results will be presented.

Elizabeth Carlson
University of Nebraska
elizabeth.carlson@huskers.unl.edu

MS159

Estimating Key Parameters for a Fracture Propagation Model using Data Assimilation

Abstract not available.

Humberto C. Godinez
Los Alamos National Laboratory
Applied Mathematics and Plasma Physics
hgodinez@lanl.gov

MS159

Model Reduction using Graph Networks

Reduced modeling is becoming increasingly necessary to yielding fast results for complicated, expensive models. Although many models rely on a similar modeling methodology for reduction, ie taking a matrix system and reducing it while maintaining of the same properties as the larger system, there are other avenues. A large-scale computational model for transport was developed at Los Alamos National Laboratory that uses discretized meshes for the forward solution. A reduced model was created using graph networks to simplify the problem. We propose to model the uncertainty in quantities of interest, such as pressure, by using the graph model. We propose to test our uncertainties against well-established results of the high-fidelity model as proof of trust for use in larger and unknown transport problems.

Jaime A. Lopez-Merizalde
Tulane University
jlopez8@tulane.edu

MS159

Continuous Data Assimilation from Scattered Spatial Observations in Time Dependent PDEs

We introduce one new continuous data assimilation algorithm based on the AOT data assimilation algorithm, which is using the feedback control from the PDE and weighted least square interpolation. We have tested our algorithm on a number of test problems, including 1D KPP-Burgers' equation, 1D Kuramoto-Sivashinsky equation, and 2D shallow water equations.

Tong Wu, James Hyman
Tulane University
twu2@tulane.edu, mhyman@tulane.edu

Humberto C. Godinez, Vitaliy Gyrya
Los Alamos National Laboratory
Applied Mathematics and Plasma Physics
hgodinez@lanl.gov, vitaliy_gyrya@lanl.gov

MS160

Multilevel Multifidelity Approaches for Uncer-

tainty Quantification in Cardiovascular Modeling

For hemodynamic cardiovascular models, standard approaches for uncertainty quantification face significant challenges due to the large number of uncertain inputs and high computational cost of realistic three-dimensional simulations. We propose an approach utilizing multilevel multifidelity (MLMF) estimators to improve the accuracy of our quantities of interest (QoIs) while maintaining reasonable computational cost. This is achieved by leveraging three model fidelities, each with varying resolution levels, to rigorously quantify the variability in hemodynamic outputs. We demonstrate this framework on healthy and diseased models of aortic and coronary anatomies. Uncertainties in parameters for the material properties of the vessels, the fluid, and the boundary conditions are investigated. The performance of the MLMF estimators is measured by the computational cost to obtain a chosen accuracy for global and local QoIs. We see significant reduction in the total computational cost necessary for all QoIs as compared to traditional uncertainty quantification methods. As expected, global quantities such as outlet pressure and flow show larger reductions than local quantities, such as those relating to wall shear stress, as the latter rely more heavily on the highest fidelity model evaluations. Similarly, healthy models show larger reductions than diseased models. In all cases, MLMF estimators make uncertainty quantification feasible for constrained computational budgets.

Casey M. Fleeter
Stanford University
cfleeter@stanford.edu

Gianluca Geraci
Sandia National Laboratories, NM
ggeraci@sandia.gov

Daniele E. Schiavazzi
University of Notre Dame
dschiavazzi@nd.edu

Andrew Kahn
UCSD
akahn@ucsd.edu

Alison Marsden
School of Medicine
Stanford University, USA
amarsden@stanford.edu

MS160

Data Assimilation on Lumped Parameter Models for Diastolic Heart Failure

It is estimated that about one-third of all deaths due to heart failure are caused by diastolic ventricular dysfunction. This pathology affects the ability of the left ventricle to relax during filling. However, it has no effect on indicators such as the systemic blood pressures or cardiac output, which makes this condition difficult to diagnose. In this context, a reliable predictor of diastolic left ventricular dysfunction is secondary pulmonary hypertension, i.e., an abnormal but reversible increase of the pulmonary arterial pressure which can be measured only though invasive right heart catheterization. In this study, in lieu of an invasive procedure, we leverage the hemodynamic consistency of a differential circulation model to predict the pulmonary arterial pressure in adults from a collection of non-invasive clinical data. This is a numerical exercise involving various aspects of forward and inverse uncertainty analysis. Specifically, we investigate the physiological admissibility of the model under healthy and heart failure conditions as well as the parameter identifiability and sensitivities. This is done in an effort to reduce the dimensionality of the estimation problem. Bayesian parameter estimation is performed on a cohort of 84 patients. We also discuss the use of machine learning classifiers to constrain parameter realizations associated with realistic physiologic response and to detect pulmonary hypertension based on assimilated model parameters.

Karlyn Harrod
University of Notre Dame
Applied and Computational Mathematics and Statistics
kharrod@nd.edu

Alison Marsden
School of Medicine
Stanford University, USA
amarsden@stanford.edu

Daniele E. Schiavazzi
University of Notre Dame
dschiavazzi@nd.edu

MS160

Shape Based Effective Estimation for Cardiovascular Models

A major challenge in the context of cardiac modeling lies in using the large amount of data available on the system to circumvent the lack of absolute modeling ground truth, since every system considered is in fact patient-specific, with possibly non-standard conditions associated with a disease. Here data assimilation originally considered in environmental sciences has now diffused into the life science modeling community and in particular the cardiac modeling community. In this talk, we will present a control oriented strategy based on the observer theory where data driven feedback terms control the trajectory of the simulated system to track the target real trajectory and jointly identify some parameters. The feedback design is clearly dependent of the type of model considered – mechanics, reaction diffusion, wave or transport systems – but also closely related to the nature of the data considered. There, if we do not want to rely on approximative data post-processing, it is a momentous challenge to be able to use the data as they are essentially delivered which are most of the time only images of the patterns or shapes of the underlying physical quantities. For instance, we will show examples of such observer design when we have access only to the front propagation of the cardiac action potential, or we only image the shape of the myocardium domain deforming during a heart beat.

Philippe Moireau
INRIA

philippe.moireau@inria.fr

MS160

Adding Constraints to Bayesian Inverse Problems

Using observation data to estimate unknown parameters in computational models is broadly important. This task is often challenging because solutions are non-unique due to the complexity of the model and limited observation data. However, the parameters or states of the model are often known to satisfy additional constraints beyond the model. Thus, we propose an approach to improve parameter estimation in such inverse problems by incorporating constraints in a Bayesian inference framework. Constraints are imposed by constructing a likelihood function based on the fitness of the solution to the constraints. The posterior distribution of the parameters conditioned on (1) the observed data and (2) satisfaction of the constraints is obtained, and the estimate of the parameters is given by the maximum a posteriori estimation or posterior mean. Both equality and inequality constraints can be considered by this framework, and the strictness of the constraints can be controlled by constraint uncertainty denoting our confidence on its correctness. Furthermore, we extend this framework to the ensemble Kalman filter method, where the constraint is imposed by re-weighing the ensemble members based on the likelihood function. A synthetic model is presented to demonstrate the effectiveness of the proposed method. The numerical simulations show that imposing constraints using the method presented improves identification of the true parameter solution among multiple local minima.

Jiacheng Wu
UC Berkeley
wujiacheng@berkeley.edu

Jian-Xun Wang
University of Notre Dame
jwang33@nd.edu

Shawn Shadden
University of California, Berkeley
shadden@berkeley.edu

MS161

Supporting Advanced Finite Element Methods in the Deal.II library

Realistic simulators for continuum mechanics problems today need to be able to deal both with complex physics, large numbers of unknowns, and complex discretization and solver schemes. I will review the support for all of these things in the generic finite element library deal.II (see http://www.dealii.org) and illustrate them with applications built on deal.II. In particular, I will mention support for complex discretizations on CPUs, GPUs, and using matrix-free operations, and illustrate the limitations of these approaches using examples from codes that solve real-world cases.

Wolfgang Bangerth
Colorado State University
bangerth@colostate.edu

MS161

Balancing the Numerical and Parallel Performance for Reservoir Simulations

The overall performance of a PDE-based simulator depends on two factors: the algorithmic efficiency of the numerical scheme chosen and the parallel efficiency of the software implementation. Since aspects from the two factors may influence each other's performance, a suitable balance between the two is important. The focus of this talk is on the OPM framework of oil reservoir simulation, for which the computational core is to solve the black-oil model: a coupled system of nonlinear PDEs. Due to large variations in the geological properties of a reservoir, the sparse matrix that arises from discretizing the coupled PDEs exhibits a strong heterogeneity in its nonzero values. These reflect the strength of coupling between the degrees of freedom. It is thus necessary to consider this heterogeneity in the unstructured mesh partitioning process, typically translated to partitioning a graph with weighted edges. Particularly, we study the impact of different strategies of edge weighting on both the numerical and parallel performance. The ordering of the degrees of freedom, which also affects both sides, is studied in addition. Our purpose is to shed some light on a suitable mesh partitioning and ordering methodolgy, which is also relevant beyond the context of reservoir simulation. The issue of how to allow users of OPM to inject such flexibility into the existing software framework is also discussed.

Andreas Thune, Xing Cai
Simula Research Laboratory
1325 Lysaker, Norway
andreast@simula.no, xingca@simula.no

MS161

The SUNDIALS Suite of Time Integrators and Nonlinear Solvers: Preparing for Exascale Computing

SUNDIALS is a suite of robust and scalable solvers for systems of ordinary differential equations, differential-algebraic equations, and nonlinear equations. The suite consists of six packages: CVODE(S), ARKode, IDA(S), and KINSOL. Each package is built on a common vector, matrix, linear solver, and nonlinear solver API allowing for application-specific and user-defined data structures and solvers, encapsulated parallelism, and algorithmic flexibility. In this presentation we will overview the capabilities of the SUNDIALS suite and discuss recent changes and current efforts as part of the DOEs Exascale Computing Program and FASTMath Institute to enable time integrators for exascale architectures. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

David J. Gardner, Carol S. Woodward
Lawrence Livermore National Laboratory
gardner48@llnl.gov, woodward6@llnl.gov

Daniel R. Reynolds
Southern Methodist University
Mathematics
reynolds@smu.edu

Alan Hindmarsh
Lawrence Livermore National Laboratory
Center for Applied Scientific Computing
alanh@llnl.gov

Slaven Peles
Lawrence Livermore Natinoal Laboratory
peles2@llnl.gov

Cody J. Balos
Lawrence Livermore National Laboratory
balos1@llnl.gov

MS161

Walberla: A General Purpose Software Framework for Massively Parallel Simulations

Simulations are gradually becoming an attractive alternative to laboratory experiments to study complex physical processes, as they are in many cases faster and cheaper. To obtain reliabel and expressive results, however, a great level of detail has to be incorporated into such simulations, rendering them inherently compute intensive. In this talk, we will present the open-source HPC framework waLBerla (www.walberla.net) that has been carefully designed for massively parallel multi-physics simulations running on the largest supercomputers. Starting as a C++ framework for CFD simulations based on the lattice Boltzmann method, it has continuously developed to be suitable for all kinds of numeric codes, like rigid particle dynamics and phase-field methods. Applications range from flows inside and around complex geometries over studies of dune formation in riverbeds to additive manufacturing processes. Vital cornerstones for these applications are features like adaptive mesh refinement and dynamic load balancing. We present illustrating examples and discuss the respective challenges posed to the framework by the users and lessons learnt through experience. To support the rapidly evolving variety of hardware architectures and method variants, an automatic code generation front-end written in Python has been developed. Sustained code quality is maintained through a rigorous continuous integration (CI) pipeline, featuring automated built environments and unit tests.

Christoph Rettinger
Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU)
christoph.rettinger@fau.de

Martin Bauer
University of Erlangen-Nuremberg
Department of Computer Science
martin.bauer@fau.de

Sebastian Eibl
Friedrich-Alexander-University Erlangen-Nuremberg
sebastian.eibl@fau.de

Christian Godenschwager, Florian Schornbaum
Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU)
christian.godenschwager@fau.de,
florian.schornbaum@fau.de

Harald Köstler
Universität Erlangen-Nürnberg
Department of Computer Science 10 (System Simulation)
harald.koestler@fau.de

Ulrich Rüde
Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU)
ulrich.ruede@fau.de

MS162

Recent Developments in Spectral Element Methods for Turbulent Flows

We present recent developments in application of the spectral element method (SEM) for the simulation of turbulent flows. The SEM is a high-order weighted residual technique for the solution of partial differential equations that combines local tensor-product-based element representations with globally-unstructured assembly. The local representation admits fast operator evaluations with memory counts that depend only on the total number of grid points, n, independent of the representation order, N. The work complexity is only O(Nn), with very small constants. For the incompressible Navier-Stokes equations, the pressure solve presents the stiffest substep and efficient preconditioning of the associated Poisson problem is of paramount importance. We describe a hybrid Schwarz-multigrid strategy for solving this problem that has strong-scaled to over a million MPI ranks for engineering problems in complicated domains. We then explore the potential of new low-order preconditioning strategies that employ algebraic multigrid solvers. Finally, in an effort to further economize on costs across a broad array of applications, we consider Schwarz-based extensions to the SEM discretization of the Navier-Stokes equations and examine the potential of reduced-order models applied to turbulent heat transfer.

Paul Fischer
University of Illinois at Urbana-Champaign
fischerp@illinois.edu

MS162

High-order Steady State Solvers for Multiphysics Applications

This talk will present high-performance high-order simulation techniques for steady-state solution of the incompressible Navier-Stokes. We explore several approaches to steady-state solution of the incompressible Navier-Stokes equations using the spectral element method (SEM) as the spatial discretization. Strategies include pseudo-timestepping, Jacobi-free Newton-Krylovmethods, and full Newton with exact Jacobians. GMRES is used to solve for the divergence-free updates at each Newton step. Hybrid-Schwarz multigrid is used to solve for pressure within each GMRES iterate. I will present performance results for a

variety of 2D and 3D applications, including steady-state Reynolds-averaged Navier-Stokes solutions based on a new $k-\omega$ formulation.

MiSun Min
Argonne National Laboratory
Mathematics and Computer Science Division
mmin@mcs.anl.gov

MS162

Low-memory Implicit Solvers for the Sparse Line-DG Method with Kronecker-SVD Preconditioning

High-order methods such as the discontinuous Galerkin (DG) method have received significant attention over the last decade, for example in the area of LES simulation of turbulent flows. However, for high polynomial degrees these methods tend to have very large stencils, since they connect all DOFs inside each element. We will demonstrate how naturally sparse schemes such as the Line-DG method can be used to fundamentally improve the ratio between computation and memory access, and therefore obtain excellent scaling to future generations of computer architectures. We focus on the case of implicit time-stepping, where we also use a low-memory Kronecker-SVD technique to precondition our Krylov methods. We apply our methods to relevant applications such as conservative re-mapping for Lagrangian methods.

Per-Olof Persson
Dept. of Mathematics
University of California, Berkeley
persson@berkeley.edu

Will Pazner
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
pazner1@llnl.gov

MS162

Geometry, Meshing and Mesh Adaptation of High-order Curved Meshes

To maintain the rate of convergence of high-order methods in simulations over curved domains the mesh entities must be curved to the boundary with a sufficient order of geometric approximation. The application of high-order methods on problems with complex geometric domains requires effective tools for the creation and adaptation of curved meshes. This presentation will overview the issues associated with the generation and adaptation of high-order curved meshes starting from detailed CAD models. The first step is manipulation of, and interaction with, the CAD geometry as required to create the geometry to be meshed, the analysis geometry. The second step is the generation of the initial curved mesh while the third step is the ability to adapt that mesh based on a posteriori information from multiple sources including the discretization errors and element shape quality in the case of deforming meshes. Specific consideration will be given to consideration of the local mesh modification operations needed to support curved mesh adaptation when the mesh geometry is higher order.

Morteza Hakimi, Kazem Kamran, Onkar Sahni
Rensselaer Polytechnic Institute
m.hakimi.siboni@gmail.com, kazemkamran@gmail.com, sahni@rpi.edu

Mark S. Shephard
Rensselaer Polytechnic Institute
Scientific Computation Research Center
shephard@rpi.edu

Blair Downie, Rocco Nastasia, Saurabh Tendulkar
Simmetrix Inc.
downie@simmetrix.com, nastar@simmetrix.com, saurabh@simmetrix.com

Mark Beall
Simmetrix, Inc.
mbeall@simmetrix.com

MS163

H/P Adaptive Entropy Stable Methods for the Compressible Navier-Stokes Equations: Curvilinear Grids

Our focus is on the solution of the compressible Euler and Navier-Stokes (NS) equations on curvilinearly mapped grids. First, note that both the continuous Euler and NS equations are equipped with an incomplete L_2 stability theory in the form of a secondary conservation law: the conservation of entropy. (Albeit, incomplete in that positivity of density and temperature are assumed). Our approach is to mimic the continuous stability proofs at the discrete level by focusing exclusively on discrete operators satisfying the diagonal norm, summation-by-parts (SBP) property. Because SBP operators mimic the continuous integration-by-parts property, semi-discrete proofs of nonlinear stability can be directly constructed. Indeed, the discrete stability proofs follow the continuous proofs in a one-to-one manner. This combination has previously been used to construct entropy stable schemes that allow for h/p adaptivity on Cartesian grids. The extension of these ideas to curvilinear grids requires great care to ensure nonlinear stability, conservation, free-stream preservation and design order accuracy for arbitrary polynomial orders.

Mark H. Carpenter
NASA Langley Research Center
mark.h.carpenter@nasa.gov

David C. Del Rey Fernandez
NASA Langley Research Center
National Institute of Aerospace, Hampton
dcdelrey@gmail.com

Matteo Parsani
King Abdullah University of Science and Technology (KAUST)
matteo.parsani@kaust.edu.sa

Travis Fisher
Sandia National Labs

tcfishe@sandia.gov

Jared Crean, Jason E. Hicken
Rensselaer Polytechnic Institute
creanj@rpi.edu, hickej2@rpi.edu

MS163

Conservative Split form Summation-by-parts Discretization of the Euler Equations on Unstructured Grids

Many split form schemes are more robust than standard divergence form discretization schemes for two main reasons. Firstly, they reduce aliasing errors introduced by the quadratic and cubic nonlinear convective terms in the compressible Navier-Stokes equations. Furthermore, they can discretely satisfy secondary equations such as the kinetic energy equation. To date, no split formulation of the Euler equations that conserves (or dissipates) mathematical entropy has been constructed. In some cases, however, it has been shown that split forms can be as robust as entropy-stable schemes while being computationally more efficient. In this work, a robust, high-order, element-wise conservative semi-discrete scheme applicable to unstructured grids is constructed and analyzed for the compressible Euler equations. The key ingredient is spatial discretization operators which satisfy the summation-by-parts property. These operators allow split form discretization of conservation laws without sacrificing element-wise conservation in the Lax-Wendroff sense. Numerical investigations demonstrate the robustness, accuracy, conservation, and efficiency properties of the scheme.

Siavosh Shadpey
University of Toronto, Canada
siavosh@ualberta.ca

David Zingg
professor and director
University of Toronto, Institute for Aerospace Studies
dwz@oddjob.utias.utoronto.ca

MS163

High-order Entropy Stable Spectral Collocation Schemes for the Navier-Stokes Equations on Moving Deforming Grids

We develop new high-order entropy stable spectral collocations schemes for the unsteady 3-D Navier-Stokes equations on moving deforming grids. To take into account the grid motion and deformation, we use an arbitrary Lagrangian-Eulerian (ALE) formulation to map the Navier-Stokes equations to a fixed reference system of coordinates. The proposed scheme is constructed by using the skew-symmetric form of the Navier-Stokes equations, which are discretized by using summation-by-parts spectral collocation operators that preserve the conservation properties of the original governing equations. Furthermore, the governing equations and the metric coefficients are approximated such that the geometric conservation laws (GCL) are satisfied with the machine accuracy on both static and dynamic grids. To make the scheme entropy stable, a new entropy conservative flux has been derived for the 3-D Euler and Navier-Stokes equations on dynamic unstructured hexahedral grids. The new flux preserves the design order of accuracy of the original spectral collocation scheme and guarantees the entropy conservation on moving deforming grids. We present numerical results demonstrating the entropy stability and design order of accuracy of the new schemes on moving deforming grids.

Nail Yamaleev
Old Dominion University
nyamalee@odu.edu

David C. Del Rey Fernandez
NASA Langley Research Center
National Institute of Aerospace, Hampton
dcdelrey@gmail.com

Jialin Lou
Old Dominion University
jlou@odu.edu

Mark H. Carpenter
NASA Langley Research Center
mark.h.carpenter@nasa.gov

MS164

An Efficient and Accurate Parallel Simulator for Streamer Discharges in Three Dimensions

Abstract not available.

Lin Bo
National University of Singapore
linbo@u.nus.edu

MS164

Band-Toeplitz Preconditioners for Nonsymmetric Toeplitz Systems

Abstract not available.

Sean Hon
University of Oxford
seanyshon@gmail.com

MS164

ALSVINN: A Fast Multi-GPGPU Finite Volume Solver with a Strong Emphasis on Reproducibility

Abstract not available.

Kjetil Olsen Lye
ETH
kjetil.lye@sam.math.ethz.ch

MS164

A Robust Solver Based on Multiplicative Schwarz Methods for Isogeometric Discretizations

Abstract not available.

Ivaro Pé de la Riva, Carmen Rodrigo, Francisco José Gaspar

University of Zaragoza
apedelariva@unizar.es, carmenr@unizar.es, fjgaspar@unizar.es

MS165

Multiscale Model Reduction for Wave Equations

It is important to develop fast yet accurate numerical methods for seismic wave propagation to characterize complex geological structures and oil and gas reservoirs. However, the computational cost of conventional numerical modeling methods, such as finite-difference method and finite-element method, becomes prohibitively expensive when applied to very large models. We propose a Generalized Multiscale Generalized Multiscale Finite-Element Method (GMsFEM) for elastic wave propagation in heterogeneous, anisotropic media, where we construct basis functions from multiple local problems for both boundaries and the interior of a coarse node support or coarse element. The application of multiscale basis functions can capture the fine scale medium property variations, and allows us to greatly reduce the degrees of freedom that are required to implement the modeling compared with conventional finite-element method for wave equation, while restricting the error to low values. The research is partially supported by the Hong Kong RGC General Research Fund (Project: 14304217) and CUHK Direct Grant for Research 2017-18.

Eric Chung
The Chinese University of Hong Kong
Department of Mathematics
tschung@math.cuhk.edu.hk

MS165

Hyperbolic Homogenization

There has been significant work on numerical homogenization methods for problems governed by the wave equation, even for cases without scale separation. However, these methods typically derive from the study of elliptic problems, exploiting the concept of G-convergence to produce approximate wave equations. The problems which arise to define the approximate system are inherently global but approximated via localization on patch neighborhoods. In this talk we propose methods specifically designed for hyperbolic systems, which in particular exploit the fundamental features of local domain-of-dependence and causality. The mathematical framework is that of reduced order modeling (ROM) and the tools are essentially algebraic. Specifically, we consider the direct approximation of a fine grid solution operator, S_h, by a computationally-discovered coarse grid subspace, spanned by V_H, leading to a reduced evolution matrix $V_H^T S_h V_H$.

Thomas M. Hagstrom
Southern Methodist University
Department of Mathematics
thagstrom@smu.edu

MS165

Efficient Density Estimation in Noisy PDEs and Uncertainty Propagation

The effect of model uncertainties and noise on a model output is often best described by its probability density function (PDF). Although density estimation is a common uncertainty-quantification (UQ) task, the adequacy of approximation methods (surrogates) for density estimation has scarcely been analyzed before. We first show that standard spectral methods in uncertainty propagation, such as generalized polynomial chaos (gPC), sometimes fail to approximate the PDF even in the case of one-dimensional noise. Therefore, we developed a novel spline-based algorithm for this task. Our method offers significant advantages over existing methods for density estimation, primarily a guaranteed convergence rate which is polynomial in the sampling resolution. This convergence rate is better than that of standard statistical density-estimation methods (such as histograms and kernel density estimators) given an input noise of moderate dimension. Furthermore, our spline-based approximation outperforms spectral methods when the sample size is small and the quantity of interest has sharp-gradients regions. Finally, we present applications of our algorithm for problems in nonlinear optics and fluid dynamics.

Amir Sagiv
Tel Aviv University
amirsagiv@mail.tau.ac.il

Gadi Fibich
Tel Aviv University
School of Mathematical Sciences
fibich@tau.ac.il

Adi Ditkowski
Department of Applied Mathematics
Tel-Aviv University, Israel
adid@post.tau.ac.il

MS165

Rational Krylov Subspaces and Phase-preconditioning for Model Reduction of Wave Equations

In many applications like approximation of matrix functions or dynamical systems, rational Krylov subspaces show superior approximation qualities over polynomial Krylov subspaces. Especially, if the spectrum of the function and the domain where the function is approximated are well separated from each other, rational approximants are known to converge fast. This, however, is generally not the case when approximating discretizations of wave equations with a reduced order model. Nevertheless, reduced order models based on rational Krylov subspaces still perform well for applications with very lossy or resonant structures. In case the considered application is dominated by waves traveling over long distances the order of the rational function needed to approximate a response grows proportional to the travel time in accordance with the Nyquist sampling limit. In this talk we combine rational Krylov subspaces with asymptotic methods to obtain a projection-based model order reduction method that can handle strong scatterers as well as long travel times. The resulting method shifts the computational cost from

the poorly scalable forward solvers to the evaluation of large inner products which are embarrassingly parallelizable. We show that incorporating asymptotic methods into a projection-based framework relaxes the spatial discretization required and allows sub-Nyquist sampling in the frequency domain.

Jorn Zimmerling
University of Michigan
jzimmerl@umich.edu

Rob F. Remis
Delft University of Technology
Faculty of Elect Engineering, Mathematics, Computer Scienc
r.f.remis@tudelft.nl

Vladimir Druskin
Worcester Polytechnic Institute
vdruskin@gmail.com

Mikhail Zaslavsky
Schlumberger-Doll Research
mzaslavsky@slb.com

MS166

Ultra High Resolution Topology Optimization: Brute Force or Smart Discretizations?

The optimal topology of large structural systems has until now been concerned with the design of individual parts and not that of complete assemblies. Following recent advances in numerical algorithms tailored for large scale structural optimization, this limitation has now be circumvented, mainly by utilizing powerfull HPC system [Aage, N., et.al., Topology optimization using PETSc: An easy-to-use, fully parallel, open source topology optimization framework. SMO, 51:565-572, 2015.]. The design approach has been demonstrated on both aircraft, bridge and ship design problems, resulting in noticeable performance enhancement. However, the increase in design resolution comes at a great cost in terms of the needed computational power. Therefore, it is interesting, if not paramount, to pursue alternative design representation schemes that allows for ultra high resolution optimal designs without the need for large computing resources. This talk will present both the brute force topology optimization approach, in which the governing PDEs are solved by classical methods, i.e. Krylov methods and multigrid preconditioning [Aage, N., et.al. Giga-voxel computational morphogenesis for structural design. Nature, 52:84-86, 2017.], as well as novel homogenization based projection schemes. The latter allows for the solution of the design problem at low resolution, and by parameter extraction, one can represent the design using orders of magnitude finer design representations.

Niels Aage
Technical University of Denmark
naage@mek.dtu.dk

MS166

Topology Optimization of Structures Subject to Snapping Behaviour

In this talk we discuss methods which is intended to be used for of design structures that are expected to snap. One example of such structures are multi-stable micro-flexures that buckle to perform digital computations. Since such structures inevitably function under finite strains we model the material by general finite strain elasticity. The balance equations are solved using the finite element method in a total Lagrangian setting along with Newton-Raphson iterations. To trace the load path and to be able to pass singular points we make use of path following technique. To find the optimal material layout we associate one design variable with each element such that material and void can be represented. The optimization problem is to find a material distribution such that the objective is minimized while fulfilling the constraints. In our application the objective is to find a layout that gives a stable state in terms of energy level while having distinct states, i.e. the difference in deformation between the stable states should exceed a given threshold. We also impose a constrains on the available mass of the device. To solve the optimization problem we use mathematical programming and in particular we use the Method of Moving Asymptotes. The gradients required to form the convex approximations are established via the adjoint sensitivity approach. To form a well-posed problem we regularize the optimization problem via the use of a PDE filter.

Mathias Wallin
Division of Solid Mechanics
Lund University
mathias.wallin@solid.lth.se

Daniel Tortorelli
Lawrence Livermore National Laboratory
Univerisity of Illinios at Urbana-Champaign
tortorelli2@llnl.gov

MS166

A Density Gradient Approach to Topology Optimization under Design-dependent Boundary Loading

The paper proposes a density gradient based approach to topology optimization under design-dependent boundary loading. In the density-based topology optimization method, we impose the design dependent loads through spatial gradient of the density. We transform design-dependent boundary loads into a volume form through volume integral of density gradient. In many applications where loadings only need to be exerted on partial boundary, we introduce an auxiliary loading density to keep track of the loading boundary. During the optimization, the loading density is updated by tracking the changes of the physical density in the vicinity of the loading boundary at previous iteration. The proposed approach is easy to implement and computationally efficient. In addition, by adding more auxiliary density fields, the proposed approach is applicable to multiple design-dependent loads. To prevent the intersection of different loading boundaries, a Heaviside projection based integral constraint is developed. Both heat conduction problems under convection loading and elastic problems under hydrostatic pressure loading are

presented to illustrate the effectiveness and efficiency of the method.

Cunfu Wang
University of Wisconsin-Madison
cwang392@wisc.edu

Xiaoping Qian
Department of Mechanical Engineering
University of Wisconsin-Madison
qian@engr.wisc.edu

MS167

A High Performance Algorithm for X-ray Image Deblurring

Image deblurring via deconvolution can be formulated as a hierarchical Bayesian inverse problem, and numerically solved by Markov Chain Monte Carlo (MCMC) methods. Numerical solution is difficult because (a) inconsistent assumptions about the data outside of the field of view of the image lead to artifacts near the boundary, and (b) the Bayesian inverse problem is high-dimensional for high-resolution images. The numerical MCMC framework I present addresses these issues. Boundary artifacts are reduced by reconstructing the image outside the field of view. Numerical difficulties that arise from high-dimensions are mitigated by exploiting sparse problem structure in the prior precision matrix.

Jesse Adams
Department of Energy
adamsjj@nv.doe.gov

MS167

Recent Advances in Lifting Factorization for Discrete Wavelet Transforms

Wavelet transforms form an important alternative to methods like Principal Component Analysis or Fourier transforms and are particularly useful for image-like data owing to their joint time-frequency localization, good approximation properties, and fast digital realizations. They have been used in image coding standards like the FBI/NIST WSQ Fingerprint Specification, NASA/CCSDS Recommendation 122.0-B-1 for Space Data Systems, and the ISO/IEC 15444-x (JPEG2000) standards. The latter two employ efficient cascade-form decompositions, known as lifting factorizations, for the filter banks that make up discrete wavelet transforms. The author has developed an approach to causal lifting factorization based on the theory of linear Diophantine equations over polynomial rings, which guarantees existence and uniqueness of causal factorizations satisfying polynomial degree-reducing inequalities. This enables a new factorization scheme, the Causal Complementation Algorithm (CCA), that provides an alternative to the noncausal scheme of Daubechies and Sweldens based on the Laurent polynomial Extended Euclidean Algorithm (EEA). The CCA uses a generalization of polynomial division that ensures existence and uniqueness of quotients whose remainders satisfy divisibility constraints. The CCA is more general than the causal variant of the EEA, generating factorizations not obtainable using the causal EEA, and it yields all degree-reducing causal lifting factorizations of a given filter bank.

Christopher M. Brislawn
Los Alamos National Laboratory
cbrislawn@yahoo.com

MS167

Quantifying Spatially Varying Blur in X-ray Radiographic Imaging Systems

In complex imaging systems, resolution can vary spatially throughout the image. Quantitatively evaluating the spatial variation and developing principled methods to mitigate its effect are important problems in high-energy x-ray imaging systems supporting dynamic materials experiments conduced at Nevada National Security Site. In particular, systems that employ a scintillating crystal that converts x-rays to visible light are pushing the thickness of the crystal to higher and higher levels to meet experimental requirements. This introduces a spatially varying blur that, in certain circumstances, can be on the order of (or even dominate) blur due to the radiographic spot size. This work provides methods for quantifying the effect of that blur and estimating a spatially varying impulse response of the system with a rigorous quantification of uncertainty.

Kevin Joyce
Nevada National Security Site
Signals Processing and Applied Mathematics
joycekt@nv.doe.gov

MS168

Discussion on Reproducibility of Network Algorithms

In the final session of the minisymposium, we will organize a discussion among the speakers and the audience on the issue of reproducibility in the context of network analysis. With the talks as a starting point, we will explore the issues of how stochasticity in network algorithms and the non-determinism of parallel algorithms lead to sometimes non-repeatable results. We will also solicit ideas to measure the amount of non-determinism, how to bridge between between "mathematically optimal results" and "ground truth" and how to tackle the joint effects of parallelism, stochastic algorithms and dyanmic networks.

Sanjukta Bhowmick
Department of Computer Science
University of Nebraska, Omaha
sbhowmick@unomaha.edu

MS168

Dynamic Graph Clustering with Variable Time Resolution

Clustering graphs of massive scale relies heavily on stochastic clustering algorithms, such as the local agglomerative method of Louvain clustering, which have results which are inherently difficult to reproduce due to the partially random nature of the output. This is exacerbated in the case

of dynamic graph sequences which experience evolution of the underlying communities: during the transition period between two underlying models stochastic clustering algorithms may have high variance due to the confusion between the two possible models. As part of the DyGT (Dynamic Graph Tool) project we have lowered the variance of stochastic clustering algorithms through the use of an ensemble and handled the effect of evolving communities using a smoothing process. Secondly, the choice of time resolution when creating the sequence of snapshots in the dynamic graph can have a dramatic effect on the output of clustering algorithms. Time slice resolution, e.g. daily or weekly snapshots of a dynamic network, is a property usually decided arbitrarily during the process of collecting or interpreting the data. Despite the typically casual selection of the time resolution parameter this decision can have dramatic effects on the output of clustering algorithms. To deal with the issue of time resolution we have explored methods which adaptively select non-uniform time slices from the data in a way which best represents the community evolution.

Timothy La Fond
Lawrence Livermore National Lab
lafond1@llnl.gov

MS168

Reproducibility in Parallel Graph Algorithms

The rapidly growing number of large network analysis problems has led to the emergence of many parallel and distributed graph processing libraries—one survey in 2014 identified over 80. Since then, the landscape has evolved; some packages have become inactive while more are being developed. Furthermore, performance characteristics often vary wildly between packages. This complexity makes selecting the optimal package for a given problem infeasible. To mitigate this complexity, our work automates the collection of performance data for graph processing systems and uses this data to predict the performance of graph processing algorithms based on features of the input graph. We ensure fairness in experimental measurements inspecting timing calls, stopping criterion, and graph representation (e.g. directed/undirected, simple/multigrpah). We also predict performance using both regression models and by binary classification, labeling packages as well-performing or poorly-performing. We demonstrate our approach on six graph processing packages: GraphMat, the Graph500, the Graph Algorithm Platform Benchmark Suite, GraphBIG, Galois, and PowerGraph and four algorithms: PageRank, Single Source Shortest Paths, Triangle Counting, and Breadth First Search. Given a dataset, our method can either estimate execution time or suggest an implementation and thread count. Our method correctly identifies well-performing configurations in 97% of test cases.

Samuel D. Pollard
University of Oregon
spollard@cs.uoregon.edu

MS168

Parallel Dynamic Networks

We present shared-memory parallel algorithms for updating the properties of dynamic networks. In particular, we will introduce a three step template for creating scalable, parallel algorithms for updating dynamic networks and demonstrate how this template can be to create algorithms for updating Connected Components(CC), Minimum Weighted Spanning Tree (MST) and Single Source Shortest Path (SSSP). We will present empirical results that show that these algorithms are scalable and the time for updating is faster than recomputing with Galois, a state-of the art software.

Sriram Srinivasan
University of Nebraska at Omaha
sriram882004@gmail.com

MS169

Efficient Discovery of Heterogeneous Treatment Effects via Anomalous Pattern Detection

We propose Treatment Effect Subset Scan (TESS), a new method for discovering which subpopulation in a randomized experiment is most significantly affected by a treatment. We frame this challenge as a pattern detection problem where we efficiently maximize a nonparametric scan statistic over subpopulations. Furthermore, we identify the subpopulation which experiences the largest distributional change as a result of the intervention, while making minimal assumptions about the intervention's effects or the underlying data generating process. In addition to the algorithm, we demonstrate that the asymptotic Type I and II error can be controlled, and provide sufficient conditions for detection consistency—i.e., exact identification of the affected subpopulation. Finally, we validate the efficacy of the method by discovering heterogeneous treatment effects in simulations and in real-world data from a well-known program evaluation study.

Edward McFowland
University of Minnesota
mcfowland@umn.edu

MS169

Queueing Theory in the Age of Smartphone Technology: Dynamics and Computation

Many service systems provide real-time information to their customers via smartphones with the goal of reducing the customers' anxiety of the unknown. However, the information might be unreliable or not given in real-time. In this talk, we show how to prove fluid and diffusion limit theorems for a state dependent in nite server queueing model where customers choose which queue to join by a generalized customer choice model and where the information about the queue length is updated in discrete intervals. We compare our queueing model with periodic updates against queues that update constantly, but are delayed by a constant. We also show using data from Disneyland that giv-

ing customers information via smartphones may not be a smart decision.

Jamol Pender
Cornell University
jjp274@cornell.edu

MS169

Resource Allocation in Sequential Adaptive Clinical Trials: Exploration versus Exploitation

Adaptive clinical trials for new drugs or treatment options promise significant benefits to both the trial sponsor and the patients, but complicate resource allocation. We focus on Phase 3 group sequential adaptive trials, which allow for early trial termination at an interim analysis point due to either proven benefit or futility. First, we build a simulation model to analyze drug misclassification risk in current sequential testing practices. Then, we build a stochastic dynamic programming model to analyze the impact of interim analyses on the candidate drugs misclassification risk, time-to-market, and expected profit from commercialization. The resource allocation decision in this setting (i.e., the number of patients to enroll in each period and the decision to terminate or continue the trial at each interim analysis) is characterized by endogenous uncertainty, and a trade-off between the incentive to establish that the drug is effective early on (exploitation), due to a time-decreasing market revenue, and the benefit from collecting some information on the drugs efficacy early on the trial (exploration), prior to committing a large budget. We present structural properties of an optimal resource allocation policy and perform a numerical study utilizing realistic data. Finally, we propose an augmented sequential testing methodology based on our research findings (joint work with Ebru K. Bish and Niyousha Hosseinichimeh)

Alba Rojas-Cordova
Southern Methodist University
alba@smu.edu

MS170

Layers of Low-rank Couplings for Large-scale Bayesian Inference

We present a framework for the greedy approximation of high-dimensional Bayesian inverse problems through the composition of low-dimensional transport maps. The action of each transformation is confined to a low-dimensional subspace of the parameters. The algorithm works on a sequence of "residuals", given by recursively pulling back the posterior distribution through the sequence of computed maps. At each step, the algorithm identifies (i) a relevant sub-space of the residual and (ii) a low-dimensional transformation between the restriction of the residual onto this sub-space and a standard normal. The map is computed by minimizing the Kullback-Leibler divergence over a class of parametric transformations. The algorithm will be showcased on a range of inference problems arising in spatial statistics.

Daniele Bigoni
Massachusetts Institute of Technology
dabi@mit.edu

Olivier Zahm
Univ. Grenoble Alpes, Inria
olivier.zahm@inria.fr

Alessio Spantini, Youssef M. Marzouk
Massachusetts Institute of Technology
spantini@mit.edu, ymarz@mit.edu

MS170

Use of the Bayesian Approximation Error ... proach to Account for Model Discrepancy: ... Robin Problem Revisited

We address the problem of accounting for model disc ancy by the use of the Bayesian approximation error (B approach in the context of inverse problems. In many verse problems when one wishes to infer some primary rameter of interest there are other secondary parame which are also uncertain. In the standard Bayesian deterministic) approach such nuisance parameters are ther inverted for or are ignored (perhaps by assignmer some nominal value). However, it is well understood the ill-posedness of general inverse problems means t do not handle modelling errors well. The BAE appro has been developed as an efficient means to approxima pre-marginalize over nuisance parameters so that one systematically incorporate the effects of neglecting t secondary parameters at the modelling stage. We moti the method through an application to the Robin prob governed by the Poisson equation

Ruanui Nicholson
Engineering Sceience
University of Auckland, New Zealand
rnic052@aucklanduni.ac.nz

Noemi Petra
University of California, Merced
npetra@ucmerced.edu

Jari Kaipio
Department of Mathematics
University of Auckland
jari@math.auckland.ac.nz

MS170

Efficient Sampling from High-dimensional Dis butions using Low-rank Tensor Surrogates

High-dimensional distributions are notoriously difficul sample from, particularly in the context of PDE strained Bayesian inverse problems. In this talk, we present general purpose samplers based on low-rank te surrogates in the tensor-train (TT) format, a methodo that has been exploited already for many years for able, high-dimensional function approximations in qu tum chemistry. In the Bayesian context, the TT surro is built in a two-stage process. First, we build a su gate of the entire PDE solution in the TT format, u a novel combination of alternating least squares and TT cross algorithm. It exploits and preserves the b diagonal structure of the discretised operator in stocha collocation schemes, requiring only independent PDE

lutions at a few parameter values, thus allowing the use of existing high performance PDE solvers. In a second stage, we approximate the high-dimensional likelihood function also in TT format. Due to the particular structure of the TT surrogate, we can build an efficient inverse Rosenblatt (or cumulative) transform that only requires a sampling algorithm for one-dimensional conditionals. The overall computational cost of the sampler grows only linearly with the dimension. For sufficiently smooth prior distributions of the input random fields, the ranks required for accurate TT approximations are moderate, leading to significant computational gains.

Robert Scheichl
University of Heidelberg
r.scheichl@uni-heidelberg.de

Karim Anaya-Izquierdo, Sergey Dolgov
University of Bath
Department of Mathematical Sciences
k.anaya-izquierdo@bath.ac.uk, s.dolgov@bath.ac.uk

Colin Fox
University of Otago, New Zealand
fox@physics.otago.ac.nz

MS170

Multi-Index and Multi-level Markov Chain Monte Carlo in MUQ2

A common approach for solving Bayesian inference problems is the Markov Chain Monte Carlo (MCMC) algorithm. In many applications, models based on partial differential equations (PDEs) are employed, requiring a costly numerical PDE solution in every single step of the chain. Further, many samples are then discarded, particularly for high-dimensional parameters. This clearly leads to a massive computational effort. In analogy to multilevel PDE solvers, the Multilevel Markov Chain Monte Carlo (MLMCMC) method exploits structure within the underlying PDE. Here, most of the sampling is shifted to simpler coarse approximations of the original PDE model, requiring only few samples of full accuracy and significantly reducing cost. A further extension to MLMCMC is the Multiindex Markov Chain Monte Carlo (MIMCMC) method which introduces multiple dimensions of model refinement. We give an algorithmic view of MLMCMC and MIMCMC, as well as a description of a new implementation leveraging both the high-performance PDE library DUNE (https://www.dune-project.org/) as well as a new extension to the MIT Uncertainty Quantification library MUQ (http://muq.mit.edu/). The former is an established modular library for the solution of large-scale PDEs in high performance computing environments, while the latter provides an easy-to-use framework to handle the statistical part of the methods. We demonstrate the effectiveness of the new implementation on a number of applications.

Linus Seelinger
Universität Heidelberg
Institut für Wissenschaftliches Rechnen
linus.seelinger@iwr.uni-heidelberg.de

Ole Klein
Interdisciplinary Center for Scientific Computing
Heidelberg University
ole.klein@iwr.uni-heidelberg.de

Robert Scheichl
University of Heidelberg
r.scheichl@uni-heidelberg.de

MS171

Multidisciplinary High Performance Data Analysis Forum: summary and recommendations

The Multidisciplinary High Performance Data Analysis for Societal Challenges Forum held in Saint Girons (France) on June 17-19, 2018 was organized by the Informatics Research Institute of Toulouse (IRIT) and the Lawrence Berkeley National Laboratory (LBNL), with the support of the LabEx CIMI (Centre International de Mathmatiques et Informatique de Toulouse). The participants included decision-makers in research laboratories and academy, all faced with problems related to societal challenges. The forum had an international dimension with the participation of representatives from the USA, Japan and Europe. The planning of the forum was motivated by the realization that the digital transformation is of great concern for public and private research organisations. This forum will be convened yearly, such that participants will be able to take stock, share experiences, and propose dynamic approaches for organizing R&D communities. The first Forum sought to foment discussions about ongoing or yet to-be-implemented mechanisms for data collection and processing related to the societal challenges. It is important to assess their potential impact, particularly on the management of research teams working on societal challenges. The outcome of the first forum is summarized in recommendations that will be presented during the talk.

Michel Dayde
CNRS, IRIT
michel.dayde@irit.fr

Pierre-Henri Cros
IRIT - University of Toulouse - CNRS
pierre-henri.cros@irit.fr

Osni A. Marques
Lawrence Berkeley National Laboratory
Berkeley, CA
oamarques@lbl.gov

MS171

Obtaining Performance from a Julia-Implementation of Trilinos Data Librairies

Julia is a high level language designed to match the performance of code written in C++. Because Julia is a high level language, optimizing Julia code differs from optimizing C++ code. To study Julia's ability to perform in large-scale, distributed settings, Trillinos's Petra Object Model was implemented in Julia. The Julia version of the Petra Object Model was able to run as fast as ePetra, and out-perform an existing Julia distributed linear algebra library, in an implementation of the power method to com-

pute eigenvalues.

Neil Lindquist
St. Johns University
nslindquist@csbsju.edu

MS171

Combining Extreme Computing and Big Data for Future Machine Learning

Machine Learning applications have experienced spectacular improvements in recent times, thanks to accelerators (GPUs) that can better process larger available data. Nevertheless, at the moment only very few of those applications use supercomputers. It is expected that the announced exascale architectures and the future accessible data will open the road to strategic applications with strong societal impacts. Nevertheless, we face several difficult challenges, starting with efficient programming paradigms for exascale computing that could be used to tackle future applications. Therefore, it is important to begin to evaluate the required whole ecosystem, to efficiently develop extreme scale intelligent machine learning, and related milestones. Efficiently combining extreme scale computing and big data should be the first requirement to guide research on new methods for intelligent applications. In this talk, we survey some of those challenges and summarize some ideas to set a road map to extreme scale future machine learning, based first on combining high parallel and distributed computing with big data science computing. We list several machine learning methods that we think could be adapted to extreme scale computers, and we anticipate that linear algebra will still be essential in that context.

Serge Petiton
University of Lille
CNRS
serge.petiton@univ-lille1.fr

Kesheng Wu
NERSC, Lawrence Berkeley National Lab
kwu@lbl.gov

Osni A. Marques
Lawrence Berkeley National Laboratory
Berkeley, CA
oamarques@lbl.gov

MS171

A Scalable Randomized Singular Value Decomposition with Multiple Sketches for Big Data Analysis

Singular value decomposition (SVD) is an essential tool to perform dimension reduction and low-rank approximation. When the data size is big, randomized SVD with single or multiple sketches can nicely approximate the leading singular values and singular vectors with certain mild assumptions. The rapid growth in the size of matrices further increases the need for developing such large-scale SVD algorithms. Here, we consider a Monte Carlo type integrated SVD algorithm based on multiple random sketches. The proposed integration algorithm takes multiple random sketches and then integrates the results obtained from the multiple sketched subspaces. While the integrated SVD can achieve higher accuracy and lower stochastic variations, the computational cost, however, is higher due to the multiple sketches and integration process. We will discuss how multiple level parallelism can shorten the elapsed runtime of these operations. We modify the algorithms and data structures to increase the scalability and reduce communication costs. Numerical results suggest that the proposed approach can approximate the leading singular pairs of huge matrices and achieve near linear scalability with respect to the matrix size and number of nodes on a hybrid CPU-GPU computer cluster. This is a joint work with Ting-Li Chen and Su-Yun Huang at the Institute of Statistical Science, Academia Sinica, Dawei D. Chang, and Mu Yang at the Institute of Applied Mathematical Sciences, National Taiwan University.

Weichung Wang
National Taiwan University
Institute of Applied Mathematical Sciences
wwang@ntu.edu.tw

MS172

Leveraging Seismic Modeling with Mixed Precision Arithmetic

Seismic modeling aims at generatingsynthetic seismograms bysimulating the propagation of seismic wave in a givengeological medium. Seismic modeling is used for the design of seismic acquisition geometriesand plays an essential role in seismic imaging, inversion and interpretation. Among the numerous approaches to seismic modeling, direct methods based on approximating the geological model by a numerical mesh and solving the discretized wave equation are of particular interest. Finite Difference Time Domain (FDTD) methods are usually used to solve the wave equation. Although FDTD allows for significant vectorization on modern multi-core architectures, the algorithm is limited by an inherently low arithmetic intensity and its performance is bound by the available memory bandwidth. Since memory bandwidth grows slower than compute, the performance of stencil kernels will not scale with increasing compute density. To increase the arithmetic intensity we explore the use of lower floating point precision to store the data involved in seismic modeling. The computations are done in single precision. For the velocity stress formulation, we study the impact of using half precision for each of the parameters and show that by selecting the right parameters we may obtain important speedup with limited impact on the produced seismograms. Using lower floating point precision also reduces the memory requirement for every shot-point improving the total throughput of a given system.

Rached Abdelkhalak
KAUST
rached.abdelkhalak@kaust.edu.sa

Hatem Ltaief
King Abdullah University of Science & Technology (KAUST)
hatem.Ltaief@kaust.edu.sa

Kadir Akbudak

Bilkent University
kadir.cs@gmail.com

Vincent Etienne, Thierry Tonellot
ARAMCO
vincent.etienne@aramco.com,
thierrylaurent.tonellot@aramco.com

David E. Keyes
King Abdullah University of Science and Technology (KAUST)
david.keyes@kaust.edu.sa

MS172

High-performance Asynchronous Execution of the Reverse Time Migration for the Oil and Gas Industry

We take the performance of the Reverse Time Migration (RTM) method to a new level by relying on a dynamic runtime system to schedule the various tasks of the RTM (e.g., stencil computation kernel, Perfectly Matched Layer computations, I/O operations, image condition calculations, etc.) on systems equipped with hardware accelerators. The overall RTM application translates into an out-of-order execution, which opens up new opportunities to further overlap expensive and non-critical operations, such as I/O, with tasks which belong to the critical path, such as compute-bound GPU stencil kernel during the forward/backward modeling. Idle time is then reduced, while load balancing is achieved through work stealing on each node.

Amani Alonazi
KAUST
amani.alonazi@kaust.edu.sa

Hatem Ltaief
King Abdullah University of Science & Technology (KAUST)
hatem.Ltaief@kaust.edu.sa

Issam Said
PMC Univ Paris 06 and CNRS UMR 7606, LIP6
isaid@nvidia.com

Rached Abdelkhalak
KAUST
rached.abdelkhalak@kaust.edu.sa

David E. Keyes
King Abdullah University of Science and Technology (KAUST)
david.keyes@kaust.edu.sa

MS172

The ARM HPC Ecosystem for Oil and Gas Applications

Improving three-dimensional models of the subsurface is vital for oil exploration. This relies on a better understanding of complex geological structures to fully characterize the reservoir and optimize the drilling process. For this purpose, seismic imaging techniques play a major role by processing recorded signals from wave propagation. Due to the implementation of advanced numerical methods and the growing size of data available to build these models, the oil and gas industry is facing major HPC challenges.
In recent years, AArch64 architectures have gained traction in the HPC community. This is the case for Cavium ThunderX2 processor which is now used by major HPC vendors. Moreover, several ongoing efforts rely on Arm-based systems including the Japanese Post-K Computer. On the software ecosystem side, GCC, Clang and Arm HPC compilers are available to generate optimized code on Arm-based platforms.
From an application perspective, popular numerical kernels at the heart of seismic imaging algorithms may be leveraged by features such as the improved memory bandwidth or the upcoming Scalable Vector Extension (SVE). In this talk, we will therefore discuss opportunities for key oil and gas kernels to benefit from Arm architectures from a co-design perspective.

Fabrice Dupros
BRGM
ORLEANS, FRANCE
Fabrice.Dupros@arm.com

Conrad Hillairet
ARM
conrad.hillairet@arm.com

MS172

Evaluation of Compression Algorithms and Memory Hierarchies Impact Within Time Reversal Algorithms

Seismic imaging is a well-known data intensive process. The prestack seismic data are now reaching the petabytes scale and the 3D seismic images are also becoming extremely large when sorting common angle gathers for later amplitude variation analysis. In this work, we first evaluate the need and the impact of wavefield compression with respect to the wave equation implementation and to the next-generation memory/storage hardware. For those tests, we use an isotropic acoustic finite difference code. Based on this comparison of present compression algorithms, we evaluate analytically the impact on time reversal algorithm as reverse time migration (RTM) in where the management of the snapshots needed for the wavefields correlation will drastically impact the performance. To address this concern of snapshot management, different workarounds as optimal checkpointing (Symes, 2007), boundary only snapshots (Dussaud, 2009) and random velocity boundaries (Clapp, 2009) have been implemented to mitigate the IO demands at the cost of extra computations and extra memory space requirements. Regardless which tricks are employed, one main question remains: should we keep as much data as possible in memory, write it on disk, or recompute the wavefields when the data is required - or any combination of these three approaches. Compression of the wavefields is obviously becoming mandatory and has a direct impact on the computer architecture.

Philippe Thierry
Intel Corp
philippe.thierry@intel.com

sunny gogar, Mathieu Gontier, Alberto Villarreal, Yvan Zavarzin, Sergey Khlystov, Leonardo Borges
intel
sunny.l.gogar@intel.com, mathieu.gontier@intel.com, alberto.villarreal@intel.com, ivan.zavarzin@intel.com, sergey.khlystov@intel.com, leonardo.borges@intel.com

Essam Algizawy
Brightskies Inc
essam.algizawy@brightskiesinc.com

MS173

Analysis of Equity Markets: A Graph-theory Approach

In this paper, we develop a characterization of the structure of the US Stock Market by studying how correlations between the various stocks and sectors of the market fluctuate. Through this characterization, we hope to identify the strongest of stocks, and sectors and thus identify which investments are safest. This analysis will allow us to provide an alternate investment strategy for those wishing to avoid long term risk in equity markets. This is done using a correlation based graph representing the stock market. The central finding of this study is that transportation sector, a subset of the industrial sector, is the best indicator of the markets fluctuation. Mentors: Lofti Hermi, Florida International University lhermi@fiu.edu; Tova Brown, tovesmath@gmail.com

Joshua Abrams
University of Arizona
rabrams12@math.arizona.edu

José Celaya-Alcalá
Brown University
Jose_Celaya-Alcala@Brown.edu

Drew Baldwin, Ryan Gonda, Zhaoren Chen
University of Arizona
drewbaldwin@email.arizona.edu, gondar@email.arizona.edu, zhaorenchen@email.arizona.edu

MS173

Multiresolution Methods for Convolutional Neural Networks

Convolutional neural networks (CNN) are widely used for speech, image, and video recognition due to their benchmark performance. However, little theory exists for designing CNNs and CNNs typically depend explicitly on the resolution of the input data. In a CNN the convolution operators acting on an image can be related to a linear combination of differential operators which gives a continuous understanding of CNNs. Multiscale methods are used to efficiently solve PDEs using a family of fine and coarse grids. The continuous understanding of CNNs provides a way to implement multiscale methods on the convolution operators of a CNN. This can be used to efficiently handle images of different resolutions and to train on computationally cheaper lower resolutions. Preliminary results show that first training on lower resolutions and then applying multiscale methods results in the network weights converging to a different, more optimal local minima than without the use of multiscale methods. Mentor: Ruthotto, Lars, Emory University Department of Mathematics, lruthotto@emory.edu

Evan Scope Crafts
Emory University
escopec@emory.edu

MS173

Evaluating Tangling in High-order Meshes with Tangent Vectors

The generation of high-order meshes is essential to implementing high order numerical methods. High-order meshes more accurately represent geometric domains with fewer elements and are overall less computationally expensive. While high-order meshes seem like an effective solution to this problem, their curvilinear nature introduces a potential for self-tangling which may be difficult to quantify. One of the most commonly used metrics to evaluate high-order mesh quality is the Scaled Jacobian, which measures an elements distortion with respect to a straight-sided element. Even though it is one of the most popular metrics used to determine tangling, it has been shown to fail in cases of overly distorted elements. For this reason, a new metric for evaluating mesh tangling is required. In this presentation, I will propose a new criteria for determining whether a triangular high-order mesh element is tangled. For each pair of adjacent edges in a given element, my method uses the sign of the cross-product of these edges tangent vectors to determine whether tangling occurs at that node. I will present several examples which demonstrate the use of my method to detect tangling in two-dimensional elements. I will then discuss implications of the method and suggestions for how it may be applied to further research in quality metrics for high-order elements. Mentor: Suzanne Shontz, University of Kansas, shontz@ku.edu

Myra Dotzel, Mike Stees
University of Kansas
myradotzel@ku.edu, mstees@ku.edu

MS173

Calibration of the Ross Recovery Theorem to Real-world Data, and Tests of its Practical Value

The Ross Recovery Theorem challenges the commonly held thought that derivative prices do not contain useful predictive information about the distribution of financial variables (such as an equity index). The theorem recovers, under certain hypotheses on the character of the market, the subjective probability distribution of an equity index from current derivative prices. In this paper, building on the method of Backwell for extracting state prices from option prices, we develop a strategy for combining option data with the Recovery Theorem to estimate the subjective distribution. Using real-world data, we then investigate whether the Recovery Theorem yields predictive information and has practical value, concluding that the answer might be no. Mentor: Robert Kohn, Courant Institute of

Mathematical Sciences, kohn@cims.nyu.edu

Ling Lan
New York University
ll3178@nyu.edu

Zhengxu Li
Courant Institute of Mathematical Sciences, New York Univers
zhengxu.li@nyu.edu

MS173

Reducing Computation Time of Image Deblurring Algorithms Using Wavelets

Image deblurring is an ill-posed inverse problem and regularization techniques are needed to solve it to obtain meaningful solutions. In this talk, we will present a method that relies on only performing the deblurring on the blurred image compressed by using wavelets, to significantly reduced the computational cost involved. While some information is lost by ignoring the highest frequency portions of the compressed image, this research demonstrates that the details lost are small relative to the amount of noise present in the high frequency that were left out. in 1-dimension, but through future work may also be applied to 2- and 3-dimensional scenes. Mentor: Malena Espanol, University of Akron, mespanol@uakron.edu

Jonathan Wittmer
University of Akron
jdw112@zips.uakron.edu

MS174

Numerical Modeling and Simulation of Parachute Inflation by Tracking an Immersed Elastic Interface in Incompressible Flow

A mesoscale spring model based on Rayleigh-Ritz analysis is used to mimic the fabric surface as an elastic membrane in parachute simulation. The elastic interface structure is coupled with fluid solver through the impulse method under the front tracking framework. We will also discuss several challenging problems in this multi-physics system including turbulence modeling, fabric collision, parachutist coupling, and computational parallelization. We will show numerical proof of convergence, verification and validation of numerical components, and the programming design for the simulations of different air-delivery assemblies.

Xiaolin Li
Department of Applied Math and Stat
SUNY at Stony Brook
xiaolin.li@stonybrook.edu

Tengbo Yang, Brandon Balletine
Stony Brook University
Stony Brook, NY 11794
tengbo.yang@stonybrook.edu, brandon.balletine@stonybrook.edu

MS174

A Moment-of-Fluid Method for Diffusion Equations on Irregular Domains with Application to Problems in Multi-material Systems

Keywords: finite volume method; multi-material Moment-of-Fluid method; thermal boundary layer; rate of heat transfer.

ABSTRACT

A new numerical method is developed for the solution of the diffusion problem in a system of several materials with complex boundaries and discontinuous diffusion coefficients. Some of the complex boundaries might make up thin filamentary regions. The key features of such diffusion problems exist in approximating the heat transfer rate in: (a) The micro-region of the initial stages of nucleate boiling (b) the thin methane hydrate layer (c) the initial stages of freezing and (d) the initial stages of melting. The Moment-of-fluid (MOF) procedure is employed to reconstruct all the material(s) interface(s) and it has the following properties: (i) MOF enables accurate reconstruction of any number of materials; (ii) MOF reconstruction captures thin filamentary region without resorting to the adaptive mesh refinement (AMR) procedure;(iii) MOF is a volume preserving reconstruction. Our new method is tested on prototype multi-material diffusion problems which demonstrate the potential of our method to be an enabling technology in predicting heat transfer in materials and manufacturing processes which are of technological importance. It is a simple and robust method which is capable of preserving the energy presentation relation for each material.

Yang Liu
Florida State University
yliu@math.fsu.edu

Mark Sussman
Department of Mathematics
Florida State University
sussman@math.fsu.edu

Yongsheng Lian
Department of Mechanical Engineering
University of Louisville
yongsheng.lian@louisville.edu

M. Yousuff Hussaini
Department of Mathematics
Florida State University
yhussain@fsu.edu

MS174

Dynamics of Particles Trapped at Complex Fluid Interfaces

We numerically investigate the effect of finite interfacial deformations on the force and torque experienced by a single particle straddling a fluid interface under the influence of an external force. We focus on the case where the interface is devoid of any adsorbed species (e.g. surfactants) and where the interfacial distortions may be caused by the presence of either an external field (e.g. gravity) or the nonspherical shape of the floating particle (e.g. ellipsoid, cylinder). We use a diffuse interface approach based on

the Cahn-Hilliard theory to model the two-phase flow dynamics while the moving fluid-solid boundary, due to particle motion, is taken into account thanks to an Arbitrary Lagrangian Eulerian (ALE) formulation of the governing equations. Drag coefficients, both in translation and rotation (anisotropic particles), are computed for a variety of fixed contact angles, fluid viscosities and interfacial deformation amplitudes. We first validate our numerical implementation with previous simulations and theoretical works which mostly assumed a flat interface and small Reynolds numbers. We then present our predictions for deformed interfaces and explore a regime where inertial effects and time-dependent changes of the interface shape are important.

Mingfeng Qiu, James Feng, Jean-Christophe Loudet
Department of Mathematics
University of British Columbia
mqiu@math.ubc.ca, james.feng@ubc.ca, jcloudet@math.ubc.ca

MS175

Optimizing Artificial Viscosity using Machine Learning

Lagrangian-based algorithms for hydrodynamics employ artificial viscosity to control spurious oscillations that develop behind shocks. The simplest form of artificial viscosity introduces two additional terms into the discrete equations, which are proportional to the divergence of the velocity field (linear artificial viscosity), and to square of the divergence (quadratic artificial viscosity), respectively. However, in general, there is no systematic procedure dictating how to choose these two proportionality constants (control parameters). Moreover, in addition to post-shock ringing, incorrect parameter choices can lead to negative density profiles and other unphysical behavior in solutions. In this presentation, we develop a data-driven approach for automating artificial viscosity parameter optimization using several machine learning-based strategies. To illustrate, we compare the space of optimal values across a range of different shock intensities in Riemann problems.

Jason Albright
LANL
ejalbright@lanl.gov

Mikhail Shashkov, Nathan Urban
Los Alamos National Laboratory
shashkov@lanl.gov, nurban@lanl.gov

MS175

Isotropic Adaptive Mesh Refinement on Polyhedral Unstructured Meshes in 1, 2 and 3 Dimensions

I present an adaptive mesh refinement (AMR) capability that operates on polyhedral domain-decomposed meshes in 1, 2 and 3 dimensions and staggered Arbitrary Lagrangian-Eulerian (ALE) multi-material hydrodynamics. A suite of capabilities are required to support ALE-AMR including (1) mesh pre-refinement where the initial, coarse mesh is pre-refined based on user-defined criteria as an initial condition to an AMR simulation; (2) Lagrangian hydrodynamics solver for fluids and solids; (3) Lagrange-plus-remap or indirect ALE which requires smoothing of the mesh after the Lagrange step and remapping fields from unsmoothed to smoothed meshes; (4) dynamic load balancing (DLB) where the decomposition of the mesh onto multiple processors is adjusted dynamically to maintain equal per-core load; (5) mesh re-ordering which renumbers mesh indexes to manage indirect addressing cost as unstructured mesh connectivity is continuously altered; (6) an error indicator model that identifies where, when and how much the resolution should be altered; (7) the dynamic (de-)refinement of the polyhedral mesh based on error indicators; and (8) remapping of fields during (de-)refinement. The presentation will focus on the unstructured polyhedral mesh (re)generation problem, error indicators and coupling of AMR+ALE+DLB. Results for several test cases including Sod, Sedov, Triple-point, Rayleigh-Taylor and shocked helium bubble will be presented and compared to constant resolution results.

Mack Kenamond
LANL
kenamond@lanl.gov

MS175

The Distribution-Based Remapping of the Nodal Mass and Momentum Between Arbitrary Meshes for Staggered Arbitrary Lagrangian-Eulerian Hydrodynamics

We present a new distribution-based method for remapping nodal mass and momentum between arbitrary source (Lagrangian) and arbitrary target (rezoned) meshes for indirect staggered arbitrary Lagrangian-Eulerian hydrodynamics. The method is based on the following ideas: we first define the cell-centered momentum and mass on the source mesh and conservatively remap those cell-centered quantities from source to target meshes using an intersection-based remap. Next, we use a local constrained optimization approach for each cell of the target mesh to conservatively distribute cell mass and momentum between nodes of the cell. The new method is efficient, conservative, accurate and bound preserving and can be used for remapping between two arbitrary meshes which may have completely different connectivity. Using the new method we present results for several test problems: Sod, Sedov, Triple-point and Taylor-Green vortex. We compare our results with the standard flux-based remapping.

Mikhail Shashkov
Los Alamos National Laboratory
shashkov@lanl.gov

Mack Kenamond
LANL
kenamond@lanl.gov

MS175

A Higher Order Approximate Static Condensation Method for Multi-material Diffusion Problems

We introduce a higher order approximate static condensation method for the diffusion problem with discontinu-

ous diffusion coefficients. The method allows for a general polygonal mesh which is unfitted to the material interfaces. Moreover, the interfaces can be discontinuous across the mesh edges as typical for numerical reconstructions using the volume or moment-of-fluid methods. We apply a mimetic finite difference method to solve local diffusion problems and use P_1 (mortar) edge elements to couple local problems into the global system. The properties of the resulting algebraic system are discussed. It is demonstrated that the method is second order accurate on smooth solutions and performs well for problems with high contrast in diffusion coefficients. Experiments also show the robustness with respect to position of the interface against the underlying mesh.

Alexander Zhiliakov
University of Houston
Department of Mathematics
alex@math.uh.edu

Daniil Svyatskiy
Los Alamos National Laboratory
dasvyat@lanl.gov

Maxim A. Olshanskii
Department of Mathematics
University of Houston
molshan@math.uh.edu

Evgeny Kikinzon, Mikhail Shashkov
Los Alamos National Laboratory
kikinzon@lanl.gov, shashkov@lanl.gov

MS176

Discovering Unknown Physics and Enforcing Known Symmetries and Constraints with Machine Learning,

Accurate and efficient reduced-order models are essential to understand, predict, estimate, and control high-dimensional nonlinear systems in physics and chemistry. These models should ideally be generalizable, interpretable, and based on limited training data. This talk will explore the sparse identification of nonlinear dynamics (SINDy) approach to uncover interpretable reduced-order models for unsteady physics. First, we will discuss how it is possible to enforce known constraints, such as energy conserving quadratic nonlinearities in fluid dynamics, to essentially "bake in" known physics. Next, we will demonstrate that higher-order nonlinearities can approximate the effect of truncated modes, resulting in more accurate models of lower order than Galerkin projection. Finally, we will discuss the use of intrinsic measurement coordinates to build nonlinear models, circumventing the well-known issue of continuous mode deformation associated with methods based on the proper orthogonal decomposition. This approach will be demonstrated on several relevant systems in fluid dynamics with low-dimensional dynamics.

Steven Brunton, Nathan Kutz
University of Washington
sbrunton@uw.edu, kutz@uw.edu

Jean-Christophe Loiseau
ParisTech
loiseau.jc@gmail.com

Bernd Noack
LIMSI-CNRS
Technical University Braunschweig
bernd.noack@limsi.fr

MS176

Theory-guided Data Science: A New Paradigm for Scientific Discovery from Data

This talk will introduce theory-guided data science, a novel paradigm of scientific discovery that leverages the unique ability of data science methods to automatically extract patterns and models from data, but without ignoring the treasure of knowledge accumulated in scientific theories. Theory-guided data science aims to fully capitalize the power of machine learning and data mining methods in scientific disciplines by deeply coupling them with models based on scientific theories. This talk will describe several ways in which scientific knowledge can be combined with data science methods in various scientific disciplines such as hydrology, climate science, aerospace, and chemistry. To demonstrate the value in combining physics with data science, the talk will also introduce a novel framework for combining deep learning methods with physics-based models, termed as physics-guided neural networks, and present some preliminary results of this framework for an application in lake temperature modeling. The talk will conclude with a discussion of future prospects in exploiting latest advances in deep learning for building the next generation of scientific models for dynamical systems, where theory-based and data science methods are used at an equal footing.

Anuj Karpatne
University of Minnesota
karpa009@umn.edu

MS176

A Machine Learning Guided Stochastic Modelling of Biomolecule Systems

Computational modelling of biomolecular systems is centered around projecting the high-dimensional conformation dynamics onto a low-dimensional collective-variable space. In this talk, I will introduce a data-driven based model reduction framework for such systems. In particular, based on the Mori-Zwanzig formalism, the projected dynamics is casted into the generalized Langevin Equation (GLE). To numerically approximate the system, we further develop a variational inference algorithm to construct the free energy landscape, as well as data-driven approach to construct the memory kernel of GLE. The constructed models naturally characterize the non-local correlation and fluctuations arising from smaller scale interactions. The method is demonstrated in challenging problems including the non-equilibrium transition dynamics in physics, engineering and biological systems.

Lei Wu
School of Mathematical Sciences

Peking University, China
leiwu@pku.edu.cn

Huan Lei
Pacific Northwest National Laboratory
huan.lei@pnnl.gov

MS176

Physics-informed Learning for Multiscale Systems

A probabilistic deep learning framework incorporating physical constraints as prior knowledge is developed for uncertainty quantification in dynamic multiscale/multiphysics models. Encoding structured information into a learning algorithm enhances its ability to extract the informational content of the data thus yielding superior performance for inference and decision making in the Small Data regime. Two complementary developments will be introduced. Firstly, physics is incorporated at multiple levels of our architecture using novel physics-driven loss functions as well as PDE residuals to encourage stable, realizable learning. The proposed paradigm includes both physical parameters and latent variables learned using data as coarse grained variables. Secondly, an approach will be introduced that adapts machine learning tools to physical models rather than just using the former on data generated by the latter. Such an approach employs physical models by computing residuals without requiring any expensive forward simulations, and is able to retain the full physical structure that characterizes the underlying dynamic process. We consider application of these developments to forward and inverse problems.

Nicholas Zabaras
University of Notre Dame
nzabaras@gmail.com

Phaedon-Stelios Koutsourelakis
Technische Universität München
Germany
p.s.koutsourelakis@tum.de

Paris Perdikaris
University of Pennsylvania
pgp@seas.upenn.edu

Jianxun Wang
University of Notre Dame
jwang33@nd.edu

MS177

Data-driven Discovery of Nonlinear Dynamics

A major challenge in the study of dynamical systems is that of model discovery: turning data into models that are not just predictive, but interpretable. The recent sparse identification of nonlinear dynamical systems (SINDy) method is successfully able to discover nonlinear governing equations from data. However, SINDy requires full state measurements of the governing variables of the system, which may be unavailable in many applications. Additionally, the problem becomes computationally expensive for systems with multiple time scales. Here we present recent work for addressing such challenges.

Kathleen Champion
University of Washington, Department of Applied Mathematics
kpchamp@uw.edu

J. Nathan Kutz
University of Washington, Seattle
Dept of Applied Mathematics
kutz@uw.edu

Steven Brunton
University of Washington
sbrunton@uw.edu

MS177

Machine Learning Methods for Fluid Dynamics on Manifolds: Physics-Informed Gaussian Process Regression (PI-GPR) and PIV

Recent developments in machine learning methods provide some promising approaches for making inferences in the sciences and engineering. A common central challenge is to fuse information from multiple sources, knowledge domains, or even levels of resolution. For problems in fluid mechanics and related experimental measurements, we develop a class of Gaussian Process Regression (GPR) machine learning methods. To enhance inferences of the global velocity fields of hydrodynamic flows from localized measurements, we formulate GPR methods that make use of the principles of fluid mechanics and prior knowledge of material properties. We then demonstrate our approaches both in the domain of standard 3D hydrodynamics flows and in the context of 2D thin-film hydrodynamic flows on curved surfaces. We discuss applications of these techniques in the context of engineering problems in fluid mechanics, measurements in biophysics, and other problems involving contributions from the geometry of fields that reside on manifolds.

Ben J. Gross
UC Santa Barbara
bgross09@yahoo.com

Paul J. Atzberger
University of California-Santa Barbara
atzberg@gmail.com

MS177

Model Compression and Fast Prediction for LSTM

LSTM has been widely used for NLP applications, such as language modeling and machine translation. However, the prediction time and model size of LSTM is proportional to the vocabulary size, which can grow to hundres of thousands. In this talk, I will discuss a novel compression method to reduce the LSTM model size by block-wise low-rank approximation, and another algorithm to improve the prediction time by exploiting the clustering structure of the output vectors of LSTM cells.

Cho-Jui Hsieh

University of California, Davis
chohsieh@ucdavis.edu

MS177

Enforcing Constraints for Interpolation and Extrapolation in Generative Adversarial Networks

Generative Adversarial Networks (GANs) are becoming popular choices for unsupervised learning. At the same time there is a concerted effort in the machine learning community to expand the range of tasks in which learning can be applied as well as to utilize methods from other disciplines to accelerate learning. With this in mind, we suggest ways to enforce given constraints in the output of a GAN both for interpolation and extrapolation. The constraints can be both algebraic and differential and respect the game-theoretic setup of GANs. Illustrative examples will be presented.

Panos Stinis
Pacific Northwest National Laboratory
panagiotis.stinis@pnnl.gov

MS178

Auxiliary Space Preconditioning for Mixed Finite Element Discretizations of Richards Equation

We seek to solve the steady state solution of Richard's Equation in three dimensions using HAZMATH, a finite element discretization and solver package written by Ludmil Zikatanov, Xiaoze Hu, and James Adler. We use a standard Picard iteration to solve the steady state problem and discretize using $RT0-P0$ finite elements with standard preconditioners. We then use the same Picard linearization and consider an auxiliary space preconditioner using EAFE to precondition and solve each Picard iteration. Physically realistic simulations are computed and analyzed and nonlinear time dependence is also considered.

Juan Batista
Department of Mathematics
The Pennsylvania State University
jxb1131@psu.edu

MS178

Enrichment for Multilevel Solvers

We have incorporated into the PETSc libraries tools for enriching the coarse basis of a multilevel solve, using a multilevel eigensolver to produce basis functions adapted to the level solves. We will discuss the tradeoffs encountered in maintaining sparsity and construct solvers for challenging problems with highly variable coefficients.

Matthew Knepley
Univ. of Buffalo
knepley@buffalo.edu

MS178

An Adaptive Unsmoothed Aggregation Multigrid Method Based on Path Cover

We propose a path cover adaptive algebraic multigrid (PC-αAMG) method for solving linear systems of weighted graph Laplacians and can also be applied to discretized second order elliptic partial differential equations. The PC-αAMG is based on unsmoothed aggregation AMG (UA-AMG). To preserve the structure of smooth error down to the coarse levels, we approximate the level sets of the smooth error by first forming vertex-disjoint path cover with paths following the level sets. The aggregations are then formed by matching along the paths in the path cover. In such manner, we are able to build a multilevel structure at a low computational cost. The proposed PC-αAMG provides a mechanism to efficiently re-build the multilevel hierarchy during the iterations and leads to a fast nonlinear multilevel algorithm. Traditionally, UA-AMG requires more sophisticated cycling techniques, such as AMLI-cycle or K-cycle, but as our numerical results show, the PC-αAMG proposed here leads to nearly optimal standard V-cycle algorithm for solving linear systems with weighted graph Laplacians. Numerical experiments for some real world graph problems also demonstrate PC-αAMG's effectiveness and robustness, especially for ill-conditioned graphs.

Junyuan Lin
Tufts University
junyuan.lin@tufts.edu

MS178

Robust Multilevel Preconditioners for a New Stabilized Discretization of the Poroelastic Equation

In this talk, we present block preconditioners for a stabilized discretization of the poroelastic equations developed recently. The discretization is proved to be well-posed with respect to the physical and discretization parameters, and thus provides a framework to develop preconditioners that are robust with respect to such parameters as well. We construct these preconditioners for both the stabilized discretization and a perturbation of the stabilized discretization that leads to a sparser overall problem. Numerical tests confirm the robustness of the block preconditioners with respect to the physical and discretization parameters. If time permits we will also discuss the development of monolithic multigrid methods for these problems.

Peter Ohm, James H. Adler, Xiaozhe Hu
Tufts University
peter.ohm@tufts.edu, james.adler@tufts.edu,
xiaozhe.hu@tufts.edu

Ludmil Zikatanov
Pennsylvania State University
ludmil@psu.edu

Carmen Rodrigo, Francisco José Gaspar
University of Zaragoza
carmenr@unizar.es, fjgaspar@unizar.es

MS179

High-order Discontinuous Galerkin Discretization for Multi-physics Simulations of Plasmas

Plasma simulations require capturing rich physics on a

wide range of scales. What is more, the mean time between particle collisions can be large in comparison to the time scales of interest and the assumption of thermodynamic equilibrium used in fluid simulations is no longer valid. A discretization of the full phase space is then necessary. High-order discontinuous Galerkin methods from the finite element family are well suited for the discretization of the Vlasov-Maxwell and the Boltzmann-Maxwell systems. This work presents a modal discontinuous Galerkin implementation in the Gkeyll plasma framework. Emphasis is placed on an algorithm utilizing modal Serendipity basis, precise initialization of different species on various meshes, and an efficient implementation of boundary conditions capturing surface physics.

Petr Cagas
Aerospace and Ocean Engineering
Virginia Tech
pcagas@vt.edu

Ammar Hakim
Princeton Plasma Physics Laboratory
ahakim@pppl.gov

James Juno
IREAP
University of Maryland
jjuno@umd.edu

Bhuvana Srinivasan
Aerospace and Ocean Engineering
Virginia Tech
srinbhu@vt.edu

MS179

Output-based Mesh Optimization for Embedded Discontinuous Galerkin Methods

We present a mesh optimization algorithm for accurately predicting functional quantities of interest from finite-dimensional solutions to partial differential equations. The target discretization is the embedded discontinuous Galerkin (EDG) method, which augments the standard DG method with additional unknowns on faces, edges, and nodes: these become the only globally coupled unknowns in the system. Optimization of the computational mesh for output prediction must account for the contribution of the face state approximation to the numerical error in the solution, as well as the computational cost of additional faces created by refinement. We present such an optimization approach, driven by discrete adjoints, for steady-state fluid dynamics simulations governed by the compressible Navier-Stokes equations. We also compare the resulting optimized meshes and error/cost balance to the standard DG discretization to assess the benefits of EDG in a mesh-adaptive setting.

Krzysztof Fidkowski
University of Michigan
kfid@umich.edu

Guodong Chen
Aerospace Engineering department
University of Michigan, Ann Arbor
cgderic@umich.edu

MS179

High-order Hybridized Discontinuous Galerkin Method and a Scalable Solver for Incompressible Magnetohydrodynamics

We present a high order hybridized discontinuous Galerkin (HDG) method and a scalable solution strategy for the incompressible resistive Magnetohydrodynamics. Incompressible resistive Magnetohydrodynamics presents challenges in the form of incompressibility constraints on both velocity and magnetic fields and requires special care in solution strategies for obtaining scalability. In particular we will focus on block preconditioning strategies in the context of linear systems arising from HDG discretizations and show how we can get scalable results. The block preconditioners requires efficient solution strategies for small subsystems and we will compare AMG and a multilevel solver for efficient solution of these subsystems. We will show various numerical results including the island coalescence problem to verify our proposed approach.

Sriramkrishnan Muralikrishnan
Department of Aerospace Engineering
The University of Texas at Austin
sriramkrishnan@utexas.edu

Stephen Shannon
The University of Texas at Austin
stephenshannon1120@gmail.com

Tim Wildey
Optimization and Uncertainty Quantification Dept.
Sandia National Laboratories
tmwilde@sandia.gov

Tan Bui
The University of Texas at Austin
tanbui@ices.utexas.edu

John N. Shadid
Sandia National Laboratories
jnshadi@sandia.gov

MS180

An Immersogeometric FSI Framework for the Modeling and Simulation of Transcatheter Heart Valves

Transcatheter heart valves (THVs) have emerged as a minimally invasive alternative to surgical bioprosthetic heart valves therapy. THVs offer advantages such as less post-operative pain, faster rehabilitation, and better pressure gradients. However, issues such as paravalvular leakage, leaflet fatigue, and valve migration limit the widespread use of THV in the younger population, especially due to the lack of data concerning its long-term performance and durability. In this work, we develop a novel computational fluid-structure interaction (FSI) framework for the modeling and simulation of THVs. To account for physiological realism, methods are proposed to model and couple the main components of the system, including the arte-

rial wall, blood flows, heart valve leaflets, and the stent. For a better integration between design and analysis, the computational framework is developed based on the isogeometric and immersogeometric analysis. The proposed FSI framework is applied to study the performance of various transcatheter valve designs and the effectiveness of the models will be discussed. The application of the proposed method to other medical devices such as the endograft will also be demonstrated.

Ming-Chen Hsu
Department of Mechanical Engineering
Iowa State University
jmchsu@iastate.edu

Michael C. H. Wu, Heather Muchowski
Iowa State University
mikechwu@iastate.edu, heatherm@iastate.edu

MS180

The Dynamic Augmented Lagrangian Method: Application to Immersogeometric Heart Valve Analysis and Implementation using FEniCS

This talk covers the development and application of a family of numerical methods for boundary (and coupling) condition enforcement in time-dependent partial differential equation systems. We call these methods dynamic augmented Lagrangian (DAL) approaches. In particular, we apply DAL to the problem of fluid–structure interaction (FSI) in heart valves, and discuss a variety of connections to existing numerical methods for FSI, while highlighting new stabilization techniques and numerical analysis that emerge from the DAL concept. A key feature of DAL-based FSI is that it is straightforward to implement using existing software elements for fluid and shell structure simulation. We support this claim by reviewing an implementation of heart valve FSI analysis based on the FEniCS finite element toolchain, augmented with the recent open-source library tIGAr for isogeometric analysis (IGA). Using IGA for the structure subproblem allows us to directly immerse a geometrical model of the valve into a background fluid discretization, which we refer to as immersogeometric analysis. By leveraging FEniCS and tIGAr, immersogeometric heart valve FSI simulation can be implemented within a few hundred lines of Python code.

David Kamensky
Brown University
david_kamensky@brown.edu

Ming-Chen Hsu
Department of Mechanical Engineering
Iowa State University
jmchsu@iastate.edu

Yue Yu
Department of Mathematics, Lehigh University
yuy214@lehigh.edu

Yuri Bazilevs
School of Engineering
Brown University
yuri_bazilevs@brown.edu

MS180

Coupled Peridynamic and Subsurface Flow Models for Environmentally-assisted Crack Growth in Concrete Structures

Damage evolution in concrete infrastructure can be strongly influenced by the hydrology of the surrounding environment. Developing high-fidelity models of such systems is challenging because it requires integration of physics and solution strategies from multiple disciplines. In this presentation, we discuss our current effort to couple nonlocal peridynamic solid-mechanics models with subsurface flow and transpose models from the geophysics community. Our goal is a multiphysics formulation that allows for coupling of the Peridigm peridynamics code with the PFLOTRAN code for multiphase flow and transport. This requires extension of the peridynamic model to incorporate environmental effects, and modification of the flow model to capture the effects of an evolving fracture network. Water saturation levels, as determined by the flow model, provide input to the peridynamic bond-failure rule. Mechanical response, including crack initiation and propagation, is then captured by the peridynamic model and the resulting discrete fracture network geometry is passed back to the flow model. The result is a multiphysics model incorporating mechanisms of environmentally-assisted fracture, branching and intersecting fracture networks, and mechanistic multiphase flow.

David Littlewood, Jessica Rimsza, Jennifer Frederick
Sandia National Laboratories
djlittl@sandia.gov, jrimsza, jmfrede@sandia.gov

Reese Jones
Co-author
rjones@sandia.gov

MS180

Computational Methods for Fluidstructure Interaction with Applications

Abstract not available.

Padmanabhan Seshaiyer
George Mason University
pseshaiy@gmu.edu

MS181

Fast Recovery from Node Failures for the Parallel Preconditioned Conjugate Gradient Method

We compare and refine exact and heuristic fault-tolerance extensions for the *preconditioned conjugate gradient* (PCG) and the *split preconditioner conjugate gradient* (SPCG) methods for recovering from failures of compute nodes of large-scale parallel computers. In the *exact state reconstruction* (ESR) approach, which is based on a method proposed by Chen (2011), the solver keeps extra information from previous search directions of the (S)PCG solver, so that its state can be fully reconstructed if a node fails unexpectedly. ESR does not make use of checkpointing or

external storage for saving dynamic solver data and has only negligible computation and communication overhead compared to the failure-free situation. In exact arithmetic, the reconstruction is exact, but in finite-precision computations, the number of iterations until convergence can differ slightly from that of the failure-free case due to rounding effects. We perform experiments to investigate the behavior of ESR in floating-point arithmetic and compare it to the heuristic *linear interpolation* (LI) approach by Langou et al. (2007) and Agullo et al. (2016), which does not have to keep extra information and thus has lower memory requirements. The experiments show that ESR, on average, has essentially zero overhead in terms of additional iterations until convergence, whereas the LI approach incurs much larger overheads on average.

Markus Levonyak, Carlos Pachajoa
University of Vienna
markus.levonyak@univie.ac.at,
carlos.pachajoa@univie.ac.at

Wilfried N. Gansterer
Department of Computer Science
University of Vienna
wilfried.gansterer@univie.ac.at

MS181

Adapting Scientific Software and Designing Algorithms for Next Generation GPU Computing Platforms

Over the past decade, the massively parallel hardware architecture of GPUs has enabled them to become the de facto accelerator of choice to address high performance computing needs across a broad range of scientific and technical computing domains. The adaptation of existing software and algorithms to so-called heterogeneous computing systems composed of a mix of CPU and GPU processors poses many challenges to implementors, which have been addressed in part through a decade of improvements to programming abstractions, languages, and libraries. State-of-the-art systems such as ORNL Summit and the NVIDIA DGX-2 bring significant hardware advances that have begun to outpace progress the programming side, and one can forecast a number of likely directions for ongoing hardware evolution that will create additional needs for further programming abstractions, language features, and algorithm design methodologies. This talk will discuss past and ongoing experiences with the adaptation of technical and scientific software for GPU-accelerated platforms and the challenges they embody, and will describe challenges associated with emerging and future hardware capabilities and the their impact on algorithm and software design.

John Stone
Beckman Institute
Univ of Illinois at Urbana-Champaign
johns@ks.uiuc.edu

MS181

Neural-inspired Computing at Sandia Labs Enabling and Performing Advanced Computation

To continue to meet computational demands amidst the slowing of Moore's Law as well as other scaling factors, researchers have turned towards alternative computing approaches to reduce power-per-computation metrics and to develop more efficient algorithms and architectures. Taking inspiration from the brain, neural-inspired algorithms and architectures have emerged as a promising approach to both enable and perform advanced computation. At Sandia Labs we have focused upon exploring what computational advantages neural-inspired concepts such as operating in a spiking paradigm may enable. Spiking neural networks are innately parallel and comprised of event-driven lightweight computational units which in aggregation can compute sophisticated functions. Additionally, spiking neurons operate in a temporal regime offering complexity advantages. Beyond simply using neural networks as universal function approximators, we show that spiking neural circuits can be used to compute a suite of computations including canonical optimizations, fundamental machine learning algorithms, and scientific calculations. And finally, we discuss how coupling neural-inspired algorithms with emerging low-power neuromorphic architectures provides a path forward to achieve requirements for next-generation computing platforms.

Craig Vineyard
Sandia National Laboratories, USA
cmviney@sandia.gov

MS181

Scientific Computing in a Changing Landscape

In the early 2000s, due to constraints on economical heat dissipation, clock speeds of single-core CPUs could no longer be increased, which marked the adoption of multi-core CPUs, together with a paradigm shift to algorithms specifically designed for parallel architectures. About 15 years into this architectural cycle and on its way to exascale performance, the computing industry finds itself at the confluence of technical difficulties that cast doubt on its ability to sustain this architectural model beyond the exascale capability. These difficulties are driving the hardware industry and computational scientists to develop application-specific chips and to look beyond silicon-based technology (e.g., quantum computing, physical annealing, neuromorphics, etc.), with a continued emphasis on raw processing power while addressing concerns about energy efficiency. Hardware specialization will likely redefine the development of computational algorithms over the next two decades for a wide range of important applications: large-scale PDE-based problems, artificial intelligence, computational chemistry, and optimization problems, to name just a few. The pressure to decrease time to solution or improve simulation fidelity, or both, for these applications will continue unabated. Sitting between hardware designers and application users, the computational science & engineering community will play a pivotal role to commercialize future computing technologies.

Laurent White
ExxonMobil
laurent.white@exxonmobil.com

Dimitar Trenev
ExxonMobil Research and Engineering
dimitar.trenev@exxonmobil.com

MS183

Model-order Reduction for the Optimization of Distribution Grids with Battery Control

We focus on the optimization of power networks at the distribution grid level. At this level, newly developed renewable energy sources such as wind turbines or photovoltaics are modeled as time-dependent PQ-buses that need to be included in the classical steady-state solution of the nonlinear power flow equations. Moreover, we include storage units (batteries) at a few grid nodes with controllable active power. The flexibility of the batteries is used to control the overall line loading of the distribution grid and to react on the changes of power demand and generation in time, which can exceed the physical dimensions the lines were designed for due to more sources and higher loads (e-mobility). Our numerical experiments demonstrate on the one hand the benefits of battery usage with respect to grid stability and, on the other hand, we present model-order reduction techniques for the more efficient numerical solution of the resulting constrained optimization problem.

Manuel Baumann
Max Planck Institute
for Dynamics of Complex Technical Systems
baumann@mpi-magdeburg.mpg.de

Sara Grundel
MPI Magdeburg
grundel@mpi-magdeburg.mpg.de

MS183

Newton-Krylov Methods for Solving the AC Load Flow Equations

We discuss the Newton-Krylov method to solve the AC load flow equations in a power system. These equations yield the nodal voltage amplitude and phase angle such that the line power flows are congruent with the location of the loads and generators. Newton-Krylov methods deliver the computational efficiency required to tackle problems in the design of future power systems. Such problems are challenging either due to scale or due to the repetitive nature. An example of the former is the modeling of the pan-European net with high resolution. Examples of the latter are contingency and optimization problems. Newton-Krylov also allow to distribute computations in problems in which the propriety of data is important to preserve. In this talk we illustrate both the aspect of efficiency and of distributed computation.

Domenico Lahaye
Delft Institute for Applied Mathematics
d.j.p.lahaye@tudelft.nl

MS183

Hierarchical Modeling and Simulation of Power Networks

In recent years, the increasing integration of renewable energy sources into the existing power networks have lead to several challenges. In particular, the increased dynamics and distributed structure of the systems make control and optimization tasks difficult, and the current approach to describe the dynamics of power grids using stationary models is not sufficient [A. van der Schaft and T. Stegink, Perspectives in modeling for control of power networks, Annual Reviews in Control 41 2106 119–132]. To address these problems, we propose a hierarchy of power network models, using port-Hamiltonian descriptor systems [C. Beattie, V. Mehrmann, H. Xu and H. Zwart, Port-Hamiltonian descriptor systems, arXiv:1705.09081v2, 2017]. Such port- and energy-based systems have many good properties, for example they are more robust to perturbations, and well suited for discretization, model order reduction and moduralization, allowing the application of simpler or more detailed models to different parts of the network or intervals of time. The systems are analyzed and numerical experiments are presented to illustrate the performance of numerical methods applied to these systems.

Riccardo Morandin
Technical University of Berlin
morandin@math.tu-berlin.de

MS183

Four Mathematical Formulations of the (Optimal) Power Flow Problem and Their Impact on the Performance of Solution Methods

The power system is designed and organized in a way that the total power consumption and the total power generation are in balance at all time. In order to keep this balance, grid operators perform the Power Flow (PF) and the Optimal Power Flow (OPF) computations along with other tools in system planning, operation and real-time control of the grid. Furthermore, the PF computation determines the steady- state behavior of the network whereas the OPF computation is used to obtain the optimal operational state of the electrical power system while satisfying system constraints and control limits. Especially in the Smart Grid (SG), where due to dynamic renewable energy sources rapid changes occur, these problems need to be solved efficiently. Moreover, grid operators need the fast, scalable and robust PF and OPF solvers at the same time because the size and complexity of the network are increasing in the SG. In this talk, I will explain that depending on the formulations of power flow equations (power or current balance) and specification of the coordinates (Polar or Cartesian), PF and OPF problems can be formulated mathematically in four different ways for a single physical formulation. Furthermore, I will address how different mathematical formulations result in different performances for the solution methods such as the Newton power flow method for the PF computation and the Interior Point Method for the OPF computation.

Baljinnyam Sereeter
Delft Institute for Applied Mathematics
b.sereeter@tudelft.nl

Kees Vuik, Cees Witteveen
Delft University of Technology
c.vuik@tudelft.nl, c.witteveen@tudelft.nl

MS184

Entropy Stable Schemes for Multicomponent Flows

The development of robust numerical schemes is a pacing research item in turbulence. In this context, the family of Entropy Stable (ES) schemes pioneered by Tadmor in 1987 has gained a lot of attention over the last decade. ES schemes have the interesting property of ensuring that the second principle of thermodynamics is never violated, regardless of whether the discontinuities encountered are physical or the pure product of under-resolution. In this talk, we consider the use of ES schemes in multicomponent flows simulations. A high-order ES finite volume scheme is developed and tested on a number of non-reacting and reacting test problems.

Ayoub Gouasmi
University of Michigan
Department of Aerospace Engineering
gouasmia@umich.edu

Karthik Duraisamy
University of Michigan Ann Arbor
kdur@umich.edu

Scott Murman
NASA Ames Research Center
scott.murman@nasa.gov

MS184

Entropy Stable Discretizations of Compressible Flows in Thermochemical Non-equilibrium

Entropy stable numerical methods for compressible flow have been demonstrated to exhibit better robustness than purely linearly stable methods and need less overall artificial dissipation for long simulations in subsonic and transonic flows. In this work we seek to extend these benefits to flows which require Thermochemical Non-Equilibrium. We first derive an entropy function which symmetrizes the compressible flow equations for multicomponent, reacting systems in thermochemical nonequilibrium with up to three temperatures. Based on this entropy function we develop an affordable, entropy-conservative inviscid flux function in conservation form. In addition, we discuss the impact of diffusion model selection on symmetrization and entropy stability, considering rigorous models of irreversible thermodynamics as well as simplified models of greater practical interest.

Michael Hansen, Travis Fisher
Sandia National Labs
mahanse@sandia.gov, tcfishe@sandia.gov

MS184

A Fully-Discrete Kinetic Energy Preserving and Entropy Conservative Scheme for Compressible Flows

The construction of entropy conservative/stable numerical schemes for conservation laws, such as the compressible Euler equations, is crucial to single out a physically relevant weak solution. Furthermore, a faithful representation of kinetic energy is of key importance for compressible flows, especially for direct numerical simulation of turbulence. In this talk, we propose a fully-discrete finite difference scheme for the Euler equations, which is both entropy conservative and kinetic energy preserving. This is achieved by using an implicit Crank-Nicolson type scheme, and the evaluation of the numerical fluxes at a carefully crafted time-averaged vector of entropy variables. When used in conjunction with the Navier-Stokes equation, a suitable discretization of the viscous fluxes leads to a kinetic energy preserving and entropy stable scheme for the viscous model. Several numerical results are present to demonstrate the performance of the proposed scheme.

Deep Ray
EPFL-SB-MATH-MCSS
deep.ray@epfl.ch

Praveen Chandrashekar
TIFR Centre for Applicable Mathematics
Bangalore
praveen@tifrbng.res.in

MS184

Entropy Stability and the Computation of Entropic Measure-valued Solutions

Entropy stability plays an important role in the dynamics of non-linear systems of conservation laws. We revisit the general framework of numerical entropy stability for finite-difference/finite-volume approximations of such nonlinear systems. Our approach is based on comparing numerical viscosities with those of entropy conservative schemes. We demonstrate this approach with host of high order entropic schemes. In particular, this paradigm serves as the building block for a class of ENO-based entropic schemes of arbitrarily high-order of accuracy, called TeCNO schemes.
Recent evidence shows that entropic solutions need not be unique in more than one-spatial dimension. Instead, entropic measure-valued solutions can be interpreted in an averaged sense, as an ensemble average in configuration space.Numerical experiments provide a remarkable evidence for the effectiveness of the TeCNO schemes in the computation of such entropy measure valued solutions.
Joint work with U. Fjordholm, R. Kappeli and S. Mishra.

Eitan Tadmor
University of Maryland
USA
tadmor@cscamm.umd.edu

MS185

Fluxo: An Open Source Parallel Split-Form DG Solver for Compressible Navier-Stokes and Resistive MHD Equations

The class of Discontinuous Galerkin (DG) methods is a

promising candidate for large scale simulations of nonlinear problems found in fluid dynamics or magnetohydrodynamics. DG methods combine high order accuracy and flexible unstructured meshes. Elements couple only to direct face neighbors, via numerical fluxes. This locality makes DG well suited for parallelization. The DG spectral element method (DGSEM) is an efficient variant for three-dimensional curved unstructured meshes with hexahedral elements. Integration and interpolation points (Gauss-Lobatto) are collocated, leading to a tensor-product structure of the element-local operator that significantly reduces the cost for high polynomial degrees. However, aliasing errors will occur when collocation methods are applied to nonlinear problems, which can lead to simulations crashing for low resolutions. More recently, following discrete stability proofs for DGSEM, so-called split forms lead to a significantly increased robustness. They require more operations but preserve the tensor-product structure using the same data. In this talk, we present the open-source code FLUXO (github.com/project-fluxo/fluxo). We will discuss the efficient implementation of the robust split forms on curved three-dimensional elements, especially regarding vectorization on Intel KNL and Skylake. We also give details about the parallelization concept to overlap non-blocking communication with local computations.

Florian Hindenlang
Max-Planck Institute for Plasma Physics, NMPP division Garching
florian.hindenlang@ipp.mpg.de

Eric Sonnendrücker
Max-Planck-Institut f{̈u}r Plasmaphysik, NMPP division, and
Technische Universit{̈a}t M{̈u}nchen
eric.sonnendruecker@ipp.mpg.de

Gregor Gassner, Andrew R. Winters
Mathematical Institute
University of Cologne
ggassner@math.uni-koeln.de, awinters@math.uni-koeln.de

MS185

Performance Portable Implementation of High Order DG Methods through Code Generation

Explicit SIMD vectorization is one of the key challenges in achieving good float- ing point performance on modern HPC platforms. Typically, compiler-based auto- vectorization is only applicable if the problem exhibits a favorable structure. However, domain knowledge may be used to identify alternate sources of parallelism to be used for SIMD vectorization at the cost of writing code, that is tuned specifically to the PDE problem at hand and to the target architecture. We solve the maintainability issue of such codes through a code generation approach: Explicitly vectorized finite element as- sembly kernels to be used with the discretization framework dune-pdelab are generated from a DSL (UFL) describing the finite element assembly problem. We demonstrate the power and flexibility of the approach for high order DG methods on hexahedra, ex- ploiting the tensor product structure of basis functions and quadrature formulae through sum factorization. Performance studies for the Intel Skylake microarchitecture using the AVX512 instruction set are shown.

Dominic Kempf, René Heß, Steffen Müthing
IWR, Heidelberg University
dominic.kempf@iwr.uni-heidelberg.de,
rene.hess@iwr.uni-heidelberg.de,
steffen.muething@iwr.uni-heidelberg.de

Peter Bastian
Interdisciplinary Center for Scientific Computing
University of Heidelberg
peter.bastian@iwr.uni-heidelberg.de

MS185

Vectorisation for High-order Simplicial Elements

The use of SIMD vectorisation in combination with high-order methods is now well-established, with many development efforts now focusing on matrix-free formulations of key finite element operators, which avoid the memory bandwidth constraints of per-element matrices. In particular, attention has been drawn to how tensor products of one dimensional shape functions, when combined with the operator-reduction technique of sum-factorisation, can utilise wide SIMD units to attain near-peak performance on modern CPU architectures. So far however, only quadrilateral and hexahedral elements have been considered under this regime. Although these elements offer exceptional performance, from a practical perspective, it is difficult at best to construct all-hexahedral meshes for complex geometries. Unstructured meshes of triangles, tetrahedra and prisms therefore offer an alternative route to avoid this issue, but in the classical spectral element formulation using Lagrange interpolants, this leads to a non-tensorial expansion in shape functions, thereby losing the performance gains found by using sum-factorisation. In this talk, we demonstrate how an alternative choice of basis can be combined with vectorisation to construct an efficient matrix-free discretisation of elliptic operators for mixed element types. We present the results of a performance analysis on modern CPU architectures that demonstrates the efficacy of this scheme when applied to complex geometries.

David Moxey
Department of Aeronautics
Imperial College London
d.moxey@exeter.ac.uk

Roman Amici
Scientific Computing and Imaging Institute
University of Utah
amicir@gmail.com

Mike Kirby
School of Computing
University of Utah
kirby@cs.utah.edu

MS185

SIMD Vectorization for High-order Matrix-free Fi-

nite Element Computations in CFD

In our group, we have recently extended a solver for the incompressible Navier-Stokes equations (NSE) to the compressible NSE with explicit time stepping. The solver makes use of the framework for matrix-free evaluation of finite element operators, contained in the deal.II finite element library. One of the central ingredients is the SIMD vectorization over several cells or faces, based on C++ wrapper classes around intrinsics. Previous performance analysis has mostly concentrated on simple cases, such as the SIMD-vectorized evaluation of the Laplace operator. A more complex discrete operator for compressible NSE is the focus of this presentation: scalar and vectorial conserved quantities as well as additional derived quantities are processed for multiple elements at once due to explicit vectorization, which leads to higher pressure on the caches. Published preliminary results have shown that performance independent of the polynomial degree can be maintained for degrees up to 6. I will present a node level performance analysis including roofline models and cache level transfer measurements on modern CPU hardware (incl. Skylake) for corner cases where the size of the cell local data exceeds the capacity of a cache level. I will also show strong scaling results proving the parallel efficiency of our high-order matrix-free CFD code despite of increased granularity due to vectorization.

Peter Munch, Niklas Fehn, Wolfgang A Wall
Technical University of Munich
munch@lnm.mw.tum.de, fehn@lnm.mw.tum.de,
wall@lnm.mw.tum.de

Martin Kronbichler
Technical University of Munich
Institute for Computational Mechanics
kronbichler@lnm.mw.tum.de

MS186

Probabilistic Numerics and Applications to Roundoff Error

Probabilistic numerical methods are numerical methods which output probability measures which are designed to describe, in a structured way, the level of uncertainty in the output. To date, most of these measures have been designed to capture the discretization error of the method, but another important source of uncertainty which has thus far received less attention is roundoff error. This talk will introduce probabilistic numerical methods and discuss the ways in which they have been, and could be, constructed to describe roundoff error.

Jon Cockayne
The University of Warwick
j.cockayne@warwick.ac.uk

Chris Oates
Newcastle University
Chris.Oates@newcastle.ac.uk

Tim Sullivan
University of Warwick
tim.sullivan@warwick.ac.uk

Mark Girolami
Imperial College London
m.girolami@imperial.ac.uk

Han Cheng Lie
Free University of Berlin / Zuse Institute Berlin
hlie@math.fu-berlin.de

MS186

Probabilistic Error Analysis for Inner Products

The objective is a probabilistic model for the floating point errors in an inner product computation of two real n-vectors. We consider the relative forward error, and present probabilistic perturbation and roundoff error bounds. The perturbations are represented as independent, discrete, bounded random variables, and the derivation of the perturbation bound is based on Azuma's inequality for independent random variables. In contrast, the roundoff errors, although also represented as discrete, bounded random variables, are not required to be independent. The derivation of the roundoff error bounds is based on constructing a Martingale that mirrors the particular order of computations, followed by an application of Azuma's inequality for Martingales. Numerical experiments confirm that our bounds are much more realistic (often by several orders of magnitude) than existing deterministic worst-case bounds – even for small vector dimensions n and stringent success probabilities.

Ilse Ipsen
North Carolina State University
Department of Mathematics
ipsen@ncsu.edu

MS186

A New Approach to Probabilistic Roundoff Error Analysis

Probabilistic approaches to roundoff error analysis were proposed as long ago as the 1960s to obtain error bounds that are in better agreement with the errors observed on average than the traditional, worst-case bounds. In this talk, we reassess the relevance of this probabilistic approach to roundoff error analysis in the light of large-scale, low-precision computations, while, at the same time, formalizing its conclusions by computing rigorous probabilistic bounds based on a minimal number of assumptions. Our main contribution is to formally prove the well-known rule of thumb that the constants in the traditional bounds in the analysis of inner products (and other algorithms based on them) can be replaced by their square root. Our result applies to a very wide class of algorithms, including essentially all numerical linear algebra algorithms.

Nicholas J. Higham
School of Mathematics
The University of Manchester
nick.higham@manchester.ac.uk

Theo Mary
Universite de Toulouse, UPS-IRIT

theo.mary@manchester.ac.uk

MS186

Stochastic Analysis and Correction of Floating Point Errors in Monte Carlo Simulations

In this talk we will show how the floating point errors in the simulation of SDEs (stochastic differential equations) can be modelled as stochastic. Furthermore, we will show how these errors can be corrected within a multilevel Monte Carlo approach which performs most calculations with low precision, but a few calculations with higher precision. The same procedure can also be used to correct for errors in converting from uniform random numbers to approximate Normal random numbers. Numerical results will be generated on both CPUs (using single/double precision) and GPUs (using half/single precision).

Oliver Sheridan-Methven
University of Oxford
oliver.sheridan-methven@maths.ox.ac.uk

Mike Giles
University of Oxford, United Kingdom
Mike.Giles@maths.ox.ac.uk

MS187

Hamilton Jacobi Formulation of Vlasov Poisson and a Method of Lines Transpose Approach

Abstract not available.

Andrew J. Christlieb
Michigan State Univerity
Dept. of Comp. Math., Sci & Engr.
christli@msu.edu

William A. Sands
University of Akron
sandswi3@msu.edu

MS187

Neutron Kinetics in Liquid-fueled Nuclear Reactors

Abstract not available.

Kathryn Huff
University of Illinois at Urbana–Champaign
kdhuff@illinois.edu

MS187

Assessment of the Lagrange Discrete Ordinates Equations for Three-dimensional Neutral Particle Transport

The Lagrange Discrete Ordinates (LDO) equations, developed by Ahrens as an alternative to the traditional discrete ordinates formulation, have been implemented in Denovo, a three-dimensional radiation transport code developed by Oak Ridge National Laboratory. The LDO equations retain the formal structure of the classical discrete ordinates equations typically used in radiation transport calculations but treat particle scattering in a different way. Solutions of the LDO equations have an interpolatory structure such that the angular flux can be naturally evaluated at directions other than the discrete ordinates used in arriving at the solutions, and the ordinates themselves may be chosen in a strategic way for the problem under consideration. Of particular interest is that the LDO equations have been shown to mitigate ray effects at increased angular resolutions. We present scalar flux solutions of the LDO equations for a small number of test cases of interest and compare the results against flux solutions generated using standard quadrature types as well as against experimental flux measurements from a benchmark test case.

Kelly L. Rowland
National Energy Research Scientific Computing Center (NERSC)
kellyrowland@lbl.gov

Cory Ahrens
X Theoretical Design Division, Primary Physics Group
Los Alamos National Laboratory
coryahrens@gmail.com

Steven Hamilton
ORNL
hamiltonsp@ornl.gov

Rachel Slaybaugh
UC Berkeley
Dept. of Nuclear Engineering
slaybaugh@berkeley.edu

MS187

Piecewise P_N Approximations for Multidimensional Radiative Transfer

Linear transport equations are ubiquitous in many application areas, including as a model for neutron transport in nuclear reactors and the propagation of electromagnetic radiation in astrophysics. The main computational challenge in solving the linear transport equations is that solutions live in a high-dimensional phase space that must be sufficiently resolved for accurate simulations. The three standard computational techniques for solving the linear transport equations are the implicit Monte Carlo, discrete ordinate (S_N), and spherical harmonic (P_N) methods. We focus on the P_N approximation, which is based on expanding the part of the solution that depends on velocity direction (i.e., two angular variables) into spherical harmonics. A big challenge with the P_N approach is that the spherical harmonics expansion does not prevent the formation of negative particle concentrations. In this work, we introduce an alternative formulation of the P_N approximation that hybridizes aspects of both P_N and S_N. Although our basic scheme does not guarantee positivity of the solution, the new formulation allows for the introduction of local limiters that can be used to enforce positivity. We first develop our scheme and limiting strategy on the one-dimensional linear transport equations. We then show how to extend this idea to the multidimensional case using unstructured grids in phase space. The resulting scheme is validated on

several standard test cases.

Minwoo Shin
Iowa State University
mws@iastate.edu

James A. Rossmanith
Iowa State University
Department of Mathematics
rossmani@iastate.edu

MS188

HIGAMod Approximation for the Stokes Equations in Patient-specific Geometries

In the field of hemodynamics, numerical models have evolved to account for the demands in speed and accuracy of modern diagnostic medicine. Methods have incorporated different reduction techniques to perform, with the same level of precision, the computation of solutions in the constraints of time and computation power available in most diagnostics centers.In this context, we studied the Hierarchical Model (HiMod) reduction technique, a reduction procedure to describe phenomena characterized by an intrinsic directionality, combined with Isogeometric Analysis (HigaMod).In hemodyamics, such a directionality is provided by the blood flow.Hierarchical model reduction showed a significant improvement in reducing the computational power and simulation time, while ensuring enough reliability in modeling the problem at hand.Recently, we are applying HigaMod reduction to the approximation of the Stokes equations in patient-specific geometries.This represents the actual goal of this communication.In more detail,after setting suitable procedures to reconstruct generic geometries of the patients starting from multiple 2D X-ray coronary arteriogram projections, we resort to HigaMod to simulate the blood flow in coronary and cerebral arteries. The main goal is to numerically assess the improvements provided by the isogeometric hierarchical discretization with respect to standard 1D models in terms of computational efficiency, without waiving the accuracy typical of (full) 3D models.

Yves Barbosa
Mox - Politecnico di Milano
yvesantonio.brandes@polimi.it

MS188

A Localized Reduced-order Modeling Approach for PDEs with Bifurcating Solutions

Although Reduced-order modeling (ROM) has been successfully used in many settings, ROMs built specifically for the efficient treatment of PDEs having solutions that bifurcate as the values of input parameters change have not received much attention. In such cases, the parameter domain can be subdivided into subregions, each of which corresponds to a different branch of solutions. Popular ROM approaches such as proper orthogonal decomposition (POD), results in a global low-dimensional basis that does not respect the often large differences in the PDE solutions corresponding to different subregions. We developed and tested a new ROM approach specifically aimed at bifurcation problems. In the new method, the k-means algorithm is used to cluster snapshots so that within cluster snapshots are similar to each other and are dissimilar to those in other clusters. This is followed by the construction of local POD bases, one for each cluster. The method also can detect which cluster a new parameter point belongs to, after which the local basis corresponding to that cluster is used to determine a ROM approximation. Numerical experiments show the effectiveness of the method both for problems for which bifurcation cause continuous and discontinuous changes in the solution of the PDE.

Martin Hess
SISSA, International School for Advanced Studies
Trieste, Italy
mhess@sissa.it

Alessandro Alla
PUC-Rio, Department of Mathematics
Rio de Janeiro, Brazil
alla@mat.puc-rio.br

Annalisa Quaini
Department of Mathematics, University of Houston
quaini@math.uh.edu

Gianluigi Rozza
SISSA University, Italy
Trieste, Italy
grozza@sissa.it

Max Gunzburger
Florida State University
Department of Scientific Computing
mgunzburger@fsu.edu

MS188

An Approximated Minimum Residual Reduced Basis Method for the Parameterized Incompressible Navier-Stokes Equations

We present a reduced basis technique for the parameterized incompressible Navier-Stokes equations. The key features of the approach are (i) an approximate minimum residual approach based on the introduction of an empirical test space [Taddei, *An offline/online procedure for dual norm calculations of parameterized functionals: empirical quadrature and empirical test spaces*, submitted], and (ii) a residual-based weak-Greedy strategy for the construction of a reduced space to simultaneously estimate velocity and pressure. We present a theoretical discussion to motivate the approach; in particular, we provide a complete stability analysis for a linearized problem. We further present several numerical results to show the effectivity of the proposed approach.

Tommaso Taddei
INRIA Bordeaux
taddei.tommaso@gmail.com

MS188

Reduced Order Methods for Parametric Optimal Flow Control in Coronary Bypass Grafts: Patients

Specific Data Assimilation and Geometrical Reconstruction

Bypass grafting is a common surgical treatment for diseases causing arterial stenosis or aortic coarctation, etc. In this talk, we aim at presenting a complete pipeline from clinically provided data to optimal control on blood flow in coronary artery bypass grafts, to investigate the haemodynamics in grafts and their corresponding arteries. The proposed framework is assimilated with patient-specific geometrical and physiological data, and thanks to flow control, proficiently bridges clinical measurements with numerical simulations. Furthermore, we have incorporated reduced order methods to achieve cost-efficient and reliable solution of such parametric problems. Some numerical results, verifying our methodology, will be discussed.
This work is in collaboration with Piero Triverio (University of Toronto) and Laura Jiménez-Juan (Sunnybrook Health Sciences Centre, Toronto).

Zakia Zainib, Francesco Ballarin
SISSA International School for Advanced Studies
zzainib@sissa.it, fballarin@sissa.it

Gianluigi Rozza
SISSA University, Italy
Trieste, Italy
grozza@sissa.it

MS189

Goodness-of-fit Testing of Copulas using Quasi-Monte Carlo Methods

Multivariate models with dependent variables are popular in financial industry, geostatistics, hydrology, insurance mathematics, medicine, and reliability theory. Simulation of copulas can be done by Monte Carlo or quasi-Monte Carlo methods. Goodness-of-fit tests can be used to find the best simulation algorithms for copulas. Since most of the goodness-of-fit tests are designed for univariate distributions, dimension reduction or bootstrapping is normally used in existing tests. We introduce a new goodness-of-fit test based on the collision test and low-discrepancy sequences, and present numerical results to compare its efficiency with some current tests.

Yiran Chen
Department of Mathematics
Florida State University
yc15t@my.fsu.edu

MS189

QMC Methods for Elliptic PDEs Driven by White Noise

When solving partial differential equations driven by additive spatial white noise, the efficient sampling of white noise realizations can be challenging. In this talk we focus on the efficient sampling of white noise using quasi-random points in a finite element method and multilevel Quasi Monte Carlo (MLQMC) setting. This work is an extension of previous research on white noise sampling for MLMC. We express white noise as a wavelet series expansion that we divide in two parts. The first part is sampled using quasi-random points and contains a finite number of terms in order of decaying importance to ensure good QMC convergence. The second part is a correction term which is sampled using standard pseudo-random numbers. We show how the sampling of both terms can be performed in linear time and memory complexity in the number of mesh cells via a supermesh construction. Furthermore, our technique can be used to enforce the MLQMC coupling even in the case of non-nested mesh hierarchies. We demonstrate the efficacy of our method with numerical experiments.

Matteo Croci
Mathematical Institute
University of Oxford
matteo.croci@maths.ox.ac.uk

Mike Giles
University of Oxford, United Kingdom
Mike.Giles@maths.ox.ac.uk

Marie E. Rognes
Simula Research Laboratory
meg@simula.no

Patrick E. Farrell
Mathematical Institute
University of Oxford
patrick.farrell@maths.ox.ac.uk

MS189

The Start of Community Supported Quasi-Monte Carlo Software

Abstract not available.

Fred J. Hickernell
Illinois Institute of Technology
Dept of Applied Math & Ctr for Interdisc Sci Comput
hickernell@iit.edu

MS189

Computing Multivariate Normal Variance Mixture Distributions with Quasi-Monte Carlo Methods

An efficient method for computing the distribution function and density of multivariate normal variance mixtures is presented. The method is based on previous work of A. Genz and F. Bretz for computing the distribution function of a multivariate normal and t distribution but is generalized to arrive at an efficient randomized quasi-Monte Carlo algorithm suited for all multivariate normal variance mixture distributions whose radial part has a tractable generalized inverse. It is demonstrated through numerical experiments that this algorithm is fast even for dimensions as high as 1000. Furthermore, applications include the evaluation of normal variance mixture copulas as well as a likelihood-based fitting procedure for normal variance mixture models. Software has been developed and is provided in the R package nvmix.

Erik Hintz, Marius Hofert, Christiane Lemieux
Faculty of Mathematics
University of Waterloo

erik.hintz@uwaterloo.ca, marius.hofert@uwaterloo.ca, clemieux@uwaterloo.ca

MS190

Boundary Element Methods for High-intensity Focused Ultrasound Simulations

High-Intensity Focused Ultrasound (HIFU) is a promising technology for the non-invasive treatment of a range of modalities. However, medical treatment planning is essential for the effective application of HIFU. In this talk we discuss linear high-frequency wave models for treatment planning in the lower abdomen, which are implemented through coupled systems of integral equations. We consider sound-hard and penetrating models for bones and tissue, and discuss the effectiveness of different preconditioning strategies at realistic frequencies.

Timo Betcke, Pierre Gelat
University College London
t.betcke@ucl.ac.uk, p.gelat@ucl.ac.uk

Seyyed Reza Haqshenas
University College London
Department of Mechanical Engineering
s.haqshenas@ucl.ac.uk

Nader Saffari
Department of Mechanical Engineering
University College London
n.saffari@ucl.ac.uk

MS190

On Accurate Evaluation of Oscillatory Integrals and Functions

Our goal is to construct efficient functional representations of oscillatory integrals (or functions) within a finite, user-supplied accuracy. For example, instead of a Filon-type approach of using polynomials to approximate a mildly oscillatory phase and amplitude of the classical oscillatory integral, we use appropriate changes of variables and robust methods for nonlinear approximation of functions via exponentials. The resulting functional representations allow us to evaluate the integral with the prescribed accuracy and with the computational cost that does not depend or depends only mildly on the size of the parameter responsible for the oscillatory behavior. We will discuss our approach and some of its applications.

Gregory Beylkin
University of Colorado, Boulder
beylkin@colorado.edu

MS190

The Windowed Green Function Method and Other Specialized High-frequency Techniques in Computational Electromagnetism

We present fast spectral electromagnetic solvers that address some of the main difficulties associated with the simulation of realistic engineering problems in the frequency- and time-domain. Based on a novel windowing approach in conjunction with fast high-order methods for evaluation of integral operators and new approaches for accurate Fourier expansion of non-periodic functions, these algorithms can solve, with high-order accuracy, problems of electromagnetic and acoustic propagation and scattering for large and complex three-dimensional geometries such as aircraft, submarines, and ships; they can be applied to challenging fluid-dynamic and seismic problems; and they can be used to design and optimize optical and photonic devices and metamaterials of very large electrical size. A variety of applications will be presented which demonstrate the significant capabilities inherent in the new methods, as well as the improvements these algorithms can provide, over other methods, in generality, accuracy, and speed.

Oscar P. Bruno
California Institute of Technology
obruno@caltech.edu

MS191

Dictionary-based Reduced Basis Method via Random Sketching

We propose a dictionary-based manifold learning method for solving parameter-dependent linear systems of equations. For each value of the parameter, the solution is approximated by a projection onto a subspace spanned by several vectors from a set of candidate basis vectors (so-called dictionary). The solution vector and parameter-dependent reduced subspace are selected such that the associated residual is minimum. From an algebraic point of view, the proposed approach can be formulated as a parameter-dependent sparse least-squares problem. Solving such a problem with standard tools may require very heavy computations and be numerically unstable. The aforementioned issues can be easily circumvented with random sketching technique. It is shown that the solution can be accurately (with high probability) estimated from a random sketch (introduced in [Balabanov and Nouy 2018]) associated with the dictionary. The sketch can be efficiently precomputed during the offline stage in any computational environment and then operated with online. A rigorous analysis of the dimension of the random sketch required to obtain a given level of accuracy is presented. Finally, we propose an efficient greedy-like procedure for the dictionary generation. Since all our algorithms require computations only on a small sketch but not the dictionary itself, therefore maintaining and operating with large data sets in the offline stage can be easily avoided.

Oleg Balabanov
Ecole Centrale de Nantes
Polytechnic University of Catalonia
olegbalabanov@gmail.com

Anthony Nouy
Ecole Centrale Nantes
anthony.nouy@ec-nantes.fr

MS191

Greedy Controllability of Reduced-order Linear

Dynamical Systems

Often a dynamical system is characterized by one or more parameters describing physical features of the problem or geometrical configurations of the computational domain. As a consequence, by assuming that the system is controllable, corresponding to different parameter values, a range of optimal controls exists. The goal of the proposed approach is to avoid the computation of a control function for any instance of the parameters. The greedy controllability [M. Lazar, E. Zuazua *Greedy controllability of finite dimensional linear systems Automatica*, (Journal of IFAC) Volume 74 Issue C, Pages 327-340, 2016.] consists in the selection of the most representative values of the parameters that allows a rapid approximation of the control function for any desired new parameter value, ensuring that the system is steered to the target within a certain accuracy. By proposing the Reduced Basis method [J.S.Hesthaven, G. Rozza, B. Stamm, *Certified Reduced Basis Methods for Parametrized Partial Differential Equations*, SpringerBriefs in Mathematics, 2016.] in this framework, the computational costs are drastically reduced and the efficiency of the greedy controllability approach is significantly improved.

Giulia Fabrini
University of Konstanz
giulia.fabrini@uni-konstanz.de

Laura Iapichino
TU Eindhoven
l.iapichino@tue.nl

Martin Lazar
University of Dubrovnik
martin.lazar@unidu.hr

Stefan Volkwein
Universität Konstanz
Department of Mathematics and Statistics
stefan.volkwein@uni-konstanz.de

MS191

Reduced-order Models: Convergence Between Data and Simulation

How to revisit computational schemes using new data-based modeling approaches in order to reduce computational costs, improve the accuracy of simulation codes in situations where the physical model is incomplete, allow realignment with respect to experimental tests...? This is a widespread issue in industry since despite advances in high-performance computing, numerical models remain heavy and their exploitation to explore the parametric space is prohibitive. One approach allowing considerable acceleration of this exploration is model reduction from data. A reduced-order model (ROM) is based on an initial offline sampling phase to collect data. The dataset is then used to define a reduced approximation space and in the online phase, the PDE solution is sought in the reduced space. Each of the steps of a ROM implies modelling choices about how to effectively sample the solution space, how to define the approximation space from the sampling, how to determine the solution that best approximates the PDE in this space. These modelling choices rely on ideas from numerical analysis, geometry and optimisation: residual minimisation, optimal transport theory, interpolation on geodetics. Through the study of some realistic applications in industry and medicine, we will discuss the implementation of these reduced-order methods to the equations of Navier-Stokes, compressible Euler, BGK, concluding on the perspectives of convergence between data and EDP models.

Angelo Iollo
Institut de Mathématiques de Bordeaux
angelo.iollo@math.u-bordeaux1.fr

Michel Bergmann
Universite Bordeaux
michel.bergmann@inria.fr

Andrea Ferrero
Politecnico di Torino
andrea_ferrero@polito.it

Sebastien Riffaud
INRIA
University of Bordeaux
sebastien.riffaud@inria.fr

Edoardo Lombardi, Angela Scardigli, Haysam Telib
Optimad
edoardo.lombardi@optimad.it, angela.scardigli@optimad.it, haysam.telib@optimad.it

MS191

Adaptive Multiscale and Asynchronous Optimization Methods for Large-scale PDE Paramteter Estimation

Estimating parameters of Partial Differential Equations (PDEs) is of interest in a number of applications such as geophysical and medical imaging. Parameter estimation is commonly phrased as a PDE-constrained optimization problem that can be solved iteratively using gradient-based optimization. The problem is extremely computationally expensive as the PDEs must be solved numerous times before the model is reconstructed with sufficient accuracy. In this work, we introduce two approaches for reducing the computational burdens of large-scale PDE parameter estimation problems. The first approach aims at reducing the computational costs per PDE solve. To this end, we use an adaptive Multiscale Finite Volume (MSFV) technique which employs a parameter-dependent projection onto a nested coarse mesh in the optimization scheme. The second approach is an asynchronous Alternating Direction Method of Multipliers (ADMM) that aims at allowing massively parallel inversions with minimal latency and communication. Here, we present a weighting scheme based on the uncertainties of the model that accelerates the convergence of consensus ADMM. We exemplarily show that both approaches lead to computational savings for a 3D single-physics and multiphysics PDE parameter estimation problem.

Samy Wu Fung
Emory University
400 Dowman Dr, 30322 Atlanta, GA, USA

samy.wu@emory.edu

Lars Ruthotto
Department of Mathematics and Computer Science
Emory University
lruthotto@emory.edu

MS192

A Direct Filter Method for Parameter Estimation

The parameter estimation problem is one of the most important research topics in data assimilation. In this presentation, we shall introduce a novel parameter estimation method in which we consider the parameter as the state process in a nonlinear filtering problem and use the state model that contains the parameter to construct a pseudo-observation. We name this approach the direct filter method due to the fact that we use nonlinear filter algorithms to estimate parameters directly without estimating the state model as part of the solution in the nonlinear filtering problem.

Feng Bao
Florida State University
bao@math.fsu.edu

MS192

Bayesian Inference with Nonparametric Likelihood Functions

Abstract not available.

Shixiao W. Jiang
Institute of Natural Sciences
Shanghai Jiao-Tong University
suj235@psu.edu

John Harlim
Pennsylvania State University
jharlim@psu.edu

MS192

When is High-dimensional State Estimation Hard or Easy?

Abstract not available.

Chris Snyder
Mesoscale and Microscale Meteorology Division
NCAR
chriss@ucar.edu

MS192

A Unified Framework for Transportation Particle Filters

Abstract not available.

Peter Jan van Leeuwen
University of Reading and Colorado State University
p.j.vanleeuwen@reading.ac.uk

MS193

Optimal Actuator Selection for Airfoil Separation Control using Empirical Data

A systems-theoretic approach is presented for determining the optimal actuator location for separation control from numerical and experimental fluids data. Input-output data from a set of candidate locations is used to determine the optimal location among the set that can drive the system output to an arbitrary value with minimal control energy. The optimal actuator selection framework is formulated for generality and can be readily applied within the context of other systems, both stable and unstable. The approach is demonstrated using lift and separation-angle response data from high-fidelity numerical fluids simulations of an airfoil. Body-force actuation at six localized positions on the upper-surface of the airfoil are evaluated. The results of this study indicate that the optimal actuator locations for controlling lift and separation angle are not identical.

Supported by AFOSR grants FA9550-16-1-0392 and FA9550-17-1-0252.

Maziar S. Hemati
Department of Aerospace Engineering and Mechanics
University of Minnesota
mhemati@umn.edu

Debraj Bhattacharjee
University of Minnesota
bhatt234@umn.edu

Bjoern Klose
San Diego State University
bklose@sdsu.edu

Gustaaf Jacobs
Department of Aerospace Engineering
San Diego State University
gjacobs@mail.sdsu.edu

MS193

Flow Reconstruction using Manifold Learning

The problem of estimating the state of a physical system is ubiquitous in science. However measurements are almost always limited so that the high-dimensional state cannot be observed and the associated mathematical problem is ill-posed. Popular workarounds include dimension reduction and regularization by imposing some structure to the class of elements in which the estimation is sought. In this work, we rely on a purely data-driven approach and learn the map between extended measurements and the nonlinear manifold the state vector lies on. Specifically, we use time-embedding to address the non-Markovianity of the raw measurements. Combined with multi-kernel learning, it results in high-dimensional measurement features. The state vector nonlinear manifold is approximated using Diffusion Map and the map from measurement features to the estimated state is then the solution to a Sylvester equation. The methodology is illustrated with the estimation

of a fluid flow field from a few wall-mounted sensors.

Lionel Mathelin
LIMSI - CNRS, France
mathelin@limsi.fr

Srikanth Derebail Muralidhar, Berengere Podvin
LIMSI-CNRS
muralidhar@limsi.fr, podvin@limsi.fr

MS193

Physically-informed Bayesian Learning of Linear Embeddings for Fluid Dynamics Problems

Koopman operator theory offers an elegant framework to analyze nonlinear dynamical systems via the discovery and use of linear embeddings. In this work, we present a Bayesian deep learning framework for linear embeddings in a continuous dynamical system form. A key aspect of this approach involves the utilization of the underlying governing equation with data sampling. Moreover, we propose an analysis algorithm to select relevant Koopman eigenmodes that spans the state for EDMD and KDMD algorithms. Application examples are presented and compared with discrete and deterministic counterparts.

Shaowu Pan
University of Michigan, Ann Arbor
shawnpan@umich.edu

Karthik Duraisamy
University of Michigan Ann Arbor
kdur@umich.edu

MS193

Using Machine Learning for the Development of Closure Models for Multiphase Flows

Direct numerical simulations of multiphase flows and other complex dynamic heterogeneous systems, where the governing continuum equations are solved accurately for large systems, so that every statistical quantity can be computed, have reached a level of maturity where the question of what to do with the results is urgent. Such simulations produce unprecedented amount of data that challenges traditional approaches. Helping improve tools to predict the behavior of large-scale industrial processes is one obvious application. Models of such processes generally rely on reduced order models where the full governing continuum equations have been averaged of filtered to retain only large-scale motion. The filtering does, however, lead to a loss of information and the equations contain unknown terms that need to be closed by relating them to the averaged flow. In many cases those closure terms also depend on some properties describing the state of the unresolved flow that must be evolved along with the averaged properties. Traditionally, closure relationships have been developed using experimental data, dimensional analysis and simplified modeling. We will discuss how the use of machine learning is providing additional tools and opportunities, including for the development of closing relationships for relatively simple laminar flows, for determining the influence of various variables in turbulent flows, and for classifying more complex flows with topology changes.

Gretar Tryggvason
Johns Hopkins University
gtryggv1@jhu.edu

MS194

A Relaxed Physical Factorization Preconditioner for Mixed Finite Element Coupled Poromechanics

In this talk, we present a Relaxed Physical Factorization (RPF) preconditioner for the iterative solution of coupled poromechanics equations discretized by mixed finite elements. The preconditioner is obtained by using a proper splitting of the block system and setting a relaxation parameter α. An automatic procedure is presented for the optimal selection of α. Then, the RPF setup involves the computation of two local preconditioners for the inner blocks depending on α. Numerical experiments in both theoretical benchmarks and real-world poromechanical applications are presented to verify the theoretical properties, performance, and robustness. The proposed approach appears to be a promising alternative to existing techniques for large-size and ill-conditioned problems.

Matteo Frigo
Department ICEA
University of Padova, Padova, Italy
matteo.frigo.3@studenti.unipd.it

Nicola Castelletto
Stanford University
Department of Energy Resources Engineering
castelletto1@llnl.gov

Massimiliano Ferronato
University of Padova
ferronat@dmsa.unipd.it

MS194

Robust Preconditioners for the Biot's Model

Poroelasticity models the processes of coupled deformable porous media flow which is crucial in many applications. An essential component, and usually the most time-consuming part of simulating coupled PDEs, is solving the large-scale and ill-conditioned linear systems of equations arising from the discretization of the Biots model. In this work, we generalize the traditional framework of multilevel preconditioners on saddle point systems for the poroelasticity and develop effective preconditioners that are robust with respect to the physical and discretization parameters. Preliminary numerical experiments are presented to support the theory and demonstrate the robustness of our preconditioners.

James H. Adler
Tufts University
james.adler@tufts.edu

Francisco José Gaspar
University of Zaragoza
fjgaspar@unizar.es

Xiaozhe Hu, Peter Ohm
Tufts University
xiaozhe.hu@tufts.edu, peter.ohm@tufts.edu

Carmen Rodrigo
University of Zaragoza
carmenr@unizar.es

Ludmil Zikatanov
Pennsylvania State University
ludmil@psu.edu

MS194

Flexible Monolithic Solvers for Dynamic Poroelastic Wave Propagation

The higher-order mixed-multiphysics space-time finite element discretisation of the dynamic Biot-Allard equations including elastic wave propagation in a fluid-saturated porous media

$$\partial_t((1-\phi)\rho_s v + \phi\rho_f q) - \nabla\cdot\sigma(u,p) = f\,,$$

$\sigma(u,p) = C\epsilon(u) - pI$, $\epsilon(u) = (\nabla u + \nabla u^T)/2$, coupled fluid flow

$$\partial_t(c_0 p + \alpha\nabla\cdot u) + \nabla\cdot q = s$$

and $q = -K(\nabla p - \rho_f g - \partial_t v)$, results in huge linear systems with block structure. The unknowns are the displacement u, velocity v, fluid-pressure p and fluid flux q. The linear systems are solved in a fully-coupled monolithic way by applying the flexible GMRES iterative linear system solver. We present details on our efficient preconditioning strategy consisting of the application of multiple-step fixed-stress iterations in a multigrid in time setting. The performance of the automatically tuned preconditioning strategy and implementation in our distributed-parallel solver suite DTM++ for the deal.II library is analysed carefully with several challenging numerical experiments with relevance to physical problems.

Uwe Koecher
Helmut-Schmidt-Universitat
Hamburg, Germany
koecher@hsu-hamburg.de

MS194

Multiscale Computation of Multiphase Fluid Dynamics at the Pore Scale

Multiphase flow in porous media is pervasive in geoscientific applications including geologic CO2 storage, hydrocarbon recovery, and geothermal energy. They are also relevant in manufacturing of fuel cells and batteries. In order to achieve engineering control over the fluid dynamics in these materials, it is important to be able to optimize for the properties of the injected fluids and/or the microstructure of the porous medium. Geologic porous media pose additional challenges due to the inherent uncertainty associated with the subsurface. Direct numerical simulation of the governing physics (Navier-Stokes) does, in principle, provide a means to conduct such optimizations. However, in practice, simulations are prohibitively expensive limiting the applicability of current Digital Rock Physics (DRP) workflows. In this work, we present an adaptive multiscale framework designed specifically for the pore scale, which can achieve orders of magnitude speedup in single- and two-phase flow calculations. Unlike other order-reduction techniques (e.g., pore-network modeling), the approach allows one to estimate and control prediction errors. The framework is amenable to scalable parallelism and multiphysics (or hybrid) computations.

Yashar Mehmani
Energy Resources Engineering
Stanford University, Stanford, CA, USA
ymehmani@stanford.edu

Hamdi Tchelepi
Stanford University
Energy Resources Engineering Department
tchelepi@stanford.edu

MS195

A New Undergraduate Engineering Degree in Computational Engineering at UT Austin

Abstract not available.

Clint Dawson
Institute for Computational Engineering and Sciences
University of Texas at Austin
clint@ices.utexas.edu

Noel Clemens
UT Austin
clemens@mail.utexas.edu

MS195

CMDA: Math Modeling, Data Science, and HPC in Undergraduate Education

In Spring 2015, Virginia Tech launched a new undergraduate major in Computational Modeling and Data Analytics (CMDA), based on ten new courses developed by a team of faculty from Computer Science, Mathematics, Physics, and Statistics. The first students graduated with the CMDA degree in May 2017, and the program has continued to grow: as of January 2018, more than 400 Virginia Tech students are enrolled in the CMDA major. This talk will survey the CMDA curriculum, then address the process behind the curriculum design and implementation.

Serkan Gugercin
Virginia Tech
Department of Mathematics
gugercin@vt.edu

MS195

Applied and Computational Mathematics: BYU's New Degree for 21st Century Discovery and Innovation

We present BYU's Applied and Computational Math program, which provides students with a rigorous foundation in mathematics, statistics, and computation. This program was redesigned from the ground up with the hopes

of attracting top employers and graduate programs and thus becoming a national model for modern applied mathematics education. The first year of the program devotes rigorous study to the design, analysis, and optimization of algorithms, and gives students an arsenal of mathematical and statistical tools to explore the performance, complexity, and accuracy of algorithms. The second year focuses on the art and science of mathematical modeling, which gives students the ability to connect the real world with abstract mathematics and numerical simulation. In terms of their technical education, the students become well versed in numerical analysis, scientific computing, data processing, relational databases, advanced programming concepts, software development, big data analytics, and scientific visualization. Having recently concluded its 5th year, we have learned a lot about how to build a program, how to adapt to the needs of the students and the changes in technology, and how to adjust the program in real time to serve the needs of employers and other stakeholders.

Jeffrey Humpherys
Brigham Young University
jeffh@math.byu.edu

MS195

Experiences in CSE Graduate Education: 12 Years of AICES

The Aachen Institute for Advanced Study on Computational Engineering Science (AICES Graduate School) was established in 2006 in the framework of the German Excellence Initiative. AICES researchers conduct interdisciplinary research at the interface between mathematics, computer science, and engineering, and the graduate school itself is a collaborative effort of more than 25 institutes from 8 academic departments. AICES focuses on (1) supporting early career researchers, particularly junior research group leaders, and (2) offering a doctoral program for bachelor and master students enabling a shortened and attractive path to the doctorate. In this talk, we will present our experiences on interdisciplinary graduate education, and give an overview of AICES successes and lessons learned on aspects such as recruitment, outreach, communication, student advising and supervision, as well as collaboration with other academic institutions and industry.

Karen Veroy
Graduate School AICES
RWTH Aachen University
veroy@aices.rwth-aachen.de

MS196

Higher Order Globally Constraint-preserving FVTD and DGTD Schemes for Time-dependent Computational Electrodynamics

Abstract not available.

Dinshaw Balsara
Notre Dame University
Physics Department
Dinshaw.S.Balsara.1@nd.edu

MS196

High-order Bound-preserving Discontinuous Galerkin Methods for Compressible Miscible Displacements in Porous Media on Triangular Meshes

In this talk, we develop high-order bound-preserving (BP) DG methods for the coupled system of compressible miscible displacements on triangular meshes. We consider the problem with a multi-component fluid mixture and the concentration of the jth component, c_j should be between 0 and 1. There are three main difficulties. Firstly, c_j does not satisfy a maximum-principle. Therefore, we cannot apply the numerical techniques of X. Zhang and C.-W. Shu directly. The main idea is to apply the positivity-preserving techniques to all c_j's and enforce $\sum_j c_j = 1$ simultaneously to obtain physically relevant approximations. By doing so, we have to treat the $\frac{dp}{dt}$ as a source in the concentration equation and choose suitable fluxes in the pressure and concentration equations. Secondly, we use interior penalty DG methods for the concentration equations, and the first-order numerical fluxes are not easy to construct. So, we will construct second-order BP schemes and then combine the second- and high-order fluxes to obtain a new one which further yields positive numerical cell averages. Finally, we cannot apply the slope limiter because c_j's are not the conservative variables. Moreover, we cannot simply set the upper bound of each c_j to be 1. Therefore, a suitable limiter for multi-component fluid will be introduced. Numerical experiments will be given to demonstrate the high-order accuracy and good performance of the numerical technique.

Nattaporn Chuenjarern, Yang Yang, Ziyao Xu
Michigan Technological University
nchuenja@mtu.edu, yyang7@mtu.edu, ziyaox@mtu.edu

MS196

An Alternative Formulation of Discontinous Galerkin Schemes for Solving Hamilton-Jacobi Equations

In this talk, we will introduce an alternative formulation of discontinuous Galerkin (DG) schemes for approximating the viscosity solutions to nonlinear Hamilton-Jacobi (HJ) equations. The main difficulty in designing DG schemes lies in the inherent non-divergence form of HJ equations. One effective approach is to explore the elegant relationship between HJ equations and hyperbolic conservation laws: the standard DG scheme is applied to solve a conservation law system satisfied by the derivatives of the solution of the HJ equation. In this work, we consider an alternative approach to directly solving the HJ equations, motivated by a class of successful direct DG schemes by Cheng et al. [J. Comput. Phys., v223 (2009); J. Comput. Phys., 268(2014)]. The proposed scheme is derived based on the idea from the central-upwind scheme by Kurganov et al. [SIAM J. Sci. Comput., v23 (2001)]. In particular, we make use of precise information of the *local speeds of propagation* at discontinuous element interface with the goal of adding adequate numerical viscosity and hence naturally

capturing the viscosity solutions. A collection of numerical experiments is presented to demonstrate the performance of the method for solving general HJ equations with linear, nonlinear, smooth, non-smooth, convex, or non-convex Hamiltonians.

Wei Guo
Texas Tech University
weimath.guo@ttu.edu

MS196

Local Discontinuous Galerkin Method for Convection-diffusion Equations on Overlapping Meshes

In this talk, we will introduce a new local discontinuous Galerkin (LDG) method on overlapping meshes.The scheme is stable and has optimal error accuracy. It is well known that the traditional maximum-principle-preserving (MPP) LDG method is only available up to second-order accuracy. With the new scheme, we can construct third-order MPP LDG method. Numerical experiments will be given to demonstrate the stability and accuracy of MPP LDG method.

Jie Du
Tsinghua University
jdu@math.tsinghua.edu.cn

Yang Yang
Michigan Technological University
yyang7@mtu.edu

Eric Chung
The Chinese University of Hong Kong
Department of Mathematics
tschung@math.cuhk.edu.hk

MS197

Simulation of a Sparse Direct Solver on Heterogeneous Systems using Starpu and Simgrid

Due to the continual evolution of hardware trends, the high performance community has considered high-level programming paradigms to design scientific libraries. In the particular model of task-based programming, one of these popular high-level approaches, the accurate modeling of the tasks they are composed of is essential for many components such as scheduling algorithms, tuning or performance predictions. While most studies have so far considered homogeneous architectures or regular codes, we propose to model an irregular algorithm on an heterogeneous architecture. We illustrate our discussion with the qr_mumps sparse direct solver, based on a QR multifrontal factorization, running on top of the StarPU runtime system and the SimGrid simulation framework, on multicore processors enhanced with two types of modern GPU accelerators.

Emmanuel Agullo
INRIA
emmanuel.agullo@inria.fr

Alfredo Buttari
CNRS-IRIT-Université de Toulouse, France
alfredo.buttari@enseeiht.fr

Ian Masliah
INRIA bordeaux
ian.masliah@inria.fr

MS197

Exploiting Nested Task-based Parallelism in the Factorization of Hierarchical Matrices

Hierarchical matrices (H-matrices) lie in-between the dense and sparse scenarios. Therefore, it is natural to tackle the LU factorization of H-Matrices via a task-parallel approach, which has reported successful results in the recent past for related linear algebra problems. Concretely, in this talk we leverage some recent features in the OmpSs-2 programming model, such as support for weak operands and early release of dependencies, to considerably improve the parallel efficiency of H-LU factorizations arising from boundary element methods. Discovering the data-?ow parallelism intrinsic to the operation at execution time, via the analysis of data dependencies based on the memory addresses of the tasks operands, is especially challenging for H-matrices, as the data structures vary in dimension during the execution. We overcome this issue by decoupling the data structure from that used to detect dependencies.

Rocio Carratalá-Sáez
UNIVERSITAT JAUME I
rcarrata@uji.es

Enrique S. Quintana-Ortí
Universidad Jaume I
Spain
quintana@icc.uji.es

MS197

Exploiting Parameterized Task-graph in Sparse Direct Solvers

Task-based programming models have been widely studied in the context of dense linear algebra, but remains less studied for the more complex sparse solvers. In this talk, we will present the use of two different programming models: Sequential Task Flow from StarPU, and Parameterized Task Graph from PaRSEC to parallelize the factorization step of the PaStiX sparse direct solver. We will present how those programming models have been used to integrate more complex and finer parallelism to take into account new architectures with many computational units. Efficiency of such solutions on homogeneous and heterogeneous architectures with a spectrum of matrices from different applications will be shown. We will also present how such solutions enable, without extra cost to the programmer, better performance on irregular computations such as in the block low-rank implementation of the solver.

Mathieu Faverge
Bordeaux INP - Labri - Inria
mathieu.faverge@inria.fr

Gregoire Pichon
INRIA
gregoire.pichon@inria.fr

Pierre Ramet
Bordeaux University - INRIA
pierre.ramet@inria.fr

MS197

Hierarchical Algorithms on Hierarchical Architectures

Some algorithms achieve optimal arithmetic complexity with low arithmetic intensity (ops/Byte), or possess high arithmetic intensity but lack optimal complexity, while others, such as Fast Multipole and its H-matrix algebraic generalizations, realize a combination of optimal complexity and high intensity. Implemented with task-based dynamic runtime systems, these hierarchical methods also have potential for relaxed synchrony, which is important for future energy-austere architectures, since there may be significant nonuniformity in processing rates of different cores even if task sizes can be controlled. We describe modules of KAUST's Hierarchical Computations on Manycore Architectures (HiCMA) software toolkit that illustrate these features and are intended as building blocks of more sophisticated applications, such as matrix-free higher-order methods in optimization. HiCMA's target is hierarchical algorithms on emerging architectures, which have hierarchies of their own that generally may not align with those of the algorithm. Some modules of this open source project have been adopted in the software libraries of major vendors. We describe what is currently available and some motivating applications.

Kadir Akbudak
Bilkent University
kadir.cs@gmail.com

Ali M. Charara
KAUST
ali.charara@kaust.edu.sa

David E. Keyes
King Abdullah University of Science and Technology (KAUST)
david.keyes@kaust.edu.sa

Hatem Ltaief, Aleksandr Mikhalev
King Abdullah University of Science & Technology (KAUST)
hatem.Ltaief@kaust.edu.sa, aleksandr.mikhalev@kaust.edu.sa

MS198

Exploiting Ridge Structure in Chance Constrained Design under Uncertainty

In engineering applications it is important to create a design that is robust in the face of uncertainty of operating conditions, material properties, and realizations of the design. As a design which is robust to all possible uncertainties may be too conservative, an alternative is to allow the failure criteria to be violated with a small probability; this converts deterministic constraints on the design into chance constraints. These chance constraints pose a significant challenge for optimization. Here we discover low-dimensional *ridge structure* in the constraint functions for a multiphysics model of a jet nozzle: each constraint is well-approximated by a low-order polynomial of one or two linear combinations of the input variables. By exploiting this ridge structure, we are able to convert these chance constraints into an easier robust constraints over a low-dimensional set. Moreover this optimization problem can be converted into a simple linear program if the objective and constraint functions are approximated by one-dimensional ridges. Although approximations have been in made in this simplification, most can be made conservative such that the resulting solution satisfies the original constraints. This approach provides a compelling methodology for chance constrained optimization problems where constraints are expensive to evaluate and derivatives are unavailable.

Jeffrey M. Hokanson
University of Colorado Boulder
jeffrey.hokanson@colorado.edu

Paul Constantine
University of Colorado Boulder
Department of Computer Science
pconstan@mines.edu

MS198

Low-rank Tucker Decomposition of Large Tensors using TensorSketch

Many real datasets have more than two dimensions and are therefore better represented using tensors, or multi-way arrays, rather than matrices. In the same way that methods such as the singular value decomposition can help in the analysis of data in matrix form, tensor decompositions are important tools when working with tensor data. As multidimensional datasets grow larger and larger, there is an increasing need for methods that can handle them, even on modest hardware. One approach to the challenge of handling big data, which has proven to be very fruitful in the past, is the use of randomization. We present two algorithms for computing the Tucker decomposition of a tensor which incorporate random sketching. Our algorithms, which only need a single pass of the data, are suitable when the decomposition we seek is of low rank. We test our algorithms on sparse synthetic data and compare them to multiple other methods. We also apply one of our algorithms to a real dense 38 GB tensor representing a video and use the resulting decomposition to correctly classify frames containing disturbances.

Osman Asif Malik
University of Colorado, Boulder
osman.malik@colorado.edu

Stephen Becker
University of Colorado at Boulder

stephen.becker@colorado.edu

MS198

Data Compression Based on Orthogonal Transforms,Impact on Post-processing of Large Data Sets

Beyond compression of large data sets for efficient storage and handling, lies the question of error incurred by truncation of the original information. Issues such as accuracy and stability dictate the necessities of computation at scale, however, data analysis and checkpointing pose lower requirements. We present a highly parallelizable compression strategy tailored to the underlying numerical discretization of the simulation, here spectral element method. Subsequently, we provide an extensive study of I/O speedup, impact on the statistical analysis of compressed fields, resilience, visualization as well as the building of reduced order models. This encompassing study leads to the conclusion that in the field of computational fluid dynamics the overhead of simulation data can be drastically reduced, up to 90% in data savings, for most commonly used data processing tools.

Oana Marin
Argonne National Laboratory
oanam@anl.gov

MS198

Scenario Selection for Reliable Operations and Planning of Large Power Grids

Intermittency, uncertainty, and variability of generation and load present challenges to the reliable and cost-effective operation and expansion of power grids. These challenges are exacerbated by increases in renewable energy generation, demand response, and cyber security events. Successful operations and expansions that address these challenges will require new algorithmic approaches designed to account for intermittency, uncertainty, and variability distributed across networks. Exascale computing capabilities will enable the direct consideration of high-fidelity operational models, such as security-constrained AC optimal power flow, in an algorithmic framework that includes variability and uncertainty. In this talk, we will focus on renewable energy, and wind energy in particular, within the context of two-stage stochastic optimization. In our problem formulation, uncertainty is represented using scenarios via the sample average approximation (SAA). Choosing input scenarios for the SAA that are adequate tests of grid robustness, and also possess realistic spatial, temporal, and physical relationships is a difficult task. We propose an importance sampling approach to drawing scenarios from high-fidelity simulation or historical data. Ideally, scenarios drawn in such a manner will provide superior representations of spatial, temporal, and physical relationships. Simulations on test systems show promising results.

Matthew J. Reynolds
National Renewable Energy Laboratory
matthew.reynolds@nrel.gov

MS199

Nested Solvers for Improved Scalability of Domain Decomposition Methods

Domain decomposition and multigrid methods are some of the most versatile solvers for large-scale computations. However, in the strong-scaling regime, they sometime exhibit poor scalability because of an increasing amount of communications and low arithmetic intensities. In this talk, I'll present different ways to improve the efficiency of different solvers based on domain decomposition or multigrid methods. These improvements rely on the use of appropriate nested solvers or subspace recycling. Some results obtained on large number of processes will be showcased in the context of fluid dynamics and radiative transfer. Open-source implementations using PETSc, SLEPc, and HPDDM will also be quickly explained.

Pierre Jolivet
CNRS
pierre.jolivet@enseeiht.fr

MS199

A New Configurable Preconditioner for Additive Subspace Correction Methods in PETSc

Small block overlapping, and non-overlapping, Schwarz methods are theoretically highly attractive as multilevel smoothers for a wide variety of problems that are not amenable to point relaxation methods. Examples include monolithic Vanka smoothers for Stokes, overlapping vertex-patch decompositions for $H(\mathrm{div})$ and $H(\mathrm{curl})$ problems, along with nearly incompressible elasticity and augmented Lagrangian schemes. While it is possible to manually program these different schemes, their use in general purpose libraries has been held back by a lack of generic, composable interfaces. We present a new approach to the specification and development such additive Schwarz methods in PETSc that cleanly separates the topological space decomposition from the discretisation and assembly of the equations. Our preconditioner is flexible enough to support overlapping and non-overlapping additive Schwarz methods, and can be used to formulate line, and plane smoothers, Vanka iterations, amongst others. We illustrate these new features with some examples utilising the Firedrake finite element library.

Lawrence Mitchell
Durham University
lawrence.mitchell@durham.ac.uk

Patrick E. Farrell
Mathematical Institute
University of Oxford
patrick.farrell@maths.ox.ac.uk

Rob C. Kirby
Baylor University
robert_kirby@baylor.edu

Matthew G. Knepley

University at Buffalo
Department of Computer Science and Engineering
knepley@gmail.com

MS199

A Scalable Solver for the Stationary Navier-Stokes Equations in 3D Based on Augmented Lagrangian and Subspace Correction Methods

In Benzi & Olshanskii (SIAM SISC 28 (2006)) a preconditioner of augmented Lagrangian type was presented for the two-dimensional stationary incompressible Navier–Stokes equations that exhibits convergence almost independent of Reynolds number. It relies on a highly specialized multigrid scheme involving a custom prolongation operator. In this talk we present two extensions of their work. Firstly, we adapt their scheme to three dimensions. This requires a particular choice of finite element for the velocity space and a modification of the prolongation operator. Secondly, we have implemented the solver in Firedrake, expressing both the prolongation and the smoothing operators using the newly developed interface for subspace correction methods in PETSc. We present numerical results running on up to 25 000 cores demonstrating convergence independent of the mesh and nearly independent of the Reynolds number.

Florian Wechsung
University of Oxford
wechsung@maths.ox.ac.uk

MS199

BDDC and FETI-DP Methods in PETSc

In this talk, we will present the implementations, within the PETSc library, of the Balancing Domain Decomposition by Constraints (BDDC) and of the Finite Element Tearing and Interconnecting Dual Primal (FETI-DP) methods, and accessible through a variety of open-source finite element libraries like FEniCS and MFEM. Connections with the methods of subspace corrections will be also provided. Large-scale numerical results will be shown in order to assess the quality and the generality of the implementation, providing results ranging from heterogeneous elliptic PDEs to indefinite linear systems.

Stefano Zampini
King Abdullah University of Science and Technology
stefano.zampini@kaust.edu.sa

MS200

Sparse Polynomial Approximation on Irregular Domains

The approximation of smooth, high-dimensional functions via sparse polynomial expansions has received significant attention in the last five to ten years. Such approaches enjoy several benefits. For instance, using the theory of compressed sensing, one can show that such techniques mitigate the curse of dimensionality to a substantial extent. However, the majority of works consider approximating functions defined over tensor-product domains. Yet many problems arising in applications involve functions defined over irregular domains. This leads to significant practical and theoretical challenges. In this talk, I will first overview these challenges, and then elaborate several new, compressed sensing-based strategies for attaining high-accuracy polynomial approximation on irregular domains.

Ben Adcock
Simon Fraser University
ben_adcock@sfu.ca

Daan Huybrechs
Department of Computer Science
University of Leuven
daan.huybrechs@cs.kuleuven.be

Mohsen Seifi
Simon Fraser University
smseifi@sfu.ca

MS200

Learning Functions with Low Complexity in High Dimensions

Many data sets in image analysis and signal processing are in a high-dimensional space but exhibit a low-dimensional structure. We are interested in building efficient representations of these data for the purpose of compression and inference. In the setting where a data set in R^D consists of samples from a probability measure concentrated on or near an unknown d-dimensional manifold with d much smaller than D, we consider two sets of problems: low-dimensional geometric approximations to the manifold and regression of a function on the manifold. In the first case, we construct multiscale low-dimensional empirical approximations to the manifold and give finite-sample performance guarantees. In the second case, we exploit these empirical geometric approximations of the manifold and construct multiscale approximations to the function. We prove finite-sample guarantees showing that we attain the same learning rates as if the function was defined on a Euclidean domain of dimension d. In both cases our approximations can adapt to the regularity of the manifold or the function even when this varies at different scales or locations.

Wenjing Liao
Georgia Institute of Technology
wliao60@gatech.edu

Mauro Maggioni, Stefano Vigogna
Johns Hopkins University
mauro@math.jhu.edu, svigogna@math.jhu.edu

MS200

Analysis of Sparse Recovery for Legendre Expansions using Envelope Bound

We provide novel sufficient conditions for the uniform recovery of sparse Legendre expansions using $\ell_2$minimization, where the sampling points are drawn according to orthogonalization (uniform) measure. So far, conditions of the form $m \leq K^2 s \times$log factors have been relied on to determine the minimum number of samples

$m$that guarantees successful reconstruction of s-sparse vectors when the measurement matrix is associated to an orthonormal system. However, in case of sparse Legendre expansions, the uniform bound K of Legendre systems is so high that these conditions are unable to provide meaningful guarantees. In this talk, we present an analysis which employs the envelop bound of all Legendre polynomials instead. We prove a new recovery guarantee for$s$-sparse Legendre expansions, $m \leq s^2 \times$log factors, which is independent of the polynomial subspace. Arguably, this is the first recovery condition established for orthonormal systems without assuming the uniform boundedness of the sampling matrix. Along the way, we derive simple criteria to detect good sample sets. Our numerical tests show that sets of uniformly sampled points that meet these criteria will perform better recovery on average.

Hoang A. Tran
Oak Ridge National Laboratory
Computer Science and Mathematics Division
tranha@ornl.gov

Clayton G. Webster
University of Tennessee and Oak Ridge National Laboratory
webstercg@ornl.gov

MS200

Sparse Recovery and Outlier Detection for Dependent Data

Learning non-linear systems from noisy, limited, and/or dependent data is an important task in data-based learning across various scientific fields including statistics, engineering, computer science, mathematics, etc. One of the key ideas is to learn an unknown generating function from a set of input-output pairs, which can be rephrased as finding an approximation to a high-dimensional function. Without additional restrictions or structures, this learning task is ill-posed; however, knowledge on structure or behavior of the unknown function can make the task well-posed. In this work, we study the problem of learning nonlinear functions from identically distributed but not independent data that is corrupted by outliers and noise. The learning problem is written as a parameter estimation problem where we incorporate both the unknown coefficients and the corruptions in a basis pursuit denoising framework. The main contribution of our paper is to provide a reconstruction guarantee for the associated ℓ_1-optimization problem where the sampling matrix is formed from dependent data. We prove that the sampling matrix satisfies the null space property, provided that the data is compact and satisfies a suitable concentration inequality.

Giang Tran
University of Waterloo
Department of Applied Mathematics
giang.tran@uwaterloo.ca

Rachel Ward
Department of Mathematics
University of Texas at Austin
rward@math.utexas.edu

Hayden Schaeffer
Carnegie Mellon University
schaeffer@cmu.edu

MS201

Theoretical Aspects of Nonlocal Helmholtz Decomposition of a Vector Field

Nonlocal theories have been introduced in the mechanics of solids where the propagation of cracks and other discontinuities hinder the use of classical differential operators. The study of integral operators has been central to the analysis of solutions to nonlocal systems. By replacing classical differential operators with integral operators, nonlocal frameworks allow the consideration of solutions with little to no regularity (L^2-level). In this talk, we will discuss preliminary well-posedness results for nonlocal material science models where the interaction kernel of the integral operator is weakly singular and motivate the development of theory for Helmholtz-type decompositions of vector fields in the nonlocal setting. The work presented is part of the recent Women in Mathematics of Materials (WIMM) Workshop at University of Michigan.

Marta D'Elia
Sandia National Laboratories
mdelia@sandia.gov

Cynthia Flores
California State University
Channel Islands
cynthia.flores@csuci.edu

Xingjie Li
University of North Carolina, Charlotte
xli47@uncc.edu

Petronela Radu
University of Nebraska-Lincoln
Department of Mathematics
pradu@math.unl.edu

Yue Yu
Department of Mathematics, Lehigh University
yuy214@lehigh.edu

MS201

Investigation of Dispersion Relation of the 1D Quasi-nonlocal Coupling for Nonlocal and Local Mechanics

In this talk, we explore the dispersion relation of 1D quasi-nonlocal coupling for the peridynamics and local mechanics models. We propose a finite difference numerical method for the spatial and temporal discretizations, and study the dispersion relation on both continuous and discrete settings. The imaginary part of the dispersion relation on the transition region is proved to be of order of horizon size, whereas other existing coupling methods are at most of order of one over horizon size. Hence, the quasi-nonlocal coupling is dynamically consistent and convergent to the local mechanics with first order speed. Several numerical

tests are performed to confirm the theoretical findings.

Helen Li
University of North Carolina, Charlotte
mathslxj@gmail.com

MS201

Theoretical and Numerical Aspects for Doubly Nonlocal Problems on Bounded Domains

Abstract not available.

Petronela Radu
University of Nebraska-Lincoln
Department of Mathematics
pradu@math.unl.edu

MS201

Numerical Aspects of Nonlocal Helmholtz Decomposition of a Vector Field

Important applications in diffusion, elasticity, fracture propagation have benefited from the introduction of nonlocal models. However, this advantage is counterbalanced by the fact that the theory of nonlocal calculus is still being developed. For instance, the Helmholtz decomposition describes a vector field in terms of its divergence-free and rotation-free components. Such a description is valuable in the analysis of vector fields since some important properties like incompressibility and vorticity can be studied on the components directly, which makes the Helmholtz decomposition one of the fundamental theorems in fluid dynamics. However, in nonlocal calculus both the theory and the computation of Helmholtz decomposition remind empty. In this work we provide a Helmholtz decomposition in the nonlocal setting, and develop a numerical solver in the context of the discretization method based on a meshfree quadrature rule [N.Trask, H.You, Y.Yu, M.L. Parks, A meshfree quadrature rule for non-local mechanics, CMAME, 2018.]. To demonstrate the applicability of the developed Helmholtz decomposition, we solve for the divergence-free and rotation-free components with proper Dirichlet type boundary conditions, and numerically verified the convergence of the developed approach. To show the connection with the local theory, the asymptotically property of the nonlocal solution to the corresponding components from the local Helmholtz decomposition will also be investigated.

Yue Yu
Department of Mathematics, Lehigh University
yuy214@lehigh.edu

Marta D'Elia
Sandia National Laboratories
mdelia@sandia.gov

Cynthia Flores
California State University
Channel Islands
cynthia.flores@csuci.edu

Xingjie Li
University of North Carolina, Charlotte
xli47@uncc.edu

Petronela Radu
University of Nebraska-Lincoln
Department of Mathematics
pradu@math.unl.edu

MS202

Finite Element and Finite Difference Modeling and Convergence for State Based Peridynamic Fracture

We analyze finite element and finite difference approximations for a class of nonlocal fracture models. The nonlocal model is initially elastic but beyond a critical strain the material softens with increasing strain. This model is formulated as a state-based perydynamic model using two potentials: one associated with hydrostatic strain and the other associated with tensile strain. We show that the dynamic evolution is well-posed for deformations that belong to either Hölder and Sobolev spaces. However as the scale of nonlocality ϵ tends to zero the evolutions converge to a sharp fracture evolution in SBD. We fix the length scale of nonlocality and the discrete approximations are shown to converge to the Hölder or Sobolev evolution at a rate controlled by the Hölder or Sobelev space that the evolution belongs to. The semi-discrete approximations are found to be stable with time. Numerical examples are given to illustrate the numerical methods. This is joint work with Prashant Jha.

Robert P. Lipton
Department of Mathematics
Louisiana State University
lipton@math.lsu.edu

MS202

Subsurface Applications for Peridynamics

Peridynamics is a nonlocal reformulation of continuum mechanics that is suitable for representing fracture and failure, see [S. A. Silling, Reformulation of elasticity theory for discontinuities and long-range forces, J. Mech. Phys. Solids 48, (2000), 175-209., S. A. Silling, Reformulation of elasticity theory for discontinuities and long-range forces, J. Mech. Phys. Solids 48, (2000), 175-209.] and the references therein. Better understanding and control of the subsurface is important to the energy industry for improving productivity from reservoirs. We motivate and explore two relevant subsurface applications for peridynamics. The first involves solving inverse problems in heterogeneous and fractured media, which may be useful in characterizing subsurface stress-state conditions [D. Turner, B. van Bloemen Waanders and M.L. Parks, Inverse problems in heterogeneous and fractured media using peridynamics, Journal of Mechanics of Materials and Structures, 10(5), pp. 573-590, 2015.]. The second involves the study of fracture initiation and growth from propellant-based stimulation of a wellbore [R. Panchadhara, P.A. Gordon, and M.L. Parks, Modeling propellant-based stimulation of a borehole with peridynamics, International Journal of Rock Mechanics and Mining Sciences, 93, pp. 330-343, 2017.]. Simple models

and proof-of-concept numerical studies are presented.

Michael L. Parks
Sandia National Laboratories
mlparks@sandia.gov

MS202

Foundations of Material Anisotropy in Classical Elasticity and Peridynamics

Peridynamics is a nonlocal reformulation of classical continuum mechanics, suitable for material failure and damage simulation. Originally, this nonlocal theory was presented as a bond-based theory, for which the material response of an isotropic medium is limited by a fixed Poisson's ratio. To overcome this limitation, the state-based theory was developed. Applications in peridynamics to date cover a wide range of engineering problems; however, the majority of those applications employ isotropic material models. Only recently, a limited number of anisotropic peridynamic models were developed. In this talk, we will discuss various concepts related to material anisotropy in classical elasticity, and survey the different classes of anisotropic material models in classical linear elasticity. We will then show how these material anisotropy concepts extend to peridynamics and present anisotropic peridynamic models. We will show a classification and a hierarchy of those models, and we will discuss their relation to classical elasticity models, as well as restrictions arising from a bond-based interaction assumption.

Pablo Seleson, Jeremy Trageser
Oak Ridge National Laboratory
selesonpd@ornl.gov, trageserje@ornl.gov

Max Gunzburger
Florida State University
Department of Scientific Computing
mgunzburger@fsu.edu

MS202

Asymptotically Compatible Foundations for Nonlocal Mechanics

Nonlocal models for solid mechanics offer a natural functional setting to handle problems involving discontinuities, as occur in fracture mechanics and material science problems involving interfaces. While such models have attained popularity, relatively little work has been done establish discretizations of these models with rigorous mathematical underpinnings. We present an optimization-based quadrature technique that builds upon ideas from meshfree approximation and may be used to develop strong-form discretizations of nonlocal models that preserve the so-called asymptotically compatible limit. This technique remedies the issues with boundary conditions and internal interfaces which have traditionally plagued nonlocal models for mechanics. We present applications of this work to material science problems at the mesoscale involving anomalous transport and fracture.

Nathaniel Trask
Brown University
Division of Applied Mathematics
nat.trask@gmail.com or natrask@sandia.gov

MS203

Predicting Genomic Variation in Error-prone Data Regimes

Each time a cell divides and its DNA is duplicated, there is an opportunity for the introduction of errors such as insertions and deletions. Both the introduction and accumulation of genomic variation may lead to speciation events or to development of genetic diseases (e.g., cancer). As genome sequencing costs decrease, the volume of sequencing data has resulted in the need for advanced mathematical and computational methods. We explore mathematical models in error-prone genomic data regimes to improve predictions of genomic variation in organisms. We predict genomic variation between members of the same species in two ways. First, we develop a constrained- optimization framework using gradient-based methods and apply inheritance assumptions to constrain our solution with a sparsity-promoting ℓ_1 penalty to detect structural variants. Second, we propose machine and deep learning models in this context to allow for a generalization of relatedness among individuals.

Mario Banuelos
California State University, Fresno
mbanuelos22@csufresno.edu

MS203

Assessing Bias in the Prediction of Mortality in the ICU

Efforts to model risk of mortality in hospitalized patients dates back to over half a century. These models have evolved from methods relying in subjective variable selection (i.e., SAPS and APACHE scores), to machine learning (ML) and recurrent neural network (RNN) algorithms that are now possible due to the explosion in digital clinical data available for analysis. In fact, it has been shown that models based on ML outperform baselines based on logistic regression or face validity. In parallel, there is an increasing awareness on the presence of bias and the potential disproportionate impact that ML can have on groups that are already socially disadvantaged. If there were bias in the mortality prediction in the ICUs, we may be amplifying societal inequalities that already exist, by for example, overseeing early signals associated to underrepresented groups. We will use data derived from the publicly available Medical Information Mart for Intensive Care (MIMIC-III) to compare the diagnostic ability and specificity of different mortality ICU predictors among different ethnical and gender groups. Results from our models will inform health care policies and may allow to identify areas of research or improvement in patient care.

Lara Reichmann
Data Institute
University of San Francisco

lreichmann@usfca.edu

MS203

Classification of Short ECG Readings: TDA Informed Machine Learning

Atrial Fibrillation is a common heart condition affecting nearly thirty-three million people worldwide. Atrial Fibrillation can lead to erratic heart behavior that can cause blood clots, heart failure, and even death. It is difficult to diagnose due to distortion in the cardiac rhythm caused by spasms. In this talk, we develop a novel approach to the PhysioNet/Computing in Cardiology 2017 challenge which was focused on classifying electrocardiogram (ECG) readings and accurately detect Atrial Fibrillation. While the leading models used deep learning, random forests and medically in- formed feature engineering, we use persistent-homology based features fed to a simple random forest model to classify ECGs and show that this approach achieves similarly accurate scores for detecting Atrial Fibrillation.

David T. Uminsky
University of San Francisco
Department of Mathematics
duminsky@usfca.edu

MS203

A Consistent Density-based Clustering Algorithm and its Application to Microstructure Image Segmentation

Data clustering is a fundamental task for discovering patterns in data, and is central to machine learning. Often, a data set is assumed to live in a manifold and be sampled according to a probability measure. Then the clusters can be defined as peaks in the sampled probability density, and a clustering algorithm would need to identify the peaks in the density to compute the clusters. Some of the challenges in this approach include the non-uniform sampling of the density and the bridges between peaks of the density. To solve these problems, we propose a new clustering algorithm that divides the clustering problem into three steps: picking a good threshold on the sample density to separate the peaks, clustering the superlevel set at the chosen threshold, and classifying the remaining points. We explain the key details of these steps, and provide theoretical assurances on the performance. As an important application, we show how to apply this method to segment microstructure images by considering the images as a point-cloud of image patches. We present results on 2D microscopic images of various materials.

Marilyn Vazquez
ICERM
vazquez.marilyn.y@gmail.com

MS204

Consistent Parameter Distributions in High-dimensional Spaces: Built One Dimension at a Time

We consider a density-free "consistent" Bayesian approach that can be applied to problems where observed distributions for scalar output quantities are potentially obtained and inverted asynchronously in time. The basic idea is that we sequentially update weights on prior samples in a parameter space so that the push-forwards of the empirical distribution functions match the distributions of data. This approach exploits implicit low-dimensional structure that each output quantity of the model imparts onto the input parameter space. This is, in a sense, an analog of data assimilation that is more like "functional assimilation for distributions" where instead of trying to update statistics (like the mean and variance) at each time, we actually update an entire distribution.

Troy Butler
University of Colorado Denver
troy.butler@ucdenver.edu

MS204

Hierarchical Off-diagonal Low-rank Approximation for Hessians in Bayesian Inference

We address the problem of quantifying the uncertainty in the solution of inverse problems governed by PDE-based models within the framework of Bayesian inference. Computing the general solution of the inverse problem—i.e., the posterior probability density—is challenging due to potentially expensive to solve forward models and the high dimensionality of the uncertain parameters (which are discretizations of parameter fields in the model). In this talk, we focus on problems that exhibit local sensitivity of data to parameters and build a hierarchically off-diagonal low-rank approximation for the Hessian (of the cost function or log posterior); exploiting the structure of these problems has the potential to alleviate the computational cost for solving large-scale statistical inverse problems. We apply this Hessian approximation to Newton-CG type methods (as a preconditioner), and to build a Gaussian approximation of the posterior for the inference of the basal sliding coefficient field in an ice sheet model.

Tucker Hartland
University of California, Merced
thartland@ucmerced.edu

Ruanui Nicholson
Engineering Sceience
University of Auckland, New Zealand
rnic052@aucklanduni.ac.nz

Noemi Petra
University of California, Merced
npetra@ucmerced.edu

MS204

Mitigating the Cost of PDE-constrained Bayesian Inverse Problems Using Dimensionality Reduction and Machine Learning

Given a hierarchy of reduced-order models to solve the inverse problems for quantities of interest, each model with varying levels of fidelity and computational cost, a machine learning framework is proposed to improve the mod-

els by learning the errors between each successive levels. Hierarchical reduced-order models are derived from higher-fidelity models using approaches such as simplifying physics assumptions, using coarser grids, alternative basis expansions, and looser residual tolerances. Each reduced-order model is a statistical model generating rapid and reasonably accurate solutions to new parameters, and are typically formed using expensive forward solves to find the reduced subspace. These approximate reduced-order models speed up computational time but they introduce additional uncertainty to the solution. By statistically modeling errors of reduced order models and using training data involving forward solves of the reduced order models and the higher fidelity model, we train a deep neural network to learn the error between successive levels of the hierarchy of reduced order models thereby improving their error bounds. The training of the deep neural network occurs during the offline phase and the error bounds can be improved online as new training data is observed.

Sheroze Sheriffdeen
UT Austin
sheroze@ices.utexas.edu

MS204

Low-rank Structure in Optimization-based Sampling

Optimization-based samplers are sampling algorithms that employ optimization to generate a proposal sample. These methods use the proposal sample as an independent proposal in Metropolis–Hastings or as a biasing distribution in importance sampling. In the context of MCMC, they provide a global non-Gaussian proposal mechanism. Each proposal sample is more expensive to obtain than those from MCMC algorithms with Gaussian proposals, but their cost is offset by the quality of the proposal. Randomize-then-optimize (RTO) is an optimization-based sampler that is dimension independent. We present a scalable implementation of RTO that stores an SVD of the Jacobian of the forward model. When the rank of the SVD is truncated, more directions of the proposal becomes distributed according to the prior. For linear forward models, RTO's truncated proposal distributions are exactly the same as Gaussian posterior approximations that use the likelihood informed subspace (LIS). Hence, RTO's proposal can be thought of as a non-Gaussian generalization of LIS-type proposals. We also formulate RTO's proposal mechanism in function space to explain its dimension independent sampling behavior. This formulation leads to a natural way to correlate samples between different discretizations of the parameter field which can be used in a multi-level/multi-fidelity setting.

Zheng Wang, Youssef M. Marzouk
Massachusetts Institute of Technology
zheng_w@mit.edu, ymarz@mit.edu

Tiangang Cui
Monash University
tiangang.cui@monash.edu

MS205

A Parallel Time Integration Approach without Re-discretization in Time Direction

Many parallel time integration methods apply re-discretization with enlarged time step width to propagate solution information to later time steps with low calculation cost. This re-discretization tends to create unstable time integration problem and causes the performance degradation depending on the time step width. Therefore, we investigated the approach that uses Jacobi matrices of the original problem to propagate solution to future times steps fast without re-discretization. This talk introduces our approach and performance results for simple non-linear time step simulation.

Akihiro Fujii, Teruo Tanaka
Kogakuin University
fujii@cc.kogakuin.ac.jp, teru@cc.kogakuin.ac.jp

Takeshi Iwashita
Hokkaido University
iwashita@iic.hokudai.ac.jp

MS205

Recent Trends and Challenges for High Performance Sparse Linear Algebra

The path to scalable performance on next generation platforms must include design and implementation for accelerated computations, including addressing all phases of computation and data movement. In this presentation, we will provide an overview of the current challenges and approaches being used by several teams of sparse linear algebra algorithm and software developers. We will discuss design requirements, implementation strategies and use of new approaches for reducing data bandwidth requirements. We end the presentation with a few key challenges for the future, as heterogeneity becomes even more prevalent.

Michael A. Heroux
Sandia National Laboratories
maherou@sandia.gov

MS205

H-Matrix Framework for Many-core Processors

Hierarchical matrices ($\mathcal{H}$-matrices) are an approximation technique for dense matrices, such as the coefficient matrix of the boundary element method (BEM). An $\mathcal{H}$-matrix is expressed by a set of low-rank approximated and small dense sub-matrices, each of which has various ranks. The use of $\mathcal{H}$-matrices reduces the required memory footprint of dense matrices from $O(N^2)$ to $O(N\log N)$ and is suitable for manycore processors that have relatively small memory capacities compared to traditional CPUs. However, existing parallel adaptive cross approximation (ACA) algorithms, which are low-rank approximation algorithms used to construct $\mathcal{H}$-matrices, are not designed to exploit many-core processors in terms of load balancing. In existing parallel algorithms, the ACA process is independently applied

to each sub-matrix. The computational load of the ACA process for each sub-matrix depends on the sub-matrix s rank; however, the rank is defined after the ACA process is applied. This makes it difficult to balance the load. We propose load-balancing-aware parallel ACA algorithms for $\mathcal{H}$-matrices that focus on many-core processors. We implemented the proposed algorithms into $\mathcal{H}$ACApK, which is an open-source $\mathcal{H}$-matrix library originally developed for CPU-based clusters. The proposed algorithms were evaluated using BEM problems on an NVIDIA Tesla P100 GPU (P100) and an Intel Xeon Broadwell processor.

Tetsuya Hoshino, Akihiro Ida
The University of Tokyo
hoshino@cc.u-tokyo.ac.jp, ida@cc.u-tokyo.ac.jp

MS206

Smoothing Data Movement Between Ram and Storage for Reverse Time Migration

Simulation systems often produce a massive amount of data that cannot fit on the aggregate fast memory of the compute nodes, and they also require to read back these data for computation. As a result, I/O data movement can be a bottleneck in large-scale simulations. Advances in memory architecture have made it feasible and affordable to include multiple heterogeneous storage media in a single compute node. A typical workstation nowadays contains several HHDs and SSDs while an advanced one might contain a high-throughput NVMe. However, while adding additional and faster storage media increases I/O bandwidth, it pressures the CPU, as it becomes responsible for managing and moving data between these layers of storage. Simulation systems are thus vulnerable to being blocked by I/O operations. The system, proposed in this research, demonstrate a general and versatile method for overlapping I/O with computation that helps to ameliorate the strain on the processors through asynchronous access. We illustrate our system on a prefetching strategy for a specific Reverse Time Migration (RTM) application, minimizing the impact on the computational kernel. On representative workloads and problem sizes from seismic imaging, our system shows that coupling both storage and memory systems and using the knowledge of data access patterns can decrease the time spent blocking for I/O by more than 70 %, thereby improving the entire execution time by nearly three-fold.

Tariq Alturkestani
King Abdullah University of Science and Technology
tariq.alturkestani@kaust.edu.sa

Thierry Tonellot
ARAMCO
thierrylaurent.tonellot@aramco.com

Hatem Ltaief
King Abdullah University of Science & Technology (KAUST)
hatem.Ltaief@kaust.edu.sa

Vincent Etienne
ARAMCO
vincent.etienne@aramco.com

David E. Keyes
King Abdullah University of Science and Technology (KAUST)
david.keyes@kaust.edu.sa

MS206

Automation of Velocity Analysis with Recurrent Neural Networks

Can seismic data processing and imaging be automated with machine learning? Due to the shear complexity of this problem, the answer remains elusive. Still, part of the solution can be gleaned from smaller, simpler sub-problems. The machine learning community has fully embraced this strategy—the MNIST dataset being one of the most successful examples. In this talk, we argue that velocity analysis has the potential to become such a canonical example for the seismic community. More precisely, we define the problem as the estimation of NMO and interval velocity from CMP gathers in the approximation of a flat layered Earth. This problem is small enough so that training can be performed with modest resources and is still directly applicable to a stacking approach of seismic imaging. We investigate estimating the NMO velocity with a deep, recursive neural net. The architecture of the network mimics traditional semblance analysis: steps such as semblance computation and picking are replaced with different layers of convolutional, recursive and recurrent neural nets. We show that the precision of the predicted velocity profiles using the network is favorable compared to traditional semblance analysis. We additionally show that the neural net can identify primary reflections. In summary, the NMO velocity estimation problem possess all the necessary attributes to become a testbed for new machine learning approaches to seismic processing.

Gabriel Fabien-Ouellet
Polytechnique Montreal
gabriel.fabien-ouellet@polymtl.ca

Rahul Sarkar
Stanford University
rahul@sep.stanford.edu

MS206

Fossil Segmentation and Classification using 3D U-NET Based Architecture

We present a novel Neural Network Architecture able to segment and classify microfossils in 3D rock micro CT-scans. Microfossils are very important in analyzing and age-dating rocks as well as in paleoenvironmental reconstruction. Hence, oil and gas, mining, engineering industries as well as environmental and general geology are constantly looking for the most effective way to analyse microfossils. In the last few years Artificial Intelligence and especially Deep Learning techniques have been applied widely to image recognition processes to analyse microfossils. These methods have been growing in application thanks to the exponential growth of computational power and large amount of data available. However the analysis of Volumetric Data with Deep Learning methods has only been recorded in the Medical Imaging field and re-

lied mostly on Vnet and 3D-Unet Neural Network Architectures. Through our work, we extend the use of Volumetric Deep Learning to the detection, segmentation and classification of microfossils in 3D micro CT-scans. Our approach is based on a tuned 3D-Unet Neural Network. We show that our 3D technique outperforms advanced 2D segmentation approaches using Mask R-CNN and the Unet Architecture followed by 3D interpolation. In this talk, we showcase the systems development phases, from synthetic volumes with simple geometric shapes to real 3D rock micro CT-scans, and their corresponding performances and accuracies.

Amine Kerkeni, Youssef Ben Dhieb, Karim Beguir
InstaDeep
ak@instadeep.com, y.bendhieb@instadeep.com, kb@instadeep.com

Noomane Keskes, Maxime Bellay
Total
noomane.keskes@total.com, maxime.bellay@total.com

MS206

Deep Learning in an Industrial Context: an Introduction to the NVIDIA Ecosystem

Deep Learning (DL) is a discipline of Machine Learning (ML) that is making an impact across diverse industries, including medical imaging, financial services, security, transportation, and others. Applications of Deep Learning have produced dramatic results, enabling new opportunities by introducing new compute models and data-driven approaches to problem solving. Applied to several industries, Oil and Gas in particular in this talk, these data-driven approaches are complementary tools for physical-based modeling, simulation, and inversion, but also are changing the downstream business by automating repetitive tasks. With the help of a powerful compute architecture and a complete end-to-end software stack, NVIDIA is providing an efficient platform that helps developers and Data Scientist get up to speed quicker and improve their productivity. In this talk we present the implication of NVIDIA in transforming the Oil and Gas industry through some case studies while emphasizing the hardware and software stack that democratizing Deep Learning amongst major industries.

Issam Said
PMC Univ Paris 06 and CNRS UMR 7606, LIP6
isaid@nvidia.com

MS207

Modelling Moving Contact Lines in Multi-phase Fluids

The moving contact line problem is a classical problem in fluid mechanics. The difficulty stems from the fact that the classical Navier-Stokes equation with no-slip boundary condition predicts a non-physical singularity at the contact line with infinite rate of energy dissipation. Many modified continuum models have been proposed to overcome this difficulty. Though they all succeed in removing the singularity, it is not clear whether they describe the real physics near the contact line region. In this talk, we will discuss how the continuum theory, molecular dynamics and the more recently developed multiscale techniques can be combined to give us a better understanding of the fundamental physics of the moving contact line and formulate simple and effective models. We also illustrate how this model can be used to analyze the behavior of the apparent contact angle, hysteresis and other important physical problems for the moving contact line.

Weiqing Ren
National University of Singapore and IHPC, A-Star
matrw@nus.edu.sg

MS207

Domain Decomposition Solvers and Preconditioners for the Implicit Closest Point Method

The numerical treatment of surface intrinsic elliptic PDEs presents several interesting challenges not present in flat space. The Implicit Closest Point Method, (ICPM) is an embedding method well-suited to solving such problems on fairly general surfaces. At the core of the ICPM lies an extension operator bringing surface bound information into the embedding space. With this extension operator, the surface PDE can be solved by standard methods in a neighbourhood of the surface. In this talk, the positive Helmholtz equation, $(c - \Delta_S)u = f$, is considered. The ICPM discretization leads to a non-symmetric system with some complex eigenvalues and weak non-locality, reducing the utility of iterative solvers and motivating the development of strong preconditioners. Optimized restricted additive Schwarz (ORAS) methods are well suited to this task and are formulated herein. The obtained solvers and pre-conditioners are compared with standard RAS and unpreconditioned GMRES, showing a reduction the required iterations in all tested problems.

Ian May
Mathematics
Simon Fraser University
mayianm@sfu.ca

Ronald Haynes
Memorial University of Newfoundland
St. John's, NL, Canada
rhaynes74@gmail.com

Steven Ruuth
Mathematics
Simon Fraser University
sruuth@sfu.ca

MS207

Stabilized IEQ/SAV Approach for Gradient Flow System with Strong Spatial Anisotropy

We consider numerical approximations for gradient flow models with strong anisotropy by taking the anisotropic Cahn-Hilliard/Allen-Cahn equations with their applications to the faceted pyramids on nanoscale crystal surfaces and the dendritic crystal growth problems, as special examples. The main challenge of constructing nu-

merical schemes with unconditional energy stabilities for these type of models is how to design proper temporal discretizations for the nonlinear terms with the strong anisotropy. We combine the recently developed IEQ/SAV approach with the linear stabilization approach, where some linear stabilization terms are added. These terms are shown to be crucial to remove the oscillations caused by the anisotropic coefficients, numerically. The novelty of the proposed schemes is that all nonlinear terms can be treated semi-explicitly, and one only needs to solve some coupled/decoupled, but linear equations at each time step. We further prove the unconditional energy stabilities rigorously and present various 2D and 3D numerical simulations to demonstrate the stability and accuracy.

Xiaofeng Yang
University of South Carolina
xfyang@math.sc.edu

MS207

An Interface-preserving Level-set Method for Interfacial Flows with Contact Lines

It is well-known that the level-set methods have poor mass conservation properties, mostly due to the shift of the zero level set during the reinitialization process. In this talk, we present an interface-preserving discontinuous Galerkin method to solve the Hamilton-Jacobi equation for the level-set reinitialization. This reinitialization method essentially causes no mass loss as long as the interface curvature can be resolved by the computational mesh. More importantly, it allows for an easy implementation of the artificial boundary conditions that arise when the interfaces intersect the domain boundaries at non-90 degree angles. We solve the level-set and flow equations by finite element methods, developed based on the finite-element library deal.II. A generalized Navier condition is adopted to relax the moving contact line singularity. We first compute a rising bubble for code validation as well as an illustration of the mass conservation performance. We then compute the advancing menisci in a Poiseuille flow, which are validated against Coxs hydrodynamic theory. We will also present some preliminary results on a non-iterative model for contact-angle hysteresis.

Jiaqi Zhang, Pengtao Yue
Virginia Polytechnic Institute and State University
zjiaqi@vt.edu, ptyue@vt.edu

MS208

High Order Direct ALE Schemes on Moving Voronoi Meshes with Changing Connectivity

In this talk we present a novel high order accurate both Finite Volume (FV) and Discontinous Galerkin (DG) scheme on ALE Voronoi meshes *regenerated* at any time step from a set of moving generator points. The Voronoi tessellations are obtained through an AREPO-type code (see Springel, MNRAS, 2010), which rapidly rebuilds a new mesh exploiting the previous one and efficiently redistributes the moving elements among the parallel CPU cores. Then, the old and new elements associated to the same generator points are connected in space–time to construct the so-called space–time control volumes, whose bottom and top faces may also be polygons of different types; moreover, even degenerating *sliver* elements are incorporated in order to fill the space–time holes raised by the changing connectivity (as suggested in Gaburro et al., CAF, 2017). The final ALE FV-DG scheme is obtained by a novel redesign of the high order accurate (both in space and time) direct ALE $P_N P_M$ scheme of Dumbser (JCP, 2008) and Boscheri et al. (JCP, 2014) which has been adapted to Voronoi and sliver space–time elements. A large set of different numerical tests shows the robustness and efficiency of the code. Moreover, due to the space-time conservation formulation of the direct scheme, the geometric conservation law (GCL) is automatically satisfied by construction and, thanks to the careful treatment of the holes, the method is locally and globally conservative.

Elena Gaburro
University of Trento
elena.gaburro@unitn.it

Walter Boscheri
University of Bolzano
walter.boscheri@unibz.it

Michael Dumbser
University of Trento
Italy
michael.dumbser@ing.unitn.it

Christian F. Klingenberg
Wurzburg University, Dept. of Mathematics
Germany
klingen@mathematik.uni-wuerzburg.de

MS208

Gradient-based Nonlinear Stabilization for Finite Element Discretizations of Conservation Laws

We propose a unified approach to gradient-based stabilization and bound-preserving limiting for continuous Galerkin discretizations of hyperbolic conservation laws. Building on recent advances in the analysis and design of edge-based algebraic flux correction schemes for singularly perturbed convection-diffusion problems, we derive artificial diffusion operators that generate nonlinear high-order stabilization in smooth regions and enforce discrete maximum principles everywhere. The correction factors for antidiffusive element or edge contributions are defined in terms of nodal gradients that vanish at local extrema. The proposed limiting strategy is linearity-preserving and provides Lipschitz continuity of constrained terms. Moreover, the proposed methodology leads to a well-defined limiter for antidiffusive corrections depending on time derivatives. The use of nonlinear high-order time integrators makes it possible to guarantee positivity preservation without any restrictions on the time step size.

Dmitri Kuzmin
Dortmund University of Technology

kuzmin@math.uni-dortmund.de

MS208

On a High-order Lagrangian Discontinuous Galerkin Hydrodynamic Method

We present a new high-order discontinuous Galerkin (DG) Lagrangian hydrodynamic method for curvilinear cells, which have edges that can bend. The specific volume, velocity, and specific total energy fields within a cell are represented with a high-order Taylor-series polynomial. The discontinuity in the polynomials at the cell boundary is addressed by solving a multi-directional Riemann problem at the surface nodes of the cell and at additional locations along the edges so that the surface integration is exact for the polynomial order. To ensure robust mesh motion, the reconstructed density at the cell boundary includes a correction term that is a function of the difference between the modal density field and a density field calculated using a subcell mesh. A new limiting approach is presented that permits high-order solutions in smooth regions and enforces monotonicity near shocks. The accuracy and robustness of the new high-order DG Lagrangian hydrodynamic method is demonstrated by simulating a diverse suite of test problems covering smooth flows and strong shocks.

Nathaniel Morgan
LANL
nmorgan@lanl.gov

Xiaodong Liu
Los Alamos National Laboratory
Computational Physics Division
xliu@lanl.gov

Donald E. Burton
LANL
burton@lanl.gov

MS209

Learning Strategies for Generating Low-rank Representations of Physical Systems

In this talk we consider the problem of learning dynamical systems using streaming data obtained during online operations. We discuss several approaches that include stochastic gradient descent, gradient descent in functional space, and Bayesian filtering. We focus on learning models that enforce certain physical constraints, and we partially achieve these goals through the use of symmetry enforcing techniques based on Lagrangians and Hamiltonians. To develop fast and scalable methods, we perform learning over the space of low-rank models. Finally, we explore the effects of partial observability on the reconstruction error.

Alex Gorodetsky
University of Michigan
goroda@umich.edu

MS209

Physics-aware and Sparse Coarse-graining of Multiscale Dynamics

High-dimensionality as well as disparity of scales are two common challenges in the area of computational physics. To overcome these issues, we describe a Bayesian perspective to coarse-graining of multiscale dynamics. This is similar to finding a low-dimensional, coarse-grained system that is physically inspired and that is maximal predictive for the full system and its dynamics. It is obvious, that the quality of such an approximation depends on how well the full system and its dynamics can be represented by the low-dimensional surrogate. Data generated from the original, full-order system are used to calibrate the low-dimensional model using Stochastic Variational Inference. In order to reduce the computational complexity of the low dimensional surrogate, special attention is given to a sparse representation. Moreover, the coarse-grained model is formulated in a way that it fulfills existing physical constraints.The usage of such a surrogate system introduces an additional uncertainty due to the unavoidable information loss that takes place. We seek to quantify this uncertainty and produce probabilistic predictions.

Sebastian Kaltenbach
Technical University of Munich
sebastian.kaltenbach@tum.de

Phaedon-Stelios Koutsourelakis
Technische Universität München
Germany
p.s.koutsourelakis@tum.de

MS209

Learning Coarse-grained Models from Molecular Dynamics

We present a novel learning framework which consistently incorporates the underlying physics while bypassing a major drawback of most modern, data-driven approaches i.e. the availability of Big Data. The generation of a sufficiently large training dataset poses a computationally demanding problem and the complete coverage of the atomistic configurational space is not guaranteed. As a result, extrapolative capabilities of such models are limited. We propose a novel objective based on a reverse variational formulation which fully incorporates the physical constraints or governing equation. Optimizing the proposed objective relies on the evaluation of the force field but does not require extensive MD simulation data. The resulting, generative coarse-grained model serves as an efficient and physically-motivated surrogate model for predicting atomistic configurations and estimating relevant properties. The formulation is probabilistic and accounts for epistemic uncertainty. The models uncertainty is consistently propagated resulting into probabilistic predictions of observables.

Markus Schoeberl
Technical University of Munich
m.schoeberl@tum.de

Nicholas Zabaras
University of Notre Dame
nzabaras@gmail.com

Phaedon-Stelios Koutsourelakis
Technische Universität München
Germany
p.s.koutsourelakis@tum.de

MS209

LES Wall Modeling via Physics-informed Neural Networks

While data-based approaches were found to be useful for sub-grid scale (SGS) modeling in Reynolds-averaged Navier-Stokes (RANS) simulations, there have not been many attempts of using machine learning (ML) techniques for wall modeling in large-eddy simulations (LES). LES differs from RANS in many aspects. For one thing, LES is scale resolving. For another, LES is in and of itself a high-fidelity tool. Because datasets of higher fidelity are in general not frequently accessible or available, this poses additional challenges to data-based modeling in LES. In this work, we discuss how the above-noted challenges may be addressed when taking a data-based approach to wall modeling. We will also show the necessity of incorporating physics insights in model inputs, i.e. using inputs that are inspired by the vertically integrated thin boundary layer equations and the eddy population density scalings. We will show that the inclusion of above physics-based considerations would enhance extrapolation capabilities of a neural network to flow conditions that are not within the train data. Being cheap-to-evaluate and using only channel flow data at $\mathrm{Re}_\tau = 1,000$, the trained networks are found to capture the law of the wall at arbitrary Reynolds numbers and outperform the conventional equilibrium model in a non-equilibrium flow.

Xiang Yang, Suhaib Zafar
Penn State University
xzy48@engr.psu.edu , sxz311@psu.edu

Jian-Xun Wang
University of Notre Dame
jwang33@nd.edu

Heng Xiao
Dept. of Aerospace and Ocean Engineering, Virginia Tech
hengxiao@vt.edu

MS210

Training Domain Expert Neural Networks with Weak Supervised Learning

Beginning with the deep learning revolution in computer vision and natural language research, deep learning is now permeating at every corner of science and technology. An inherent limitation in adapting deep learning models into new domains is the necessity to train models from scratch, which is particularly time consuming, and further compounded if many different tasks are to be predicted. In this presentation, we will examine how historical research in feature-engineering by domain experts can be used to develop a deep learning domain expert model that can be quickly adapted to predict other tasks that it was not trained on originally. We demonstrate this example in the chemistry domain, by predicting a wide range of physical, biological, and materials properties while using only a small set of labeled data. Our results not only show better performance relative to state-of-the-art models developed using engineered features, but also improved performance relative to analogous deep learning models trained in the conventional supervised learning approach.

Garrett Goh
Pacific Northwest National Labratory
garrett.goh@pnnl.gov

MS210

Neural Network Surrogates of PDE-based Dynamical Systems, Application to Ice Sheet Dynamics

One of the main objectives of climate modeling is to provide reliable probabilistic projections of the sea-level rise over the next decades to centuries. Sea-level rise is mainly due to the ocean thermal expansion and the mass loss of Greenland and Antarctic ice sheets. Accurate prediction of the evolution of the total mass of ice sheets, our quantity of interest, requires the solution of very expensive computational models for solving the ice sheet dynamics that feature a large number of unknowns and uncertain parameters. The high dimensionality of the parameter space and the high computational cost of the forward model make uncertainty quantification analysis extremely challenging. We present a surrogate model built using Recurrent Neural Networks for approximating the parameter-to-prediction map of a transient physical problem. The RNN surrogate is trained using the output of our high-fidelity computational model. The surrogate is then used to quantify the uncertainties of our quantity of interest in a multifidelity approach, where the inexpensive low-fidelity surrogate model is evaluated much more frequently than the high-fidelity computational model in order to maintain accuracy while significantly reducing the computational cost.

Mauro Perego
CCR Sandia National Laboratories
mperego@sandia.gov

MS210

Interpretability and Representation of Data in Recurrent Neural Networks

Recurrent Neural Networks (RNNs) allow for real-time computational capabilities such as sensor fusion, real-time scene understanding, navigation and natural language processing. We study how the information is represented in these networks and how it is being shaped during the training process. In particular, we study the distribution of the underlying graph in typical RNNs such as GRU, LSTM, SeqtoSeq, Random RNN, and classify dominant patterns, which contribute to RNN ability to encode and decode inputs.

Eli Shlizerman
University of Washington, Seattle
shlizee@uw.edu

MS210

Multilevel Adaptive Reduction of Data for Large

Scale Scientific Simulation

We develop a mathematical framework that enables the adaptive compression of data while controlling the loss as measured in a variety of norms. The same technique can be used to ensure that user-supplied quantities of interest are preserved to within a user-prescribed tolerance. The compression algorithms are based on a stable decomposition we call the multilevel splitting. We introduce this splitting and give efficient algorithms for its computation on both structured tensor product grids and general unstructured meshes. In addition to the theory underlying the method we discuss the use of the method on a range of example applications, including an autocatalytic reaction simulation, experimental data obtain from a magnetic confinement experiment, and a simulation of turbulent flow along a rectangular channel. In each case, we consider one or more relevant quantities of interest and reduce the data so as to preserve these quantities. The techniques are gathered in a data reduction package, MGARD, which has been integrated into the I/O system ADIOS.

Ben Whitney
Division of Applied Mathematics
Brown University
ben_whitney@brown.edu

MS211

Efficient Operator-coarsening Multigrid Schemes for Local Discontinuous Galerkin Methods

An efficient *hp*-multigrid scheme is presented for local discontinuous Galerkin (LDG) discretizations of elliptic problems, formulated around the idea of separately coarsening the underlying discrete gradient and divergence operators. We show that traditional multigrid coarsening of the primal formulation leads to poor and suboptimal multigrid performance, whereas coarsening of the flux formulation leads to optimal convergence and is equivalent to a purely geometric multigrid method. The resulting operator-coarsening schemes do not require the entire mesh hierarchy to be explicitly built, thereby obviating the need to compute quadrature rules, lifting operators, and other mesh-related quantities on coarse meshes. We show that good multigrid convergence rates are achieved in a variety of numerical tests on 2D and 3D uniform and adaptive Cartesian grids, as well as for curved domains using implicitly defined meshes and for multi-phase elliptic interface problems with complex geometry. Extension to non-LDG discretizations is briefly discussed.

Daniel Fortunato, Chris H. Rycroft
Harvard University
dfortunato@g.harvard.edu, chr@seas.harvard.edu

Robert Saye
Lawrence Berkeley National Laboratory
rsaye@lbl.gov

MS211

Local Fourier Analysis of BDDC-like Algorithms

Local Fourier analysis is a commonly used tool for the analysis of multigrid and other multilevel algorithms, providing both insight into observed convergence rates and predictive analysis of the performance of many algorithms. In this talk, we adapt local Fourier analysis to examine variants of two- and three-level BDDC algorithms, to better understand the eigenvalue distributions and condition number bounds on these preconditioned operators. This adaptation is based on a new choice of basis for the space of Fourier harmonics that greatly simplifies the application of local Fourier analysis in this setting. The local Fourier analysis is validated by considering the two dimensional Laplacian and predicting the condition numbers of the preconditioned operators with different sizes of subdomains. Several variants are analyzed, showing the two- and three-level performance of the "lumped' variant can be greatly improved when used in multiplicative combination with a weighted diagonal scaling preconditioner, with weight optimized through the use of LFA.

Yunhui He
Memorial University of Newfoundland
yh6171@mun.ca

Jed Brown
University of Colorado Boulder
jed.brown@colorado.edu

Scott Maclachlan
Department of Mathematics and Statistics
Memorial University of Newfoundland
smaclachlan@mun.ca

MS211

Randomized Solvers

The rise of randomized algebraic multigrid methods using the concept of low stretch trees, has been one the biggest developments of the field in the 21st century. In this talk, we cover recent developments in this area.

John Urschel
Massachusetts Institute of Technology
urschel@mit.edu

MS211

Computing the Eigenvalues of the Dirichlet-to-Neumann Map for Indefinite Helmholtz Equation

We introduce an efficient method for computing the eigenvalues of the Dirichlet-to-Neumann (DtN) map associated with the Helmholtz equation. Potential applications are found in the computation of refraction indices of obstacles in inverse scattering. We use periodic extensions and restrictions to subspaces to solve the constant coefficient Helmholtz equation efficiently. The proposed Fourier method, combined with an eigensolver, results in an efficient and clear approach for computing the eigenvalues of DtN map

Yangqingxiang Wu
Department of Mathematics
The Pennsylvania State University

yzw137@psu.edu

MS212

A New Adaptive Block ILU Preconditioning Technique with Application in High-order Simulations of Three-dimensional Aircrafts

Toward developing strong solvers, in this work, an effective, efficient, and parallelizable preconditioning technique is proposed for Newton-Krylov methods. The proposed methodology can be viewed as the generalization of the implicit line preconditioning technique by extending the groups of implicitly-solved unknowns from lines to blocks. In this approach, the blocks are formed based both the strength of connections between unknowns, and the ILU algorithm is used to obtain an approximate or exact factorization for each block. The blocks can be adaptively updated during the solution process and the flexibility in choosing the size of them allows for tailoring of local linear and/or nonlinear smoothers to block sizes which are suitable for current and emerging hardware architectures. To demonstrate the effectiveness of the proposed preconditioner in Newton-Krylov solvers, we have utilized it in second- and third- order steady-state Reynolds-averaged Navier-Stokes (RANS) solutions for three-dimensional aircraft configurations such as the NASA common research model (CRM). For the studied cases, it is shown that in comparison with the traditional ILU(k) method, the proposed preconditioner requires significantly less memory, especially for the higher-orders, and it is more efficient.

Behzad Ahrabi
Department of Mechanical Engineering
University of Wyoming
brezaahr@uwyo.edu

Dimitri Mavriplis
Department of Mechanical Engineering
University of Wyoming
mavripl@uwyo.edu

MS212

Enabling Large-scale Wind Farm Simulations Using An Hp-Adaptive Discontinuous-Galerkin Method and Overset Grids

To accurately predict the physics of a wind farm many individual wind turbines are modeled including their wake interactions. This problem is made difficult due to complex geometry, rotating blades, and downstream wake effects. To address this, our computational methodology uses an overset mesh framework to enable specialized flow solvers in different regions of the flow. The near-body solver is a robust finite-volume scheme designed for unstructured meshes, viscous boundary layers, and implicit time stepping. Whereas the background flow solver is an *hp*-adaptive discontinuous-Galerkin (DG) finite element method which is designed to be highly efficient and accurate. The DG solver is connected to the p4est AMR library which is used for *h*-adaption and automatically adapts to the wind turbine wakes. Also *p*-adaption is used to increase efficiency by using low-order cells near overset regions and high-order cells in the wake. High-order cells are more computationally efficient and accurate per degree of freedom compared to low order cells, they are also less of a burden on the AMR library and the overset mesh library. To demonstrate the efficacy of this approach the weak-scalability of the computational framework is demonstrated using 6, 12, 24, 48, and 96 wind turbine set-ups including the 48 turbine wind plant known as Lillgrund. The largest case consisting of 96 wind turbines and a total of 385 overset grids are run on 44k cores at a weak scaling efficiency of 86%.

Michael Brazell
University of Wyoming
mbrazell@uwyo.edu

Dimitri Mavriplis
Department of Mechanical Engineering
University of Wyoming
mavripl@uwyo.edu

MS212

Wind Turbine Simulation with High-order Hybridized Discontinuous Galerkin Methods

In this study, we investigate the structure of wind turbine wakes and their interaction between two turbines. Since the wake effects can cause power loss, it is important to predict the wake structure accurately and efficiently in a large-scale wind energy project. In order to simulate a three-dimensional wind turbine model, we employ the higher-order hybridized discontinuous Galerkin method (HDG) for the Navier-Stokes equations. Compared to the discontinuous Galerkin (DG) method, the HDG method allows to drastically reduce the degrees of freedom (DOF) by introducing a trace unknown that is defined only on the skeleton of the mesh. The HDG method does not only reduce the DOF but also flexibly couple the non- matching grid for handling the rotating wind turbine by introducing curved mortar interfaces between the rotating mesh and the stationary mesh. To enhance the computational efficiency, we apply the adaptive mesh refinement technique to near the wind turbine. Numerical results for various examples demonstrate that our methods accurately capture the wake structure with low dissipation.

Shinhoo Kang
Department of Aerospace Engineering
The University of Texas at Austin
sh.kang@utexas.edu

Tan Bui-Thanh
The University of Texas at Austin
tanbui@ices.utexas.edu

MS212

Nonhydrostatic Atmospheric Models using Dynamically Adaptive Continuous and Discontinuous Galerkin Methods

We give a report on the NUMA framework, which is a unified continuous and discontinuous Galerkin framework for nonhydrostatic atmospheric models. The code uses tensor product tetrahedral (spectral) elements and is capable of

both regional and global simulations. In this talk, we will describe the NUMA components as well as discuss recent developments to extending the framework for use with non-conforming adaptive grids in 3D. A particular emphasis of the talk will be on stability and conservation on curvilinear and non-conforming meshes.

Jeremy E. Kozdon
Department of Applied Mathematics
Naval Postgraduate School
jekozdon@nps.edu

Francis X. Giraldo
Naval Postgraduate School
fxgirald@nps.edu

Michal Kopera
Boise State University
michalkopera@boisestate.edu

Lucas Wilcox
Department of Applied Mathematics
Naval Postgraduate School
lwilcox@nps.edu

MS213

Analysis of An Atomistic Model for Anti-plane Fracture with Sequential Coupling

In this talk I will present a model for an anti-plane crack defect posed on a lattice supplied with far-field boundary condition from either continuum linearised elasticity or some higher-order continuum theory. In particular, I will discuss existence, local uniqueness and stability of solutions for small loading parameters and further comment on qualitatively sharp far-field decay estimates. If time permits, I will also touch upon the corresponding bifurcation problem, which occurs as we increase the loading parameter beyond a critical value.

Maciej Buze
The University of Warwick
m.buze@warwick.ac.uk

MS213

Large-scale Embedding Quantum Simulations of Twisted Graphene Layers.

High-level many-body quantum methods (full configuration interaction (FCI), coupled cluster (CC), and density matrix renormalization group (DMRG)) can deliver the accurate solutions for a small system. For a large system, the computational cost can be significantly reduced by embedding those high-level models, as impurities, into low-level methods (eg. Hartree-fock (HF) method, and density functional theory (DFT)). In this talk, one of quantum embedding approaches, density matrix embedding theory (DMET), will be firstly applied to the large supercell of low angle twisted graphene layers.

Xiaojie Wu
University of California, Berkeley
wxj6000@gmail.com

MS213

A Neumann-type Boundary Condition for Peridynamics and its Coupling with Fluids

In this work we consider a 2D peridynamic model with a finite nonlocal horizon parameter d which characterizes the range of nonlocal interactions, and its interaction with the surounding fluids. In the nonlocal IDEs the boundary conditions should be defined in a nonlocal way, namely, on a region with non-zero volume outside the surface, while in fluid–structure interaction applications the physical boundary conditions are typically provided on a sharp co-dimension one surface. In this work we present a new nonlocal Neumann-type constraint as an analog to the traction load applied on a sharp interface. Moreover, with rigorous mathematical analysis we show that the peridynamic model with this nonlocal constraint converges to the corresponding local problem with the local traction load with second order convergence as d-¿0. This convergence rate is optimal considering the O(d2) difference from the nonlocal equation to its local limit. As a numerical verification and application of the new constraint, we have implemented the nonlocal boundary condition in the context of an optimization based meshfree quadrature rule [N. Trask, H. You, Y. Yu, M. L. Parks, A meshfree quadrature rule for non-local mechanics, CMAME, 2018.] and coupled it with a numerical solver based on div-conforming B-splines for incompressible flow, to capture the fluid-induced damage on materials.

Yue Yu
Department of Mathematics, Lehigh University
yuy214@lehigh.edu

Huaiqian You
Lehigh University
huy316@lehigh.edu

Nathaniel Trask
Brown University
Division of Applied Mathematics
nat.trask@gmail.com or natrask@sandia.gov

Ming-Chen Hsu
Department of Mechanical Engineering
Iowa State University
jmchsu@iastate.edu

MS213

Understanding the Tunable Adhesion and Wetting of Wrinkled Surfaces via a Lattice Model

In this talk, I will present some recent developments of lattice models for solids to handle structures with complicated geometries and their applications to understand the tunable adhesion and wetting on wrinkled surfaces. I first show that the lattice model can be derived from finite element methods for neo-Hookean solids under large deformation. By coupling the lattice model and nonlocal interactions, we elucidate the nature of enhancing and weakening the adhesion due to surface wrinkles. We further integrate the

lattice model with smoothed-particles hydrodynamics to investigate the static and dynamical wetting of a droplet on wrinkled surfaces. Through these examples, we can see that the simplicity and adoptability of the lattice model open possibilities to develop novel numerical platforms for simulating multiphysics and multiscale problems via integrating it with other modeling techniques.

Teng Zhang
Syracuse University
tzhang48@syr.edu

MS214

Fast and Flexible Bayesian Inference for Low-rank Tensor Problems

Low-rank tensor methods are a common approach to large-scale or multimodal problems. Low-rank assumptions can speed up computation and increase accuracy in tasks ranging from high-dimensional integration to multimodal regression. The main challenge in applying low-rank tensor methods is the determination of tensor rank a priori. The common solution is to use variational Bayesian inference. Traditional variational inference requires strong assumptions and uses inflexible approximating families. We show how to apply fast and flexible variational inference to a wide class of low-rank tensor problems, and show why the flexibility of our method increases accuracy.

Cole Hawkins, Zheng Zhang
University of California, Santa Barbara
colepshawkins@gmail.com, zhengzhang@ece.ucsb.edu

MS214

Low-rank Tensor Methods for Optimal Control of Fractional PDEs

Modelling, simulation and optimization of systems governed by fractional and nonlocal PDEs have attracted considerable interest in recent literature. The nonlocality of the operator leads to a dense discretized system. This presents a major difficulty for standard solvers, since the solution of large, dense problems is computationally very expensive. We present a method to overcome these for a fractional Laplacian control problem by exploiting well-known approximation results for the underlying operator and extending them to the control case. In the discrete case on tensor grids, these properties lead to a low Kronecker rank structure, which allows a cheap computation of matrix-vector products. We use this structure to construct a direct solver for the optimal control problem, whose computational complexity is independent of the problem dimension. We show our novel approach to outperform existing standard methods.

Gennadij Heidel
Trier University
heidel@uni-trier.de

Venera Khoromskaia
Max Planck Institute for Mathematics in the Sciences
vekh@mis.mpg.de

Boris Khoromskij
Max-Planck Institute, Leipzig
bokh@mis.mpg.de

Volker Schultz
Trier University
volker.schulz@uni-trier.de

MS214

Optimal Solutions to Complex Tensor Approximation Problems almost Always Exist

We show that over the complex field, a variety of tensor approximation problems including the best rank-r, best symmetric rank-r, best skew-symmetric rank-r, best r block terms, best k-sparse-plus-rank-r, best G-rank-r where G is any tensor network, best rank-r completion, etc, almost always have an optimal solution. This is well known to be false over the real field.

Lek-Heng Lam, Yang Qi
University of Chicago
lekheng@galton.uchicago.edu,
yangqi@galton.uchicago.edu

Mateusz Michalek
Max Planck Institute
mateusz.michalek@mis.mpg.de

MS214

The Space-time Problem with Model Reduction for Traveling Waves

Large scale spatiotemporal data are ubiquitous across many fields of science and engineering, specifically in fluids, atmospheric science, and neurobiology. Data reduction methods aim to allow big data to be more quickly evaluated by first finding a lower dimensional basis to approximate the data. One of the most popular ways to find a low-rank basis is the Singular Value Decomposition (SVD), which reveals the optimal rank r orthonormal basis for any matrix. This framework naturally generalizes to multivariate data via tensor decompositions. However, SVD fails to provide a compact representation where traveling waves are present because the SVD is inherently a space-time separation of variables. This necessitates a data-driven method to decompose and reduce spatiotemporal systems with multiple traveling waves. In this work, we investigate alternative approaches to dimensionality reduction that are designed specifically to separate and represent interpretable traveling wave structures in low rank. We investigate several formulations of an optimization-based separation and demonstrate on example systems which pose additional challenges of non-periodicity, nonlinearity, and changing wave speed.

Ariana Mendible, Steven Brunton
University of Washington
mendible@uw.edu, sbrunton@uw.edu

J. Nathan Kutz
University of Washington, Seattle
Dept of Applied Mathematics

kutz@uw.edu

MS216

Hyperbolic Stochastic Galerkin Formulations to Model Fluctuations in Supply Networks

Since transmission system operators assume nearly constant consumption, there is a need to assess the risk of stochastic fluctuations. Typically, simulating stochastic processes by a Monte-Carlo method is computationally not feasible. Recently, the representation of stochastic processes by orthogonal polynomials has gained interest. Applications of this procedure have been proven successful for diffusion and kinetic equations. Results for hyperbolic balance laws, which model e.g. gas and water flow, are still partial. A problem is posed by the fact that hyperbolicity of the stochastic system is lost in general. We present a stable hyperbolic stochastic Galerkin formulation that accounts for these stochastic fluctuations.

Stephan Gerster, Michael Herty
RWTH Aachen University
gerster@igpm.rwth-aachen.de, herty@igpm.rwth-aachen.de

MS216

A Graph-based Framework for Load Flow Analysis of Multi-carrier Energy Systems

Multi-carrier energy systems (MES) have become more important over the years as the need for efficient, reliable and low carbon energy systems increases. In MES, different energy carriers, such as electricity and heat, interact with each other leading to one combined energy network. Conventional load flow models for the single-carrier networks have been widely studied. However, they are not able to capture the full extend of the coupling. Recently, different models for multi-carrier networks have been proposed, either using the energy hub concept or using a case specific approach. Although the energy hub concept can be applied to a general MES, it does not state how the graphs of single-carrier networks can be combined into one multi-carrier network. We present a general load flow model framework for MES, based on graph and network theory. We introduce coupling nodes and dummy links to combine the single-carrier networks into a multi-carrier network. Based on this framework, we discuss the effect of coupling on the system of non-linear equations. The framework is tested on a small MES consisting of gas, electricity, and heat. Different coupling models were used for comparison with case studies available in literature. The tests showed that our framework is applicable to a general network, and that it generalizes both the energy hub concept and the case specific approaches.

Anne S. Markensteijn, Kees Vuik
Delft University of Technology
A.S.Markensteijn@tudelft.nl, c.vuik@tudelft.nl

MS216

Fast Numerical Methods for Gas Network Modeling and Simulation

We study the modeling and simulation of gas pipeline networks, with a focus on fast numerical methods for the simulation of transient dynamics. The obtained mathematical model of the underlying network is represented by a nonlinear differential algebraic equation (DAE). With our modeling, we reduce the number of algebraic constraints, which correspond to the $(2, 2)$ block in our semi-explicit DAE model, to the order of junction nodes in the network, where a junction node couples at least three pipelines. We can furthermore ensure that the $(1, 1)$ block of all system matrices including the Jacobian is block lower-triangular by using a specific ordering of the pipes of the network. We then exploit this structure to propose an efficient preconditioner for the fast simulation of the network. We test our numerical methods on benchmark problems of (well-)known gas networks and the numerical results show the efficiency of our methods.

Yue Qiu
Max Planck Institute for Dynamics of Complex Tech. Systems
qiu@mpi-magdeburg.mpg.de

Sara Grundel
MPI Magdeburg
grundel@mpi-magdeburg.mpg.de

Martin Stoll
Chemnitz University of Technology
martin.stoll@mathematik.tu-chemnitz.de

Peter Benner
Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany
benner@mpi-magdeburg.mpg.de

MS216

Rigorous MINLP Optimization of Power-to-methane Processes

Renewable energy sources, such as wind, solar or water power, are undergoing strong fluctuations due to changing environmental conditions. The corresponding fluctuation of the energy supply causes new challenges to handling the energy grid. Offset between supply and demand for energy requires buffering via storage of energy, not only on a daily but on a seasonal scale. For this, in particular, the gas net is an attractive option. Energy can be stored in form of a chemical energy carrier, methane. Methane is produced from surplus renewable energy via water electrolysis and carbon-dioxide methanation. The resulting process connects the energy and natural gas grid, allowing for a more flexible handling of power fluctuations. The drawback of this technology is, however, losses accompanying every conversion step throughout the process. For an efficient coupling of energy and natural gas grid optimization of the conversion process is required. A structural optimization of the process on site of an anaerobic digester, acting as a carbon dioxide source, is proposed. The chemical conversion process is formulated as a MINLP and optimized with respect to exergetic efficiency via a branch-and-bound approach. The process model is given by a superstructure,

considering different process alternatives. A challenge is here is, in particular, the handling of rigorous unit level models. They include partial differential equations for dynamic models such as pressure swing adsorption.

Jennifer Uebbing
Max Planck Institute for Dynamics of Complex Technical Sys.
uebbing@mpi-magdeburg.mpg.de

Sebastian Sager
Otto-von-Guericke Universität Magdeburg, Germany
sager@ovgu.de

Kai Sundmacher
Max Planck Institute for Dynamics of Complex Technical Sys.
sundmacher@mpi-magdeburg.mpg.de

MS217

Entropy-stable, High-order Discretizations based on Continuous Summation-by-parts Operators

This talk will describe an entropy-stable (semi) discretization of the Euler equations based on continuous summation-by-parts (SBP) operators. These continuous SBP discretizations assemble global operators from element operators, and they are the SBP analog of continuous Galerkin finite-element methods. Consequently, a stabilization method is needed to suppress high-frequency oscillations and ensure optimal convergence rates. To this end, we present a high-order stabilization that ensures entropy is dissipated. Furthermore, the stabilized discretizations have favorable spectral radii when used in conjunction with optimized SBP first-derivative operators. Numerical experiments demonstrate the accuracy and efficiency of the continuous SBP discretizations relative to more popular discontinuous SBP discretizations.

Jason E. Hicken
Rensselaer Polytechnic Institute
Assistant Professor
hickej2@rpi.edu

MS217

Convex Limiting for Nonlinear Hyperbolic Systems with Source Terms

We introduce an approximation technique for nonlinear hyperbolic systems with sources that is invariant domain preserving. The low order method is discretization independent provided elementary symmetry and skew-symmetry properties are satisfied by the scheme. The method is formally first-order accurate in space. Then, we introduce a second-order method which could violate the invariant domain properties of the PDE. Then, the second order method is corrected by a limiting technique that we call convex limiting. After limiting, the resulting method satisfies all the invariant domain properties that are imposed by the user and is formally second-order accurate. The two key novelties are that limiting is done by enforcing bounds on quasiconcave functionals and the bounds that are enforced on the solution at each time step are necessarily satisfied by the low-order approximation.

Bojan Popov
Department of Mathematics
Texas A&M University
popov@math.tamu.edu

MS217

Fully Discrete Entropy Stable and Kinetic Energy Preserving Discontinuous Galerkin Methods

We present a space-time discontinuous Galerkin (DG) method for systems of non-linear hyperbolic conservation laws that is entropy conservative (for smooth solutions) or entropy stable (for discontinuous solutions). The resulting numerical scheme is fully discrete and provides a bound on the mathematical entropy according to its initial condition and boundary conditions. The crux of the method is that discrete derivative approximations in space and time are summation-by-parts (SBP) operators in order to mimic results from the continuous entropy analysis. Importantly, the presented space-time method does not assume any exactness of quadrature in the variational forms that arise in the DG method. Typically, the development of entropy stable numerical schemes is done on the semi-discrete level ignoring the temporal dependence. In this talk we describe how an entropy stable DG method in time is similar to the spatial discrete entropy analysis, but there are important differences. For the compressible Euler equations, the preservation of kinetic energy is also of interest. Again, construction of kinetic energy preserving (KEP) schemes is typically done on the semi-discrete level. We generalize the KEP condition from Jameson to the space-time framework and provide the temporal components for both entropy stability and kinetic energy preservation. The properties of the presented space-time DG method are validated through numerical tests for the compressible Euler equations.

Gregor Gassner, Andrew R. Winters
Mathematical Institute
University of Cologne
ggassner@math.uni-koeln.de, awinters@math.uni-koeln.de

David Flad
University of Stuttgart
flad@iag.uni-stuttgart.de

Gero Schnuecke, Lucas Friedrich
Mathematical Institute
University of Cologne
gschnuec@math.uni-koeln.de, lfriedri@math.uni-koeln.de

MS218

Shannon & Shrdinger: Statistical Algorithms and Architectures

We describe the interplay between statistical physics, quantum annealing, Bayesian inference and Boltzmann machines to provide insights into areas where quantum computation and quantum simulation will play a key role in solving data-driven real-world problems. We present an overview of the D-Wave 2000Q quantum computer archi-

tecture and discuss several novel optimization and machine learning algorithms that have been designed to leverage current and future D-Wave quantum computers.

Aaron Lott
D-Wave Systems, Inc.
palott@gmail.com

MS218

Non-von Neumann Computing using Networks of Optical Parametric Oscillators

Combinatorial optimization problems are central in numerous important application areas, including operations and scheduling, drug discovery, finance, circuit design, sensing, and manufacturing. There is a long history of attempts to find alternatives to current von Neumann-computer-based methods for solving such problems, including neural networks, DNA computing, and most recently adiabatic quantum computation and quantum annealing.
Networks of coupled optical parametric oscillators (OPOs) are an alternative physical system, with an unconventional operating mechanism, for solving the Ising problem, which is an NP-hard optimization problem. We have realized a fully-programmable 100-spin Ising machine using a network of OPOs, and with it can solve many different Ising problem instances. Our design supports all-to-all connectivity among the implemented spins via a combination of time-division multiplexing and measurement feedback.
In this talk I will describe our work on constructing Ising machines using OPO networks with feedback, and will present the experimental results from our first prototype system.

Peter McMahon
Stanford University
pmcmahon@stanford.edu

MS218

Leveraging Future Computing Platforms: An Open Discussion

We will close the minisymposium with an open discussion on the future of scientific computing. While uncertainty is the only certainty there is, identifying the promising current and near-future hardware trends and developing the algorithmic modifications and innovations needed to leverage them is of great importance to the advancement of computational science and engineering. Audience participation is encouraged.

Dimitar Trenev
ExxonMobil Research and Engineering
dimitar.trenev@exxonmobil.com

MS219

Algorithm Based Error Analysis for Mixed Precision Matrix Factorizations.

Many modern computer architectures now support low precision arithmetics and exhibit sky-high performance when it comes to performing flop-intensive operations such as matrix multiplications, a keystone in numerical linear algebra and scientific computing. This talk will give some elements to understand how to make a smart use of half precision arithmetic on modern accelerators in order to perform matrix factorizations efficiently and accurately. Focus will be made on tracking and monitoring roundoff errors using algorithmic tools as well as data- and architecture-aware error analysis.

Pierre Blanchard
School of Mathematics
The University of Manchester
pierre.blanchard00@gmail.com

MS219

Bayesian Analysis of the Effects of Lower Precision Arithmetic in Inverse Problems

The Bayesian framework for inverse problems models all unknown parameters as random variables, randomness being a way to convey uncertainty in their value, and described in terms of probability densities. The Bayesian solution of an inverse problem is the posterior density of the unknown of primary interest: a priori belief about the solution and errors in the data determine the shape of the posterior, hence the variability that can be expected of the solution. Evaluation of the forward model, linking the unknown of interest to the observations in lower precision arithmetic may be particularly attractive when the problems has to be solved in real time or for computationally costly forward models. In this talk we will propose ways to account for the effects of lower precision arithmetic on the posterior, to decide when the problem is amenable to lower precision computation and to compensate for the ensuing roundoff error by treating it as some form of modeling error.

Daniela Calvetti
Case Western Reserve Univ
Department of Mathematics, Applied Mathematics and Statistic
dxc57@case.edu

MS219

Approximate Arithmetic - A Hardware Perspective

Computer architecture is changing, and numerical computation with it. This talk will introduce FPGAs - "Lego circuits" - and their application in numerical computing, with a focus on the potential and possibilities for customisation of numerical representation on FPGAs. These devices are being rapidly adopted in industry, with low-power deep neural networks a current key numerical driver. Customisation options range from "just enough floating point precision" to exact real arithmetic, with many other options in between. There are two major pitfalls: (a) placing the burden of custom arithmetic design on each individual programmer or (b) squandering the potential of customisation using only standard components. I will propose a roadmap to overcome these pitfalls, based on formal reasoning and automated synthesis of numerical circuits.

George Constantinides
Imperial College London

g.constantinides@imperial.ac.uk

MS219

Automated Backward Error Analysis

Backward Error Analysis (BEA) is fundamental in estimating uncertainty and stability of numerical algorithms. Numerical analysts have applied BEA manually, in the paper-and-pencil style, as it lacks an automated solution. This work presents an automated BEA. The computed backward error results are then used to compute the condition number. Experiments are conducted on Intel x87 FPU instructions and GNU C Library functions. Our results demonstrate the effectiveness and efficiency of the proposed analysis. To the best of our knowledge, this is the first successful attempt for automated BEA.

Zhoulai Fu
Department of Computer Science
IT University of Copenhagen
Zhoulai@ucdavis

MS220

A Third Order Accurate Wave Propagation Algorithm for Hyperbolic Partial Differential Equations

I will present recent work on the extension of LeVeque's wave propagation algorithm, a widely used finite volume method for hyperbolic partial differential equations, to a third order accurate method. The resulting method shares main properties with the original method, i.e. it is based on a wave decomposition at grid cell interfaces and limiting is introduced in the form of a wave limiter.

Christiane Helzel
Department of Mathematics
Ruhr-University-Bochum
christiane.helzel@hhu.de

MS220

Fully Discrete Energy Decay Finite Difference Schemes for Poisson-Nernst-Planck Equations

The Poisson-Nernst-Planck (PNP) model has a long history in description of electro-diffusion, semiconductor and electrochemical systems. Recently this model has been widely applied to the study of biological transport especially the ionic channel structure. When designing a numerical scheme for the PNP model, it would be desirable to preserve the three main properties of the solution: mass conservation, non-negativity, and energy dissipation. In this work, we introduce implicit finite difference schemes for the PNP equations that are able to preserve these properties. In particular, we prove that the fully discrete schemes decay energy under a generic boundary condition. Numerical examples in multi-D and for multiple species will be presented.

Xiaodong Huang
Purdue University
huan1178@purdue.edu

Jingwei Hu
Department of Mathematics, Purdue University
jingweihu@purdue.edu

MS220

A Positive Asymptotic Preserving Scheme for Linear Kinetic Transport Equations

We present a positive and asymptotic preserving numerical scheme for solving linear kinetic, transport equations that relax to a diffusive equation in the limit of infinite scattering. The proposed scheme is developed using a standard spectral angular discretization and a classical micro-macro decomposition. The three main ingredients are a semi-implicit temporal discretization, a dedicated finite difference spatial discretization, and realizability limiters in the angular discretization. Under mild assumptions on the initial condition and time step, the scheme becomes a consistent numerical discretization for the limiting diffusion equation when the scattering cross-section tends to infinity. The scheme also preserves positivity of the particle concentration on the space-time mesh and therefore fixes a common defect of spectral angular discretizations. The scheme is tested on the well-known line source benchmark problem with the usual uniform material medium as well as a medium composed from different materials that are arranged in a checkerboard pattern. We also report the observed order of space-time accuracy of the proposed scheme.

Paul Laiu
Oak Ridge National Laboratory
laiump@ornl.gov

Martin Frank
Karlsruhe Institute of Technology
martin.frank@kit.edu

Cory Hauck
Oak Ridge National Laboratory
hauckc@ornl.gov

MS220

Efficient and Highly Accurate Semi-Lagrangian Discontinuous Galerkin Method for Convection-diffusion Problems

We propose to organically combine semi-Lagrangian discontinuous Galerkin (SLDG) method to convection terms with local DG (LDG) scheme to diffusion terms for convection-diffusion problems. In particular, we apply a weak formulation of the SLDG method for the convection term (Cai, Guo and Qiu, JSC, 2017), and use high order implicit Runge-Kutta method for a LDG discretization of the diffusion term along characteristics . The proposed scheme is shown to be mass conservative, high order accurate in both space and in time, and highly efficient due to large time stepping sizes allowed from the semi-Lagrangian and implicit nature of time discretization. The scheme can be straightforwardly extended to 2D problems without operator splitting in the truly multi-D SLDG framework previously proposed. The performance of the scheme will be showcased by several classical linear test problems such as rigid body rotation, as well as incompressible Navier-Stokes

equations.

Jingmei Qiu
University of Daleware
jingqiu@udel.edu

MS221

Sampling and Clustering on the Pod-Grassmann Manifold

One of the main issue in multi-parametric reduced order modeling is how to perform a sampling of the input parameter space. In this study, we will focus on Proper Orthogonal Decomposition to reduce the solution space. Many previous studies have been concerned with error estimators to sample the input parameter space. The drawback of such approaches is that there are no guaranties to obtain an appropriate, quasi-uniform, sampling of the solution space. In this study, our goal is thus to compute equidistant POD subspaces covering the whole input parameter space. In a first approach, we will perform an iterative sampling of the input parameter space based on a criterion linked to a metric computed from the Principal Angles (or geodesic on the Grassmann manifold). In a second approach, starting from a given sampling, equidistant POD subspaces can be computed from a snapshots clustering using an appropriate metric. Simple examples for Burgers and two-dimensional Navier-Stokes problems will be given.

Michel Bergmann
Universite Bordeaux
michel.bergmann@inria.fr

Angelo Iollo
Institut de Mathématiques de Bordeaux
angelo.iollo@math.u-bordeaux1.fr

MS221

Learning a Hierarchy of Context-aware Low-fidelity Models for Adaptive Multifidelity Monte Carlo Estimation

Traditional model reduction aims to construct low-cost low-fidelity models to replace computationally expensive high-fidelity models for speeding up computations. However, in multifidelity methods, low- and high-fidelity models are used together and so the primary purpose of low-fidelity models is supporting computations with the high-fidelity models rather than approximating and replacing them. In this presentation, we introduce context-aware low-fidelity models that are explicitly constructed for being used together with high-fidelity models in multifidelity settings. We discuss the adaptive multifidelity Monte Carlo (AMFMC) method that learns context-aware low-fidelity models for the estimation of statistics of high-fidelity model outputs. The context-aware low-fidelity models are constructed by quasi-optimally trading off adapting the low-fidelity models–to improve their deterministic approximation quality–with sampling the models–to reduce the statistical error. Our analysis shows that the quasi-optimal computational effort to spend on improving the low-fidelity models is bounded, which means that low-fidelity models can become too accurate for multifidelity methods, which is in stark contrast to traditional model reduction. Numerical results demonstrate that our context-aware low-fidelity models outperform traditional model reduction approaches by orders of magnitude if low-fidelity models are used in multifidelity settings.

Ionut-Gabriel Farcas
Technical University of Munich
Chair of Scientific Computing in Computer Science
farcasi@in.tum.de

Benjamin Peherstorfer
Courant Institute of Mathematical Sciences
New York University
pehersto@cims.nyu.edu

MS221

PGD-based Computational Vademecums for Parameterized Flows

Computational fluid dynamics simulations are ubiquitous in nowadays automotive and aerospace industry. Some numerically evaluated quantities of interest (e.g. drag and lift) match reasonably well experimental results. However, parametric studies are still unreasonably time and computer demanding in practical problems of interest for engineers. Typical parameters affecting flow features are user-defined data, such as boundary conditions or the prescribed domain. Here the proper generalized decomposition (PGD) is employed to obtain explicit parametric approximations, starting from an efficient separated representation of a hybridizable discontinuous Galerkin (HDG) solution. The resulting strategy inherits all advantages of HDG for both low and high-order spatial discretizations [M. Giacomini, A. Karkoulias, R. Sevilla, and A. Huerta. A superconvergent HDG method for Stokes flow with strongly enforced symmetry of the stress tensor arXiv:1802.09394 (2018)]. In particular, for problems of industrial interest, low-order, i.e. face-centered finite volumes, can be employed for fast and robust computations [R. Sevilla, M. Giacomini, and A. Huerta. A face-centred finite volume method for second-order elliptic problems Int. J. Numer. Methods Eng. 115(8), pp. 986-1014 (2018)], while high-fidelity solutions can be obtained by means of high-order approximations [R. Sevilla, and A. Huerta. HDG-NEFEM with degree adaptivity for Stokes flows to appear in J. Sci. Comput. (2018)].

Matteo Giacomini
Universitat Politècnica de Catalunya
matteo.giacomini@upc.edu

Ruben Sevilla
Swansea University
r.sevilla@swansea.ac.uk

Antonio Huerta
Universitat Politècnica de Catalunya
antonio.huerta@upc.edu

MS221

Goal-oriented Model Reduction of Parametrized

Nonlinear PDEs in Aerodynamics

We introduce a goal-oriented model reduction framework for rapid and reliable solution of parametrized nonlinear partial differential equations (PDEs) in aerodynamics. Our framework builds on the following ingredients: a discontinuous Galerkin (DG) finite element (FE) method, which provides stability for convection-dominated problems; reduced basis (RB) spaces, which provide rapidly convergent approximation of the solution manifold; the dual-weighted residual (DWR) method, which provides effective output error estimates; hyperreduction of the primal residual, adjoint residual, and output forms by the empirical quadrature procedure (EQP), which provides online-efficient evaluation of the nonlinear forms while providing quantitative control of hyperreduction errors; and output-based adaptive computation of the snapshots. In the offline stage, a reduced model is constructed using a weak greedy algorithm which simultaneously identifies appropriate FE spaces, RB spaces, and EQPs; in the online stage, the reduced model provides rapid output predictions and effective error estimates. We demonstrate the framework for a parametrized drag prediction problem in aerodynamics modeled by the Reynolds-averaged Navier-Stokes equations.

Masayuki Yano
University of Toronto
myano@utias.utoronto.ca

MS222

Efficient Implementation of Communication-avoiding One-step Ader-Dg Schemes for Nonlinear Hyperbolic Systems of Balance Laws Using AVX 512 Instructions

In this talk we present a new and highly efficient implementation of arbitrary high order accurate ADER discontinuous Galerkin (DG) finite element schemes for the solution of general nonlinear systems of hyperbolic balance laws. ADER-DG is a novel, genuinely communication-avoiding family of algorithms, which achieves high order of accuracy in time not via the standard multi-stage Runge-Kutta (RK) time discretization like most other DG schemes, but at the aid of an element-local predictor stage. This predictor stage is arithmetically very intensive and can be fully carried out within the L2 cache of modern Intel Scalable CPUs. ADER-DG schemes only require one MPI communication per time step, instead of s communication steps of an s-stage RK scheme. Furthermore, the algorithm is naturally cache blocking and can be very efficiently vectorized at the aid of the AVX 512 SIMD instructions of modern Intel CPUs. We make use Intel AVX 512 intrinsics and for evaluating some dense tensor-matrix products, some lines of the algorithm are directly coded in assembler. We show a series of computational examples, ranging from the equations of nonlinear elasticity over the equations of relativistic MHD to the novel FO-CCZ4 formalism of the Einstein equations. We show preliminary results obtained on an Intel Xeon Scalable test platform.

Francesco Fambri
University of Trento
francesco.fambri@unitn.it

MS222

Optimizing Performance for Portable Generic Finite Element Interfaces

One of the challenges with high-order finite element and spectral element methods is that a global sparse matrix is no longer a good representation of a high-order linear operator, both with respect to the FLOPs needed for evaluation and the memory transfer needed for a matrix-vector multiply. libCEED is an extensible library that provides a portable algebraic interface and optimized implementations suitable for high-order operators with extensible solvers on hetrogeneous architectures. We discuss performance optimizations targeting CPU architectures using element vectorization, SIMD intrinsics, and LIBXSMM. libCEED performance is compared to native implementations in production software, such as MFEM and Nek5000. We provide examples relevant to a variety of applications.

Jeremy Thompson
University of Colorado Boulder
jeremy.thompson@colorado.edu

MS222

Experiences with Code Generation for High-order ADER-DG Schemes in SeisSol

SeisSol is a software package for simulating seismic wave propagation in heterogeneous media, as well as simulating dynamic rupture physics. The ADER-DG scheme of SeisSol boils down to matrix chain multiplications of small matrices. As all sizes and sparsity patterns of the matrices are known a priori, we have a lot of optimization potential, which we exploit with automatic code generation. We have shown that the code generation ansatz leads to a high single node performance, up to about 50 percent of peak performance, on recent architectures such as Intel's Haswell or Knights Landing architecture. In the minisymposium, we present an overview of our code generation ansatz as well as our hybrid OpenMP/MPI parallelization scheme and our experiences with recent many-core processors such as the Intel Scalable processor. Furthermore, we present our efforts to extend code generation to novel applications which involve tensor contractions.

Carsten Uphoff
Technische Universität München
uphoff@in.tum.de

Michael Bader
Technical University of Munich
bader@in.tum.de

MS223

Automatic Bayesian Cubature with Quasi-Monte Carlo Sampling

Automatic cubatures provide approximations to multidimensional integrals that satisfy user-specified error tolerances. For multidimensional problems, the sampling den-

sity is fixed, but the sample size, n, is determined automatically. Bayesian cubature postulates that the integrand is an instance of a stochastic process. Prior information about the mean and covariance of this process is used to form data-driven error bounds. However, the process of inferring the mean and covariance governing the stochastic process from n integrand values involves computing matrix inverses and determinants, which are in general time-consuming $O(n^3)$ operations. Our work employs low discrepancy data sites and matching kernels that lower the computational cost to $O(n \log n)$. The confidence interval for the Bayesian posterior error is used to choose n automatically to satisfy the user-defined error tolerance. This approach is demonstrated using rank-1 lattice sequences and shift-invariant kernels. We develop a concept called *Fast transform kernel* to enable this technique usable with other matching kernels and data sites. This is a joint work with my doctoral adviser Prof. Fred J. Hickernell.

Jagadeeswaran R
Department of Applied Mathematics
Illinois Institute of Technology
jrathin1@hawk.iit.edu

MS223

Mc and Qmc Algorithms for the Calculation of Dempster-Shafer Belief Functions

Belief function theory, also known as Dempster-Shafer theory or evidence theory, gives us a general framework for quantifying, representing, and managing uncertainty, and it is widely used in several applications from artificial intelligence to accounting. The belief function theory provides tools to combine several sources' opinions (belief functions), among which, Dempster's rule of combination is the most commonly used. One of the main drawbacks of using Dempster's rule to combine belief functions is that it is #P-complete, which limits the application of Dempster's rule to small number of belief functions. We will introduce some Monte Carlo and quasi Monte Carlo algorithms that approximate Dempster's rule of combination, and present examples that illustrate how Monte Carlo methods can extend the scope of the applications for belief function theory.

Nima Salehy
Department of Mathematics
Florida State University
nimasalehy@math.fsu.edu

MS223

An Adaptive Quasi-Monte Carlo Method for Bayesian Inference with User-specified Error Tolerance

Bayesian inference is based on observed data plus prior beliefs. Computing the expected value of a parameter via Bayesian inference involves the numerical approximation of the quotient of two intractable integrals. In this paper, an adaptive quasi-Monte Carlo(QMC) method is proposed to evaluate this quotient to a user-specified error tolerance. The method is illustrated by a logistic regression model. The efficiency of the computation using two different sampling distributions and their combination is studied.

Kan Zhang
IIT
Illinois Institute of Technology
kzhang23@hawk.iit.edu

MS224

Applying Integral Transforms using Nonoscilatory Phase Functions

It has long been known that certain second order differential equations, such as those defining the Jacobi polynomials and the Bessel functions, admit nonoscillatory phase functions. We will discuss a theorem which states that essentially all second order differential equations with smooth coefficients have this property, Moreover, we will describe a fast method for the numerical calculation of nonoscillatory phase functions and describe its use in the rapid numerical evaluation of certain special functions and in the rapid application of various discrete Sturm-Liouville eigenfunction transforms.

James Bremer
UC Davis
bremer@math.ucdavis.edu

MS224

Fast Algebras on Optimized Butterfly Structures

Butterfly-based direct integral equation solvers represent promising techniques to rapidly solve high-frequency scattering and propagation problems that involve many RHSs or ill-conditioned systems. These solvers rely on butterfly-compression of blocks in the forward and inverse discretized IE operators and fast randomized algebras (e.g., addition, multiplication, inversion, and Schur update of given butterflies into new butterflies) to construct reduced-complexity solvers. However, current butterfly solvers are not optimal due to redundant rank structures of existing butterfly algorithms. This work proposes an algebraically optimized butterfly structure that gets rid of the rank redundancy and opens new possibilities to construct optimal-complexity butterfly solvers.

Yang Liu
Computational Research Division, Lawrence Berkeley National
liuyangzhuan@lbl.gov

MS224

Recent Advances in Butterfly Algorithms

Butterfly algorithms generalize the data structure and compression inherent in the Fast Fourier Transform to oscillatory kernels other than complex exponentials. They have found various applications from special function transforms (e.g. large-scale spherical harmonic transforms) to high-frequency wave propagation problems. In this talk, we will discuss recent advances in computing butterfly-compressed forms of oscillatory integral operators, as well

as their role in high-frequency wave problems.

Michael O'Neil
New York University
oneil@cims.nyu.edu

MS224

A Unified Framework for Oscillatory Integral Transforms: When to use Nufft Or Butterfly Algorithms?

This talk introduces a nearly linear scaling unified framework for evaluating the matvec $g = Kf$, where K is the discretization of an oscillatory integral transform $g(x) = \int K(x,\xi)f(\xi)d\xi$ with a kernel function $K(x,\xi) = \alpha(x,\xi)e^{2\pi i\Phi(x,\xi)}$, where $\alpha(x,\xi)$ is a smooth amplitude function , and $\Phi(x,\xi)$ is a piecewise smooth phase function with $O(1)$ discontinuous points in x and ξ. This unified framework is based on either the non-uniform fast Fourier transform (NUFFT) or the butterfly factorization (BF), together with an $O(N)$ fast algorithm to determine whether NUFFT or BF is more suitable.

Haizhao Yang
National University of Singapore
matyh@nus.edu.sg

MS225

Highly Effective A-Posteriori Error Estimation for Approximation of Nonlinear Problems

Quantification of the error induced by numerical approximation schemes is an important aspect for certification of the numerical results. We show that a classical error estimation technique for nonlinear problems based on the Calloz-Rappaz framework can be considerably improved by introducing an additional auxiliary linear approximation problem. The resulting error bounds can be shown to be rigorous and highly effective. The application of this technique in reduced basis methods, yields rigorous and tight error bounds for very general linear and nonlinear problems such as parametric PDEs, ODEs, vectorial or matrix equation systems.

Bernard Haasdonk, Andreas Schmidt, Dominik Wittwar
University of Stuttgart
haasdonk@mathematik.uni-stuttgart.de,
schmidta@mathematik.uni-stuttgart.de,
dominik.wittwar@mathematik.uni-stuttgart.de

MS225

Some New Results on PBDW, Application to Time-dependent Problems

We shall present in this presentation some new results and methods in the frame of data assimilation based on model reduction approaches. The approach is a generalization of the parameterized-background data-weak approach to variational data assimilation (PBDW) introduced in [Maday, Patera, Penn, Yano, Int J Numer Meth Eng, 2015] dealing with the treatment of data polluted with random noise and also time-dependent problems, the time providing a new direction of assimilation.

Yvon Maday
Universite Pierre et Marie Curie
and Brown university
maday@ann.jussieu.fr

MS225

Model Reduction of Systems of Hyperbolic Conservation Laws

Hyperbolic conservation and balance laws govern a large variety of phenomena of interest from supersonic air-flow to astronomical hydrodynamics to coastal hazards such as tsunamis and storm surge. Uncertainty in these flows is a natural result of these systems being complex and difficult to observe fully and often involve goals that require data assimilation, sensitivity studies and parameter estimation. As is often common in these applications the use of high-fidelity forward models in conjunction with traditional uncertainty techniques often yields a computationally intractable problem. Instead reduced order models are often constructed to reduce this overhead. However the hyperbolic PDEs that govern these applications are notoriously impervious to traditional dimension reduction techniques and therefore reduced order model construction. In this talk we will discuss these difficulties and outline a new approach to constructing the reduced order models, named displacement interpolation, that leverages optimal transport maps and Radon transforms.

Kyle T. Mandli
Columbia University
Applied Physics and Applied Mathematics
kyle.mandli@columbia.edu

Donsub Rim
Columbia University
dr2965@columbia.edu

MS225

Certified Reduced Basis Methods for Variational Data Assimilation

In order to approximate the state of a physical system, data from physical measurements can be incorporated into a mathematical model to improve the state prediction. Discrepancies between data and models arise, since on the one hand, measurements are subject to errors and, onthe other hand, a model can only approximate the actual physical phenomenon. In this talk, we present a model order reduction method for (an interpretation of) the 3D- and 4D-VAR methods of variational data assimilation for parametrized partial differential equations. The classical 3D- and 4D-VAR methods make informed perturbations in order to find a state closer to theobservations while main physical laws described by the model are maintained. For the 3D-VAR method, we take inspiration from recent developments in state and parameter estimation and analyse the influence of the measurement space on the amplification of noise. Here, we prove a necessary and sufficient condition for the identification of a good measurement space which can, in turn, be used for a stability-based selection of measurement

functionals. For both 3D- and 4D-VAR we propose a certified reduced basis (RB) method for the estimation of the model correction, the state prediction, the adjoint solution, and the observable misfit. Finally, we introduce different approaches for the generation of the RB spaces suited for different applications, and present numerical results testing their performance.

Sebastien J. Boyaval
Laboratoire Saint-Venant
Ecole des Ponts ParisTech
sebastien.boyaval@enpc.fr

Martin Grepl, Mark Kaercher, Nicole Nellesen
RWTH Aachen University
grepl@igpm.rwth-aachen.de,
kaercher@aices.rwth-aachen.de,
nicole.nellesen@rwth-aachen.de

Karen Veroy
AICES
RWTH Aachen University
veroy@aices.rwth-aachen.de

MS226

New Strategies for Inference with Multilevel Monte Carlo

Abstract not available

Kody Law
University of Manchester
kodylaw@gmail.com

MS226

Nonlinear Filtering with Local Couplings

We introduce a class of structure-exploiting nonlinear filters for high dimensional state-space models. The idea is to transform a forecast ensemble into samples from the current filtering distribution via a sequence of local (in state-space) nonlinear transport maps, computed mostly via low-dimensional convex optimization. Construction of the maps is regularized by leveraging potential structure in the filtering problem—e.g., decay of dependence, approximate conditional independence, and local likelihoods—thus extending notions of localization to nonlinear updates. Many square-root ensemble filters can be interpreted as special instances of the proposed framework when we restrict our attention to linear transformations. We consider applications to chaotic dynamical systems.

Youssef M. Marzouk, Alessio Spantini
Massachusetts Institute of Technology
ymarz@mit.edu, spantini@mit.edu

Ricardo Baptista
MIT
rsb@mit.edu

MS226

Localization for MCMC - Sampling High Dimensional Posterior Distributions with Local Structure

We investigate how ideas from covariance localization in numerical weather prediction can be used in Markov chain Monte Carlo (MCMC) sampling of high-dimensional posterior distributions arising in Bayesian inverse problems. To localize an inverse problem is to enforce an anticipated "local' structure by (i) neglecting small off-diagonal elements of the prior precision and covariance matrices; and (ii) restricting the influence of observations to their neighborhood. For linear problems we can specify the conditions under which posterior moments of the localized problem are close to those of the original problem. We explain physical interpretations of our assumptions about local structure and discuss the notion of high dimensionality in local problems, which is different from the usual notion of high dimensionality in function space MCMC. The Gibbs sampler is a natural choice of MCMC algorithm for localized inverse problems and we demonstrate that its convergence rate is independent of dimension for localized linear problems. Nonlinear problems can also be tackled efficiently by localization and, as a simple illustration of these ideas, we present a localized Metropolis-within-Gibbs sampler. Several linear and nonlinear numerical examples illustrate localization in the context of MCMC samplers for inverse problems.

Matthias Morzfeld
University of Arizona
Department of Mathematics
mmo@math.arizona.edu

Xin T. Tong
National University of Singapore
Department of Mathematics
mattxin@nus.edu.sg

Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu

MS226

Data Assimilation for Stochastic Advection by Lie Transport Euler SPDE

We present an effective high dimensional data assimilation methodology for the damped and forced incompressible 2D Euler fluid flows driven by stochastic advection by Lie transport (SALT) type noise. SALT type stochastic partial differential equations (SPDE) for GFD were introduced by [Holm, Proc Roy Soc, 2015]. According to [Holm, Proc Roy Soc, 2015] and [Cotter et al., 2017], the principles of transformation theory and multi-time homogenisation imply a physically meaningful, data-driven approach for decomposing the fluid transport velocity into its drift and stochastic parts, for a certain class of fluid flows. This results in a stochastic parameterisation for the fluid transport and gives us SPDE models that preserve circulation. A numerical methodology for implementing this velocity decomposition was developed in [Cotter et al., 2018] and applied to the 2D Euler equations in consideration, on a simply connected domain with no-penetration bc. Successful uncertainty quantification results were obtained thus forming the ground work for this data assimilation study. Our

SPDE model is prescribed on coarse resolutions, whereas the original deterministic partial differential equations are prescribed on fine resolutions. We consider two stochastic filtering data assimilation problems. In the first problem the observations come from a single realisation of the SPDE. In the second problem the observations correspond to the fine scale PDE system.

Colin J. Cotter
Imperial College London
Department of Mathematics
colin.cotter@imperial.ac.uk

Dan Crisan
Imperial College
d.crisan@imperial.ac.uk

Darryl D. Holm, Wei Pan
Imperial College London
d.holm@imperial.ac.uk, wei.pan@imperial.ac.uk

Igor Shevchenko
Imperial College
i.shevchenko@imperial.ac.uk

MS227

Modal Decomposition for Fluid-structure Interaction, and its Application to Flag Flapping

We focus in this talk on data-driven modal decomposition methods for fluid-structure interaction problems, where the flow and immersed-body dynamics are coupled to one another. Despite this coupling, modal decompositions are typically performed on either the fluid or the structure separately, which can obscure driving physical mechanisms in the omitted quantity. We present a framework for simultaneously decomposing the fluid and the structure, allowing for fluid-structure correlations to directly inform the resulting modes. As part of this formulation, we present a physically meaningful norm by which to perform the decomposition. Our framework can be used directly with standard techniques such as proper orthogonal decomposition, dynamic mode decomposition, and spectral proper orthogonal decomposition. We then consider the canonical problem of two-dimensional flow past a flapping flag, and use our coupled fluid-structure formulation to identify the mechanisms driving the well known bifurcation from limit-cycle flapping to chaotic flapping.

Andres Goza
University of Illinois at Urbana-Champaign
ajgoza@gmail.com

Tim Colonius
Division of Engineering and Applied Science
California Institute of Technology
colonius@caltech.edu

MS227

Data-driven Feedback Control Strategies for Unsteady Flows

Unsteady fluid flows are high-dimensional characterized by nonlinear interactions and controlling them is challenging and computationally expensive. We develop data-driven networked-oscillator (model-based) and cluster-based (model-free) strategies for feedback control of unsteady fluid flows. The networked-oscillator-based analysis models the nonlinear interactions by tracking amplitude and phase perturbations among modal oscillators. The linear network model is utilized to design a model-based feedback controller that suppresses the modal amplitudes, leading to drag reduction. We also introduce a model-free self-learning feedback control strategy that leverages unsupervised clustering for in-situ learning and optimization of coarse-grained control laws to manipulate turbulent post-stall flows over an airfoil in high-fidelity simulations. The approach partitions baseline flow trajectories into clusters, which correspond to coarse-grained phases in a low-dimensional feature space. A feedback control law is sought for each cluster state through iterative evaluation and downhill simplex search to minimize power consumption in flight. Re-routing the flow trajectories modifies the baseline Markov transition network amongst the clusters to achieve aerodynamically favorable states with control enabling the flows to reach a low-drag state. Supported by the US Air Force Office of Scientific Research (Grant FA9550-16-1-0650)

Aditya G. Nair
Florida State University
Mechanical Engineering
agn13@my.fsu.edu

Chi-An Yeh
Florida State University
cy13d@my.fsu.edu

Eurika Kaiser
University of Washington
eurika@uw.edu

Bernd Noack
LIMSI-CNRS
Technical University Braunschweig
bernd.noack@limsi.fr

Steven Brunton
University of Washington
sbrunton@uw.edu

Kunihiko Taira
Florida State University
Mechanical Engineering
ktaira@seas.ucla.edu

MS227

Feedback Control of Nonlinear PDEs using Data-efficient Reduced Order Models Based on the Koopman Operator

In this talk we present a data driven reduced order modeling approach for control of nonlinear PDEs which relies on the Koopman operator. We construct a bilinear surrogate model via linear interpolation between two Koopman operators corresponding to constant controls. Using

a recent convergence result for Extended Dynamic Mode Decomposition, convergence of the reduced order model based control problem towards the true optimum can be guaranteed if the control system depends linearly on the input. The resulting feedback controller is used to control the flow around a cylinder governed by the Navier-Stokes equations.

Sebastian Peitz
Paderborn University
speitz@math.uni-paderborn.de

Stefan Klus
Freie Universitaet Berlin
stefan.klus@fu-berlin.de

MS227

Identifying and Modeling Coherent Structures in Turbulent Flows using Spectral Proper Orthogonal Decomposition and Resolvent Analysis

The chaotic nature of turbulent flows poses a formidable challenge for data-driven methods meant to identify coherent, low-rank flow dynamics. Whereas each trajectory of a turbulent flow evolves differently, popular data-driven methods typically either utilize only a single trajectory (e.g., dynamic mode decomposition, DMD) or altogether ignore the dynamic nature of the flow in favor of a purely statistical approach (e.g., proper orthogonal decomposition, POD). In this talk, I will discuss an alternative method called spectral proper orthogonal decomposition (SPOD). While SPOD goes back to the early work of Lumley (Stochastic Tools in Turbulence, 1970), I will present new results demonstrating that SPOD combines the advantages of POD and DMD for statistically stationary turbulent flows – each SPOD mode represents a structure that is dynamic in the same sense as DMD modes, but that also accounts for and optimally describes the statistical variability of the flow. Lastly, I will show that SPOD is closely related to an equation-based modeling technique known as resolvent analysis. This connection enables data-free prediction of the turbulent flow structures educed by SPOD and provides additional insight into their physics.

Aaron Towne
Stanford University
towne@umich.edu

MS228

A Parallel Multigrid Reduction Framework for Coupled Multiphase Poromechanics

Algebraic multigrid (AMG) is a powerful tool for solving large sparse linear systems that result from discretizations of partial differential equations (PDEs). Using AMG, one can take advantage of the efficiency and scalability of multigrid methods without the need of constructing an explicit hierarchy of grids as in geometric multigrid. However, similar to other multigrid methods, its applicability is typically limited to scalar elliptic PDEs, despite ongoing efforts to adapt it to general non-symmetric matrices. Today, simulation requires solving coupled multi-physics problems, in which different processes are modeled by PDEs with different characteristics (i.e. elliptic, parabolic, and hyperbolic). Thus, there is a need for a robust and general method for solving the linear systems that arise in discretizations of these PDEs. We present our efforts in developing an algebraic framework based on multigrid reduction that is suited for systems of PDEs. In this framework, the decoupling between the equations is done algebraically through defining appropriate interpolation and restriction operators. One can then employ existing solvers for each of the decoupled blocks or design a new solver based on knowledge of the physics. We will demonstrate the applicability of our framework for the case of multiphase flow coupled with poromechanics. We show that the framework is flexible to accommodate a wide range of scenarios, as well as efficient and scalable for large problems.

Quan M. Bui, Daniel Osei-Kuffuor
Lawrence Livermore National Laboratory
bui9@llnl.gov, oseikuffuor1@llnl.gov

MS228

A Coupled Multi-physics Preconditioning Technique Based on Block Sparse Approximate Inverses

This talk presents a general approach for preconditioning the block Jacobian matrix arising from the linearization of coupled multi-physics discretized problems. The objective is to provide a fully algebraic framework, based on approximate inverses, which can be employed as a starting point for the development of specialized algorithms exploiting unique features of the specific problem at hand. The main idea is to separate the Jacobian solution in two phases. First the different physics are decoupled, than each of them is solved independently one from the other. To decouple two different diagonal blocks, we define a decoupling operator based on some extensions of the theory of block sparse approximate inverses. Both theory and practical consideration are deeply discussed. To solve each diagonal block, it is enough to use some of the many tools already available in the literature. The proposed approach is implemented for two multi-physics applications, namely the simulation of a coupled poromechanical system and the mechanics of fractured media. The numerical results obtained in experiments taken from real-world examples are used to analyze and discuss the properties of the preconditioner.

Andrea Franceschini
Stanford University
af1990@stanford.edu

Nicola Castelletto
Stanford University
Department of Energy Resources Engineering
castelletto1@llnl.gov

Carlo Janna, Massimiliano Ferronato
Dept. ICEA - University of Padova
carlo.janna@unipd.it, massimiliano.ferronato@unipd.it

Hamdi Tchelepi
Stanford University
Energy Resources Engineering Department

tchelepi@stanford.edu

MS228

Block Preconditioners for Incompressible Magnetohydrodynamics Problems

We consider preconditioning techniques for a mixed ?nite element discretization of an incompressible magnetohydrodynamics (MHD) problem. Upon discretization and linearization, a 4-by-4 nonsymmetric block-structured linear system needs to be (repeatedly) solved. One of the principal challenges is the presence of a skew-symmetric term that couples the ?uid velocity with the magnetic ?eld. We propose a preconditioner that exploits the block structure of the underlying linear system, utilizing and combining e?ective solvers for the mixed Maxwell and the Navier-Stokes subproblems. We perform a spectral analysis for an ideal version of the preconditioner, and develop and test a practical version of it. Large-scale numerical results in two and three dimensions validate the e?ectiveness of our approach.

Michael P. Wathen
The University of British Columbia
mwathen@cs.ubc.ca

Chen Greif
Department of Computer Science
The University of British Columbia
greif@cs.ubc.ca

Dominik Schoetzau
Mathematics Department
University of British Columbia
schoetzau@math.ubc.ca

MS228

Multiphysics Simulations with PFLOTRAN and PETSc on Novel Manycore Computer Architectures

As the high-performance computing community pushes towards the exascale horizon, power and heat considerations are driving the increasing importance and prevalence of fine-grained parallelism in new computer architectures. High-performance computing centers have become increasingly reliant on GPGPU accelerators and 'manycore' processors such as the Intel Xeon Phi, and 512-bit SIMD registers have even been introduced in the latest generation of Intel's mainstream Xeon server processors. The high degree of fine-grained parallelism and more complicated memory hierarchy considerations of such 'manycore' processors present several challenges to existing scientific simulation software, and particularly for complex multiphysics codes that incorporate models of many different physical and chemical processes. Here, we consider how the massively parallel, open-source hydrologic flow and reactive transport code PFLOTRAN—and the underlying Portable, Extensible Toolkit for Scientific Computation (PETSc) library on which it is built—can best take advantage of such architectures. We will discuss some key features of these novel architectures and our code optimizations and algorithmic developments targeted at them, and present experiences drawn from working with a wide range of benchmark problems from PFLOTRAN as well as other PETSc-based multiphysics applications on these architectures.

Richard T. Mills
Argonne National Laboratory
rtmills@anl.gov

MS229

Industrial Strength Computing in a Python Universe

The Python programming language, created in 1991 and gaining in popularity ever since, offers an appealing means of programming high level constructs in a very succinct way. Many mathematical tools based on Python have matured during the last 15 years, including SciPy, NumPy, Matplotlib, Jupyter, Anaconda, and others. This richness of these tools, taken collectively, provides an appealing universe in which very large numbers of diverse mathematical calculations can be performed in an industrial setting. Particularly appealing in this paradigm is Pythons ability to interface quickly and effectively with large numbers of other tools, making it possible to deploy mathematical algorithms and applications wherever they are needed. This talk will explore some of the many use cases for this approach and discuss the knowledge and skills needed by practitioners to perform effectively in this environment.

Thomas A. Grandine
Applied Mathematics
The Boeing Company
thomas.a.grandine@boeing.com

MS229

High Performance Computing at BP

BP is one of the worlds largest oil and gas companies, with revenues over $240 M US annually and daily production over 3.6 M barrels. Technology has been critical to our success, and High Performance Computing is a critical enabler to our research efforts in seismic imaging, reservoir simulation, rock physics, and computational fluid dynamics. We will discuss our history and how our team structure has evolved. We will discuss the skills in our team, and how we manage workforce development. We will describe the challenges we must address to continue to satisfy BP's research goals. And we will describe the breakthroughs and Business impact of our research.

Keith Gray
BP
Keith.Gray@bp.com

MS229

Perspectives from a DOE Lab on CSE Education and Workforce Needs

The US Department of Energy National Labs are among the largest employers of PhDs in Computational Science & Engineering (CSE). This is evidenced by their high participation in SIAM's CSE conference series. There are a

variety of CSE jobs at the labs, and I'll describe some different career paths as well as how work is allocated. There are certain technical skills that can really differentiate a prospective job candidate, so I'll discuss what these are and how to get them while you're still in school. Most importantly, there are several characteristics that can be developed and make someone a good match for the a career in CSE and especially the labs - ability to work in interdisciplinary teams, strong oral and written communication skills, excitement for learning new areas, and technical flexibility.

Tamara G. Kolda
Sandia National Laboratories
tgkolda@sandia.gov

MS229

Past, Present and Future of CSE Education at ICES

The Institute for Computational Engineering and Sciences (ICES) at The University of Texas at Austin has long been recognized as a pioneer in CSE education. ICES has trained hundreds of graduate students in unique interdisciplinary programs that emphasize the integration of applied mathematics, high performance scientific computing, and applications drawn from diverse fields in engineering, science, geoscience, medicine and business. This talk will reflect on two decades of accomplishments under the leadership of J. Tinsley Oden, and then discuss ongoing and future plans to evolve and extend ICES' educational impact. This includes growing emphasis and connections with data science, computer science and computational medicine, as well as an increased emphasis on undergraduate education.

Karen E. Willcox
Massachusetts Institute of Technology
kwillcox@ices.utexas.edu

MS230

Adaptive Radial Basis Function Methods for PDEs

In our previous research, we developed adaptive radial basis function WENO/ENO methods for solving hyperbolic problems. The adaptivity was carried out by optimizing the shape parameter of the new basis. The optimization is based on the cancelation of the leading truncation error term and the formula was given by finite difference methods. In this talk, we present how to generalize the optimization technique to higher order for solving differential equations. Then we extend the idea to PDEs. Several numerical examples will be presented.

Jiaxi Gu
SUNY Buffalo
jiaxigu@buffalo.edu

Jae-Hun Jung
SUNY at Buffalo
jaehun@buffalo.edu

MS230

An Implicit Sparse Grid Discontinuous Galerkin Method for High Dimensional Reaction-diffusion Equations

In this talk, we will introduce a class of implicit sparse grid discontinuous Galerkin method based on multiwavelets on nested grids. In order to relax the CFL constraint due to the high order spatial derivatives, krylov implicit integration factor method is used as temporal solver. Numerical examples in 2D and 3D cases are provided to show the performance of our proposed method.

Yuan Liu
Mississippi State University
yliu@math.msstate.edu

Yingda Cheng
Michigan State University
ycheng@msu.edu

Shanqin Chen
Department of Mathematical Sciences
Indiana University South Bend
chen39@iusb.edu

Yong-Tao Zhang
Department of Applied and Computational Math and Statistics
University of Notre Dame
yzhang10@nd.edu

MS230

Accessing Superconvergence for Accurate Multi-resolution Analysis of Discontinuous Galerkin Approximations

Passing numerical information from fine scales to coarse scales and vice-versa is a key ingredient in multi-scale/multi-resolution information. It is desirable to isolate the numerical errors in passing information between scales. One way to do this is to exploit the multi-scale structure of the underlying numerical approximation while taking advantage of the hidden accuracy of the approximation. This talk will focus on exploiting superconvergence to obtain more accurate multi-resolution analysis. Specifically, we concentrate on enhancing the quality of passing of information between scales by implementing the Smoothness-Increasing Accuracy-Conserving (SIAC) Filtering combined with multi-wavelets. This includes the ability to use a one-dimensional filter to extract information for multi-dimensional data. Although this talk presents the details of the SIAC filter using the standard discontinuous Galerkin method, these techniques are easily extendable to other types of data.

Jennifer K. Ryan
University of East Anglia
Mathematics Institute, Heinrich-Heine University

jennifer.ryan@uea.ac.uk

MS230

Symmetric Direct Discontinuous Galerkin Method for Elliptic Interface Problems

We first review recent development of direct discontinuous Galerkin (DDG) method and its variation, the symmetric DDG method on second order elliptic equations. We then discuss our studies of symmetric DDG method to linear variable coefficient elliptic interface problems with none zero solution jump and none zero flux jump interface conditions. Optimal order error estimate under energy norm is obtained with P^k polynomial approximations. A sequence of numerical examples are carried out to illustrate the optimal $(k+1)th$ order convergence. We focus on the case that mesh partition is aligned with the curved interface.

Jue Yan
Dept. of Mathematics
Iowa State University
jyan@iastate.edu

MS231

Preconditioning using Rank-structured Sparse Matrix Factorization

We present recent results with a sparse direct solver and preconditioner based on approximate LU factorization. The fill in the triangular factors is compressed using hierarchical matrices, a data-sparse format which uses matrix partitioning and low-rank decompositions. We compare Hierarchically Semi-Separable (HSS), Hierarchically Off-Diagonal Low-Rank (HODLR) and Block Low-rank (BLR) formats. Low-rank approximations are computed with adaptive randomized projection/sampling or with Adaptive Cross Approximation (ACA). The talk focuses on recent performance and scalability improvements. We evaluate usage of the SLATE dense linear algebra library, with OpenMP task scheduling, as a replacement for LAPACK and ScaLAPACK within the solver.

Pieter Ghysels
Lawrence Berkeley National Laboratory
Computational Research Division
pghysels@lbl.gov

MS231

A 2D Task-based Fan-both Factorization Algorithm for Sparse Symmetric Matrices

In this work, we present a new 2D data distribution of a sparse symmetric matrix, which balances the workload and the number of nonzero entries explicitly. We have incorporated low-overhead task-based Cholesky and LDLt algorithms in the symPACK solver using this new distribution, implemented using a custom runtime environment based on the UPC++ communication framework. We analyze the performance achieved during the numerical factorization as well as that of the solution phase.

Mathias Jacquelin
Lawrence Berkeley National Lab
mjacquelin@lbl.gov

Esmond G. Ng
Lawrence Berkeley National Laboratory
egng@lbl.gov

MS231

Task-based Sparse Direct Solver for Symmetric Indefinite Systems

Many applications in science and engineering require the solution of large sparse linear systems of equations. For solving such problems, direct methods are frequently employed because of their robustness, accuracy and usability as black-box solvers.As modern architectures become more and more complex, with an increasing number of cores per chip, a deeper memory hierarchy and the integration of accelerators such as GPUs, it becomes all the more challenging to exploit the potential performance of such machines for sparse matrix factorization algorithms especially in the context of symmetric indefinite systems. Although significant efforts has gone into positive-definite systems, little progress has been reported in the much harder indefinite case. One major advance for tackling these problems is the design of the APTP (a posteriori threshold pivoting) strategy that has been implemented in the SSIDS solver and proven to be both efficient on multicore architectures compared to the state-of-the-art direct solvers. In this talk, we present the DAG-based solver SpLDLT that relies on a APTP strategy and uses the StarPU runtime system for implementing it parallel version. We show the benefits of our approach for exploiting heterogeneity in the the context of GPU-accelerated multicore systems.

Florent Lopez
Rutherford Appleton Laboratory
florent.lopez@enseeiht.fr

Iain Duff
Science & Technology Facilities Council, UK
and CERFACS, Toulouse, France
iain.duff@stfc.ac.uk

MS231

Slate: Multilevel Tasking in Dense Linear Algebra Libraries

The objective of the Software for Linear Algebra Targeting Exascale (SLATE) project is to provide fundamental dense linear algebra capabilities to the US Department of Energy and to the high-performance computing (HPC) community at large, and ultimately to replace the venerable Scalable Linear Algebra PACKage (ScaLAPACK). Dynamic task scheduling is a cornerstone of the SLATE design, and the OpenMP standard is SLATEs technology of choice in that regard. SLATE relies on the full range of OpenMP capabilities, including data-dependent tasking, nested tasking, and task priorities. SLATE dispatches work to both CPUs and GPUs in the form of tasks, and all SLATE communications are issued as tasks.

Asim YarKhan
Department of Electrical Engineering and Computer

Science
University of Tennessee, Knoxville
yarkhan@eecs.utk.edu

Mark Gates
Computer Science Department
University of Tennessee
mgates3@icl.utk.edu

Piotr Luszczek
University of Tennessee
luszczek@icl.utk.edu

Jakub Kurzak
University of Tennessee, Knoxville
kurzak@icl.utk.edu

Jack J. Dongarra
University of Tennessee, Oak Ridge National Laboratory, USA
dongarra@icl.utk.edu

MS232

Pass-efficient Matrix Algorithms for Lossy Data Compression

The future of high-performance computing, specifically the next generation of Exascale computers, will presumably see memory capacity and bandwidth fail to keep pace with data generation. Current strategies proposed to address this bottleneck include the omission of large fractions of data, as well as the incorporation of *in situ* compression algorithms to avoid overuse of memory. To ensure that post-processing operations are successful, this must be done in a way that ensures that a sufficiently accurate representation of the solution is stored. Moreover, in situations in which the input/output system becomes a bottleneck in analysis, visualization, etc., the number of passes made over the input data must be minimized. In the interest of addressing this problem, this work focuses on the application of pass-efficient matrix decompositions to high-dimensional simulation data from turbulent particle-laden flows, including a novel single-pass algorithm for computing interpolative decompositions. The algorithms are easily parallelized and are anticipated to be incorporated directly into PDE solvers, enabling effective *in situ* compression of large-scale simulation data.

Alec M. Dunton
University of Colorado
Department of Applied Mathematics
alec.dunton@colorado.edu

Lluis Jofre
Center for Turbulence Research, Stanford University
jofre@stanford.edu

Gianluca Iaccarino
Stanford University
Mechanical Engineering
jops@stanford.edu

Alireza Doostan
Department of Aerospace Engineering Sciences
University of Colorado, Boulder
Alireza.Doostan@Colorado.EDU

MS232

A Low-rank, Bi-fidelity Approximation for Linear Bayesian Inference

A significant challenge for the Bayesian inverse problem of systems with high-dimensional uncertainty is the calculation of the posterior covariance, as it requires a possibly dense matrix inversion. Recently there have been several works focused on the low-rank estimation of the posterior covariance, which results in a computational cost that scales linearly in the number of unknown parameters. However, in this case, the equivalent of many high-fidelity model solves are still required, and thus this method may incur a large computational cost. In this work we consider further cost reduction by introducing the use of a low-fidelity model to form a low-rank, bi-fidelity approximation to the posterior covariance. We will present the formulation of this method as well as numerical examples demonstrating the achieved cost reduction.

Hillary Fairbanks
Lawrence Livermore National Laboratory
fairbanks5@llnl.gov

Alireza Doostan
Department of Aerospace Engineering Sciences
University of Colorado, Boulder
Alireza.Doostan@Colorado.EDU

MS232

Data Reconstruction for Computational Fluid Dynamics using Deep Neural Networks

As modern computational efforts reach exascale, hardware and software faults will increasingly cause difficulty in completing simulations. These faults often lead to loss of data in portions of the domain and require significant memory consumption for frequent checkpoints and large resimulation efforts. The ability to reconstruct data to fill the gaps without resorting to checkpoints has the potential to increase simulation performance and resilience. We use convolutional neural networks (CNN) for spatial reconstruction of the flow solution for simulations with gappy data, referred to in the machine learning community as inpainting or infilling. In contrast to POD, we assume that the current gappy data is the only available data for inpainting. This assumption is relevant to large simulations where it is computationally expensive to reload data from the file system. This approach avoids eigenmode decompositions which may restrict the applicability or translation to new configurations. The method proposed is evaluated by performing inpainting for two canonical flows: laminar flow past a cylinder and homogeneous isotropic turbulence (HIT). These results are compared to inpainting with Gaussian Process regression (GPR). For HIT, the energy spectrum error is approximately four times lower for inpainting with CNN. Across the range of mask sizes, the normalized error in kinetic energy and enstrophy is below 10% for the CNN and around 20% for the GPR reconstruc-

tion.

Marc Henry de Frahan
National Renewable Energy Laboratory
marc.henrydefrahan@nrel.gov

Ryan King
NREL
ryan.king@nrel.gov

Ray W. Grout
National Renewable Energy Laboratory
ray.grout@nrel.gov

MS232

Dimensional Analysis and Reduced Order Models for Large Scale Simulations of Particle-laden Turbulence in a Radiation Environment

Data-driven modeling is emerging as an attractive tool in scientific computation, particularly for overcoming sampling expenses. However, arbitrary models lack physical interpretability, limiting their use for scientific insight. In this work, we augment modeling with physical interpretability by appealing to dimensional analysis: Through a simple transform, subspace dimension reduction can be equated to dimensionless numbers. This observation can be leveraged to seek meaningful dimensionless numbers, provide a priori bounds on intrinsic dimensionality, and identify missing information. We present several examples in the context of large-scale simulations of particle-laden turbulent flow in a radiation environment.

Zachary del Rosario
Stanford University
zdr@stanford.edu

Lluis Jofre, Gianluca Iaccarino
Stanford University
Mechanical Engineering
jofre@stanford.edu, jops@stanford.edu

MS233

A Posteriori Error Estimators of Recovery Type

In this talk, we will present our recent progresses on a posteriori error estimators of recovery type. These include the improved Zienkiewicz-Zhu estimator for higher order elements and the equilibrated estimator for discontinuous (nonconforming and DG) elements.

Zhiqiang Cai
Purdue University
Department of Mathematics
caiz@purdue.edu

MS233

On the Equations of Poroelasticity

The equations of poroelasticity can be used to model and predict the mechanical behavior of fluid-infiltrated porous media which is significant since many natural substances, for example, rocks, soils, clays, shales, biological tissues, and bones, as well as man-made materials, such as, foams, gels, concrete, water-solute drug carriers, and ceramics are all elastic porous materials (hence poroelastic). After giving an overview of some problems in poroelasticity and their mathematical analysis I will describe finite element based numerical methods for efficiently and accurately approximating solutions of model problems in poroelasticity, and the available a-priori error estimates.

Amnon J Meir
Southern Methodist University
ajmeir@mail.smu.edu

MS233

Adaptive FEM for Helmholtz Equation with High Wave Number

Adaptive linear FEM based on the a posteriori error estimator of residual type for the Helmholtz equation with high wave number is considered. Convergence and quasi-optimality of the adaptive algorithm are proved under the condition that k^3h^2 is sufficiently small. It is shown that the a posteriori error estimator actually characterizes the interpolation error rather than the error of the FE solution in the preasymptotic range. Numerical tests are proposed to verify the theoretical findings.

Haijun Wu
Department of Mathematics
Nanjing University
hjw@nju.edu.cn

MS233

Direct Sampling Methods for General Inverse Problems in Optimal Control Framework

In this talk we shall review the recent developments in direct sampling methods for both nonlinear wave-type and non-wave-type inverse problems. General motivations, principles and justifications are discussed for the choices of several key ingredients of direct sampling methods in an optimal control framework, and numerical experiments are presented for several representative inverse problems. This is a joint work with Yat Tin Chow (UC Riverside) and Kazufumi Ito (NCSU).

Jun Zou
The Chinese Univ. of Hong Kong
zou@math.cuhk.edu.hk

MS234

Ridge Approximations and Active Subspaces

Many multivariate functions that arise in computational science models – i.e., the maps from model input parameters to output observables – exhibit strong off-axis anisotropic dependence on the parameters. An active subspace is the span of important directions in the model input space, and a ridge function naturally exploits this anisotropy for approximation (e.g., response surface modeling). I will present recent work on discovering important directions and exploiting them (via dimension reduction) to reduce the cost of parameter studies in computational

science models.

Paul Constantine
University of Colorado Boulder
Department of Computer Science
pconstan@mines.edu

Jeffrey M. Hokanson
University of Colorado Boulder
jeffrey.hokanson@colorado.edu

Rachel Ward
Department of Mathematics
University of Texas at Austin
rward@math.utexas.edu

Armin Eftekhari
Alan Turing Institute
armin.eftekhari@gmail.com

MS234

Learning Large-scale Sparse Graphical Models: Theory, Algorithm, and Applications

Learning models from data has a significant impact on many disciplines, including computer vision, medical imaging, social networks and signal processing. In the network inference problem, one may model the relationships between the network components through an underlying inverse covariance matrix. The sparse inverse covariance estimation problem is commonly solved using an l1-regularized Gaussian maximum likelihood estimator, known as graphical lasso. Despite the popularity of graphical lasso, its computational cost becomes prohibitive for large data sets. In this talk, we will develop new notions of sign-consistent matrices and inverse-consistent matrices to obtain key properties of graphical lasso and prove that although the complexity of solving graphical lasso is high, the sparsity pattern of its solution has a simple formula if a sparse graphical model is sought. We will prove under mild assumptions that the graphical lasso estimator can be retrieved by soft-thresholding the sample covariance matrix and solving a maximum determinant matrix completion (MDMC) problem, and describe a Newton-CG algorithm to efficiently solve the MDMC problem. We will illustrate our results in different case studies.

Somayeh Sojoudi
University of California, Berkeley
Electrical Engineering and Computer Sciences
sojoudi@berkeley.edu

Salar Fattahi, Richard Zhang
University of Berkeley
fattahi@berkeley.edu, ryz@berkeley.edu

MS234

Learning Interaction Laws in Agent Dynamics from Observations

In different disciplines, inferring the laws of interaction of agents in complex dynamical systems from observational data is a fundamental challenge. We propose a non-parametric statistical learning approach to estimate the governing laws of distance-based interactions, with no reference or assumption about their analytical form, from data consisting trajectories of interacting agents. We demonstrate the effectiveness of our learning approach both by providing theoretical guarantees, and by testing the approach on a variety of prototypical systems in various disciplines.

Sui Tang, Fei Lu, Mauro Maggioni
Johns Hopkins University
stang@math.jhu.edu, feilu@math.jhu.edu, mauro@math.jhu.edu

MS234

Identifying and Elimination of Data Corruption

One of the issues in function approximation is constructing an accurate model when data are corrupted by unexpected errors. The unexpected corruption errors are different from the standard observation noise, in the sense that they can have much larger magnitude and in most cases are not probabilistic. By focusing on overdetermined case, we prove that corruption errors can be effectively eliminated by using several sparse promoting methods such as l0, l1,lp (p in (0,1)) minimization. Furthermore, the corrupted data can be identified and separated from the approximation errors by k-means clustering.

Dongbin Xiu, Jun Hou
Ohio State University
xiu.16@osu.edu, hou.345@osu.edu

MS235

Applications of Data Science and Analytics in the Defense Industry and Challenges

The field of data science, to include artificial intelligence and machine learning, holds tremendous promise for enhancing our national security and protecting the warfighter. However, there are many challenges that need to be overcome prior to applying these technologies within the defense industry. Even though data science, AI and machine learning provide promise within the military context, the field also presents unique challenges to this domain in terms of the nascency in which the ilities are currently addressed: predictability, reliability, vulnerability, maintainability, debug-ability, attackability, explainability. This talk will cover the potential and emerging application areas of data science and analytics within the defense industry, as well as open gaps and challenges. I will also discuss my own experiences leading data science and analytics projects, delivering data analytics products to customers, and discuss lessons learned.

Veronica Bloom
Northrop Grumman
veronica.bloom@ngc.com

MS235

Data Analytics is Imperative for Business Success

Why are Companies Struggling?

In their 2018 global markets outlook report, T. Rowe Price added a new section on technology innovation and disruption recognizing the increasing importance technology plays across all industries. They site big data analytics and AI among the key trends driving value and disruption. Companies leading the charge with these technological trends have had outsized success, today this group includes the five largest companies by market cap and the first trillion dollar companies. These powerful trends have enabled relatively new entrants to disrupt a myriad of well-established areas including finance, transport, retail, automotive, hospitality, entertainment and even space exploration. The companies that successfully leverage these trends are growing stronger and getting the bulk of the benefits from innovation, while the rest are struggling with the technologies and falling behind. To successfully incorporate analytics, companies need to overcome technical issues such as analysis depth, metrics source and definition, data and pipeline quality, integration and distribution channel filters. They must overcome cultural barriers at the leadership and working levels, and work through issues that include managing their talent and hiring university graduates from programs that are lacking. This talk discusses the success barriers at many otherwise well established and well run companies that are impacting their competitiveness, and proposes steps for moving forward.

Mario R. Garzia
Diga Group
mariogarzia@digagroup.com

MS235

The Technical and Organizational Challenges of Data Science

In October 2012 shortly after I began my career in the data science space the Harvard Business Review published the article Data Scientist: The Sexiest Job of the 21st Century and generated an enormous amount of buzz about the field. Since then, data science has matured: technical skill sets required to do the work are better defined and specializations are emerging. However, the field is still evolving and how data science is used by an organization can vary greatly. In such a dynamic and broadly defined field, it has been my experience that data scientists need to have a wide range of technical skills augmented by soft skills in order to be successful. In this talk, I will share my experience working as a predictive modeler, data scientist, and software developer across the insurance, energy, and financial industries. I will provide an overview of the mathematics used in each of my jobs, give examples of practical challenges I have faced in my work, and discuss the balance of hard and soft skills required in the field.

Catherine Micek
3M
cmicek@mmm.com

MS235

Being Rational: Weighted Adaptive Approximation of Scientific Data

Scientific datasets, generated from large scale simulations or experiments, are becoming increasingly challenging to transfer, analyze, and store. There is a need for methods to transform these datasets into compact representations that facilitate their downstream management and analysis, and ideally model the underlying scientific phenomena with defined numerical fidelity. In this work, we borrow ideas from computer-aided geometric design (CAGD) to generalize parametric curve and surface fitting to higher dimensions, volumes and hypervolumes, in order to approximate multivariate discrete scientific data. Specifically, we use nonuniform rational B-spline functions (NURBS) not only to compress input data points, but also to enable further analysis directly on the continuous fitted model, without the need for decompression. Traditionally, NURBS fitting in CAGD systems is formulated as a linear least squares system where output control point locations are solved for given a set of weights, knot locations, and domain parameterization. Additional NURBS refinement is mostly manual and interactive (user driven). We investigate several methods to adaptively determine optimal control point locations, weights, and knot vector that accurately approximate given input points to prespecified accuracy. We evaluate our methods, in terms of compression factors, achieved accuracy, and computational cost, on both synthetic smooth functionals and real-world scientific data.

Youssef Nashed
Mathematics and Computer Science Division
Argonne National Laboratory
ynashed@anl.gov

Tom Peterka, Vijay Mahadevan
Argonne National Laboratory
tpeterka@mcs.anl.gov, mahadevan@anl.gov

Iulian Grindeanu
Mathematics and Computer Science Division
Argonne National Laboratory
iulian@anl.gov

MS236

Finite Element Approximation of Fractional Minimal Surfaces

Fractional minimal surfaces, which can be interpreted as a non infinitesimal version of classical minimal surfaces, were introduced by Caffarelli, Roquejoffre and Savin in 2010. These surfaces are motivated by the convergence of a nonlocal threshold dynamics approximation governed by a jump lévy process to moving fronts. An interesting feature of these fractional problems is the emergence of stickiness phenomena: in general, nonlocal minimal surfaces develop discontinuities across the boundary of the domain. In this talk, we address the graph fractional minimal surface problem on bounded domains. This problem can be reinterpreted as a Dirichlet problem for a nonlocal, nonlinear, degenerate operator. We will present a finite element scheme that is able to capture the sticky behavior at the boundary of the domain and can be proved to be convergent without any regularity assumptions on the solution.

Numerical experiments further illustrate the method's capabilities.

Juan Pablo Borthagaray, Ricardo Nochetto, Wenbo Li
Department of Mathematics
University of Maryland
jpb@umd.edu, rhn@umd.edu, wenboli@math.umd.edu

MS236

Improving the Efficiency of FEM Discretizations for Nonlocal Problems: a New Concept of Nonlocal Neighborhoods

Fractional and nonlocal models have gained a lot of popularity in recent years due to their ability to describe anomalous diffusion phenomena where classical partial differential equations (PDEs) fail to provide an accurate description. Several modeling and computational challenges are associated with the solution of these equations: the computational cost required by the Finite Element (FE) solution of a simple fractional/nonlocal equation can be prohibitively expensive, especially in two- or three-dimensional domains. This is due to the fact that, contrary to the case of PDEs, points in a domain interact with a neighborhood of points. Standard interaction neighborhoods are Euclidean balls around a point; in the context of FE methods this clearly creates computational challenges in terms of assembling and accuracy of the matrix of the discretized system. We propose a new concept of neighborhood, including square neighborhoods, that makes the assembling process easier and faster. We show that the new nonlocal problem is still well-posed and we provide modeling error estimates and results on the convergence to local and fractional limits. Also, using two-dimensional domains, we illustrate the theoretical results and apply our approach to image processing applications.

Marta D'Elia
Sandia National Laboratories
mdelia@sandia.gov

Christian Vollmann
University of Trier
DFG-RTG Algorithmic Optimization
vollmann@uni-trier.de

Max Gunzburger
Florida State University
Department of Scientific Computing
mgunzburger@fsu.edu

Volker Schultz
Trier University
volker.schulz@uni-trier.de

MS236

Finite Element Approximation of An Obstacle Problem for a Class of Integro-differential Operators

We study the regularity of the solution to an obstacle problem for a class of integro-differential operators. The differential part is a second order elliptic operator, whereas the nonlocal part is given by the integral fractional Laplacian. The obtained smoothness is then used to design and analyze a finite element scheme. Joint work with Andrea Bonito (Texas A&M) and Wenyu Lei (SISSA)

Abner J. Salgado
Department of Mathematics
University of Tennessee
asalgad1@utk.edu

MS236

A Fractional Calculus Framework for the Prediction of Material Failure

We develop a fractional calculus and probabilistic framework to model and simulate the failure of materials. We formulate a damage and fatigue phase-field model, in which the damage phase field is a continuous dynamical variable, fatigue is treated as a continuous internal variable, and many parameters are arbitrary or not physically measurable, treated here as random variables with a probability distribution. Considering an isothermal isotropic linear elastic material with viscous dissipation under the hypothesis of small deformations, we employ the Monte Carlo and Probabilistic Collocation methods in the forward parameter uncertainty propagation and sensitivity analysis. The underlying parameters are closely related to terms in the free-energy potential that disregard nonlocal effects, which motivates the employment of the fractional-order operators to capture the intermittent and self-similar phenomena experimentally observed in fracturing materials.

Mohsen Zayernouri
College of Engineering
Michigan State University
zayern@msu.edu

MS238

Dimension Adaptive Sparse Quadrature and Sparse Polynomial Parametrized Transport Maps for High Dimensional Bayesian Integration

To tackle the problem of *efficient* posterior integration for computing moments of quantities of interest (QoIs) in high dimensions, we previously exploited the Laplace approximation of the posterior to inform *dimension-adaptive sparse quadrature*. While dimension-independent convergence occurs for sufficiently smooth integrands, the *accuracy* of this posterior surrogate is poor for highly non-Gaussian posteriors. We address this issue of accuracy by constructing a more flexible posterior surrogate via a *transport map* between the prior and the posterior. The pushforward of the prior through this map then defines the posterior surrogate. Computing posterior moments of a QoI via the surrogate is easy: given quadrature rules for the prior, particularly tensor product quadrature and Monte Carlo (MC) methods, we map quadrature points from the prior to the posterior. These quadratures are also used to compute a KL-divergence used when finding the optimal transport map. However, these quadrature rules are *inefficient*: MC converges like $O(\sqrt{N})$ and tensor product quadrature scales exponentially with respect to dimension. In this presentation, we describe an exten-

sion to these previous works: we introduce an *efficient* and *dimension-scalable* scheme for constructing *accurate* transport map posterior surrogates for posterior integration of QoIs based on dimension-adaptive sparse quadrature and sparse polynomial parameterizations.

Joshua Chen
UT Austin
ICES
joshua@ices.utexas.edu

Peng Chen
UT Austin
peng@ices.utexas.edu

Daniele Bigoni, Youssef M. Marzouk
Massachusetts Institute of Technology
dabi@mit.edu, ymarz@mit.edu

Omar Ghattas
The University of Texas at Austin
omar@ices.utexas.edu

MS238

Marginal MCMC with Noisy Local Approximations

We consider Bayesian statistical models where one is primarily interested in a low-dimensional marginal of the joint distribution of the parameters—for instance, hyperparameters in a Bayesian latent variable model. The desired posterior marginal density is often not available analytically, and instead must be estimated using Monte Carlo. In this setting, we develop a Markov chain Monte Carlo (MCMC) scheme that directly targets the marginal of interest using local approximations. Our scheme exploits regularity in the marginal density to construct a local regression approximation from noisy pointwise density evaluations, on the fly, during MCMC sampling. Continually refining the approximation leads to an asymptotically exact characterization of the desired marginal. The decay rate of the bias due to the regression model can be balanced with that of the variance in the MCMC estimate in order to guide an ideal refinement strategy, given a finite number of MCMC steps. Our approach significantly reduces computational expense relative to both standard and pseudo-marginal MCMC.

Andrew D. Davis
MIT
davisad@mit.edu

Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu

Natesh Pillai
Statistics
Harvard
pillai@fas.harvard.edu

Aaron Smith
University of Ottawa
asmith3@math.stanford.edu

MS238

Fast Sampling of Parameterised Gaussian Random Fields

Gaussian random fields are popular models for spatially varying uncertainties, arising for instance in geotechnical engineering, hydrology or image processing. A Gaussian random field is fully characterised by its mean function and covariance operator. In more complex models these can also be partially unknown. In this case we need to handle a family of Gaussian random fields indexed with hyperparameters. Sampling for a fixed configuration of hyperparameters is already very expensive due to the nonlocal nature of many classical covariance operators. Sampling from multiple configurations increases the total computational cost severely. In this report we employ parameterised Karhunen-Loève expansions for sampling. To reduce the cost we construct a reduced basis surrogate built from snapshots of Karhunen-Loève eigenvectors. In particular, we consider Matérn-type covariance operators with unknown correlation length and standard deviation. We suggest a linearisation of the covariance function and describe the associated online-offline decomposition. In numerical experiments we investigate the approximation error of the reduced eigenpairs. As an application we consider forward uncertainty propagation and Bayesian inversion with an elliptic partial differential equation where the logarithm of the diffusion coefficient is a parameterised Gaussian random field.

Jonas Latz
Technical University of Munich
jonas.latz@tum.de

Elisabeth Ullmann
TU München
elisabeth.ullmann@ma.tum.de

MS239

Composable Multiphysics: From Solvers to Quantities of Interest

We describe efforts to enable the composition of the many facets of multiphysics computation. This critically entails the reuse of developed components, such as modeling kernels, as well as modularization and extensibility of those components. Effective multiphysics computation also requires easy experimentation of solvers. Contemporary efforts require not only considerations of the forward solve, but also of the "outer loop", incorporating aspects such as parameter sensitivity, leveraging the solution of adjoint problems, as well as uncertainty in parameters and potentially the models themselves. We illustrate the deployment of these principles through the GRINS multiphysics package, built on the libMesh finite element library. We illustrate the modularization of the forward problem and the ability to easily reuse components to construct complex models at runtime using advanced numerical methods, including adaptive mesh refinement driven by quantities of interest. We describe the incorporation of the composable solver framework in the PETSc library in order to facilitate

effective and robust solvers, including geometric multigrid leveraging mesh hierarchies. We discuss the effective computation of parameter sensitivities and the path forward to composable inverse problems within a Bayesian framework. We illustrate with several examples from thermally coupled and chemically reacting flows as well as mixed-dimensional manifold structures.

Paul Bauman
Mechanical and Aerospace Engineering
University at Buffalo, State University of New York
pbauman@buffalo.edu

MS239

preCICE – a Comprehensive Coupling Library for Large-scale Surface-coupled Multi-physics Problems

preCICE (Precise Code Interaction Coupling Environment) is a coupling library for partitioned multi-physics simulations, including, but not restricted to fluid-structure interaction and conjugate heat transfer simulations. Partitioned means that preCICE couples existing programs (solvers) capable of simulating a subpart of the complete physics involved in a simulation. This allows for the high flexibility that is needed to keep a decent time-to-solution for complex multi-physics scenarios. The software offers methods for transient equation coupling, communication, and data mapping. Adapters for well-known commercial and open-source solvers, such as OpenFOAM, SU2, or CalculiX, are available. Adapters for in-house codes can be implemented and validated in only a few weeks. preCICE is an open-source software under the LGPL3 license and available on GitHub. We present the general concept and functionality of preCICE including consistent partitioned time-stepping, iterative interface solvers for implicit coupling, data mapping, and communication with a focus on new developments such as an improved communication channel initialization that will enable dynamically changing mesh connectivity at the coupling surface in the future, stable and efficient radial basis function interpolation for data mapping, asynchronous inter-solver communication, and constistent time-steppimg based in waveform relaxation. In addition, we show examples for adapter codes and real-world applications.

Miriam Mehl
University of Stuttgart, Germany
miriam.mehl@ipvs.uni-stuttgart.de

MS239

Multiscale Modeling of Bloodflow in the Human Vasculature

Over the last few decades, we have seen an emergence of high-resolution hemodynamic simulation packages enabling personalized blood flow models. In particular, there is an increasing emphasis on the ability to track the movement of individual cells in specific regions of the human vasculature. Building a detailed, realistic model of human blood flow, however, is a formidable mathematical and computational challenge. The models must incorporate the motion of fluid, intricate geometry of the blood vessels, continual pulse-driven changes in flow and pressure, and the behavior of suspended bodies such as red blood cells. In this talk, I will discuss the development of HARVEY, a parallel fluid dynamics application designed to model hemodynamics in patient-specific geometries. I will focus on the recent introduction of a scalable fluid-structure-interaction model and the techniques introduced to enable efficient scaling on over one million cores. This code relies on the lattice Boltzmann method to capture the underlying fluid flow and the immersed boundary method to couple the fluid with the cell model. Within this framework, we are able to investigate the role of different biophysical properties in influencing a cell's trajectory through the vasculature and explore the impact of inter-cellular interactions. The computational framework's parallel performance in this setting is evaluated and future development lessons will be discussed.

Amanda Randles
Duke University
amanda.randles@duke.edu

John Gounley
Oak Ridge National Laboratory
gounleyjp@ornl.gov

MS239

Scalable Solution of Implicit / Imex Fe Continuum Plasma Physics Models

Continuum models of plasma physics systems require the solution of the governing partial differential equations describing conservation of mass, momentum, and energy, along with various forms of approximations to Maxwell's equations. The resulting systems are nonsymmetric, strongly nonlinear, and exhibit a significant range of time- and length-scales. To enable accurate and stable approximation of these systems a range of spatial and temporal discretization methods are employed. For finite element methods these include variational multiscale methods and structure-preserving approaches. For time integration two well-structured approaches are fully-implicit and implicit-explicit type methods. The requirement to accommodate disparate spatial discretizations, and allow the flexible assignment of mechanisms as explicit or implicit operators, implies a wide variation in unknown coupling, ordering, and conditioning of the implicit sub-system. Our approach to overcome these challenges has been the development of robust, scalable, and efficient fully-coupled physics-based multilevel preconditioned Newton-Krylov type iterative solvers.To discuss the structure of these algorithms, and to demonstrate the flexibility of this approach various forms of magnetohydrodynamic and multifluid electromagnetic plasma models are considered. Results are presented on robustness, efficiency, and the parallel and algorithmic scaling of the methods with weak scaling results up to 1M cores.

John Shadid
Sandia National Laboratories
Albuquerque, NM
jnshadi@sandia.gov

Roger Pawlowski

Multiphysics Simulation Technologies Dept.
Sandia National Laboratories
rppawlo@sandia.gov

Edward G. Phillips, Paul Lin
Sandia National Laboratories
egphill@sandia.gov, ptlin@sandia.gov

Sidafa Conde
Mathematics Department
University of Massachusetts Dartmouth
sconde@sandia.gov

Ray S. Tuminaro
Sandia National Laboratories
Computational Mathematics and Algorithms
rstumin@sandia.gov

Sibusiso Mabuza
Sandia National Laboratories
smabuza@sandia.gov

Luis Chacon
Los Alamos National Laboratory
chacon@lanl.gov

MS240

Optimal Design of Deterministic Lateral Displacement Devices for Cell Sorting

We solve a design optimization problem for deterministic lateral displacement (DLD) device to efficiently sort same-size red blood cells by their deformability. Such optimal designs enable rapid medical diagnoses of several diseases such as malaria since infected cells are stiffer than their healthy counterparts. DLD device consists of pillar arrays in which pillar rows are tilted and hence are not orthogonal to the columns. This arrangement sorts cells laterally depending on their deformability. Pillar cross section, tilt angle of the pillar rows and center-to-center distances between pillars define a unique device. For a given pair of cells with different deformability we seek optimal DLD designs. We fix all the parameters except the pillar cross section which we parameterize with uniform 5th order B- splines. We propose an objective function to capture efficient cell sorting. The objective function is evaluated by simulating the cell flows through a DLD using our 2D model based on a boundary integral method. We solve the optimization problem using a stochastic, derivative-free algorithm. We present several scenarios where the optimal designs can sort cells with slightly different deformability. These designs have cross sections that have features similar to a triangle.

Gokberk Kabacaoglu
University of Texas
kabacaoglu.gokberk@gmail.com

George Biros
The Institute for Computational Engineering and Sciences
The University of Texas at Austin
gbiros@acm.org

MS240

Transport in Viscous Erosion

In groundwater flow, a complex porous media structure is defined by a suspension of eroding bodies. Erosion creates conduits or channels of preferred flow direction, resulting in geometries with anisotropic permeability. In addition to the bodies eroding, they also sediment and are transported by the flow. This talk will describe a boundary integral equation method for simulating suspensions of eroding and transporting bodies. This is joint work with Nick Moore.

Bryan D. Quaife
Scientific Computing
Florida State University
bquaife@fsu.edu

Matthew N. Moore
Florida State University
nickmoore83@gmail.com

MS240

Numerical Simulations of Two-phase Flow of Electrolyte Drops

We study the deformation and breakup of an ideally polarizable, axisymmetric electrolyte drop which is freely suspended in either an electrolyte or dielectric fluid and subjected to an imposed electric field. Two cases are considered: (1) The electric potential in the drop phase is small, so that its governing equation is approximated by a linearized Poisson-Boltzmann or modified Helmholtz equation (the Debye-Hŭckel regime). An accurate and efficient boundary integral method is developed to solve the low-Reynolds-number flow problem for a dielectric surrounding fluid, in the case of arbitrary Debye layer thickness. (2) The electric potential in the drop phase is $O(1)$. A hybrid or multiscale method is developed to describe the two-phase flow of electrolyte liquids. The method is based on an asymptotic reduction in the thin Debye layer limit. This is joint work with Michael Booty, Manman Ma, and Qiming Wang.

Michael Siegel
New Jersey Institute of Technology
misieg@njit.edu

MS240

Numerical Simulations of 3D Surfactant-covered Drop Electrohydrodynamics

We study the influence of an electric field on 3D surfactant-covered drops. Surfactants (surface-active-agents) are compounds that change the surface tension between liquids and are widely used in engineering applications, in pharmaceuticals, foods and petroleum industries, and their evolution is modeled by a surface pde: the convection-diffusion equation. The dynamic of the drops is governed by the Stokes equations, while the leaky dielectric model is used to study the influence of the external electric field. We present a numerical method based on a boundary integral

formulation, able to simulate drops with different viscosities and close interactions, adaptive in time and able to handle substantial drop deformations. A global representation of the variables using spherical harmonics expansions and a specialized quadrature method for the singular and nearly singular integrals that appear in the formulation are used to ensure the spectral accuracy of the method. The numerical method is validated against theoretical and experimental results and compared with the available numerical literature.

Chiara Sorgentone
KTH Royal Institute of Technology
sorgento@kth.se

Anna-Karin Tornberg
KTH
akto@kth.se

MS241

A Fredholm Operator Approach to Clamped Plate Problems

In the Kirchoff-Love theory, the vertical displacement of a homogeneous, thin clamped plate under pure bending satisfies the biharmonic equation in a two dimensional domain with Dirichlet boundary conditions (the value and normal derivative are specified along the boundary). Because the boundary conditions are of a different order (one is a derivative of the solution and the other is not), it is difficult to derive integral representations of the solution which result in a second kind integral equation for the unknowns. We show that existing integral representations for the related Stokes equation can be bootstrapped to provide a representation for this problem which is well behaved, even on domains with high curvature and on multiply connected domains. Certain eigenvalues and eigenmodes of the differential equation are also of interest, as they correspond to buckling modes of the elastic plate. We present a robust framework for computing these eigenvalues.

Travis Askham
University of Washington
askhamwhat@gmail.com

MS241

Asymptotic Analysis for the Close Evaluation Problem

When using the boundary integral equation method to solve boundary value problems for linear, elliptic partial differential equations, the evaluation of layer potentials near the boundary is challenging to compute because they are nearly singular integrals. To study this close evaluation problem, we determine the leading-order asymptotic behavior of the double-layer potential for Laplace's equation as the distance away from the boundary vanishes. By doing so, we obtain an asymptotic approximation given by the Dirichlet data at the boundary point nearest the interior evaluation point plus a nonlocal correction. We present numerical methods to compute this asymptotic approximation in two and three dimensions. The efficiency and accuracy of the asymptotic approximation is demonstrated through several examples. These examples show that the asymptotic approximation is useful because it is accurately approximates the close evaluation of the double-layer potential while requiring only modest computational resources.

Camille Carvalho
UC Merced
ccarvalho3@ucmerced.edu

Shilpa Khatri
School of Natural Sciences
University of California, Merced
skhatri3@ucmerced.edu

Arnold D. Kim
Applied Mathematics Unit
University of California, Merced
adkim@ucmerced.edu

MS241

Close Interactions of Particles in Stokes Flow

A Nystrom discretization of a second kind integral equation is used to compute the hydrodynamic interactions between rigid particles immersed in a viscous fluid. As particles get closer, their interactions become increasingly difficult to accurately resolve. This has two causes since close interactions both yield near singularities in the integrals in the equation, and introduce a sharp peak in the double layer density on each particle surface around any point close to contact. In this talk, we address both issues. For the first, we present error estimates and introduce a Quadrature by Expansion (QBX) method for axisymmetric rigid particles, using pre-computation for efficiency with very low storage requirements thanks to the axisymmetry. For the second part, we investigate the interplay between temporal and spatial resolution errors, once the quadrature errors have been made negligible by the use of the QBX method. This will also give us the means to evaluate the effects of different approximations introduced to overcome the above mentioned problems.

Anna-Karin Tornberg
KTH
akto@kth.se

MS241

An Efficient Method for Nearly Singular Line Integrals in Two and Three Dimensions

Layer potentials over one-dimensional integration domains are commonly observed in integral equation methods in two dimensions, where the integration typically is over the boundary of the computational domain. However, they also appear in slender body theory for three-dimensional Stokes flow, where flexible fibers suspended in a viscous fluid are modeled as one-dimensional objects. Independent of problem dimension, evaluation of the layer potential near the integration domain leads to a nearly singular integral, which must be evaluated using a specialized quadrature method. For boundary integrals in two dimensions, the panel-based, kernel-split interpolatory quadrature by Hels-

ing et al. has proven to be a highly efficient and accurate method for close evaluation of layer potentials. This talk will describe a new version of this interpolatory quadrature, which under certain conditions is much more robust with respect to geometry. In addition, it can be extended to line integrals in three dimensions. This allows us to formulate a close evaluation scheme also for the layer potentials of slender body theory, which has previously not been available.

Ludvig af Klinteberg
Royal Institute of Technology (KTH)
ludvigak@kth.se

MS242

Thermal Effects in Transport of Ionic Solutions

Abstract not available

Chun Liu
Department of Applied Mathematics, Illinois Tech
Chicago, IL 60616
cliu124@iit.edu

MS242

Compressible Multiphase Fluid Flows

Abstract not available

Qi Wang
University of South Carolina
Beijing Computational Science Research Center
qwang@math.sc.edu

MS242

Approximation of Contact Angle Hysteresis by using Onsager Principle

Contact angle hysteresis is a complicated wetting phenomenon induced by the inhomogeneity or roughness of the solid surface. There exist stick-slip motions of the three-phase contact line between the two-phase flow interface and the solid boundary. Both numerical simulations and analytical study are very challenging, especially to quantitatively compare with physical experiments. In this talk, we will show that the Onsager principle could be used as a powerful approximation tool to study this problem.

Xianmin Xu
Institute of Computational Mathematics
Chinese Academy of Sciences
xmxu@lsec.cc.ac.cn

MS242

Cell Motility Dependence on Adhesive Wetting

Abstract not available

Yanxiang Zhao
Department of Mathematics
George Washington University
yxzhao@email.gwu.edu

MS243

Model Reduction of Dynamical Systems on Nonlinear Manifolds using Deep Convolutional Autoencoders

Nearly all model-reduction techniques project the governing equations onto a linear subspace of the original state space. Such subspaces are typically computed using methods such as balanced truncation, rational interpolation, the reduced-basis method, and (balanced) POD. Unfortunately, restricting the state to evolve in a linear trial subspace imposes a fundamental limitation to the accuracy of the resulting reduced-order model (ROM). In particular, linear-subspace ROMs can be expected to produce low-dimensional models with high accuracy only if the problem admits a fast decaying Kolmogorov n-width (e.g., diffusion-dominated problems). Unfortunately, many problems of interest exhibit a slowly decaying Komolgorov n-width (e.g., advection-dominated problems). To address this, we propose a novel framework for projecting dynamical systems onto *nonlinear trial manifolds* using minimum-residual formulations at the time-continuous and time-discrete levels; the former leads to *manifold Galerkin projection*, while the latter leads to *manifold least-squares Petrov–Galerkin projection*. Next, we show that this manifold can be efficiently computed using convolutional autoencoders from deep learning. Finally, we demonstrate the ability of the method to significantly outperform even the ideal linear-subspace ROM on benchmark advection-dominated problems, thereby demonstrating the method's ability to overcome the intrinsic limitations of linear-subspace ROMs.

Kevin T. Carlberg
Sandia National Laboratories
ktcarlb@sandia.gov

Kookjin Lee
University of Maryland College Park
koolee@sandia.gov

MS243

Physically-constrained Data-driven Correction for Reduced Order Modeling of Fluid Flows

In this talk, we present two approaches for enforcing better conservation properties for reduced order models (ROMs) of fluid flows. In the first approach, to construct the centering trajectory, we use the Stokes extension instead of the standard snapshot average. We show that the Stokes extension yields significantly more accurate results. In the second approach, we enforce physical constraints in the data-driven modeling of the ROM closure term. The constrained data-driven ROM is significantly more accurate than its unconstrained counterpart.

Muhammad Mohebujjaman
Department of Mathematics
Virginia Tech
jaman@vt.edu

Leo Rebholz
Clemson University
Department of Mathematical Sciences
rebholz@clemson.edu

Traian Iliescu
Department of Mathematics
Virginia Tech
iliescu@vt.edu

Xuping Xie
Oak Ridge National Lab
xiex@ornl.gov

MS243

Reduced Stochastic Models for Complex Fluids through Computational Information Geometry

A numerical procedure for constructing a hierarchy of homogenized models of an underlying stochastic phenomenon is introduced based upon concepts of information theory. Similar to adaptive mesh refinement, several levels of description are simultaneously advanced in time. At the finest, microscopic level the stochastic phenomenon is described by a direct numerical simulator (e.g., Brownian dynamics for complex fluids or molecular dynamics). Coarse grained versions are obtained by dynamic estimation of probability distributions describing quantities of interest at each level. The novel aspect of the procedure is the use of information functionals to define the prolongation and restriction operations that ensure inter-grid communication. Applications from complex fluids and cell biology are presented.

Sorin Mitran
University of North Carolina Chapel Hill
mitran@unc.edu

MS243

Reduced-order Model for the BGK Equation based on Pod and Optimal Transport

The BGK equation describes the dynamics of a gas flow in both hydrodynamic and rarefied regimes:

$$\frac{\partial f}{\partial t}(x,\xi,t) + \xi \cdot \nabla_x f(x,\xi,t) = \frac{M_f(x,\xi,t) - f(x,\xi,t)}{\tau(x,t)} \quad (3)$$

where f is the density distribution function representing the density of gas particles at point x, velocity ξ and time t, and M_f is the Maxwellian distribution function. In the reduced-order model, the distribution functions are approximated by a small number of basis functions Φ_n computed offline:

$$\widetilde{f}(x,\xi,t) = \sum_{n=1}^{N} a_n^f(x,t)\Phi_n(\xi) \quad \text{and} \quad \widetilde{M_f}(x,\xi,t) = \sum_{n=1}^{N} a_n^M(x,t)\Phi_n(\xi) \quad (4)$$

In the offline phase, the BGK equation is sampled by an high-dimensional model. Then, optimal transportation provides additional snapshots of the distribution functions to complete this sampling. Finally, the basis functions are built by Proper Orthogonal Decomposition. During the online phase, the BGK equation is projected onto the basis functions, leading to an hyperbolic system of partial differential equations. Moreover, this projection is modified to conserve mass, momentum and energy of the gas. We investigate the prediction of shock waves in 1D and vortices in 2D. The results show the significant reduction of the computational cost and the accuracy of the reduced-order model.

Florian Bernard
INRIA
University of Bordeaux
florian.bernard@inria.fr

Angelo Iollo
Institut de Mathématiques de Bordeaux
angelo.iollo@math.u-bordeaux1.fr

Sebastien Riffaud
INRIA
University of Bordeaux
sebastien.riffaud@inria.fr

MS244

Approximability Models and Optimal System Identification

We revisit a worst-case system identification problem from an optimal-recovery perspective. The model assumption that unknown transfer functions, acquired through point evaluations, are dilations of bounded functions from Hardy spaces H_p gives way to their good approximability in H_p by polynomials of a given degree. Leveraging on recent results in data assimilation, we are then able to explicitly characterize the optimal performance and to construct optimal algorithms when $p = 2$ and $p = \infty$.

Simon Foucart
Texas A&M University
foucart@tamu.edu

MS244

An Iterative Method for Classification of Binary Data

Abstract not available

Denali Molitor
University of California, Los Angeles
dmolitor@ucla.edu

MS244

Large-scale Classification Techniques with Applications to Medical Data

Abstract not available

Deanna Needell
University of California Los Angeles
deanna@math.ucla.edu

MS244

Deep Learning for Quasi-optimal Polynomial Ap-

proximations of High-dimensional Functions

In this talk, we present and analyze rectified linear unit (ReLU) deep neural networks (DNNs) for constructing quasi-optimal polynomial approximations, applicable to a wide class of high-dimensional functions (including parameterized PDEs with both deterministic and stochastic inputs). We consider several cases of multidimensional functions, and our proofs reveal sharp asymptotic error estimates in which we achieve sub-exponential convergence rates with respect to the total complexity of the DNN. Computational evidence complements our theory and shows the advantage of our generalized methodologies compared to existing approaches and current published results.

Clayton G. Webster
University of Tennessee and Oak Ridge National Laboratory
webstercg@ornl.gov

MS245

Using Discriminative Graphical Models for Insurance Recommender Systems

Recommender systems have become extremely important to various types of industries where customer interaction and feedback is paramount for the success of the business. For companies that face changes that arise with ever-growing markets, providing product recommendations to new and existing customers is a challenge. Furthermore, it is important to have an algorithm which is descriptive, scalable, agnostic to missing features, and robust in providing these recommendations. Directed graphical models meet all these demands; however, if the dimensionality of the features is high, structure learning and inference can become computationally prohibitive. In this talk, we propose an algorithm with some novel aspects to learn the structure of a graphical model (e.g., Bayesian network), which considerably speeds up both training and inference runtimes with respect to standard Bayesian structure learning approaches, while achieving similar accuracy. We also show that this approach produces more accurate predictions than a state-of-the-art matrix factorization algorithm in the absence of complete evidence on several insurance-related datasets. We will also will present initial results obtained by extending this idea to a deep-learning-based framework.

Teja Kanchinadam
American Family Insurance
AMFAM tk@gmail.com

MS245

Dealing with Uncertainties in Investment Robo-advice

An investment Robo-Advice prototype is presented to demonstrate the value of uncertainties in machine learning prediction and inference. When onboarding new customers, required financial inputs are estimated using known personal demographic and socioeconomic information. A generative data completion engine is applied to predict unknown inputs from conditional distributions, and the associated uncertainties are used to govern Chatbot interventions for inputs validation. To simulate investment outcomes in complex market conditions, a Deep Bayesian Learning model is applied to predict success rates of investment goals, and prediction uncertainties are used to support casual inference and trade-off decision making towards goals.

Shi Yu
Vanguard
shi.yu@hotmail.com

MS245

Mapping Advanced Knowledge: the Lessons from Economic-climate and Macro-financial Modeling

We discuss the fundamental uncertainty resulting from reducing complexity of big data and big theory into a particular domain of knowledge within which the primary actors such as decision-making individuals or entities operate. Development of advanced tools and metrics allowing to distill and map knowledge between different domains is stressed. A new technique for uncertainty quantification in highly non-linear stochastic models using Hamilton-Jacobi-Bellman partial differential equations (HJB PDEs) to compute Chernoff entropy is presented.

Victor Zhorin
Independent Researcher
vzhorin@uchicago.edu

MS246

Deep Neural Networks for Data-driven Turbulence Models

Machine learning methods and in particular neural networks have generated significant enthusiasm in the last years. This interest reaches beyond the ML community itself into other fields of science and engineering. Since these methods can provide approximations to non-linear functions by learning from data without a-priori assumptions, they are particularly attractive for the generation of subspace models for multiscale problems. In this presentation, we present a novel data-based approach to turbulence modeling for Large Eddy Simulation by deep learning via neural networks. We first define the exact closure terms including the discretization operators and generate training data from direct numerical simulations of decaying homogeneous isotropic turbulence. We then present the design and training of artificial neural networks based on local convolution filters to predict the underlying non-linear mapping from the coarse grid quantities to the closure.
All investigated networks are able to generalize from the data and learn approximations with a cross correlation of up to 47% and even 73% for the inner elements, leading to the conclusion that the current training success is merely data-bound. We further show that selecting both the coarse grid primitive variables as well as the coarse grid LES operator as input features significantly improves training results. Finally, we show how to construct a stable and accurate LES model from the learned closure terms.

Andrea D. Beck
Institute of Aerodynamics and Gas Dynamics

University of Stuttgart
andrea.beck@iag.uni-stuttgart.de

David Flad
University of Stuttgart
flad@iag.uni-stuttgart.de

Claus-Dieter Munz
Institut für Aerodynamik und Gasdynamik (IAG)
munz@iag.uni-stuttgart.de

MS246

Convolutional Neural Networks for Direct Prediction of Aerodynamic Flowfields

This paper explores using Convolutional Neural Networks(CNNs) to predict steady flows around airfoils, given only the shape of the airfoil and the free stream flow characteristics, in the form of Angle of Attack(AoA) and Reynolds Number(Re). The network uses an Encoder-Decoder architecture to take as input a Signed Distance Function form of the airfoil shape. The Reynolds number and angle of attack are taken as numerical inputs after the encoding process.Using training data already given to us, we first investigate the effectiveness of CNNs in predicting flows along each of the aforementioned axes(airfoil shape, AoA, Re) independently, i.e uniaxially. We then expand our objective to obtain a network which generalizes well to new,unseen airfoil shapes and flow conditions, while calculating the force coefficients on the airfoils. The flows that the network is designed to predict are in the incompressible regime, and predictions of both flows and force coefficients are compared to solutions of the RANS equations for the same free stream flow characteristics, using the SA turbulence model for closure. We examine sources of error in the methodology, investigating the effect of image and output resolution on the relative accuracy of the solution compared to the RANS solver.

Saakaar Bhatnagar
University of Michigan
saakaarb@umich.edu

Shaowu Pan
University of Michigan, Ann Arbor
shawnpan@umich.edu

MS246

Deep Learning for Data-driven Modeling and Analysis of Fluid Flows

Deciphering dynamical models and spectral characteristics of fluid flows from limited data have gained renewed interest in recent times with the rise in popularity of Koopman operator-based frameworks. The Koopman operator represents a symmetric Markov linear time-invariant (LTI) approximation of the dynamics using snapshots of data and can be used for prediction and analysis. A key ingredient to the success of such data-driven modeling approximations is the appropriateness of the map that transforms the input flow state to the 'feature space where the linear Markovian representation is physically viable. In this work, we leverage deep learning algorithms to deduce transformation maps from training data and ultimately the LTI Koopman operator. In particular, we present (i) a deep Koopman network (DKN) that learns the observable and the transition operator simultaneously and (ii) a deep autoencoder network (DAN) that accomplishes the same sequentially. We compare the efficacy of these architectures to the class of Dynamic Mode Decomposition (DMD)-based Koopman approximation techniques in the context of both long-term predictions from limited data as well as their ability to accurately capture the spectral characteristics of canonical fluid flows. In addition, we present recent progress on dynamic online adaptation of such models to streaming data.

Balaji Jayaraman, Shivakanth Chary, Chen Lu
Oklahoma State University
balaji.jayaraman@okstate.edu, shivakanthchary.puligilla@okstate.edu, chen.lu@okstate.edu

MS246

Data-driven Filter Estimation for the Sub-grid Modelling of Kraichnan Turbulence

In this study, we demonstrate the use of artificial neural networks as optimal maps which are utilized for the convolution and deconvolution of coarse-grained fields to account for sub-grid scale turbulence effects. We demonstrate that an effective eddy-viscosity is characterized by our purely data-driven large eddy simulation framework without the explicit utilization of phenomenological arguments. In addition, our data-driven framework does not require the knowledge of true sub-grid stress information during the training phase due to its focus on estimating an effective filter and its inverse so that grid-resolved variables may be related to direct numerical simulation data statistically. Through this we seek to unite the structural and functional modeling strategies for modeling non-linear partial differential equations using reduced degrees of freedom. Both *a priori* and *a posteriori* results are shown for the Kraichnan turbulence case in addition to a detailed description of validation and testing. Our findings indicate that the proposed framework approximates a robust and stable sub-grid closure which compares favorably to the Smagorinsky and Leith hypotheses for capturing theoretical kinetic-energy scaling trends in the wavenumber domain.

Romit Maulik
Oklahoma State University
romit.maulik@okstate.edu

Omer San
Mechanical and Aerospace Engineering
Oklahoma State University
osan@okstate.edu

Adil Rasheed
SINTEF Digital
adil.rasheed@sintef.no

Prakash Vedula
University of Oklahoma

pvedula@ou.edu

MS247

Adaptive Multifidelity Contour Estimation for Reliability-based Design Optimization

This work develops a method to adaptively refine multifidelity surrogates for contour location in reliability-based design optimization (RBDO). RBDO can be computationally prohibitive due to numerous evaluations of the expensive high-fidelity models to estimate the reliability of the system in each optimization iteration. In this work, the high-fidelity model is replaced by a cheaper-to-evaluate adaptively refined surrogate for the reliability estimation. The past optimization iterations are reused as an information source and combined with the high-fidelity data to build a multifidelity surrogate. We develop one-step lookahead adaptive sampling algorithms for contour location to efficiently refine the multifidelity surrogate.

Anirban Chaudhuri
Massachusetts Institute of Technology
anirbanc@mit.edu

Alexandre Marques
Massacheusetts Institute of Technology
noll@mit.edu

Karen E. Willcox
Massachusetts Institute of Technology
kwillcox@ices.utexas.edu

MS247

Breaking the Curse of Dimensionality for PDE-constrained Optimization under High-dimensional Uncertainty

It is a key challenge to solve PDE-constrained optimization problems under high-dimensional uncertainty due to the curse of dimensionality by most deterministic approximation methods or the slow convergence by statistical (Monte Carlo) methods. In this talk, we present sparse polynomial (Taylor, Legendre, and Hermite) approximations for such problems, and show theoretical results on the dimension-independent convergence rates which depend only on the sparsity or regularity of the control objective with respect to the uncertain parameter, and not on the nominal parameter dimensions. We consider two cases for the uncertain parameter, infinite-dimensional lognormal random field and random field parametrized by uniformly distributed random variables.

Peng Chen
UT Austin
peng@ices.utexas.edu

Omar Ghattas
The University of Texas at Austin
omar@ices.utexas.edu

MS247

Robust Optimization of PDE-constrained Problems using Second-order Models

We present a framework for the robust optimization of nonlinear PDE-constrained problems with ellipsoidal uncertainty sets. The framework follows a deterministic worst-case approach, which leads to a difficult min–max formulation. By using second-order Taylor approximations of the involved objective and constraint functions with respect to the uncertain parameters, the inner maximization problems are reduced to trust-region subproblems, thereby combining efficiency with reasonable accuracy in many applications. We consider two alternatives for the resulting approximated robust counterpart: a nonsmooth formulation in terms of the maximum-value functions of the inner problems, and a smooth MPCC formulation. In the case of a highly nonlinear parameter dependence, we utilize a steering strategy for the expansion points of the Taylor approximations to increase the accuracy. High-dimensional uncertainties are addressed with matrix-free methods. We apply the method to various shape optimization problems and present numerical results.

Philip Kolvenbach
Technische Universitaet Darmstadt
kolvenbach@mathematik.tu-darmstadt.de

Stefan Ulbrich
Technische Universitaet Darmstadt
Fachbereich Mathematik
ulbrich@mathematik.tu-darmstadt.de

MS247

Sparse Optimal Control of PDEs under Uncertainty

We study optimal control problems under uncertainty with a sparsity-enhancing objective function. Sparsity of controls is achieved by incorporating the L^1-norm of the mean of the pointwise squared controls in the objective which leads to a shared sparsity structure of stochastic optimal controls. To solve the corresponding nonsmooth optimization problem, we propose an iterative reweighting algorithm. The method is based on a reformulation of the problem and iterates over a reweighting function, which is only defined over the physical space and thus avoids sampling of the random space. Combined with low-rank operator approximations, this results in a monotone first-order method. To accelerate the method, we introduce a reduced formulation which only depends on the reweighting function and derive a novel preconditioned Newton conjugate gradient method. The shared sparsity structure of the optimal controls and the performance of the algorithms are studied numerically using control problems governed by the Laplace and Helmholtz equations.

Chen Li
New York University
cl3898@nyu.edu

Georg Stadler
Courant Institute for Mathematical Sciences
New York University

stadler@cims.nyu.edu

MS248

Gluon: A Communication Optimizing Framework for Distributed Heterogeneous Graph Analytics

This talk introduces a new approach to building distributed-memory graph analytics systems that exploits heterogeneity in processor types (CPU and GPU), partitioning policies, and programming models. The key to this approach is Gluon, a communication-optimizing substrate. Gluon enables these programs to run on heterogeneous clusters and optimizes communication in a novel way by exploiting structural and temporal invariants of graph partitioning policies. To demonstrate Gluons ability to support different programming models, we interfaced Gluon with the Galois and Ligra shared-memory graph analytics systems to produce distributed-memory versions of these systems named D-Galois and D-Ligra, respectively. To demonstrate Gluons ability to support heterogeneous processors, we interfaced Gluon with IrGL, a state-of-the-art single-GPU system for graph analytics, to produce D-IrGL, the first multi-GPU distributed-memory graph analytics system. Our experiments were done on CPU clusters with up to 256 hosts and roughly 70,000 threads and on multi-GPU clusters with up to 64 GPUs. The communication optimizations in Gluon improve end-to-end application execution time by 2.6 on the average. D-Galois and D-IrGL scale well and are faster than Gemini, the state-of-the-art distributed CPU graph analytics system, by factors of 3.9 and 4.9, respectively, on the average.

Roshan Dathathri
University of Texas at Austin
roshan@cs.utexas.edu

Keshav Pingali
The University of Texas at Austin
pingali@cs.utexas.edu

MS248

A Model Driven Sparse Gemm on GPUs

General sparse matrix-matrix multiplication (SpGEMM) multiplies a sparse matrix A with a sparse matrix B and generates a resulting sparse matrix C. This is an essential building block in a number of applications such as algebraic multigrid methods, shortest path algorithms and Markov cluster algorithms. As a result, fast algorithms for parallel SpGEMM received much attention. The basic way to calculate SpGEMM is the row-by-row method proposed by Gustavson that multiplies each row of A with the whole matrix B to obtain the corresponding row of C. So the SpGEMM computation becomes a combination of a number of sparse vector-matrix multiplications, i.e., the so-called sparse accumulators. Because the rows are independent of each other, they can be easily parallelized on modern many-core processors. In this paper, we propose three novel register-aware SpGEMM algorithms for three representative sparse accumulators, i.e., sort, merge and hash, respectively. We fully utilise the GPU registers to fetch data, finish computations and store results out. In the experiments, our algorithms deliver excellent performance on a benchmark suite including 205 sparse matrices from the SuiteSparse Sparse Matrix Collection. Specifically, on Nvidia Pascal P100 GPU, our three register-aware sparse accumulators achieve on average 2.0x (up to 5.4x), 2.6x (up to 10.5x) and 1.7x (up to 5.2x) speedups over their original implementations in libraries bhSPARSE, RMerge and NSPARSE, respectively.

Xin He
Institute of Computing Technology
Chinese Academy of Sciences
hexin2016@ict.ac.cn

Guangming Tan
Institute of Computing Technology,
Chinese Academy of Sciences
tgm@ict.ac.cn

Junmin Xiao
Chinese Academy of Sciences
xiaojunmin@ict.ac.cn

MS248

CVR: Efficient SpMV Vectorization on X86 Processors

Sparse Matrix-vector Multiplication (SpMV) is an important computation kernel widely used in HPC and data centers. The irregularity of SpMV is a well-known challenge that limits SpMV's parallelism with vectorization operations. Existing work achieves limited locality and vectorization efficiency with large preprocessing overheads. To address this issue, we present the Compressed Vectorization-oriented sparse Row (CVR), a novel SpMV representation targeting efficient vectorization. The CVR simultaneously processes multiple rows within the input matrix to increase cache efficiency and separates them into multiple SIMD lanes so as to take the advantage of vector processing units in modern processors. Our method is insensitive to the sparsity and irregularity of SpMV, and thus able to deal with various scale-free and HPC matrices. We implement and evaluate CVR on an Intel Knights Landing processor and compare it with five state-of-the-art approaches through using 58 scale-free and HPC sparse matrices. Experimental results show that CVR can achieve a speedup up to 1.70x (1.33x on average) and a speedup up to 1.57x (1.10x on average) over the best existing approaches for scale-free and HPC sparse matrices, respectively. Moreover, CVR typically incurs the lowest preprocessing overhead compared with state-of-the-art approaches.

Xu Liu
College of William & Mary
xl10@cs.wm.edu

Biwei Xie
Institute of Computing Technology
xiebiwei@ict.ac.cn

MS248

Enabling On-the-fly Storage Format Prediction and

Optimization for SpMV

Sparse matrix vector multiplication (SpMV) is an important kernel in many applications and is often the major performance bottleneck. The storage format of sparse matrices critically affects the performance of the SpMV. Our work exploits machine learning techniques to assist in selecting the best format. We first present how to effectively bridge the gap between deep learning and the special needs of the HPC problem through a set of techniques on matrix representations, deep learning structures, and cross-architecture model migrations. The proposed solution cuts format selection errors by two thirds, and improves SpMV performance by 1.73 on average over the state-of-the-art work. We further explore methods to minimize the runtime overhead of prediction and format conversion as ignoring them frequently makes the predictions sub-optimal or inferior. To address the overhead issue, we present a novel two-stage lazy-and-light method, and a set of regression models for learning the influence of the overhead and the benefit of the new format on the overall program performance. The proposed scheme helps control the risks in the format predictions and significantly improve the overall format conversion benefits.

Weijie Zhou, Xipeng Shen
North Carolina State University
wzhou9@ncsu.edu, xshen5@ncsu.edu

MS249

Eigenvector-based Centrality Measures in Multilayer Networks – Theory and Computation

Complex networks are used to model several different types of interactions among various entities. Being so versatile, they provide a useful tool to understand today's world and answer question about the underlying phenomenon or structure. Often in practice the question that needs answered is: which is the most important entity in the network? Several measures of importance, or centrality, have been introduced over the years. Among these, one of the most popular ones builds on the idea that the importance of a node stems from that of its neighbors. Mathematically this results in considering the Perron eigenvector of suitable graph matrices. In this talk, we describe how to generalize this simple but effective idea to the more complicated setting of multi-layer networks, where entities interact on different levels and the higher-order structure is captured by tensors rather than matrices.

Francesca Arrigo
University of Strathclyde
francesca.arrigo@strath.ac.uk

Francesco Tudisco
Dept. of Mathematics
University of Padua
tudisco.francesco@gmail.com

Antoine Gautier
Saarland University
ag@cs.uni-saarland.de

MS249

Tensor Canonical Correlation Analysis - A Probabilistic Method and Estimation Algorithm

Canonical correlation analysis (CCA) is a popular multivariate analysis technique that finds the linear relationship between two data sets. Recent technologies such as neuroimaging and remote sensing generate data in form of multi-dimensional arrays or tensors. Classic CCA is insufficient for dealing with tensor data due to the multi-dimensional structure and ultra high-dimensionality. We present tensor CCA, a technique that discovers linear relationship between two tensor data sets while respecting the spatial information. We delineate various population models and propose efficient and scalable estimation algorithms that have global convergence guarantees. Simulation studies illustrate the performance of our method.

Eric Chi
Department of Statistics
North Carolina State University
eric_chi@ncsu.edu

Hua Zhou
UCLA
huazhou@ucla.edu

Eun Jeong Min
Department of Biostatistics
University of Pennsylvania
mineunj@pennmedicine.upenn.edu

MS249

Stochastic Gradient Descent for Large-scale Generalized CP Tensor Decomposition

Tensor decomposition is a well-known tool for multiway data analysis, and the authors have recently proposed a generalized canonical polyadic (GCP) version of tensor decomposition that allows alternative loss functions such as logistic loss for binary data or Huber loss for robust estimation. The standard CP tensor decomposition, which uses squared error loss, has special structure so that the computational cost for fitting large-scale sparse tensors is proportional to the number of nonzeros in the tensor. In contrast, GCP lacks this structure and requires *dense* computations for sparse inputs. To address this limitation, we propose to an approach based on stochastic gradient descent (SGD) for efficient calculation of GCP for both sparse and dense large-scale tensors. We provide an efficient, pragmatic method for factoring very large-scale tensors using momentum-based SGD, compare our method to non-stochastic approaches, and show its scalability to large-scale problems.

Tamara G. Kolda
Sandia National Laboratories
tgkolda@sandia.gov

David Hong
University of Michigan

dahong@umich.edu

Jed A. Duersch
University of California Berkeley
jaduers@sandia.gov

MS249

Stable Tensor Neural Networks

We propose a tensor neural network framework that offers an exciting new direction in machine learning. Using our design, we can store data in high-dimensional structures and extract multidimensional correlations otherwise latent in traditional algorithms. Our network architecture is based on the t-product (Kilmer and Martin, 2011), an algebraic formulation to multiply tensors via circulant convolution which inherits mimetic matrix properties. In this study, we demonstrate that our tensor neural network architecture is a natural high-dimensional extension of traditional neural networks. Then, we expand on recent work (Haber and Ruthotto, 2017) to analyze the stability of our tensor neural network. Haber and Ruthotto interpret deep neural networks as discretizations of nonlinear differential equations and introduce stable neural networks which promote superior generalization. Motivated by their framework, we examine the stability of tensor differential equations from which we introduce a stable tensor neural network architecture. We illustrate the advantages of stability and demonstrate the potential of tensor neural networks with numerical experiments.

Elizabeth Newman
Tufts University
e.newman@tufts.edu

Lior Horesh
IBM Research
lhoresh@us.ibm.com

Haim Avron
Tel Aviv University
haimav@post.tau.ac.il

Misha E. Kilmer
Mathematics Department
Tufts University
misha.kilmer@tufts.edu

MS250

Accelerating Therapeutics for Opportunities in Medicine

Accelerating Therapeutic opportunities in Medicine (ATOM) is a consortium, which aims to accelerate drug discovery using large scale computational simulations and modeling to guide the design process in order to reduce the number of costly experiments required validation. Efforts are underway to build a computational framework for predicting drug-like properties of molecules using a large collection of experimental pharmaceutical data supplemented with molecular simulations. Models are being designed and scaled to screen large collections of compounds using DOE high performance computers. A particularly challenging aspect of the project is to develop models that accurately predict on novel parts of the chemical space. The talk will focus on strategies for representing statistical model uncertainty and use of the models latent space to evaluate domain of applicability metrics to determine, which new molecules can be evaluated using existing models and where new experimental data is needed.

Jonathan Allen
Lawrence Livermore National Laboratory
allen99@llnl.gov

MS250

Opportunities for Polypharmacology Using Machine Learning

Drugs have polypharmacological properties; a given drug interacts with many different proteins and a given protein interacts with multiple drugs. This can be exploited to improve drug efficacy and prevent drug resistance. Many diseases have multiple genetic determinants, and individual genetic determinants may be involved in multiple diseases. Furthermore, protein function and expression are controlled by a regulatory network of other proteins. When targeted therapies work initially, patients often develop resistance due to secondary mutations or compensation from other parts of the underlying biological network. This illustrates the potential benefits of establishing computational polypharmacology methods discovering drugs that intentionally target multiple proteins for a beneficial therapeutic result. Conversely, many adverse drug reactions result from drugs interacting with non-therapeutic off-targets (unintended interactions). Animal studies during preclinical trial are not always a good indication of these adverse interactions in humans, and such adverse effects are generally not discovered until a drug has reached clinical trial or is already on the market. With the number of different proteins in humans and the genetic variations observable in the population, a full understanding of all possible interactions through experiments and clinical testing alone is not feasible, making computational investigations particularly useful and relevant.

Sally R. Ellingson
University of Kentucky
sally@kcr.uky.edu

MS250

Leveraging Large Scale Molecular Dynamics Simulations and Deep Learning for Binding Activity Models

Understanding a candidate drug molecules activity in the human body is critical for developing safer patient targeted therapies. To accomplish this task, increased modeling accuracy of molecular interactions between a broader array of human proteins is needed. Drug-protein binding is a complex dynamic process and there is a limited number of high quality public experimental data available, making it difficult to build a comprehensive collection of activity models. Efforts at Lawrence Livermore National Laboratory are underway to develop a human protein small molecule interaction atlas using molecular dynamics simulations. How-

ever, simulations are too computationally costly to scale to evaluate the millions of molecules that will be needed. Improvements are being made to integrate simulation data with experimental data to build new structure based deep learning models that approach the accuracy of expensive simulations while scaling to evaluate more molecules across a larger range of protein targets.

William D. Jones
University of Colorado, Boulder
jones289@llnl.gov

MS250

Feature Selection in Biomolecular Models

Protein-protein interactions regulate many essential biological processes and play an important role in health and disease. The process of experimentally characterizing protein residues that contribute the most to protein-protein interaction affinity and specificity is laborious. Thus, developing models that accurately characterize hotspots at protein-protein interfaces provides important information about how to drug therapeutically relevant protein-protein interactions. In this work, we combined the KFC2a protein-protein interaction hotspot prediction features with Rosetta scoring function terms and interface filter metrics. A 2-way and 3-way forward selection strategy was employed to train support vector machine classifiers, as was a reverse feature elimination strategy. From these results, we identified subsets of KFC2a and Rosetta combined features that show improved performance over KFC2a features alone.

Julie Mitchell
Oak Ridge National Laboratory
mitchelljc@ornl.gov

MS251

The Rational Krylov Toolbox

The Rational Krylov Toolbox for MATLAB is a collection of scientific computing tools based on rational Krylov techniques. In this talk, we provide an overview of the toolbox, mainly focusing on the most recent updates. The core of the toolbox has been generalized to the block rational Arnoldi algorithm. Utility functions have been added to convert between barycentric, RKFUN, and Newton representations of rational interpolants. We will also discuss a connection between block rational Krylov spaces and nonlinear eigenvalue problems. The toolbox is available at http://rktoolbox.org/.

Steven Elsworth, Stefan Guettel
The University of Manchester
steven.elsworth@manchester.ac.uk, stefan.guettel@manchester.ac.uk

MS251

NEP-PACK: A Julia Package for Nonlinear Eigenproblems

We present an open-source library for nonlinear eigenvalue problems (NEPs), designed for scientists working on algorithm development, high-performance computing, as well as specific NEP-applications. The package is constructed to provide easy access to many state-of-the-art algorithms and benchmark problems. Problem transformations and modifications such as deflation and variable transformations are natively supported by the package. The software is implemented in Julia, and exploits the multiple dispatch system and parametric types for efficiency. The multiple dispatch allows us to naturally incorporate problem specific structures into the algorithms. The interface to access NEP-data is constructed with algorithm neutrality in mind such that performance and other algorithm properties can be compared and simulations can be reproduced. The interface consists of several equivalent computational quantities, which are related, and can be computed, by matrix function relations. This is joint work with E. Ringh, G. Mele, P. Upadhyaya, M. Bennedich.

Elias Jarlebring
KTH Stockholm
eliasj@kth.se

MS251

NLEVP: A Collection of Nonlinear Eigenvalue Problems

NLEVP is a collection of nonlinear eigenvalue problems presented in the form of a MATLAB toolbox. The collection serves both to illustrate the tremendous variety of applications of nonlinear eigenvalue problems and to provide representative problems for testing, tuning, and benchmarking of algorithms and codes. The aim of this talk is to present our new release of NLEVP. We will give a description of the new problems and discuss their challenges.

Francoise Tisseur, Gian Maria Negri Porzio
The University of Manchester
francoise.tisseur@manchester.ac.uk, gianmaria.negriporzio@manchester.ac.uk

MS251

CORK++: Compact Rational Krylov Methods for Nonlinear Eigenvalue Problems

We introduce CORK++, which is a fully templated high-performance C++ library for nonlinear eigenvalue problems. The compact rational Krylov (CORK) framework is a generic class of numerical methods for solving both small- and large-scale nonlinear eigenvalue problems. CORK methods rely on rational approximation and are characterized by a uniform and simple representation of structured linearization pencils. By fully exploiting the structure of these linearization pencils and by representing the Krylov subspace in a compact form, CORK++ allows for efficiently solving large-scale nonlinear eigenvalue problems with millions of degrees of freedom.

Roel Van Beeumen
Lawrence Berkeley National Laboratory

rvanbeeumen@lbl.gov

MS252

Algebraic Multigrid-based Approaches for Preconditioning Unassembled Operators in H(curl)

Preconditioning in a "matrix-free" setting where the matrix entries are not explicitly available is challenging even for standard diffusion problems, but the problems are made worse for problems in $H(\text{curl})$. A standard approach for H^1 problems is the low-order refined technique, in which a low-order mesh is formed on the nodal points of the high-order problem. We show here why this approach does not work in $H(\text{curl})$, because the large nullspace of the operator prevents the low-order space from being spectrally equivalent to the high-order space. Instead, we present a matrix-free variant of the auxiliary space Maxwell solver which operates by standard diffusion solves on auxiliary spaces coupled to the original problem by specially constructed interpolation operators. We show how the auxiliary solves and the connection operators can be performed in a fast, tensorized, matrix-free manor that is much faster than using fully assembled preconditioners.

Andrew T. Barker
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
barker29@llnl.gov

Tzanio Kolev
Lawrence Livermore National Laboratory
kolev1@llnl.gov

MS252

Approximate Fast Diagonalization for the Spectral Element Method

Fast diagonalization methods (FDMs) for the Poisson problem were introduced in 1964 by Lynch, Rice, and Thomas. Based on tensor-product decompositions, they are strictly applicable only to separable problems. However, they are also known to be effective preconditioners for nearly separable operators. Pazner and Persson (2018) have shown how to generate separable operators that are optimal approximations (in the Frobenius norm) to more general operators. A significant challenge arises, however, if the system is not well approximated by a rank-2 tensor. This is a common scenario in 3D, where even for the Poisson problem the discrete system matrix can take the form A = Cz x By x Ax + Bz x Ay x Cx + Az x Cy x Bx. The difficulty is that one cannot find an single set of eigenvectors to simultaneously diagonalize (A*,B*,C*) unless these matrices commute. We present a means to circumvent this requirement that has exactly the same work complexity as the commonly separable case and, to within a constant, the same storage complexity. We apply this fast preconditioner as a smoother in the context of p-multigrid for spectral element solutions of the steady-state advection and Navier-Stokes equations.

Pablo Brubeck, Paul Fischer
University of Illinois at Urbana-Champaign
pablo bbrubeck@protonmail.com, fischerp@illinois.edu

MS252

Efficient Solvers and Tensor-product Preconditioners for the Implicit Time Integration of Discontinuous Galerkin Methods

We present a matrix-free solver for the implicit time integration of very-high-order discontinuous Galerkin discretizations. This solver is based on the iterative solution of the large, block-sparse Jacobian matrices by means of the GMRES method. The linearized operator is applied in linear time per degree of freedom using a sum-factorization technique. A block Jacobi-type preconditioner is constructed by finding optimal (in Frobenius norm) approximations to the element-wise matrix blocks that are written as the sum of Kronecker products. These approximations can be formed in linear time per degree of freedom in two spatial dimensions by means of a matrix-free low-tensor-rank SVD procedure. The blocks are then inverted using a Schur form technique that is similar to classical matrix diagonalization. This preconditioner is applied to a wide range of test problems, including convection-diffusion, Euler equations, and Navier-Stokes. Asymptotic improvements to runtimes are observed.

Will Pazner
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
pazner1@llnl.gov

Per-Olof Persson
Dept. of Mathematics
University of California, Berkeley
persson@berkeley.edu

MS252

Acceleration of Tensor-product Operations for High-order Finite Element Methods

Abstract not available

Tim Warburton
Virginia Tech
tcew@vt.edu

MS253

Error Analysis of ZFP Compression for Floating-point Data

Compression of floating-point data will play an important role in high-performance computing as data bandwidth and storage become dominant costs. Lossy compression of floating-point data is powerful, but theoretical results are needed to bound its errors when used to store look-up tables, simulation results, or even the solution state during the computation. In this talk, we address the round-off error introduced by ZFP, a state-of-the-art lossy compression algorithm. The stopping criteria for ZFP depends on the compression mode specified by the user; either fixed rate, fixed accuracy, or fixed precision [Lindstrom, ZFP version 0.5.3]. While most of our discussion is focused on the fixed precision mode of ZFP, we establish a bound on the er-

ror introduced by all three compression modes. Following the presentation of these error bounds, numerical tests are provided to demonstrate the accuracy of the established bounds.

James D. Diffenderfer
University of Florida
jdiffen1@ufl.edu

MS253

Stability Analysis of Inline ZFP Compression for Floating-point Data in Iterative Methods

ZFP [Lindstrom and Isenburg, Fast and efficient compression of floating-point, IEEE Transactions on Visualization and Computer Graphics (2014)], a state-of-the-art lossy compression algorithm, can easily be used inline during numerical simulations due to the built in local structure of the algorithm. ZFP decomposes a solution state into 4d blocks, which are then compressed and decompressed independently. If the advancement operator of an iterative method has local structure, for a particular value ZFP can decompress only the data needed for the update, while the remaining data can remain in a compressed state. In a numerical simulation, the solution state already contains traditional errors, e.g., floating-point round-off, truncation error, and iteration error. The information that is lost during ZFP compression may represent the traditional errors, however, any additional error caused by ZFP will contaminate the current iterate. The goal of this work is to analyze the stability of using ZFP in fixed precision inline in time-evolving iterative methods.

Alyson Fox
Lawrence Livermore National Laboratory
fox33@llnl.gov

MS253

Error Analysis of ZFP Compression in Multigrid Methods

ZFP is a state-of-the-art lossy compression algorithm that can easily be used inline during numerical simulations due to the inherent locality of the algorithm. Using ZFP fixed-rate mode, which compresses a block to a fixed number of bits, ZFP can be implemented as a new data structure similar to a double precision array. Multigrid methods update the current iterate by approximating the error on coarse grids. Depending on the application, e.g. using multigrid as a preconditioner, the solution using ZFP compressed arrays can significantly reduce the storage cost while maintaining the required accuracy. The information that is lost during ZFP compression may represent the traditional errors however, any additional error caused by ZFP could contaminate the solution. It is important to understand if the error from ZFP compression overwhelms other sources of error. The goal of this talk is to analyze the stability of using ZFP in multigrid methods. We use previous error bounds to establish a bound for the error for fixed-rate ZFP when used in multigrid methods. Furthermore, we analyze the relative error for an adaptive rate multigrid algorithm.

Avary Kolasinski
The University of Kansas
avaryk@ku.edu

MS253

Beyond IEEE: Next-generation Floating-point Formats

Today, data movement is with few exceptions the primary performance barrier in high-performance computing. This problem is exacerbated by most applications' reliance on high-precision, bandwidth-consuming numerical types, as single precision either lacks the dynamic range or precision required by many computations. Born out of an era when computation rather than data movement was the dominating cost, the IEEE floating-point format provides a one-size-fits-all solution to balancing dynamic range, precision, and special values that often is wasteful and involves complex hardware logic to handle special cases like NaNs, signed zeros, and subnormals. We present a modular framework for novel number representations that maximize precision where it is most needed and that remove the redundancies and complexities associated with IEEE. Our framework allows new number representations to be built via composition of independent concepts. We show how a variety of representations, such as Gustafson's *Posits*, logarithmic number systems, and IEEE-like representations, can be expressed in our framework, as well as novel number systems that generalize the Elias universal codes for positive integers to the reals. Finally, we present empirical results that show how these new number representations increase—often by orders of magnitude—the accuracy of numerical computations, from basic arithmetic to linear algebra, PDE solvers, and physics mini-applications.

Peter Lindstrom, Scott Lloyd
Lawrence Livermore National Laboratory
pl@llnl.gov, lloyd23@llnl.gov

Jeffrey A. Hittinger
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
hittinger1@llnl.gov

MS254

The Derivation of Fluctuating Heat Conduction Models from Molecular Dynamics

Nowadays, there has been rapidly growing interest in modeling heat transport at the microscopic scale. Both theoretical and experimental results show that nano-systems exhibit different heat conduction properties from that in macro-systems. At microscopic scale, the observable quantities also carry substantial fluctuations. In this talk, heat conduction models are derived directly from the molecular dynamics (MD) setting. By selecting the local averaged energy as the coarse-grained (CG) variables, we apply different projection formalism to derive the reduced models that exhibit non-locality in space and time. In sharp contrast to conventional energy transport models, this derivation yields stochastic dynamics models and makes connections to nonlinear SPDEs. We discuss the approximation of the non-local term and ensure the correct statistics of the solution by adding appropriate additive and multiplicative

noise. As an example, we also compute the thermal conductivity using a CG description and provide a comparison with non-equilibrium MD simulations and Green-Kubo formula.

Weiqi Chu
PSU
wzc122@psu.edu

Xiantao Li
Pennsylvania State University
xxl12@psu.edu

MS254

Coarse-graining the Overdamped Langevin Equation for a Potential with Entropic Barriers via the Mori-Zwanzig Formalism

In this talk we discuss the use of the Mori-Zwanzig formalism to study SDE of overdamped Langevin type, and in particular, study a benchmark problem which may give some hints towards the relationship between memory and good choices of reaction coordinate.

Thomas Hudson
The University of Warwick, UK
t.hudson.1@warwick.ac.uk

MS254

Data-driven Modelling and Parameterization of Generalized Langevin Equation

Computational modelling of multiscale systems is centered around projecting the dynamics within the high-dimensional fully-resolved space onto a low-dimensional collective-variable space. In this talk, I will introduce a data-driven based model reduction framework for such systems. In particular, based on the Mori-Zwanzig formalism, the projected dynamics is casted into the generalized Langevin Equation (GLE). To numerically approximate the system, we further develop a variational inference algorithm to construct the free energy landscape, as well as a data-driven approach to construct the memory kernel of GLE. The constructed models naturally characterize the non-local correlation and fluctuations arising from smaller scale interactions. The method is demonstrated in challenging problems such as the non-equilibrium transition dynamics in physics, engineering and biological systems.

Huan Lei
Pacific Northwest National Laboratory
huan.lei@pnnl.gov

MS254

Coarse-graining Strategies for Multi-scale Partial Differential Equations using the Mori-Zwanzig Approach

The Variational Multiscale Method (VMS) is used as a formal starting point to separate resolved and unresolved scales in a continuous Galerkin setting. Following the VMS projection, the Mori-Zwanzig formalism is used to integrate out the small scale variables from the large scale equation. We examine the impact of different memory approximations on stabilization and closure and compare results to existing implicit and explicit Large Eddy Simulation models. Results for Burgers turbulence, 2-D and 3-D homogenous turbulence using the various models will be discussed.

Aniruddhe Pradhan
University of Michigan, Ann Arbor
anipra@umich.edu

Karthik Duraisamy
University of Michigan Ann Arbor
kdur@umich.edu

MS255

Problems in Underwater Acoustics

Recent advances in computational capabilities and sensor technologies present many new opportunities and mathematical challenges for underwater acoustics applications. This talk will provide the basic framework for capturing, processing and analyzing acoustic data in the ocean and outline key applications and corresponding contemporary challenges that remain. Special attention will be paid to the complex nature of the ocean environment and how applied mathematical will play a central role in moving these fields forward.

Kevin P. Bongiovanni
Raytheon BBN Technologies, USA
bongi@raytheon.com

MS255

Localization in a Random Medium Using Lucky Techniques

Ocean random medium effects cause distortions to signal wave fronts that are detrimental to detecting and localizing signals with sonar arrays. Motivated by lucky imaging, we show in real underwater acoustic data that even in environments where the signals are distorted on average at longer time scales, intermittent lucky scintillations or brief moments occur regularly during which time the signal wave fronts are relatively undistorted and can be exploited for detection and localization.

Ivars Kirsteins
NUWC
ivars.kirsteins@navy.mil

MS255

Shallow Water Inversion with Sequential Filtering and Linearization

We are investigating the challenges of inversion in shallow water focusing on multipath arrivals at short range. We combine a particle filter with a linearized model relating arrival times to experiment geometry and the environment and obtain estimates of unknown parameters efficiently and effectively.

Eliza (Z.-H.) Michalopoulou

Department of Mathematical Sciences
New Jersey Institute of Technology
michalop@njit.edu

MS255

Maximal Invariants and the Singular Value Decomposition in Detection Theory

In this paper we frame several representative detection problems in the multivariate normal model to explore when singular values of a data matrix or eigenvalues of a covariance matrix arise as maximal invariants under a transformation group that leaves the detection problem invariant. Generalized likelihood ratios are then functions of these maximal invariants.

Louis Scharf
Colorado State University
scharf@colostate.edu

MS256

Eigenvalue Problems in Mechanical Engineering

We will describe the source of the various eigenvalue problems that arise in Mechanical Engineering ranging from real symmetric ordinary eigenvalue problems to real unsymmetric second order eigenvalue problems. We will describe the algorithms in use to solve these problems including AMLS and Block Shift and Invert Lanczos. Finally we will provide an overview of our efforts to solve real symmetric generalized eigenvalue problems with hundreds of millions of rows on distributed memory computers with thousands of processes.

Roger Grimes
Livermore Software Technology Corporation
grimes@lstc.com

MS256

Substructuring Eigensolver Capabilities in Siemens PLM Software with Applications in Dynamic Analyses

The solution of eigenvalue problems is one of most important tasks in industrial applications such as analysis of Noise, Vibration, and Harness (NVH) of vehicles. For large-scale problems with millions of degrees of freedom, substructuring methods have merged as common approaches to reduce the hardware resources and computational run time. One such method developed by Siemens PLM is the RDMODES recursive substructuring method in SC Nastran. Using distributed memory parallel capability, the RDMODES method has achieved a factor of ten or more speedup in NVH examples. We will present three different applications to demonstrate the capabilities of the RDMODES method, namely NVH analysis, linear contact analysis and external superelement analysis. In those analyses, we discuss the possibility of parallel implementations, demonstrate the efficiency of the method, and address the difficulties and challenges we faced. We will also present a scalability study on a single node multi-core workstation and a multiple node cluster to show that the RDMODES method is suitable for a variety of hardware available today.

Ben-Shan Liao
SIEMENS PLM Software
ben-shan.liao@siemens.com

MS256

Nonlinear Eigenvalue Problems: Applications and Challenges from Mechanical Engineering

We review techniques that have recently been proposed for the solution of nonlinear eigenvalue problems. These rely on rational approximation of the nonlinearity and linearization of the resulting rational eigenvalue problem. A class of applications arises from mechanical engineering where the model contains nonlinear damping materials or porous materials. The finite element discretization leads to problems that are easily solved by the CORK method with AAA approximation. We also discuss open problems related to the boundary element method. The same type of methods can be used as fast solver for frequency response functions and for model order reduction. We explain why these methods are reliable for these applications.

Karl Meerbergen
K. U. Leuven
karl.meerbergen@cs.kuleuven.be

MS256

FISS: A Scalable Parallel Eigenvalue Solver for Large-scale Structural Analysis

We present a new distributed parallel eigensolution software FISS for large-scale real-symmetric generalized eigenvalue problems that arising in structural analysis. FISS is a new special implementation of the Sakurai-Sugiura method (SSM) for structural analysis and is designed to take advantage of next-generation many-node computing environments by fully exploiting the inherent coarse-grained parallelism of SSM. In order to assess the potential of FISS for state-of-the-art real industrial applications, we show a performance evaluation on up to 2048 nodes (139K cores) of Oakforest-PACS (KNL cluster) with a practical model of 60 millions DOF.

Yasunori Futamura
University of Tsukuba
futamura@cs.tsukuba.ac.jp

Tetsuya Sakurai
Department of Computer Science
University of Tsukuba
sakurai@cs.tsukuba.ac.jp

Sakujiro Hatazawa
MSC Software, Ltd.
sakujiro.hatazawa@mscsoftware.com

Tei Yamato
Mitsubishi Heavy Industries, Ltd.

tei_yamato@mhi.co.jp

MS257

The GANESO Software (Gas Network Simulation and Optimization)

This presentation discusses an industrial appliation concerning stationary and transient simulation, and optimization of gas transportation networks. GANESO is now an operational tool currently used by the Spanish companies REGANOSA TSO and REGANOSA SERVICES to simulate and optimize gas networks around the world. Concerning mathematical techniques, it implements efficient algorithms for solving mixed-integer programming problems that allow solving optimization problems for gas networks with more that 500 nodes in a few minutes. The transient model consists of the full Euler equations for incompressible flows of real gas mixtures, including friction and gravity effects. It also involves the computation of gas composition. This model is solved by using state of the art finite volume methods specifically developed for this kind of problems.

Julio González Diaz
University of Siantiago de Compostela
julio.gonzalez@usc.es

Alfredo Bermúdez de Castro
Universidad de Santiago de Compostela
Spain
mabermud@usc.es

MS257

The European Network Eu-maths-in and its Efforts for Digital Twins

The European service network for mathematics in industry and innovation was established in 2013 by a joint effort of the European Mathematical Society (EMS) and the European Consortium for Mathematics in Industry (ECMI). EU-MATHS-IN aims to leverage the impact of mathematics on innovations in key technologies by enhanced communication and information exchange between and among the involved stakeholders on a European level. It is establishing a dedicated one-stop shop to coordinate and facilitate the required exchanges in the field of application-driven mathematical research and its exploitation for innovations in industry, science and society. In the past 2 years, EU-MATHS-IN has created an industrial core team, organised workshops and together written vision documents to present MSO (mathematical modelling, simulation and optimisation) as a key enabling technology for industry in Europe. A mathematics consultation has been caarried out in Europe, policy makers have been visited, and right now actions are carried out to position mathematics within the new framework of EuroHPC, a joint undertaking established in September 2018. In the talk, more details will be given about these efforts at the European level, and how we envision achieving true digital twinning in a joint effort between mathematicians and industry.

Volker Mehrmann
Technische Universität Berlin
mehrmann@math.tu-berlin.de

Wil Schilders
Eindhoven University of Technology
w.h.a.schilders@TUE.nl

MS257

Towards Digital Twins: Merging Data and Model-based Approaches

The combinatorial explosion of product complexity driven by technical advances, novel business models, and ecosystems requires a paradigm shift in todays product development, production, and operation. Digital twins are one of the key technologies to master this complexity. Digital twins integrate all data, models, and other information of a physical asset generated along its life cycle that leverage business opportunities. This information is leveraged to predict and optimize performance. To this purpose simulation methods and / or data-based methods are used. The digital twin approach will make simulation widely available to everyone providing enhanced decision support improved throughout the entire life-cycle. Furthermore, it will closely link the two worlds of data-based approaches and simulations, to make prediction and optimizations based on all available information. This will require new computational concepts and corresponding algorithms. In this talk, we will discuss the concept of digital twins, provide several examples combining data based and model-based concepts and close with a discussion of open technical challenges from an industrial point of view.

Dirk Hartmann
Siemens AG
hartmann.dirk@siemens.com

Lucia Mirabella
Department of Mathematics and Computer Science
Emory University
lucia.mirabella@bme.gatech.edu

MS257

PDE-informed Surrogate Modeling for Fluid-mechanics Applications

A novel approach for constructing parametric solutions for PDEs is proposed. Multifidelity uncertainty quantification has gained much attention in recent years. Using high-fidelity solvers to perform UQ for high- dimensional problems is often infeasable. A multi-fidelity approach uses a combination of low and high- fidelity solvers to significantly decrease the computational cost of performing UQ. In our approach a neural network is used to generate a mapping between a low and high-fidelity solver, which can be used as a post- processing tool for enhancing low-fidelity solutions in the future. The neural network structure and training procedure in our approach results in a neural network that is capable of enhancing low-fidelity solutions outside the training set, which removes the need for a high-fidelity solver after training, as the enhanced low- fidelity solutions are perfectly suited for constructing surrogate models or extracting proper statistical moments. Our approach is demonstrated on a range of test-cases, ranging from the

linear advection equation to the non-linear Navier-Stokes equations.

Yous van Halder
CWI
y.van.halder@cwi.nl

MS258

Asynchronous Finite Element Simulation of Coastal Inundation

The numerical simulation of hurricane storm surge plays a key role in crafting effective responses to these extreme events. In this talk, we will present recent work on the development of a task-based parallelization of a discontinuous Galerkin finite-elements solver for the shallow water equations using HPX. One key aspect of these simulations is the modeling of inundation as a hurricane makes landfall. Due to computational load differences between dry and wet regions of the mesh, this flooding incurs dynamic load imbalances. In order to sustain machine utilization, we will present dynamic load balancing strategies and preliminary results utilizing HPX's active global address space (AGAS).

Maximilian Bremer
University of Texas at Austin
max@ices.utexas.edu

Hartmut Kaiser
Louisiana State University
hkaiser@cct.lsu.edu

Clint Dawson
Institute for Computational Engineering and Sciences
University of Texas at Austin
clint@ices.utexas.edu

MS258

OpenKIM: Reliable Interatomic Models for Multiscale Simulations

Multiscale and Atomistic simulations of materials using empirical interatomic potentials (also called "forcefields", or more generically "models") play a key role in realistic scientific and industrial applications. The Open Knowledgebase of Interatomic Models project (https://openkim.org) includes an automated user-extendable framework for testing the predictions of models for a host of material properties. Visualization tools have been developed to compare model predictions to help select the most appropriate one for a given application. Verification checks ensure the integrity of the models. Models in OpenKIM that conform to the KIM application programming interface (KIM API) can be seamlessly used with several major molecular simulation codes. Although KIM's objective is to have all models conform to the KIM API, this is not always immediately possible or practicable. So, OpenKIM also supports "Simulator Models" (SMs), i.e., models that are currently only available within a single molecular simulation code. SMs are treated just like any other model in OpenKIM so that results for verification checks, material properties predictions, and visualizations are available for SMs as well. This talk will describe the OpenKIM project and how the testing framework can assist materials researchers.

Ryan S. Elliott
University of Minnesota
relliott@umn.edu

MS258

Computing in the Kitchen Sink: Computational Stability of a Circular Hydraulic Jump

A water jet incident on a planar surface is known to generate a circular standing shock wave known as a hydraulic jump. Many people are familiar with this phenomenon in the everyday setting of the kitchen sink. We investigate the stability of a circular hydraulic jump in the shallow water model without friction, viscosity, or surface tension, using large-scale finite volume computations with PyClaw.

David I. Ketcheson
CEMSE Division
King Abdullah University of Science & Technology
david.ketcheson@kaust.edu.sa

MS258

Incorporating Computational Science and Engineering into Traditional STEM Education

Abstract not available

Claire Ralph
Cornell University
clairer@caltech.edu

MS259

Asynchronous Multigrid Methods

We define asynchronous multigrid methods where an approximate solution to a discretized PDE may be updated by multiple grids asynchronously. In particular, we investigate the following components: (1) the smoothing on each grid can be performed by multiple processors asynchronously, (2) smoothing is performed simultaneously on each grid like in additive multigrid, and (3) restriction and prolongation between the finest grid and coarser grids may be performed in parallel and asynchronously. The starting point for asynchronous multigrid methods are additive multigrid methods that can operate as solvers rather than only as preconditioners. We specifically start with the AFACx method for composite grids by McCormick, Quinlan, Lee, and Philip, but with full refinement, and the additive variant of the standard multiplicative multigrid algorithm by Vassilevski and Yang.

Edmond Chow, Jordi Wolfson-Pou
School of Computational Science and Engineering
Georgia Institute of Technology
echow@cc.gatech.edu, jwp3@gatech.edu

MS259

Communication-avoiding and Pipelined Krylov

Solvers in Trilinos

As part of DOE's Exascale Computing Project (ECP), we recently added communication-avoiding and pipelined Krylov solvers to the Trilinos software library. We will discuss how we implemented these solvers using Trilinos components, summarize lessons learned from implementations in other libraries, and present performance results on problems relating to ECP applications.

Mark Hoemmen
Sandia National Laboratories
mhoemme@sandia.gov

Ichitaro Yamazaki
University of Tennessee, Knoxville
iyamazak@icl.utk.edu

Erik G. Boman
Center for Computing Research
Sandia National Labs
egboman@sandia.gov

Jack J. Dongarra
University of Tennessee, Oak Ridge National Laboratory, USA
dongarra@icl.utk.edu

MS259

Recent Developments in the Symbolic Phase of Sparse Direct Methods

The direct solution of a sparse symmetric positive definite system of linear equations typically has two phases: symbolic and numeric. The symbolic phase is as important as the numeric phase, since the symbolic phase is where we permute the matrix to reduce fill in the numeric phase, identify the sparsity structure of the triangular factor, and set up the data structure for efficient numerical factorization. In the case of symmetric positive definite matrices, the symbolic phase depends solely on where the nonzero entries are in the original matrix, and hence does not involve numerical operations. In this talk, we will discuss some of the recent advances in the symbolic phase for a symmetric positive definite matrix, which have resulted in significant improvements in the efficiency of the numerical factorization.

Mathias Jacquelin
Lawrence Berkeley National Lab
mjacquelin@lbl.gov

Esmond G. Ng
Lawrence Berkeley National Laboratory
egng@lbl.gov

Barry Peyton
Dalton State College
bpeyton@daltonstate.edu

MS259

Low Synchronization GMRES Algorithms

Communication-avoiding and pipelined variants of Krylov solvers are critical for the scalability of linear system solvers on future exascale architectures. We present low synchronization variants of iterated classical (CGS) and modified Gram-Schmidt (MGS) algorithms that require one and two global reduction communication steps. Derivations of low synchronization iterated CGS algorithms are based on previous work by Ruhe. Our main contribution is to introduce a lower triangular solver compact WY form of MGS that results in a backward stable GMRES algorithm that requires only one global synchronization per iteration. The reduction operations are overlapped with computations and pipelined to optimize performance. Further improvements in performance are achieved by accelerating GMRES BLAS-2 operations on GPUs.

Stephen Thomas
Scientific Computing Division
National Center for Atmospheric Research
stephethomas@gmail.com

Katarzyna Swirydowicz
National Renewable Energy Lab
Golden CO
katarzyna.swirydowicz@nrel.gov

Julien Langou
University of Colorado Denver
julien.langou@ucdenver.edu

Shreyas Ananthan
National Renewable Energy Laboratory
shreyas.ananthan@nrel.gov

Yang Ulrike
Lawrence Livermore National Laboratory
yang11@llnl.gov

MS260

Inverting a Diverse Set of Spatially Sparse Diagnostic Measurements to Reconstruct a Complete 3D Picture of the Plasma Equilibrium in Fusion Devices

Due to the harsh environment and limitations of diagnostic port access in fusion reactors, it is not possible to measure every quantity at all locations. Using Bayesian principles and a model of the plasma equilibrium, the missing information including information not directly measured by diagnostic systems can be determined. V3FIT is a fully 3D equilibrium reconstruction code that models diagnostics synthetically using a variety of methods including Gaussian processes. By minimizing the difference between synthetic and experimental measurements, V3FIT determines the optimal parameters of the equilibrium model. Using a best unbiased linear estimate, V3FIT quantifies the uncertainty of the reconstructed result by propagating the experimental error from signal space to parameter space.

From this most probable model of the equilibrium, a complete picture of the plasma state is formed.

Mark Cianciosa
Oak Ridge National Laboratory
cianciosamr@ornl.gov

J.D Hanson
Auburn University
hansojd@auburn.edu

E.C. Howell
Tech-X Corporation
ehowell@txcorp.com

S.K. Seal, S.P. Hirshman
Oak Ridge National Laboratory
sealsk@ornl.gov, hirshmansp@ornl.gov

MS260

Using Large-eddy Simulations and Feature-basedData Assimilation to EstimateEnvironmental Variables of Simplified Cloud and Rain Models

Shallow clouds in the Earth system cover large regions of the oceans and are important for climate predictability. We study emergent behavior of marine stratocumulus cloud desks as they transition between weakly reflective open-cell, and highly reflective closed-cell structures. It was shown in previous work that these transitions exhibit predator prey type oscillations with the rain acting as a predator of clouds. Using this insight, a model of a single delay differential equation has been constructed. We are concerned with the question: how should one pick these parameters and how well are these parameters constrained by large-eddy simulations? We take a Bayesian approach: the prior distribution is based on a linear stability analysis and the likelihood is constructed by mapping model outputs to features derived from the LES. The feature we use is a typical cloud oscillation. Prior and likelihood distributions jointly define a posterior distribution over the model parameters. The posterior distribution describes our knowledge of the parameters given the model and the LES data. The result is a simplified representation of the cloud cycles in terms of a DDE with stochastic parameters. We find a physically relevant distribution of parameters and a good fit of the calibrated model to the features derived from the LES data. These results provide new insight into complex cloud and rain interactions and might be useful for representing these systems in climate models.

Spencer C. Lunderman
University of Arizona
Department of Mathematics
Lunderman@math.arizona.edu

MS260

Assimilating Data for Adaptive Multi-modal Sensing in Nuclear Security

Multi-modal sensing can be used to characterize large-scale experimentation, by measuring so-called "patterns of life' around the experiments. At the U.S. Department of Energy's Nevada National Security Site, we engage in dynamic materials experiments, and small, cheap multi-sensors can measure quantities like magnetic fields, ground accelerations, and acoustic signals to characterize the experiments and assess whether the human behaviors around the facilities uniquely determine what's going on scientifically. Compact, low-cost sensors platforms provide multiple data streams of relatively low-quality data, and fusing the different signatures into a single, coherent description of the science being performed requires computational methods not generally applied in this regime. In this work we will present results demonstrating emergent features in low-quality, multi-modal sensing – which is to say features that are evident only when the different modalities are analyzed in concert, as opposed to individually or serially – and show these emergent features are capable of uniquely identifying certain classes of experimentation.

Aaron B. Luttman, Derek Constantino
Nevada National Security Site
luttmaab@nv.doe.gov, constadj@nv.doe.gov

Christopher Stewart, Bethany Goldblum
University of California Berkeley
clstewart@berkeley.edu, bethany@nuc.berkeley.edu

Tori Hoff
Embry-Riddle Aeronautical University
hofft1@my.erau.edu

Margaret Hoeller
Illinois Institute of Technology
mhoeller@hawk.iit.edu

MS260

Multilevel Markov Chain Monte Carlo Methods for Seismic Source Inversion

We focus on implementing a multi-level sampling strategy using Markov Chain Monte Carlo (MCMC) methods to recover the probability distribution of the spatial location of the epicenter of an earthquake in a bounded domain. We introduce a hierarchy of levels $\ell = 0, \ldots, L$, for which there corresponds a discretization parameter $\{h_\ell\}_{\ell=0}^{L}$ of the underlying PDE modeling the seismic event, with an associated posterior distribution $\{\pi_\ell\}_{\ell=0}^{L}$. At each level we simultaneously generate two MCMC chains $\chi_{\ell,\ell}$ and $\chi_{\ell,\ell-1}$ of length N_ℓ, targeting π_ℓ and $\pi_{\ell-1}$. The crux of the algorithm lies on the construction of the proposal distribution that is being used to generate the chains. To do so, we propose to use an independent sampler strategy where the candidate state for $\chi_{\ell,\ell}$ and $\chi_{\ell,\ell-1}$ is drawn from an estimated density $\hat{\pi}_{\ell-1}$, built from the collected samples $\chi_{\ell-1,\ell-1}$ at the previous level. Experimental results suggest that the benefits provided by this algorithm are two-fold; the majority of the computations are performed at a coarse discretization level, at a cheaper cost, and additionally, experimental results have also shown that the samples generated at higher levels are (almost) independent and identically distributed. Lastly, we discuss a non-asymptotic bound for the total cost of the algorithm, measured in terms of the mean square error.

Juan Pablo Madrigal Cianci

University of New Mexico
juan.madrigalcianci@epfl.ch

Fabio Nobile
EPFL, Switzerland
fabio.nobile@epfl.ch

Raul F. Tempone
Mathematics, Computational Sciences & Engineering
King Abdullah University of Science and Technology
raul.tempone@kaust.edu.sa

MS261

Empowering Farmers with Affordable Digital Agriculture Solutions

Data-driven techniques help boost agricultural productivity by increasing yields, reducing losses and cutting down input costs. However, these techniques have seen sparse adoption owing to high costs of manual data collection and limited connectivity solutions. Our solution, called FarmBeats, an end-to-end IoT and AI platform for agriculture that enables seamless data collection from various sensors, cameras and drones. Our system design explicitly accounts for weather-related power and Internet outages, which has enabled six month long deployments in two US farms. In this talk, we will describe the FarmBeats system, and also outline the AI challenges we are currently addressing for outdoor as well as indoor agriculture.

Ranveer Chandra
Microsoft Research, Redmond, WA USA
ranveer@microsoft.com

MS261

Computational Advances and Challenges in Data-enabled Life Sciences: An Overview

The field of life sciences has by and large become a data-driven field, where data generated from a plethora of high-throughput technologies are driving scientific innovation and biological discovery. Biological databases are being inundated with data from a rich variety of omic technologies (e.g., genomic, proteomic, transcriptomic), offering insights into the molecular states (both static and dynamic) of individuals. In addition, a new wave of high-throughput sensing and imaging technologies is starting to be used to monitor an array of physiological and functional traits alongside environmental parameters, generating longitudinal datasets. In this talk, we will visit some of the key computational challenges and related advances in the analysis of complex life sciences datasets. More specifically, we will focus on the modeling and representation of such multi-dimensional biological data, and some recent mathematical and computational techniques that have emerged in the implementation of analytics of such representations. Applications of these emerging techniques to two different use-cases namely, plant sciences in agriculture and antibiotic stewardship in hospitals will be used as examples to illustrate the potential and opportunities offered by this rapidly evolving field.

Ananth Kalyanaraman
Boeing Centennial Chair, Associate Professor
School of EECS, Washington State University
ananth@wsu.edu

Bala Krishnamoorthy
Department of Mathematics
Washington State University-Vancouver
bkrishna@math.wsu.edu

Patrick Schnable
Iowa State University
schnable@iastate.edu

Eric Lofgren
Paul G. Allen School for Global Animal Health
Washington State University
eric.lofgren@wsu.edu

MS261

Persistent Homology in Metagenomics

Many problems in genomics are combinatorial in nature, as opposed to geometric, where subtle relationship between the objects are of importance. In metagenomics, reads from many different organisms are obtained en masse and one of the primary tasks is to identify the organisms (using a reference database). But these microorganisms could be very similar in sequence posing a challenge to the algorithms. I will talk about two abstract mapping of the problem to different filtered complexes and the use of persistent homologies to solve the metagenomic problem at hand.

Laxmi Parida
IBM T.J. Watson Research
parida@us.ibm.com

MS261

Computational Challenges and Opportunities in Crop Phenotyping

To overcome some of the myriad challenges facing sustainable crop production we are seeking to develop statistical models that will predict crop performance in diverse agronomic environments. Crop phenotypes such as yield and drought tolerance are controlled by genotype, environment (considered broadly) and their interaction (GxE). As a consequence of the next generation sequencing revolution genotyping data are now available for a wide diversity of accessions in each of the major crops. The necessary volumes of phenotypic data, however, remain limiting and our understanding of molecular basis of GxE is minimal. To address this limitation, we are collaborating with engineers to construct new sensors and robots to automatically collect large volumes of phenotypic data. New sensors and high-throughput, high-resolution, field-based phenotyping systems will be described. Some of these technologies will be introduced within the context of the Genomes to Fields Initiative.

Patrick Schnable
Iowa State University

schnable@iastate.edu

MS262

Hybrid Methods for Steady Fractional Advection Diffusion Reaction Equations

We apply a hybrid method for the fractional Poisson equation with fractional Laplacian on arbitrary bounded domains. Unlike integer-order equations, extra efforts are needed to resolve the boundary singularity, especially when we solve these problems on domains with complex boundaries. We use Monte Carlo methods for solutions at grid points near the boundary and finite difference methods for solutions at grid points away from the boundary. Numerical results will be presented.

Yuchen Dong, Zhongqiang Zhang
Worcester Polytechnic Institute
ydong@wpi.edu, zzhang7@wpi.edu

MS262

Machine Learning of Linear Fractional Differential Equations

Data-driven discovery of "hidden physics" has recently been approached by embedding the discovery problem into a Gaussian Process regression of spatial data, treating unknown equation parameters as hyperparameters of a modified "physics informed" Gaussian Process kernel. This kernel includes the parametrized differential operators applied to a prior covariance kernel. We extend this framework to linear space-fractional differential equations. Our methodology is compatible with a variety of fractional operators and stationary covariance kernels, including the Matern class. The fractional physics-informed GP kernel is given by d-dimensional Fourier integral formulas amenable to generalized Gauss-Laguerre quadrature. The method allows for discovering fractional-order PDEs for systems characterized by heavy tails or anomalous diffusion, which are of increasing prevalence in physical and financial domains. Second, a single fractional-order archetype allows for a derivative of arbitrary order to be learned, with the order itself being a parameter in the regression. Thus, the user is not required to assume a "dictionary" of derivatives of various orders, and directly controls the parsimony of the models being discovered. We illustrate on several examples, including fractional-order interpolation of advection-diffusion and modeling relative stock performance in the S&P 500 with alpha-stable motion via a fractional diffusion equation.

Mamikon Gulian, Maziar Raissi
Brown University
mamikon_gulian@brown.edu, maziar_raissi@brown.edu

Paris Perdikaris
University of Pennsylvania
pgp@seas.upenn.edu

George E. Karniadakis
Brown University
Division of Applied Mathematics
george_karniadakis@brown.edu

MS262

Lagrangian Solution of Fractional Advection-dispersion Equations

Various fractional derivative models have been proposed to capture complex transport processes observed in natural systems, and a consistent and simple numerical solver is needed to approximate these fractional-order PDEs. This presentation introduces the particle tracking based, fully Lagrangian solver that we developed to approximate the fractional-order advection-dispersion equations (fADEs), including the spatial fADE with variable coefficients, the time fADE with the index between 0 and 2, the multiscaling (vector) fADE with an arbitrary mixing measure, the scaling limit of the coupled continuous time random walk, the spatiotemporal fADE with discrete interface of transport properties, and small-scale reactive transport mixed with large-scale anomalous diffusion. Testing examples show that each of these models with the unified Lagrangian solver may find real-world applications for anomalous transport in geological media. Analysis further shows that 1) the multi-scale heterogeneity nature of the geological features motivates various types of anomalous transport that may be quantified by the fractional PDEs, and 2) many geological targets can be conceptualized as individual particles, including water packages, solute particles (in all three phases), minerals, sediment, soil, and rocks. Therefore, the grid-free particle tracking solver for the fractional-order PDEs may provide an ideal numerical tool to quantify complex dynamics in Earth Sciences.

Yong Zhang
University of Alabama
yzhang264@ua.edu

MS262

The Variable-order Fractional Differential Equation Driven by the White Noise: Mathematical and Numerical Analysis

We study the nonlinear variable-order fractional differential equation driven by the white noise. The well-posedness and the regularity of its solution is investigated.

Xiangcheng Zhen
University of South Carolina
xz3@math.sc.edu

MS263

Nonequilibrium Dynamics of Vesicles: Comparison with Experiments

Vesicles exposed to external fields and fluid flow show a wide variety of deformation mechanisms and dynamics, such as trembling, tumbling, tank-treading, and electric-field induced alignment. Many prior works assume that either the vesicle is nearly spherical in which case (semi-)analytic techniques can be used or start with an already extended vesicle. Here, a numerical model is developed which allows for an apparent area increase in the vesicles via a sub-optical/visible area exchange. Applying an exter-

nally driven shear for or electric field will allow for direct comparison between numerical and experimental results. Future work will explore the use of the numerical model to obtain material properties, such as bending rigidity and membrane capacitance.

Afsoun Rahnama Falavarjani, David Salac
University at Buffalo - SUNY
afsounra@buffalo.edu, davidsal@buffalo.edu

MS263

Torque-dipolar Micro-swimmers: Modeling, Circling Behavior and Large-scale Simulations

We present a new model for micro-swimmers that takes into account the counter-rotation of the body and flagella, as seen in motile bacteria or spermatozoa. The disturbance fluid flow of one such swimmer now contains a torque-dipole singularity in addition to the well-known force-dipolar singularity. We show that this head-and-flagella counter-rotation gives rise to clock-wise circling at no-slip walls just as observed in experiments of bacteria on surfaces. We discuss the scattering behavior of spermatozoa in a forest of cylindrical pillars, confirmed also by new experiments. Last, we show large scale and fast simulations of thousands of such swimmers that interact with each-other through direct collisions as well as through the fluid.

Enkeleida Lushi
Flatiron Institute
Simons Foundation
lushi@njit.edu

MS263

Rigid Particle Brownian Dynamics without Green's Functions

We introduce a method to perform large scale Brownian dynamics simulations on dense suspensions of rigid particles without any need to explicitly construct a Green's function or mobility operator. The Rigid-Body Fluctuating Immersed Boundary Method (RB-FIBM) we present incorporates physical domains with arbitrary boundary conditions and arbitrary shapes for the rigid particles. The RB-FIBM is used to investigate the wave patterns formed in a quasi–2D colloidal lattice which is hydrodynamically forced across a commensurate substrate potential in a slit channel.

Brennan Sprinkle
Northwestern University
brennan@nyu.edu

Amneet P.S. Bhalla
San Diego State University
asbhalla@sdsu.edu

Aleksandar Donev
Courant Institute of Mathematical Sciences
New York University
donev@courant.nyu.edu

Neelesh Patankar
Northwestern University
neeleshpatankar@gmail.com

MS263

Simulating Suspensions of Passive and Active Filaments

Microscopic flexible filaments and rigid fibres suspended in and interacting with fluids arise in many industrial processes and biological contexts. In suspension, they form complex fluids, where the microstructure affects the behaviour at the macroscale. Simulations of multiple filaments motion require resolving a complex low Reynolds number fluid-structure interaction problem that couples the motion and deformation of all filaments with one another through their hydrodynamic interactions. Torques from filaments bending or twisting give rise to numerical stiffness, and conditions such as inextensibility give rise to constraints that must always be satisfied. Computations of such suspensions are typically undertaken using one of many existing methods e.g. drag-based resistive force theory or slender-body theory each with a level of approximation. The hydrodynamic interaction computation is often quite developed and rapid, however the numerical treatment of the elastic aspects of the problem are less developed. We present a comprehensive, scalable methodology for simulating filament dynamics which employs Lagrange multipliers and implicit time integration to handle these elastic aspects. This framework seamlessly combines with many existing methods used to resolve the hydrodynamics. We will provide a clear description of the method and generalise it for fully 3D simulations that allow for both filament bending and twisting using quaternions.

Adam Townsend
Imperial College London
adam.townsend@imperial.ac.uk

Simon Scholler, Timothy Westwood, Eric Keaveny
Imperial College
simon.schoeller14@imperial.ac.uk,
t.westwood16@imperial.ac.uk, e.keaveny@imperial.ac.uk

MS264

High-performance Implementation of High-order Discontinuous Galerkin Methods

Matrix-free implementation of high-order finite element methods using sum factorization offers the possibility bypassing the memory bottleneck while at the same time reducing the number of operations substantially. In this talk we particularly address three issues arising with this approach. 1) How to achieve a substantial fraction of peak-performance of modern multi-core processors with wide SIMD instructions through a combination loop fusion and splitting techniques. 2) How to achieve performance portability across different PDEs and hardware platforms using code generation techniques. 3) How to effectively precondition elliptic problems with highly varying coefficients using a combination of matrix-free block smoother and low-order subspace correction. We will present numerical results of our implementation within the DUNE-PDELAB software

framework on scalar problems and systems of PDEs as well as on various hardware platforms including AVX512.

Peter Bastian
Interdisciplinary Center for Scientific Computing
University of Heidelberg
peter.bastian@iwr.uni-heidelberg.de

MS264

Efficient Discretizations for Exascale Applications

Efficient exploitation of exascale architectures requires rethinking of the numerical algorithms used in many large-scale applications. These architectures favor algorithms that expose ultra fine-grain parallelism and maximize the ratio of floating point operations to energy intensive data movement. One of the few viable approaches to achieve high efficiency in the area of PDE discretizations on unstructured grids is to use matrix-free/partially-assembled high-order finite element methods, since these methods can increase the accuracy and/or lower the computational time due to reduced data motion. In this talk we report on recent work in the Center for Efficient Exascale Discretizations (CEED, http://ceed.exascaleproject.org), a co-design center in the US Exascale Computing Project that is focused on the development of next-generation discretization software and algorithms to enable a wide range of finite element applications to run efficiently on future hardware. CEED is a research partnership involving 30+ computational scientists from two US national labs and five universities, including members of the Nek5000, MFEM, MAGMA, OCCA and PETSc projects. Topics of discussion will include recent progress in CEED packages and applications, new miniapps, benchmarks and API libraries developed in the project, and our efforts in scalable unstructured adaptive mesh refinement, matrix-free linear solvers and high-order data analysis and visualization.

Tzanio Kolev
Lawrence Livermore National Laboratory
kolev1@llnl.gov

MS264

High-performance Matrix-free Finite Element Programming with the DEAL.II Library

I will present the matrix-free operator evaluation capabilities for high order finite elements contained in the deal.II finite element library. Our implementations target the setting of fast quadrature for cell and face integrals that are applicable to general meshes and both linear and nonlinear partial differential equations. I demonstrate the algorithm selection for achieving high application performance on modern cache-based architectures, guided by performance analysis, and where the choices differ from GPU implementations, by the example of the CEED bakeoff problems. We have recently concentrated on two particular aspects of the operator evaluation, namely the memory transfer tradeoff between evaluating metric terms and variable coefficients on the fly versus pre-computation, and the memory transfer of loading the input and output vectors with the associated MPI ghost exchange. The limited memory bandwidth of CPUs makes this issue particularly pressing. The performance of operator evaluation has reached a performance level where many downstream solvers, such as multigrid smoothers, explicit time stepping, or conjugate gradient solvers, now spend more time on vector operations rather than the matrix-vector product. This situation necessitates more holistic software concepts, such as merging vector operations with the operator evaluation.

Martin Kronbichler
Technical University of Munich
Institute for Computational Mechanics
kronbichler@lnm.mw.tum.de

MS264

From the Variational Formulation to High Performance Kernels in NGsolve

In this talk we present recent development in our finite element package NGSolve. One aspect is a the Python-frontend to formulate mulit-physics problems in a native mathematical way. A second aspect is the definition of arbitrary and variable order finite elements for scalar as well as vectorial function spaces. We discuss our implementation based on C++11 lambda functions, which allows to seperate the mathematical definition, and high performance implementation using vectorization.

Joachim Schoeberl
TU Wien
joachim.schoeberl@tuwien.ac.at

MS265

High-order Methods for Second-order Stochastic Differential Equations with Rough Noises

We investigate the strong convergence order of piecewise linear finite element methods for a class of semilinear elliptic equations with additive spatial fractional noise. We present a spectral numerical models for the fractional Brownian motion (fBm) with hurst index $H \in (0, 1)$. Taking the suitable number of modes for spectral approximation of the fractional noise, we show that the strong convergence order of the finite element approximation is h^{H+1}, where h is the element size. Numerical results confirm our prediction for elliptic problems.

Wanrong Cao
Southeast University of China
wanrongcao@gmail.com

Zhongqiang Zhang
Worcester Polytechnic Institute
zzhang7@wpi.edu

MS265

Efficient Numerical Methods for the Generalized Langevin Equation

In this talk I discuss ergodic properties and numerical integration schemes for systems which can be described by a quasi-Markovian generalized Langevin equation. Traditionally, in thermal equilibrium, one assumes (i) the forces in the generalized Langevin equation (GLE) are given as

the gradient of a potential and (ii) a fluctuation-dissipation relation holds between stochastic and dissipative forces; these assumptions ensure that the system samples a prescribed invariant Gibbs-Boltzmann distribution for a specified target temperature. In this talk I will relax these assumptions, incorporating non-stationary noise and temperature parameters and allowing nonconservative force fields, for which the form of the stationary state is typically not known a priori. I will discuss theoretical issues such as the existence of stationary states and ergodic properties of such systems. I will also introduce a class of numerical integrators which allow efficient simulation of such systems. Under certain technical assumptions geometric ergodicity for the respective discrete dynamics is shown and an asymptotic error analysis of the discretisation error in ergodic averages is provided. The analysis is supported by numerical experiments.

Matthias Sachs
Duke University, Mathematics
msachs@math.duke.edu

Benedict Leimkuhler
University of Edinburgh
b.leimkuhler@ed.ac.uk

MS265

Minimum Action Method for Systems with Constant Time Delays

In this work, we develop a minimum action method to capture the most probable small-noise-induced transition path in dynamical systems subject to constant time delays. We minimize the Freidlin-Wentzell action functional, which will be discretized by finite elements. Adaptive meshes are used to deal with the ill-posedness of the problem and a penalty method is used to deal with the time delay. Numerical results will be presented.

Xiaoliang Wan
Louisiana State University
Department of Mathematics
xlwan@math.lsu.edu

MS265

Finite Element Methods for the Stochastic Cahn-Hilliard Equation with Gradient-type Noise

We consider the stochastic Cahn-Hilliard equation with gradient-type multiplicative noises which is motivated by the stochastically perturbed Hele-Shaw flow, and propose a fully discrete finite element method. We will discuss strong convergence with rates. Numerical experiments are provided to gauge the performance of the method and to study the interplay of the geometric evolution and gradient-type noises. This is joint work with Xiaobing Feng and Yukun Li.

Yi Zhang
University of North Carolina at Greensboro
y_zhang7@uncg.edu

MS266

Advancements in Spatial Resolution of D-Bar EIT Reconstructions for Human Thoracic Imaging

Electrical Impedance Tomography (EIT) is a technique which shows promise for use in many imaging applications. The 2D D-bar reconstruction method provides a direct solution to the underlying inverse problem, which is severely ill-posed and unstable in the presence of noise. Recent advancements in D-bar techniques have resulted in images with greatly improved spatial resolution over older methods. This talk will provide an overview of some of the latest advances in D-bar imaging, applied to human thoracic data.

Melody Alsaker
Gonzaga University
Department of Mathematics
alsaker@gonzaga.edu

MS266

Dynamic Inverse Problems in Medical Imaging

In the last years, the interest for imaging dynamic processes has increased a lot because in a lot of biomedical applications the target is non-stationary and the measurements time-dependent. However, the standard approach is to solve the reconstruction problem frame by frame using the standard static regularization approaches. Recently, some new approaches have been proposed which take the temporal evolution of the target into account. In this talk, we propose a quadratic spatio-temporal regularization which can be efficiently solved using a low rank approximation of the forward operator. We illustrate our approach in real data experiments for magnetic particle imaging (MPI). MPI is a new imaging modality which is capable of capturing fast dynamic processes in 3D volumes, based on the non-linear response of the magnetic particles to an applied magnetic field.

Christina Brandt
Universität Hamburg
Fachbereich Mathematik
christina.brandt@uni-hamburg.de

MS266

Limited Angle Tomography: Inpanting in Phase Space by Deep Learning

Limited angle geometry is still a rather challenging modality in computed tomography (CT). Compared to the standard filtered back-projection (FBP), regularization-based methods, combined with iterative schemes, help in removing artifacts but still cannot deliver satisfactory reconstructions. Based on the result that limited tomographic datasets reveal parts of the wavefront (WF) set in a stable way and artifacts from limited angle CT have some directional property, we propose a method that combines, in the phase space, the information coming from the visible part of the WF set and "inpaints" the invisible one by learning it with a convolutional neural network (CNN) architecture.

The WF set information is accessed by using the directional features of shearlets combined with a compressed sensing formulation, which is well suited to derive visible and invisible coefficients. Compared to other recently proposed deep learning strategies for (limited data) CT, our method provides a superior performance, an (heuristic) understanding of why the method works, providing a more reliable approach especially for medical applications.

Tatiana A. Bubba
University of Helsinki
Dept. of Mathematics and Statistics
tatiana.bubba@helsinki.fi

MS266

Learned Model-based Reconstructions for Accelerated Limited-view Photoacoustic Tomography

Recent advances in deep learning for tomographic reconstructions have shown great potential to create accurate and high quality images with a considerable speed-up. In this work we present a deep neural network that is specifically designed to provide high resolution 3D images from restricted photoacoustic measurements. The network is designed to represent an iterative scheme and incorporates gradient information of the data fit to compensate for limited view artefacts. Due to the high complexity of the photoacoustic forward operator, we separate training and computation of the gradient information. A suitable prior for the desired image structures is learned as part of the training. The resulting network is trained and tested on a set of segmented vessels from lung CT scans and then applied to in-vivo photoacoustic measurement data.

Andreas Hauptmann
University College London
Department of Computer Science
a.hauptmann@ucl.ac.uk

MS267

System-theoretic Model Order Reduction for Classes of Nonlinear Systems

System-theoretic model order reduction (MOR) methods consider dynamical systems with inputs and outputs. Inputs are often control functions, outputs model measurements or derived quantities. The advantage of system-theoretic methods over other MOR techniques like reduced basis methods, proper orthogonal/generalized decomposition techniques, and data-driven approaches is that they neither require training inputs, nor sampling or discretization of the input functions. The obtained reduced-order models are independent of the applied input, as long as feasible input functions are chosen. For linear time-invariant systems, the most common system class for modeling control systems, system-theoretic MOR methods are well-known and used in practice since the 1980ies. This includes balanced truncation and variations thereof, as well as rational interpolation techniques. For nonlinear systems, system-theoretic MOR methods have been developed during the last decades as well, often trying to generalize methods for linear systems. They often failed to be applicable to large-scale problems due to their overwhelming computational demands. Recently, there has been significant progress for certain classes of nonlinear systems, including in particular (quadratic-)bilinear and polynomial systems. Methods related to system balancing as well as rational interpolation can now be applied to large-scale problems. We will give a survey of these recent developments, and present some new results.

Peter Benner
Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany
benner@mpi-magdeburg.mpg.de

MS267

Interpolation-based $\mathcal{H}_2$ Optimal Approximation for Quadratic-bilinear Systems

Many nonlinear systems can be represented by *quasi*-linear parameter-varying (LPV) models, where the parameters depend on both external signals and system variables. In this talk, we focus on *quasi*-LPV models where the parameters only depend on system variables. Reducing such kind of systems are often difficult because of the state-dependent behavior of the parameters. Hence, instead of considering the model order reduction problem in the *quasi*-LPV setting, we return to the nonlinear system representation. By augmenting the state space, many such systems can be transformed into the quadratic-bilinear (QB) systems format. Based on the Volterra series analysis, the transfer function concept can be generalized to QB systems. We propose a definition of the moments according to the generalized transfer functions and a scheme to interpolate the 0^{th}-order moments. We also generalize the the $\mathcal{H}_2$ norm definition to QB systems. Minimizing the $\mathcal{H}_2$ norm difference between the full-order and the reduced-order models results in an interpolation variable selection strategy. The 1^{st}-order moments are interpolated automatically in the $\mathcal{H}_2$ optimal approximation setting. In such a way, the interpolation-based model order reduction problem is tackled by solving the corresponding generalized Sylvester equations iteratively. Numerical examples are tested to demonstrate the proposed method.

Xingang Cao
TU Eindhoven
x.cao.1@tue.nl

Joseph Maubach
Dept. Mathematics and Computer Science
Eindhoven University of Technology
j.m.l.maubach@tue.nl

Siep Weiland, Wil Schilders
Eindhoven University of Technology
s.weiland@tue.nl, w.h.a.schilders@TUE.nl

MS267

A New Perspective on ADI-based Balanced Trunction

We will consider model order reduction for stable linear time-invariant systems

$$\dot{x} = Ax + Bu, \quad y = Cx$$

with real, large and sparse system matrices. In particular, A is a square $n \times n$ matrix, B is rectangular $n \times m$, and C is $p \times n$. Among the many existing model order reduction methods our focus will be on (approximate) balanced truncation. The method makes use of the two Lyapunov equations

$$A\mathfrak{P} + \mathfrak{P}A^T = -BB^T,$$

and

$$A^T\mathfrak{Q} + \mathfrak{Q}A = -C^TC.$$

The solutions $\mathfrak{P}$ and $\mathfrak{Q}$ of these equations are called the controllability and observability Gramians, resp.. Balanced Truncation transforms the LTI system into a balanced form whose controllability and observability Gramians become diagonal and equal, together with a truncation of those states that are both difficult to reach and to observe. One way to solve these large-scale Lyapunov equations is via the Cholesky factor-alternating direction implicit (CFADI) method which provides a low rank approximation to the exact solution matrix $\mathfrak{P}$, $\mathfrak{Q}$ resp.. We will present and analyze a system of ODEs, whose solution for $t \to \infty$ is the Gramian $\mathfrak{P}$. We will observe that the solution evolves on a manifold and will characterize numerical methods whose approximate low-rank solution evolves on this manifold as well. This will allow us to give a new interpretation of the ADI method.

Heike Fassbender, Christian Bertram
TU Braunschweig
h.fassbender@tu-braunschweig.de, ch.bertram@tu-braunschweig.de

MS267

H-Infinity Balanced Truncation for Feedback Control of Flow Problems

The control, in particular the stabilization, of flows is of high interest in practical applications and a field of ongoing research. As recent theoretical and simulation results have shown, Riccati-based feedback is a suitable tool for that purpose at moderate Reynolds numbers. For the implementation as digital controllers, the resulting regulators are generally of too high order. Therefore, model order reduction techniques can be used for the construction of low-order controllers. A balancing-related method directly aiming at observer-based designs is the $\mathcal{H}_\infty$ balanced truncation. In contrast to the conventional LQG-based approach, the $\mathcal{H}_\infty$ balanced truncation provides more robust results that can compensate for uncertainties and approximation errors. In our work, we generalize the standard method from [D. Mustafa, D.Glover. Controller reduction by $\mathcal{H}_\infty$-balanced truncation. IEEE Trans. Autom. Control 36 (6) (1991) 668–682.] to the case of incompressible Navier-Stokes equations. Also, we provide implementations of the underlying numerical procedures that deal with the arising large-scale sparse systems with algebraic constraints from the area of flow problems.

Peter Benner
Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany
benner@mpi-magdeburg.mpg.de

Jan Heiland
Max Planck Institute, Magdeburg
heiland@mpi-magdeburg.mpg.de

Steffen W. R. Werner
Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg
werner@mpi-magdeburg.mpg.de

MS268

Adjoint-based Sensitivity Analysis and Error Estimation for Scale-resolving Turbulent Flow Simulations

Adjoint methods are powerful engineering design tools for computational fluid dynamics. Adjoints have mostly been used for steady flow simulations such as Reynolds-averaged Navier-Stokes (RANS) solvers. Recently, engineers have started to use scale-resolving simulations like Large-eddy simulations (LES) for flows with unsteady separation and jets. These simulations model flows more accurately than RANS and capture the chaotic dynamics inherent in turbulence, making the application of adjoints challenging. This talk discusses the use of adjoint methods for scale-resolving simulations. Adjoint computations of time-averaged objective sensitivities, error estimates, and extreme event predictions will be presented. Sensitivity analysis and error estimation is conducted with non-intrusive least-squares shadowing (NILSS) because the conventional adjoint method computes large, unusable sensitivities for scale-resolving simulations [P. Blonigan. Adjoint sensitivity analysis of chaotic dynamical systems with non-intrusive least squares shadowing. Journal of Computational Physics, 348:803826, Nov. 2017]. The extreme event prediction approach uses a conventional adjoint since the objective is instantaneous [P. Blonigan, M. Farazmand, and T. Sapsis. Are extreme dissipation events predictable in turbulent fluid flow? Submitted to Physical Review Fluids, July 2018]. Results are presented for the minimal flow unit for near-wall turbulence, a channel flow through a truncated domain.

Patrick J. Blonigan
NASA Ames Research Center
USRA
pblonig@sandia.gov

MS268

Grid-adaptation for Chaotic Multi-scale Simulations As a Verification-driven Inverse Problem

The grid-spacing has a direct effect on both the numerical and the modeling errors in coarse-grained simulations of multi-scale problems (e.g., large eddy simulations of turbulence). We make the argument that it is impossible to estimate where errors are introduced in such simulations with absolute certainty, and thus that the problem of finding an “optimal’ adapted grid must be framed as a verification-driven problem. After posing this new problem, one possible general approach to solving it is proposed. This is then tested and demonstrated on a modified Kuramoto-Sivashinsky equation and the turbulent flow over a backward-facing step.

Johan Larsson

University of Maryland
Department of Mechanical Engineering
jola@umd.edu

Siavash Toosi
University of Maryland
stoosi@umd.edu

MS268

Nilsas: Adjoint Sensitivity Analysis for Chaos via Computing Adjoint Shadowing Directions

We develop the NILSAS algorithm, which performs adjoint sensitivity analysis of chaotic systems via computing the adjoint shadowing direction. NILSAS can be implemented with little modification to existing adjoint solvers; its computational cost scales with the dimension of the unstable subspace, but is independent of the number of parameters.

Angxiu Ni
University of California, Berkeley
niangxiu@gmail.com

MS268

Model Reduction and Error Control in Chaotic Flow Simulations

Turbulent flows are characterized by chaotic variations in state variables and are commonly found in many applications such as jet engine mixing and flow over bluff bodies. Large Eddy Simulations (LES) of these chaotic flows have already proven to be useful and beneficial to the design process. However, LES is resource and time intensive. Application of output-based methods for error estimation and mesh adaptation would decrease the cost of these chaotic simulations while still retaining their overall accuracy. However, a direct application of unsteady adjoint-based methods is not possible due to the flows inherent sensitivity to the initial conditions and exponential growth of the corresponding adjoint solutions. We present the Hyper-Reduced Order Modeling - Least Squares Shadowing (HROM-LSS) method that combines model reduction principles and adjoint sensitivity techniques for chaotic 2D Navier-Stokes simulations. All primal solutions are solved using the discontinuous Galerkin finite element method. Results of HROM-LSS for flow over a NACA 0012 airfoil at high Reynolds number show promise for this combined technique and has shown to outperform the LSS method for adjoint calculations and error estimation.

Yukiko S. Shimizu, Krzysztof Fidkowski
University of Michigan
ykmizu@umich.edu, kfid@umich.edu

MS269

On Soft Errors in the Conjugate Gradient: Sensitivity and Robust Numerical Detection

The conjugate gradient (CG) method is the most widely used iterative scheme for the solution of large sparse systems of linear equations when the matrix is symmetric positive definite. Although more than sixty year old, it is still a serious candidate for extreme-scale computation on large computing platforms. On the technological side, the continuous shrinking of transistor geometry and the increasing complexity of these devices affect dramatically their sensitivity to the natural radiation, and thus diminish their reliability. One of the most common effects produced by natural radiation is the single event upset which consists in a bit-flip in a memory cell producing unexpected results at application level. Consequently, the future computing facilities at extreme scale might be more prone to errors of any kind including bit-flip during calculation. These numerical and technological observations are the main motivations for this work, where we first investigate through extensive numerical experiments the robustness of CG to bit-flips in its main computationally intensive kernels, namely the matrix-vector product and the preconditioner application. We further propose numerical criteria to detect the occurrence of such faults; we assess their robustness through extensive numerical experiments.

Emmanuel Agullo
INRIA
emmanuel.agullo@inria.fr

Siegfried Cools
University of Antwerp
siegfried.cools@uantwerp.be

Luc Giraud
Inria
luc.giraud@inria.fr

Wim Vanroose
Antwerp University
wim.vanroose@uantwerpen.be

Emrullah Fatih Yetkin
Istanbul Technical University, Informatics Institute
fatih.yetkin@khas.edu.tr

MS269

Exascale Resilience Strategies for Transient Solvers

The solution of time-dependent partial differential equations arises in a wide range of application areas, for example, fluid dynamics. High-fidelity industrial flows in complex geometries require large-scale computational resources and are one of the drivers towards exascale. However, to maintain the usefulness and efficiency of these tools when run using hundreds of thousands of nodes, they need to be made tolerant of frequent hardware failures. We present our latest efforts to address this challenge, combining the use of the user-level failure mitigation (ULFM) extensions to MPI, multi-level lightweight checksum-checkpoints of arbitrary size, and minimally intrusive approaches to augment existing transient solvers with fault tolerance capabilities. We describe the algorithms and their performance characteristics, and illustrate these techniques through examples using the Nektar++ spectral/hp element framework.

Chris Cantwell
Imperial College London
c.cantwell@imperial.ac.uk

Allan Nielsen
EPFL
allan.nielsen@epfl.ch

David Moxey
Department of Aeronautics
Imperial College London
d.moxey@exeter.ac.uk

Jan S. Hesthaven
EPFL
jan.hesthaven@epfl.ch

Spencer Sherwin
Department of Aeronautics
Imperial College London
s.sherwin@imperial.ac.uk

MS269

Physics-based Checksums for Silent-error Detection in PDE Solvers

We discuss techniques for efficient local detection of silent data corruption in parallel scientific computations, leveraging physical quantities such as momentum and energy that may be conserved by discretized PDEs. The conserved quantities are analogous to algorithm-based fault tolerance checksums for linear algebra but, due to their physical foundation, are applicable to both linear and nonlinear equations and have efficient local updates based on fluxes between subdomains. These physics-based checksums enable precise intermittent detection of errors and recovery by rollback to a checkpoint, with very low overhead when errors are rare. We present applications to both explicit hyperbolic and iterative elliptic (unstructured finite-element) solvers with injected memory bit flips.

Maher Salloum
Sandia National Laboratories
Quantum and Digital Systems Research Department
mnsallo@sandia.gov

Jackson Mayo, Rob Armstrong
Sandia National Laboratories
jmayo@sandia.gov, rob@rob-armstrong.net

MS269

Correcting Uncorrectable Data Faults Using Mathematics and Application Context

Uncorrectable faults in extreme-scale systems are increasing in scope and complexity as we approach exascale. Resilience efforts addressing this kind of fault are of great importance, and a mathematical analysis can be of great use in such an effort. We have extended a method that corrects hardware error detection and correction (EDAC) contextually, and have focused the current effort on data errors in an approach suitable to a real application of interest. In this talk, we will stress the mathematics behind this method and its application to the system, and will discuss the interplay between experimentation, simulation and mathematical theory in developing such a technique.

Laura Monroe, Nathan A. DeBardeleben
Los Alamos National Laboratory
lmonroe@lanl.gov, ndebard@lanl.gov

MS270

Reviewing of High-order Teno Schemes for Hyperbolic Conservation Laws

In this talk, we review the recently proposed high-order TENO schemes for hyperbolic conservation laws: (1) the novel shock-capturing TENO concept for compressible fluid simulations [Fu et al., Journal of Computational Physics 305 (2016): 333-359]. (2) the new dissipation and dispersion controlling strategy for very-high-order TENO schemes [Fu et al., Journal of Computational Physics 349 (2017): 97-121]. (3) implicit large eddy simulations with a high-order TENO scheme [Fu et al., Communications in Computational Physics, 2018, Accepted, in press]. (4) a new framework for high-order TENO reconstruction [Fu et al., Journal of Computational Physics, 2018, Accepted, in press].

Lin Fu
Center for Turbulence Research
Stanford University
linfu@stanford.edu

MS270

A Multi-species, Multi-fluid Model for Simulating Plasma Interpenetration

Multifluid effects, such as plasma interpenetration and species separation, are important in the plasma dynamics of inertial confinement fusion (ICF) and high energy-density physics (HEDP) experiments. Traditional multi-species hydrodynamic codes fail to capture these effects due to a single velocity/momentum representation for all fluid species, and therefore, the multifluid Euler equations are needed, where a distinct set of mass, momentum, and energy equations are solved for each species. In this talk, we describe EUCLID, a high-order, AMR-capable, conservative finite-difference code for multifluid plasma dynamics. We present our governing equations, including a computationally-tractable electron model for our applications where the Debye lengths are much smaller than the characteristic domain length. We also discuss our numerical method as well as the numerical issues resulting from the electrostatic and collisional source terms representing the coupling between the distinct plasma streams, especially in regions where one species is several orders of magnitude sparser than the other species. Finally, we present results where we verify and validate EUCLID with ICF and HEDP experiments. *This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract No. DE-AC52-07NA27344 and funded by the LDRD Program at LLNL under project tracking code 17-ERD-081. LLNL-ABS-757670

Debojyoti Ghosh
Lawrence Livermore National Laboratory
Center for Applied Scientific Computing

ghosh5@llnl.gov

Thomas Chapman, Richard Berger
Lawrence Livermore National Laboratory
chapman29@llnl.gov, berger@llnl.gov

Andris Dimits
Lawrence Livermore National Labs
dimits1@llnl.gov

Jeff Banks
Rensselaer Polytechnic Institute
banksj3@rpi.edu

MS270

High Resolution Schemes for Shock-Capturing in Turbulent Flows

Compressible turbulent flows in many engineering and science applications, such as high-speed propulsion systems, inertial confinement fusion and supernovae evolution, involve shock waves and multi-material interfaces which interact with the flow and dynamically impact the turbulent mixing. Numerical simulations of such problems require special attention to accurately capture the multi-scale spatial and temporal dynamics. This requires high-order and high-resolution numerical schemes which robustly capture shock waves, material interfaces and thin shear layers. We have developed new high-resolution weighted compact schemes(WCHR)1 for this purpose and carried out comparisons with some other high-order schemes such as WCNS-JS, WCNS-Z, WCNS-LD, and high-order artificial properties treatment of Cook and Cabot. The test problems range from classical 1-d test cases including a multi-scale version of the Shu-Osher entropy wave shock interaction problem, 2-d shock-vortex interaction, and 3-d Taylor-Green problem and compressible isotropic turbulence with shocklets. The results demonstrate the ability of the new scheme to localize the required numerical dissipation only to the regions containing discontinuities and preserve wideband accuracy in smooth flow regions.

Sanjiva Lele
Stanford University
Aeronautics and Astronautics
lele@stanford.edu

Akshay Subramaniam, Man Long Wong
Stanford University
akshays@stanford.edu, wongml@stanford.edu

MS271

Uncertainty Quantification of Global Climate Models Using MCMC with Gaussian Process Emulators

One source of uncertainty in climate change simulations comes from the choice of parameters used to characterize unresolvable processes in Global Climate Models (GCMs). The models are poorly constrained by data, and no mathematical tools to tune such parameters are routinely used. This work focuses on the application of Bayesian inverse problem algorithms to climate models. The model tested is an idealized GCM with a relatively simple parameterization scheme for moist convection. The convection parameterization contains 2 parameters: a reference relative humidity, and a relaxation timescale. The goal is to estimate these parameters from synthetic data generated with the GCM itself (i.e., in a perfect-model setting) and to quantify uncertainty in the estimated parameters. The computational cost of GCMs prohibits standard Bayesian inversion approaches such as MCMC. We use ensemble Kalman inversion for parameter estimation, followed by MCMC using Gaussian Process emulators for uncertainty quantification. The emulators are trained during the ensemble Kalman inversion, minimizing their marginal computational costs. We demonstrate that this approach yields reliable estimates of the parameters and their uncertainties, with the relative humidity parameter more tightly constrained than the relaxation time.

Emmet Cleary
Caltech
ecleary@caltech.edu

Tapio Schneider
California Institute of Technology
tapio@caltech.edu

Andrew Stuart
Computing + Mathematical Sciences
California Institute of Technology
astuart@caltech.edu

MS271

Accelerating Physics-constrained Bayesian Inverse Problems Using Inaccurate Models and Data-driven Learning

The solution of physics-constrained Bayesian inverse problems is always reliant on a numerical model of the system which can make predictions with sufficient accuracy. For complex physical simulations however, the ability to infer the unobserved causes is often severely restricted or precluded by the excessive numerical cost of repeated solutions of a highly-complex and high-resolution forward model. While many surrogate-based techniques exist for the acceleration of Bayesian inverse problems in lower dimensions, it remains an open problem for high-dimensional settings. To this end we propose a framework which employs a cheap, lower-resolution model of the physical process in combination with data-driven machine learning tools to construct a multi-fidelity posterior. While loss of information is inevitable, bias and increased epistemic uncertainty can be correctly quantified and incorporated into the posterior in a fully Bayesian framework. We will demonstrate the feasibility of our approach using a nonlinear PDE and high dimensionality of the unknown parameter values.

Maximilian Koschade
Technical University of Munich
maximilian.koschade@tum.de

Phaedon-Stelios Koutsourelakis
Technische Universität München
Germany

p.s.koutsourelakis@tum.de

MS271

A Bayesian Framework for Assessing the Strength Distribution of Composite Structures with Random Defects

Due to their very good strength to weight ratio, composite materials make up over 50% of recent aircraft constructions. The materials are manufactured from very thin fibrous layers ($\sim 10^{-4}$m) and even thinner resin interfaces ($\sim 10^{-5}$m). To achieve the required strength, a particular layup sequence of orientations of the anisotropic fibrous layers is used. During manufacturing, small localised defects in the form of misaligned fibrous layers can occur in composite materials. I will introduce a Markov Chain Monte Carlo algorithm, which derives the stochastic distribution of such wrinkle defects from image data. The defects are parameterized by stochastic random fields defined using Karhunen-Loéve modes inferred from misalignment data extracted from B-Scan data using a modified Multiple Field Image Analysis. Further, I will discuss using the GENEO coarse space as a surrogate model for the fine-scale displacement and stress fields. GENEO computes generalised eigenvectors of the local stiffness matrices on the overlapping subdomains and builds an approximate coarse space by combining the smallest energy eigenvectors on each subdomain via a partition of unity.

Anne Reinarz
Technical University of Munich
reinarz@in.tum.de

Robert Scheichl
University of Heidelberg
r.scheichl@uni-heidelberg.de

Tim J. Dodwell
University of Exeter
College of Engineering, Mathematics & Physical Sciences
t.dodwell@exeter.ac.uk

Linus Seelinger
Universität Heidelberg
Institut für Wissenschaftliches Rechnen
linus.seelinger@iwr.uni-heidelberg.de

MS271

A New Approach for Bayesian Inversion using Sequential Piecewise Polynomial Chaos Expansion

The Bayesian framework for solving inverse problems has mostly relied on Markov-chain Monte Carlo sampling in the past. While this approach is proven to theoretically yield a sample distributed according to the Bayesian posterior distribution, it is well known that in practice it can require considerable tuning efforts and excessive forward model evaluations. We present a novel, sampling free approach for Bayesian inversion. It is based on initially constructing a piecewise polynomial chaos expansion (P-PCE) of the likelihood function over the prior domain. This P-PCE is computed by sequentially building sparse PCEs on subdomains of the prior using the likelihood residual. Following the P-PCE computation, it becomes possible to utilize results from the spectral likelihood expansion (SLE) approach to compute posterior moments and general expectations under the posterior distribution by mere post-processing of the P-PCE coefficients in a sampling-free manner.

Paul-Remo Wagner
ETH Zürich
wagner@ibk.baug.ethz.ch

Stefano Marelli
Chair of Risk, Safety and Uncertainty Quantification
Institute of Structural Engineering, ETH Zurich
marelli@ibk.baug.ethz.ch

Christos Lataniotis, Bruno Sudret
ETH Zurich
lataniotis@ibk.baug.ethz.ch, sudret@ibk.baug.ethz.ch

MS272

Efficient Solver Composition with High-order Methods

We present an extensible low-level library that provides a versatile algebraic interface and optimized implementations suitable for high-order operators: libCEED. This library aims to overcome the challenges in high-order methods that use global sparse matrices as operator representations and enables portable performance through optimal memory transfers and FLOPs for operator evaluation. We investigate operator composition and design of coupled solvers in the context of atmospheric modeling providing examples of the usage of libCEED with PETSc.

Valeria Barra, Jed Brown, Jeremy Thompson
University of Colorado Boulder
valeria.barra@colorado.edu, jed.brown@colorado.edu, jeremy.thompson@colorado.edu

MS272

Collaborative Multiphysics, Multiscale Simulation using Moose

Multiphysics, by its very nature, is best achieved through collaboration. Researchers often have deep knowledge only within one discipline. Attempting to tie multiple disciplines together within a single simulation often means reaching out to peers to collaborate. However, when researchers all have very different codebases and workflows collaboration can become difficult. The open-source Multiphysics Object Oriented Simulation Environment (MOOSE: mooseframework.org) framework presents an ideal platform for collaboration by allowing domain scientists to work independently within their specializations and then smoothly tie their efforts together to achieve high-fidelity, scalable, multiphysics simulations. MOOSE has many features enabling multiphysics collaboration including: a unique object structure, execution of multiple simultaneous MOOSE-based calculations with data transfers between them, and a collaboratively built set of basic physics modules. These features have accelerated many fruitful multiphysics endeavors including: multiscale material science, full-core nuclear reactor simulation and geo-

physical simulation. An overview of how the design of MOOSE accelerates collaboration along with multiple examples of extended collaboration will be presented.

Derek R. Gaston
Idaho National Laboratory
derek.gaston@inl.gov

Cody J. Permann
Advanced Modeling and Simulation
Idaho National Laboratory
cody.permann@inl.gov

David Andrs, John Peterson, Andrew Slaughter
Idaho National Laboratory
david.andrs@inl.gov, jw.peterson@inl.gov, andrew.slaughter@inl.gov

Fande Kong
Department of Modeling and Simulation
Idaho National Laboratory
fdkong.jd@gmail.com

Robert Carlsen, Alex Lindsay, Richard Martineau, Brian Alger, Jason Miller
Idaho National Laboratory
robert.carlsen@inl.gov, alexander.lindsay@inl.gov, richard.martineau@inl.gov, brian.alger@inl.gov, jason.miller@inl.gov

MS272

Coupled Finite Element/Boundary Element Methods for Fast High Intensity Focused Ultrasound Treatment Modelling

High-intensity focused ultrasound (HIFU) is a promising non-invasive, non-radioactive technology for the ablation of tumours. By focusing energy on a small target region the tissue temperature in the target region can be elevated such that with sufficient treatment time the tissue is destroyed. Challenges in planning HIFU treatments in the abdomen include the presence of a large number of scatters, nonlinear effects and the presence of transport mechanisms for heat. We consider the development of fast and robust methods for coupling boundary element and finite element methods with the objective of simulating HIFU treatment. The nonlinear near-target region is modelled using finite element methods, and regions away from the target are modelled using boundary element methods. This is a step towards accurate methods that are fast enough to be clinically relevant, and which will hopefully support patient treatment planning in the future. The presented methods and approaches are implemented in the open source FEniCS (https://fenicsproject.org) and Bempp (https://bempp.com/) libraries.

Timo Betcke
University College London
t.betcke@ucl.ac.uk

Sam Groth
University of Cambridge
spg32@cam.ac.uk

Garth Wells
Department of Engineering
University of Cambridge
gnw20@cam.ac.uk

MS272

Albany: A Trilinos-based Multi-physics Partial Differential Equation Research Tool Created using the Agile Components Code Development Strategy

The focus of this talk is the Albany multi-physics finite element code, and Sandias Agile Components code development strategy, of which Albany is a chief incarnation. The Agile Components strategy enables the rapid development of parallel, efficient multi-physics tools by requiring that: (1) application codes be built primarily from modular pieces (independently developed libraries, e.g., Trilinos), abstract class hierarchies, and template-based generic classes, and (2) projects both leverage and contribute to a comprehensive set of software components, consisting of libraries, interfaces, software quality tools, and demonstration applications. Codes like Albany created using Agile Components are born scalable, fast, robust, performance-portable, and equipped with a quality software infrastructure. This talk will describe the overall design of Albany in the context of the Agile Components approach, and highlight the specifics of several of its key features, including: (1) template-based generic programming, (2) automatic differentiation utilities for calculating Jacobians and enabling beyond-forward analyses requiring sensitivities/adjoints, (3) the Model Evaluator interface for specifying complex nonlinear PDEs, (4) flexible interfaces to 2 adaptive mesh libraries, and (5) performance portability using Kokkos. I will present results showcasing these key capabilities for some Albany multiphysics problems from the domains of climate and solid mechanics.

Irina K. Tezaur
Sandia National Laboratories
ikalash@sandia.gov

Andrew Salinger
CSRI
Sandia National Labs
agsalin@sandia.gov

Glen Hansen, Daniel Ibanez
Sandia National Laboratories
gahanse@sandia.gov, daibane@sandia.gov

MS273

Contact-free Simulations of Rigid Particle Suspensions

In many composite materials, rigid fibers are distributed throughout the material to tune the mechanical, thermal, and electric properties of the composite. The orientation and distribution of the fibers play a critical role in the properties of the composite. Many composites are processed as a liquid molten suspension of fibers and then solidified, holding the fibers in place. Once the fiber orientations are known, theoretical models exist that can predict properties of the composite. Modeling the suspended fibers in the liq-

uid state is important because their ultimate configuration depends strongly on the flow history during the molten processing. We build upon recent advances in boundary integral equations to develop a robust, accurate, and stable method that simulates fibers of arbitrary shape in a planar flow. To maintain fiber separation, a repulsion force and torque are added when required. This repulsion force is free of tuning parameters and is determined by solving a sequence of linear complementarity problems to ensure that the configuration does not have any overlap between fibers. Numerical experiments demonstrate the stability of the method for concentrated suspensions.

Lukas Bystricky
Florida State University
lb13f@my.fsu.edu

Bryan D. Quaife
Scientific Computing
Florida State University
bquaife@fsu.edu

MS273

Boundary Integral Computation via the Polynomial Atlas Method

We describe the polynomial atlas method for boundary integral computations. In this method we parameterize a surface embedded in R^3 using technology from basic differentiable manifold theory: charts, transition functions, partition of unity, and so forth. All of these functions are represented as finite-rank bivariate Chebyshev polynomials. By representing the surface with smooth functions as opposed to a cloud of points or a discrete quadrature rule, we gain the ability to numerically discretize integrals on arbitrarily refined grids. We present several test cases and describe the convergence of the numerical method.

William H. Mitchell
Macalester College
wmitchel@macalester.edu

MS273

Accurate Simulation of Surfactant-covered Droplets in Wall-confined Two-dimensional Stokes Flow

The study of droplet dynamics in micro-scale flows is a field of growing interest, fuelled by the development of miniaturised equipment in biology and chemistry. On such small scales surface forces dominate volume forces and the flow can be described by the linear Stokes equations. We simulate the deformation of droplets in two dimensions using a boundary integral formulation discretised with a Nyström method. If no special treatment is employed, the numerical errors of such methods grow large when computing the solution close to any interface in the domain. We employ a specialised method of numerical integration to avoid these issues and obtain high accuracy in the entire domain, also for drops in close proximity to each other and to the boundaries of the domain. A method for simulating multiple deforming droplets in different channel geometries is presented. The method handles both non-unity viscosity ratios and the inclusion of surface active agents (surfactants) on drop interfaces. Through examples it is shown how both channel geometry and surfactants affect the droplet deformation. The robustness of the method is demonstrated through challenging examples of closely interacting drops and walls.

Sara Pålsson
KTH Royal Institute of Technology
sarapal@kth.se

Anna-Karin Tornberg
KTH
akto@kth.se

MS274

The State-of-the-art of Solvers for Quadratic Eigenvalue Problems

The quadratic eigenvalue problem (QEP) is a special yet most-studied case of nonlinear nonlinear eigenvalue problems that arise in solid mechanics, acoustics and coupled structural acoustics among others. We will present the state-of-the-art of algorithms for computing full and partial eigenpairs of the QEP, along with a numerical evaluation using a large set of QEPS arising from practical applications.

Zhaojun Bai
Departments of Computer Science and Mathematics
University of California, Davis
bai@cs.ucdavis.edu

Ding Lu
University of Geneva
ding.lu@unige.ch

Tianyi Lu
School of Mathematical Sciences
Fudan University
17110840004@fudan.edu.cn

Yangfeng Su
School of Mathematical Sciences
Fudan University, China
yfsu@fudan.edu.cn

MS274

Accelerating Commercial FEA Software through Advanced Computational Technologies

The progress made in the area of Automated Multi-Level Substructuring (AMLS), including GPU acceleration of the AMS (which is the proprietary SIMULIA implementation of the AMLS), the AMS-based substructure generation procedures, and the AMS-based coupled acoustic-structure eigenvalue extraction procedure is discussed. The AMS in Abaqus is accelerated using GPU devices on SMP nodes with multiple processors. The AMS eigensolver is implemented using efficient dynamic parallel scheduler to achieve not only the parallel scaling of the AMS, but also significant reduction of memory required by the AMS. The conventional substructure generation procedure, which in-

cludes projecting the substructure system matrices onto the substructure modal space, takes significant computational time. The AMLS eigensolution procedure can be combined with the substructure generation to address drawbacks of the conventional substructure generation procedure. Fixed-interface, free-interface, and mixed-interface substructures are generated within the AMS eigensolver. If required, the necessary orthogonalization of the dynamic modes is efficiently done on the reduced AMLS subspace. A coupled structural-acoustic problem is modeled using structural, acoustic, and interface elements. Transformation for converting of the original nonsymmetrical coupled structural-acoustic eigenvalue problem to the symmetric eigenproblem with positive semi-definite matrices suitable for solution using the AMS is presented.

Vladimir Belsky
SIMULIA-Dassaults systemes
vladimir.belsky@3ds.com

Mintae Kim
SIMULIA Dassault Systemes
mintae.kim@3ds.com

Mikhail Belyi
SIMULIA-Dassaults systemes
mikhail.belyi@3ds.com

MS274

The Accelerated ACMS (Automated Component Modal Synthesis) Method: Extracting Parallel Efficiency through Shared Memory Processing

Automated Component Modal Synthesis (ACMS) is the high performance, highly scalable core solver in MSC Nastran Modal Frequency Response Analysis. This MSC Nastran solution technology is heavily used by the worldwide automotive industry. ACMS employs a multilevel substructuring algorithm first introduced in MSC Nastran in the year 2000 and which has been under continuous improvement since. The current implementation of ACMS solver in MSC Nastran employs high level shared-memory multiprocessing parallelism. The core solver operates under a sophisticated dynamic parallel scheduling scheme to provide parallel load balancing, improving scalability. ACMS performs domain decomposition of the global stiffness and mass matrices using a modified nested dissection graph partitioner. After constructing a multilevel reduction tree, component modal synthesis is performed at each node in the tree. Examples are shown from auto industry of performance and scalability advances on current production models in recent years. We present results demonstrating how todays requirements were presented originally, and todays results having met those challenges. Next, we report on new challenges resulting from the automotive industry, and how the collaborative effort of MSC Software intends to meet those challenges. Some preliminary examples of future requirements and solutions are presented.

David Ehrlich, Peter Schartz
MSC Software Corporation
david.ehrlich@mscsoftware.com, peter.schartz@mscsotware.com

MS274

Lanczos Method for Buckling Eigenvalue Problems of Singular Pencils

We consider computing a few desired eigenpairs of the buckling eigenvalue problem $Kx = \lambda K_G x$, where K is symmetric positive semi-definite and K_G is symmetric indefinite. Furthermore, the pencil $K - \lambda K_G$ is singular, namely, K and K_G share a non-trivial common null space. To use an industrial strength shift-invert Lanczos method, we face two critical issues: (1) The shift-invert operator $(K-\sigma K_G)^{-1}$ may not exist or be extremely ill-conditioned. (2) The use of the semi-inner product induced by K may lead to the growth of the components of the Lanczos vectors lying in the null space of K to overwhelm the components in the row-space of K, and cause permanent loss of information and the failure of the method. We will present our approaches to resolve these issues and demonstrate the efficacy of our approaches by industrial examples.

Chao-Ping Lin
University of California, Davis
cplin@ucdavis.edu

Huiqing Xie
East China University of Science and Technology
hqxie@ecust.edu.cn

Zhaojun Bai
Departments of Computer Science and Mathematics
University of California, Davis
bai@cs.ucdavis.edu

MS275

Maximum Principle Preserving Exponential Time Differencing Schemes for the Nonlocal Allen-Cahn Equation

The nonlocal Allen-Cahn (NAC) equation is a generalization of the classic Allen-Cahn equation by replacing the Laplacian with a parameterized nonlocal diffusion operator, and satisfies the maximum principle as its local counterpart. In this talk, we develop and analyze first and second order exponential time differencing (ETD) schemes for solving the NAC equation, which unconditionally preserve the discrete maximum principle. The fully discrete numerical schemes are obtained by applying the stabilized ETD approximations for time integration with the quadrature-based finite difference discretization in space. We derive their respective optimal maximum-norm error estimates and further show that the proposed schemes are asymptotic compatible, i.e., the approximate solutions always converge to the classic Allen-Cahn solution when the horizon, the spatial mesh size and the time step size go to zero. We also prove that the schemes are energy stable in the discrete sense. Various experiments are performed to verify these theoretical results and to investigate numerically the relationship between the discontinuities and the nonlocal parameters.

Qiang Du

Columbia University
Department of Applied Physics and Applied Mathematics
qd2125@columbia.edu

Lili Ju
University of South Carolina
Department of Mathematics
ju@math.sc.edu

Xiao Li
Beijing Computational Science Research Center
xiaoli@csrc.ac.cn

Zhonghua Qiao
Department of Applied Mathematics
The Hong Kong Polytechnic University
zhonghua.qiao@polyu.edu.hk

MS275

A Time Adaptive Rescaling Method for Computing Interface Problems

Abstract not available

Shuwang Li
Department of Applied Mathematics
Illinois Institute of Technology
sli@math.iit.edu

MS275

The SAV Approach for Gradient Flows and Positivity Preserving Schemes for Poisson-Nernst-Planck Equations

We first present in this talk the scalar auxiliary variable (SAV) approach for a large class of gradient flows. The technique is not restricted to specific forms of the nonlinear part of the free energy, it leads to linear and unconditionally energy stable second-order (or higher-order with weak stability conditions) schemes which only require solving decoupled linear equations with constant coefficients. Hence, these schemes are extremely efficient as well as accurate. However, the SAV approach does not usually preserve maximum principle or positivity. We then present a set of numerical schemes for the Poisson–Nernst–Planck equations which are uniquely solvable, preserve the positivity and unconditionally energy dissipative.

Jie Shen
Purdue University
Department of Mathematics
shen7@purdue.edu

MS275

Diffusion Generated Methods for Target-valued Maps

A variety of tasks in inverse problems and data analysis can be formulated as the variational problem of minimizing the Dirichlet energy of a function that takes values in a certain target set and possibly satisfies additional constraints. These additional constraints may be used to enforce fidelity to data or other structural constraints arising in the particular problem considered. We will present diffusion generated methods for solving this problem for a wide class of target sets and prove stability and convergence. We will give examples of how these methods can be used for the geometry processing task of finding Dirichlet partitions, constructing smooth orthogonal matrix valued functions, and solving inverse problems for target valued maps. This is joint work with Braxton Osting.

Dong Wang
University of Utah
dwang@math.utah.edu

MS276

Online Adaptive Basis Refinement and Compression for Reduced-order Models

Reduced-order models (ROMs) can provide a fast approximation to many-query problems in scenarios where the underlying model is too expensive to solve. However, while ROM simulations are typically fast, there is no guarantee that the output they produce is accurate, as their accuracy depends on the training data used to construct the low-dimensional basis. To combat this problem, we propose a mechanism for enriching low-dimensional bases *a posteriori* in a way which eventually guarantees convergence to the full-order model, but without requiring any additional full-order solves. The mechanism is based on the previously proposed h-refinement approach for ROMs, but improves upon previous work in two crucial ways. First, our new method allows refinement to be performed with respect to any orthogonal basis allowing the refinement mechanism to be tailored to the physics of the underlying problem. And second, we provide a fast algorithm for compressing enriched bases via an efficient online POD whose operation count is independent of the underlying full-order-model dimension. These two features allow our techniques to act both as an effective failsafe mechanism for reduced order models, and also as a fast online way of enriching or correcting reduced-order models, thereby enabling them to better resolve the physics of the underlying problem and even compensate for inadequate training.

Philip Etter
Stanford University
paetter@stanford.edu

Kevin T. Carlberg
Sandia National Laboratories
ktcarlb@sandia.gov

MS276

A Variational Approach to Closure of Nonlinear Dynamical Systems Based on Small-scale Parameterizations

In this talk, we discuss a new approach to deal with the parameterization problem of small spatial scales by large ones for nonlinear stochastic PDEs. The approach is variational in nature, and relies on manifolds that aim to provide in a mean square sense approximate small-scale parameterizations. We will highlight a simple semi-analytic approach to determine such manifolds based on backward-forward

auxiliary systems. We will then illustrate the approach on some simple academic models.

Honghu Liu
Virginia Tech, Department of Mathematics
hhliu@vt.edu

Mickael Chekroun
University of California, Los Angeles
Department of Atmospheric and Oceanic Sciences
mchekroun@atmos.ucla.edu

James McWilliams
Atmospheric Sciences
UCLA
jcm@atmos.ucla.edu

Shouhong Wang
Department of Mathematics
Indiana University
showang@indiana.edu

MS276

A Pod Based Artificial Compression Scheme for the Incompressible Navier-Stokes Equations

In recent years there has been a growing interest in incorporating pressure data into reduced order models for incompressible flows. Due to the lack of fulfillment of the discrete inf-sup condition, a naive incorporation of pressure basis functions will lead to spurious oscillations in the reduced order solution. Some popular approaches for overcoming this issues include the Pressure Stabilization Petrov Galerkin approach, supremizer stabilization, or penalty schemes. Depending on the application these methods can be too expensive to implement, or require parameters that have a dependency which may ruin the offline/online decomposition of the reduced order approach. In an attempt to overcome these issues, we present a new artificial compression based reduced order model known as AC-ROM. AC-ROM is notable for its computational efficiency and convergence even when the basis does not fulfill the discrete inf-sup condition.

Michael Schneier
Department of Mathematics
University of Pittsburgh
mschneier89@gmail.com

Traian Iliescu
Department of Mathematics
Virginia Tech
iliescu@vt.edu

William Layton, Victor P. Decaria, Michael McLaughlin
University of Pittsburgh
wjl@pitt.edu, vpd7@pitt.edu, mem266@pitt.edu

MS276

Incremental POD Mode Algorithm for Fluids with Easily Computable Error Bounds

We discuss an incremental algorithm for proper orthogonal decomposition (POD) computations. Specifically, we develop an incremental matrix SVD algorithm with respect to a weighted inner product for POD computations of data arising from Galerkin-type simulation methods for time dependent PDEs. The algorithm initializes and efficiently updates the POD eigenvalues and modes during the time stepping in a PDE solver without storing the simulation data. Also, the algorithm returns an easily computed error bound. We demonstrate the effectiveness of the algorithm using finite element computations for fluid flows.

John Singler
Missouri S&T
Mathematics and Statistics
singlerj@mst.edu

MS277

Title Not Available

Abstract not available

Rich Archibald
Oak Ridge National Laboratory
archibaldrk@ornl.gov

MS277

Geometric Insights into Spectral Clustering by Graph Laplacian Embeddings

Abstract not available

Franca Hoffmann
Caltech
franca.hoffmann@gmail.com

Nicolas Garcia Trillos
Brown University
nicolas_garcia_trillos@brown.edu

Bamdad Hosseini
California Institute of Technology
bamdadh@caltech.edu

MS277

Nonlinear Transformation via Reversible ResNets for Dimension Reduction in Function Approximation

We consider the dimension reduction problem in approximating high-dimensional functions. There are various existing methods, including sliced inverse regression and active subspace, reduce the input dimension by constructing a linear transform of the original coordinates. In this work, we developed a new method for generating nonlinear transforms of the original coordinates. Specifically, we define the target nonlinear transform as a reversible ResNet, which guarantees the necessary bijective property. The objective function for training the ResNet is exclusively defined such that a well-trained transform can align most of the transformed coordinates with the nonlinear isosurface of the target function. As such, the function is only sensitive to a few transformed coordinates that are perpendicular to the isofurface. Several test examples using high-dimensional

function will be given to demonstrate the effectiveness of the new method.

Guannan Zhang
Oak Ridge National Laboratory
zhangg@ornl.gov

MS278

Covariant Neural Networks for Learning Structured Data

Deep neural network architectures have proved to be remarkably successful at some of the classical applied machine learning tasks, such as object recognition, face detection, and a variety of natural language tasks. However, in some domains where the learning task has more structure, such as learning to simulate physical systems or learning from graphs, "brute force" deep learning approaches are less likely to be successful. We argue that in many of these cases, the key is to take the natural invariances of the data into account, and discuss some recent developments in the field of *covariant neural networks* that achieve this goal. This talk is partially sponsored by American Family Insurance and eMALI.IO Ltd.

Risi Kondor, Brandon Anderson
University of Chicago
risi@cs.uchicago.edu, brandona@uchicago.edu

MS278

Non-IID Model Learning via Graph Neural Networks

Most of the more commonly used machine learning methods assume that training and test data comes from independent and identically distributed (IID) random variables. However, this assumption is strongly violated when the data points available to learn/test the model are highly interdependent. Two examples of this scenario are: when the data exhibits temporal or spatial correlations and when the task is to learn relative characteristics between data points (e.g., ranking or attribute learning). In this talk we show how emerging ideas in graph neural networks can yield a solution to various problems that broadly fall under this characteristics. More specifically, we show interesting results for the relative attribute learning problem from images. This setting, naturally benefits from exploiting the graph of dependencies among the different relative attributes of images, especially when only partial ordering is provided at training time. Our experiments show that this simple framework is effective in achieving competitive accuracy with specialized methods for both relative attribute learning and binary attribute prediction.

Glenn Fung
American Family Insurance
GFUNG@amfam.com

MS278

Understanding Climate-vegetation Interactions in Global Rainforests through a GP Model-tree Analysis

The question about the future of rainforests is of immense importance with the recent severe Amazonian droughts in 2005 and 2010 and ongoing drought since 2000 in the Congo region. Contradicting research has claimed that these forests are resilient to such extreme climatic events. A significant reason behind these diverse conclusions is the lack of a holistic spatio-temporal analysis of the remote sensing data available for these regions. Small scale studies that use statistical correlation measure and simple linear regression to model the climate-vegetation interactions have suffered from the lack of complete data representation and the use of simple (linear) models that fail to represent physical processes accurately, thereby leading to inconclusive or incorrect predictions about the future. In this talk we present a genetic programming (GP) based approach called symbolic regression for discovering nonlinear equations that govern the vegetation climate dynamics in the rainforests. Expecting micro-regions within the rainforests to have unique characteristics compared to the overall general characteristics, we use a modified regression-tree based hierarchical partitioning of the space and build a nonlinear GP model for each partition. Based on the partitioning of the observed data points over years, we can conclude that in the absence of adequate precipitation, the trees adopt to reach a different steady state and recover as soon as precipitation is back to normal.

Anuradha Kodali
Allstate Corporation
akoda@allstate.com

Kamalika Das
University of Maryland, Baltimore County
kdas1@cs.umbc.edu

MS278

Contextually-sensitive Named Entity Disambiguation

Named Entity Disambiguation (or Entity Linking), the task of predicting link probability between a specific text mention and candidate entities in a knowledge base (KB) with high accuracy, is one of the major challenges for deciphering massive and diverse content in real-world in petabyte scale. It is challenging because mentions are often very ambiguous, and it often requires the understanding of local and global context. In an experiment with real-world corpus from Facebook, a naïve system of entity linking of text mentions would have 85% of error rate without a sophisticated disambiguation system. We have developed a system that takes both text context features from the local context and global entity coherence features into account. The text context features are used to measure the similarity between candidate entity and the derived textual context. The system involves a joint learning system of word embeddings and entities, and multiple predictive models for disambiguation, pruning and salience detection. The system is highly accurate (.82 F1-measure) and is currently deployed to the production system of Facebook.

Jeonghee Yi, Kyle Yan, Jingfei Du, Javad Dousti, Bi Xue
Facebook, Inc.

jeonghee.yi@gmail.com, xyan@fb.com, jingfeidu@fb.com, dousti@fb.com, bixue@fb.com

MS279

Controlling Oscillations in High-order Accurate Methods Through Artificial Neural Networks

High-order methods for approximating solutions to hyperbolic systems of conservation laws, need to be carefully treated near discontinuities to avoid Gibbs oscillations. In the context of discontinuous Galerkin (DG) schemes, several methods have been developed to control spurious oscillations, such as limiting the local approximating polynomials or adding shock capturing terms. However, most existing methods require the prescription of problem-dependent parameters, which are usually determined empirically. A non-optimal choice of these parameters can either lead to the re-appearance of Gibbs oscillations, or the loss of accuracy in smooth regions. In this talk, we propose a new data-driven approach to overcome this bottleneck. In particular, we train artificial neural networks (ANNs) using supervised learning, which are then used as a black-box to detect numerical discontinuities and prescribe artificial viscosity required in the shock capturing terms. The advantage of the proposed strategy is that it is free from problem-dependent parameters, computationally efficient, and can easily be integrated into existing code frameworks. Several numerical results are presented to demonstrate the robustness of the networks in the framework of Runge-Kutta DG schemes.

Jan S. Hesthaven, Niccolo Discacciati
EPFL
jan.hesthaven@epfl.ch, niccolo.discacciati@epfl.ch

Deep Ray
EPFL-SB-MATH-MCSS
deep.ray@epfl.ch

Jian Yu
School of Aeronautic Science and Engineering
Beihang University
yuj@buaa.edu.cn

MS279

Exploiting Sparsity for Modeling and Control of Dynamical Systems

High-dimensional, nonlinear, multi-scale phenomena, such as turbulence, are ubiquitous. The lack of simple equations and unprecedented amount of available high-fidelity data are leading to a paradigm shift in how we interact with complex nonlinear systems. Leading approaches stem from data-driven methods which have the potential to discover new mechanisms, models and control laws and are driven by the tremendous advances in computing power, new sensors and infrastructures, and advanced algorithms in machine learning. In this talk, I will discuss recent advances in data-driven, equation-free architectures leveraging sparsity-promoting techniques for the modeling and control of dynamical systems. One direction is connected to Koopman operator theory, which has emerged as a principled framework to obtain linear embeddings of nonlinear dynamics, enabling the estimation, prediction and control of strongly nonlinear systems using standard linear techniques. A data-driven architecture is presented for the identification of Koopman eigenfunctions using sparse regression and polynomial expansions, based on the partial differential equation governing the infinitesimal generator of the Koopman operator. In addition, I will discuss work related to statistical modeling in fluids and how to exploit sparsity in dynamical systems for sensing. The presented work is demonstrated on Hamiltonian systems and different high-dimensional systems from fluids.

Eurika Kaiser, Steven Brunton
University of Washington
eurika@uw.edu, sbrunton@uw.edu

J. Nathan Kutz
University of Washington, Seattle
Dept of Applied Mathematics
kutz@uw.edu

Bernd Noack
LIMSI-CNRS
Technical University Braunschweig
bernd.noack@limsi.fr

Andreas Spohn
Institute PPRIME, ENSMA
andreas.spohn@ensma.fr

MS279

Data-driven Discovery of Governing Physical Laws and Their Parametric Dependencies

We present a data-driven method for the discovery of parametrically dependent differential and partial differential equations, thus allowing one to disambiguate between the underlying evolution equations and their parametric dependencies. Group sparsity is used to ensure parsimonious representations of observed dynamics in the form of a parametric ODE or PDE, while also allowing the coefficients to have arbitrary time series, or spatial dependence. This work builds on previous methods for the identification of constant coefficient ODEs/PDEs, expanding the field to include a new class of equations which until now have eluded machine learning based identification methods. We show that group sequentially thresholded ridge regression outperforms group LASSO in identifying the fewest terms in the PDE along with their parametric dependency. The method is demonstrated on four canonical models with and without the introduction of noise.

J. Nathan Kutz
University of Washington, Seattle
Dept of Applied Mathematics
kutz@uw.edu

Samuel Rudy
Applied Mathematics
University of Washington
shrudy@uw.edu

Steven Brunton
University of Washington

sbrunton@uw.edu

MS279

Machine Learning Algorithms for Explicit Algebraic RANS Turbulent Closures: A Galilean Invariant Approach

The maximal hot gas temperature in gas turbine engines is continuously increasing to further improve the thermodynamic cycle efficiency. Thus, temperatures are nowadays far beyond the melting point of the turbine blade material; development of sophisticated cooling schemes for gas turbine blades is therefore an always current topic. Computational Fluid Dynamics (CFD) can complement or reduce the experiments needed for the design of complex cooling passages. Transient simulations of wall heat transfer such as Direct Numerical Simulation (DNS) and Large Eddy Simulation (LES) are prohibitively expensive in terms of computational times and not feasible today for high Reynolds numbers. For this reason, turbulence models for the Reynolds-Averaged Navier Stokes (RANS) equations, though partially based on ideas from more than a century ago, still undergo continuous development. A means of improving RANS turbulence closures is the replacement of empirical constants or correlations with data-driven functional forms derived from machine learning algorithms. We focus our investigation on the development of constitutive equations for the Reynolds stress tensor and the scalar flux vector. Neural networks are trained to reproduce the correct values of the model scalars at each position in the flow domain. The resulting models show very good accuracy and generalize well in flow conditions different from these of the training cases.

Corrado Sotgiu, Bernhard Weigand
University of Stuttgart
corrado.sotgiu@itlr.uni-stuttgart.de,
bernhard.weigand@itlr.uni-stuttgart.de

MS280

Large-scale Optimization in Deep Learning for PDE Representation

We investigate the use of deep neural networks as surrogates for PDE parameter-to-observable maps. Specifically, we consider deterministic PDE maps with uncertain input parameters. Our target applications are stochastic PDE constrained optimization in settings where nonlinear PDE solves are prohibitively costly. Neural networks provide a surrogate framework that is scalable in terms of the input parameter dimension. We use second order inexact Newton methods that are scalable in the control parameter dimension, and also ideal for avoiding saddle points which are ubiquitous in neural network training. Numerical experiments for Helmholtz problems will be shown.

Tom O'Leary-Roseberry
The University of Texas at Austin
The Institute for Computational and Engineering Sciences
tom@ices.utexas.edu

Joshua Chen
UT Austin
ICES
joshua@ices.utexas.edu

Peng Chen
UT Austin
peng@ices.utexas.edu

Umberto Villa
University of Texas at Austin
uvilla@wustl.edu

Omar Ghattas
The University of Texas at Austin
omar@ices.utexas.edu

MS280

Efficient PDE-constrained Optimization under Uncertainty using Adaptive Model Reduction and Sparse Grids

This work introduces a framework for accelerating optimization problems governed by partial differential equations with random coefficients by leveraging adaptive sparse grids and model reduction. Adaptive sparse grids perform efficient integration approximation in a high-dimensional stochastic space and reduced-order models reduce the cost of objective function and gradient queries by decreasing the complexity of primal and adjoint PDE solves. A trust region method that allows for inexact gradient and objective evaluations manages these two sources of inexactness and ensures global convergence. Numerical results show the proposed method is an order of magnitude less expensive than existing methods when applied to the model problem of optimal stochastic flow control of the incompressible Navier-Stokes equations.

Matthew J. Zahr
Lawrence Berkeley National Laboratory
University of California, Berkeley
mjzahr@lbl.gov

MS280

An Adaptive Sample-based Approximation Approach for Stochastic Inverse Problems

In this work, we adopt a framework based on a Gibbs posterior for updating belief distributions for inverse problems governed by PDEs. In contrast with regular Bayesian methods where distributions of noise are assumed to be known exactly, the Gibbs posterior update does not require a likelihood function. Hence, no exact model for the noise is needed. Instead, the Gibbs posterior is applicable where the unknown parameters are connected to the data through a loss function. We employ a sample based discretization to approximate the continuous prior distribution, which after applying the Gibbs update results in an explicit formula for the posterior weights associated with each sample. To control the number of samples and provide efficient approximations to the posterior, we borrow ideas from sequential Monte Carlo methods to adaptively add samples that cluster within the support of the posterior in a sequential manner. To manage the cost of propagating an increasing number of samples through the loss function,

we employ a local reduced basis method to build efficient surrogate models. We demonstrate the performance of our approach through several numerical examples.

Zilong Zou
Duke University
zilong.zoou@duke.edu

Harbir Antil
George Mason University
Fairfax, VA
hantil@gmu.edu

Sayan Mukherjee
Duke University
sayan@stat.duke.edu

Wilkins Aquino
Professor
Duke University
wa20@duke.edu

MS281

Communication-avoiding Sparse Matrix Algorithms for Large Graph and Machine Learning Problems

Machine learning (ML) and graph analytics have proven to be powerful tools for analyzing business, social and scientific data sets. Data from scientific domains where ML and graph methods excel are often sparse with many missing entries. Many ML methods rely on numerical optimization algorithms, which themselves are based on sparse matrix operations. Sparse linear algebra software is needed for ML due to sparse datasets or the need to enforce output (and model) sparsity for avoiding overfitting and increasing interpretability. Additionally, a vast majority of graph algorithms can be cast into the language of sparse matrices as exploited by the GraphBLAS effort. The challenges of ML and graph methods for science problems include extreme-scale data volumes and data rates, necessitating parallel algorithms that will run on exascale architectures. Due to sparsity, popular implementations of common ML and graph algorithms struggle to efficiently harness the capabilities of large-scale parallel computers. One prevalent problem is the increasingly dominant cost of communication. In this talk, I will describe a sample from our recent work on distributed-memory parallelization of prevalent ML and graph problems. I will highlight the importance of communication-avoiding sparse matrix primitives to achieve scalability in these problems.

Aydin Buluc
Lawrence Berkeley National Laboratory
abuluc@lbl.gov

MS281

Enabling Significant Storage Reduction for Deep Non-Linear Neural Nets via Efficient Self-Adaptive Lossy Compression

Sparse linear algebra kernels (S-BLAS) such as SpMV, SpGEMM, SpTrsv, and SpTranspose are the most widely-used building-blocks in many important HPC applications and emerging applications. Spare kernels often turn out to be the bottlenecks of HPC applications due to their indirect memory access and irregular computing, leading to poor memory and computing efficiency. This is one of the most fundamental challenges that HPC researchers face today and the current complex heterogeneous architectures exacerbate the poor scalability of these kernels. For example, existing optimization paradigms for single or multicore platforms primarily focus on cache performance tuning. However, the latest development such as many-core accelerators, novel memory architectures, and network design imply that scalability is the most crucial feature to harvest the compute power of the modern hardware. Therefore, conventional designs without a focus on scaling have to be redesigned to incorporate this new development in the future extreme heterogeneous environment. In this work, we focus on developing highly scalable sparse kernels on modern multi-GPU platforms. We start with SpVM using a new workload balancing strategy and two runtime scheduling methods. The experimental results show that our design achieves 1.35x to 4.33x (3.01x on average) speedup against the state-of-the-art work with superior scaling.

Shuaiwen Song
Pacific Northwest National Laboratory, USA
shuaiwen.song@pnnl.gov

Jieyang Chen
University of California, Riverside
jchen098@ucr.edu

MS281

Sparse Tensor Algebra and its Relations to Matrix and Graph Problems

This talk consists of two parts: the recent development of sparse tensor decompositions and the potential relation between tensor, matrix, and graph algorithms. Sparse tensor decompositions are critical to applications of tensor-based analytics. This talk presents some recent research of sparse tensor decompositions and tensor algebra from high performance computing view. Coordinate (COO) and Hierarchical COO (HiCOO) formats for general sparse tensors will be discussed along with parallel optimizations on diverse platforms. HiCOO improves upon COO by compressing the indices in units of sparse tensor blocks, with the goals of preserving the mode-agnostic simplicity of COO while reducing the bytes needed to represent the tensor and promoting data locality. Sparse matrices, tensors, and graphs are different descriptive methods of sparse data. We compare these methods under different application scenarios and state their advantages and disadvantages to inspire the future research.

Jiajia Li
Pacific Northwest National Laboratory
Pacific Northwest National Laboratory
jiajia.li@pnnl.gov

Jimeng Sun
College of Computing
Georgia Institute of Technology
jsun@cc.gatech.edu

Richard Vuduc
Georgia Institute of Technology
richie@cc.gatech. edu

MS281

SPARTan: Scalable PARAFAC2 for Large and Sparse Data

In exploratory tensor mining, a common problem is how to analyze a set of variables across a set of subjects whose observations do not align naturally. For example, when modeling medical features across a set of patients, the number and duration of treatments may vary in time, meaning there is no meaningful way to align their clinical records across time points. To handle such data, the state-of-the-art tensor model is the so called PARAFAC2, which yields interpretable and robust output and can naturally handle sparse data. However, its main limitation has been the lack of efficient algorithms that can handle large-scale datasets. In this work, we fill this gap by developing a scalable method to compute the PARAFAC2 decomposition of large and sparse datasets, called SPARTan. Our method exploits special structure within PARAFAC2, leading to a novel algorithmic reformulation that is both faster (in absolute time) and more memory-efficient than prior work. We evaluate SPARTan on both synthetic and real datasets, showing 22X performance gains over the best previous implementation and also handling larger problem instances for which the baseline fails. Furthermore, we are able to apply SPARTan to the mining of temporally-evolving phenotypes on data taken from real and medically complex pediatric patients. The clinical meaningfulness of the phenotypes identified, as well as their temporal evolution over time for several patients, have been endorsed by clinical experts.

Ioakeim Perros
Georgia Institute of Technology
perros@gatech.edu

Evangelos Papalexakis
University of California, Riverside
epapalex@cs.ucr.edu

Fei Wang
Cornell University
few2001@med.cornell.edu

Rich Vuduc
Georgia Institute of Technology
richie@cc.gatech.edu

Elizabeth Searles
Childrens Healthcare Of Atlanta
n/a

Michael Thompson
N/A
n/a

Jimeng Sun
Georgia Institute of Technology
jimeng.sun@gmail.com

MS282

Parallel Algorithms for Multiple Tensor-times-matrix Computation

The multiple tensor-times-matrix (multi-TTM) computation is a kernel within multiple tensor decomposition computations, including the higher-order SVD algorithm for computing a Tucker decomposition. Multi-TTM can be performed as a sequence of TTMs, which are equivalent to matrix multiplications, or it can be performed all at once, exploiting the tensor structure of the overall computation. This talk will cover a range of algorithmic technique for parallelizing multi-TTM and discuss the tradeoffs in terms of both the computational and interprocessor communication costs.

Grey Ballard
Wake Forest University
ballard@wfu.edu

MS282

Spacey Random Walks and Closely Related Tensor Eigenvector Problems for Data Analysis

Abstract not available

David F. Gleich
Purdue University
dgleich@purdue.edu

Austin Benson
Cornell University
arb@cs.cornell.edu

MS282

Tensor-tensor Products for Optimal Representation and Compression of Snapshot Data

The Tensor-tensor algebra proposed by Kilmer and Martin (LAA 2011) offers matrix algebra mimetic properties. In particular, an Eckart-Young optimality theorem was proved upon this construct. Here, we extend this result to an entire family of tensor-tensor products and show that the corresponding tensor approximation is often better than its matrix counterpart. Armed with this more complete framework surrounding these tensor-tensor products, we propose new tensor truncated SVD variants which provide further approximated representation and compression and discuss under which conditions they are considered optimal. We finalize with a numerical study demonstrating the utility of the theory.

Misha E. Kilmer
Mathematics Department
Tufts University
misha.kilmer@tufts.edu

Lior Horesh
IBM Research
lhoresh@us.ibm.com

Haim Avron
Tel Aviv University
haimav@post.tau.ac.il

Elizabeth Newman
Tufts University
e.newman@tufts.edu

MS282

Randomized Pass-efficient Algorithms for Tucker Decompositions

This talk is concerned with compression of large-scale datasets that are expensive to store and compute with but can be efficiently compressed and stored in an appropriate tensor format. In recent years, randomized matrix methods have been used to efficiently and accurately compute low-rank matrix decompositions. Motivated by this success, we focus on developing randomized algorithms for tensor decompositions in the Tucker representation. Specifically, we develop randomized versions of two well-known compression algorithms, namely, HOSVD and STHOSVD. We present a detailed probabilistic analysis of the error of the randomized tensor algorithms. We also present variants of these algorithms that tackle specific challenges posed by large-scale datasets. The first variant adaptively finds a low-rank representation that satisfies a given tolerance and is beneficial when the target-rank is not known in advance. The second variant is pass-efficient in that it requires fewer passes through the data, and it is beneficial when the tensor is difficult to load in memory. Numerically, we test the algorithms on several different datasets: synthetic test tensors that have appeared in the literature, and realistic applications involving the compression of handwriting digit samples in the MNIST database as well as images in the Yale Face dataset.

Rachel Minster, Arvind Saibaba
North Carolina State University
rlminste@ncsu.edu, asaibab@ncsu.edu

MS283

The Abstraction of 3-D Biomolecular Structure and Property Data for the Training of OpeNNdd Models

Generating and operating with large amounts of data is typically a characteristic of efforts involving neural networks. Large-sized and high-resolution datasets tend to require extensive amounts of memory and time to generate. These challenges may be minimized using various techniques involving supercomputing and data manipulation such as compression supported parallelization, abstracting information within datasets, and optimizing the methods in which data is being created and stored. The application of these techniques toward enabling the neural networks that underlie the open access machine learning platform OpeNNdd will be presented. These efforts are helping to find medicine specifically for addressing severe acute respiratory syndrome (SARS) and advancing medicine design in general.

Bryce Kroencke
American River College
brycekroencke@gmail.com

MS283

The Development and Validation of OpeNNdd Network Architecture

In the process of computationally designing medicine, physics-based equations which describe the interaction energy between a receptor and ligand are typically solved. This process can take several minutes to complete making the high throughput evaluation of medicine candidates challenging. What is sought is a method which combines speed and accuracy in predicting interaction energies so that medicine candidates may more quickly and confidently be vetted. Deep convolutional neural networks (DCNN) are emerging as a powerful tool in computational medicine design since they may estimate interaction energies at the accuracy of higher-levels of theory while being as fast as lower-levels of theory. The application of DCNNs to finding medicines against the severe acute respiratory syndrome (SARS) causing virus, using the open access machine learning platform OpeNNdd, will be discussed.

Nicholas Pavini
American River College
nickpavini@gmail.com

MS283

An Open Access Platform Empowering Research, Education, and the Crowdsourced Discovery of Medicine

Two grand challenges of our time are advancing personalized learning and engineering better medicines according to the National Academy of Engineering. Many diseases do not have therapeutic interventions while reemerging diseases are making medicines that were once effective at fighting them obsolete. Many students struggle to find inspiration, a creative outlet, and purpose in what they are learning. Efforts at creating new medicine are typically carried out by small groups of highly trained scientists working in secret. OpeNNdd is a novel platform which empowers students and citizen scientists to learn science in an engaging way and create ideas for medicine in an open access and collaborative manner. OpeNNdd hosts several neural networks which increase the speed, accuracy, and information content of medicine design efforts enabling crowdsourced medicine discovery and design on a global scale. The development and application of OpeNNdd will be discussed as well as some of the challenges and opportunities that remain.

Benjamin Samudio
Sierra College
samudioben@gmail.com

MS284

Nonlinear Eigenvalue Localization for Damping Bounds

Nonlinear eigenvalue problems arise naturally from frequency-domain analysis of models with damping. In

many cases, we do not understand the damping mechanisms as well as we would like, and the models we use are crude, but we would like to say whether a mode is "highly damped" or "lightly damped." In this talk, we discuss how nonlinear eigenvalue localization results can help answer such questions even when the model used for the analysis is known to be crude or uncertain.

David S. Bindel
Cornell University
bindel@cs.cornell.edu

MS284

The Nonlinear FEAST Algorithm

In this talk, we will present the Nonlinear FEAST (NLFEAST) algorithm to solve eigenvalue problems with the nonlinear dependence on eigenparameter. As a generalization of the Residual Inverse Iteration [A. Neumaier, *Residual inverse iteration for the nonlinear eigenvalue problem*, SIAM J. Numer. Anal. 22 (5) (1985) 914–923] with multiple shifts, the NLFEAST can be used to efficiently compute the eigenvalues enclosed in the user-defined region in the complex plane together with the associated eigenvectors. We will present the convergence analysis of the proposed method as well as various practically-motivated examples to illustrate the applicability of the new algorithm to solve general (non-polynomial) nonlinear eigenvalue problems.

Agnieszka Miedlar
University of Kansas
amiedlar@ku.edu

MS284

Automatic Rational Approximation and Linearization for Nonlinear Eigenvalue Problems

We present a method for solving nonlinear eigenvalue problems using rational approximation. The method uses the AAA method by Nakatsukasa, Séte and Trefethen to approximate the nonlinear eigenvalue problem by a rational eigenvalue problem and is embedded in the state space representation of a rational polynomial by Su and Bai. The advantage of the method compared to related techniques such as NLEIGS and infinite Arnoldi is that the rational interpolant is computed efficiently by an automatic procedure. In addition, a set-valued approach is developed that allows for building a low degree rational approximation of a nonlinear eigenvalue problem, and the method perfectly fits within the framework of the Compact rational Krylov methods (CORK and TS-CORK), which allow for the efficient solution of large scale nonlinear eigenvalue problems. Numerical examples show that the presented framework is competitive with NLEIGS, usually produces smaller linearizations with the same accuracy but with little effort by the user.

Javier Perez-Alvaro
University of Montana
javier.perez-alvaro@mso.umt.edu

MS284

Reliable Computation of Certain Exterior Eigenvalues of Large Polynomial Eigenproblems

We propose to use functions of matrices to reliably compute certain exterior eigenvalues of large polynomial eigenvalue problems. These eigenproblems are linearized into companion forms, and matrix exponentials and trignometrics are used to transform the eigenvalues of largest real/imaginary part to the dominant eigenvalues of the corresponding functions of matrices. The action of functions of matrices on vectors are approximated by rational Krylov subspace methods (RKSM) with a single or double poles, and each step of RKSM can be performed on the original problem dimension, thanks to the recent development of the compact rational Krylov methods (CORK). Preliminary results show the robustness and efficiency of the new approaches.

Fei Xue
Clemson University
fxue@clemson.edu

Roel Van Beeumen
Lawrence Berkeley National Laboratory
rvanbeeumen@lbl.gov

MS285

BDDC Preconditioners for Matrix-free Higher-order Methods

In this talk, we present a lightweight memory Balancing Domain Decomposition by Constraints (BDDC) preconditioner for higher-order finite elements. By treating each element as its own subdomain, it is possible to isolate many of the computational and preconditioning challenges to individual elements. More standard preconditioning approaches can then be used for the resulting BDDC coarse problem which involves the assembly of lower-order elements. The approach avoids the use of large dense matrices and is designed to work together with a matrix-free approach. Numerical examples are presented for a variety of problem types.

Clark R. Dohrmann
Sandia National Laboratories
crdohrm@sandia.gov

Stefano Zampini
King Abdullah University of Science and Technology
stefano.zampini@kaust.edu.sa

MS285

Uniform Preconditioner for the Mass Matrix for Hp-Fem on Triangles

It is well-known that the stiffness matrix that arises in high order finite element discretisation of elliptic PDEs is generally ill-conditioned as the polynomial order p is increased. Although the mass matrix for low order (h-version) finite elements is well-conditioned, the mass matrix for the p-

version is, like the stiffness matrix, ill-conditioned as the order p is increased. For transient problems it is necessary to invert the mass matrix in order to carry out time-stepping. We present an algorithm for preconditioning the two dimensional mass matrix which results in a preconditioned system whose condition number is independent of the polynomial order and the mesh size. Although the preconditioner is applicable to any choice of basis for the high order polynomials, we show that the preconditioner can be implemented efficiently in the case where a Bernstein polynomial basis is chosen. Numerical examples are presented illustrating the performance of the algorithm for a range of challenging applications.

Mark Ainsworth
Division of Applied Mathematics
Brown University
Mark_Ainsworth@brown.edu

Shuai Jiang
Division of Applied Mathematics, Brown Univ
shuai_jiang@brown.edu

MS285

Matrix-free Algebraic Multigrid Preconditioners

Multigrid preconditioners are very popular preconditioners because of their convergence properties and scalability. Multigrid methods can be divided into geometric multigrid (GMG) and algebraic multigrid (AMG). For GMG, different coarser grids are built and therefore, GMG are excellent matrix-free preconditioners. The challenge is that GMG cannot be easily applied when the mesh is unstructured. AMG, however, uses the entries in the matrix of the system to build coarser grids. This makes AMG more flexible when the matrix is available but it limits their use as matrix-free preconditioners. We surpass the limitation of knowing the matrix entries by basing our AMG on the spectral AMGe method. Unlike other AMG, the spectral AMGe method does not use the matrix of the system. Instead it requires the eigenvectors of the operator evaluated on parts of the domain (agglomerates). This means that, similar to GMG, a mesh is required. However the constraints on how the agglomerates are built are minimal since we do not need to rediscretize the operator on the agglomerates. While at the fine level no matrix associated to the operator is assembled, we still assemble a matrix at the coarsest level in order to use a direct solver.

Bruno Turcksin, Wayne Joubert, Andrey Prokopenko, Damien Lebrun-Grandie
Oak Ridge National Laboratory
turcksinbr@ornl.gov, joubert@ornl.gov, prokopenkoav@ornl.gov, lebrungrandt@ornl.gov

MS286

Nonlinear Error Transport for Parabolic Problems in Mixed Precision

Numerical simulations of systems of differential equations are typically implemented with double precision arithmetic, to prevent roundoff error from becoming unacceptably large. In many cases, single or even half precision is sufficiently accurate and much more computationally efficient, but there are applications for which the resulting roundoff error can be too large. Strategies based on the defect correction process can be used for improving an initial approximate solution by computing the associated error correction. In the context of linear systems, work in the literature has shown that the iterative refinement technique can be used to systematically improve the numerical solution computed at a lower precision, by solving an equation for the error correction. In this work, we investigate the practical use of low precision arithmetic and error correction strategies, for nonlinear systems of equations. We evaluate the behavior of various types of time integration methods and applications, demonstrating performance gains while maintaining accuracy.

Dylan M. Copeland
Texas A&M University
copeland11@llnl.gov

Daniel Osei-Kuffuor
Lawrence Livermore National Laboratory
oseikuffuor1@llnl.gov

MS286

Experiments with Mixed Precision Algorithms in Linear Algebra

Low-precision floating-point arithmetic is a powerful tool for accelerating scientific computing applications, especially those in artificial intelligence. Here, we present an investigation showing that other high-performance computing (HPC) applications can also harness this power.

Jack J. Dongarra
University of Tennessee, Oak Ridge National Laboratory, USA
dongarra@icl.utk.edu

MS286

Exploiting Half Precision Arithmetic in Solving $Ax = b$

We consider solving a linear system $Ax = b$, with double precision A and b, by the use of an LU factorization computed in half precision and mixed-precision iterative refinement. For many matrices arising in practice, rounding A to half precision results in underflow and overflow. We describe how to use scaling to avoid overflow and underflow in the rounding and thereby to best exploit the limited dynamic range of half precision. We illustrate the benefits on a variety of test matrices.

Nicholas J. Higham
School of Mathematics
The University of Manchester
nick.higham@manchester.ac.uk

Srikara Pranesh
University of Manchester
srikara.pranesh@manchester.ac.uk

Mawussi Zounon

Univeristy of Manchester
mawussi.zounon@manchester.ac.uk

MS286

Simulated Half-precision Implementation of Blocked QR Factorization and Graph Clustering Applications

Development of new GPUs that perform half-precision arithmetic up to 16 times faster than double-precision arithmetic motivates use of low-precision computations whenever possible. In this work, we explore half-precision QR factorization for data and graph analysis applications. Our initial findings are that block QR algorithms are able to operate in half precision more robustly than non-block techniques. Since communication-avoiding, parallelizable QR algorithms already exist for tall-and-skinny matrices, we study how those algorithms behave in half-precision. While the standard Householder QR factorization algorithms are highly unstable in half-precision, our numerical simulations show that the Tall-and-Skinny QR (TSQR) algorithm often reduces the backward error of QR factorization. When using subspace iteration for graph clustering applications, half-precision accuracy in forming the eigenspace is sufficient for clustering with high precision and recall for some small-scale benchmark problems. We simulate half-precision arithmetic in our experiments in various ways that include conversions into single-precision for computation and half-precision for storage. These results motivate detailed numerical analysis of half-precision block QR factorization both for the purposes of replacing higher-precision QR (in applications less sensitive to error) and using the half-precision versions to produce warm starts that initialize higher-precision QR factorization.

Lucia Yang
University of Colorado Boulder
lucia.yang@colorado.edu

Geoff Sanders
Lawrence Livermore National Laboratory
sanders29@llnl.gov

MS287

Are Extreme Dissipation Events Predictable in Turbulent Fluid Flows?

We derive precursors of extreme dissipation events in a turbulent channel flow. Using a recently developed method that combines dynamics and statistics for the underlying attractor, we extract a characteristic state that precedes laminarization events that subsequently lead to extreme dissipation episodes. Our approach utilizes coarse statistical information for the turbulent attractor, in the form of second order statistics, to identify high-likelihood regions in the state space. We then search within this high probability manifold for the state that leads to the most finite-time growth of the flow kinetic energy. This state has both high probability of occurrence and leads to extreme values of dissipation. We use the alignment between a given turbulent state and this critical state as a precursor for extreme events and demonstrate its favorable properties for prediction of extreme dissipation events. Finally, we analyze the physical relevance of the derived precursor and show its robust character for different Reynolds numbers. Overall, we find that our choice of precursor works well at the Reynolds number it is computed at and at higher Reynolds number flows with similar extreme events.

Patrick J. Blonigan
NASA Ames Research Center
USRA
pblonig@sandia.gov

Mohammad Farazmand
MIT
mfaraz@mit.edu

Themistoklis Sapsis
Massachusetts Institute of Techonology
sapsis@mit.edu

MS287

Simplified Gentlest Ascent Dynamics for Saddle Point with Application to Allen-Cahn in Presence of Shear

The transition states on potential energy surface belong to a class of special saddle points having only one unstable direction. They play key roles in the understanding the rare events like phase transitions and noise-induced transitions escaping a stable point. We first review our work on theoretic and numerical progresses: the gentlest ascent dynamics (GAD) and the iterative minimization algorithm. We then present the recent work of the simplified GAD (joint with Shuting Gu) for the non-gradient system which reduces the cost of GAD by half. The application of the Allen-Cahen equation subject to the shear flow is presented to investigate the influence of the shear on the morphology of the transition state.

Xiang Zhou
Department of Mathematics
City Univeristy of Hong Kong
xiang.zhou@cityu.edu.hk

Shuting Gu
City University of Hong Kong
shutinggu2-c@my.cityu.edu.hk

MS287

Importance Sampling for Slow-fast Diffusions Based on Moderate Deviations

We consider systems of slow-fast diffusions with small noise in the slow component. We construct provably logarithmic asymptotically optimal importance schemes for the estimation of rare events based on the moderate deviations principle. Using the subsolution approach we construct schemes and identify conditions under which the schemes will be asymptotically optimal. Moderate deviations based importance sampling offers a viable alternative to large deviations importance sampling when the events are not too rare. In particular, in many cases of interest one can indeed construct the required change of measure in closed form, a task which is more complicated using the large devia-

tions based importance sampling, especially when it comes to multiscale dynamically evolving processes. The presence of multiple scales and the fact that we do not make any periodicity assumptions for the coefficients driving the processes, complicates the design and the analysis of efficient importance sampling schemes. Simulation studies illustrate the theory.

Konstantinos Spiliopoulos
Boston University
Department of Mathematics and Statistics
kspiliop@bu.edu

MS287

A Generalized Theory of Rare Event Simulation for Stochastic Differential Equations

We focus on improving state-of-the-art rare event simulation methods for dynamical systems. Current methods focus on using dynamic importance sampling and particle splitting algorithms. Both approaches have been found to be improved by considering large deviations theory, in which finding biasing functions or choosing importance sets for splitting hinges upon finding solutions and subsolutions of a Hamilton-Jacobi equation. This task becomes increasingly difficult in high dimensional settings or when the attractors in the dynamical system are strongly stable. In this talk we present a general theory for constructing efficient importance sampling estimators for high dimensional stochastic linear dynamical systems. To do so, we study how rare events occur in linear systems which will aid in constructing provably efficient biasing functions. Along the way we will explore connections with importance sampling estimators constructed via large deviations theory. We also explore connections with gentlest ascent dynamics, nonnormal transient growth, and optimal time dependent modes. This will also improve understanding of how rare events occur in nonlinear dynamical systems.

Benjamin J. Zhang
Massachusetts Institute of Technology
bjz@mit.edu

Tuhin Sahai
United Technologies Research Center
sahait@utrc.utc.com

Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu

MS288

Mori-Zwanzig Formalism and Discrete-time Modeling of Chaotic Dynamics

Nonlinear dynamic phenomena often require a large number of dynamical variables for their description, only a small fraction of which are of direct interest. Reduced models that use only the relevant dynamical variables can be very useful in such situations, both for computational efficiency and insights into the dynamics. This talk concerns a discrete-time, parametric approach to model reduction based on the NARMAX (Nonlinear Auto-Regressive Moving-Average with eXogenous inputs) representation of stochastic processes [Chorin-Lu 2015]. Previous work have demonstrated the effectiveness of the NARMAX method on a variety of concrete examples. Here, I will examine the NARMAX representation within from the point of view of the Mori-Zwanzig formalism, and show that NARMAX arises from a specific choice of projections with nice properties. The ideas will be illustrated on a prototypical model of spatiotemporal chaos.

Kevin K. Lin
Department of Mathematics
University of Arizona
klin@math.arizona.edu

Fei Lu
Johns Hopkins University
feilu@math.jhu.edu

Alexandre Chorin
Department of Mathematics
University of California at Berkeley
chorin@math.berkeley.edu

MS288

Coarse-grained Modeling of Polymers in Solvent Via Mori-Zwanzig Formalism

While atomistic simulations allow for precise reconstruction of molecular structures and can capture all the atomistic details of a molecular system, they are computationally prohibitive to predict large-scale effects or long-time behaviors. When only mesoscopic structures of molecules or their coarse-grained properties are of practical interest, it may not be necessary to explicitly take into account all the atomistic details of materials. And hence, with less degrees of freedom, the mesoscopic modeling captures the observable properties of systems on larger spatial and temporal scales beyond the capacity of atomistic simulations. In this talk, we will focus on polymer solution systems and establish their coarse-grained models derived from the Mori-Zwanzig formalism. We will present our findings in the assessment of the accuracy of the coarse-grained modeling to conserve the solvent-mediated dynamical properties as well as the static properties of the underlying atomistic systems.

Wenxiao Pan
University of Wisconsin Madison
wpan9@wisc.edu

Shu Wang
University of Wisconsin-Madison
swang597@wisc.edu

MS288

Renormalized Reduced Order Models for Long Term Prediction

In recent work we have developed a renormalization framework for stabilizing reduced order models for time-dependent partial differential equations. In this talk, we will present this new renormalization method for construct-

ing reduced order models. We will then discuss its application to the open problem of finite-time singularity formation (blow-up) for the 3D Euler's equations of incompressible fluid flow. To the best of our knowledge this is the first time-dependent perturbative renormalization approach for 3D Euler which includes all the complex effects present in the Euler dynamics. For the Taylor-Green initial condition, the renormalized coefficients in the reduced order models decay algebraically with time and resolution. The renormalized reduced models are stable and we evolve them for long times. Our results for the behavior of the solutions are consistent with the formation of a finite-time singularity.

Jacob Price
University of Washington
jrprice@pugetsound.edu

Panos Stinis
Pacific Northwest National Laboratory
panagiotis.stinis@pnnl.gov

MS288

On the Estimation of the Mori-Zwanzig Memory Integral

In this talk, we will present our latest research on the Mori-Zwanzig equation. Specifically, a prior estimation of the MZ memory and its series expansion will be established through operator analysis. The new theory enables us to derive computable upper bounds and presumably convergent series expansion schemes of the MZ memory integral. Numerical applications will be presented for both linear and nonlinear dynamical systems.

Yuanran Zhu
University of California, Santa Cruz
yzhu22@ucsc.edu

MS289

Modeling the Acoustic Response of Elastic Targets in a Layered Medium Using the Coupled Finite Element/Boundary Element Method

The fluid-structure interaction technique provides a paradigm for solving scattering from elastic targets embedded in a fluid by a combination of finite and boundary element methods. In this technique, the finite element method is used to compute the target's elastic response and the boundary element method with the appropriate Green's function is used to compute the field in the exterior medium. The two methods are coupled at the surface of the target by imposing the continuity of pressure and normal displacement. This results in a boundary element equation that can be used to compute the scattered field anywhere in the surrounding environment. This method reduces a finite element problem to a boundary element one with drastic reduction in the number of unknowns, which translates to a significant reduction in numerical cost. In this presentation, the technique will be outlined, and it will be applied to compute scattering from various targets in complex ocean environments.

Ahmad Abawi
HLS Research
abawi@hlsresearch.com

Petr Krysl
Department of Structural Engineering
University of California, San Diego
pkrsyl@ucsd.edu

MS289

Identifying Acoustic Scattering Mechanisms for Tilted Solid Cylinders in Water by Comparing Ray Theory with Multi-domain Processing

We apply signal-processing methods discussed in Plotnick, Marston & Marston [J. Acoust. Soc. Am. 136, EL61–EL66, (2014)] to a relatively complicated situation of a bimetallic solid cylinder and we interpret results with guided-wave ray theory: Gipson & Marston [J. Acoust. Soc. Am. 107, 112–117, (2000)].

Timothy Daniel, Philip Marston
Washington State University
Physics & Astronomy Dept.
t.derekdaniel@gmail.com, marston.phil@gmail.com

MS289

Multistatic Localization with Uncertain Data Association

Real world data often lacks perfect knowledge of association between a signal source and its representation at a receiver site. Clutter, noise, and other (un-modeled) sources all contribute to the degradation of a clean data set. To understand the degradation in performance of a multistatic localization estimator, the application is modeled as a Maximum Likelihood (ML) problem and the underlying Cramer-Rao Lower Bound (CRLB) is developed to understand the best case estimator behavior that is achievable. Simulations of the ML estimator are conducted over a range of false-to-valid contact ratios and the departure from CRLB performance is studied.

Jamie Irza
Raytheon Company
Integrated Defense Systems
jamie.irza@raytheon.com

MS289

Improving Sonar Training Systems using WaveQ3D

Existing undersea acoustic propagation models often fail to meet the speed and accuracy needs of sonar training simulation/stimulation systems for littoral environments. In these environment, propagation is highly dependent on interactions with the ocean bottom. WaveQ3D improves the speed of undersea acoustic transmission loss calculations, without sacrificing accuracy, by implementing a 3D Gaussian beam model in the same geodetic coordinates used by the underlying environmental databases. This research was supported by the High Fidelity Active Sonar Training (HiFAST) Project at the U.S. Office of Naval Research.

Sean Reilly

Raytheon BBN Technologies
sean.m.reilly@raytheon.com

MS290

Scalable Convergence using Deflation Preconditioning for the Helmholtz Equation

Abstract not available

Vandana Dwarka
TU Delft
vandana.dwarka@gmail.com

MS290

Field-of-values Analysis of a Two-level Shifted Laplace Preconditioner for Finite Element Helmholtz Problems

One of the main tools for solving linear systems arising from the discretization of the Helmholtz equation is the shifted Laplace preconditioner, which results from the discretization of a perturbed Helmholtz problem $-\Delta u-(k^2+i\varepsilon)u=f$ where $\varepsilon>0$ is an absorption parameter. We extend the two-level shifted Laplacian first proposed in [Y. Erlangga and R. Nabben, ETNA, 31 (2008), 403-424] and further simplified in [A. Sheikh, D. Lahaye and C. Vuik, NLAA, 20 (2013), 645-662] to Helmholtz problems discretized with linear finite elements, and prove estimates of the field of values of the preconditioned system that are used to analyze the convergence of GMRES. Our main result is that GMRES applied to the two-grid preconditioned system with a shift $\varepsilon=O(k^2)$ converges in a number of iterations independent of the wavenumber k if the coarse mesh size H satisfies $Hk^2\leq C$, for some constant C depending on the domain but independent of the wavenumber k. We show in experiments that wavenumber-independent convergence of GMRES also holds for pollution-free meshes where the coarse mesh size satisfies $Hk^{3/2}\leq C$. This behaviour is sharply different to the standalone shifted Laplacian, for which wavenumber-independent GMRES convergence has been established only under the condition $\varepsilon=O(k)$ in [M.J. Gander, I.G. Graham and E.A. Spence, Numer. Math., 131 (2015), 567-614].

Luis Garcia Ramos
TU Berlin
garcia@math.tu-berlin.de

Reinhard Nabben
Institute fur Mathematik
Berlin, Germany
nabben@math.tu-berlin.de

MS290

A Hybrid Domain Decomposition and Shifted Laplacian Multigrid Preconditioner for the Elastic Helmholtz Equation

The elastic Helmholtz equation is used to model wave propagation in the frequency domain in cases where the underlying medium is elastic. Such a case arises in seismic imaging applications, where the earth's subsurface is the elastic medium. Discretizing the elastic Helmholtz equation leads to a huge indefinite linear system which is hard to solve due to its size, indefiniteness, and properties coming from its complicated physics. In this work we propose a hybrid domain decomposition and shifted Laplacian multigrid method to solve the equation in distributed computation settings at large scales. The domain decomposition approach is the most natural way to distribute a numerical solution of almost any PDE, while the shifted Laplacian multigrid is one of the most efficient methods to solve the elastic Helmholtz equation in a single computing node. Here we use the shifted Laplacian method to approximately solve the local sub-domain problems. Thanks to the relatively low memory footprint of the multigrid method, the combination of multigrid and domain decomposition allows us to use relatively large sub-domains, which leads to favorable convergence properties of the method. We show numerical experiments for elastic problems with heterogeneous media in two and three dimensions.

Eran Treister
Ben Gurion University of the Negev
erant@bgu.ac.il

MS290

Fast Contour Integral Preconditioner for Solving 3D High-frequency Helmholtz Equations

In this talk, we propose an iterative solution method for the 3D high-frequency Helmholtz equation. In a contour integration framework, the solution in certain invariant subspace is approximated by solving problems with complex shifts, and this accelerates GMRES iterations by restricting the spectrum. We construct a polynomial fixed-point iteration for solving the shifted problems, which is robust even if the magnitude of the shifts is small. Numerical tests in 3D show that $O(n^{1/3})$ matrix-vector products are needed for solving a high-frequency problem with matrix size n to high accuracy. The method has little storage requirement, can be applied to both dense and sparse linear systems, and is suitable for parallel computing.

Yuanzhe Xi
University of Minnesota
yxi26@emory.edu

Xiao Liu
Rice University
xiao.liu@rice.edu

Yousef Saad
University of Minnesota
saad@umn.edu

Maarten de Hoop
Rice University
mdehoop@rice.edu

MS291

HPC Oriented Cloud Infrastructure for Digital Twins

Implementation of a physical model based digital twin requires, in many cases plenty of offline HPC simulation of

the physical system for model reduction purposes. This needs the efficient and easy use of supercomputers or at least big clusters for the simulations. In this talk we present a workflow constructed in the MSO4SC HPC oriented cloud infrastructure, which has been developed from the European Commission funded H2020 project. MSO4SC runs singularity containers for simulations. We apply the workflow for a POD in an urban air pollution use case.

Zoltán Horváth
Department of Mathematics
Széchenyi István University, Győr, Hungary
horvathz@sze.hu

Tamás Budai, kos Kovács
Szechenyi Istvan University, Hungary
budai.tamas@sze.hu, kovacs.akos@sze.hu

MS291

What Could Be Digital Twins for the Insurance Industry? Example of the Sensibility Analysis of the Solvency Capital Requirement

The Solvency Capital Requirement (SCR) is fundamental for the insurance business and the regulators. It depends on several parameters subject to some uncertainty. Sensitivity analysis may help in the quantification of the uncertainty of parameters on the variability of the SCR. Sensitivity analysis is well adapted when the inputs are independent and for central variability. We propose to use Shapley index (adapted when inputs are dependent) and Quantile Oriented Sensitivity Analysis, more adapted to the SCR (which is indeed a quantile). We shall focus on estimation and computational issues. This scheme is related to Digital Twins in the sense that the computation of the SCR mimic the behavior of an insurance portfolio.

Veronique Maume-Deschamps
University of Lyon/AMIES
veronique.maume@univ-lyon1.fr

Kevin Elie dit Cosaque
University of Lyon / SCOR
kelie-dit-cosaque-external@scor.com

MS291

Towards Digital Tokamaks

Numerical simulations form an indispensable tool to understand the behavior of a hot plasma that is created inside a tokamak for providing nuclear fusion energy. Various aspects of tokamak plasmas have been successfully studied through the reduced magnetohydrodynamic (MHD) model. The need for more complete modeling through the full MHD equations is addressed here. A computational method for the full MHD equations is quickly presented. Several well-known tokamak-plasma instabilities are simulated to demonstrate the capabilities of the computational method. The linear growth rate of an internal kink mode and a tearing mode are benchmarked against the results of a linear MHD code. The evolution of a tearing mode and the resulting magnetic islands are simulated well into the nonlinear regime. The results are compared with predictions from the reduced MHD model. Finally, a simulation of a ballooning mode illustrates the possibility to use our method as an ideal MHD method without the need to add any physical dissipation.

Barry Koren
Department of Mathematics and Computer Science
Eindhoven University of Technology
b.koren@tue.nl

MS291

Towards Digital Twins for the Distribution Management in Energy Transport Networks

The control of energy transport through gas or power networks in Germany is more and more characterized by complex decisions that need support by digital twins. There are two main challenges. First, the increasing percentage of renewable energies leads to higher fluctuations of available energy with respect to time and location. Secondly, energy carriers are not supplied near by the locations with highest demands but at locations where gas or power is available at lowest price. We present recent developments towards a digital twin based on a hierarchical modeling and simulation approach that prognosts the load distribution and available capacities in the network. The models range from simple algebraic descriptions to hyperbolic partial differential equation systems with algebraic and stochastic constraints. We demonstrate the implemented simulations within a 24h time window for two benchmark networks. Finally, we discuss briefly current and future issues with respect to optimization tools for designing and controlling energy networks.

Caren Tischendorf
Humboldt University Berlin
tischendorf@math.hu-berlin.de

MS292

Controlled Experimentation at Scale

The Internet provides developers of connected software, including web sites, applications, and devices, an unprecedented opportunity to accelerate innovation by evaluating ideas quickly and accurately using trustworthy controlled experiments (e.g., A/B tests). From front-end user-interface changes to backend recommendation systems and relevance algorithms, from search engines (e.g., Google, Microsofts Bing, Yahoo) to retailers (e.g., Amazon, eBay, Netflix, Etsy) to social networking services (e.g., Facebook, LinkedIn, Twitter) to Travel services (e.g., Expedia, Airbnb, Booking.com) to many startups, online controlled experiments are now utilized to make data-driven decisions at a wide range of companies. The deployment and mining of online controlled experiments at scale (e.g., hundreds of experiments run every day at Bing) and deployment of online controlled experiments across dozens of web sites and applications has taught us many lessons. We provide an introduction, share real examples, key lessons, and cultural challenges.

Jordan Atlas
Analysis & Experimentation Team, Microsoft

jordanat@microsoft.com

MS292

Designing HPC Applications for Exascale Supercomputers and Beyond

I will describe hardware-software co-design in the context of high-performance computing (HPC) applications with specific examples from computational chemistry. The methods and tools for co-design of systems and processors will be discussed. Suggestions for how the the HPC application developer community can contribute to the vendor co-design process will be given.

Jeff R. Hammond
Intel Corporation
Parallel Computing Lab
jeff_hammond@acm.org

MS292

Communication-avoiding Distributed-memory Algorithms for QR Factorization

Scalable QR factorization algorithms for solving least squares and eigenvalue problems are critical given the increasing parallelism within modern machines. We provide a more general parallelization of the CholeskyQR2 algorithm. This algorithm executes over a 3D processor grid, the dimensions of which can be tuned to trade-off costs in synchronization, interprocessor communication, computational work, and memory footprint. We implement this algorithm, achieving up to a factor of $\Theta(P^{1/6})$ less interprocessor communication than any previous parallel QR implementation. Our performance study on Intel Knights-Landing and Cray XE supercomputers demonstrates that this QR factorization method can achieve better absolute performance and parallel scalability than ScaLAPACK's QR.

Edward Hutter
Illinois
hutter2@illinois.edu

MS292

Computer Security: The Science of the Impossible

Multiple problems in computer security are intractable. We give several new theorems for undecidable security problems, meaning that these problems have no algorithmic solutions. These proofs allow us to describe computer security as 'impossible'. Computer security is necessary, due to the ubiquity of modern digital computers. In spite of the algorithmic intractability of problems in computer security, we must find practical solutions to these problems. Furthermore, we would like those practical solutions to have some performance guarantees.

Brent Kirkpatrick
Intrepid Net Computing
bbkirk@intrepidnetcomputing.com

MS293

On the Development and Implementation of Optimized, High-order Time Integrators for Multi-physics Problems

Modern multi-physics applications present numerous challenges for legacy numerical methods, including the presence of multiple processes that act on distinct time scales. In many scenarios, each of these processes, as well as the couplings between them, must be tracked accurately, for an accurate and stable solution to the multi-physics problem. In this talk, we focus on our recent work in extending the ARKode time integration library to enable high-order methods for multi-physics applications. In particular, we will discuss our recent enhancements to support implicit/explicit and multi-rate splittings of the dynamical processes, enabling fine-tuning of the integration methods to the processes under consideration.

Daniel R. Reynolds
Southern Methodist University
Mathematics
reynolds@smu.edu

David J. Gardner
Lawrence Livermore National Laboratory
gardner48@llnl.gov

MS293

Devising Unconditionally Stable Multistep ImEx Schemes that Avoid Stiff Nonlinear Implicit Terms

In this talk we devise unconditionally stable multistep ImEx schemes in problems where both the implicit, and explicit terms are stiff. Unconditional stability is a desirable property for a numerical scheme as it implies the absence of a (stiff) time step restriction. One particular application where such an approach may be advantageous is in nonlinear problems, where a (simple) implicit term is taken to be a constant coefficient operator, and the stiff nonlinear terms are treated explicitly. This then bypasses the need for nonlinear solvers. We first use the new stability theory to explain the fundamental stability restrictions of the well-known semi-implicit backward differentiation formulas (SBDF). We then show, using the new theory, how to overcome the limitations of SBDF to obtain higher order schemes (r ¿ 2). Using this insight, rigorous, unconditionally stable schemes are devised for the linear variable coefficient diffusion problem. We will then use the linear results to show that they can be used to avoid the implicit treatment of nonlinear terms in some nonlinear diffusion problems. Numerical examples will be presented.

David Shirokoff
New Jersey Institute of Technology
david.g.shirokoff@njit.edu

Benjamin Seibold, Dong Zhou
Temple University

seibold@temple.edu, dzhou11@calstatela.edu

MS293

A Family of Second and Third Order Implicit-explicit Runge-Kutta Methods for Stiff Time-dependent Partial Differential Equations

Nonhydrostatic atmosphere models generate vertically propagating acoustic waves that, while being unimportant for accurate weather and climate simulations, result in a CFL restriction for standard explicit methods that is unacceptably small. Fully implicit methods can avoid this step-size restriction, but at a much higher computational cost. Implicit-explicit (IMEX) methods are a compromise that allows larger stable step-sizes than standard explicit methods and at a lower computational cost than fully implicit methods. In this talk we discuss the derivation of a family of low-storage second and third order IMEX Runge-Kutta (RK) methods for the time-integration of partial differential equations with a partitioning into stiff and non-stiff components. The explicit part of the method is constrained so that its stability region is optimal on either the imaginary or real axis and L-stability is enforced for the implicit part. The number of explicit internal stages can be set arbitrarily high to increase the size of the stability region of the explicit part of the method and the number of implicit stages can be adjusted to improve stability or reduce the number of implicit solves. A subset of these methods are implemented for the solution of the nonhydrostatic atmosphere dynamic core HOMME-NH and their performance is tested against other IMEX Runge-Kutta methods.

Andrew J. Steyer
Sandia National Laboratories
asteyer@sandia.gov

Christopher J. Vogl
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
vogl2@llnl.gov

Mark A. Taylor
Sandia National Laboratories
mataylo@sandia.gov

Oksana Guba
Sandia National Laboratories
NM, USA.
onguba@sandia.gov

MS293

Adaptive Krylov Based Time Integration Methods

The Rosenbrock-Krylov family of time integration schemes is an extension of Rosenbrock-W methods that employs a specific Krylov based approximation of the linear system solutions arising within each stage of the integrator. In this talk we will briefly introduce Rosenbrock-Krylov methods as well as present an extension of these methods to address stability questions which arise for methods making use of inexact linear system solutions. Rosenbrock-Krylov methods employing this extension show a substantial improvement in computational efficiency relative to prior implementations.

Paul Tranquilli
Virginia Polytechnic Institute and
State University
ptranq@vt.edu

MS294

Multilevel Computational Design and Analysis of Logistics Networks

Supply chains are complex supply and demand networks critical to product manufacture, support and development. Complex interdependencies and uncertainties throughout the global supply chain have the potential to increase the risk of systemic failure making the development of decision support tools critical for risk visibility, quantification, minimization, and mitigation. Although commercial-off-the-shelf (COTS) software enables supply chain simulation quantitative risk characterization remains challenging. COTS software is typically not scalable for detailed scenario planning and exploration under uncertainty let alone for design for system agility, resilience, and adaptability. The potential benefits of simulation-enabled situational awareness in supply chain are manifold: opportunities for early alerting and driver identification, inventory optimization, and reduced KPI recovery time following disruptions. We describe an approach to provide supply chain situational awareness based around a data-driven, model-based simulation capability. Multilevel Monte Carlo (MLMC) is used to orchestrate multi-fidelity stochastic supply chain simulation; scalable global sensitivity analysis (GSA) for KPI driver identification is realized by use of this multi-fidelity simulator. Supply chain risk management is accomplished by coupling (in an iterative loop) efficient simulation with optimization and reliability quantification to plan for evolving operational parameters.

Quan Long, Nai-Yuan Chiang, Thomas A. Frewen, Yiqing Lin
United Technologies Research Center
longq@utrc.utc.com, chiangn@utrc.utc.com, frewenta@utrc.utc.com, liny@utrc.utc.com

MS294

Multilevel Ensemble Kalman Filter

The ensemble Kalman filter (EnKF) is a sequential filtering method that uses an ensemble of particle paths to estimate the means and covariances required by the Kalman filter by the use of sample moments, i.e., by the MC method. EnKF is often both robust and efficient, but its performance may suffer in settings where the computational cost of accurate forward simulations of particles is high. We will present ideas and numerical results on combining MLMC and EnKF to construct the multilevel ensemble Kalman filter (MLEnKF) for settings of either finite or infinite dimensional state and observation spaces.

Haakon Hoel
Center for Uncertainty Quantification
KAUST
hakon.hoel@kaust.edu.sa

Raul F. Tempone
Mathematics, Computational Sciences & Engineering
King Abdullah University of Science and Technology
raul.tempone@kaust.edu.sa

Kody Law
University of Manchester
kodylaw@gmail.com

Fabio Nobile
EPFL, Switzerland
fabio.nobile@epfl.ch

Alexey Chernov
Institut für Mathematik
Carl von Ossietzky Universität Oldenburg
alexey.chernov@uni-oldenburg.de

MS294

Big Data-based Analytics for Supercomputer Logs: Architecture and Application

Todays supercomputers are heavily instrumented generating a huge volume of logs that are intended to boost reliability, availability and serviceability (RAS) of the system. These data are generated from multiple logging systems and sensors that each portrays status of certain components (both software and hardware). A close investigation on these data, therefore, can provide invaluable information regarding system status, performance, and resource utilization. However, the analysis of these data faces two challenges. First, the sheer volume of data is beyond what manual inspection can offer. Second the heterogeneous nature of log contents, which are also of irregular structure, i.e., combination of numbers, texts, or hexadecimal codes, are hard to apply conventional analytic workflows. To address the issues, we developed a Big Data analytics framework named LogSCAN that is specifically designed for HPC log data at the Oak Ridge National Laboratory (ORNL). In this presentation, we describe the architecture of the backend database, the analytic engine, the computation subsystem, and the frontend for a user access. We will also introduce several analytic examples that utilize LogSCAN to understand years of log data collected from ORNLs Titan supercomputer.

Byung Hoon Park
Computer Science and Mathematics Division
Oak Ridge National Laboratory
parkbh@ornl.gov

MS294

High-dimensional Bayesian Inversion with Generative and Invertible Neural Networks

For high-dimensional inverse problems with sparse and noisy observations and complex forward models, a strong prior is often necessary to allow us to compute a solution. In this work, a generative model based on generative adversarial networks is trained to represent the prior of high-dimensional input with a low-dimensional latent variable. To further characterize the uncertainty of this challenging inversion, an invertible neural network is trained between the latent variable and the sparse observations which are also low-dimensional. Inversion of channelized permeability in porous media flows is studied to show the effectiveness of the proposed method.

Yinhao Zhu, Nicholas Zabaras
University of Notre Dame
yzhu10@nd.edu, nzabaras@gmail.com

MS295

A Stein Variational Newton Method

Stein variational gradient descent (SVGD) was recently proposed as a general purpose nonparametric variational inference algorithm [Liu & Wang, NIPS 2016]: it minimizes the Kullback–Leibler divergence between the target distribution and its approximation by implementing a form of functional gradient descent on a reproducing kernel Hilbert space. In this paper, we accelerate and generalize the SVGD algorithm by including second-order information, thereby approximating a Newton-like iteration in function space. We also show how second-order information can lead to more effective choices of kernel. We observe significant computational gains over the original SVGD algorithm in multiple test cases.

Gianluca Detommaso
University of Bath
g.detommaso@bath.ac.uk

Tiangang Cui
Monash University
tiangang.cui@monash.edu

Youssef M. Marzouk, Alessio Spantini
Massachusetts Institute of Technology
ymarz@mit.edu, spantini@mit.edu

Robert Scheichl
University of Heidelberg
r.scheichl@uni-heidelberg.de

MS295

Scable Learning and Inference with Optimal Transport and Steins Method

Stein variational gradient descent (SVGD) is a generalization of gradient descent for Bayesian inference. It iteratively transports a set of particles to match the target distribution, by applying a form of functional gradient descent of KL divergence that reveals a connection to Stein's method. SVGD provides a new tool for Bayesian inference and uncertainty quantification, which inherits the efficiency of gradient-based optimization, with the flexibility of providing flexibility nonparametric particle based approximation to complex distribution. We will discuss various practical applications of SVGD, theoretical justification from various angles (including both an interacting particle view and moment matching view) and various algorithmic extensions.

Qiang Liu
University of Texas at Austin

lqiang@cs.utexas.edu

Jian Peng
University of Illinois at Urbana-Champaign
Computer Science
jianpeng@illinois.edu

Jason Lee
University of Southern California
jasonlee@marshall.usc.edu

MS295

Layered Transport Maps for Conditional Sampling

For large models, algorithms for exploring Bayesian posterior distributions (e.g., Markov chain Monte Carlo) can be prohibitively expensive for real time or time-critical applications. Transport maps, which are nonlinear random variable transformations, can be constructed offline (before collecting data) and thus have the potential to significantly reduce the online (after collecting data) cost of Bayesian inference. However, transport maps can be difficult to construct when either the data or parameters are high dimensional. In this talk, we will discuss a new approach for overcoming this with layers of lower triangular transport maps. Applications in sea ice modeling will be used to highlight the strengths and weaknesses of this approach.

Matthew Parno
US Army Corp of Engineers, USA
Matthew.D.Parno@usace.army.mil

MS296

Lattice Boltzmann Model for Time Sub-diffusion Problem in Caputo Sense

Anomalous diffusions, including subdiffusion and superdiffusion, are usually encountered in many diverse applications in science and engineering. Although many numerical methods have been proposed to study anomalous diffusion problems that are modeled by fractional advection-diffusion equations, in this paper, a fresh lattice Boltzmann (LB) model for time sub-diffusion equation in Caputo sense is proposed. Through the Chapman-Enskog analysis, the time-fractional diffusion equation can be recovered from the developed LB model. In addition, we also test the present LB model through some problems, and find that the numerical results agree well with the analytical solutions to these problems.

Rui Du
southeast university
rdu@seu.edu.cn

MS296

Spectral Methods for Fractional Advection-diffusion-reaction Equations on Smooth Domains

We propose a domain decomposition spectral method for fractional advection-diffusion-reaction equations on arbitrary smooth domains. We first consider the fractional advection-diffusion-reaction equations in a ball and then embed irregular domain into balls. Numerical results for problems on various smooth domains will be presented.

Zhaopeng Hao
Mathematical Sciences of Worcester Polytechnic Institute
zhao@wpi.edu

Marcus Sarkis
Worcester Polytechnic Institute
Instituto de Matematica Pura e Aplicada (Brazil)
msarkis@wpi.edu

Zhongqiang Zhang
Worcester Polytechnic Institute
zzhang7@wpi.edu

MS296

A Fast Temporal Second-order Difference Scheme for the Fractional Wave Equation

In this paper, a fast second-order accurate difference scheme is proposed for solving the fractional wave equation. Combing the order reduction technique with FL2-1$_\sigma$ formula, which employs sum-of-exponential approximation to the kernel function appeared in Caputo derivative, we approximate the Caputo derivative. This method can reduce the storage and computational cost significantly. The unique solvability and unconditional convergence of the resulting difference scheme are proved by energy method. The numerical examples are presented to verify the theoretical results.

Hong Sun
Jiangsu university of science and technology
sunhongzhal@126.com

MS296

Solving Fractional Partial Differential Equations using a Deep Neural Network

We approximate the solution of fractional partial differential equations using multilayer artificial neural networks. The effectiveness of this method will be demonstrated by the numerical experiments.

Hong Wang
University of South Carolina
hwang@math.sc.edu ¡hwang@math.sc.edu¿

MS297

Microswimmer Design: Pullers and Pushers from Bimetallic Rods

Micro-rods with gold and platinum parts powered by hydrogen peroxide are a well established example of artificial micro-swimmers. We show that using different ratios of gold and platinum the rods behave like pullers, pushers or symmetric swimmers. We use a Rigid Multiblob Algorithm and experiments to study the dynamics of this family of rods in shear flow and in a periodic lattice of obstacles. The nature of the swimmers (puller, symmetric or pusher) and their level of activity (Peclet number) affect the dynamics in both scenarios. At intermediate shear rates the rods show rheotaxis and swim upstream. In the lattice, the

rods orbit around the obstacles which reduces their effective diffusion at long times. The nature of a swimmer can be inferred from its trajectory in these setups if the Peclet number is high enough to overcome the Brownian motion inherent to micro-swimmers.

Florencio Balboa Usabiaga
Flatiron Institute
Simons Foundation
fbalboausabiaga@flatironinstitute.org

Quentin Brosseau
New York University
qb3@nyu.edu

Enkeleida Lushi
Flatiron Institute
Simons Foundation
lushi@njit.edu

Yang Wu
Molecular Design Institute
New York University
yw2077@nyu.edu

Leif Ristroph
Courant Institute
New York University
ristroph@cims.nyu.edu

Jun Zhang
Courant Institute
jun@CIMS.nyu.edu

Michael D. Ward
Molecular Design Institute
New York University
mdw3@nyu.edu

Michael J. Shelley
New York University
Courant Inst of Math Sciences
shelley@cims.nyu.edu

MS297

Dynamics of Colloids Above a Bottom Wall Driven by Active Torques and Forces

I will describe numerical methods to study the dynamics of suspensions of colloids sedimented above a bottom wall and driven by externally-applied forces (sedimentation) or torques (microrollers). The proximity of the boundary controls the collective dynamics of these active suspensions. Recently a new instability has been observed experimentally and numerically: the fingering of a front of suspended microrollers near a floor. A continuum model shows that this instability is linear and that the size scale selection arises only from hydrodynamic interactions between the particles and the wall. From these fingers, long-lived compact motile structures, called 'critters', can be formed just with hydrodynamic interactions. The presence of a nearby no-slip boundary strongly affects the structures emerging in sedimenting colloidal suspensions. The suspensions first forms a monolayer with a dense traveling front, which can be described using a simple one dimensional nonlocal PDE. The front then transitions into finger-like structures whose width depends on the particle size and height from the floor.

Aleksandar Donev
Courant Institute of Mathematical Sciences
New York University
donev@courant.nyu.edu

MS297

A Coarse-grained Model for Beds of Elastic Fibers

An important class of fluid-structure problems involve the dynamics of ordered arrays of immersed, flexible fibers. While specialized numerical methods have been developed to study fiber-fluid systems, they become infeasible when there are many, rather than a few, fibers present, and do not lend themselves to analytical calculation. In this talk I will introduce a coarse-grained continuum model, based on local-slender body theory for elastic fibers immersed in a viscous Newtonian fluid. We will explore some of the basic properties of these systems subjected to steady and oscillatory shear flows; and then show how qualitatively different phenomenon can emerge in some systems as the fiber density is varied. Finally, we will show how the model can be used to study pumping in beds of actuated cilia.

David Stein
Flatiron Institute
Simons Foundation
dstein@flatironinstitute.org

Michael J. Shelley
New York University
Courant Inst of Math Sciences
shelley@cims.nyu.edu

MS298

Efficient High-order Discretization with MFEM

In this talk, we present an overview of the modular finite element methods (MFEM) library (mfem.org), including its main abstraction classes, their corresponding linear algebra objects, and implementation variants. We discuss the components required for the construction and application of general finite element discretization operators and the various choices for their software implementation, e.g. as a single assembled parallel CSR matrix, or as a product of linear operators, i.e. "matrix-free" representations. We highlight the pros and cons of the various choices based on the discretization parameters such as solution space order, mesh order, choice of quadrature, etc., and present numerical illustration with MFEM examples. We also report on the progress of the ongoing efforts to implement efficiently and integrate seamlessly support for architectures with accelerators (e.g. GPUs) in the library. The behavior of the various algorithms with respect to scalability on HPC systems is also studied numerically and discussed. Prepared by LLNL under Contract DE-AC52-07NA27344.

Veselin Dobrev
Lawrence Livermore National Laboratory

dobrev1@llnl.gov

MS298

libParanumal: GPU Accelerated Scalable High-order Finite Element Flow Solvers

In this talk, we briefly introduce our recently released library, libParanumal which offers a experimental test bed used to develop scalable multi-GPU finite element simulation tools. We discuss nearly incompressible single phase and incompressible multiphase flow solvers developed using libParanumal core. In the nearly incompressible flow solver, Boltzmann equations are discretized with Hermite polynomials in velocity space yielding a first order system. A stabilized unsplit perfectly matching layer (PML) formulation and a multirate semi-analytic time stepping method are introduced to avoid excessive domain truncation and time step restriction in small relaxation times, respectively. The multiphase flow solver uses the phase field method to track the interface between immiscible fluids and utilizes an efficient splitting scheme leading to time-independent coefficient matrices for all flow variables in coupled Navier-Stokes Cahn-Hillirad equations. Accuracy and performance of the solvers for various flow regimes are illustrated through three dimensional test problems.

Ali Karakus
Argonne National Laboratory
Argonne National Laboratory
alikarakus.me@gmail.com

Noel A. Chalmers
AMD Research
noel.chalmers@gmail.com

Tim Warburton
Virginia Tech
tcew@vt.edu

MS298

Code for Efficient Extensible Discretizations LibCEED

The Center for Efficient Exascale Discretrizations comprises multiple research teams focused on the development of high-order methods for forthcoming exascale computers. In addition to distributed memory parallel computing, which yields anywhere from thousand- to million-fold speedups, the project focuses on high performance at the single node level, particularly for accelerator-based platforms. To simplify optimization on these complex nodes, the team is developing a low-level API, libCEED, that will deliver performance for the principel kernels common to high-order finite element codes. These high-order codes trade flops for data; in many applications higher order translates into fewer gridpoints and hence less data movement for the target accuracy. Here, we present recent developments in the CEED and paranumal libraries as applied to spectral element solutions of the time-dependent incompressible Navier-Stokes equations. libParanumal, developed by the Warburton group at Virginia Tech, is a vanguard application for high-order methods on GPUs. With many of the kernels (e.g., Jacobi-preconditioned conjugate gradient solvers for the Helmholtz problems) are achieving 1-2 TFLOPS per GPU on the Nvidia V100, its performance sets the bar for CEED library development. We describe how the libparanumal kernels are translated into the libCEED library and how we have extended libparanumal to support several of the solver features found in the spectral element code Nek5000.

Thilina Rathnayake
University of Illinois
thilinarmtb@gmail.com

Paul Fischer
University of Illinois at Urbana-Champaign
fischerp@illinois.edu

MS298

A Novel Approach to Using Fused Batched BLAS to Accelerate High-order Discretizations

We present a set of Batched BLAS optimizations, current efforts to standardize Batched BLAS APIs, and a Batched BLAS API extension and its use in accelerating high-order methods on new architectures. Expressing scientific computations in terms of BLAS, and in particular the general dense matrix-matrix multiplication (GEMM), has been of fundamental importance for obtaining high performance portability across architectures. We show how this can be used for high-order methods to compute tensor contractions. The approach is based on developing and exposing to users capabilities for easily fusing Batched BLAS calls. In particular, we developed device (GPU) interfaces to Batched BLAS that can be used to construct custom/application-specific high order operator assembly/evaluation routines. This improves the readability of the codes, allows low-level optimizations to be offloaded to libraries with consistent interfaces, e.g., Batched BLAS, providing performance portability, and simplifies the custom generation of new batched routines that minimize data movements. We illustrate the approach and its performance through the MAGMA backend for the CEED API library - a lightweight portable library for high order methods that allows a wide variety of applications to share a wide variety of discretization kernels, e.g., available through independent backends.

Stanimire Tomov, Ahmad Abdelfattah
University of Tennessee, Knoxville
tomov@icl.utk.edu, ahmad@icl.utk.edu

Jack J. Dongarra
University of Tennessee, Oak Ridge National Laboratory, USA
dongarra@icl.utk.edu

Azzam Haidar
NVIDIA
azzamhaidar@nvidia.com

MS299

Computational Stochastic Dynamics

Dynamical systems arising in engineering and science are

often subject to random fluctuations. The noisy fluctuations may be Gaussian or non-Gaussian, which are modeled by Brownian motion or a-stable Levy motion, respectively. Non-Gaussianity of the noise manifests as nonlocality at a macroscopic level. Stochastic dynamical systems with non-Gaussian noise (modeled by a-stable Levy motion) have attracted a lot of attention recently. The non-Gaussianity index a is a significant indicator for various dynamical behaviors. The speaker will overview recent advances in computing deterministic quantities that carry stochastic dynamical information, including random invariant manifolds, stochastic bifurcation, mean exit time, escape probability, tipping time, most probable orbits, and transition pathways between metastable states.

Jinqiao Duan
Illinois Institute of Technology
Department of Applied Mathematics
duan@iit.edu

Xiaoli Chen
Huazhong University of Science and Technology
xlchen@hust.edu.cn

MS299

Model-form Uncertainty Quantification for Predictive Modeling in Probabilistic Graphical Models

Probabilistic Graphical Models (PGM) is an important class of methods for probabilistic modeling and inference, and constitutes the mathematical foundation of modeling uncertainty in Artificial Intelligence (AI). Its hierarchical structure allows us to bring together in a systematic way statistical and multi-scale physical modeling, different types of data, incorporating in expert knowledge, correlations and causal relationships. However, due to multi-scale modeling, learning from sparse data and mechanisms without full knowledge, many predictive models will necessarily have diverse sources of uncertainty at different scales. On the other hand, traditional Uncertainty Quantification (UQ) methods mostly consider parametric approaches, e.g., by perturbing, tuning, or inferring the model parameters, which are not suitable for the aforementioned models. For this type of model-form (epistemic) uncertainty, we develop a new information-theoretic, nonparametric approach for UQ. We develop new model-form UQ indices that can handle both parametric and non-parametric PGMs, as well as small and large model/parameter perturbations in a single, unified mathematical framework and provide an envelope of model predictions. Moreover, we propose a model-form Sensitivity Index, which allows us to rank the impact of each component of the PGM, and provide a systematic methodology to go back and update the components of the computational model that underperform.

Jinchao Feng
University of Massachusetts Amherst
feng@math.umass.edu

Joshua Lansford
University of Delaware
lansford@udel.edu

Markos A. Katsoulakis
University of Massachusetts, Amherst
Dept of Mathematics and Statistics
markos@math.umass.edu

Luc Rey-Bellet
Department of Mathematics & Statistics
University of Massachusetts
luc@math.umass.edu

Dionisios Vlachos
University of Delaware
vlachos@udel.edu

MS299

Homogenization for Generalized Langevin Equations

Generalized Langevin equation (GLE) is a stochastic integro-differential equation commonly used as a model in non-equilibrium statistical mechanics to describe the dynamics of a particle coupled to a heat bath. From modeling point of view, it is desirable to work with a reduced model that could capture the essential features of the dynamics. In this talk, I will present some results from our recent work on homogenization for a class of GLEs with state-dependent coefficients. One of our main results says that in the limit, in which all the characteristic time scales of the system modeled by the GLEs vanish at the same rate, the position variable of the system converges to a homogenized process. The homogenized process is described by a stochastic differential equation containing additional drift terms induced by the noise in the considered limit. This is joint work with Jan Wehr and Maciej Lewenstein.

Soon Hoe Lim
Nordita
KTH Royal Institute of Technology and Stockholm University
shoelim@math.arizona.edu

MS299

Transitions as Rare Events in Stochastic Delayed Systems

In this work, we develop a strategy for the minimum action method (MAM) using a posteriori error estimate and generalize it to deal with systems with delays. MAM plays an important role in minimizing the Freidlin-Wentzell action functional, which is the central object of the Freidlin-Wentzell theory of large deviations for noise-induced transitions in stochastic dynamical systems as rare events. We will demonstrate the effectiveness of our method to calculate the transition probability and the most possible transition for systems with or without delays.

Jiayu Zhai
Luisiana State University
zhai@math.umass.edu

Xiaoliang Wan
Louisiana State University
Department of Mathematics

xlwan@math.lsu.edu

MS300

Joint Reconstruction in Multienergy X-ray Tomography

We propose a joint image reconstruction approach for multi-energy CT with a novel low dose measurement protocol. The central idea in joint reconstruction is to combine all the projection data into a single inverse problem and to utilize regularization models that promote some prior information on the unknown images within and across the energies. In multi-energy CT, a feasible prior assumption for the attenuation images at different energies is that they can be expected to be structurally similar in the sense that an edge (e.g. an organ boundary) that is present at one energy, is likely to be at same location and alignment with the other energies as well. In this paper, we investigate various structural similarity promoting regularization models for the joint reconstruction in multi-energy CT. We also propose a new joint regularization model based on the structure function from the structural similarity index. To reduce the radiation dose in multi-energy CT, we combine the joint reconstruction approach with a novel low dose measurement protocol, which utilizes sparse projection sampling for each energy with mutually non-overlapping projection angles between the spectral components. The proposed approach is evaluated with simulated multi-energy CT data and experimental data from a biological specimen.

Alexander Meaney
University of Helsinki
alexander.meaney@helsinki.fi

Jussi Toivanen
University of Eastern Finland
jussi.toivanen@uef.fi

Samuli Siltanen
University of Helsinki
Finland
samuli.siltanen@helsinki.fi

Ville P. Kolehmainen
Department of Applied Physics
University of Eastern Finland
ville.kolehmainen@uef.fi

MS300

Forward and Inverse Computation in Ultrasound Tomography

In Ultrasound Tomography (UST) transducers around the boundary of a medium measure the scattered acoustic waves arising from transmitted pulses emitted sequentially from the transducers. UST has medical applications in breast cancer detection and imaging, in which slices of the breast are imaged with the subject lying prone, and the breast suspended in a saline-filled vessel. The technique has the advantages of being non-ionizing, fast, and comfortable. In this talk, the effects on computation time and image quality of several input excitation patterns, transducer models, and computational implementation will be compared for several reconstruction algorithms for ultrasound tomography.

Jennifer L. Mueller
Colorado State University
Dept of Mathematics
jennifer.l.mueller@gmail.com

MS300

Classifying Stroke using Electrical Impedance Tomography

Stroke is a leading cause of death all around the world. There are two main types of stroke: ischemic (blood clot preventing blood flow to a part of the brain) and hemorrhagic (bleeding in the brain). The symptoms are the same, but treatments very different. A portable "stroke classifier" would be a life-saving equipment to have in ambulances, but so far it does not exist. Electrical Impedance Tomography (EIT) is a promising and harmless imaging method for stroke classification. In EIT one attempts to recover the electric conductivity inside a domain from electric boundary measurements. This is a nonlinear and ill-posed inverse problem. The so-called Complex Geometric Optics (CGO) solutions have proven to be a useful computational tool for reconstruction tasks in EIT. A new property of CGO solutions is presented, showing that a one-dimensional Fourier transform in the spectral variable provides a connection to parallel-beam X-ray tomography of the conductivity. One of the consequences of this nonlinear Fourier slice theorem is a novel capability to recover inclusions within inclusions in EIT. In practical imaging, measurement noise causes strong blurring in the recovered profile functions. However, machine learning algorithms can be combined with the nonlinear PDE techniques in a fruitful way. As an example, simulated strokes are classified into hemorrhagic and ischemic using EIT measurements.

Samuli Siltanen
University of Helsinki
Finland
samuli.siltanen@helsinki.fi

MS300

Debluring and Denoising of Signals with Uncertainty Quantification with Applications to Inverse Problems

Images and sounds are signals that are subject to blur and noise. The removal of these two destructive factors of a signal is crucial for the applications that the signals are used. We will see in this talk a statistical near optimal filtering method that has been used in both images and acoustical signals effectively that removes blur and noise with specific examples from the above categories.

Viktoria Taroudaki
Eastern Washington University
vtaroudaki@ewu.edu

MS301

Choosing Interpolation Points in Interpolatory

Model Reduction

In contrast to methods like balanced truncation (BT) model reduction (which is global) interpolatory model reduction is local, and requires the choice of interpolation points. For instance, generically, given any two finite dimensional LTI systems, the higher complexity one can be projected to the second one by means of an interpolatory projection. This shows the generality of this approach. The issue however in general, is how to choose interpolation points so as to achieve desired goals. The first method addressing this issue (introduced in 2008) was H2 optimal model reduction. It involves an iterative choice of interpolation points, which converges to an optimal selection. Recently, a new selection method has been suggested. It consists in selecting as many interpolation data as possible and subsequently compressing it to the required complexity. The basis for this approach is the Loewner framework for interpolatory model reduction, which constructs a high-order model of the data and then truncates it to the desired complexity. It has been shown recently that this approach compares favorably with methods like Vector Fitting, the AAA algorithm, and the CUR factorization. In addition to this method we will discuss issues arising from certain choices of interpolation points which should be avoided. Following the theoretical foundations of the proposed method the presentation will conclude with several illustrative numerical examples.

Athanasios C. Antoulas
Dept. of Elec. and Comp. Eng.
Rice University
aca@rice.edu

MS301

Data-driven Strategies for Modeling Dissipative Systems

The modeling of physical systems should take into account conservation laws, but this can be a significant challenge when models are derived directly from system response data in the absence of ancillary knowledge of internal dynamics. Observational noise and unmodeled sources or sinks can further complicate this enterprise. A data-driven modeling framework will be introduced here that yields a convex family of passive/dissipative models all of which are consistent with observed response profiles. This sets the stage for data-driven optimization using various performance metrics; we discuss the maximization of passivity margins and $\mathcal{H}_2$-minimization of frequency response error.

Christopher A. Beattie
Virginia Polytechnic Institute and State University
beattie@vt.edu

MS301

Model Reduction of Linear and Nonlinear Semi-discretized Fluid Flow Problems from Data

Model reduction is generally used to replace complex, large-scale models corresponding to time-varying processes with simpler and smaller models that are able to capture the behavior of the original process. The primary method of interest in this work is the Loewner framework. It is a data-driven method that constructs reduced order models (ROMs) directly from measurements of transfer functions (input-output mappings in the frequency domain) of the underlying system. Additionally, by compressing the usually large data set, it extracts the dominant features and eliminates the inherent redundancies. While the Loewner framework for linear systems is fairly developed, its extension to nonlinear systems is still a matter of active ongoing research. In recent years, the method has been extended to bilinear, quadratic-bilinear (QB) and linear switched systems. We study the generalization of the Loewner framework to compute ROMs of QB systems. Such systems arise in semi-discretizations of fluid flow problems, such as Burgers' or Navier-Stokes equations. We propose certain improvements that allow an efficient extension, implementation, and scaling of the method to high dimensional problems. We address such issues as stability preservation, data interpretation, and acquisition that need to be carefully dealt with when applying the method to fluid flow problems. Finally, we present numerical results that show the potential of this framework, but also highlight some open issues.

Ion Victor Gosea
Max Planck Institute Magdeburg
gosea@mpi-magdeburg.mpg.de

MS301

Learning Reduced-order Models in New Variables: Lifting and Operator Inference

We present a framework for data-driven model reduction of nonlinear systems via lifting. That is, we lift the original general nonlinear dynamics to a quadratic system via the introduction of auxiliary state variables.We then use a data-driven operator inference framework to infer reduced-order quadratic operators in a non-intrusive way. The inference step can be performed offline, enabling precomputation of the reduced operators for rapid online evaluation of the reduced system. Because our quadratic reformulation is exactly equivalent to the original system at the PDE level, our framework bypasses additional errors that would be introduced by hyper-reduction of the nonlinearity. Moreover, this framework provides a pathway to analysis of the lifted, quadratic ROM. Numerical results are presented for fluids and combustion model problems.

Elizabeth Qian, Boris Kramer, Karen E. Willcox
Massachusetts Institute of Technology
elizqian@mit.edu, bokramer@mit.edu, kwillcox@ices.utexas.edu

MS302

Space-split Sensitivity (S3) Computation of Statistics in Chaotic Dynamical Systems

In a chaotic system, the derivatives of state functions with respect to system parameters such as design or control inputs grow exponentially with time. Yet, the infinite-time average of a state function, equal to its average according to the steady-state distribution over state-space, has a bounded derivative to parameters. Computing this statisti-

cal response to infinitesimal changes in parameters without the unstable long-time evolution of instantaneous derivatives is the objective of the perturbation space-split sensitivity (S3) algorithm. Based on the fact that perturbations that lie in the stable subspace decay in time exponentially, the contribution to the overall sensitivity from them is computed similar to in non-chaotic systems. The unstable contribution is reduced by an integration-by-parts procedure to a variance-reduced Monte-Carlo sampling. Unlike Shadowing-based computation wherein the sensitivity is computed along a single trajectory which is not guaranteed to be typical, this method is provably convergent. The S3 algorithm is demonstrated on classical examples and extended to high-dimensional numerical fluid flow simulations.

Nisha Chandramoorthy, Qiqi Wang
Massachusetts Institute of Technology
nishac@mit.edu, qiqi@mit.edu

MS302

Stability, Sensitivity and Optimization of Unsteady Thermoacoustic Systems

In combustors, thermoacoustic oscillations occur when the heat released by the flame is sufficiently in phase with the acoustic pressure. These oscillations can severely damage an engine and, in some cases, shake it apart. Controlling thermoacoustic systems is challenging because they are extremely sensitive to small modifications and exhibit highly nonlinear behaviours, such as chaos. Chaos is caused by the flame saturation and the turbulent hydrodynamic field, the latter of which unpredictably modulates the phase between the heat release and pressure. We propose covariant Lyapunov vector (CLV) analysis to study the stability and sensitivity of a qualitative model of a chaotic thermoacoustic system. First, we show that eigenvalue and Floquet analyses are the limits of CLV analysis when the attractor is a fixed point and a periodic orbit, respectively. Second, we find that, if the hydrodynamic field is a (quasi) hyperbolic dynamical system, so will the chaotic thermoacoustic behaviour. Third, we apply the non-intrusive shadowing least squares method to calculate the sensitivity of the time averaged acoustic energy and Rayleigh criterion to the flame parameters. The sensitivity is then employed in a gradient-based optimization routine to find the set of parameters for which the system is stable. The method and results of this work open up new possibilities for the calculation of the stability of unsteady thermoacoustic systems with highly nonlinear behaviours.

Luca Magri
Stanford University
lm547@cam.ac.uk

Francisco Huhn
University of Cambridge
fh360@eng.cam.ac.uk

MS302

Analysis of Rare Events using Transfer Operators and Community Clustering

Abstract not available

Peter Schmid
Imperial College London
pjschmid@imperial.ac.uk

MS302

Identifying Precursors to Intermittent or Extreme Events using Conditional Space-time Proper Orthogonal Decomposition

Intermittency is a salient feature of the acoustic emissions of aircraft engines. Both experimental and numerical studies show that jet noise comprises acoustic bursts that occur in rapid succession. Despite its technical relevance, the role and origin of intermittency in the generation of jet noise are not well understood. In this talk, we present a data mining technique that unravels the complex space- and time-local behavior of the physical mechanism associated with the statistically dominant, i.e. the prototype, acoustic burst event. The method is based on a space-time formulation of proper orthogonal decomposition that is conditioned to loud events. As an example, we apply the algorithm to high-fidelity large eddy simulation data of a supersonic nearly ideally expanded hot turbulent jet. The space-time formulation allows us not only to identify the spatial structure of the average acoustic burst event but also to trace its evolution back in time to a precursor event. This has important implications for future model predictive control.

Oliver T. Schmidt
University of California, San Diego
oschmidt@ucsd.edu

Peter Schmid
Imperial College London
pjschmid@imperial.ac.uk

MS303

Title Not Available

Abstract not available

Aurelien Bouteiller
University of Tennessee, Knoxville
bouteill@icl.utk.edu

George Bosilca
University of Tennessee - Knoxville
bosilca@icl.utk.edu

MS303

Resilience for Extreme Scale Systems: Understanding the Problem

Resilience is one of the critical challenges of extreme-scale high-performance computing (HPC) systems, as component counts increase, individual component reliability decreases, and software complexity increases. Building a reliable supercomputer that achieves the expected perfor-

mance within a given cost budget and providing efficiency and correctness during operation in the presence of faults, errors, and failures requires a full understanding of the resilience problem. This talk provides an overview of the Catalog project, which develops a taxonomy, catalog and models that capture the observed and inferred fault, error, and failure conditions in current supercomputers and extrapolates this knowledge to future-generation systems. To date, this project has analyzed billions of node hours of system logs from supercomputers at Oak Ridge National Laboratory and Argonne National Laboratory.

Christian Engelmann
Oak Ridge National Laboratory
engelmannc@ornl.gov

MS303

Scalable, Efficient Fault Tolerance in Asynchronous Many Task (AMT) Programming Models

With growing scale and complexity of computational systems, HPC applications are increasingly susceptible to a wide variety of hardware and software faults. Accordingly, applications are ill-equipped to deal with the full spectrum of possible faults and often their response, particularly in synchronous programming models, is disproportionate to fault rate. Alternatively, Local Failure Local Recovery (LFLR), is based on the notion that a fault recovery that is localized around their occurrence is more scalable and efficient than a bulk response. LFLR is more amenable with an asynchronous programming model as opposed to synchronous ones. In this study, we demonstrate the efficiency and scalability of task-based fault recovery methodologies: task-replication and task-replay in an exemplar AMT runtime, Habanero-C++. The data/task semantics and API in Habanero were augmented to include functionality for the recovery techniques in a manner that incurs negligible overhead. Three representative mini-applications were implemented in Habanero to study the performance overhead of the recovery strategies at varying fault rates: 1-D explicit stencil, 3-D explicit stencil and sparse-matrix vector multiplication. Experiments with the three applications show that, with an efficient load balancing strategy, the additional cost incurred due to fault recovery (additional tasks) is proportional to rate of failed tasks.

Keita Teranishi
Sandia National Laboratories
knteran@sandia.gov

Nicole Slattengren
Sandia National Labs
nlslatt@sandia.gov

Sri Raj Paul
Rice University, USA
sriraj@rice.edu

Hemanth Kolla
Sandia National Laboratories
hnkolla@sandia.gov

Akihiro Hayashi
Rice University, USA
ahayashi@rice.edu

Jackson Mayo
Sandia National Laboratories
jmayo@sandia.gov

Vivek Sarkar
College of Computing
Georgia Institute of Technology
vivek.sarkar@cc.gatech.edu

MS303

Memory Errors on Modern Systems: An Analysis of Data over the Lifetime of Cielo

Maintaining the performance of high-performance computing (HPC) applications as failures increase is a major challenge for next-generation extreme-scale systems. Recent research demonstrates that hardware failures are expected to become more common due to increased component counts, reduced device-feature sizes, and tightly-constrained power budgets. Few existing studies, however, have examined failures in the context of the entire lifetime of a single platform. In this presentation, we discuss our analysis of failure data collected over the entire lifetime of Cielo, a leadership-class HPC system. Our analysis reveals several key findings, including: (i) Cielos memory (DRAM and SRAM) exhibited no discernible aging effects; (ii) correctable memory faults were not predictive of future uncorrectable memory faults; and (iii) continued advances will be required to ensure current failure mitigation techniques remain a viable option for future platforms.

Scott Levy, Kurt Ferreira
Sandia National Laboratories
sllevy@sandia.gov, kbferre@sandia.gov

MS304

Compact Reconstruction Using Teno

Abstract not available

James D. Baeder
Department of Aerospace Engineering
University of Maryland
baeder@umd.edu

MS304

Time-accurate and Highly-stable Explicit Operators

When computational scientists discover unanticipated temporal stiffness issues in their explicit solver, or want to improve its numerical stability properties, they need to rewrite significant portions of the code by formulating the responsible terms in an implicit manner. This is both time-consuming and a difficult process to reach high-order accuracy. We propose a novel method for achieving high-order and highly-stable numerical solutions with an explicit code, even for non-linear problems. It is based on a family of Time-Accurate and highly-Stable Explicit (TASE) operators that act as preconditioners on the stiff operators. The TASE operators are formally shown to be of any arbitrary

order, hence allowing to preserve the original order of accuracy in time of the numerical solver. They lead to an explicit time integration with stability properties similar to those of an implicit scheme, and of comparable cost, but without requiring major modifications of the solver at hand. The TASE operators only depend on the stiff operators, the numerical time step, and a single parameter that is analytically prescribed as a function of the desired order of accuracy and the solver's time advancement scheme. Application of the TASE methodology to canonical partial differential equations in which numerical or physical temporal stiffness justifies the use of an implicit scheme demonstrates that they achieve the desired accuracy and stability properties.

Maxime Bassenne
Center for Turbulence Research
Stanford University
bassenne@stanford.edu

Ali Mani
Stanford University
alimani@stanford.edu

MS304

A Fourth-order Adaptive Mesh Refinement Algorithm for the Multicomponent, Reacting Compressible Navier-Stokes Equations

Compressible reacting flow exhibits a range of dynamic scales in both space and time with the finest scales existing in only a small fraction of the total area of interest in the simulation. Most of the previous work on Adaptive Mesh Refinement (AMR) has focused on second-order finite volume discretizations and other processes using operator splitting or other second-order temporal integration approaches. In this talk, a novel strategy will be presented to achieve high-accuracy in both time and space. The spatial discretization uses a higher-order conservative finite volume treatment of advection and diffusion. The method uses an implicit/explicit (IMEX) Spectral Deferred Correction (SDC) temporal integration strategy that treats advection and diffusion explicitly while treating reactions implicitly, enabling the methodology to handle stiff reaction kinetics. Overall, the algorithm developed here makes a potentially valuable approach for DNS of reacting flows, and has been successfully applied to the simulation of reacting dimethyl ether and hydrogen jet flames.

Emmanuel Motheau
Lawrence Berkeley National Lab
Center for Computational Sciences and Engineering
emotheau@lbl.gov

Matthew Emmett, Weiqun Zhang
Lawrence Berkeley National Laboratory
Center for Computational Sciences and Engineering
mwemmett@lbl.gov, weiqunzhang@lbl.gov

Michael Minion
Lawrence Berkeley National Lab
mlminion@lbl.gov

John B. Bell
Lawrence Berkeley National Lab
Center for Computational Sciences and Engineering
jbbell@lbl.gov

MS304

Large-eddy Simulation of Separated Flows: The Role of Subgrid-scale Models

Most of the celebrated subgrid-scale (SGS) models for large-eddy simulation (LES) of turbulent flows perform less well than expected due to the lack of physically meaningful constraints, especially for the separated flows. Here, we introduce a constrained large-eddy simulation (CLES) method for simulation of wall-bounded separated flows. Formally, the SGS models are constructed in different forms within the near-wall and far-wall regions. In the far-wall region, traditional SGS models is employed, whereas in the near-wall region, the mean SGS models are constrained by prescribed Reynolds stress or/and heat flux. Several Reynolds stress (and heat flux) models, including Sparlat-Allmaras (SA) model, Menters k-ω Shear Stress Transport (SST) model, k-ω SST coupled with Gamma predictive transition model, etc. are evaluated as the model constraints. The CLES models are compared with frequently used SGS models, such as Smargorinsky model and Wall Adaptive Local-Eddy (WALE) Viscosity model, in simulations of both internal and external flows. On the one hand, The CLES method can eliminate the non-physical Log-layer mismatch phenomenon reported in hybrid RANS-LES methods, and can predict the mean velocity profile, friction force and other statistical quantities more accurately than traditional LES and hybrid RANS-LES methods; on the other hand, the practical performance of the CLES method is closely associated with the Reynolds constraint models.

Zuoli Xiao, Rui Wang
College of Engineering
Peking University
z.xiao@pku.edu.cn, pkuwangrui@pku.edu.cn

MS305

Parallel Sparse Tensor Decomposition with the Trilinos Parallel Linear Algebra Framework

Abstract not available

Karen D. Devine
Sandia National Laboratories
kddevin@sandia.gov

MS305

Tensor Decompositions on Emerging Manycore Hardware with Genten and Kokkos

We describe Genten, a new software implementing canonical polyadic (CP) tensor decompositions of sparse tensors on emerging manycore hardware. The software uses Kokkos to provide a performant implementation of the numerical kernels arising in the alternating least-squares (ALS) approach for computing CP decompositions on a variety of multicore and manycore hardware, such as multicore CPUs, Intel Xeon Phi, and Nvidia GPUs. We also de-

scribe a simple variation of the traditional coordinate storage format that reorders the traversal of tensor nonzeros in the matricized-tensor-times-Khatri-Rao product (MTTKRP) to reduce atomic-write contention and can lead to significantly improved performance with minimal increase in memory footprint. Performance of the CP-ALS decomposition is measured on several hardware platforms, including Xeon Phi and Nvidia GPUs, and performance of the MTTKRP calculation is compared to other state-of-the-art approaches available in the literature.

Eric Phipps
Sandia National Laboratories
Optimization and Uncertainty Quantification Department
etphipp@sandia.gov

Tamara G. Kolda
Sandia National Laboratories
tgkolda@sandia.gov

MS305

On Tensor Orderings for HiCOO

HiCOO is a recently proposed storage format sparse tensors. It blocks and compresses the nonzeros of the tensor into subtensors in order to speedup the classical algorithms for Candecomp/Parafac decomposition of sparse tensors. In this talk we will cover strategies to improve blocking performance in HiCOO format. This is joint work with Jiajia Li, Umit V. Catalyurek, and Rich Vuduc.

Bora Ucar
LIP, ENS-Lyon, France.
Bora.Ucar@ens-lyon.fr

MS305

On the Decomposition of Tensors that are Given Implicitly

In practice, tensor decomposition algorithms commonly try to minimize the error between the model and a given tensor. However, in some cases, the tensor is not available as such, but only implicitly as a solution of a linear system. In this talk, we discuss the algebraic and optimization-based computation of a canonical polyadic decomposition of an implicitly given tensor. An extension to tensors determined by more-dimensional null spaces is presented. We show that a variety of problems can be formulated as an implicitly given tensor decomposition and illustrate this for face recognition and finding the roots of sets of polynomial equations.

Nico Vervliet, Martijn Boussé
Department of Electrical Engineering (ESAT)
KU Leuven
nico.vervliet@esat.kuleuven.be,
martijn.bousse@kuleuven.be

Ignat Domanov
Faculty of Science
KU Leuven, Campus Kulak
ignat.domanov@kuleuven.be

Lieven De Lathauwer
KU Leuven - KULAK
lieven.delathauwer@kuleuven.be

MS306

Enabling Application Optimizations with Unified Thread-based Runtime MPC

Application developers face complex programming language features with a plethora of runtime systems on offer. Basing the runtimes associated with each language on the same basic blocks helps to unravel the difficulty of runtime stacking. For this purpose, MPC provides a user-level thread scheduler and a thread-aware allocator shared by both distributed-memory and shared-memory programming models. This common ground allows providing additional features combining best of both worlds. This presentation will highlight some extensions and optimizations available in MPC.

Julien C. Jaeger
CEA
julien.jaeger@cea.fr

MS306

Analyzing Build System Pressure for Large Applications: Experience from the US DOE ASC Program

Abstract not available

Simon D. Hammond
Scalable Computer Architectures
Sandia National Laboratories
sdhammo@sandia.gov

Paul Lin
Sandia National Laboratories
ptlin@sandia.gov

MS306

Using Just-in-time Compilation to Optimize Scientific Codes

Writing high performing scientific codes is often an exercise in learning (or guessing) how to write high-level source code in such a way that the compiler can optimize kernels. However, some important optimizations are impossible to exploit at compile time since they depend on knowing parameters that aren't determined until run time. This talk will describe a prototype just-in-time (JIT) compiler that allows us to capture run-time optimizations and substantially improve application performance. Our JIT has been used successfully for codes written in both the RAJA and Kokkos programming idioms.

David F. Richards
Lawrence Livermore National Laboratory
richards12@llnl.gov

MS306

Modernizing Compiler Design for Platform Porta-

bility

Abstract not available

Patrick McCormick, George Stelle
Los Alamos National Laboratory
pat@lanl.gov, stelleg@lanl.gov

MS307

Parallel-in-time Integration in SU2

Since its emergence roughly 50 years ago, Computational Fluid Dynamics (CFD) always has been leveraged the available computational resources and enabled the solution of progressively more complex problems in engineering and science. However, lately, there is an increasing disparity between the methodological development and the hardware. Over the last twenty years, there has been a significant increase in the levels of parallelism in the field of high-performance computing, leading to an ever-increasing number of available cores on today's supercomputers. At the same time the clock-speed per core stagnated or decreased in favor of a lower overall power-consumption. Despite these developments the fundamentals in CFD remained largely unchanged with a key objective of method design being to minimize the aggregate number of floating point operations. One of the major bottlenecks that can be identified in current CFD solvers is the sequential time marching approach because its lacks parallelism in the time dimension. Although not intuitive, a obvious solution to this problem is to simultaneously solve for multiple times steps. In this talk we want to present the latest results of our efforts of applying this idea to the solution of unsteady flow problems, by augmenting existing numerical methods to exploit the massive parallelism of current and future HPC systems.

Tim Albring
TU Kaiserslautern
tim.albring@scicomp.uni-kl.de

Stefanie Guenther
RWTH Aachen University
stefanie.guenther@scicomp.uni-kl.de

Nicolas R. Gauger
TU Kaiserslautern
nicolas.gauger@scicomp.uni-kl.de

MS307

Wall-modeled LES Simulations using the SU2 DG-FEM Solver

The simulation of turbulent flows for high Reynolds numbers using Large-Eddy Simulation (LES) can be a computationally expensive endeavor. The need to resolve a wide range of spatial and temporal scales leads to the requirement of extremely fine grids and small time steps. When the geometries involved are complex, the use of unstructured grid topologies only add to the computational complexity of the task. In this talk we discuss our efforts to decrease the cost of LES in this context through a combination of higher-order numerical discretizations, advanced time-stepping algorithms, and the use of wall-modeling techniques. These approaches have been implemented in the new higher-order, Discontinuous Galerkin Finite Element Method (DG-FEM) solver in the SU2 framework. With a combination of higher-order methods (DG) on unstructured meshes, ADER-DG time stepping, and an equilibrium wall model, we show how DG can be an effective methodology to compute such flows. The talk also discussed new developments in shock capturing techniques for DG solvers and discusses results of turbulent wall-bounded flows we have been using for validation and verification.

Edwin van der Weide
University of Twente
e.t.a.vanderweide@utwente.nl

Juan J. Alonso
Department of Aeronautics and Astronautics
Stanford University
jjalonso@stanford.edu

MS307

Advances for Incompressible Flow Calculations with SU2

While the open-source SU2 software package has a well-documented history as a compressible flow solver for high-speed aerodynamics, this presentation details the recent advances in SU2 for addressing a range of incompressible flow problems. This includes both constant and variable-density incompressible flows with heat transfer. Some potential application spaces are natural convection problems, low-speed internal flows with heat transfer, and reactive flows, to name a few. The custom formulation enables an efficient, coupled solution algorithm for the incompressible Navier-Stokes and Reynolds-averaged Navier-Stokes (RANS) equations. The method has its roots in preconditioning approaches for mixed low- and high-Mach number flows as well as in the artificial compressibility approach, which also results in a coupled solution algorithm through establishing a pressure-velocity coupling. The implementation solves the governing equations in the context of a general, unstructured finite volume code with vertex-based schemes, and it is written in C++/MPI. The solver has been demonstrated for a number of classic verification and validation cases, but as a general baseline implementation, one can extend the formulation to unsteady flows, reactive flows by way of additional scalar transport equations and/or chemistry, or adjoint-based optimization, for example. An outlook on activities in these new directions will be given.

Thomas D. Economon
Bosch
thomas.economon@us.bosch.com

MS308

Multi-fidelity Modeling of Rocket Combustion Instability

Combustion instability in liquid rocket engines (LRE) is an extremely complex phenomenon, with strongly stochastic and nonlinear physics, and time and length scales that range over many orders of magnitude. Though modern

computational capability has demonstrated the potential to beyond the empirically-based design analyses of the past, high-fidelity simulations of full-scale combustors remain out of reach, and will continue to be out of reach for engineering work flows, even on Exascale computers. While individual LES simulations of combustion instability in laboratory-scale combustors are relatively routine [Harvazinski et al., Physics of Fluids, 2015], full-scale rocket engine simulations remain too computationally intensive for design applications. Our goal is to address the gap that exists between the high-fidelity simulations of small and simplified geometries and the need for computations of large and complicated domains through reduced order modeling (ROM) for efficient and accurate predictions of combustion instability in LRE. This work will present progress towards the establishment of a multi-fidelity framework using reduced-order models (ROMs). The ROMs are trained based on high-fidelity simulations of combustion dynamics in single-injector geometry, which are then integrated and coupled with reduced-fidelity models (RFMs) for the rest of the domain in the framework to model combustion instabilities in multi-injector configuration.

Cheng Huang
University of Michigan Ann Arbor
Department of Aerospace Engineering
huangche@umich.edu

Karthik Duraisamy
University of Michigan Ann Arbor
kdur@umich.edu

Charles Merkle
Purdue University
merkle@purdue.edu

MS308

Fidelity-adaptive Combustion Modeling for Turbulent Combustion

Abstract not available

Matthias Ihme
Stanford University
Mechanical Engineering
mihme@stanford.edu

MS308

A Multi-fidelity Framework for Stability Analysis of Complex Dynamical Systems

Abstract not available

Alexandre Marques
Massacheusetts Institute of Technology
noll@mit.edu

MS308

Combustion System Model Reduction Using Pod and Neural Networks

Abstract not available

Nathan L. Mundis
ERC, Inc. @ AFRL
nmundis@miner.mst.edu

MS309

Convolutional Dictionary Learning for Inverse Problems

Inverse modeling is one of the analytical techniques that tries to facilitate the conversion of measurements into interpretable knowledge by formulating a mathematical model to explain the data and finding the parameters of the model that best fit the observations. In this work we use the convolutional dictionary learning (CDL) framework to build such a model. In a convolutional sparse representation, sums of a set of convolutions with dictionary filters (basis elements in convolutional representation) are used to construct the model. The optimization is computed over the entire signal domain, yielding representations that are very sparse both spatially and across the filter indices. We describe the CDL technique for inverse problems and illustrate its performance in a tomographic reconstruction application.

Cristina Garcia-Cardona
Los Alamos National Laboratory
cgarciac@lanl.gov

MS309

Ultrasonic Tomographic Imaging Using Physics Constrained Learning Method

Ultrasonic tomographic imaging involves inverting the sound speed distribution of an imaged medium through which the ultrasonic signals propagate before arriving at the receivers. Specifically in travel-time tomography, the estimated sound speed map minimizes the difference between the estimated and the predicted travel times. The predicted travel times are traditionally obtained via a nonlinear forward model such as ray tracing. The accuracy of the resulting sound speed estimates is affected by both the model sensitivity determined by the Jacobian and the illumination specified by the directivity pattern. We apply a deep learning technique to train the mapping between the sound speed distribution and the travel time and establish a connection between the network coefficients and the Jacobian.

Weichang Li
Aramco Research Center
weichang.li@aramcoservices.com

MS309

VelocityGAN: Subsurface Velocity Image Estimation Using Conditional Adversarial Networks

Acoustic- and elastic-waveform inversion is an important and widely used method to reconstruct subsurface velocity image. Waveform inversion is a typical non-linear and

ill-posed inverse problem. Existing physics-driven computational methods for solving waveform inversion suffer from the cycle skipping and local minima issues, and not to mention solving waveform inversion is computationally expensive. In this work, we developed a real-time data-driven technique, VelocityGAN, to accurately reconstruct subsurface velocities. Our VelocityGAN is an end-to-end framework which can generate high-quality velocity images directly from the raw seismic waveform data. A series of numerical experiments are conducted on the synthetic seismic reflection data to evaluate the effectiveness and efficiency of VelocityGAN. We not only compare it with existing physics-driven approaches but also choose some deep learning frameworks as our data-driven baselines. The experiment results show that VelocityGAN outperforms the physics-driven waveform inversion methods and achieves the state-of-the-art performance among data-driven baselines.

Youzuo Lin
Los Alamos National Laboratory
ylin@lanl.gov

MS309

Linear Regression on Imbalanced Data: New Theoretical Insights

Many real-world datasets have imbalanced features and labels, and training a machine learning model on imbalanced data is much harder than on balanced data. Although learning from imbalanced data has been intensively studied and many methods have been proposed in the past decades, theoretical understanding of this problem is limited. We study the theoretical question: Does linear regression trained on an imbalanced dataset generalize to unseen test samples? The question is equivalent to whether the inverse problem can be well solved given imbalanced data. We offer novel theoretical insights and analysis for this important problem using very simple mathematical tools. Our answer to the question is yes, linear regression generalizes, at least for the frequent samples in the imbalanced dataset. As for the rare-and-extreme samples, linear regression may or may not make a correct prediction.

Shusen Wang
Stevens Institute of Technology
wssatzju@gmail.com

MS310

Representation Learning for Large Graphs

Learning a useful graph representation lies at the heart and success of many machine learning tasks such as node and link classification, anomaly detection, and link prediction. Methods capable of learning such representations have many advantages over feature engineering in terms of cost and effort. However, the complex nature of real-world graphs provide challenges for these methods. For example, the majority of real-world networks are naturally dynamic, with complex structure, and associated with rich attribute information. This makes it challenging for representation learning methods that are capable of considering only the direct relationships among nodes. In this talk, we discuss the challenges facing graph representation learning and recent work for solving these challenges.

Nesreen Ahmed
Intel Labs
nesreen.k.ahmed@intel.com

MS310

Parallel Algorithms for Ensemble Clustering

We consider two problems in data mining: (a) consensus clustering: creating a consensus from a set of partitions of the same data and (b) meta clustering: creating a concise summary from a set of partitions of comparable, but different, data. We present combinatorial algorithms based on matching and edge cover to solve these problems efficiently. We demonstrate the use of these algorithms in image registration and immunophenotype discovery.

Ariful Azad
Lawrence Berkeley National Laboratory
azad@iu.edu

MS310

Scaling Clustering Algorithms Using Graph Theoretic Approaches

Scientific computing codes are now capable of simulating trillions of particles on petascale systems. Each timestep in such simulations generates data on the order of 10s of TBs. Summarizing and analyzing the particles are challenging, and scientists often require big data analytics solutions. In this work, we present a key machine learning kernel, named DBSCAN, to facilitate scientific insights that were previously intractable at this sheer size of datasets. We showcase how graph algorithmic techniques and data locality can be exploited to achieve massive parallelism and higher performance. We demonstrate strong and weak scaling results of trillion range cosmology and plasma physics datasets utilizing 100,000 cores. In addition, we showcase that our system is capable of performing end-to-end analysis at trillion particle scale (35+ TBs of data) in 30 minutes.

Mostofa Patwary
Baidu
patwarymostofa@baidu.com

MS310

Asynchronous Communication-efficient Algorithm for Clustering Coefficient Computation

Local Clustering Coefficient (LCC) is one of the most important queries in Graph Analytics with many application in several domains such as network mining, recommendations, cybersecurity, and functional biology among others. Distributed algorithms suffer from poor scalability due to irregular message sizes and synchronization points. In this work, we present a new asynchronous algorithm for LCC computation, namely ALCC. We demonstrate how our MPI-3 One-sided implementation combined with a cached RMA layer improves the performance of ALCC up to a

factor of 5x.

Flavio Vella
DIVIDITI
flavio.vella@unibz.it

Torsten Hoefler
ETH Zürich
htor@inf.ethz.ch

Salvatore Di Girolamo
ETH Zurich
salvatore.digirolamo@inf.ethz.ch

MS311

Deep Learning for Gust Detection from Available Wing Sensors in Unsteady Aerodynamics

Flight control decisions in aerodynamics are made based on the information supplied from available sensors. In the presence of large-amplitude external disturbances, such as incident gusts and wing maneuvers, effective control might depend on more thorough knowledge of the time-varying flow state. However, sensors are typically limited to surface-based pressure measurements, which raises the question, "Can the full flow state be estimated from these sensors?' In this work, we demonstrate the use of tools from supervised deep learning to estimate the characteristics of external flow disturbances from time-resolved surface pressure measurements obtained from an inviscid flow simulation. Using an architecture composed from a combination of convolutional and recurrent neural networks, we show that time-varying incident gust and wing maneuver parameters can be determined with good accuracy. Furthermore, these parameters can often be determined even in the presence of multiple disturbances, such as when a maneuvering wing is subjected to a gust. This capability for discerning overlapping large-amplitude disturbances contrasts notably with other approaches, which cannot contend with the non-linear interactions in the flow response.

Jeff D. Eldredge, Wei Hou
University of California, Los Angeles
Mechanical & Aerospace Engineering
eldredge@seas.ucla.edu, weihou@ucla.edu

MS311

Classifying Flow Patterns using Neural Networks

We consider the inverse problem of classifying flow patterns from local sensory measurements. This problem is inspired by the ability of various aquatic organisms to respond to ambient flow signals, and is relevant for translating these abilities to underwater robotic vehicles. Specifically, we train neural networks to classify flow patterns by relying on a single flow sensor that measures a time history of the local flow signal. We systematically investigate the network performance for four distinct types of sensors vorticity, flow velocities parallel and transverse to the direction of flow propagation, and flow speed and show that the networks trained using transverse velocity outperform other networks, even when subjected to aggressive data corruption. We then train the network to classify flow patterns in real time, using a spatially-distributed array of sensors and a single one time sensory measurement. The network, based on a handful of spatially-distributed sensors, exhibits remarkable accuracy in flow classification. These results lay the groundwork for developing learning algorithms for the dynamic deployment of sensory arrays in unsteady flows.

Mohamad Alsalman, Brendan Colvert, Eva Kanso
University of Southern California
m.alsalman92@hotmail.com, brendancolvert@gmail.com, kanso@usc.edu

MS311

A Search for Physical Insight in Machine Learning Classification of Vortex Patterns

Despite relying only on data, with no underlying model of the datasets or phenomena being labeled, machine learning methods have proven remarkably effective in a wide variety of classification tasks. Research in this field is typically measured by the ability to achieve high accuracy scores, often tested on publicly available datasets. However, in many scientific applications, achieving high classification accuracy is not the end goal. Instead, scientists often seek to gain a deeper understanding of the data they analyze. This presentation explores the application of machine learning classifiers to label vortex patterns generated by pitching and plunging plates, respectively. Using limited downstream measurements, the classifiers are able to identify the upstream motion that generates the observed vortex pattern. Rather than focus solely on the highest achievable accuracy score, we investigate how choices in data collection affect classifier accuracy, and what this might imply about the underlying physics. For instance, what variables should be measured? Are some locations in the flow field more informative than others? We present our findings and discuss our strategies for answering these questions, including different approaches for visualizing results.

Jonathan H. Tu
NSWC Carderock
jonathan.tu@navy.mil

MS311

Clustering and Classifying Vortex Wakes by Dynamical Regimes

Biological swimmers generate vortical wakes during locomotion. The dynamic evolution of these wakes can be markedly different depending on the characteristics of the swimmer and its swimming gait. In this study, we use analytical point vortex models of 2S and 2P wakes—in which the dynamic regimes of motion are known—to establish the efficacy of using machine learning to achieve various goals. First, we show an unsupervised learning strategy that processes global snapshots of the velocity field to cluster wakes according to their dynamics. Second, we show that supervised learning strategies can be tailored for the purposes of classifying wake regimes from the localized hydrodynamic signals they impart on nearby bodies. Finally, we present results and extensions of these techniques for classifying

wake regimes from noisy and streaming measurements of the hydrodynamic signal at a single point on a fish-like body.

Mengying Wang
University of Minnesota
wang5772@umn.edu

Maziar S. Hemati
Department of Aerospace Engineering and Mechanics
University of Minnesota
mhemati@umn.edu

MS312

Derivative-free Multifidelity Design Optimization under Uncertainty of a Scramjet

We present new developments for the derivative-free nonlinear constrained stochastic optimization method SNOWPAC, with demonstrations on a design optimization problem involving a scramjet. Optimization of the scramjet design requires accounting for uncertain conditions while ensuring reliable performance. Moreover, simulation codes for the scramjet and many other complex engineering systems are only available as black boxes. Optimization under uncertainty (OUU), in a derivative-free setting, is thus required. Here, SNOWPAC uses a model-based trust region approach to obtain approximate derivative information and optimize over fully linear surrogate models. Additionally, it uses Monte Carlo sampling approaches to estimate various measures of uncertainty or risk (in the objective or constraints). Sampling produces noise in the estimates, however, which imposes a lower bound on the trust region radius. To mitigate this effect, SNOWPAC employs an "outer' Gaussian process surrogate model. By biasing the pointwise sample estimates using Gaussian processes, we are able to reduce the Monte Carlo error and can shrink the trust region until we reach convergence. In this talk we present developments in the convergence theory for SNOWPAC. We also present our newest implementations of multifidelity Monte Carlo methods, and their error estimates, to evaluate typical OUU objectives and constraints. We show results for a scramjet design problem.

Friedrich Menhorn
Department of Informatics
Technical University of Munich
menhorn@in.tum.de

Michael S. Eldred
Sandia National Laboratories
Optimization and Uncertainty Quantification Dept.
mseldre@sandia.gov

Gianluca Geraci
Sandia National Laboratories, NM
ggeraci@sandia.gov

Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu

MS312

Bayesian Optimization of Combinatorial Structures

The optimization of expensive-to-evaluate black-box functions over combinatorial structures is an ubiquitous task in machine learning, engineering and the natural sciences. The combinatorial explosion of the search space and costly evaluations pose challenges for current techniques in discrete optimization and machine learning, and critically require new algorithmic ideas (NIPS BayesOpt 2017). In this talk we will propose Bayesian optimization of combinatorial structures (BOCS), that takes a novel approach to overcome these challenges. It is based on an adaptive scalable statistical model that is able to identify useful combinatorial structures even when data is scarce. BOCS' acquisition function pioneers the use of semidefinite programming to achieve efficiency and scalability. We will also discuss a comprehensive experimental evaluation that demonstrate that BOCS consistently outperforms other methods from combinatorial and Bayesian optimization.

Matthias Poloczek
Cornell University
polo.email@gmx.de

Ricardo Baptista
MIT
rsb@mit.edu

MS312

Bayesian Optimization with Expensive Integrands

We propose a Bayesian optimization algorithm for objective functions that are sums or integrals of expensive-to-evaluate functions, allowing noisy evaluations. These objective functions arise in design of engineering systems with random environmental conditions, multi-task Bayesian optimization for tuning machine learning hyperparameters, and optimization via simulation. Our method is average-case optimal by construction when a single evaluation of the integrand remains within our evaluation budget. We also show consistency of our method for objective functions that are sums. In numerical experiments, our method performs as well or better than other benchmarks across a wide range of problems.

Saul Toscano-Palmerin
Cornell University
st684@cornell.edu

Peter I. Frazier
Dept. of Operations Research and Information Engineering
Cornell University
pf98@cornell.edu

MS312

Derivative-free Robust Optimization by Outer Ap-

proximations

We develop an algorithm for minimax problems that arise in robust optimization in the absence of objective function derivatives. The algorithm utilizes an extension of methods for inexact outer approximation in sampling a potentially infinite-cardinality uncertainty set. Clarke stationarity of the algorithm output is established alongside desirable features of the model-based trust-region subproblems encountered. We demonstrate the practical benefits of the algorithm on a new class of test problems.

Matt Menickelly
Lehigh University
mmenickelly@anl.gov

Stefan Wild
Argonne National Laboratory
wild@mcs.anl.gov

MS313

The Scalability of Block Iterative Methods

We discuss new developments for the block Cimmino solver that uses a block conjugate gradient algorithm where each iteration involves the embarrassingly parallel solution of a number of independent rectangular subsystems using a parallel sparse direct method. We first describe work that we have done in designing a new solution code for symmetric indefinite matrices. This code implements the forward and backward substitution steps needed to solve the system after the matrix has been factorized, and we show that it is far more efficient and scalable than other state-of-the-art codes. This is important in the block Cimmino context because the factorizations on the subblocks are done only once but the solves are performed on each iteration. In addition our direct code is designed to run on a multicore node. We have used as a basis the ABCD code of Mohamed Zenadi from ENSEEIHT-IRIT that uses MUMPS as the direct solver. We distribute the direct solution of the subsystems over the compute nodes using MPI and then solve these subsystems on a single multicore node using our new code. We compare the performance of our new approach and code with the original ABCD code. We have also modified the code to solve sparse least squares problems and show results from this also.

Sebastien Cayrols
Rutherford Appleton Laboratory
sebastien.cayrols@stfc.ac.uk

Iain Duff
Science & Technology Facilities Council, UK
and CERFACS, Toulouse, France
iain.duff@stfc.ac.uk

MS313

Solution of Sparse Unsymmetric Systems

We compare two approaches for the solution of sparse unsymmetric equations. The first uses a parallel Markowitz threshold criterion and is a right-looking factorization that is efficient when the matrix factors are not too much denser than the original matrix, that is when there is not much fill-in in the sparse factorization. We have developed novel algorithms for pivoting and have incorporated these in the ParSHUM code. The other approach uses an ordering based on a symmetrization of the matrix and works using an assembly tree representation of the factorization with dense kernels at each node of the tree. The framework is similar to that which we have used for symmetric indefinite matrices. We discuss the implementation of these two approaches within the NLAFET Project and examine the performance of these two codes on test matrices from a power systems application. We also study the effect of first reordering the matrix to singly bordered block diagonal form using the Zoltan code from Sandia.

Iain Duff
Science & Technology Facilities Council, UK
and CERFACS, Toulouse, France
iain.duff@stfc.ac.uk

Florent Lopez
Rutherford Appleton Laboratory
florent.lopez@enseeiht.fr

Stojce Nakov
STFC-RAL
stojce.nakov@stfc.ac.uk

MS313

Preconditioned Linear Solvers in Cosmic Microwave Background Data Analysis

Studies of the Cosmic Microwave Background (CMB) anisotropies have been driving the progress in our understanding of the Universe for more than 20 years. The current and forthcoming CMB observatories are expected to deliver unprecedented insights about the Universe's beginning and evolution, producing enormous data sets of size $O(10^{15})$ and thus calling for advanced, high performance data analysis techniques and efficient algebraic solvers. We present several applications in CMB data analysis that result in linear algebraic problems or sequences of problems with matrices that, properly represented, are large but very sparse. Then we discuss the state-of-the-art algebraic solvers as well as alternative techniques, which may yield to significant savings in the solution time. This is a joint work with Laura Grigori (INRIA Paris) and Radek Stompor (Université Paris 7 Denis Diderot)

Jan Pape
Inria
Paris
jan.papez@inria.fr

MS313

Distributed Tasking in the PLASMA Numerical Library

PLASMA (Parallel Linear Algebra Software for Multicore Architectures) is an established numerical linear algebra library for shared memory multicore and manycore architectures. Since its origins around 2007, PLASMA has been an early adopter of the task-based programming model. The recent tasking features of the OpenMP standard have been

introduced in the latest version of the library. Due to the limitation of OpenMP to shared memory architectures, we have recently explored different runtime systems with support for distributed tasking. In particular, we have ported PLASMA to the StarPU runtime as well as to the recent Dynamic Task Discovery interface of PaRSEC with a minimal intrusiveness to the shared memory version. In this talk, we summarize our experience with using these two task-based programming libraries and discuss preliminary performance results.

Mawussi Zounon
Univeristy of Manchester
mawussi.zounon@manchester.ac.uk

Jakub Sistek
The University of Manchester
School of Mathematics
jakub.sistek@manchester.ac.uk

Jack J. Dongarra
University of Tennessee, Oak Ridge National Laboratory, USA
dongarra@icl.utk.edu

MS314

Optimal Bayesian Experimental Design Using Generalized Laplace Method

Bayesian experimental design is essential to improve data quality in engineering. Its application on real problems lags behind mainly due to the involved computational costs. We develop a series of methods to accelerate the computations of the utility function (expected information gain) under rigorous error control. Specifically, we extend the applicable domain of Laplace methods from the asymptotic posterior Gaussianity, to where the shape of the posterior is characterized by non-informative manifolds. The developed methodologies have been applied to various engineering problems, e.g., impedance tomography, seismic source inversion and parameter inference of combustion kinetics.

Quan Long
United Technologies Research Center
longq@utrc.utc.com

MS314

Scalable Methods for Bayesian Optimal Experimental Design Using Laplace Approximation

Optimal experimental design (OED) for Bayesian nonlinear inverse problems governed by partial differential equations (PDEs) is an extremely challenging problem. First, the parameter to be inferred is often a spatially correlated field and it leads after discretization to a high dimensional parameter space. Second, the forward model is often extremely complex and computationally expensive to solve. A common objective function for OED is the expected information gain (EIG). Naïve evaluation of EIG is unfeasible for large-scale problems due to the large number of samples required in the double-loop Monte Carlo. To overcome these difficulties, we invoke an approximation of the EIG based on the Laplace approximation of the posterior. Each evaluation of the objective function then requires computing a sample average approximation (over possible realization of the data) of the information gain (IG) between the Laplace approximation and the Gaussian prior distribution. An analytical formula for the IG is available and it involves the log-determinant and trace of the posterior covariance operator. Randomized eigensolver algorithms allows us to efficiently estimate such invariants at a cost that is independent of the dimension of the parameter space, thus allowing for a scalable evaluation of the objective function. Variational adjoint methods are then used to efficiently compute the gradient of the OED objective function.

Umberto Villa
University of Texas at Austin
uvilla@wustl.edu

Omar Ghattas
The University of Texas at Austin
omar@ices.utexas.edu

MS314

A Sparse Optimization Framework for Optimal Sensor Placement

A novel approach to optimal sensor placement for infinite dimensional Bayesian inverse problems with partial differential equations. The identification of a distributed parameter entering a partial differential equation from pointwise measurements of the associated state is considered. We formulate a suitable optimal sensor placement problem and model the distribution of the measurement sensors as a regular Borel measure on the spatial domain. This leads to a non-smooth but convex optimization problem. Results concerning its well-posedness and a suitable approximation framework are presented. The theoretical findings are illustrated by extensive numerical examples.

Daniel Walter
Technische Universität München
walter@ma.tum.de

Boris Vexler
Technische Universitaet Muenchen
Centre for Mathematical Sciences, M1
vexler@ma.tum.de

MS314

Sum-up Rounding Approaches for Design of Experiments

We present a numerical method for approximating the solution of convex integer programs stemming from optimal experimental design. The statistical setup consists of a Bayesian framework for linear inverse problems for which the direct relationship is described by a discretized integral equation. Specifically, we aim to find the optimal sensor placement from a set of candidate locations where data are collected with measurement error. The convex objective function is a measure of the uncertainty, described here by the posterior covariance matrix, of the discretized linear inverse problem solution. The resulting convex in-

teger program is relaxed producing a lower bound. An upper bound is obtained by extending the sum-up rounding approach to multiple dimensions. For this extension, we provide an analysis of its accuracy as a function of the discretization mesh size. We show asymptotic optimality of the integer solution defining the upper bound for different experimental design criteria (A, D, E - optimal), by proving the convergence to zero of the gap between upper and lower bounds as the mesh size goes to zero and the number of candidate locations goes to infinity in a fixed proportion to the total number of mesh points. The technique is demonstrated on a two-dimensional gravity surveying problem for both A-optimal and D-optimal sensor placement where our designs yield better procedures compared to random placements.

Jing Yu
university of Chicago
jingyu@galton.uchicago.edu

Mihai Anitescu
Argonne National Laboratory
Mathematics and Computer Science Division
anitescu@mcs.anl.gov

MS315

Constrained Inversion for Subsurface Imaging

Abstract not available

Jodi Mead
Boise State University
Department of Mathematics
jmead@boisestate.edu

MS315

Numerical Methods for Direct Reconstruction Methods in Electrical Impedance Tomography

Abstract not available.

Peter Muller
Department of Mathematics
Rowan University, Glassboro, NJ 08028
pmuller1187@gmail.com

MS315

A Global Reconstruction Method for a Coefficient Inverse Problem with One Measurement

We discuss in this talk our recent results on a newly developed reconstruction method for solving a coefficient inverse problem with multi-frequency scattering data. These data are generated by only one direction of the incident plane wave. The method aims to provide a good approximation for the exact solution without using any a priori knowledge of any point in a small neighborhood of that solution. This is the main advantage of the method, compared with classical approaches using optimization schemes. Numerical results are presented for both simulated data and experimental data. This is joint work in part with M. V. Klibanov, L. H. Nguyen, M. A. Fiddy and H. Liu.

Dinh-Liem Nguyen
Department of Mathematics
Kansas State University
dlnguyen@ksu.edu

MS315

A Spectral Projection Preconditioner for Solving Ill-Conditioned Linear Systems

We present a preconditioner based on spectral projection that is combined with a deflated Krylov subspace method for solving ill conditioned linear systems of equations. Our results show that the proposed algorithm requires many fewer iterations to achieve the convergence criterion for solving an ill conditioned problem than a Krylov subspace solver. In our numerical experiments, the solution obtained by the proposed algorithm is more accurate in terms of the norm of the distance to the exact solution of the linear system of equations.

Man-Chung Yeung
University of Wyoming
myeung@uwyo.edu

MS316

The AAA Algorithm for Parameterized Dynamical Systems

The barycentric representation of rational functions has been effectively used to tackle the problem of one-variable rational approximation. In particular, the AAA algorithm combines this representation with a greedy search of poles to obtain approximations that can compare with popular methods like vector fitting. In the context of data driven model reduction, the barycentric representation has been used to extend the nonparametric Loewner framework to the parametric case. Here we show an extension of the AAA algorithm to the parametric case and draw connections with the parametric Loewner framework.

Andrea Carracedo Rodriguez, Serkan Gugercin
Virginia Tech
crandrea@vt.edu, gugercin@math.vt.edu

MS316

Quadrature Based Rational Ls for Transfer Function Approximation

Transfer function approximation in the frequency domain is one of the kernel computational routines in system identification and model order reduction. Given frequency response data, the task to construct best least-squares rational approximant is accomplished in a numerically sound way by the widely used Vector Fitting, which implements the Sanathanan-Koerner iterations in barycentric form, with iterative reallocation of the poles at each iteration. The algebraic least squares approximation can be improved by selecting the sampling points and introducing weights so that the weighted least squares error approximates the error in the $\mathcal{H}_2$ norm. With an appropriate quadrature rule, such as e.g. the Clenshaw-Curtis, the method al-

most reaches the performance of the IRKA (Iterative Rational Krylov Algorithm) algorithm for model order reduction, and it can be considered an effective preprocessing for IRKA. We discuss this approach and its performances, together with other issues such as regularization, using a discrete Sobolev norm, and ill-conditioning.

Zlatko Drmac
University of Zagreb
Department of Mathematics
drmac@math.hr

Christopher A. Beattie
Virginia Polytechnic Institute and State University
beattie@vt.edu

Serkan Gugercin
Virginia Tech
Department of Mathematics
gugercin@vt.edu

MS316

Solving the Laplace Equation using Rational Functions

The AAA Algorithm introduced by Nakatsukasa, Sète, and Trefethen is a robust and efficient way of computing rational approximations to functions in the complex plane. The core idea is the use of the rational barycentric representation with adaptive choice of support points and minimization of a linearized residual. In this talk, we present how AAA can be leveraged to construct approximations to solutions of the Laplace problem. We are given a domain and real-valued boundary data on the boundary of the domain, and seek to construct a rational function whose real part closely matches the prescribed boundary data and is analytic in the domain. Our strategy for this is to first neglect the analyticity constraint, and use AAA to construct a rational function which matches the boundary data as closely as possible. We then proceed by iteratively refining approximations to the harmonic conjugate while maintaining accurate matching of the real part of the approximation with the boundary data using a similar approach to earlier work by Hochman, Leviatan, and White.

Abinand Gopal
University of Oxford
abinand.gopal@maths.ox.ac.uk

MS316

Projected Nonlinear Least Squares for H2 Model Reduction

Optimal H2 model reduction seeks to build a reduced order model of a linear, time invariant (LTI) system that minimizes the mismatch of the transfer function along the imaginary axis in the L2-norm. Here we propose a new approach to the H2 model reduction problem that only requires access to samples of the transfer function. Our approach projects the infinite dimensional optimization problem onto a finite dimensional subspace via the sampling operator that results in a nonlinear weighted least squares program. Rational approximation in a conventional pole-residue parameterization often possesses many spurious local minima; therefore, we propose Gauss-Newton initialized with the Adaptive Antoulas-Anderson (AAA) Algorithm that frequently avoids these spurious minima. Finally, we provide an iteration that improves this subspace until the suboptimal solutions converge to the optimal H2 solution. Currently, the main technique for constructing an optimal reduced order model in the H2-norm is the Iterative Rational Krylov Algorithm (IRKA), requiring access to the full order system in state-space form which is not always available in experimental settings. We compare our algorithm to two data-driven approaches: Transfer Function IRKA (TF-IRKA) and quadrature-based vector fitting (QuadVF). Since this projection-based approach is based on numerical optimization techniques, our approach has the potential to be extended to new classes of transfer function realizations.

Jeffrey Hokanson
University of Colorado Boulder
jeho8774@colorado.edu

Caleb C. Magruder
Rice University, Computational and Applied Mathematics
caleb.magruder@tufts.edu

MS318

A Hybrid Material Point and Discrete Element Method for Granular Media Modeling

Capturing the propagation of microscale physics to macroscale phenomena is intractable for many large systems. This is a major issue in granular media, wherein two extremes are often taken. In one, granular materials are modeled as a continuum, which can be simulated relatively quickly. However continuum models are not always precise and have difficulty capturing certain effects such as particle size dependence. In discrete element methods (DEM), every grain and the interactions between them are simulated. DEM is accurate but solve time scales poorly with large grain numbers. Here, we present a hybrid simulation scheme, which seeks a best-of-both-worlds solution by bridging these two approaches. A mass of granular media is partitioned into three domains: a continuum domain represented using the material point method (MPM), discrete grains using DEM, and a transition zone of both MPM and DEM that are coupled via kinematic constraints. An oracle determines which areas of the domain are MPM and which are DEM, and converts between the two. In the canonical example of silo flow, flow with a sufficiently small orifice jams, resolving length scale dependent effects. Collapse of granular columns modeled with the hybrid method compare quantitatively well with pure discrete simulation and experiments in literature. A significant speedup is seen with the hybrid method over a similar domain of pure discrete grains.

Maytee Chantharayukhonthorn
Massachusetts Institute of Technology
USA
mchantha@mit.edu

Breannan Smith
Columbia University

smithbreannan@gmail.com

Yonghao Yue
The University of Tokyo
yonghao.y@gmail.com

Peter Chen, Eitan Grinspun
Columbia University
cyc@cs.columbia.edu, eitan@grinspun.com

Ken Kamrin
MIT
kkamrin@mit.edu

MS318

CFD-DEM: Modeling the Small to Understand the Large

The ubiquitous tradeoff between fine-grained (microscopic) and coarse-grained (macroscopic) simulations of particle flows is computational speed vs. physical complexity. At one end are the industrial users of such tools, who want to predict units containing billions upon billions of particles within a reasonable time frame and with reasonable computational resources. A multiphase tool analogous to CFD for single-phase flows accurate and robust - is their goal. At the other end are academic-type users, who often add physical complexity, in the form of finer-grained simulations (e.g., Lattice-Boltzmann simulations), until additional complexity is no longer warranted. Here we consider a middle ground CFD-DEM, in which the fluid phase is solved via CFD and particles are tracked via the discrete element method (DEM) and how it can be used by industrial users today. We begin with the results of a recent, industrial consortium survey to better gauge what is meant by reasonable (time-to-solution, computational resources, etc.) for an industrial user. We then present several examples of how todays CFD-DEM tools, despite their computational challenges, can already be used in an industrial setting. Finally, we overview a multi-institution effort toward MFIX-Exa, an open-source CFD-DEM package tailored to future high-performance computation.

Christine Hrenya
University of Colorado, Boulder
USA
hrenya@colorado.edu

MS318

Blood Flow Modelling in Science and Engineering - from Single Vessels to Networks

Understanding the dynamics of red blood cells (RBCs) in complex geometries is crucial for various medical applications. For example, the distribution of RBCs in blood vessel networks affects the oxygen delivery in biological tissues. It is known that solid tumors are more susceptible to radio- and chemotherapy when they are well perfused. However, tumor vessel networks are often irregular, which leads to poor oxygen delivery. We use high-end computer simulations to investigate the distribution of RBCs in complex geometries. The model is based on the lattice-Boltzmann method for fluid flow, coupled with a finite-element method for the dynamics of deformable RBCs. Due to the scale separation of single-cell dynamics (micrometers) and the network dimensions (up to millimeters), we employ HemeLB, a highly parallelized code for fluid dynamics in sparse geometries. We show that the dynamics of RBCs in complex geometries leads to memory effects which in turn change the downstream distribution of RBCs at blood vessel bifurcations. These memory effects are particularly important when distances between consecutive bifurcations are short — such as in tumors rather than in healthy tissue. Our results are therefore relevant to the development of "renormalization" strategies for blood vessel networks to make tumors more susceptible to therapy. We propose a constitutive model for the RBC distribution in complex networks which takes the memory effect into account.

Timm Krueger, Albert Beardo Ricol, Nicholas Daly, Miguel Bernabeu
University of Edinburgh
UK
timm.krueger@ed.ac.uk, albertbeardo@gmail.com, s1530715@sms.ed.ac.uk, miguel.bernabeu@ed.ac.uk

MS318

Modeling of Active Particles and Hydrodynamics: Density Heterogeneities at Multiple Scales

Phytoplankton often encounter turbulence in their habitat. As most toxic phytoplankton species are motile, resolving the interplay of motility and turbulence has fundamental repercussions on our understanding of their own ecology and of the entire ecosystems they inhabit. The spatial distribution of motile phytoplankton cells exhibits patchiness at distances of decimeter to millimeter scale for numerous species with different motility strategies. The explanation of this general phenomenon remains challenging. Furthermore, hydrodynamic cell-cell interactions, which grow more relevant as the density in the patches increases, have been so far ignored. Here, we combine particle simulations and continuum theory to study the emergence of patchiness in motile microorganisms in three dimensions. By addressing the combined effect of turbulent flow conditions, and spatial correlations in the particle positions, we uncover a general mechanism: when motility allows cells to cross the fluid streamlines, the typical length scale associated to the small-scale turbulence selects a characteristic cell-cell interactions scale where strong patches form. Our results shed light on the dynamical characteristics necessary for the formation of patchiness, and complement current efforts to unravel planktonic ecological interactions.

Marco G. Mazza
Loughborough University
UK
m.g.mazza@lboro.ac.uk

MS319

Addressing Exascale Challenges for Molecular Dynamics in GROMACS

Molecular dynamics simulations of biomolecules are an extreme strong scaling challenge, as the system size is fixed,

but billions of time steps need to be computed sequentially. As we approach the exascale, computational hardware is becoming extremely powerful, but the demands its evolution places on the software are also extreme. On the other hand, the goal is never to compute a single number, but always involves determining properties of distributions from ensembles. In this talk, the challenges and scaling approaches in the GROMACS molecular simulation package will be discussed, with a focus on hybrid CPU-GPU acceleration.

Mark Abraham
KTH Royal Institute of Technology in Stockholm
mjab@kth.se

MS319

Scalable Solvers for Dynamic Charge Models in Molecular Dynamics Simulations

Including atom polarizability in molecular dynamics (MD) simulations is important for high-fidelity simulations. Linear solvers for charge models that are used to dynamically determine atom polarizations constitute significant bottlenecks in terms of time-to-solution and the overall scalability of polarizable and reactive force fields. In this talk, we will present novel preconditioning techniques to accelerate the iterative solvers used for several charge models – charge equilibriation (QEq), electronegativity equilization (EE), and atom-condensed Kohn-Sham approximated to second order (ACKS2), and describe a high performance implementation of these techniques in the PuReMD (Purdue Reactive Molecular Dynamics) package. We will also give a detailed analysis of how these novel preconditioners affect solver convergence rate, scaling and the overall performance. Our results show that incomplete Cholesky, incomplete LU with fine-grained parallelism, and sparse approximate inverse-based schemes tuned to produce good quality factors with relatively low construction and application costs yield significant improvements over a baseline Jacobi preconditioner. These results are significant as they enable efficient and high-fidelity molecular simulations on large computer clusters.

H. Metin Aktulga, Kurt O'Hearn, Abdullah Alperen
Michigan State University
hma@cse.msu.edu, ohearnku@msu.edu, abdllh.alprn@gmail.com

MS319

Leveraging Node-level Performance for Molecular Dynamics Through Auto-tuning

Total time to solution in molecular dynamics (MD) simulations is highly sensitive to simulation parameters, including material as well as algorithmic settings. For example, there exist multiple data storage concepts and traversal algorithms, which show advantages in different scenarios. Besides, MD settings can change over the course of the simulation, for example, the homogeneity of the particle distribution. The library AutoPas is an easy to use foundation for building arbitrary N-body simulations aiming for optimal node-level performance. Therefore, AutoPas employs auto-tuning during runtime to find the optimal combination of algorithms for the current state of the simulation and adapts when the situation changes. As a proof of concept, we integrated AutoPas into the established highly scalable MD simulation code ls1 mardyn. We showcase the library's ease of use, scalability, and to what extent ls1 mardyn can benefit from the auto-tuning process. This talk gives an overview of the AutoPas library interface, implemented algorithms, and the auto-tuning process. We are further highlighting early results of the coupling with ls1 mardyn.

Fabio A. Gratl
Technical University of Munich
seckler@in.tum.de

Steffen Seckler
Technical University of Munich
Department of Informatics
steffen.seckler@tum.de

Philipp Neumann
University of Hamburg, Germany
Scientific Computing
philipp.neumann@uni-hamburg.de

Hans-Joachim Bungartz
Technical University of Munich, Department of Informatics
Chair of Scientific Computing in Computer Science
bungartz@in.tum.de

MS319

Comparison of Load-balancing Techniques for MD Simulations with Dynamic Particle Binding

For short-range molecular dynamics simulations with inhomogeneous particle distributions that dynamically change over time, highly flexible and dynamical load balancing methods are mandatory to achieve good parallel scalability. Designing and implementing load balancing algorithms is complex, especially for existing applications which were not designed to support arbitrary domain decompositions. In this talk we present our approach to incorporate general domain decompositions and dynamic re-balancing into the existing software package ESPResSo. Our approach is minimally invasive in the sense that the implementations of models and solvers in ESPResSo remain largely unchanged. The implemented load-balancing mechanisms are based on graph-partitioning, kd-trees, space-filling curves and diffusive load exchange. In our results we show the effect of different kinds of load balancing methods applied to simple and complex scenarios. We report especially on an inhomogeneous, complex multi-scale simulation of soot particle agglomeration which we conducted together with application experts. This scenario has 100 million particles in a turbulent background flow that dynamically bind together with an elaborate model. We show that our mechanisms are capable of reducing the imbalance among processes and the total runtime of simulations and that they enable us to perform inhomogeneous simulations with large particle numbers.

Steffen Hirschmann, Dirk Pflüger
University of Stuttgart

steffen.hirschmann@ipvs.uni-stuttgart.de,
dirk.pflueger@uni-stuttgart.de

Colin W. Glass
HLRS
glassc@hsu-hh.de

MS320

Efficient PCE Representation for Estimation of Rare Events in High Dimensions

In many areas of science and engineering, it is of interest to estimate rare event probabilities involving computationally intensive numerical models. Efficient estimation of the rare event in such cases can be performed by use of surrogate models, such as polynomial chaos expansions (PCE). A major drawback of standard PCE (as with most surrogate models) is the fact that its predictive ability decreases with increase of the dimension of the problem for a fixed computational budget (i.e. number of runs of the computationally intensive model). We address this issue by introducing a linear transformation of the coordinate system of the input random variables and performing a PCE in the transformed coordinate system. Our approach is based on the method described in [R. Tipireddy, R.G. Ghanem, Basis adaptation in homogeneous chaos spaces, J. Comput. Phys. 259 (2014) 304317]. The transformation is based on identifying a set of important directions by application of an approach termed partial least squares (PLS) regression. PLS exploits the (potential) existence of low-dimensional structures underlying the covariance of the input parameters and model response [A. Hskuldsson, PLS regression methods, J. Chemom. 2 (1988) 211-228]. The directions are identified adaptively and the transformed PCE is updated using successive experimental designs obtained from intermediate levels of a sequential importance sampling method.

Iason Papaioannou
Engineering Risk Analysis Group
TU M{ü}nchen
iason.papaioannou@tum.de

Max Ehre
TU München
max.ehre@tum.de

Daniel Straub
Engineering Risk Analysis Group
TU München
straub@tum.de

MS320

From Catastrophe Principle to Strongly Efficient Importance Sampling for Heavy-tailed Rare Events

We propose a class of strongly efficient rare event simulation estimators for random walks and compound Poisson processes with a regularly varying increment distribution in a general large deviations regime. Our estimator is based on an importance sampling strategy that hinges on the heavy-tailed sample path large deviations result recently established in [1]. The new estimators are straightforward to implement and can be used to systematically evaluate the probability of a wide range of rare events with bounded relative errors. They are "universal" in the sense that a single importance sampling scheme applies to a very general class of rare events that arise in heavy-tailed systems. In particular, our estimators can deal with rare events that are caused by multiple big jumps (therefore, beyond the principle of a single big jump) as well as multidimensional processes such as the buffer content process of a queueing network. We illustrate the versatility of our approach with several applications that arise in the context of mathematical finance, actuarial science, and queueing theory. [1] Rhee, C.H., Blanchet, J., and Zwart, B., 2016. Sample Path Large Deviations for Heavy-Tailed Lvy Processes and Random Walks. arXiv preprint arXiv:1606.02795.

Chang-Han Rhee, Bohan Chen
Centrum Wiskunde & Informatica
rhee@cwi.nl, b.chen@cwi.nl

Jose Blanchet
Stanford University
jose.blanchet@stanford.edu

Bert Zwart
Centrum Wiskunde & Informatica
bert.zwart@cwi.nl

MS320

Optimal Sequential Monte Carlo Methods in Rare Event Sampling

Sequential Monte Carlo (SMC) methods, also called splitting methods, are flexible algorithms for rare event simulation that can be easily integrated into scientific modeling software. While SMC methods such as multilevel splitting and Diffusion Monte Carlo are already common in todays simulations, the efficiency of these methods depends critically on a series of user choices. Specifically, users decide when to split, what splitting functions to use, and what resampling schemes to use. The wrong choices can lead to a poor quality of estimation; hence, there is pressing need for a more complete understanding of how to select the best parameters for these algorithms. Our work identifies optimal splitting schedules, splitting functions, and resampling schemes for both multilevel splitting and Diffusion Monte Carlo. When the optimal choices are used, SMC methods are exponentially more efficient than direct sampling. Transitioning theory into practice, we present an example in which optimal SMC is used to generate rare samples of intense tropical cyclones with minimal computing power.

Robert Webber
University of Chicago
rwebber@uchicago.edu

Jonathan Weare
University of Chicago
Department of Statistics and the James Franck Institute
weare@uchicago.edu

MS320

Estimation of Failure Probabilities via Local

MCMC Subset Approximations

The forecasting of rare events is a very formidable task because an improper quantification can lead to a fatal failure. Therefore, it is a critical issue being able to efficiently evaluate the probability of rare events. The classical approaches like Monte Carlo would require a high number of samples to efficiently estimate small failures. Thus, an adaptive simulation approach called Subset Simulation was proposed which exploits intermediate failure events as a product of larger conditional failure probabilities to shift the sampling focus to a lower failure region of interest to increase efficiency. These conditional probabilities are achieved through Markov Chain Monte Carlo (MCMC) approach. The MCMC chain would require many steps to converge within an acceptable error and it is limited by the computational cost of the likelihood. Thus, we built on the work of P. Conrad et al. (2015) by introducing local approximations of a limit state function within the Metropolis-Hastings kernel for the Subset Simulation approach. The ideas are based on deterministic approximation theory, optimization, and experimental design. Following the PhD work by Andrew Davis to reduce significant approximation errors, we employ an optimal refinement rate by the poisedness measurement. We demonstrate the advantages and effectiveness of this approach against classical implementations. The comparisons have been made based on the estimation errors and the computation cost.

Kenan Šehic, Mirza Karamehmedovic
Department of Applied Mathematics and Computer Science
Technical University of Denmark
kense@dtu.dk, mika@dtu.dk

Youssef M. Marzouk
Department of Aeronautics and Astronautics
Massachusetts Institute of Technology
ymarz@mit.edu

MS321

A Fluctuating Boundary Integral Method for Brownian Suspensions

We present a fluctuating boundary integral method (FBIM) for Brownian Dynamics of suspensions of rigid particles of complex shape immersed in a Stokes fluid. We develop a linear-scaling algorithm to generate, together, both the deterministic (mean) component of the particle linear and angular velocities that arise in response to the applied forces and torques, as well as the stochastic (fluctuating) Brownian displacements that arise in response to the thermal fluctuations in the fluid. Our approach relies on a first-kind boundary integral formulation of a Stochastic Stokes Boundary Value Problem in which a random surface velocity is prescribed on the particle surface. This random surface velocity has zero mean and covariance proportional to the Green's function for the Stokes flow (Stokeslet). FBIM provides the key ingredient for time integration of the overdamped Langevin equations for Brownian suspensions of rigid particles. We demonstrate that FBIM obeys discrete fluctuation-dissipation balance by performing equilibrium BD simulations of suspensions of starfish-shaped bodies using a random finite difference temporal integrator.

Yuanxun Bao
The University of Texas at Austin
billbao@cims.nyu.edu

MS321

Domain Decomposition Based Preconditioners for Forward and Inverse Scattering Problems

We consider the use of domain decomposition-based preconditioners for the solution of the integral equation forms of the forward and inverse scattering problem. The solution of those scattering problems can be obtained with an iterative solver. However, in some regimes, the number of iterations necessary for the convergence of the method is extremely high. To mitigate this problem we propose preconditioning techniques to accelerate the iterative solvers. In particular, for the forward scattering problem, we use the Additive Schwarz (and its variations) domain decomposition-based preconditioning techniques for accelerating the iterative solver for the integral equation case. In the inverse scattering problem, the construction of the dense linear system is very costly and depends itself on the solver of the forward problem. In this case, we present an approximation of the operator that uses the forward scattering preconditioner with a low-rank correction. Our numerical results show that both strategies present substantial improvement, providing a sizable decrease in the number of iterations for some regimes and decreasing the number of computational operations necessary to obtain the solution of the problems with the same accuracy.

Carlos Borges
Institute for Computational Engineering and Sciences
University of Texas at Austin
carlos.borges@ucf.edu

George Biros
The Institute for Computational Engineering and Sciences
The University of Texas at Austin
gbiros@acm.org

MS321

A Fast Algorithmic Framework for Dense Rigid Body Suspensions in Stokes Flow

Abstract not available.

Eduardo Corona
Department of Mathematics
University of Michigan
coronae@umich.edu

MS321

Fast Multipole Method in Layered Media

In this paper, we extend fast multipole method for 3D Helmholtz kernel to general layered media. The method is based on the Taylor expansion of layered media Green's function. The key part is an efficient algorithm based on discrete complex image and recurrence formula for the calculation of the Layered media Green's function and its

derivatives which are respresented by Sommerfeld integrals. The new method has $O(N)$ time complexity. Actually, it takes only a few times more time than free space case.

Bo Wang
Department of Mathematics
Southern Methodist University
bowang@mail.smu.edu

Wei Cai
Southern Methodist University
cai@mail.smu.edu

Min Hyung Cho
Department of Mathematical Sciences
University of Massachusetts Lowell
minhyung_cho@uml.edu

Duan Chen
Department of Mathematics and Statistics
University of North Carolina at Charlotte
dchen10@uncc.edu

MS322

Bemfmm: An Extreme-scale Fmm-accelerated Solver Based on Boundary Element Method for the 3D Complex Helmholtz Equation

This work presents a highly scalable FMM-accelerated linear solver for oscillatory kernels. We develop an extreme scale boundary integral solver for wave scattering that uses FMM as a matrix-free accelerator for the underlying GMRES linear iterative solver. Our FMM-based Helmholtz kernels provides a stable treatment for hyper-singular and near-field integration points. The kernels are highly optimized for both shared- and distributed-memory, through which the potential thread- and data-level parallelism are leveraged. We provide various performance models for tuning the task-based tree traversal implementation of FMM, and develop optimal architecture-specific and algorithmic-aware partitioning, load balancing, and communication reducing mechanisms. The enhancements result in scaling the solver up to 6,144 compute nodes of a Cray XC40 with 196,608 hardware cores. We report near-optimal efficiency in the weak scalability study with respect to both the O(log P) communication complexity as well as the theoretical scaling complexity of FMM. We compute in excess of 2 billion Degrees-of-freedom on the full-scale of the Cray XC40 supercomputer. The numerical results match the analytical solution with convergence at 1.0e-4 relative 2-norm residual accuracy.

Noha Al-Harthi
King Abdullah University of Science and Technology
noha.harthi@kaust.edu.sa

Mustafa Abduljabbar
Extreme Computing Research Center
King Abdullah University of Science and Technology
mustafa.abduljabbar@kaust.edu.sa

Mohammed A. Al Farhan
KAUST
mohammed.farhan@kaust.edu.sa

Rui Chen
King Abdullah University of Science and Technology
rui.chen@kaust.edu.sa

Rio Yokota
Tokyo Institute of Technology
rioyokota@gsic.titech.ac.jp

Hakan Bagci
King Abdullah University of Science and Technology
hakan.bagci@kaust.edu.sa

David E. Keyes
King Abdullah University of Science and Technology (KAUST)
david.keyes@kaust.edu.sa

MS322

Iterative Solvers for Modeling Ocean Waves Using Helmholtz-like Equation

Propagation of waves from deep to shallow sea can be modeled by the mild-slope equation. This equation is capable of estimating transformation of (linear) waves in sea with impermeable bottom into various forms. The so-called modified time-independent (harmonic) mild-slope model is given by the PDE

$$\nabla_h(I_1\nabla_h\phi) + \left[(k^2 + \hat{i}\sigma D)I_1 + r(h)\right]\phi = 0,$$

with ϕ be the velocity potential, $\nabla_h = \left\langle \frac{\partial}{\partial x}, \frac{\partial}{\partial y} \right\rangle$, $r(h) = I_2\nabla_h^2 h + (\nabla_h h)^2 \left(\frac{\partial I_2}{\partial h} - I_3 \right)$, and with I_1, I_2, I_3 some integrals over the depth z. Introducing $\psi = \sqrt{I_1}\phi$, the above equation can be reformulated as a Helmholtz equation:

$$\nabla_h^2\psi + K_c x^2 \psi = 0, \tag{5}$$

where K_c depends nonlinearly on the wavenumber k, $r(h)$, and I_1. In this talk, we shall discuss and present iterative solution methods for Equation (1) iteratively. The equation is discretized using a finite difference method, with (non-standard) boundary conditions that mimic propagation of waves in a finite computational domain, which leads to a non-local condition. The iterative method is a Krylov method preconditioned by shifted Laplacian, which is adapted for this equation.

Yogi Erlangga, Rysbergen Tabarek
Nazarbayev University
yogi.erlangga@nu.edu.kz, rysbergen.tabarek@nu.edu.kz

MS322

Domain Decomposition Preconditioning for Frequency Domain Wave Problems

We describe one-level additive Schwarz preconditioners for the Helmholtz equation (with increasing wavenumber k), when discretized using nodal conforming finite elements of any fixed order. The preconditioners combine independent local solves (with impedance boundary conditions) on overlapping subdomains of diameter H, and

prolongation/restriction operators defined using a partition of unity. In numerical experiments we observe robust (i.e. wavenumber independent) GMRES convergence, even when H decreases to zero (as k increases). This provides a highly parallel and robust one-level domain decomposition method. We provide supporting theory for this observation by studying the preconditioner when applied to a range of absorptive problems, and including the "pure' Helmholtz case with no absorption. The talk will also contain applications of related preconditioning strategies to Maxwell and elastic wave problems in the frequency domain. The talk includes joint work with collaborators M. Bonazzoli, V. Dolean, E. Spence, P.-H. Tournier, E. Vainikko and J. Zou.

Ivan G. Graham
University of Bath
I.G.Graham@bath.ac.uk

MS322

Fast and Scalable Solvers for High-frequency Wave Propagation

In this talk, I will review recent progress on fast solvers for the high-frequency Helmholtz equation. The problem is harder than in the elliptic case, and the better answers seem to involve a decomposition into polarized waves. I will describe how the method of polarized traces fits in this framework, and how it can be adapted to lead to a complexity that is sub-linear in both the number of volume unknowns and the number of right-hand sides, in the 3D case, in a parallel environment. Joint work with Laurent Demanet, Matthias Taus, Adrien Scheuer, and Leonardo Zepeda.

Matthias Taus
University of Texas
matthias.taus@gmail.com

Leonardo Zepeda-Nunez
Lawrence Berkeley National Laboratory
lzepeda@lbl.gov

Russell Hewett
Total
russell.hewett@total.com

Laurent Demanet
Department of Mathematics, MIT
laurent@math.mit.edu

MS323

An Overlapping Decomposition and Heterogeneity-homogeneity Adaptation for Wave Propagation in Unbounded Media

Finite element methods (FEM) are widely used for bounded heterogeneous media models, and boundary element methods (BEM) are efficient for simulating wave propagation in unbounded homogeneous media. A natural medium for wave propagation comprises a coupled bounded heterogeneous domain and an unbounded homogeneous free space. Frequency-domain wave propagation models in the medium, such as the variable coefficient Helmholtz equation, include a faraway decay radiation condition (RC). The FEM wave propagation models are based on truncating the unbounded region and approximating the RC. The BEM models do not take into account of the heterogeneity in the medium. It is desirable to develop algorithms that incorporate the full physics of the heterogenous and unbounded medium wave propagation model, and avoid approximation of the RC. In this work we develop, analyze, and implement an overlapping decomposition framework for the heterogeneity-homogeneity adaptation and a high-order FEM-BEM algorithm for simulation of the Helmholtz equation in unbounded media with a spatially dependent refractive index and our method exactly satisfies Sommerfeld RC.

Mahadevan Ganesh
Colorado School of Mines
mganesh@mines.edu

Victor Dominguez
Universidad Publica de Navarra
Spain
victor.dominguez@unavarra.es

Francisco J. Sayas
Department of Mathematical Sciences
University of Delaware
fjsayas@math.udel.edu

MS323

An Efficient Domain Decomposition Method with Cross-point Treatment for Helmholtz Problems

Solving high-frequency time-harmonic scattering problems using finite element techniques is challenging, as such problems lead to very large, complex and indefinite linear systems. Optimized Schwarz domain decomposition methods (DDMs) are currently a very promising approach, where subproblems of smaller sizes are solved in parallel using direct solvers, and are combined in an iterative procedure. It is well-known that the convergence rate of these methods strongly depends on the transmission condition enforced on the interfaces between the subdomains. Local transmission conditions based on high-order absorbing boundary conditions (HABCs) have proved well suited. They represent a good compromise between basic impedance conditions (which lead to suboptimal convergence) and the exact Dirichlet-to-Neumann (DtN) map related to the complementary of the subdomain (which is expensive to compute). However, a direct application of this approach for domain decomposition configurations with cross-points, where more than two subdomains meet, does not provide satisfactory results. We present an improved DDM that efficiently addresses configurations with cross points. Noting that these points actually are corners for the subdomains, our strategy consists in incorporating a corner treatment developed for HABCs into the DDM procedure. After a presentation of the key aspects of the methods, the effectiveness of our approach is discussed with finite element results.

Axel Modave
CNRS - ENSTA - INRIA

axel.modave@ensta-paristech.fr

Xavier L. Antoine
Institut Elie Cartan Nancy (IECN)
Université de Lorraine
xavier.antoine@univ-lorraine.fr

Christophe Geuzaine
University of Liège
cgeuzaine@uliege.be

MS323

A Robust Solver for Elliptic PDEs in 3D Complex Geometries

Solving partial differential equations via boundary integral techniques has many benefits, such high-order accuracy, optimal asymptotic runtimes and the capacity to robustly handle complex geometries. Accurate evaluation of singular and near-singular layer potentials is an important step in such algorithms. Quadrature by expansion (QBX) a method that accurately approximates the solution locally near the boundary by a series expansion. We present a fast extension of this algorithm to three dimensions on complex geometries. Our algorithm is a simple extrapolation of the solution evaluated at auxiliary points in the domain and hence kernel-independent by relying purely on kernel evaluations. A fully-adaptive refinement process is used to ensure high accuracy on complex geometric domains while maintaining efficiency. We present numerical results on various geometries for the Laplace, Stokes and Navier equations.

Abtin Rahimian
University of Colorado Boulder
abtin.rahimian@colorado.edu

Matthew J. Morse
Courant Institute - New York University
mmorse@cs.nyu.edu

Denis Zorin
Computer Science Department
Courant Institute, New York University
dzorin@cs.nyu.edu

MS323

Basis Adaptation and Domain Decomposition for a Class of Stochastic Models

We build on our earlier work to extend stochastic basis adaptation and domain decomposition methods for a class of stochastic models represented by time varying stochastic partial differential equations (PDEs). Basis adaptation method developed for computing a scalar quantity of interest (QoI) cannot directly applied for time varying problems, because small approximation error at each time step is accumulated resulting in a large deviation from true solution over a period of time. To reduce these errors, we decompose the spatial domain into a set of non-overlapping subdomains and use Hilbert Karhunen loeve method for basis adaption in each subdomain. We use polynomial chaos based uncertainty quantification (UQ) methods to solve stochastic PDE in each subdomain in local stochastic basis. We solve the local solution in each subdomain independently of each other while the continuity conditions for the solution and flux across the interface of the subdomains is maintained. Neumann-Neumann based algorithm is used to compute the solution in the interior and at the interface of the subdomains. Continuity in the stochastic solution is imposed by projecting local solution in each subdomain onto a common basis. Instead of finding low dimensional basis at each time step, we adaptively choose time steps at which basis adaptation is required to avoid large deviation in the stochastic solution. We present numerical experiments in support of our proposed method.

Ramakrishna Tipireddy, Panos Stinis, Alexander Tartakovsky
Pacific Northwest National Laboratory
ramakrishna.tipireddy@pnnl.gov, panos.stinis@pnnl.gov, alexandre.tartakovsky@pnnl.gov

MS324

Bridging the Gap between Flat and Hierarchical Low-rank Matrix Formats

Matrices possessing a low-rank property arise in numerous scientific applications. This property can be exploited to provide a substantial reduction of the complexity of their LU or LDLT factorization. Among the possible low-rank matrix formats, the flat Block Low-Rank (BLR) format is easy to use but achieves superlinear complexity. Alternatively, the hierarchical formats achieve linear complexity at the price of a much more complex hierarchical matrix representation. In this talk, we propose a new format based on multilevel BLR approximations. Contrarily to hierarchical matrices, the number of levels in the block hierarchy is fixed to a given constant; we prove that this provides a simple way to finely control the desired complexity of the dense multilevel BLR factorization. By striking a balance between the simplicity of the BLR format and the low complexity of the hierarchical ones, the multilevel BLR format bridges the gap between flat and hierarchical low-rank formats. The multilevel BLR format is of particular relevance in the context of sparse (e.g. multifrontal) solvers, for which it is able to trade off the optimal dense complexity of the hierarchical formats to benefit from the simplicity of the BLR format while still achieving O(n) sparse complexity.

Patrick Amestoy
INP(ENSEEIHT)-IRIT
amestoy@enseeiht.fr

Alfredo Buttari
CNRS-IRIT-Université de Toulouse, France
alfredo.buttari@enseeiht.fr

Jean-Yves L'Excellent
INRIA-LIP-ENS Lyon
jean-yves.l.excellent@ens-lyon.fr

Theo Mary
Universite de Toulouse, UPS-IRIT

theo.mary@manchester.ac.uk

MS324

On the Inversion of BLR Matrices from Acoustic Fluid-structure Interactions

We consider the modeling of acoustic fluid-structure interactions using a boundary element method, doubly asymptotic approximations (DAA), that gives rise to several large dense matrices which we store and manipulate as block low rank (BLR) matrices. While the time stepping method requires only matrix-vector multiplies with two BLR matrices, computing these two matrices requires the inversion of a BLR matrix. Our present MPI code, BLR-MPP, solves a linear system by computing a $A = LDL^T$ or $A = LDU$ factorization. Our code is bulk synchronous. Several varieties of communication strategies are supported. The fundamental computational tasks are dense and low rank matrix-matrix and matrix-vector multiplies, dense factor tasks, and dense solve tasks with one, several and many right hand sides. The computations are driven by a program, a representation of the global task DAG that is distributed across processors as needed. There is a program for a LDL^T factorization, one for a LDU factorizations, for solves involving combinations of L, D and U. All use the same computational engine. We describe the task DAG to invert an upper block diagonal matrix U, where the diagonal blocks $U_{J,J}$ are the identity. The program that represents this task DAG is input to the BLR-MPP software, which computes the inverse of this upper block triangular BLR matrix.

Cleve Ashcraft
Livermore Software Technology Corporation
cleve@lstc.com

Francois-Henry Rouet
Livermore Software Technology Corp.
fhrouet@lstc.com

Clement Weisbecker
Livermore Software Technology Corporation
clement@lstc.com

MS324

Block Low-rank Algebraic Clustering for Sparse Direct Solvers

We will discuss challenges in building clusters for the Block Low-Rank (BLR) approach, for nodes inside separators appearing during the factorization of sparse matrices. We will illustrate limitations for methods that consider only intra-separators connectivity (i.e., k-way and recursive bisection) as well as methods focusing only on reducing the number of updates between separators. The new strategy we propose considers interactions between a separator and its children in the nested dissection. It allows reducing the computational cost of BLR, and the number of off-diagonal blocks. We demonstrate that this method enhances the BLR strategies in the sparse direct supernodal solver PaStiX, and discuss how it can be extended to low-rank formats with more than one level of hierarchy.

Gregoire Pichon
INRIA
gregoire.pichon@inria.fr

Eric F. Darve
Institute for Computational and Mathematical Engineering
Stanford University
eric.darve@stanford.edu

Mathieu Faverge
Bordeaux INP - Labri - Inria
mathieu.faverge@inria.fr

Pierre Ramet
Bordeaux University - INRIA
pierre.ramet@inria.fr

Jean Roman
INRIA
Jean.Roman@inria.fr

MS324

Distributed Memory Lattice H-Matrix Factorization

In this talk, we study a lattice H-matrix format that generalizes the BLR format by storing each of the blocks, referred to as lattices, in the H-matrix format. This new format aims to combine the parallel scalability of BLR with the near linear complexity of H-matrices. Our performance results on the distributed-memory computers demonstrate that compared with BLR, the lattice format may reduce the cost of factorization and improve the factorization performance.

Ichitaro Yamazaki
University of Tennessee, Knoxville
iyamazak@icl.utk.edu

Akihiro Ida
The University of Tokyo
ida@cc.u-tokyo.ac.jp

Rio Yokota
Tokyo Institute of Technology
rioyokota@gsic.titech.ac.jp

Jack J. Dongarra
University of Tennessee, Oak Ridge National Laboratory, USA
dongarra@icl.utk.edu

MS325

Embedded Error Estimation and Adaptive Stepsize Control for Optimal Explicit Strong Stability Preserving RungeKutta Methods

We construct a family of embedded pairs for optimal strong stability preserving explicit RungeKutta methods of order $2 \leq p \leq 4$ to be used to obtain numerical solution of spa-

tially discretized hyperbolic PDEs. In this construction, the goals include non-defective methods, large region of absolute stability, and optimal error measurement. The new family of embedded pairs offer the ability for strong stability preserving (SSP) methods to adapt by varying the step-size based on the local error estimation while maintaining their inherent nonlinear stability properties. Through several numerical experiments, we assess the overall effectiveness in terms of precision versus work while also taking into consideration accuracy and stability.

Sidafa Conde
Mathematics Department
University of Massachusetts Dartmouth
sconde@sandia.gov

Imre Fekete
Eötvös Loránd University
feipaat@cs.elte.hu

John Shadid
Sandia National Laboratories
Albuquerque, NM
jnshadi@sandia.gov

MS325

Strong-stability-preserving Runge-Kutta Methods with Downwind-biased Operators

Strong stability preserving (SSP) time integrators have been developed to preserve certain nonlinear stability properties (e.g., monotonicity, boundedness) of the numerical solution in arbitrary norms, when coupled with suitable spatial discretizations. Currently, all existing general linear methods (including Runge-Kutta and linear multistep methods) either attain small step sizes for nonlinear stability, or they are only first order accurate. In order to obtain larger step sizes spatial discretizations of PDEs that contain both upwind- and downwind-biased operators have been employed. In this talk, we review SSP Runge-Kutta methods that use upwind- and downwind-biased discretizations in the framework of perturbed Runge-Kutta methods. We show how downwinding improves the SSP properties of time-stepping methods and breaks some order barriers. In particular, we focus on implicit perturbed SSP Runge-Kutta methods that have unbounded SSP coefficient. We present a novel one-parameter family of three-step, third-order perturbed Runge-Kutta methods, for which the CFL-like step-size restriction can be arbitrarily large. The stability of this family of methods is analyzed, and we demonstrate that the desired order of accuracy is obtained for large CFL numbers.

Yiannis Hadjimichael
King Abdullah University of Science and Technology (KAUST)
hadjimy@cs.elte.hu

MS325

Strong Stability Preserving Integrating Factor Runge-Kutta Methods

Strong stability preserving (SSP) Runge–Kutta (RK) methods are often desired when evolving problems with two components that have very different time scales. Where the SSP property is needed, implicit and implicit-explicit methods are inefficient since they have very restrictive time-steps. For this reason, SSP integrating factor methods offer an attractive alternative to traditional time-stepping methods for problems with a stiff linear component and a not stiff nonlinear component. However, the strong stability properties of integrating factor RK methods have not been established. This work shows it is possible to define explicit integrating factor RK methods that preserve the strong stability properties satisfied by each of the components when coupled with forward Euler time-stepping. We define sufficient conditions for an explicit integrating factor RK method to be SSP, namely that they are based on explicit SSP RK methods with non-decreasing abscissas. We find such methods up to fourth order and up to ten stages and analyze their SSP coefficients. These methods are tested to demonstrate their convergence and to show that the SSP time-step predicted by the theory is generally sharp. On typical total variation diminishing linear and nonlinear test-cases our explicit SSP integrating factor RK methods out-perform the corresponding explicit SSP RK methods, implicit-explicit SSP RK methods, and some exponential time differencing methods.

Leah Isherwood
University of Massachusetts Dartmouth
lisherwood@umassd.edu

Sigal Gottlieb
Department of Mathematics
University of Massachusetts Dartmouth
sgottlieb@umassd.edu

Zachary Grant
University of Massachusetts, Dartmouth
zack.j.grant@gmail.com

MS325

Explicit Partitioned Methods for Interface Coupling

Partitioned algorithms enable effective simulation of physical processes comprising multiple mathematical models by using separate analysis codes for each constituent component. Examples of this include the coupling between atmosphere and ocean codes in Earth system models as well as welded contact problems between different elastic materials. Partitioned algorithms allow reuse of existing investments in simulation codes as well as the use of the code having the appropriate material models in their respective subdomains. In this talk, we present a partitioned explicit approach tailored to both advection-diffusion and elastodynamics for spatially coincident but mismatched interfaces, based on a generalized Schur complement. To obtain a Lagrange multiplier formulation that is fully compatible with explicit time integration, we start with an alternative coupling condition, which enforces the continuity of the time derivatives of the states across the interface. The dual Schur complement of the resulting monolithic problem then yields an explicit system for the interface flux, which can be used to drive the independent solution of the subdomain problems for the next time increment. The al-

gorithm will be presented along with an overview of a proof of the method's well-posedness. Numerical results for the elastodynamics and advection-diffusion equations demonstrate the stability of the formulation and its second-order rate of convergence.

Kara Peterson
Sandia Natl. Labs
kjpeter@sandia.gov

Pavel Bochev
Sandia National Laboratories
Computational Math and Algorithms
pbboche@sandia.gov

Paul Kuberry
Sandia National Laboratories
pakuber@sandia.gov

MS326

Extending the Added-mass Partitioned (AMP) Scheme for Solving FSI Problems Coupling Incompressible Flows with Elastic Beams to 3D

A new partitioned algorithm was recently developed for solving fluid-structure interaction (FSI) problems coupling incompressible flows with elastic beams undergoing finite deformations in 2D. The new algorithm, referred to as the Added-Mass Partitioned (AMP) scheme, overcomes the added-mass instability that has for decades plagued partitioned FSI simulations of incompressible flows coupled to light structures. The AMP scheme achieves fully second-order accuracy and remains stable, without sub-time-step iterations, even for very light structures when added-mass effects are strong. The stability and accuracy of the AMP scheme is validated through mode analysis and numerical experiments. In this talk, we will describe the ongoing progress and challenge on extending the AMP scheme to the 3D regime.

Longfei Li
University of Louisiana at Lafayette
longfei.li@louisiana.edu

MS326

A Fourth-order Accurate Fractional-step Scheme for the Incompressible Navier-stokes Equations on Overlapping Grids

A fourth-order accurate fractional-step scheme for the incompressible Navier-Stokes equations in two and three dimensions is discussed. The scheme is based on the velocity-pressure form of the equations. The velocity is advanced in a separate stage from the update for the pressure leading to an efficient *fractional-step* (or *split-step*) scheme. The choice of numerical boundary conditions for each stage of the fractional-step scheme is critical for the accuracy and stability of the approach. The equations are discretized in space using a fourth-order accurate finite difference approximation. The time-stepping method is either an explicit fourth-order accurate predictor-corrector method or a fourth-order accurate IMEX-BDF scheme. For robust solutions of under-resolved problems, a WENO-based artificial diffusion named BWENO is developed. Complex domain geometries are handled efficiently using composite overlapping grids. The stability and accuracy of the new scheme is verified for a variety of two and three-dimensional problems using manufactured solutions and known solutions.

Donald W. Schwendeman
Rensselaer Polytechnic Institute
Department of Mathematical Sciences
schwed@rpi.edu

MS326

A Stable Added-mass Partitioned (AMP) Algorithm for Elastic Solids and Incompressible Flows

A stable added-mass partitioned (AMP) algorithm is developed for fluid-structure interaction (FSI) problems involving viscous incompressible fluids and compressible elastic solids. In order to efficiently address complex moving geometries, the fluid and solid are discretized using deforming composite grids (DCGs). In the AMP algorithm, each domain is solved in separate stages and linked together using Robin interface conditions motivated from the solid characteristic structure. Unlike traditional partitioned FSI algorithms, the AMP scheme does not require sub-iterations to stabilize when added-mass effects are large (ie. light solids or fine meshes). The AMP scheme is shown to be stable for any mass ratio of the fluid and solid using normal-mode analysis applied to a few carefully selected model problems. Numerical examples are considered to showcase the robustness of the AMP algorithm and to demonstrate interesting FSI applications.

Daniel A. Serino
Department of Mathematical Sciences
Rensselaer Polytechnic Institute
serind@rpi.edu

MS326

A Partitioned FSI Algorithm for Rigid Bodies and Incompressible Flow Based on Potentials

We consider FSI problems that involve rigid bodies and viscous incompressible flows. A novel approach based on potentials is proposed to expose analytical added-mass and added-damping forms in a general FSI problem. Two sets of potential equations are derived, and some properties of the added-mass and added-damping matrices are proved. We further use model problems to validate this potential framework through exactly exposing their added-mass and added-damping forms. Based on this framework, we propose a partitioned FSI scheme that can be used to efficiently solve strongly coupled FSI problems in a truly partitioned fashion. Stability of the new scheme is analyzed using model problems. Finally, some benchmark problems are used to demonstrate the stability and accuracy of the scheme.

Qi Tang
Department of Mathematical Sciences
Rensselaer Polytechnic Institute

tangqimath@gmail.com

MS327

Accurate Estimation of High-Dimensional Transport Maps

Bayesian filtering algorithms for systems with nonlinear dynamics and high-dimensional states include ensemble transform filters and particle filters. Within these algorithms, it is often necessary to construct montone nonlinear transformations of non-Gaussian distributions—viewed as transport maps—from relatively few samples. We present an algorithm for accurately approximating such maps when the number of available samples is small relative to the dimension of the state. We explore the effect of different regularization methods for estimating maps under marginal and conditional independence assumptions on the distribution, or under sparsity or low-rank assumptions on the map. We analyze the resulting approximation error in the map estimate and in its associated pushforward/pullback distributions, and demonstrate the numerical performance of our map-learning algorithm in the context of model problems for nonlinear filtering in numerical weather prediction.

Ricardo Baptista
MIT
rsb@mit.edu

Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu

MS327

Langevin Monte Carlo and JKO Splitting

Algorithms based on discretizing Langevin diffusion are popular tools for sampling from high-dimensional distributions. We develop novel connections between such Monte Carlo algorithms, the theory of Wasserstein gradient flow, and the operator splitting approach to solving PDEs. In particular, we show that a proximal version of the Unadjusted Langevin Algorithm corresponds to a scheme that alternates between solving the gradient flows of two specific functionals on the space of probability measures. Using this perspective, we derive some new non-asymptotic results on the convergence properties of this algorithm.

Espen Bernton
Harvard University
ebernton@g.harvard.edu

MS327

Optimal Transport on Surfaces, Graphs, and Point Clouds

Dating back to early economic applications, the optimal transport (OT) problem has received renewed interest thanks to its applicability to problems in machine learning, computer graphics, geometry, and other disciplines. The main barrier to wide adoption of OT as a modeling tool is the expense of optimization in OT problems. In this talk, I will summarize efforts in my group to make large-scale transport tractable over a variety of domains and in a variety of application scenarios, helping transition OT from theory to practice.

Justin Solomon, Sebastian Claici, Edward Chien
MIT
jsolomon@mit.edu, sclaici@csail.mit.edu, edchien@mit.edu

Gabriel Peyré
Ceremade CNRS-Université Paris-Dauphine
peyre@ceremade.dauphine.fr

Marco Cuturi
ENSAE / CREST
marcocuturicameto@gmail.com

MS327

Conditional Density Estimation and Simulation through Optimal Transport

A methodology to estimate from samples the probability density of a random variable x conditional to the values of a set of covariates zl is proposed. The methodology relies on a data-driven formulation of the Wasserstein barycenter, posed as a minimax problem in terms of the conditional map carrying each sample point to the barycenter and a potential characterizing the inverse of this map. This minimax problem is solved through the alternation of a flow developing the map in time and the maximization of the potential through an alternate projection procedure.

Giulio Trigila
Baruch college - CUNY
giulio.trigila@baruch.cuny.edu

Esteban G. Tabak
Courant Institute
New York University
tabak@cims.nyu.edu

Michele Pavon
University of Padova
pavon@math.unipd.it

MS328

Combining in vitro Data and Mathematical Models to Understand Chemokine Biology: A Bayesian Parameter Inference Approach

All protective and pathogenic immune and inflammatory responses rely heavily on leukocyte migration and localisation. Chemokines are secreted chemoattractants that orchestrate the positioning and migration of leukocytes through concentration gradients. However, understanding how these gradients are formed, maintained and regulated remains unclear. We present a mathematical and computational framework that integrates experimental, computational and modelling approaches. It accounts for dynamic interactions between physical, biological and biochemical processes modelled through diffusion, advection, binding kinetics and cellular uptake leading to the establishment of spatio-temporal chemokine gradients. Numerical optimisation techniques and computational bayesian

methods are implemented to infer underlying parameters as well as to quantify uncertainty of the computational model. This study shows how a synergy between experimental and computational modelling approaches can be used to decipher elementary processes of chemokine transport, chemokine-cell interactions and cell migration. It provides a simpler link with experimental data from different experimental systems, thereby setting the stage for new computer-assisted approaches to gain holistic, insightful understanding and generate novel and testable hypotheses in chemokine biology.

Dimitris I. Kalogiros, Matthew Russell
University of Nottingham
dimitris.kalogiros@nottingham.ac.uk,
matthew.russell@nottingham.ac.uk

Willy Bonneuil, Jennifer Frattolin
Imperial College London
willy.bonneuil15@imperial.ac.uk,
j.frattolin@imperial.ac.uk

Mikaila J. Bandara, Francesca F. Masci, Robert Nibbs
University of Glasgow
mikaila.jayaweerabandara@glasgow.ac.uk,
francesca.fasanellamasci@glasgow.ac.uk,
robert.nibbs@glasgow.ac.uk

James E. Moore Jr.
Imperial College London
james.moore.jr@imperial.ac.uk

Bindi S. Brook
University of Nottingham
bindi.brook@nottingham.ac.uk

MS328

Healthcare by Design: How Computational Approaches Are Changing the Clinical Landscape

Chronic diseases are a source of crippling burden to society and economy. Treating these diseases, such as diabetes, stroke, asthma, remains a grand challenge in healthcare. As many of these diseases arise due to breakdown of tissues maintained by stem cells, corrective measures include autologous cell populations or tissues. Another alternative is to offer personalised courses of treatment to the patient. While the former requires a quantitative understanding of developmental milestones and how they can be exploited ex vivo, the latter requires the ability to quantitatively predict the interactions between therapeutic molecules and the human system. This continues to remain a significant gap, with profound implications to healthcare costs. Recent advances in computational multi-scale modelling, however, has made tangible inroads in both directions. These include computational methods that can capture the dynamically reciprocal interactions between cells within artificial environments as well as clinically-relevant tissue-level models that can predict the impact of biochemical therapies at the patient level. Here, I review these computational advancements and how they are ushering in a new era of healthcare interventions by impacting drug design, informing ex vivo histogenesis, and enabling patient-specific modelling.

Himanshu Kaul
University of British Columbia
himanshu.kaul@ubc.ca

Yiannis Ventikos
Mechanical Engineering
University College London
y.ventikos@ucl.ac.uk

Rod Smallwood
University of Sheffield ,UK
r.smallwood@sheffield.ac.uk

Chris Brightling
University of Leicester
ceb17@leicester.ac.uk

MS328

A Statistical Mechanical Perspective of Modeling Phase Transitions in Living Systems

Network biology is a strongly emerging paradigm in biology and biomedicine with the aim to analyze and explain a huge class of complex phenomena such as benign and malignant disease progression. They depend significantly on the large scale interaction of the biological entities and can be interpreted by means of phase-transitions. Here we present a generalization of the Ising model that reflects biological mechanisms and can be used to model cooperative phenomena in scale-free gene interaction networks. Using Markov chain, we observe that phase transition behavior in such a system is strongly influenced by properties unique to our generalized Ising-gene system resulting in a so-called phase transition of first order. Our numerical observations are validated by Mean-Field Theory for a Barabasi-Albert Network. We present conditions under which a E.Coli Gene Regulatory Network undergo phase transition in response to stress. We also present alternate methods that can approximate the network behavior.

Jeyashree Krishnan
Aachen Institute for advanced study in Computational Enginee
RWTH Aachen University
krishnan@aices.rwth-aachen.de

Reza Torabi
Department of Physics and Astronomy, University of Calgary
reza.torabi@ucalgary.ca

Edoardo Di Napoli
AICES, RWTH Aachen
dinapoli@aices.rwth-aachen.de

Andreas Schuppert
AICES/JRC, RWTH Aachen University

schuppert@aices.rwth-aachen.de

MS328

Modeling of Disease Progression in Cancer

Modeling the parameters of multistep carcinogenesis is key to a better understanding of cancer progression, biomarker identification and the design of individualized therapy. The dynamics of cancer progression is known to be a complex process driven by stochastic cascades of genomic instabilities and breakdown of physiological control mechanisms. Hence one-parameter concepts for cancer evolution prediction in individual patients had limited success. Complex models representing non-linear interactions between the relevant parameters driving cancer progression are expected as the method to foster precision medicine in cancer therapies. Today the main focus of disease modeling is on mechanistic modeling, where biological mechanisms are represented by explicit equations enabling insight into the cooperation of all relevant mechanisms. Data-driven modeling aims to compensate this by use of advanced statistical and machine learning methods. As parameters of disease progression models can rarely be assessed in patients in clinical standard procedures, a remaining challenge is the matching of patients to the model parameters when disease history is not known. Hence multiple independent monitoring patterns are required for individualized predictions. Using chronic myeloid leukemia as a paradigm for hierarchical disease evolution we show that combined population dynamic modeling and CML patient biopsy genomic analysis enables improved patient prognosis.

Andreas Schuppert
AICES/JRC, RWTH Aachen University
schuppert@aices.rwth-aachen.de

MS329

A Discontinuous Galerkin Method for the Aw-rascle Traffic Flow Model on Networks

In this talk we consider the second-order Aw-Rascle (AR) model for traffic flow on a network, and propose a discontinuous Galerkin (DG) method for solving the AR system of hyperbolic PDEs with appropriate coupling conditions at the junctions. For the proposed method we apply the Lax-Friedrichs flux, and for comparison, we use the first-order Lighthill-Whitham-Richards (LWR) model with the Godunov flux. Coupling conditions are also required at the junctions of the network for the problem to be well-posed. As the choice of coupling conditions is not unique, we test different coupling conditions for the Aw-Rascle model at the junctions. Numerical examples are provided to demonstrate the high-order accuracy, as well as comparisons between the first-order LWR model and the second-order AR model. The ability of the AR model to capture the capacity drop phenomenon is also explored.

Joshua Buli
University of California
jbuli001@ucr.edu

Yulong Xing
The Ohio State University
xing.205@osu.edu

MS329

High-order Discontinuous Galerkin Schemes for Water Mass Transport on Curved Surface

The particular water-mass transport equation that we consider is the thin-film equation. The thin-film equation results from an asymptotic limit of the Navier-Stokes equations and describes how a thin-film of water flows over a surface. The thin-film equation contains both a convection and diffusion term and thus is most often handled with operator splitting. We describe how to use an explicit Runge-Kutta DG method for solving the convection equation and an implicit DG method for solving the diffusion equation. This implicit problem is particularly ill-conditioned, therefore we demonstrate how a multigrid method can ensure fast convergence that scales well with the size of the model.

Caleb D. Logemann
Iowa State University
logemann@iastate.edu

MS329

A Discontinuous Galerkin Method for Gradient Flows

We consider a class of time-dependent second-order partial differential equations with a gradient flow structure. The problem is governed by a decaying entropy. The solution usually corresponds to a density distribution, hence positivity (non-negativity) is expected. This class of equations covers important cases such as Fokker-Planck type equations and aggregation models, which have been studied intensively in the past decades. In this talk, we propose a high-order discontinuous Galerkin method for these equations. If the interaction potential is not involved, or the interaction is defined by a smooth kernel, our semi-discrete scheme admits an entropy inequality on the discrete level. Furthermore, by applying the positivity-preserving limiter, our fully discretized scheme produces non-negative solutions for all cases under a time step constraint. Numerical examples are given to confirm the high-order accuracy for smooth test cases and to demonstrate the effectiveness for preserving long time asymptotics.

Zheng Sun
Brown University
zheng_sun@alumni.brown.edu

Jose Carrillo
Department of Mathematics
Imperial College London
carrillo@imperial.ac.uk

Chi-Wang Shu
Brown University
Div of Applied Mathematics

shu@dam.brown.edu

MS329

High-order Bound-preserving Discontinuous Galerkin Method for Wormhole Propagation Model

Wormhole propagation model is a set of PDEs which arises in petroleum engineering. This model is used to describe the increase of porosity and the distribution of acid concentration in carbonate reservoir under erosion of injected acid. The important physical features of porosity and acid concentration include their boundedness between 0 and 1, as well as the monotone increasing for porosity. The issue of how to keep these properties in the simulation is crucial to the robustness of the numerical algorithm. In our project, we propose a high-order bound-preserving DG method to keep these important physical properties. The main technique is to introduce a new variable r to replace of the original acid concentration variable and use a high-order slope limiter to keep a polynomial upper bound which changes over time for r. Furthermore, the high-order accuracy is attained by the flux limiter. In this talk, we present some numerical examples to show the high-order accuracy for all variables, boundness for porosity and acid concentration, and monotone-increasing for porosity under our method.

Ziyao Xu
Michigan Technological University
ziyaox@mtu.edu

MS330

Scaling Limits of Current Machine Learning Frameworks and How to Extend Them

There are many similarities between machine learning and PDE-constrained optimization problems leading researchers to suggest that such methods can be rewritten using frameworks such as TensorFlow. However, for memory intensive calculations, such as back-propagation in full-waveform inversion, performance and scaling issues become apparent. In the case of adjoint-state methods, computer memory is quickly exhausted, and other methods such as revolve and compression are employed. Additionally, while finite-difference stencil operators may be implemented as a convolutional neural network, such a generic description overlooks optimizations that developed for stencil codes. To alleviate these issues we extend TensorFlow with Devito - a Python-based DSL for implementing high-performance finite-difference PDE solvers. The computational cost is dominated by the numerical solution of PDE's and their adjoints. Therefore, a great deal of effort is invested in aggressively optimizing the performance of these solvers for different computer architectures. By embedding a DSL within Python and making heavy use of SymPy, a symbolic mathematics library, we make it possible to develop finite-difference simulators quickly using a syntax that strongly resembles the mathematics. The Devito compiler reads this code and applies a wide range of analysis to generate highly optimized and parallel code.

Gerard J. Gorman
Imperial College London
g.gorman@imperial.ac.uk

MS330

Machine Learning in Seismic Imaging: from Low-fidelity to High-fidelity

While modern-day wave-based inversion has been responsible for major breakthroughs in the field seismic imaging, the successes are only as good as our ability to simulate waves. Herein, lies a dillema because wave phenomena in the earth are highly complex and often out reach even in cases where we have the physics right but where the numerics are too demanding or parameterizations unfit for inversion. By recognizing intrinsic one-to-one similarities between physics-based time-stepping and (residual) deep convolutional neural networks, we propose a new hybrid network architecture to deal with our inability to capture and/or accurately numerically model the physics. Application of this idea to the problem of numerical dispersion demonstrates a remarkable ability of deep convolutional networks to map low-fidelity dispersed simulations to high-fidelity ones. We feel that this apparently generalizable capability of CNNs opens a new approach to simulate waves.

Felix Herrmann, Ali Siahkoohi, Mathias Louboutin
Georgia Institute of Technology
felix.herrmann@gatech.edu, alisk@gatech.edu, mlouboutin3@gatech.edu

MS330

Imaging for Inverse Scattering Using Learned Preconditioners

The problem of estimating the material properties of an object by measuring the wavefields it scatters is prevalent in a variety of applications including ground penetrating radar, seismic imaging, and microscopy. We consider in this talk the nonlinear wave-equation formulation that accounts for multiple wave scattering, and focus particularly on the reflective tomography regime where inverting the wave-equation presents a computational bottleneck. The complexity arises from the ill-conditioning of the Helmholtz operator in the reflection tomography regime when the values of the wavenumber and the material contrast ratio are very large. In this context, we explore learning methods to construct preconditioners for both the forward and adjoint Helmholtz operators and illustrate the usefulness of these techniques in reconstructing the images of highly scattering objects.

Hassan Mansour
Mitsubishi Electric Research Laboratories
mansour@merl.com

Ajinkya Kadu
Mathematical Institute, Utrecht University, The Netherlands

a.a.kadu@uu.nl

MS330

A Kernel Method for Full Waveform Inversion

Full waveform inversion aims to estimate subsurface medium parameters from seismic measurements. It is typically cast as a PDE-constrained optimization problem which can be solved using the well-known reduced-space method. It has been shown to be advantageous to relax the constraints and solve a joint parameter-state estimation problem. The challenge here is to device an algorithm that avoids having to store and update the full state. In this talk I discuss a formulation of this problem in a Reproducing Kernel Hilbert Space, where the Representer Theorem can be used to yield a finite-dimensional optimization problem whose dimension is dictated by the size of the data, which is typically much smaller than the size of the discretized state.

Tristan van Leeuwen
Mathematical Institute
Utrecht University
t.vanleeuwen@uu.nl

MS331

Combining Preconditioner Updates with Krylov Subspace Recycling for Sequences of Linear Systems

Preconditioners are generally essentially for fast convergence in the iterative solution of linear systems of equations. Computing a good preconditioner can be expensive and when solving a sequence of many systems, we want to avoid the potentially high cost of recomputing a new preconditioner too often. In this case, it can be advantageous to recycle preconditioners (or update and reuse a previous preconditioner) for subsequent matrices in the sequence. We discuss a simple and effective method that we refer to as the Sparse Approximate Map, or SAM update, which is independent of the preconditioner type and quality. Recycling Krylov subspaces from previous systems is a complementary method for reducing computational cost. Rather than discarding the spaces generated for one system, we can select a subspace (or the recycle space) to use for another, closely related system in order to speed up convergence. We examine recycling preconditioners along with recycling subspaces. We provide results for several applications including hydraulic tomography, topology optimization, the nonlinear convection-diffusion equation, and model reduction. Convergence analysis will be explored.

Arielle K. Carr
Virginia Tech
arielle5@vt.edu

MS331

Using Krylov Subspace Information for Updating Preconditioners

For many applications, we need to solve a long sequence or large group of related linear systems. Since a good preconditioner is both essential for fast convergence and expensive to compute, we would like to update preconditioners at low cost. In this presentation, we present methods to update preconditioners, for example, sparse approximate maps or low-rank updates to the identity, using information that becomes available during the Krylov iteration.

Eric De Sturler
Virginia Tech
sturler@vt.edu

MS331

Polynomial Preconditioned GMRES to Reduce Communication in Parallel Computing

Polynomial preconditioning with the GMRES minimum residual polynomial has the potential to greatly reduce orthogonalization costs, making it useful for communication reduction. We implement polynomial preconditioning in the Belos package from Trilinos and show how it can be effective in both serial and parallel implementations. We further show it is a communication-avoiding technique and is a viable option for large-scale parallel computing.

Jennifer A. Loe
Baylor University
jennifer_loe@baylor.edu

Heidi K. Thornquist
Sandia National Laboratories
hkthorn@sandia.gov

Erik G. Boman
Center for Computing Research
Sandia National Labs
egboman@sandia.gov

MS331

Krylov Methods for Rank-one Updates of Linear Equations and Eigenvalues

While matrices of low rank are increasingly useful, little has been done on iterative methods for rank-one perturbations of large matrices. We propose Krylov methods for both linear equations and eigenvalues in this situation. First, once eigenvalues and eigenvectors of an original matrix have been computed with Arnoldi, the iteration can be continued on a rank-one update in order to find eigenpairs of the updated matrix. Two methods are given for linear equations, one continuing a GMRES-DR iteration and then using subspace recycling (deflation of eigenvalues). The second approach uses the Sherman-Morrison formula and needs two phases with recycling. Examples will show these methods can take advantage of the rank-one change and be more efficient than starting from scratch.

Ron Morgan
Baylor University
Ronald_Morgan@baylor.edu

Mark Embree
Department of Mathematics
Virginia Tech

embree@vt.edu

MS332

Intel Vector Math Libraries and Their Features, Usage, and Performance Improvements

Abstract not available.

Ahmet Akkas
Intel
ahmet.akkas@intel.com

Cristina Anderson, Andrey Kolesov, Robin Gilbert
Intel Corporation
., ., .

Marius Cornea
Intel
marius.cornea@intel.com

MS332

Performance Portable SIMD Approach Implementing Block Line Solver For Coupled PDEs

Efficient vectorization is essential to achieve high performance in modern computing architectures. Developers in general rely on a compiler's capability of auto vectorization. However, many complex applications often do not benefit from the auto vectorization technique. In this talk, we show almost perfectly vectorized block line solver implementation for CPU and GPU. The line solver requires to solve many instances of block tridiagonal matrices. Our approach is to use a compact data layout interleaving elements of multiple block matrices. This enables efficient (almost perfect) vectorization for small block matrix computations, which does not work well with conventional compiler auto vectorization techniques. On GPUs, we use the same compact data layout for coalescing memory accesses and each thread in a warp performs the same scalar operations forming an equivalent SIMD operations. Providing a unified SIMD abstraction, we demonstrate performance portable block line solver that can efficiently exploit architecture specific SIMD features.

Kyungjoo Kim
Sandia National Laboratories
kyukim@sandia.gov

MS332

Ensemble Propagation for Efficient Uncertainty Quantification of Mechanical Contact Problems

A new approach called embedded ensemble propagation has been introduced to improve the efficiency of sampling-based uncertainty quantification methods on emerging architectures. This approach consists of simultaneously propagating a subset of samples, called ensemble, through a high-fidelity model instead of propagating one sample at a time. It improves memory access patterns, enables sharing of information from sample to sample, reduces message passing latency, and improves opportunities for vectorization. The embedded ensemble propagation has, however, two challenges. First, the impact of these improvements on the efficiency of the code depends on the ensemble divergence, whereby individual samples within an ensemble must follow different code execution paths. Second, the interleaved memory layout of matrices of ensemble type makes the use of standard BLAS implementation inefficient when applied on ensemble type. In this presentation, we will assess the efficiency of the embedded ensemble propagation applied to the GMRES method to solve unsymmetrical thermomechanical problems. Such an iterative solver introduces ensemble divergence that we explicitly managed with masked assignments and logical reduction functions and requires BLAS functions for ensemble types such as the GEMV. We have studied the GEMV in the context of the ensemble propagation and how numerical parameters of the algorithm are influenced by the ensemble size.

Kim Liegeois, Romain Boman
University of Liege
kim.liegeois@ulg.ac.be, r.boman@ulg.ac.be

Eric Phipps
Sandia National Laboratories
Optimization and Uncertainty Quantification Department
etphipp@sandia.gov

Maarten Arnst
Universite de Liege
maarten.arnst@ulg.ac.be

MS332

A Performance Portable SIMD Scalar Type for Effective Vectorization Across Heterogeneous Architectures

Emerging computer architectures use highly-capable vector units with varying, long vector lengths. Effectively using these vector units is a key to achieving high performance on these architectures. As architectures evolve with different capabilities (e.g. Single Instruction Multiple Threads (SIMT) vs traditional Vector Processing Units (VPUs)), vector code portability becomes increasingly important. A unified, portable, Single Instruction Multiple Data (SIMD) primitive, proposed in this work, allows intrinsics-based vectorization on traditional CPUs and manycore architectures such as Intel Knights Landing, and also facilitates SIMT based execution on GPU architectures. The primitive allows users to develop a portable code irrespective of the vectorization style of the underlying architecture. The new SIMD primitive is evaluated using three different kernels from real-life scientific applications, including batched GEMM, a complex PDE assembly and sparse matrix multiple vector multiplication on different architectures. Experiments show that using portable SIMD primitives not only makes the code portable, but also gives a speed-up up to 7.8x on Intel KNL for kernels, which do not get auto-vectorized, and up to 1.5x on KNL and NVIDIA P100 GPUs for kernels, which are already executing in SIMD / SIMT fashion, for double precision data compared to previous state-of-the-art of the kernels.

Damodar Sahasrabudhe
University of Utah
damodars@sci.utah.edu

Eric Phipps
Sandia National Laboratories
Optimization and Uncertainty Quantification Department
etphipp@sandia.gov

Siva Rajamanickam
Sandia National Laboratories
srajama@sandia.gov

Martin Berzins
Scientific Computing and Imaging Institute
University of Utah
mb@sci.utah.edu

MS333

Boundary Conditions for Tempered Fractional Diffusion

Many problems in hydrology and turbulence modeling involve transient anomalous diffusion, which exhibits super- and/or sub-diffusive scaling for some range of times before transitioning to classical diffusive scaling for larger times. Transient anomalous diffusion can be modeled by tempered space-fractional diffusion equations (TSFDEs). As tempered fractional models are inherently nonlocal, the concept of boundary conditions must be reinterpreted in this context. In this talk, we present explicit and implicit Euler schemes for TSFDEs with reflecting or absorbing boundary conditions, where reflecting boundaries are implemented so that the Euler schemes conserve mass.

Anna Lischke
Brown University
anna_lischke@brown.edu

James F. Kelly
Dept. of Statistics & Probability, Michigan State Universi
kellyja8@stt.msu.edu

Mark Meerschaert
Dept. of Statistics and Probability
Michigan State University, USA
mcubed@stt.msu.edu

MS333

A Splitting Method for the Fast Computation of Nonlocal Diffusion and Peridynamics Models in Heterogeneous Media

Abstract not available.

Xiaochuan Tian
University of Texas at Austin
xtian@math.utexas.edu

MS333

Surface Energy Approach in Modeling of Nano-objects

It is known that surface energy plays an increasingly important role in modeling of objects at nano-scales. In this talk, surface energy model proposed by Steigmann and Ogden will be applied to fracture problems and contact problems in two-dimensional setting. The boundary conditions on the banks of the crack and on the semi-plane connect the stresses and the derivatives of the displacements. The mechanical problems are reduced to systems of singular integro-differential equations which are further reduced to the systems of equations with logarithmic singularities. The existence and uniqueness of the solution for almost all the values of the parameters is proved. Additionally, it is shown that introduction of the surface mechanics into the fracture problems and contact problems leads to the size-dependent equations. A numerical scheme of the solution of the systems of singular integro-differential equations is suggested, and the numerical results are presented for different values of the mechanical and the geometric parameters.

Anna Zemlyanova
Department of Mathematics
Kansas State University
azem@ksu.edu

MS333

Collocation Methods for the Integral Fractional Laplacian with Applications

Abstract not available.

Yanzhi Zhang
Missouri University of Science and Technology
zhangyanz@mst.edu

MS334

Certifying Accuracy and Uniqueness in Ill-conditioned Imaging Problems

Abstract not available.

Stephen Becker, James Folberth
University of Colorado at Boulder
stephen.becker@colorado.edu,
james.folberth@colorado.edu

MS334

Storage-optimal Convex Low Rank Semidefinite Programming

Abstract not available.

Madeleine R. Udell, Lijun Ding
Cornell University
madeleine.udell@gmail.com, xsjdlj@gmail.com

MS334

Manipulating Neural Networks with Inverse Problems

Abstract not available.

Tom Goldstein
University of Maryland
Department of Computer Science

tomg@cs.umd.edu

MS334

Fast Methods for Low-rank Semidefinite Optimization

Abstract not available.

Yifan Sun
University of British Columbia, Vancouver
ysun13@cs.ubc.ca

MS335

Asynchronous Model-based-search for Automatic Performance Tuning.

We will present SuRF an asynchronous model-based autotuning search method that consists of sampling a small number of input parameter configurations and progressively fitting a surrogate model over the input-output space. The asynchronous aspect allows the search to avoid waiting for all the evaluation results before proceeding to the next iteration. As soon as an evaluation is finished, the data is used to retrain the surrogate model, which is then used to bias the search toward the promising configurations. The framework is designed to operate in the master-worker computational paradigm, where one master node fits the surrogate model and generates promising input configurations and worker nodes perform the computationally expensive evaluations and return the outputs to the master node. We will discuss the results from case studies on autotuning SuperLU library and QMCPACK application.

Paul D. Hovland
Argonne National Laboratory
MCS Division
hovland@anl.gov

MS335

MultiVerse-Kriging: A Transfer and Multitask-learning Based Approach to Autotuning

Autotuning, or the problem of finding the optimal parameters of a software that will maximize a certain performance criteria (computation time, memory usage, energy consumption, ...), can be considered as a black-box optimization problem. No analytical formulation of the objective function nor any information on its derivative are available. Given a set of values for the parameters, it is possible to evaluate the objective function only through a costly run of the software. Several algorithms exist in the literature that target this problem, among which are the surrogate-based black-box optimization algorithms. Their main idea is to build a cheap to evaluate model of the objective function and try to optimize this model instead of the original function. One such approach is Kriging. Our work is a generalization of Kriging to the case where not only one optimization problem has to be solved, but where several and similar problems are to be. It lies within the transfer- and multitask-learning frameworks, as we use the knowledge acquired while solving one problem to help solve the other problems simultaneously.

Wissam M. Sid-Lakhdar
ENS Lyon
wissam@lbl.gov

MS335

From Matrix Multiplication to Deep Learning: Are We Tuning the Same Kernel?

Deep neural networks are often implemented on GPUs. This talk will show the commonalities between neural layers in deep neural networks for image recognition. Further similarity will be provided between direct convolution and matrix-multiplication GEMM kernel. In our approach, autotuning can be used for both and we can correlate them along multiple axis with performance being the most prominent one. Analysis of the results will be provided as well.

Yaohung Tsai
University of Tennessee, Knoxville
ytsai2@vols.utk.edu

MS335

On the High Performance Implementation of Quaternionic Matrix Operations

In the context of high performance scientific computing, contemporary implementations of complex and real valued level-3 BLAS routines provide the archetypal example of microarchitecture optimised software in the modern era. However, certain scientific applications such as relativistic electronic structure software can drastically reduce their memory and floating point operation footprint in the exploitation of hypercomplex domains such as the quaternion numbers. In this work is presented HAXX (Hamilton's Quaternion Algebra for CXX): a novel software infrastructure which targets the high performance implementation of quaternionic linear algebra tuned for modern computing architectures. Preliminary numerical experiments are presented to demonstrate the efficacy of quaternionic linear algebra as a replacement for complex linear algebra in certain scientific applications.

David Williams-Young
Lawrence Berkeley National Laboratory
dbwy@lbl.gov

MS336

New Models for Streaming Graphs

Focusing on parallel algorithm design and implementation, we formalize a practical model for graph analysis on streaming data. In this model, a massive graph undergoes changes from an input stream of edge insertions and removals. The model supports concurrent updating of the graph while algorithms execute concurrently on the dynamic data structure. The talk introduces a concept of validity: an algorithm is valid if the output is correct for a graph consisting of the initial graph with some subset of concurrent changes. Practical examples of this model are given for valid implementations of breadth first search,

connected components, PageRank, and triangle counting, all useful graph kernels in real-world applications. This is joint work with E. Jason Riedy and Chunxing Yin.

David A. Bader
Georgia Institute of Technology
bader@cc.gatech.edu

MS336

Diffusive Load Balancing of Hypergraphs for Partitioning Evolving Unstructured Mesh Applications

Maintaining performance in simulations with complex information dependencies running on massively parallel computers require computational work to be evenly distributed throughout the simulation while minimizing communication. For evolving unstructured mesh applications, fast repartitioning techniques are required that can maintain the balance as the workload changes. We present EnGPar, a hypergraph partition improvement tool, and utilize its diffusive load balancing method on unstructured mesh partitions used in computational fluid dynamics applications. Results are provided on partitions up to 512Ki processes on an IBM BlueGene/Q.

Gerrett Diamond
Rensselaer Polytechnic Institute
diamog@rpi.edu

Mark S. Shephard
Scientific Computing Research Center
Rensselaer Polytechnic Institute
shephard@rpi.edu

MS336

Local Algorithms for Hierarchical Dense Subgraph Discovery

Finding the dense regions of a graph and the relations among them is a fundamental problem in network analysis. Core and truss decompositions reveal dense subgraphs with hierarchical relations. The incremental nature of algorithms for computing these decompositions and the need for global information at each step of the algorithm hinders scalable parallelization and approximations since the densest regions are not revealed until the end. In a previous work, Lu et al. proposed to iteratively compute the h-indices of neighbor vertex degrees to obtain the core numbers and prove that the convergence is obtained after a finite number of iterations. This work generalizes the iterative h-index computation for truss decomposition as well as nucleus decomposition which leverages higher-order structures to generalize core and truss decompositions. In addition, we prove convergence bounds on the number of iterations. We present a framework of local algorithms to obtain the core, truss, and nucleus decompositions. Our algorithms are local, parallel, offer high scalability, and enable approximations to explore time and quality trade-offs. Our shared-memory implementation verifies the efficiency, scalability, and effectiveness of our local algorithms on real-world networks.

A. Erdem Sariyuce
Sandia National Laboratories
erdem@buffalo.edu

Seshadhri C
University of California, Santa Cruz
sesh@ucsc.edu

Ali Pinar
Sandia National Labs
apinar@sandia.gov

MS336

Fast Linear Algebra-based Triangle Analytics with Kokkos Kernels

Triangle-based analytics are important in many larger data-analytics based applications, from dense neighborhood discovery to link recommendations. Previously, we developed a highly efficient linear algebra-based algorithm for these triangle-based computations (e.g., triangle counting, k-truss). We implemented this in the Kokkos Kernels library and recently extended this to use Cilk. We present an overview of the algorithm/implementations and present relevant triangle-based results on multi-core architectures, demonstrating that our linear algebra-based approach is competitive with other methods.

Abdurrahman Yasar
Georgia Institute of Technology
School of Computational Science and Engineering
ayasar@gatech.edu

Michael Wolf, Jonathan W. Berry
Sandia National Laboratories
mmwolf@sandia.gov, jberry@sandia.gov

Umit V. Catalyurek
Georgia Institute of Technology
School of Computational Science and Engineering
umit@gatech.edu

MS337

Tensor Computations: Efficiency Or Productivity?

While in mathematics and physics tensors have been first-class citizens for more than a century, from the perspective of high-performance computing they have been mostly neglected until very recently. In scientific applications, the computation of tensor operations often was, and still is, crudely translated into n-fold nested loops in combination with scalar operations. This approach has the advantage of being straightforward, and allows to incorporate simplifications due to the physics of the problem; however, due to the lack of locality, it comes at the cost of severe performance hits. For this reason, we are witnessing a rapid increase in interest for the development of high-performance libraries and tools. In this talk, we present tensor operations that arise in disciplines such as materials science and computational biology, and provide a summary classification based on the different approaches used to achieve high-performance implementations. On opposite ends of the spectrum we place operations that can be efficiently cast (possibly with significant effort) as high-performance matrix operations, and operations that instead require their

own algorithms and libraries.

Paolo Bientinesi
RWTH Aachen University
pauldj@aices.rwth-aachen.de

MS337

On Optimizing Distributed Tucker Decomposition for Sparse Tensors

We study the problem of constructing the Tucker decomposition of sparse tensors on distributed memory systems via the HOOI method. The scheme used for distributing the input tensor among the processors critically influences the HOOI execution time. Prior work has proposed different distribution schemes: an offline scheme based on sophisticated hypergraph partitioning method and simple, lightweight alternatives that can be used real-time. While the hypergraph based scheme typically results in faster HOOI execution time, being complex, the time taken for determining the distribution is an order of magnitude higher than the execution time of a single HOOI iteration. Our main contribution is a lightweight distribution scheme, which achieves the best of both worlds. We show that the scheme is near-optimal on certain fundamental metrics associated with the HOOI procedure and as a result, near-optimal on the computational load (FLOPs). Though the scheme may incur higher communication volume, the computation time is the dominant factor and as the result, the scheme achieves better performance on the overall HOOI execution time. Our experimental evaluation on large real-life tensors (having up to 4 billion elements) shows that the scheme outperforms the prior schemes on the HOOI execution time by a factor of up to 3x. On the other hand, its distribution time is comparable to the prior lightweight schemes and is typically lesser than the execution time of a single HOOI iteration.

Jee Choi
IBM TJ Watson
jwchoi@us.ibm.com

MS337

HPC Formulations of Optimization Algorithms for Tensor Completion

Tensor completion is a powerful tool used to estimate or recover missing values in multi-way data. It has seen great success in domains such as product recommendation and healthcare. Tensor completion is often accomplished via low-rank sparse tensor factorization, a computationally expensive non-convex optimization problem which has only recently been studied in the context of parallel computing. Here, we study three optimization algorithms that have been successfully applied to tensor completion: alternating least squares (ALS), stochastic gradient descent (SGD), and coordinate descent (CCD++). We explore opportunities for parallelism on shared- and distributed-memory systems and address challenges such as memory- and operation-efficiency, load balance, cache locality, and communication. Among our advancements are a communication-efficient CCD++ algorithm, an ALS algorithm rich in level-3 BLAS routines, and an SGD algorithm which combines stratification with asynchronous communication. We also show that introducing randomization during ALS and CCD++ can accelerate convergence. We evaluate our parallel formulations on real datasets on a modern systems and demonstrate speedups through 16384 cores. These improvements reduce time-to-solution from hours to seconds on real-world datasets. We show that after our optimizations, ALS is advantageous on small-to-moderate parallel systems, while both ALS and CCD++ provide the lowest time-to-solution on large-scale distributed systems.

Shaden Smith
University of Minnesota
shaden@cs.umn.edu

MS337

Accelerating Alternating Least Squares for Tensor Decomposition by Pairwise Perturbation

The alternating least squares algorithm for CP and Tucker decomposition is dominated in cost by the tensor contractions necessary to set up the quadratic optimization subproblems. We introduce a novel family of algorithms that uses perturbative corrections to the subproblems rather than recomputing the tensor contractions. This approximation is accurate when the factor matrices are changing little across iterations, which occurs when alternating least squares approaches convergence. We provide a theoretical analysis to bound the approximation error, leveraging a new notion of the tensor condition number. Our numerical experiments demonstrate that the proposed pairwise perturbation algorithms are easy to control and converge to minima that are as good as alternating least squares. The performance of the new algorithms shows improvements of 1.3-2.7X with respect to state of the art alternating least squares approaches for various model tensor problems and real datasets on 1 and 16 Intel Skylake nodes of the Stampede2 supercomputer.

Edgar Solomonik
University of Illinois at Urbana-Champaign
solomon2@illinois.edu

Linjian Ma
UC Berkeley
linjian@berkeley.edu

MS338

CIG Perspectives on 14 Years of Sustaining Software and their Communities in Geodynamics

Established in 2005 with funding from the National Science Foundation, the Computational Infrastructure for Geodynamics (CIG, geodynamics.org) is a partnership between solid-earth geophysics and computational science. CIG advances geophysics and related fields of research by developing and disseminating high quality, open source scientific software. Developing sustainable scientific software is both a technical and social challenge. In response, CIG has grown into a community of practice supporting users and developers in deployment of the best practices most effective for our community and supporting the peer interactions necessary to sustain them. CIGs Minimum, Stan-

dard, and Target Software Best Practices provides guidelines to developers who wish to share their code. However, to succeed and be sustainable, a software development project requires champions willing to invest in building a group of users through workshops, tutorials, and hackathons who can grow into user-developers and eventually, code maintainers. CIG Software Training Best Practices establishes guidelines for both instructors and participants to maximize the benefits from CIGs training and events. Multi-day hackathons, (10 days) in the ASPECT software community have proven effective for growing both the code and the developer base, fostering peer mentoring relationships and new collaborations, and preparing the geophysics scientific workforce to be expert users and to contribute to scientific software.

Lorraine J. Hwang
University of California Davis
ljhwang@ucdavis.edu

Louise Kellogg
University of California, Davis
lhkellogg@ucdavis.edu

MS338

Planning to Make Research Software More Sustainable via a US Research Software Sustainability Institute (URSSI)

Many research advances have been possible thanks to the use of software. This software, also known as "research software", has become essential to progress in science and engineering. The researchers who develop the software are experts in their discipline, but not necessarily experts in software engineering. Consequently, they may lack sufficient understanding of the practices that make software development easier, and make the software more robust, reliable, maintainable, and sustainable. The authors of this talk are the leaders of a conceptualization (planning) project that is working with these researchers and software engineering experts to understand how the research community can best collaborate to design and maintain better software with lower effort, including by bringing more members of underrepresented groups into the developer community. This understanding will allow the developers and others to use the software over long periods of time. We are conducting several workshops and a survey to gather and understand the diverse community's needs, and disseminating this information widely via newsletters and a website, http://urssi.us. The primary deliverable of this project is a design and strategic plan for a US Research Software Sustainability Institute (URSSI) that will serve as a community hub and will provide services to help researchers create improved, more sustainable software. This software in turn will accelerate the progress of research.

Daniel S. Katz
University of Illinois at Urbana-Champaign
d.katz@ieee.org

Jeffrey C. Carver
Department of Computer Science
University of Alabama
carver@cs.ua.edu

Sandra Gesing
University of Notre Dame
sandra.gesing@nd.edu

Karthik Ram
University of California Berkeley
karthik.ram@gmail.com

Nicholas Weber
University of Washington
nmweber@uw.edu

MS338

Enabling a Culture of Developer Productivity and Software Sustainability

The DOE Exascale Computing Project (ECP) has a core mission to move beyond current HPC capabilities toward a capable exascale computing ecosystem that accelerates scientific discovery. The Interoperable Design of Extreme-scale Application Software (IDEAS)-ECP is a multi-laboratory team that addresses the organizational and cultural changes ultimately required to enable software developer productivity and software sustainability—as a key aspect of increasing overall scientific productivity. This presentation describes how CSE software development and sustainability could be further enhanced by adopting an approach for fostering team and developer productivity within domains and across the exascale ecosystem. Through the creation of new techniques, customization of software engineering methodologies, utilization of productivity and sustainability improvement planning (PSIP), and community development, the IDEAS-ECP team seeks to facilitate the development of improvement processes toward a productive and sustainable ecosystem that scales the success intrinsic to a few teams to the ECP community at large.

Elaine M. Raybourn
Sandia National Laboratories
emraybo@sandia.gov

MS338

Is It a Project Or a Business? Perspectives on the Consideration of Sustainability

Grant funded software projects provide their recipients an incomplete view of the ecosystem in which they will need to survive to become sustainable. Specifically, proposal review panels and program officers are not the users of the software, nor can their guidance provide help outside of de-risking the initial software investment. Grant recipients should view themselves as recipients of seed money and behave much like an entrepreneurial organization, where the initial risk analysis is performed by an investor, but it is still entirely up to the founding team to deploy resources such that the project is a success in its intended marketplace. A useful decomposition of such behavior can be found in the business model canvas [Osterwalder, Alexander, and Yves Pigneur. Business model canvas.], which provides a template to 9 key aspects of operating a business. These 9 aspects can be grouped into 3 major categories: i) feasibility to deliver the solution, ii) attractiveness of

the solution to its intended audiences, and iii) economic viability of delivering the solution. Grant recipients tend to be highly focused on the first of these 3 major categories; yet, it has the least to do with the long term sustainability of a project. These concepts will be discussed in general terms, along with specific examples drawn from the speakers experiences in technology commercialization, nanoHUB, HUBzero, and the Science Gateways Community Institutes Incubator program.

Michael Zentner
Purdue University
mzentner@cbios.com

MS339

Numerical Methods for Wave Scattering from Multilayered Media

In this talk, numerical methods for wave scattering from layered media will be discussed. First, for the periodically patterned layered media in 2- and 3-D, instead of using well-known quasi-periodic Green's function, the free-space Green's function with a periodizing algorithm is applied to overcome slow convergence and Wood anomaly issues of the quasi-periodic Green's function. Then, the large linear system from multilayered media is efficiently solved using the Schur complement and the block tri-diagonal matrix solver. Numerical results for both 2- and 3-D will be presented. Secondly, for the planar-layered media, layered media Green's function method is used. The layered media Green's function has some advantage as it includes all the jump interface conditions at the cost of computing Green's function itself. Therefore, the free-space numerical methods work for layered media with a minimal modification. The heterogeneous fast multipole method will be introduced as a fast solver

Min Hyung Cho
Department of Mathematical Sciences
University of Massachusetts Lowell
minhyung_cho@uml.edu

MS339

Integral Equation Methods for Computing Resonances in Subwavelength Metallic Structures

The extraordinary electromagnetic field enhancement in subwavelength metallic structures is partially related to the resonances induced by tiny structures or plasmons. In this talk, I will introduce integral equation approaches for obtaining the resonances in two-dimensional structures. Both perfect electric conductors (PECs) and plasmonic metals will be discussed. For PECs, the integral equations lead to a resonance condition formulated by nonlinear equations, which can be solved very efficiently with an initial guess obtained from the asymptotic analysis. For plasmonic metals, a new integral equation formulation and a strategy for obtaining initial guess will be presented.

Junshan Lin
Auburn University
jzl0097@auburn.edu

MS339

A High-order Perturbation of Surfaces Method for Vector Electromagnetic Scattering by Doubly Layered Periodic Crossed Gratings

Abstract not available.

David P. Nicholls
University of Illinois at Chicago
davidn@uic.edu

MS339

Boundary Integral Equation Methods for Metasurface Scattering

Metasurfaces are planar metamaterial slabs of subwavelength thickness composed by a large number of optical elements that can be engineered to manipulate the direction, amplitude, phase, and polarization of light. Typical inverse metasurface designs are based on fast low-order ray-optics approximations and/or effective generalized sheet transition conditions that capture the microscale optical properties of the metasurface. Accurate full-wave electromagnetic wave simulations of metasurfaces, in turn, pose significant challenges to standard PDE solvers as they often lead to intractably large discrete problems. In this talk, I will discuss two approaches to partially overcome these challenges. The first approach is a boundary integral equation framework, based on the generalized sheet transition conditions, to assess the accuracy and compute corrections to standard ray-optics approximations used in metasurface design. The second approach is an efficient widowed Green function method based on a second-kind single-trace formulation, for the full-wave simulation of all-dielectric metasurfaces.

Carlos Perez-Arancibia
Massachusetts Institute of Technology
cperezar@mit.edu

Raphael Pestourie
Harvard University
pestourie@g.harvard.edu

Steven Johnson
MIT
stevenj@math.mit.edu

MS340

Neutrino Transport with Discontinuous Galerkin Methods

Neutrino transport is an essential ingredient in models of core-collapse supernovae and compact binary mergers, and its inclusion in numerical models poses several challenges. Since neutrino-matter interactions occur under non-equilibrium conditions, a description rooted in kinetic theory is warranted. In addition to the high dimensionality of phase space, relativity, a wide range in spatial and temporal scales, non-linear neutrino-matter coupling, and consistent lepton number and four-momentum exchange are some challenges that must be addressed. In this talk,

we give a brief overview of ingredients that pose some computational and mathematical challenges. We also discuss progress on the use of Discontinuous Galerkin methods to tackle some of the challenges in the context of moment models.

Eirik Endeve
Oak Ridge National Laboratory
endevee@ornl.gov

MS340

Numerical Methods of Fully Special and General Relativistic Boltzmann Neutrino Transport

Neutrinos play an important role in core-collapse supernovae (CCSNe) and binary neutron star mergers. Energy, momentum and lepton transfer by neutrino radiation are crucial for their matter-dynamics and nuclear composition. The consistent treatment of neutrino transport is required to uncover the dynamics of these phenomena, predict their observational signals (neutrinos, gravitational waves and electromagnetic waves) and prove physical state in extremely hot and dense matter. The Boltzmann equation is the fundamental equation for neutrino transport, which dictates the time evolution of the statistical behavior of neutrinos in 6-dimensional phase space. Although general relativity (GR) and special relativity (SR) should be taken into account appropriately, it is quite complex and there are many issues to handle them in numerical simulations. In particular, the consistent treatment of advection and collision term in relativistic Boltzmann equation has a challenge. We overcame the problem by developing a novel numerical method with the Lagrangian-remapping technique. In this conference, I will introduce the essence of this technique, and also present our recent results of multi-dimensional CCSNe simulations with full Boltzmann neutrino transport.

Hiroki Nagakura
California Institute of Technology
hirokin@princeton.edu

MS340

Three-dimensional Monte Carlo Neutrino Transport in Neutron Star Mergers

The interactions between neutrinos and other matter is the dominant mechanism for energy and lepton number transport in core-collapse supernovae and neutron star mergers. The general-relativistic Boltmann equation that governs the transport of neutrinos in these systems is exceedingly expensive to solve numerically, making approximate transport methods the primary tool for many astrophysical simulations. I will talk about a steady-state general relativistic Monte Carlo neutrino transport algorithm for arbitrary multi-dimensional discretized spacetimes. I will present the first multi-dimensional comparison of full Boltzmann neutrino transport methods in the context of core-collapse supernovae, and demonstrate that Sedonu and the discrete ordinates code of Nagakura, Shumiyoshi, and Yamada converge to the same result. I will also demonstrate the difficulties associated with using a more approximate two-moment method in the context of neutron star mergers, and will propose some improvements specific to this unique environment.

Sherwood Richers
North Carolina State University
sricher@ncsu.edu

MS341

Sparse Approximation of Data-driven Polynomial Chaos Expansions and their Applications in UQ

In this talk, we will discuss collocation method via compressive sampling for recovering arbitrary Polynomial Chaos expansions (aPC). Our approach is motivated by the desire to use aPC to quantify uncertainty in models with random parameters. The aPC uses the statistical moments of the input random variables to establish the polynomial chaos expansion and can cope with arbitrary distributions with arbitrary probability measures. To identify the aPC expansion coefficients, we use the idea of christoffel sparse approximation. We present theoretical analysis to motivate the algorithm. Numerical examples are also provided to show the efficiency of our method.

Ling Guo, Yongle Liu
Shanghai Normal University
lguo@shnu.edu.cn, 1000378341@smail.shnu.edu.cn

Tao Zhou
Institute of Computational Mathematics
Chinese Academy of Sciences
tzhou@lsec.cc.ac.cn

MS341

A Multi-fidelityModel Reduction Based on Variable-separation for Uncertainty Quantification

Abstract not available.

Lijian Jiang
Los Alamos National Lab
ljjiang@tongji.edu.cn

MS341

Random Batch Methods (RBM) for Interacting Particle Systems

We develop random batch methods for interacting particle systems with large number of particles. These methods use small but random batches for particle interactions, thus the computational cost is reduced from $O(N^2)$ per time step to $O(N)$, for a system with N particles with binary interactions. On one hand, these methods are efficient Asymptotic-Preserving schemes for the underlying particle systems, allowing N-independent time steps and also capture, in the $N \to \infty$ limit, the solution of the mean field limit which are nonlinear Fokker-Planck equations; on the other hand, the stochastic processes generated by the algorithms can also be regarded as new models for the underlying problems. For one of the methods, we give a particle number independent error estimate under some special interactions. Then, we apply these methods to some representative problems in mathematics, physics,

social and data sciences, including the Dyson Brownian motion from random matrix theory, Thomsons problem, distribution of wealth, opinion dynamics and clustering. Numerical results show that the methods can capture both the transient solutions and the global equilibrium in these problems. This is a joint work with Lei Li and Shanghai Jiao Tong University and Jian-Guo Liu at Duke University.

Shi Jin
Shanghai Jiao Tong University, China and the
University of Wisconsin-Madison
sjin@wisc.edu

MS341

Efficient Minimum Action Methods for Transition Path Computing in Noise-driven Non-gradient Systems

In this talk, we will present several efficient minimum action methods(MAM) for transition path computing, including tMAM (a MAM using optimal linear time scaling), a MAM based on Hamilton-Jacobi equation for system with degenerate driven noise, a Laguerre spectral MAM for transition paths escaping from a asymptotically stable fixed point, and a MAM for calculating transition paths escaping from a given stable limit cycle. The mathematical formulations,numerical discretization details as well as numerical results for several typical applications will be presented. References: 1. X. Wan and H. Yu, A dynamic-solverconsistent minimum action method: With an application to 2D NavierStokes equations, J. Comput. Phys. 331(2017) 209-226 2. X. Wan, H. Yu and J. Zhai, Convergence Analysis of a finite element approximation of minimum action methods, SIAM J. Numer. Anal. 56(3):1597-1620 3. L. Lin, H. Yu and X. Zhou, Quasi-potential Calculation and Minimum Action Method for Limit Cycle, arXiv:1806.04031

Haijun Yu
LSEC, Institute of Computational Mathematics
Academy of Mathematics and Systems Science, China
hyu@lsec.cc.ac.cn

MS342

Turbulence-chemistry Interactions for Lean Premixed Flames

Turbulence has well-known impacts on the large scale topology and wrinkling of a premixed flame, as well as on the internal structure of the flame through local, time-dependent strain and curvature. Less well-understood is the influence of turbulence on the chemical reaction pathways. Understanding this interaction is imperative to the development of improved models for turbulence-chemistry interactions used for simulation of realistic engine configurations. This work intends to answer these open questions in the modeling community. The data sets of lean premixed hydrogen, methane and n-dodecane generated using PeleLM, detailed in Aspden et al. (*Proc. Combust. Inst.*, 2015, *Combust. Flame*, 2016 and *Proc. Combust. Inst.*, 2017), are analyzed for this work. Laminar flame calculations for unstretched flames, stretched flames and perfectly stirred reactors are also performed to compare and contrast the results obtained from turbulent flames. In addition, laminar flame calculations for unstretched and stretched flames are performed using different transport models such as unity Lewis number, constant Lewis number and constant diffusivity with mixture averaged transport. These results are compared with DNS to understand the shift in the reaction rate profiles due to altered mechanism of species transport due to molecular and turbulent diffusion. Together, all this analysis provides a basis for the development of improved turbulence-chemistry models.

Debolina Dasgupta
Georgia Tech
ddasgupta3@gatech.edu

Marcus Day
Lawrence Berkeley National Laboratory
msday@lbl.gov

Andrew Aspden
Newcastle University
andrew.aspden@newcastle.ac.uk

Timothy Lieuwen
Georgia Institute of Technology
tim.lieuwen@aerospace.gatech.edu

MS342

Pelelm Overview: Existing and in-Progress Features and Plans for AMR Low Mach Number Combustion Simulation on the Desktop and at the Exascale

PeleLM is a code for evolving chemically reacting low Mach number flows with block-structured adaptive mesh refinement (AMR). The code features an iterative time step that efficiently couples together advection, diffusion and chemical reactions, each potentially evolving with different time scales, all while evolving the numerically conservative discretization across multiple levels of mesh refinement on a manifold where the equation of state is always satified. PeleLM is being used to study the details of turbulence-chemistry interaction in both premixed and diffusion flame configurations, from smal-scale laboratory experiments to large-scale pool fires and wildfires, while it is being simultaneously improved and extended to incorporate new physics, and to run efficiently on emerging high-performance computing hardware platforms. In this talk, I will outline the current capabilities of PeleLM and discuss improvements to the code that are under development and supported by DOE's Exascale Computing Project. These improvements include domains of arbitrary 3D geometries, electric field effects on charged particles generated during the combustion process, and incorporation of real gas and closed-chamber pressurization effects. I also discuss further extensions to the PeleLM time stepping scheme to support high-order methods, and implementation strategies used in PeleLM for many-core and GPU-based large-scale computing platforms.

Marcus Day
Lawrence Berkeley National Laboratory

msday@lbl.gov

MS342

A Numerical Strategy for Unsteady Simulation of Low Mach Number Reacting Flows Subject to Electric Fields

Experiments have demonstrated that electric fields can be employed in combustion applications to control the flame stabilization and reduce pollutant emissions. However, the multi-scale nature of the physical processes has so far prevented the development of efficient numerical methods, able to further analyze the interactions between turbulent flows, combustion kinetics and electric fields. We propose an algorithm for low Mach number combustion that incorporates the chemical production of charged species and their transport induced by electric field. The strategy relies on the multi-implicit spectral deferred correction (MISDC) approach that allows for a strong and efficient coupling of the stiff drift/diffusion and reaction terms with the slower advection terms. Along with the momentum, energy and species transport equations, a Poisson equation is solved to provide a consistent electric field. The scale separation between the light electrons and the heavier charged ions is exploited to overcome the stringent time-stepping constraint due to the electron dielectric relaxation. We will first present the details of the method and demonstrate its capabilities on several of one-dimensional cases. We will present results of the algorithm in the context of the low Mach number solver, PeleLM, applied to an experimental burner configuration.

Lucas Esclapez
Lawrence Berkeley National Laboratory
lesclapez@lbl.gov

Valentina Ricchiuti
University of California Irvine, MAE Department
Lawrence Berkeley National Lab
vricchiuti@lbl.gov

Marcus Day
Lawrence Berkeley National Laboratory
msday@lbl.gov

MS342

Direct Numerical Simulation of Mixture Formation and Ignition under Multi-injection Diesel Relevant Conditions

Heavy-duty transportation and unmanned aerial vehicle (UAV) sectors rely on diesel engines that operate under intermittent low-temperature combustion conditions. To meet pollutant emission standards and reduce noise and excessive peak pressures, multi-injection strategies are implemented where, for example, a small amount of fuel ('pilot') is injected ahead of the main injection. While there have been experimental investigations, a detailed and fundamental understanding of the mixing and ignition processes is still missing. We will present Direct Numerical Simulations (DNS) of multi-injection ignition and combustion under UAV conditions. In particular, we will elucidate the effect of the pilot on the ignition of the main injection. The simulations are performed with the block-structured low-Mach adaptive mesh refinement (AMR) code PeleLM. PeleLM relies on the spectral deferred correction method to couple chemistry and transport and a density-weighted approximate projection method to satisfy the velocity constraint imposed by continuity. With the AMR framework, a parameter space is explored beyond what is currently achievable in terms of pressure, Reynolds number and fuel complexity. Numerical and algorithmic challenges will be presented in pushing DNS towards engine-relevant conditions that are characterized by a large separation of scales to resolve complex fuel chemistry-turbulence interactions in the context of large-scale AMR simulations.

Martin Rieth
Sandia National Laboratories
mrieth@sandia.gov

Marcus Day
Lawrence Berkeley National Laboratory
msday@lbl.gov

Jacqueline Chen
Combustion Research Facility
Sandia National Laboratories
jhchen@sandia.gov

MS343

Data-driven Linear Parameter Varying Control of Wind Farms

Abstract not available.

Jennifer Annoni
NREL
jennifer.annoni@nrel.gov

MS343

Nonintrusive Nonlinear Model Reduction via Machine-learning Approximations to Low-dimensional Operators

Although projection-based reduced-order models (ROMs) for parameterized nonlinear dynamical systems have demonstrated exciting results across a range of applications, their broad adoption has been limited by their intrusivity: implementing such a reduced-order model typically requires significant modifications to the underlying simulation code. To address this, we propose a method that enables traditionally intrusive reduced-order models to be accurately approximated in a nonintrusive manner. Specifically, the approach approximates the low-dimensional operators associated with projection-based reduced-order models (ROMs) using modern machine-learning regression techniques. The only requirement of the simulation code is the ability to export the velocity given the state and parameters; this functionality is used to train the approximated low-dimensional operators. In addition to enabling nonintrusivity, we demonstrate that the approach also leads to very low computational complexity, and can realize simulation times comparable to that of typical hyper-reduction methods (e.g., DEIM, GNAT). We demonstrate the effectiveness of the proposed technique on several problems in

fluid mechanics.

Zhe Bai
University of Washington
zhbai@uw.edu

Kevin T. Carlberg, Liqian Peng
Sandia National Laboratories
ktcarlb@sandia.gov, pengliqian@gmail.com

Steven Brunton
University of Washington
sbrunton@uw.edu

MS343

Data-driven Dynamical Systems Modeling

Projection-based model reduction has been successfully employed in prominent applications. The projection-based methods require access to internal dynamics, which is not always accessible. Rather, one might have access to abundant input/output measurements in the form of transfer function evaluations or time-domain simulations. We will discuss how one can employ this available data to construct high-fidelity data-driven dynamical reduced models. We will discuss both the interpolation (Loewner) and least-squares (Vector Fitting) frameworks based on transfer function samples and will present applications ranging from circuit simulations to estimating dispersion relations in structural components.

Serkan Gugercin
Virginia Tech
Department of Mathematics
gugercin@vt.edu

MS343

Operator Inference on Manifolds for Learning Physically Consistent Models from Data

Our task is learning reduced models when only snapshot data of the large-scale full systems are available, and the system operators are unknown (nonintrusive model reduction). Our approach is based on operator inference that learns reduced-model operators by inferring them from snapshot data via a least-squares approach. However, inferring the reduced operators from data alone typically leads to reduced models without guarantees on their structure and properties (physics). We propose to impose structure on the reduced models by restricting the operator inference to matrix manifolds (data + physics). For example, inferring the operators on the matrix manifold of symmetric negative definite matrices guarantees an asymptotically stable reduced model under certain assumptions on the full system. We discuss efficient techniques for inference on matrix manifolds and demonstrate our manifold-based operator inference on numerical results that show that structure-preserving data-fit models can outperform classical data-fit and machine learning techniques without structure preservation.

Nihar Sawant
New York University
ns4361@nyu.edu

Benjamin Peherstorfer
Courant Institute of Mathematical Sciences
New York University
pehersto@cims.nyu.edu

MS344

Algorithmic Regularization for Machine Learning

Kernel methods provide a principled way to perform non linear learning. They enjoy optimal statistical properties and in a least squares setting they reduce to solving a linear system. However, in large scale problems computational requirements make them impractical. In this talk we will see how different algorithmic choices allow to design solvers which provide robust as well as resource efficient solutions. To minimize computational footprint, ideas from statistics, randomized linear algebra and convex optimization will be used, and we will see how these algorithmic choices control the generalization properties of the solutions.

Luigi Carratino
Universita' di Genova
luigi.carratino@dibris.unige.it

Alessandro Rudi
INRIA
alessandro.rudi@inria.fr

Lorenzo Rosasco
Instituto Italiano di Technologia and
Massachusetts Institute of Technology
lrosasco@mit.edu

MS344

On the Information Capacity of Neural Networks

Thomas Cover (1965) and David MacKay (2003) showed that the maximum information capacity of a perceptron is 2 bits per parameter. I extend their results to arbitrary networks of threshold neurons, allowing the analytical determination of the capacity of contemporary (deep) neural networks. Applications of this result include the capability comparison of different architectures (e.g. deep vs wide) and the possibility of sizing networks for improved experimental design. Demo at: http://tfmeter.icsi.berkeley.edu

Gerald Friedland
UC Berkeley
Lawrance Livermore National Lab
fractor@eecs.berkeley.edu

MS344

Learning-based Predictive Models: A New Approach to Integrating Large-scale Simulations and Experiments

We will describe a large research effort at Lawrence Livermore National Laboratory (LLNL) aimed at using recent advances in deep learning, computational workflows, and computer architectures to develop an improved predictive model the learned predictive model. Our goal is to first train deep neural network models on simulation data to

capture the theory implemented in advanced simulation codes. Later, we improve, or elevate, the trained models by incorporating experimental data. The training and elevation process both improves our predictive accuracy and provides a quantitative measure of uncertainty in such predictions. We will present work using inertial confinement fusion research and experiments at the National Ignition Facility (NIF) as a testbed for development. We will describe advances in machine learning architectures and methods necessary to handle the challenges of ICF science, including rich, multimodal data (images, scalars, time series) and strong nonlinearities. We will also cover state-of-the-art tools that we developed to steer both physics simulation and model training. We have used these tools to produce 100 million ICF simulations, including 4.8 billion multi-view and multi-channel images using the Sierra supercomputer. We will describe our ongoing efforts to exploit the GPU-rich architecture of Sierra to develop a learned predictive models based on this massive data set.

Brian Spears
Lawrence Livermore National Laboratory
spears9@llnl.gov

MS344

Bayesian Inference Using Adaptive Gaussian Process Models with Application to Cosmology

The field of cosmology in the recent years has made considerable use of the statistical predictive models to emulate the results of the sophisticated nonlinear codes that simulate the physics of the intergalactic medium. Predictive surrogate models, such as Gaussian process models, are used to interpolate the results of a relatively small number of accurate simulations for the purposes of uncertainty analysis and parameter inference. The high computational cost of the cosmological simulations necessitates a thoughtful approach to the selection of the training input points for the construction of the Gaussian process surrogates. In this talk we present a novel adaptive approach to the construction of the Gaussian process models for the purposes of solving inverse problems in cosmology in a Bayesian framework. Our approach relies on the Bayesian optimization methods to select training input points that provide maximum improvement in the fit of the surrogate model to the observational data. We illustrate our approach on an example inverse problem for the astrophysical parameters using the state-of-the-art simulations of the Universe and the so-called Lyman-alpha flux power spectrum as the summary statistics.

Timur Takhtaganov, Zarija Lukic
Lawrence Berkeley National Laboratory
tat@lbl.gov, zarija@lbl.gov

Juliane Mueller
Lawrence Berkeley National Lab
julianemueller@lbl.gov

Dmitriy Morozov
Lawrence Berkeley National Laboratory
dmorozov@lbl.gov

MS345

Optimal Experimental Design for Constraint Inverse Problems

We address the challenging problem of optimal experimental design (OED) of constrained inverse problems. We consider two OED formulations that allow to reduce the experimental costs by minimizing the number of measurements. The first formulation assumes a fine discretization of the design parameter space and uses sparsity promoting regularization to obtain an efficient design. The second formulation parameterizes the design and seeks optimal placement for these measurements by solving a small-dimensional optimization problem. We consider both problems in a Bayes risk as well as an empirical Bayes risk minimization framework. For the unconstrained inverse state problem, we exploit the closed form solution for the inner problem to efficiently compute derivatives for the outer OED problem. The empirical formulation does not require an explicit solution of the inverse problem and therefore allows to integrate constraints efficiently. A key contribution is an efficient optimization method for solving the resulting, typically high-dimensional, bilevel optimization problem using derivative-based methods. To address the growing computational complexity of empirical Bayes OED, we parallelize the computation over the training models. Numerical examples and illustrations from tomographic reconstruction, for various data sets and under different constraints, demonstrate the impact of constraints on the optimal design and highlight the importance of OED for constrained problems.

Matthias Chung
Department of Mathematics
Virginia Tech
mcchung@vt.edu

Lars Ruthotto
Department of Mathematics and Computer Science
Emory University
lruthotto@emory.edu

Julianne Chung
Department of Mathematics
Virginia Tech
jmchung@vt.edu

MS345

Large-scale Optimal Experimental Design for Bayesian Nonlinear Inverse Problems

Abstract not available.

Peng Chen
UT Austin
peng@ices.utexas.edu

Omar Ghattas
The University of Texas at Austin
omar@ices.utexas.edu

Umberto Villa
University of Texas at Austin
uvilla@wustl.edu

Keyi Wu
University of Texas Austin
keyiwu@math.utexas.edu

MS345

Efficient A-Optimal Experimental Design for Bayesian Linear Inverse Problems Using Randomized Estimators and Reweighted ℓ_1 Minimization

We tackle the problem of finding optimal sensor placements for large-scale Bayesian linear inverse problems. This can be formulated as an Optimal Experimental Design (OED) problem, which involves optimizing certain design criteria over the design variables. Specifically, we consider the A-optimal criterion, which involves the computation of the trace of the posterior covariance matrix. We also propose a modified A-optimal criterion that is easier to evaluate. We use structure exploiting randomized matrix methods to develop accurate and efficient methods for computing these design criteria and their gradients. The optimization problem for finding optimal designs strikes a balance between minimizing the design criterion and the number of sensors deployed. To control the sparsity of the sensor placements, we use an approach based on reweighted ℓ_1 minimization. We present illustrative numerical results for a model contaminant source identification problem, where the inverse problem seeks to recover the initial condition of a contaminant plume using discrete measurements of the contaminant in space and time.

Elizabeth Herman
North Carolina State University
eaherman@ncsu.edu

Alen Alexanderian
NC State University
alexanderian@ncsu.edu

Arvind Saibaba
North Carolina State University
asaibab@ncsu.edu

MS345

Bayesian Optimal Experimental Design in Non-Submodular Settings: Batch Algorithms and Guarantees

Objectives for Bayesian optimal experimental design can be viewed as set functions that are monotone but typically non-submodular. We propose and analyze batch greedy heuristics for these optimization problems. Reinterpreting the classical greedy algorithm using the MM principle, we propose a new class of methods that exploit modular bounds. For linear inverse problems, these bounds can be computed efficiently by exploiting low rank structure. We show how these new ideas relate to leverage scores and volume sampling.

Jayanth Jagalur Mohan, Youssef M. Marzouk
Massachusetts Institute of Technology
jagalur@mit.edu, ymarz@mit.edu

MS346

Manifold Sampling for Optimization of Composite Nonsmooth Functions

We develop a manifold sampling algorithm for the minimization of a nonsmooth composite function $f = \psi + h \circ F$ when ψ is smooth with known derivatives, h is a known, nonsmooth, function, and F is smooth but expensive to evaluate. The trust-region algorithm classifies points in the domain of h as belonging to different manifolds and uses this knowledge when computing search directions. Since h is known, classifying objective manifolds using only the values of F is simple. All cluster points of the sequence of iterates from manifold sampling are Clarke stationary; this holds although points evaluated by the algorithm are not assumed to be differentiable and when only approximate derivatives of F are available. Numerical results show that manifold sampling using zeroth-order information about F is competitive with algorithms that employ exact subgradient values from ∂f.

Jeffrey Larson
Argonne National Laboratory
jmlarson@anl.gov

Matt Menickelly
Lehigh University
mmenickelly@anl.gov

Stefan Wild
Argonne National Laboratory
wild@mcs.anl.gov

MS346

Derivative-free Optimization of Computationally Expensive Functions with Noisy Data

Computational simulation models are used in many branches of science to study physical phenomena. The accuracy of these simulations usually depends on the choice of the simulation model parameters and it is assessed by comparing the simulation output to observation data. These observation data are often noisy, and thus the question becomes how to assess the accuracy of the simulation, and thus how to define the objective function when optimizing the parameters in the simulation. Moreover, the simulations we consider are computationally expensive to run, and thus, in an optimization setting, the goal is to find the optimal parameters within as few expensive evaluations as possible. In this presentation, we explore how ensembles of surrogate models can be used to approximate the errors between the simulation and realizations of the observations. We use a separate surrogate model for each realization of the observations and we combine this information to guide our iterative sampling procedure. We explore different objective function formulations and sampling criteria, and we compare them in a numerical study.

Juliane Mueller
Lawrence Berkeley National Lab

julianemueller@lbl.gov

MS346

A Derivative-free Approach for Complex Problems

In this talk, we propose a derivative free algorithm for a class of difficult black-box optimization problems in optimal design. Simulation codes are usually considered when describing both objective functions and constraints in those problems. Furthermore, variables often need to satisfy some additional known relations. We describe the algorithm and its theoretical properties. Furthermore, we report numerical results obtained on a real-world problem in ship design optimization.

Danny D'Agostino
University of Rome
danny.dagostino@uniroma1.it

Matteo Diez
CNR-INSEAN, The Italian Ship Model Basin
matteo.diez@cnr.it

Giampaolo Liuzzi
CNR - IASI
liuzzi@iasi.cnr.it

Stefano Lucidi
University of Rome
lucidi@dis.uniroma1.it

Riccardo Pellegrini
CNR-INSEAN
National Research Council-Marine Technology Research Inst.
riccardo.pellegrini@insean.cnr.it

Francesco Rinaldi
University of Padova
rinaldi@math.unipd.it

Andrea Serani
CNR-INSEAN
National Research Council-Marine Technology Research Inst.
andrea.serani@insean.cnr.it

MS346

Quantile Adaptive Search and Partition-based Algorithms for Global Optimization

Complex systems are often modeled with computer simulations, both deterministic and stochastic, and a decision-maker may desire a set of near-optimal solutions using relatively few function evaluations. We present a partition-based random search optimization algorithm, called Probabilistic Branch and Bound (PBnB), that strives to identify solutions within a target quantile, such as solutions with performance metrics in the best 10% overall. Our approach is valid for black-box, ill-structured, noisy function evaluations, involving both integer and real-valued decision variables. Finite-time probability bounds are derived on the maximum volume of incorrectly maintained or incorrectly pruned regions. Thus, the user has a statistical quantification of the output. For example, with probability greater than 0.9, the final maintained sub-region is inside the target level set with the volume of incorrectly maintained points less than 2% of the volume of the initial set. We also present a complexity analysis of an abstracted method, called Quantile Adaptive Search (QAS), that samples from a series of nested level sets associated with decreasing quantile levels. The analysis of QAS provides a bound on the expected number of iterations that are needed to reach a specified value above the global minimum, which, under certain conditions, increases linearly in dimension. This analysis is used to motivate variations of PBnB that attempt to realize the ideal performance of QAS.

Zelda B. Zabinsky
University of Washington
Industrial Engineering
zelda@u.washington.edu

David Linz
University of Washington
ddlinz@uw.edu

MS347

Derivative Based Global Sensitivity Analysis and Dimension Reduction for Models with Function-valued Outputs

We address global sensitivity analysis for models with high-dimensional inputs and functional outputs. We propose derivative based global sensitivity measures (DGSMs) for such models, and derive a link between these functional DGSMs and generalized Sobol' indices for functional outputs. Low-rank approximations of functional QoIs, through their Karhunen–Loeve (KL) representation, are used for output dimension reduction. Moreover, these low-rank KL representations facilitate efficient computation of informative upper bounds on functional Sobol' indices, which can be used to detect unimportant inputs. Adjoint-based gradient computation is used to ensure computational cost of computing the functional DGSMs does not scale with the input parameter dimension. We illustrate our approach with a nonlinear ODE model for cholera epidemics and for problems of porous medium flow through heterogeneous media.

Helen Cleaves
North Carolina State University
hlcleave@ncsu.edu

Alen Alexanderian
NC State University
alexanderian@ncsu.edu

Hayley Guy, Ralph Smith
North Carolina State University
hguy@ncsu.edu, rsmith@ncsu.edu

Meilin Yu
University of Maryland, Baltimore County

mlyu@umbc.edu

MS347

Global Sensitivity Analysis for Multilevel Scramjet Computations

The development of scramjet engines is an important research area for advancing hypersonic and orbital flights. Attaining optimal engine designs is challenging, requiring accurate and computationally affordable flow simulations as well as uncertainty quantification (UQ). In particular, exploring the stochastic space induced by uncertain model parameters, which can easily reach tens of dimensions or more, while demanding high-fidelity and resolved large eddy simulations (LES) of the supersonic reactive turbulent flow, quickly become prohibitive. In this work, we investigate global sensitivity analysis (GSA) of such a multiscale physical system for the purpose of identifying and fixing unimportant uncertain inputs in order to reduce the stochastic dimension. We do this across multiple levels of grid resolution and under sparse datasets of different sizes, and explore the tradeoff and agreement among them.

Xun Huan, Cosmin Safta, Zachary Vane, Guilhem Lacaze
Sandia National Laboratories
xhuan@umich.edu, csafta@sandia.gov, zvane@sandia.gov, gnlacaz@sandia.gov

Joseph C. Oefelein
Sandia National Laboratories
Combustion Research Facility
oefelei@sandia.gov

Habib N. Najm
Sandia National Laboratories
Livermore, CA, USA
hnnajm@sandia.gov

MS347

Probability-space Modeling for Global Sensitivity Analysis

We present a probability-space model for global sensitivity analysis, to compute various Sobol' indices when only input-output data are available. A surrogate model is built to approximate the joint probability density function of the inputs and outputs of the system. The learned probability model is then used to compute various Sobol' indices. Three probability models are investigated for this purpose: Gaussian copula, Gaussian mixture, and a new Gaussian mixture copula. This approach can perform global sensitivity analysis with both independent and dependent variables.

Sankaran Mahadevan
Vanderbilt University
sankaran.mahadevan@vanderbilt.edu

Zhen Hu
University of Michigan-Dearborn
zhennhu@umich.edu

MS347

Exploiting Low-rank Structure for Sensitivity Analysis in Earth System Models

Uncertainty Quantification studies using state-of-the-art earth system models are challenged by the large number of parameters in these models and the high computational cost required for each model evaluation. Previous sensitivity analysis studies of the land component of the Energy Exascale Earth System Model (E3SM) highlighted nonlinear input-output dependencies, which limited the accuracy of surrogate models constructed for subsequent uncertainty quantification and propagation. Here, we present a low-rank functional tensor train decomposition approach for constructing sparse surrogate models to simulate the input-output maps in the E3SM land model and uncover interactions between model components and the parameters that control them. We explore a set of functional representations over the combined stochastic space including both the model parameters and physical space by considering models spanning several adjacent land cells. We compare the efficiency of the approach with low-rank methods spanning the stochastic space only and with sparse regression approximations in a Bayesian setting.

Cosmin Safta, Khachik Sargsyan
Sandia National Laboratories
csafta@sandia.gov, ksargsy@sandia.gov

Daniel Ricciuto
Oak Ridge National Laboratory
ricciutodm@ornl.gov

John D. Jakeman
Sandia National Labs
jdjakem@sandia.gov

Alex Gorodetsky
University of Michigan
goroda@umich.edu

MS348

Applications of Quantum Computing in Computational Chemistry

Quantum chemistry has often been referred to as the killer app for quantum computers. Since quantum chemistry utilizes the principles of quantum mechanics to simulate the molecular properties, a quantum computer is often considered to have an innate advantage in performing these simulations. Current methods used to perform electronic structure calculations on quantum hardware are rigorous and sound, but from an application standpoint, still are lacking, not only in accuracy, but also in quantum resource use. Phase estimation and variational quantum eigensolvers are the two main methods for determining the expectation value or eigenvaluves of the electronic structures, but not much attention has been given to the Ansatz formulation, which are based on methods from quantum computing and not quantum chemistry. New theories such as the qubit mean-field and qubit coupled cluster represent a theoreti-

cal milestone understanding these simulations. Interesting phenomena arises when applying these theories that help demonstrate the theoretical links between quantum circuit operations and quantum chemistry theory. For example, removing entanglement gates results in the Hartree-Fock equivalent with the exception of spin and electron number conservation. We intend to describe the methodologies and consequences in simulating molecules on quantum hardware such as Rigetti Computings 19Q-Acorn, and D-waves 2000Q quantum annealing. .

Scott N. Genin
OTI Lumionics Inc.
scott.genin@otilumionics.com

Ilya Ryabinkin, Artur Izmaylov
University of Toronto
i.ryabinkin@utoronto.ca, artur.izmaylov@utoronto.ca

MS348

Material Simulation on D-Wave

Simulating electronic properties of molecules is an important task for different areas of both research and industry. Due to the exponential scaling of the problem space, solving such problems exactly becomes rapidly unfeasible on classical computers. As electronic structure problems are quantum mechanical problems, quantum devices are naturally predestinated for such calculations. In this talk, I will present electronic structure calculation results done on current quantum computers. In particular, I will show how to utilize a quantum computer to find the energy surfaces of small molecules. This will be illustrated by examples of molecular hydrogen and lithium hydride. Moreover, I will discuss the feasibility of doing such calculations for larger molecules.

Michael Streif, Florian Neukart
Volkswagen
michael.streif@volkswagen.de, florian.neukart@vw.com

MS348

Optimization of Industrial Problems

Many industrial problems depend on excellent optimization. We used the optimization capabilities of the D-Wave 2000Q system for safe and optimal control of Automated Guided Vehicles (AGVs) in industrial plants, minimal congestion for tsunami evacuation, and clustering of data

Masayuki Ohzeki
Tohoku University
mohzeki@tohoku.ac.jp

MS348

Solving Industrial Network-optimization Problems on a Quantum Computer

Many problems in the allocation of resources in telecommunication networks are computationally hard. In this talk we present Ising Hamiltonian formulations for several important operational problems, such as bandwidth allocation and the half-duplex mesh network problem, and demonstrate how they can be solved using a D-Wave 2000Q quantum annealing processor. We find that near optimums are reached quickly using a small number of anneal cycles. We also consider the implications of future quantum annealing devices for network optimisation.

Steve P. Reinhardt
D-Wave Systems, Inc.
sreinhardt@dwavesys.com

Sheir Yarkoni
D-Wave Systems Inc.
syarkoni@dwavesys.com

MS350

The Schwarz Alternating Method for Concurrent Multiscale in Finite Deformation Solid Mechanics

The Schwarz alternating method is a well-established technique for the solution of elliptic PDEs by means of domain decomposition. On the other hand, multiscale methods are becoming prevalent in the analysis of solid mechanics problems in which it is necessary to bridge time or length scales to account for fast or small processes. Multiscale methods are of most relevance in the simulation of damage, failure or crack initiation where the information needs to flow to and from the different scales to capture these phenomena reliably. Here we advance the Schwarz alternating method as a means to couple different scales concurently. Specifically, the Schwarz method avoids the use of Lagrange multipliers or gradients that afflict other coupling methods. We discuss our formulation and its implementation in the open-source Albany finite element platform developed at Sandia National Laboratories and demonstrate the use of the method as an effective approach for concurrent coupling in finite deformation solid mechanics in quasi-static and dynamic simulations.

Alejandro Mota, Irina K. Tezaur, Coleman Alleman
Sandia National Laboratories
amota@sandia.gov, ikalash@sandia.gov,
callema@sandia.gov

Gregory Phlipot
California Institute of Technology
gphlipot@caltech.edu

MS350

Reducing the Resonance Error in Numerical Homogenization Through Parabolic Micro Problems

Over the last two decades there has been a growing interest in the development of numerical homogenization methods for multiscale problems. They rely on the solution of elliptic micro problems over representative volume elements (RVEs) of size $\delta = \mathcal{O}(\varepsilon)$, where ε is the characteristic length of the fine scale heterogeneities. A common issue is the presence of a resonance error, which is proportional to the ratio ε/δ. Reducing it through the increment of the RVE size is computationally expensive. Therefore, finding coupling strategies between the macro- and micro-scale problems capable of improving such convergence rate is a major issue. Many techniques addressing this problem have

been proposed, but several issues remain open. Inspired by a parabolic approach proposed in the context of stochastic network homogenization [J.C. Mourrat, Efficient methods for the estimation of homogenized coefficients, *Found. Comp. Math.*], we propose a parabolic micro problem as a tool for homogenizing oscillating coefficients [A. Abdulle, D. Arjmand, E. Paganoni, Parabolic micro problems for reducing the resonance error in numerical homogenization, preprint]. This model does not rely on any periodicity assumption, and can hence be used also in the quasi-periodic and stochastic settings. A faster convergence rate of the resonance error compared to the rate obtained for standard elliptic micro models can be shown.

Edoardo Paganoni, Assyr Abdulle, Doghonay Arjmand
EPFL
edoardo.paganoni@epfl.ch, assyr.abdulle@epfl.ch, doghonay.arjmand@epfl.ch

MS350

Domain Decomposition Method for High Dimensional Stochastic Systems

With the growth of computational resources available, demand on the accuracy and speed of simulation is steeply increasing. Accordingly, accounting for the variabilities in the simulations plays a pivotal role in taking informed engineering decisions. In this regard Monte Carlo simulation (MCS) and spectral stochastic finite element (SSFEM) are two computationally efficient methods. However both, that is MCS and SSFEM in the existing way are not capable of handling truly large problems which are of practical interest.In this talk we will discuss a computationally efficient method for uncertainty quantification (UQ) of problems with large stochastic dimensionality. First we will discuss a few mathematical results on the domain shape independence of the Karhunen Loeve (KL) expansion. This will be followed by derivation of a set of mathematical bounds — dependent on the domain size — on the eigenvalues in the KL expansion. Based on these mathematical results a FETI based hybrid stochastic finite element algorithm for UQ of problems with large stochastic dimensionality will we discussed. Further, to quantitatively ascertain the parallel performance of the proposed algorithm, we perform scalability studies of the algorithm in a distributed computing environment. We will conclude the talk with a few remarks on future direction of the discussed work.

Srikara Pranesh
University of Manchester
srikara.pranesh@manchester.ac.uk

Debraj Ghosh
Department of Civil Engineering
Indian Institute of Science
dghosh@iisc.ac.in

MS350

Robust Model Reduction for High-contrast Problems

Major progress has been made recently to make preconditioners robust with respect to any variation of coefficients inside and/or across the subdomains. A reason for this success is the adaptive selection of primal constraints based on localized generalized eigenvalue problems. In this talk we discuss how to transfer this technique to the field of discretizations. Given a target accuracy, we design a robust model reduction by delocalizing multiscale basis functions and establish a priori energy error estimates with such target accuracy with hidden constants independently of the coefficients. For the analysis, the regularity of the solution is not used, only that the right-hand side be in L2.

Marcus Sarkis
Worcester Polytechnic Institute
Instituto de Matematica Pura e Aplicada (Brazil)
msarkis@wpi.edu

Alexandre L. Madureira
Laboratorio Nacional de Computacao Cientifica
Departamento de Matematica Aplicada e Computacional
alm@lncc.br

MS351

A Highly-accurate Numerical Method for the Solution of the Two-dimensional Allen-Cahn Equation in Non-convex Polygons

Abstract not available

Ping Lin
University of Dundee
Scotland
plin@maths.dundee.ac.uk

Emine Celiker
University of Lincoln
eceliker@dundee.ac.uk

MS351

A Diffuse Domain Approach for Two Phase Flow in Complex Geometries

In this talk we will present a quasi-incompressible NSCH model for two-phase flows with variable density. This model will be coupled with a diffuse domain approach to mimic the fluid flow in a complex domain (DD-NSCH). The original complex physical boundary is now characterized by a DD variable that has a thin transition layer across the boundary. Validations, including the asymptotic analysis and numerical test, are presented to show that the DD-NSCH system converges to the original NSCH model as the transition layer of DD variable becomes thin enough. Several numerical examples that involve the fluid flows in complex domain will be presented to show the capability of the DD-NSCH system.

Zhenlin Guo
University of California - Irvine
zhenling@math.uci.edu

MS351

Phase-field Simulations of Contact Angle Hystere-

sis

When a fluid interface slides on a solid substrate, the advancing and receding contact angles may differ, which is known as contact angle hysteresis. We will present a novel boundary condition for contact angle hysteresis based on a simple modification of the phase-field wall energy relaxation model. By exploring the surface potential resulted from the mixing and wall energies, the advancing, receding, and pinning of contact line can be directly identified without the explicit knowledge of contact angle or contact-line velocity. Our boundary condition pins the contact line automatically if the microscopic contact angle falls between the advancing and receding values. Once the contact line moves, the correct microscopic contact angle is picked up and the contact-line dynamics with a single-valued static contact angle is recovered. For a 2D drop adhering to a wall under a pressure-driven creeping flow, our phase-field results agree very well with the boundary-integral results in literature. Based on the simulations of advancing menisci between parallel plates, we come up with guidelines on the choice of relaxation parameter to achieve desired pinning effect and slip length. In the end we will present 3D simulations on sliding drops on an inclined plane.

Pengtao Yue
Virginia Polytechnic Institute and State University
ptyue@vt.edu

MS351

Thermodynamically Consistent Phase-field Models and Their Numerical Approximations

In this talk, I will first discuss the generalized Onsager principle for deriving hydrodynamic theories (equations) for multiphase complex fluids. The generalized Onsager principle combined with variational principle would give a dissipative system that the total energy is decreasing in time. Then I will introduce the newly proposed energy quadratization approach to develop efficient, accurate and energy stable numerical approximations for dissipative systems. In the end, a few applications in would be discussed to demonstrate the effectiveness of the modeling and numerical tools.

Jia Zhao
Utah State University
jia.zhao@usu.edu

MS352

Interpolatory Model Reduction and Control of Quadratic Systems

The control of systems with quadratic nonlinearities appears in a wide range of applications. Controlling fluids is a notable example. Of course, to compute feedback control laws for these systems requires accurate reduced models that capture the input/output characterization of the system. Furthermore, the computation of stabilizing feedback controls are also a computational challenge that can be aided by reduced models. In this talk, we present our progress towards a feedback control methodology for these problems based on the interpolatory model reduction framework.

Jeff Borggaard
Virginia Tech
Department of Mathematics
jborggaard@vt.edu

MS352

Data-driven Reduced Order Modeling for Parametrized Problems

A data-driven reduced basis (RB) method is proposed for parametrized nonlinear problems, steady or unsteady. This method requires the offline preparation of a database comprising full-order solutions at different time-parameter locations. A reduced basis is extracted from the full-order data by the proper orthogonal decomposition (POD), and the maps between the time-parameter values and the projection coefficients onto the RB space are approximated as Gaussian process regression (GPR) models. For time-dependent problems, the singular-value decomposition (SVD) is utilized to extract the principal components in the data of the projection coefficients, and the regression functions are represented as the linear combinations of several tensor products of two Gaussian processes, one of the time-mode and the other of the parameter-mode. Via direct outputs from the regression models, the reduced-order solutions at new time-parameter locations can be recovered rapidly in the online stage. Featuring a non-intrusive nature, the offline and online stages are completely decoupled in the proposed scheme. The full-order solutions in the database can be generated from a 'black-box' solver, and the direct outputs from GPR models ensure efficient online evaluations for multi-query and real-time applications. The effectiveness of the scheme is illustrated by some non-trivial numerical examples.

Mengwu Guo
Ecole Polytechnique Fédérale de Lausanne
mengwu.guo@epfl.ch

Jan S. Hesthaven
EPFL
jan.hesthaven@epfl.ch

MS352

Hierarchical Solvers for Parametric Problems

The computational advantages provided by model reduction techniques play a crucial role when solving parametric problems. For instance, a rapid simulation of parametric problems becomes crucial for optimization, data assimilation, parameter estimation. In this communication, we focus on the parametric generalization of Hierarchical Model (HiMod) reduction techniques, to downscale parametric models associated with phenomenona exhibiting a preferential direction of flow. In particular, we propose two different approaches which merge HiMod with a Proper Orthogonal Decomposition (HiPOD methods) [1], and with a Reduced Basis approach (HiRB method), respectively. In both the cases, the idea is to employ HiMod discretization during the offline step, in order to contain the computational burden typical of this phase.

Thus, after introducing the basic HiMod framework, we focus on HiPOD/HiRB approaches and we verify the corresponding performances on some benchmark configurations. Finally, a cross-comparison between HiPOD and HiRB methods will be provided. [1] D. Baroli, C.M. Cova, S. Perotto, L. Sala and A. Veneziani. Hi-POD solution of parametrized fluid dynamics problems: preliminary results. In Model Reduction of Parametrized Systems. Series: MS&A Springer, P. Benner, M. Ohlberger, A.T. Patera, G. Rozza, K. Urban Eds., Chapter 15, pp. 235-254, 2017

Simona Perotto
MOX - Modeling and Scientific Computing
Dipartimento di Matematica
simona.perotto@polimi.it

MS352

Reduced-order Models for Advection-dominated Systems

In many industrial applications, model order reduction (MOR) techniques are employed to reduce the computational effort of multi-query simulations. Most MOR methods are based on a linear Galerkin ansatz which is, however, observed to be ineffective for systems where high-gradient structures (e.g. shocks) propagate through the computational domain. These kinds of dynamics are encountered in numerous applications including chemical engineering and climatology. In recent years, new MOR methods emerged which are tailored for systems whose dynamics are dominated by advection. One example is the shifted proper orthogonal decomposition (shifted POD) which approximates a snapshot matrix by a linear combination of modes shifted in space. In contrast to the classical POD, each mode is assigned a time-dependent shift which allows to track the possibly multiple advection speeds in the system. In this talk we present a new method for constructing a reduced-order model (ROM) via projection of the full-order model (FOM) based on shifted modes. In contrast to classical MOR methods, the proposed projection is nonlinear. The unknowns of the nonlinear ROM are the time-dependent amplitudes and shifts of the modes obtained in the offline phase. Moreover, we show how to achieve an efficient offline/online decomposition in the case that the FOM is linear. We illustrate the approach by means of numerical examples.

Philipp Schulze
Institute of Mathematics
TU Berlin
pschulze@math.tu-berlin.de

MS353

Microstructurally-induced Properties and Degradation in Lithium-ion Batteries

Rechargeable Lithium-Ion batteries, LIBs, are an energy storage technology that permeates many aspects of modern life. The technology is currently being miniaturized and integrated into new solid-state based gadgets while simultaneously being scaled up into large stationary energy storage containers as a new building block of an energy grid that aims to support alternative forms of energy. In all these cases, the starting chemistry and the microstructure that results the specific processing technique that generates LIBs are directly responsible for the power and energy density delivered on every discharge cycle, the short term Coulombic cycling capability, and sets the stage, the initial conditions, for its long term degradation and ultimate failure. In this paper, we quantify the effects of processing-induced microstructural features, such as particle size and morphological anisotropy, porosity and its associated tortuosity and reactivity constitutive properties of the individual components, as well as the effect of the internal mesostructure of the cathode and anode materials that integrate these devices on the resultant power response and its degradation mechanisms. By combining existing published experimental data and modeling and simulation we demonstrate the effects of the processing-induced meso- and microstructural features on the degradation of battery materials, such as graphite and lithium anodes, and LiCoO2, and LiMnO2-based cathodes.

Aniruddha Jana, Abhas Deva, Alfonso Campos, Lucas D. Robins
Purdue University
ajana@purdue.edu, adeva@purdue.edu,
campos26@purdue.edu, robin198@purdue

Ilenia Battiato
Stanford University
ibattiat@stanford.edu

Edwin Garcia
Purdue University
redwing@purdue.edu

MS353

Predictive Probabilistic Graphical Models for Energy Materials

We discuss a new class of Probabilistic Graphical Models (PGM) in energy research including models for materials design for catalysis, fuel cells and batteries. Since their introduction, PGMs have proved to be a fundamental mathematical concept for modeling uncertainty and causal relationships in Artificial Intelligence. Here the proposed PGMs need to involve multi-scale physicochemical mechanisms and incorporate available data, e.g. from electronic structure calculations or observational/experiment data. More specifically, the hierarchical structure of Probabilistic Graphical Models (PGM), allows us to bring together in a systematic way both statistical and multi-scale modeling, as well as available data, which here are typically sparse or incomplete, along with expert knowledge, correlations and causality relationships between model components and data. Furthermore, the process of building such models will necessarily involves numerous sources of uncertainty/error arising from different components of the PGM, corresponding to either data or modeling errors. In this direction, we also introduce information-based Uncertainty Quantification (UQ) methods capable to assess and improve the predictive ability of our proposed computa-

tional models.

Markos A. Katsoulakis
University of Massachusetts, Amherst
Dept of Mathematics and Statistics
markos@math.umass.edu

Eric J. Hall
RWTH Aachen University
hall@uq.rwth-aachen.de

MS353

Optimal Design of Nanoporous Materials for Electrical Storage Devices

Unique macroscopic properties of nanoporous metamaterials stem from their microscopic structure. Optimal design of such materials is facilitated by mapping a material's pore-network topology onto its macroscopic characteristics. This is in contrast to both trial-and-error experimental design and design based on empirical relations between macroscopic properties, such as the often-used Bruggeman formula that relates a material's effective diffusion coefficient to its porosity. We use homogenization to construct such a map in the context of materials design that maximizes energy/power density performance in electrochemical devices. For example, effective diffusion coefficients and specific surface area, key macroscopic characteristics of ion transport in a hierarchical nanoporous material, are expressed in terms of the material's pore structure and, equally important, ion concentrations in the electrolyte and externally applied electric potential. Using these microscopic characteristics as decision variables, we optimize the macroscopic properties for two two-dimensional material-assembly templates and several operating conditions. The latter affect the material's performance through formation of an electrical double layer at the fluid-solid interfaces, which restricts the pore space available for ion transport.

Xuan Zhang
GEIDCO
zhang.xuam@gmail.com

Daniel M. Tartakovsky
Stanford University
tartakovsky@stanford.edu

MS353

Homogenization-informed Convolutional Neural Networks for Identification of Effective Properties of Energy Storage Devices

Batteries are electrochemical energy storage devices that exhibit physico-chemical heterogeneity on a continuity of scales. As such, battery systems are amenable to mathematical descriptions on a multiplicity of scales that range from atomic to continuum. First, we present a new method to assess the veracity of macroscopic models of lithium-ion batteries. Macroscopic models treat the electrode as a continuum and are often employed to describe the mass and charge transfer of lithium since they are computational tractable and practical to model the system at the cell scale. Yet, they rely on a number of simplifications and assumptions that may be violated under given operating conditions. We use multiple-scale expansions to upscale the pore-scale Poisson-Nerst-Planck equations and establish sufficient conditions under which macroscopic dual-continua mass and charge transport equations accurately represent pore-scale dynamics. For the first time, applicability conditions of macroscopic models in terms of electric Péclet and Damköhler are defined. Then, we show how convolutional neural networks can be used in the context of homogenization framework to calculate effective transport properties of Li-ion batteries from XCT and SEM images of battery electrode microstructures.

Ilenia Battiato
Stanford University
ibattiat@stanford.edu

Ross M. Weber
Stanford University
Energy Resources Engineering
rweber3@stanford.edu

MS354

Implementation of Sparse FFT with Structured Sparsity

For general frequency sparse signals a first deterministic combinatorial Fourier algorithm for estimating the best k-term Fourier representation for a given frequency sparse signal was introduced in 2010 (see Iwen, Combinatorial Sublinear-Time Fourier Algorithms. Found. Comput. Math., 10(3):303-338). It relies heavily on the Chinese Remainder Theorem and combinatorial concepts, and has a runtime that is sublinear in the input length and scales quadratically in the sparsity k. Adapting ideas used therein, we develop a deterministic sparse Fourier algorithm for input functions where we generalize the setting of a short frequency support to a *structured frequency support*, especially allowing a support consisting of multiple long intervals. Our method then achieves a runtime that is sublinear in the input length and scales sub-quadratically in the sparsity, which, for the class of functions considered, has so far only been achieved by randomized techniques. Furthermore, our algorithm maintains an ℓ^2/ℓ^1 error guarantee of the same type as the error guarantees of many compressive sensing methods.

Sina Bittens
University of Göttingen
sina.bittens@mathematik.uni-goettingen.de

Ruochuan Zhang, Mark Iwen
Michigan State University
zhangr12@msu.edu, markiwen@math.msu.edu

MS354

Discrete Sparse Fourier Transforms: Faster Stable Implementations with Guarantees

We consider Sparse Fourier Transform (SFT) algorithms for approximately computing the best s-term approximation of the Discrete Fourier Transform (DFT) $\hat{\mathbf{f}}$ of any given input vector $\mathbf{f} \in \mathbb{C}^N$ in just $O(s \log?N)^{O(1)}$ -time using only

a similarly small number of entries of $\mathbf{f}$. When $s \ll N$ the runtimes of these methods are significantly faster than the $O(N \log^c N)$ runtimes of the fastest standard compressive sensing and fast Fourier transform methods. In particular, we present a deterministic SFT algorithm which is guaranteed to always recover a near best s-term approximation of the DFT of any given input vector $\mathbf{f}$ in $O(s^2 \log^{5.5}?(N))$-time. Unlike previous deterministic results of this kind, our deterministic result holds for both arbitrary vectors $\mathbf{f}$ and vector lengths N. In addition to these deterministic SFT results, we also develop several new publicly available randomized SFT implementations for approximately computing $\hat{\mathbf{f}}$ from $\mathbf{f}$ using the same general techniques. The best of these new implementations is shown to outperform existing discrete sparse Fourier transform methods with respect to both runtime and noise robustness for large vector lengths N.

Mark Iwen, Sami Merhi, Ruochuan Zhang
Michigan State University
markiwen@math.msu.edu, merhisam@math.msu.edu, zhangr12@msu.edu

MS354

Getting Best Performance of Memory Bandwidth Limited Algorithms with Intel MKL

This talk is devoted to the functionality provided by Intel®Math Kernel Library (Intel®MKL) to support Linear Algebra computations. Modern multi-/many-core architectures have certain limitations for memory bandwidth limited operations such as sparse matrix vector multiplication and Fast Fourier Transforms. Depending on the type of problem to solve, it is still possible to find a smart way to minimize the impact of these limitations and explore benefits to their full potential in order to tackle problems of extreme sizes. This talk provides a unique outlook on the underlying optimizations and features that enable the best performance of Sparse Linear Algebra and FFT functionality on Intel architectures.

Alexander Kalinkin
Intel Corporation
alexander.a.kalinkin@intel.com

MS354

Implementation of Parallel 3-D Real FFT with 2-D Decomposition on Intel Xeon Phi Clusters

In this talk, we propose an implementation of parallel three-dimensional real fast Fourier transforms (FFTs) with two-dimensional decomposition on Intel Xeon Phi clusters. The proposed parallel three-dimensional FFT algorithm is based on the conjugate symmetry property for the discrete Fourier transform (DFT) and the multicolumn FFT algorithm. We show that a two-dimensional decomposition effectively improves performance by reducing the communication time for larger numbers of MPI processes. We also present a computation-communication overlap method that introduces a communication thread with OpenMP. Performance results of three-dimensional real FFTs on Intel Xeon Phi clusters are reported.

Daisuke Takahashi
Center for Computational Sciences
University of Tsukuba
daisuke@cs.tsukuba.ac.jp

MS355

An Overview of Nystrom (and Not-So-Nystrom) High-order Surface Quadratures for Fast Solvers

The boundary integral equation approach is natural and efficient for solving linear BVPs where a fundamental solution is known; this covers a wide variety of PDE such as Laplace, (modified) Helmholtz, Maxwell, Stokes, elastostatics/dynamics, etc. Its benefits include the number of unknowns scaling only as the complexity of the surface data, conditions at infinity "for free", well-conditioned linear systems upon discretization, and fast iterative/direct solvers exploiting off-diagonal low-rank structure. A chief obstacle to the universal and automatic use of such solvers to achieve many accurate digits in 3D is the difficulty of high-order surface quadratures for the relevant singular kernels, and accurate potential evaluation near surfaces. We overview methods of the Nystrom flavor, including auxiliary nodes, partitions of unity, and QBX, discuss high-order approaches to corners and edges, and 3D-specific challenges such as aspect ratio. If time, we mention our recent work on a (non-second-kind but singular-quadrature-free) alternative, the method of fundamental solutions.

Alex H. Barnett
Dartmouth College and Simons Foundation
abarnett@flatironinstitute.org

MS355

An Efficient and High Order Accurate Solution Technique for Three Dimensional Elliptic Partial Differential Equations

The Hierarchical Poincaré-Steklov (HPS) method is a high order accurate discretization technique that comes with an efficient direct solver. The method is robust, even for highly oscillatory solutions, and can easily handle variable coefficient PDEs. The nested dissection-based direct solver can easily be accelerated to have linear complexity by exploiting the fact that the intermediate dense matrices are amenable to fast linear algebra. For two dimensional problems, this yields an elliptic PDE solution technique that can solve a problem with over 100 million unknowns with two hours of precomputation and 30 seconds per solve on a desktop computer. Unfortunately, out of the box fast linear algebra techniques are not sufficient for substantial acceleration for three dimensional problems. In this talk, we present a new version of the HPS method which uses potential theory techniques to gain further acceleration of the direct solver. Numerical results will illustrate the performance of the method.

Natalie N. Beams
University of Illinois at Urbana-Champaign
Natalie.Beams@rice.edu

Adrianna Gillman
Rice University
Department of Computational and applied mathematics

adrianna.gillman@rice.edu

MS355

Integral Equation Methods for the Heat Equation in Moving Geometries

Many problems in physics and engineering require the solution of the heat equation in moving geometry. Integral representations are particularly appropriate in this setting since they satisfy the governing equation automatically and, in the homogeneous case, require the discretization of the space-time boundary alone. In this talk, we present a new method for solving the 2D heat equation in moving geometry that makes use of a spatially adaptive mesh, a new version of the fast Gauss transform that allows for volume and boundary sources, and a new hybrid asymptotic/numerical method for local-in-time quadrature.

Jun Wang
Flatiron Institute
Simons Foundation
jwang@flatironinstitute.org

Leslie Greengard
Simons Foundation and Courant Institute
New York University
greengar@cims.nyu.edu

MS355

An Adaptive Technique for 2D Boundary Integral Equation.

Many problems in science and engineering can be recast as solving integral equations with elliptic kernels. For some applications, the geometry and the underlying physics may lead to difficulty in discretization. For such problems, uniform discretization, even with large number of nodes, may not be able to generate accurate enough solutions. The talk will introduce an adaptive discretization technique for 2D boundary integral equations that considers both the geometric complexity and the underlying physics and identifies the part of geometry that requires additional refinement. A fast direct solver is built to work with this technique for the computational efficiency. Numerical results will demonstrate the performance of this adaptive discretization technique over a collection of problems.

Yabin Zhang
Dept. of Computational and Applied Mathematics
Rice University, Houston, TX
YZ89@RICE.EDU

MS356

Extracting Actionable System-application Performance Factors

Actionable modeling of resource utilization and performance-impacting conditions depends on the ability to quickly and reliably extract factors that are indicators of conditions of interest. Identification of such factors is difficult in HPC systems where the complexity of interactions is such that relationships between data and performance must be extracted via exploration, experimentation, and inference. Relatedly, the utility of discovered factors is limited when the presentation of conditions with respect to possible performance impact is unclear or limited in generality or accuracy. We have been developing and deploying an end-to-end infrastructure for generation of meaningful resource utilization characterizations about shared resource contention which includes data collection, analysis, presentation and feedback. We utilize continuous data reflective of system and application demands on memory, network, filesystems etc. to develop job-centric and system-wide characterizations during run time. We present experiences in this work on a variety of production and testbed systems, highlighting details in the extraction, assessment, quantification, and presentation of conditions of interest and resultant impact.

James Brandt, Ann Gentile
Sandia National Laboratories
brandt@sandia.gov, gentile@sandia.gov

Jonathan Cook
New Mexico State University
joncook@nmsu.edu

MS356

Holistic Performance Diagnosis for HPC Workflows

Holistic HPC Workflow Performance Diagnosis encompasses the monitoring and analysis of performance problems that span across traditionally separated aspects of an HPC effort, requiring us to monitor both systems and applications. In this talk, we describe our experiences conducting a performance analysis of a drought prediction code developed at Portland State. This effort motivates the need for new tools and methods for Holistic Performance Analysis. We detail the particular challenges encountered, and describe our proposed visualization approach and the prototype under development.

Karen L. Karavanic
Portland State University
karavan@pdx.edu

MS356

Performance Analysis of a Multi-application Plasma Surface Interactions Workflow

We recently conducted increasingly realistic and complex simulations of the impact of a burning plasma on the surface of plasma-facing components as part of the US Department of Energy's Office of Fusion Energy Sciences' FY2018 Plasma Materials Interaction theory milestone experiments. Our workflow involved the Fractal TRIDYN ion collision simulator, the Global Impurity Transport (GITR) impurity migration simulator, and the Xolotl cluster dynamics simulator, coupled using the Integrated Plasma Simulator (IPS) workflow management software. Early in the year, we devised a simple, non-intrusive approach for collecting and visualizing workflow performance data that extracts data from the IPS event log, converts it to an Open Trace Format file, and visualizes it using the Vampir visualization tool. Applying this approach yielded a

near-immediate benefit: upon seeing the performance visualization from a test run, we identified a workflow misconfiguration for the GITR component that caused it run longer than expected. Fixing the mis-configuration led to a 28× improvement in simulation time over a 48-hour workflow run. Later in the year, we developed a more sophisticated performance data collection approach that can monitor activity within codes and that can better associate workflow activities with the physical resources on which they run.

Philip C. Roth
ORNL
rothpc@ornl.gov

Ane Lasa, Sophie Blondel, Timothy Younkin
University of Tennessee-Knoxville
aesquisa@utk.edu, sblondel@utk.edu,
tyounkin@vols.utk.edu

Jon Drobny
University of Illinois at Urbana-Champaign
drobny2@illinois.edu

Mark Cianciosa, Wael R. Elwasif
Oak Ridge National Laboratory
cianciosamr@ornl.gov, elwasifwr@ornl.gov

Davide Curreli
University of Illinois at Urbana-Champaign
dcurreli@illinois.edu

John Canik
Oak Ridge National Laboratory
canikjm@ornl.gov

David E. Bernholdt
Oak Ridge National Laboratory
Computer Science and Mathematics Division
bernholdtde@ornl.gov

Brian Wirth
University of Tennessee-Knoxville
bdwirth@utk.edu

MS356

Challenges and Opportunities in Capturing, Modeling, and Actively Managing Power and Energy Usage on HPC Systems

In upcoming HPC systems, both power and energy consumption will face strict limits, imposed by technological, political and facility constraints. In order to still enable efficient execution of applications, we will have to actively manage these limited resources in a way that matches both system and application needs. In particular, we need to identify where either power or energy is needed to make progress and then direct power/energy to these sections of the execution. This requires the ability to monitor and observe power and energy utilization and the ability to correlate this with application context, providing the foundation for the creation of highly accurate models based on this observational data in order to drive future decisions. In this talk I will present the current state of the art at the Leibniz Supercomputing Centre in Munich, which is one of the leading institutions in terms of power and energy management, and which provides a comprehensive measurement system. Further, I will discuss ongoing activities, through an open and international group of researchers, towards a fully integrated power stack that allows us to close the loop and turn measurements into insights and then into actions optimizing the system. This will form one the pillars in breaking the power and energy wall we are facing and to continue to scale the capability of our HPC systems.

Martin Schulz
TU-Munich
schulzm@in.tum.de

MS357

Lanczos Bidiagonalization, Orthogonal Polynomials, and Singular Value Distributions

Abstract not available.

Anthony P. Austin
Argonne National Laboratory
austina@anl.gov

MS357

A Contour Intergral-based Stochastic Estimator for Eigenvalue Counts of Generalized Eigenproblems

In this talk, we present a contour integral-based stochastic estimator for eigenvalue counts of large-sparse generalized eigenproblems. In the algorithm of the estimator, one needs to solve inner linear systems with respect to quadrature points of the discretized contour integral. In this study, we consider using Krylov subspace methods for solving the inner linear systems. We show that, in conjunction with some Krylov subspace methods, the contour integral-based estimator can be regarded as a certain polynomial-type filter for the spectrum. We also show the efficiency of the estimator using practical numerical examples.

Yasunori Futamura
University of Tsukuba
futamura@cs.tsukuba.ac.jp

Takahiro Yano, Akira Imakura, Tetsuya Sakurai
Department of Computer Science
University of Tsukuba
yano@mma.cs.tsukuba.ac.jp, imakura@cs.tsukuba.ac.jp,
sakurai@cs.tsukuba.ac.jp

MS357

Beyond Automated Multilevel Substructuring: Domain Decomposition with Rational Filtering

This talk discusses a rational filtering domain decomposition technique for the solution of large and sparse symmetric generalized eigenvalue problems. The proposed technique is purely algebraic and decomposes the eigenvalue problem associated with each subdomain into two

disjoint subproblems. The first subproblem is associated with the interface variables and accounts for the interaction among neighboring subdomains. To compute the solution of the original eigenvalue problem at the interface variables we leverage ideas from contour integral eigenvalue solvers. The second subproblem is associated with the interior variables in each subdomain and can be solved in parallel among the different subdomains using real arithmetic only. Compared to rational filtering projection methods applied to the original matrix pencil, the proposed technique integrates only a part of the matrix resolvent while it applies any orthogonalization necessary to vectors whose length is equal to the number of interface variables. In addition, no estimation of the number of eigenvalues located inside the interval of interest is needed. Numerical experiments performed in distributed memory architectures illustrate the competitiveness of the proposed technique against rational filtering Krylov approaches.

Vasileios Kalantzis, Yuanzhe Xi
University of Minnesota
kalan019@umn.edu, yxi26@emory.edu

Yousef Saad
Department of Computer Science
University of Minnesota
saad@cs.umn.edu

MS357

The EVSL Package for Symmetric Eigenvalue Problems

A number of applications require the computation of tens of thousands of eigenvalues of very large matrices. For these problems, it is imperative to take advantage of spectrum slicing strategies whereby different "slices" of the spectrum are extracted independently. The presentation will begin by describing a general approach for spectrum slicing based on polynomial filtering. This approach can be quite efficient in the situation where the matrix-vector product operation is inexpensive and when a large number of eigenvalues is sought. Polynomial filtering can be combined with the Lanczos algorithm with and without restarts, as well as with subspace iteration. An alternative to polynomial filtering that is generating a growing interest is a class of methods that exploit filtering by rational functions. Good representatives of this general approach are the FEAST eigensolver and the Sakurai-Sugiura algorithm. Here we will argue that the standard Cauchy integral–based approach can be substantially improved upon – especially when iterative solvers are involved. These two classes of techniques have recently been implemented in a code named EVSL (for eigenvalues slicing library) and the presentation will end with the latest updates to the code and our progress with its (forthcoming) parallel version.

Yousef Saad
University of Minnesota
saad@umn.edu

Ruipeng Li
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
li50@llnl.gov

Yuanzhe Xi
University of Minnesota
yxi26@emory.edu

MS358

Exponential Propagation in Time Integration.

Over the past decades, exponential integration emerged as a numerical technique that offers significant computational savings for integration of large scale stiff systems. Similar to implicit approach, exponential propagation can be used to construct a fully exponential time integrator or used together with explicit or implicit treatment of certain terms of the system. In this talk we explore different ways in which a time integrator that uses exponential propagation can be constructed. We discuss fully exponential methods such as exponential Rosenbrock or EPIRK integrators as well as hybrid methods such as implicit-exponential (IMEXP) schemes. Different approaches to constructing exponential propagation-type integrators such as Runger-Kutta or polynomial interpolation frameworks will be presented. The key to designing an efficient exponential time integrator is speeding up computation of exponential-like functions of large matrices. We will discuss some of the most efficient approaches to constructing and implementing such methods and present EPIC software package that enables practitioners to test and to use these techniques.

Mayya Tokman
University of California, Merced
School of Natural Sciences
mtokman@ucmerced.edu

Tommaso Buvoli
Department of Applied Mathematics, University of Washington
Seattle, WA 98195 USA
tbuvoli@ucmerced.edu.

MS358

Partitioned Exponential Methods for Stiff Differential Equations

Mixed methods of time integration, such as implicit-explicit (IMEX) methods, have long been used to solve semi-discretized partial differential equations composed of a stiff and a non-stiff part. As an example, in reaction-diffusion equations, the stiff part, usually, is the diffusion term and is dealt with using the implicit method, and the reaction term with the explicit method. In the event that the reaction term is also stiff, IMEX methods may no longer be computationally efficient; the explicit method will face severe timestep restriction due to the stiffness originating from the reaction term. A number of authors have addressed this issue by way of building methods that treat one of the terms implicitly and integrating the other term using an exponential integrator (see IMEXP, implicit-exponential, and IIF, implicit-integration factor, methods). Exponential integrators, as has been established through the work of many people, are more stable than explicit methods on stiff problems and can be more efficient than implicit methods. The question then arises: can we substitute exponential methods in place of implicit methods so as

to derive new classes of partitioned methods? In this presentation, we derive, construct and analyze such methods, namely, exponential-exponential and exponential-explicit methods.

Mahesh Narayanamurthi
Virginia Tech, USA
maheshnm@vt.edu

Adrian Sandu
Virginia Polytechnic Institute
and State University
sandu@cs.vt.edu

MS358

A Class of Multirate Infinitesimal GARK Schemes

Differential equations arising in many practical applications are characterized by multiple time scales. Multirate time integration seeks to solve them efficiently by discretizing each scale with a different, appropriate time step, while ensuring the overall accuracy and stability of the numerical solution. In a seminal paper Knoth and Wolke (1998) proposed a hybrid solution approach: discretize the slow component with an explicit Runge-Kutta method, and advance the fast component via a modified fast differential equation. The idea led to the development of multirate infinitesimal step (MIS) methods in (Wensch et al., 2009). Guenther and Sandu (2016) explained MIS schemes as a particular case of multirate General-structure Additive Runge-Kutta (MR-GARK) methods. The hybrid approach offers extreme flexibility in the choice of the numerical solution process for the fast component. This work constructs a new family of multirate infinitesimal GARK schemes (MRI-GARK) that extends the hybrid dynamics approach of in multiple ways. Order conditions theory and stability analyses are developed, and practical explicit and implicit methods of up to order four are constructed. Numerical results confirm the theoretical findings. We expect the new MRI-GARK family to be most useful for systems of equations with widely disparate time scales, where the fast process is dispersive, and where the influence of the fast component on the slow dynamics is weak.

Adrian Sandu
Virginia Polytechnic Institute
and State University
sandu@cs.vt.edu

MS358

High-order Operator-Splitting Methods for Cardiac Tissue Simulations

Cardiac tissue simulations often use the bidomain and monodomain models to describe the electrophysiology of cardiac tissue. These models take the form of multi-scale reaction-diffusion partial differential equations that couple the dynamic behaviour on the cellular scale with that on the tissue scale. The systems of differential equations associated with these models are large and strongly non-linear, but they also have a distinct structure due to their multi-scale nature. Operator-splitting methods attempt to take advantage of this structure to efficiently produce numerical solutions. The focus of this presentation is on operator-splitting methods with order higher than two. Such methods require backward time integration in each operator and historically have been considered unstable for solving deterministic parabolic systems. The stability and performance of operator-splitting methods of up to order four to solve the bidomain and monodomain models are demonstrated on several examples arising in the field of cardiovascular modelling.

Raymond J. Spiteri
University of Saskatchewan, Canada
Department of Computer Science
spiteri@cs.usask.ca

MS359

Goal Oriented Sensor Placement for Large-scale Inverse Problems Under Model Uncertainty

We present scalable optimal experimental design (OED) methods for large-scale Bayesian linear inverse problems. Our approach is based on using efficient randomized matrix methods for fast computation of OED objective and its gradient, and on sparsifying penalty functions to enforce sparsity of the sensor placements. Additionally, we will discuss OED for prediction, i.e., goal oriented OED, and OED for inverse problems that are governed by models with uncertainties other than those in the inversion parameters, i.e., OED under uncertainty. Goal oriented OED and OED under uncertainty require appropriate reformulations of traditional OED criteria that will be discussed in the talk, along with numerical methods that are tailored to the resulting problem formulations. Illustrative numerical results will be provided in the context of sensor placement for optimal recovery of the initial state in an advection-diffusion problem.

Alen Alexanderian
NC State University
alexanderian@ncsu.edu

Ahmed Attia
Argonne National Laboratory
attia@anl.gov

Noemi Petra
University of California, Merced
npetra@ucmerced.edu

Arvind Saibaba
North Carolina State University
asaibab@ncsu.edu

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

MS359

An Optimal Experimental Design Framework for Adaptive Inflation and Covariance Localization for

Ensemble Filters

Ensemble Kalman Filter (EnKF) is the main ensemble-based data assimilation algorithm for large-scale applications including numerical weather prediction. Inflation, and covariance localization are both necessary for almost any EnKF implementation, to avoid filter divergence. While fixed, e.g. space-independent, covariance inflation and localization can increase the reliability of the estimated covariance matrices, in many cases it may result in information loss. Tuning the space-time inflation and localization parameters is a hard problem, and has to be done adaptively and automatically. In this talk, we present a new variational framework for adaptive inflation and covariance localization in EnKF. We also present computational results using a standard Lorenz-96 experiment.

Ahmed Attia
Argonne National Laboratory
attia@anl.gov

Emil M. Constantinescu
Argonne National Laboratory
Mathematics and Computer Science Division
emconsta@mcs.anl.gov

MS359

Assimilating Data in Models with Stochastic Parameters

In this talk we introduce a statistical treatment of inverse problems with stochastic terms such as SPDEs. When standard Monte-Carlo techniques are applied, the solution of the forward problem is given by a distribution represented numerically by an ensemble of simulations. We hereby develop a strategy to solve such inverse problems by expressing the objective as finding the closest forward distribution that best explains the distribution of the observations in a certain metric. We propose to use metrics typically employed in forecast verification.

Emil M. Constantinescu
Argonne National Laboratory
Mathematics and Computer Science Division
emconsta@mcs.anl.gov

MS359

Adaptive Multivariate Schur Product Localization in Ensemble-based Filters

Undersampling in multivariate Monte-Carlo methods such as ensemble Kalman filters is an area of much static heuristic correction. Schur-product based localisation is a theoretically valid method for correcting undersampled covariance matrices. We present a Bayesian approach to adaptive-in-time Schur-product based localisation for a variant of the Ensemble Kalman Filter, and validate it on a series of toy geophysical models.

Andrey Popov
Virginia Polytechnic Institute and State University
apopov@vt.edu

Adrian Sandu
Virginia Polytechnic Institute
and State University
sandu@cs.vt.edu

MS360

Lubricated Immersed Boundary Method Simulations of Elastic Bodies in Near Contact

Fluid-mediated near contact of elastic structures is a recurring theme in biofluids, e.g. the flow of red blood through blood vessels, in which near-contact occurs at multiple levels including both cell-cell and cell-wall interactions, and the bundling of bacterial flagella in E. Coli. The thin fluid layers that arise during near-contact are difficult to resolve by standard computational fluid dynamics methods based on uniform fluid grids. Motivated by this fluid-structure interaction problem, we have developed an immersed boundary method that uses elements of lubrication theory as a subgrid model to resolve the thin fluid layers between immersed boundaries. In contrast to methods based on adaptive mesh refinement, our approach does not impose additional restrictions on the timestep. We have applied this lubricated immersed boundary method to 2D flows of increasing complexity, including suspensions of cells and rotating flagellar cross-sections. We find that our method gives accurate results and reduces numerical artifacts.

Thomas Fai
Brandeis University
tfai@brandeis.edu

MS360

Slippery Multicomponent Vesicles

Multicomponent vesicles are composed of cholesterol that combine with saturated lipids to form energetically stable domains on the vesicle surface. The presence of different lipid species lead to varying material properties, such as bending rigidity, produce a rich variety of dynamics as seen in experiments. In this work, a three-dimensional model is developed to study multicomponent vesicle dynamics in the presence of an externally driven fluid. The domains on the membrane experience an effective velocity, differing from the surrounding fluid velocity due to the molecular diffusivity of the lipids. Unlike prior modeling efforts, this effective velocity is now considered. The membrane surface is modeled using a two-phase Cahn-Hilliard equation using a level set/closest point method while the membrane is coupled to the surrounding fluid using the energy variation approach. The dynamics observed by this predictive model, influence of material properties and impact of the surface effective velocity will be discussed. We envision applications in measuring surface properties of biological cells and manufacturing of designer vesicles.

Prerna Gera
University of Wisconsin-Madison
pgera@math.wisc.edu

David Salac
University at Buffalo - SUNY
davidsal@buffalo.edu

Saverio E. Spagnolie
University of Wisconsin-Madison
Department of Mathematics
Spagnolie@math.wisc.edu

MS360

Multiphysics Computational Modeling of the Chick Embryonic Heart

Current models of the embryonic heart have focused on a single physical aspect: fluid dynamics, mechanics, contractility or electrophysiology. The interplay between these physical processes is crucial for our understanding of cardiogenesis. Some studies have started looking at the interplay between these functions experimentally but, as contractility and blood flow have proved difficult to uncouple, the combined effects are difficult to interpret. Computational models can be used to investigate such a complex system. Within this project, we use imaging techniques on the chick embryo to create a computational model of the heart at various stages of development considering both electromechanical coupling and hemodynamics. We use tetrahedral finite elements to solve the electromechanical coupling, while the fluid-structure interaction is handled by an immersed boundary finite difference/finite element algorithm. These computational models, used in combination with biological data on cell adaptation, have the potential to revolutionize the field of cardiac development and our understanding of congenital heart disease. Such computational models can help to unlock biological processes offering different perspectives than those provided by experiments.

Simone Rossi
Department of Mathematics
University of North Carolina at Chapel Hill
simone.rossi@unc.edu

Michael Bressan
Department of Cell Biology and Physiology
University of North Carolina at Chapel Hill
michael_bressan@med.unc.edu

Boyce E. Griffith
Departments of Mathematics and Biomedical Engineering
University of North Carolina at Chapel Hill
boyceg@email.unc.edu

MS360

Computational Mean-field Modeling of Confined Active Fluids

We will present a novel computational method for simulating confined active fluids, described by a continuum mean-field representation. It is constructed as a stable hybrid Finite Volume/Finite Difference method on adaptive Quadtree and Octree grids. The confining geometry is capture via a level-set representation, allowing for virtually any configuration to be considered. Simulations in simple two-dimensional domains are presented to verify and validate the method. More challenging examples will also be discussed, illustrating the full capabilities of our approach. In particular, we will examine flows in two-dimensional lattices and non-trivial three-dimensional geometries.

Maxime Theillard
University of California, Merced
mtheillard@ucmerced.edu

David Saintillan
Mechanical and Aerospace Engineering
UCSD
dsaintillan@eng.ucsd.edu

MS361

Verification of Uncertain Pomdps via Lyapunov Functions and Barrier Certificates

Abstract not available.

Mohamadreza Ahmadi
University of Texas
mrahmadi@utexas.edu

MS361

Robust Reconstruction of Flow Fields from Limited Measurements

In many applications it is important to estimate the structure of a flow field from limited and possibly corrupt measurements. Many current methods in flow estimation use least squares regression to reconstruct the flow field, which searches for the minimum-energy solution that is consistent with the measured data. However, this approach is known to be prone to overfitting and is sensitive to noise. To address these challenges, we instead seek a sparse representation of the data in a library of examples rather than a minimum-energy solution. Sparse representation has been widely used in image recognition and reconstruction examples, and is well-suited to structured data with limited measurements, corruption, and outliers. We demonstrate sparse representation for flow reconstruction using various fluid data sets, including vortex shedding past a cylinder at low Reynolds number, a mixing layer, and two geophysical flows. In particular we find considerable improvements in overall estimation accuracy and noise robustness, compared with least squares methods such as gappy POD. Sparse representation is a promising framework for extracting useful information from complex flow fields with realistic measurements.

Jared L. Callaham
University ofWashington
jc244@uw.edu

Kazuki Maeda, Steven Brunton
University of Washington
maeda@uw.edu, sbrunton@uw.edu

MS361

Kernel Koopman Spectral Analysis for Nonlinear Dynamical Systems

Spectral analysis of Koopman operators has attracted attention as a powerful tool for extracting global modal descriptions of nonlinear dynamical systems without explicit

knowledge about the governing equations. Here, we describe a data-driven method for obtaining those approximations through calculating the spectral decompositions of transfer operators in reproducing kernel Hilbert spaces. Also, we further describe its extension with vector-valued kernels for extracting dynamics of interactions among observables. We show several empirical examples of applications of these methods to synthetic and real-world data.

Yoshinobu Kawahara
ISIR, Osaka University
ykawahara@sanken.osaka-u.ac.jp

MS361

Clustering of Time Series via Dynamic Mode Decomposition and the Matrix Pencil Method

We present a novel algorithm for extracting features from sequences of multidimensional observations for classification purposes. Our method is based on the independently developed Matrix Pencil method and the Dynamic Mode Decomposition (with delay-embedding). These reduced-order least-squares approaches fit damped sinusoidals to time series and produce estimates of their frequencies and magnitudes. These estimates, equipped with an appropriate metric, can be used as features for machine learning and classification. Instead of fitting one series at a time, however, we suggest fitting all of the data at once. This approach yields a small number of frequencies shared by all of the sequences and a corresponding set of amplitudes for each sequence individually. The resulting vectors of amplitudes can then be used as features for any standard clustering technique. For illustration purposes, we apply this method to regions of different grain orientation in a Transmission Electron Microscopy image and use a constrained version of hierarchical agglomerative clustering to successfully identify those regions.

Leonid Pogorelyuk
Princeton University
leonidp@princeton.edu

Clarence Rowley
Princeton University
Department of Mechanical and Aerospace Engineering
cwrowley@princeton.edu

MS362

On the Best Approximation of the Hierarchical Matrix Product

The multiplication of matrices is an important arithmetic operation in computational mathematics. In the context of hierarchical matrices, this operation can be realized by the multiplication of structured block-wise low-rank matrices, resulting in an almost linear cost. However, the computational efficiency of the algorithm is based on a recursive scheme which makes the error analysis quite involved. We propose a new algorithmic framework for the multiplication of hierarchical matrices. It improves currently known implementations by reducing the multiplication of hierarchical matrices towards finding a suitable low-rank approximation of sums of matrix-products. We propose several compression schemes to address this task. As a consequence, we are able to compute the best-approximation of hierarchical matrix products. A cost analysis shows that, under reasonable assumptions on the low-rank approximation method, the cost of the framework is almost linear with respect to the size of the matrix. Numerical experiments show that the new approach produces indeed the best-approximation of the product of hierarchical matrices for a given tolerance. They also show that the new multiplication can accomplish this task in less computation time than the established multiplication algorithm without error control.

Jürgen Dölz
Technische Universität Darmstadt
doelz@gsc.tu-darmstadt.de

Helmut Harbrecht
Universitaet Basel
Department of Mathematics and Computer Science
helmut.harbrecht@unibas.ch

Michael D. Multerer
Università della Svizzera italiana
michael.multerer@usi.ch

MS362

Hierarchical Matrices in Stable Radial Basis Function Interpolation

The most straightforward way, RBF-Direct, to calculate Gaussian based RBF interpolants leads to dense, ill-conditioned systems of equations. On the positive side, the system matrix can be efficiently approximated by rank-structured matrix formats, and preconditioners can typically be constructed efficiently as well. (Analytic) basis transformations by the name of RBF-QR or RBF-GA lead to dense, better conditioned interpolation matrices but may be less amendable to approximation by rank-structured matrices and the efficient construction of (additional) preconditioners, hence leading to significantly higher costs. In this talk, we discuss approximation errors in relation to computational costs of the underlying methods.

Sabine Le Borne
Hamburg University of Technology
Department of Mathematics
leborne@tuhh.de

MS362

Inversion of Rank-structured Matrices: Which Formulas Should You Use?

Abstract not available.

Per-Gunnar Martinsson
University of Oxford
pgm@ices.utexas.edu

MS362

Accelerated Interpolative Decompositions by Gen-

eral Proxy Point Methods

Interpolative decomposition (ID) is a popular low-rank form widely used in structured matrix representations such as butterfly factorizations and hierarchical matrices. While dealing with a kernel matrix block, e.g., $K(X_0, Y_0)$ with a kernel function $K(x, y)$ and two point clusters X_0, Y_0, variant proxy point methods have been proposed to accelerate the ID approximation compared to the more expensive QR-based algebraic approach. In this talk, we summarize the general proxy point method that converts the ID approximation problem of $K(X_0, Y_0)$ to that of $K(X_0, Y_p)$ with a proxy point set Y_p that is manually selected. A rigorous error analysis of the proxy point method is provided. Furthermore, based on the error analysis, we discuss the selection of proxy point set Y_p, which should adapt to the kernel function and also to the domains containing X_0 and Y_0.

Xin Xing
School of Mathematics
Georgia Institute of Technology
xxing33@gatech.edu

Edmond Chow
School of Computational Science and Engineering
Georgia Institute of Technology
echow@cc.gatech.edu

MS363

Towards Robust Methods for Practical Aero Dynamic Flows: a Review of Recent Activities at NASA LaRC on Methods with the Summation by-parts Property

In this talk, we give a brief review of the summation-by-parts (SBP) concept and position it as a matrix analysis framework for the development of modern stable and conservative numerical methods. We then give an overview of ongoing research activities and collaborations at NASA LaRC including h-p entropy stable algorithms, fully discrete entropy stable algorithms, error estimates for SBP methods, decoupled SBP operators, and a number of other topics. The presentation will be rounded out by discussing future research directions.

David C. Del Rey Fernandez
NASA Langley Research Center
National Institute of Aerospace, Hampton
dcdelrey@gmail.com

Mark H. Carpenter
NASA Langley Research Center
mark.h.carpenter@nasa.gov

MS363

An Algorithmic Search for Optimally Efficient High-order Summation-by-parts Operators

This presentation will investigate the potential for developing Summation-by-Parts (SBP) operators that are optimal with respect to computational efficiency. The generalization of the SBP definition to arbitrary nodal distributions and element-type implementation has connected various high-order methods, including nodal discontinuous Galerkin (DG) methods and flux reconstruction methods. A framework for understanding the benefits and trade-offs between these different methods has not been developed; however, the flexibility of generalized SBP operators enables the possibility for developing high-order discretizations of hyperbolic systems which are optimized with respect to specific performance properties. The properties of various operators with the SBP property, including some nodal DG operators, will be investigated in order to elucidate critical operator parameters which most affect computational efficiency. Within this context, an algorithm for optimizing SBP operators through novel objective functions will be outlined, and, in addition, new SBP operators which have been optimized with respect to specific properties identified as critical for efficiency will be presented. A comparison to other operators with the SBP property through various test cases will be performed in order to characterize the performance of existing and novel SBP operators.

Kevin J. Mattalo, Pieter Boom, David W. Zingg
University of Toronto Institute for Aerospace Studies
kevin.mattalo@mail.utoronto.ca,
pieter.boom@mail.utoronto.ca,
dwz@oddjob.utias.utoronto.edu

MS363

Stable and Accurate Filtering Procedures

Errors are inevitably present in numerical simulations, even when the computations are well resolved, and are predominantly introduced at high wave numbers. A plethora of finite difference stencils have been presented with reduced dispersion error for various ranges of wave numbers within the SBP-SAT framework. However, yet, no difference stencil is accurate in the vicinity of $\xi = \pi$. Instead, artificial dissipation operators of filters designed to remove high wavenumbers from the computational domain may be applied. Stable and accurate artificial dissipation operators for SBP-SAT based schemes was developed already in [K. Mattsson, M. Svärd, J. Nordström, Stable and accurate artificial dissipation, J. Sci. Comput. 21(1) (2004) 5779]. For filtering procedures, it was shown recently in [J. Nordström, V.Linders, Journal of Computational Physics, 364 (2018) 95110] the stability problems may occur. More precisely, successful filtering may include a delicate balance between the need to remove high frequency oscillations (filter often) and the need to avoid possible growth (filter seldom). In this presentation we will discuss this filtering problem and propose a remedy.

Jan Nordström
Department of Mathematics
Linkoping University
jan.nordstrom@liu.se

Tomas Lundquist
University of Cape Town

tomas.lundquist@uct.ac.za

MS363

Flexural Wave Propagation in Ice Shelfs

Flexural gravity waves in ice sheets can arise when very long wave length ocean waves interact with large floating ice sheets in the antarctic. Simulating such waves can be of interest because they allow us to gain a insight into the quality of the ice and how a changing climate might affect the ice sheet. The waves propagate through interaction between stiff ice and the water beneath the ice. One example of a model for such a system is

$$\varepsilon\phi_{tt} = \nabla^2\phi, \quad \bar{x} \in \Omega_{\text{water}},$$

$$aw_{tt} + \nabla^2 b\nabla^2 w + w + \gamma\phi_t = 0, \quad \bar{x} \in \Gamma_{\text{ice}},$$

$$\hat{n}\cdot\nabla\phi = w_t, \quad \bar{x} \in \Gamma_{\text{ice}},$$

where ϕ is a velocity potential of the water, w is the height of the ice and ε, a, b and γ are dimensionality parameters. Due to changing requirements and continuing exploration of models for the ice-water system, we require the numerical methods we derive for this problem to be highly robust and flexible. Being a wave propagation problem we also demand high order of accuracy. Using the summation by parts framework for both space and time dimensions we are able to derive higher order schemes that meet these requirements, including robust coupling of the ice and water equations using simultaneous approximation terms. In this talk we will present methods for deriving such schemes. We will also verify the robustness and accuracy as well as show example solutions.

Jonatan Werpers
Department of Information Technology
Uppsala University
jonatan.werpers@it.uu.se

Ken Mattsson
Uppsala University, Sweden
ken.mattsson@it.uu.se

MS364

Benchmarking and Modeling Sparse Memory Accesses for Heterogeneous Systems

Abstract not available.

Patrick Lavin
Computational Science & Engineering
Georgia Institute of Technology
plavin3@gatech.edu

MS364

Novel Architectures for Applications in Data Science and Beyond

Abstract not available.

Jason Riedy
Georgia Institute of Technology
School of Computational Science and Engineering
jason.riedy@cc.gatech.edu

Jeffrey Young, Tom Conte
Georgia Institute of Technology
jyoung9@gatech.edu, conte@cc.gatech.edu

MS364

A Software Framework for Spiking Neuromorphic Computing Systems

Abstract not available.

Catherine Schuman
Oak Ridge National Laboratory
schumancd@ornl.gov

MS364

Scientific Computer Architecture Beyond Moores Law

Abstract not available.

John Shalf, David Donofrio
Lawrence Berkeley National Laboratory
jshalf@lbl.gov, ddonofrio@lbl.gov

MS365

An Adaptive Stochastic Discontinuous Galerkin Method for Reliable Uncertainty Propagation

We present a reliable and efficient computational framework for combined model uncertainty and discretization error quantification of aerodynamics simulations. Specifically, we present an adaptive solution method for stochastic partial differential equations that (i) propagates the effect of prescribed model parameter uncertainties to the quantity of interest and (ii) effectively controls the discretization errors associated with the propagation process. Our framework builds on a high-order discontinuous Galerkin (DG) method, element-wise localized polynomial chaos expansions (PCEs), the dual-weighted residual (DWR) error estimate, and a spatio-stochastic anisotropic adaptation strategy. The DG formulation provides the flexibility to use approximation spaces with elementally varying PCE orders to exploit localized spatio-stochastic solution structures. The DWR method provides an effective discretization error estimate for the quantity of interest and element-wise localized error indicators to guide adaptive refinement. An anisotropic error sampling strategy identifies the optimal anisotropic refinement in physical and/or stochastic spaces to minimize the discretization error for a given computational cost. We present a priori error bounds for our spatio-stochastic approximation and then demonstrate the effectiveness of the formulation for compressible flows with uncertainties in flow conditions.

Geoff Donoghue
Institute for Aerospace Studies
University of Toronto
geoff.donoghue@mail.utoronto.ca

Masayuki Yano
University of Toronto

myano@utias.utoronto.ca

MS365

Efficient Matrix-free Solvers for the Incompressible Navier-Stokes Equations

We present efficient high-order discontinuous Galerkin (DG) solvers for the incompressible Navier-Stokes equations. A distinctive feature of this DG discretization is the use of consistent divergence and continuity penalty terms that render the discretization robust for turbulent flows. For efficient time integration, the approach is based on well-known projection methods. The resulting (linear) systems of equations are solved by state-of-the-art iterative solution techniques with robust multigrid preconditioners achieving optimal computational complexity. A combined geometric/polynomial/algebraic multigrid approach is presented as an efficient solution algorithm for problems on complex geometries. The main algorithmic component determining the computational efficiency is the evaluation of discretized finite element operators. For high-order methods to be competitive in under-resolved situations, an efficient matrix-free implementation exploiting sum-factorization techniques on quadrilateral/hexahedral elements (as opposed to memory-heavy matrix-based methods) is inevitable. Our marix-free implementation is highly optimized for modern multicore computer architectures with Flop/Byte ratio significantly larger than one. We demonstrate excellent performance characteristics of our implementation with a throughput for matrix-free operator evaluation measured in degrees of freedom processed per second that is almost independent of the polynomial degree.

Niklas Fehn, Wolfgang A. Wall
Technical University of Munich
fehn@lnm.mw.tum.de, wall@lnm.mw.tum.de

Martin Kronbichler
Technical University of Munich
Institute for Computational Mechanics
kronbichler@lnm.mw.tum.de

MS365

An Optimization-based Discontinuous Galerkin Approach for High-order Accurate Shock Tracking

We introduce a high-order accurate, nonlinearly stable numerical framework for solving steady conservation laws with discontinuous solution features such as shock waves. The method falls into the category of a shock tracking or r-adaptive method and is based on the observation that numerical discretizations that support discontinuities along element faces, e.g., discontinuous Galerkin methods, can perfectly represent discontinuities and provide appropriate stabilization through numerical fluxes. The difficulty lies in aligning element faces with the unknown discontinuity. The proposed method recasts a discretized conservation law as a PDE-constrained optimization problem whose solution is a (curved) mesh that tracks the discontinuity and the solution of the discrete conservation law on this mesh. The objective function is a discontinuity indicator that monotonically approaches a minimum as element faces approach the shock surface and the discretized conservation law on a parametrized domain defines equality constraints. A full space optimization solver is used to simultaneously converge the state vector and mesh to their optimal values to ensure the solution of the discrete PDE is never required on meshes that are not aligned with discontinuities, which provides nonlinear stability. Using a number of canonical problems, we demonstrate framework tracks discontinuities closely with curved mesh elements and provides accurate solutions on extremely coarse meshes.

Marzieh Mirhose
Department of Aerospace and Mechanical Engineering
University of Notre Dame
malireza@nd.edu

Matthew J. Zahr
Lawrence Berkeley National Laboratory
University of California, Berkeley
mjzahr@lbl.gov

Per-Olof Persson
Dept. of Mathematics
University of California, Berkeley
persson@berkeley.edu

MS365

Efficient Methods for Solving Implicit Space-time Multi-physics Systems

For the past several years an unstructured high-order finite-element capability has been developed at NASA Ames Research Center for the simulation of complex fluid problems involving flow separation, shock-boundary layer interaaction, transition, etc. The compressible fluid solver uses a fully implicit space-time entropy-stable discontinuous-Galerkin formulation. This talk presents recent improvements to the family of matrix-free tensor-product preconditioners developed for solving this system. Further, recent work extending the finite-element framework to a general monolithich multi-physics solver for FSI, chemistry, combustion, etc. is presented. Results and discussion of the performance implications of solving a high-order implicit multi-physics problem are presented.

Scott Murman, Laslo Diosady
NASA Ames Research Center
scott.murman@nasa.gov, laslo.diosady@nasa.gov

MS366

Exploiting Node-level Performance in Sparse Linear Algebra

Over the last years, we have observed not only an explosion in hardware parallelism, but also a growing mismatch between the arithmetic performance of processors in terms of the number of floating point operations per second (FLOPS) on the one side, and the memory performance in terms of how fast data can be brought into the computational elements (memory bandwidth) on the other side. As a result, more and more applications can utilize only a fraction of the available compute power as they are waiting for the required data. With memory operations

being the primary energy consumer, data access is pivotal also in the resource balance and the battery life of mobile devices. In this talk I will introduce a disruptive paradigm change with respect to how scientific data is stored and processed in computing applications.

Hartwig Anzt
University of Tennessee, USA
hanzt@icl.utk.edu

MS366

Optimizing Single Node Performance of Parallel Trajectory Reconstruction

Parallelizing software to execute on multi-core central processing units (CPUs) and graphics processing units (GPUs) can prove to be very challenging, and for some fields outside of Computer Science, this transition comes with new issues. For example, memory limitations can require modifications to code not initially developed to run on GPUs. The parallel programming techniques implemented on a Monte Carlo based approach on trajectory reconstruction in this work include: directive based approaches OpenMP and OpenACC for multi-core CPUs and GPUs, respectively, and CUDA for NVIDIA GPUs. Large matrix operations are the most common use of GPUs, which are not present in this algorithm; however, the natural parallelism of independent trajectories in Monte Carlo simulations is exploited. Benchmarking data are presented comparing execution times of the software for single-thread CPUs, multi-thread CPUs with OpenMP, and multi-thread GPUs using OpenACC and CUDA. These data were collected using Intel Sandy Bridge CPU and NVIDIA Tesla K40 GPU hardware on the Pleiades Supercomputer cluster at the National Aeronautics and Space Administration (NASA) Ames Research Center (ARC), local Intel Broadwell and Xeon Phi 7210 CPU nodes, and local NVIDIA Tesla P100 GPUs.

Anthony Williams, Justin Green
NASA Langley Research Center
robert.a.williams@nasa.gov, justin.green@nasa.gov

MS366

Hybrid Parallelization of the Stabilized Finite Element Library in FUN3D

Compute clusters are evolving away from compositions of many nodes that contain a few processors with a small amount of shared memory towards compositions of fewer nodes that contain many processors and a large amount of shared memory. This transition in computational architecture necessitates the development hybrid parallelization implementations of scientific and engineering software to of efficiently and flexibly utilize distributed and shared resources (processors and memory). The Streamlined Upwind Petrov-Galerkin (SUPG) stabilized finite-element (SFE) discretization library within the FUN3D unstructured-grid flow solver utilizes MPI to coordinate operations across distributed resources and OpenMP to coordinate operations using shared resources. The hybrid parallelization of SFE was improved by examining shared memory operations required for two bottle necks, namely the assembly of the coefficient matrix and the linear solves. Algorithmic choices and implementation trade-offs uncovered by profiling refactorings of the coefficient matrix assembly and linear solver procedures will be discussed. The impacts on computational cost (memory, CPU hours, and wall time) for production class simulations and implications for scaling to leadership class simulations will be discussed.

Stephen L. Wood
NASA Langley Research Center
Aerosciences Branch, High Performance Computing Incubator
stephen.l.wood@nasa.gov

William Anderson
NASA Langley Research Center
Aerosciences Branch and High-Performance Computing Incubator
william.k.anderson@nasa.gov

MS367

Auto-tuning of Preconditioners with Deep Learning

Preconditioner selection for sparse iterative methods is one of critical tuning processes to solve linear equations in numerical simulation of scientific and engineering problems. In this research, a novel auto-tuning (AT) method with deep learning (DL) for the selection of preconditioners are proposed. We use color feature images by input sparse matrices for learning dataset of DL to predict the best preconditioners for a sparse iterative library (GMRES solver with restarting). According to performance evaluation, proposal method establishes more than 80% accuracy with learning dataset by using almost all usable matrices from the Florida sparse matrix collection.

Takahiro Katagiri
Information Technology Center
Nagoya University
katagiri@cc.nagoya-u.ac.jp

Kenya Yamada
Graduate School of Informatics, Nagoya Univsersity
yamada@hpc.itc.nagoya-u.ac.jp

MS367

The Big Data Approach to Autotuning of GPU Kernels Using the BONSAI Toolset

High-end computing relies increasingly on machines with large numbers of GPUs. At the same time, GPU kernels are commonly produced in the process of automated software tuning. The premise of the Benchtesting OpeN Software Autotuning Infrastructure (BONSAI) project is to produce a software infrastructure for deploying large performance tuning sweeps across many GPUs. Currently, BONSAI provides two components to facilitate this objective: (1) the LANguage for Autotuning Infrastructure (LANAI) to generate and prune a large search space and (2) the Distributed ENvironment for Autotuning using Large Infrastructure (DENALI) to launch large tuning sweeps on a massive number of GPUs. This presentation contains a

brief introduction of LANAI, which is an older BONSAI component, and a more detailed discussion of DENALI, which was implemented more recently.

Jakub Kurzak
University of Tennessee, Knoxville
kurzak@icl.utk.edu

Piotr Luszczek
University of Tennessee
luszczek@icl.utk.edu

Mark Gates
Computer Science Department
University of Tennessee
mgates3@icl.utk.edu

Yaohung Tsai, Ahmad Abdelfattah
University of Tennessee, Knoxville
ytsai2@vols.utk.edu, ahmad@icl.utk.edu

Jack Dongarra
Innovative Computing Laboratory
University of Tennessee
dongarra@cs.utk.edu

MS367

Model-driven Auto-tuning for Optimizing Tensor Contractions

Tensor contractions are higher dimensional analogs of matrix-matrix multiplication. In contrast to the availability of highly tuned libraries on all platforms for matrix-matrix multiplication, the number of possible shapes and sizes of operand tensors for a tensor contraction makes it challenging to create an efficient static library that can efficiently perform arbitrary tensor contractions. This talk will describe experiences with a domain-specific compiler for synthesizing efficient tensor contraction code, where auto-tuning in used in conjunction with model-driven pruning of the vast configuration space.

Ponnuswamy Sadayappan, Jinsung Kim, Rui Li, Aravind Sukumaran-Rajam
Ohio State University
psaday@gmail.com, kim.4232@osu.edu, li.6263@osu.edu, sukumaranrajam.1@osu.edu

Sriram Krishnamoorthy
Pacific Northwest National Laboratory
sriram@pnnl.gov

MS367

Use of Machine Learning Technologies for Performance Engineering

The complexity of HPC system architectures are increasing. Even today, expert programmers' enormous effort is required to exploit the system performance, it is intractable to manually optimize every application code. Recently, machine learning is attracting attention because it can replace some of important tasks that have been done relying on experts' knowledge and experiences. In this talk, I will introduce our approaches to effective use of machine learning technologies to help performance engineering by mapping a performance engineering task to another task, on which machine learning models are expected to successfully work.

Hiroyuki Takizawa
Tohoku University
takizawa@tohoku.ac.jp

MS368

Meshless Methods for Manifolds: Hydrodynamics of Curved Fluid Interfaces and Related Applications

We discuss recent advances in the development of meshless methods for solving partial differential equations on general manifolds. We present a discretization framework based on Generalized Moving Least Squares (GMLS) and exterior calculus formulation of the pdes.Motivated by applications arising in soft condensed matter physics and biophysics, we show how our approaches can be used to solve hydrodynamic equations on curved fluid interfaces. We also present convergence results comparing our GMLS methods with a spectral solver in the case of radial manifolds. We then show some advantages of our GMLS methods demonstrating the capability to handle quite general manifold topologies and to adapt numerical resolution. We conclude by presenting results for hydrodynamic interactions in drift-diffusion dynamics of particle inclusions within curved fluid interfaces showing some of the important roles played by topology and geometry.

Ben Gross, Christian Bueno
UC Santa Barbara
bengross@ucsb.edu, chrsbueno@ucsb.edu

Nathaniel Trask
Brown University
Division of Applied Mathematics
nat.trask@gmail.com or natrask@sandia.gov

Paul Atzberger
UC Santa Barbara
atzberg@ucsb.edu

MS368

Well-balanced Meshless Methods for the Shallow-water Equations

In this talk, I will present recent work on meshless discretizations for the nonlinear shallow-water equations and discuss their application to tsunami models. Meshless methods are attractive for modelling both coastal and river flooding, due to the unstructured nature of real-world topographic data, but standard meshless finite-difference approaches fail even for very simple test cases. This motivates the development of a mimetic scheme for the spatial derivative and averaging operators to realize well-balanced meshless discretizations. We couple this approach with a radial basis function based extrapolation method for the inundation model, giving an accurate and efficient simula-

tion scheme for the shallow-water equations.

Jörn Behrens
Dept. of Mathematics, University of Hamburg
joern.behrens@uni-hamburg.de

Alex Bihlo, Rüdiger Brecht, Scott Maclachlan
Department of Mathematics and Statistics
Memorial University of Newfoundland
abihlo@mun.ca, rbrecht@mun.ca, smaclachlan@mun.ca

MS368

The Fast Radial Basis Functions Orthogonal Gradients (RBF-OGr) Method for Solving PDEs on Arbitrary Surfaces

The RBF-OGr method was introduced in [Piret, 2012] to discretize differential operators defined on arbitrary manifolds defined exclusively by a point cloud. The method was designed to take advantage of the meshfree character of RBFs, which offers the flexibility to represent complex geometries in any spatial dimension while providing a high order of accuracy. A large limitation of the original RBF-OGr method was its large computational complexity, which greatly restricted the size of the point cloud. In this talk, a fast version of the RBF-OGr method will be introduced. This latest algorithm makes use of the RBF-Finite Difference (RBF-FD) technique for building sparse differentiation matrices discretizing continuous differential operators such as the Laplace-Beltrami or the surface biharmonic operators.

Cecile M. Piret
Michigan Technological University
cmpiret@mtu.edu

MS368

RBF-FD for PDEs on Manifolds: Stable and Stagnation-free Formulations

We present a new method for the solution of PDEs on 2D manifolds embedded in 3-dimensions using stable scale-free radial basis function (RBF) interpolation. Much like other recent RBF methods, our method involves augmenting polyharmonic spline (PHS) RBFs with polynomials to generate RBF-Finite Difference (RBF-FD) formulas. These polynomial basis functions are obtained using least orthogonal interpolation (LOI) on each RBF-FD stencil to obtain local restrictions of polynomials in 3-dimensional space to stencils on 2D manifolds. The resulting RBF-LOI method uses Cartesian coordinates, uses only nodes on the manifold, does not require tangent plane projections, and appears robust to stagnation errors. We show convergence rates for both advection and diffusion PDEs, and present some applications motivated by biology.

Varun Shankar
School of Computing
University of Utah
shankar@cs.utah.edu

Akil Narayan
University of Utah
akil@sci.utah.edu

Mike Kirby
University of Utah
School of Computing
kirby@cs.utah.edu

MS369

Coordinate-wise Descent Methods for the Full Configuration Interaction Calculation

The full configuration interaction (FCI) calculation computes the smallest eigenvalue of a sparse symmetric matrix. While the computational challenge of the FCI calculation lies in the fact that the matrix size grows exponentially with respect to the number of orbitals/electrons in the system. In this work, the coordinate-wise descent methods are considered for such problems based on a reformulation of the edge eigenvalue problem as a non-convex optimization problem. Numerical examples of physical systems demonstrate the efficiency and provide benchmarks of the FCI calculation.

Yingzhou Li
Duke University
yingzhouli0417@gmail.com

Jianfeng Lu
Mathematics Department
Duke University
jianfeng@math.duke.edu

Zhe Wang
Duke University
zhe.wang3@duke.edu

MS369

Bayesian Inference of Molecular and Atomic Properties: Representations and Regression Techniques

Unlike image and volumetric data, which are discretized representations of continuous signals, molecules and atoms could often be modeled more naturally as inherently discrete structures such as trees and graphs. In this presentation, I will examine some of the difficulties that arise when applying machine learning techniques to discrete structures like molecules, which typically possess variable dimensionalities and non-linear structures. I will then introduce a family of marginalized graph kernels for molecules, i.e. functions that compute pairwise similarity, and how they can be used to construct non-parametric Bayesian inference models to learn and predict atomictic properties. Demonstration will be given on how the method could predict the atomization energy with up to 1 kcal/mol accuracy and significantly less training effort.

Yu-Hang Tang
Lawrence Berkeley National Laboratory
tang@lbl.gov

MS369

Overview of Machine Learning Techniques

Machine learning has become an important tool for enabling and accelerating materials discovery and molecu-

lar science.It allows scientists to leverage existing scientific data and knowledge to make predictions about new properties of materials and chemical systems or characterize the relationship between miscroscopic features and experimental observables. In this talk, I will give an overview of molecular and electronic structure problems being solved by machine learning techniques, models and descriptors that are suitable for existing machine learning tools, and challenges and opportunities for using machine learning to tackle more difficult problems.

Chao Yang
Lawrence Berkeley National Lab
cyang@lbl.gov

MS369

Deep Learning for Multi-scale Molecular Modelling

We introduce a series of deep learning based methods for molecular modeling at different scales. We see model construction and data exploration as two important aspects of this task. In terms of model construction, we introduce the Deep Potential scheme based on a many-body potential and inter-atomic forces generated by a carefully crafted deep neural network trained with *ab initio* data. We show that the proposed scheme provides an efficient and accurate protocol for a variety of systems, including bulk materials and molecules, and, in particular, for some challenging systems like a high-entropy alloy system. We further show how this scheme is generalized to the context of coarse-graining. In terms of data exploration, we present a new scheme called reinforced dynamics for enhanced sampling and efficient learning, using ideas from reinforcement learning and deep learning. Applications are discussed for both macro-molecules and condensed systems.

Linfeng Zhang
Princeton University
linfengz@princeton.edu

Weinan E
Princeton Univesity
Department of Mathematics
weinan@math.princeton.edu

Roberto Car
Princeton University
rcar@princeton.edu

MS370

Scaling and Deepening Tensor Decompositions and Applications

Tensor decompositions are a powerful data analysis technique for discovering multiway patterns in complex real-world data and are increasingly being used in multiple application domains. We present ENSIGN as a ready-to-use solution for data scientists seeking to apply large-scale tensor analysis to practical applications in high-performance computing and cloud environments. ENSIGN provides optimized implementations of a variety of static and streaming tensor decomposition methods, some of which are uniquely supported in ENSIGN. We have architected ENSIGN to scale on current computing platforms and also adapt to new platforms with heterogeneous processors. Scalable performance results from optimized sparse tensor data structures, memory-efficient computations, and novel techniques for improving parallelism, data locality, load balance, and asynchrony in the execution of tensor computations. We have successfully demonstrated the application of ENSIGN in a variety of domains, including, cybersecurity, geospatial analysis, and bioinformatics, and the need for supporting an ensemble of methods (with varying properties related to model, sparsity, and non-negativity) to use across applications. We continue to expand ENSIGN with new methods for increasing the breadth and depth of tensor analysis applications. ENSIGN also provides a Python API for tensor decomposition methods to enable seamless integration with mainstream data science libraries such as scikit-learn, SciPy, and more.

Muthu M. Baskaran
Reservoir Labs
baskaran@reservoir.com

MS370

A Multi-dimensional Morton-ordered Block Storage for Mode-oblivious Tensor Omputations

We propose a blocked tensor storage where the order of the blocks are determined by the Morton space-filling curve and apply it to the problem of dense tensor-vector multiplication (TVM). Due to blocking and the cache-oblivious memory accesses on the input vector and the output tensor induced by the Morton curve, we achieve high performance that is furthermore unaffected by the mode in which the vector is applied. A naive algorithm is up to 18% slower than when using the new data structure. This is on-par and often even improves upon an unfold + optimised BLAS2 algorithm, with peaks of up to 18% performance gains for our new data structure. Also, and perhaps more importantly, the use of the Morton curve reduces the sample standard deviation of performance between different TVM modes by up to 89% for blocked tensor-vector multiplication and by up to 71% versus non-blocked variants. These last numbers convincingly show the mode-oblivious nature of our new data structure.

Filip Pawlowski
Huawei
filip.pawlowski1@huawei.com

Bora Ucar
LIP, ENS-Lyon, France.
Bora.Ucar@ens-lyon.fr

Albert-Jan N. Yzelman
Huawei Technologies France
albertjan.yzelman@huawei.com

MS370

Performance Portable Parallel CP-APR Tensor Decompositions

Sparse tensor decomposition is a tool that allows analysts to explore a compact representation (low-rank mod-

els) of high-dimensional data sets, expose patterns that may not be apparent in the raw data, and extract useful information from the large amount of initial data. In this work, we consider decomposition of sparse count data using CANDECOMP-PARAFAC Alternating Poisson Regression (CP-APR) for HPC systems. Unlike the Alternating Least Square (ALS) version, CP-APR algorithm involves non-trivial constraint optimization of nonlinear and nonconvex function, which contributes to the slow adaptation to high performance computing (HPC) systems. The recent studies by Kolda et al. suggest multiple variants of CP-APR algorithms amenable to data and task parallelism together, but their parallel implementation involves several challenges due to the continuing trend toward a wide variety HPC system architecture and its programming models. We address this problem using Kokkos, Sandia's parallel programming framework, to design a performance portable parallel CP-APR algorithms. In the presentation, we will discuss our implementation of CP-APR algorithms and their performance on multiple computing platforms.

Keita Teranishi
Sandia National Laboratories
knteran@sandia.gov

MS370

Talk Title Not Available

Abstract not available.

Rich Vuduc
Georgia Institute of Technology
richie@cc.gatech.edu

MS371

Acknowledging Scientific Software to Ensure the Future and Legacy of Scientific Research

Access to research software is foundationally important to both the future and the heritage of scientific research. Deep intellectual contributions are being made by people building software that is increasingly unable to be decoupled from data as new research methods trade off exact numeric determinism for model performance. But digital research artifacts like code present new challenges to traditional scholarly communication models and digital preservation practices; versioning and authorship in these contexts are fluid. It is therefore all the more crucial that practices are adopted to enable people to share these complex, distributed, dynamic tools as easily as they share articles. It is equally important that researchers are encouraged to create these valuable resources and that their work is acknowledged as essential. A brief overview of emerging best practices in this context will be presented along with concrete actions people can take to make their code more open, citable, and persistent to ensure the legacy of their work.

Daina Bouquin
Harvard Center for Astrophysics
daina.bouquin@cfa.harvard.edu

MS371

Supporting and Sustaining Open Source Software Development: the Commons Perspective

While software is only slowly becoming explicitly valued as a product and enabler of research, open source software projects are absolutely critical to modern science. In recent years, the mismatch between the academic weight placed on software versus the *actual* value of software has become glaringly obviously. In part because of this, many important open source projects are undergoing a sustainability crisis. The NIH Data Commons, an effort to create a sustainable ecosystem of platforms and services to enable biomedical data analysis, is confronting exactly this question: how do we build, maintain, and evolve critical open source software infrastructure? I will present this as a problem in sociotechnical systems studies, reference a panoply of intriguingly relevant literature, and reframe the underlying sustainability question as one of community: that is, *sustainability* depends critically on the formation of a *community*, and moreover that this community should be initiated, grown, and evolved according to the design principles for sustainable common pool resource management. This perspective yields some interesting points of consideration for community engagement and governance, while also highlighting a number of significant conceptual mismatches with academic approaches to sustainability.

C. Titus Brown
University of California Davis
ctbrown@ucdavis.edu

MS371

Contemporary Peer Code Review in Research Software

Contemporary peer code review is a lightweight, asynchronous method for ensuring high-quality code. Research results in traditional software engineering and open-source software engineering have shown the clear benefits peer code review provides to software quality and maintainability. The quality increase results from the focused review of the code to identify areas in need of improvement. The improved maintainability arises from the fact that developers begin writing code in a more readable fashion to enable the peer-review process (a result we have seen in our own studies). By writing code that is more readable and easier to understand, developers also make that code more maintainable over time. While this practice has been shown to be beneficial to help developers identify and remove faults from code, it is underutilized in scientific software. In this talk, I will provide a brief overview of contemporary peer code review. Then I will report on results from our efforts at developing and delivering contemporary peer code review tutorials to scientific software audiences.

Jeffrey C. Carver
Department of Computer Science
University of Alabama
carver@cs.ua.edu

MS371

Supporting Continuous Integration at Large-scale HPC Centers

Continuous integration (CI) is a critical component of mod-

ern software development. CI systems run automated tests and checks on every commit to a project, and code is only merged if it passes these tests. CI automation allows developers to work faster, with confidence that the current version of their code can is tested and ready to run at any time. Despite these benefits, deploying CI at large, multi-user HPC sites has traditionally been difficult. Most CI systems are designed for small development teams, where data sharing is allowed and expected. Security resrictions at many HPC centers require that CI jobs run by one user *not share data with other users' jobs. Batch support in production CI systems is currently weak, and strong security measures like two-factor authentication prevent externaly hosted CI tools from launching jobs at the HPC center automatically. Under the U.S. Exascale Computing Project (ECP), we are working with Onyx Point to enhance the security of the widely used GitLab CI system for use in DOE. We are using setuid runners to enforce isolation among jobs, and we are building support for batch systems. This presentation will describe the security challenges we faced and the design decisions we adopted to mitigate them. This work will, for the first time, allow scientific software developers to test their code on HPC machines, just as they would in the cloud.*

Todd Gamblin
Lawrence Livermore National Laboratory
gamblin2@llnl.gov

Rob Neely
Lawrence Livermore National Lab
neely4@llnl.gov

David Montoya
LANL
dmont@lanl.gov

MS372

Computational Tools for Periodic and Non-periodic Nanophotonic Devices Design

This talk presents numerical solvers for the design of periodic and non-periodic, electromagnetically large, metamaterial devices. Integral equations methods will be discussed in the context of periodic problems and a novel general strategy to overcome Wood anomalies will be presented. Numerical results will be demonstrated showing the high-order convergence throughout the spectrum. Additionally, a scalable FFT-based methodology as well as optimization techniques will be discussed to produce fast and efficient computational tools to simulate and design large, non-periodic, nanophotonic devices.

Agustin Fernandez Lado
Caltech
agfernan@caltech.edu

Oscar P. Bruno
California Institute of Technology
obruno@caltech.edu

MS372

Recent Advances in High-order Simulations for Electromagnetics

This talk will present high-performance high-order simulation techniques in solving the Maxwell's equations for analyzing optical systems that are required to perform massive parameter studies for phase analysis to predict optimal structure. Considering a new class of high-index metamaterials for ultrathin flat lenses and graphene-based two-dimensional materials, we will present on the challenges required to accurately predict and analyze the optical properties of these new materials in order to reduce development costs and risks arising from complicated fabrication processes and experimental limitations. we will discuss robust and rapidly convergent numerical methods based on spectral element discontinuous Galerkin approaches and demonstrate simulations of nanoresonators in a disk array with phase analysis.

Mi Sun Min, Yu-Hsiang Lan
Argonne National Laboratory
mmin@mcs.anl.gov, seanlan1994@gmail.com

MS372

Discontinuous Galerkin Methods for Electronics Calculations in Solar Cells

Predicting the efficiency of a thin film solar cell requires two calculations. First it is necessary to solve a photonics problem to predict the electric field due to the suns radiation throughout the device. Then, using this field, the current produced by the device can found by solving an appropriate electronics model. Here we use the drift diffusion model and the Hybridizable Discontinuous Galerkin scheme to discretize the electronics problem. Our first attempt produced a high order and robust solver, but we did not obtain optimal convergence rates. By taking different order polynomial spaces for the electric potential and concentration of holes and electrons, we are now able to establish optimal convergence rates. This is joint work with Gang Chen and Yangwen Zhang.

Peter B. Monk
Department of Mathematical Sciences
University of Delaware
monk@math.udel.edu

MS372

A Non-overlapping Domain Decomposition Method for Simulating Localized Surface Plasmon Resonances: High Accuracy Numerical Simulation and Analyticity of Solutions

It is important to engineers and scientists alike to simulate scattering returns of electromagnetic radiation from bounded obstacles. In addition, such simulations must be of both surpassing accuracy and high fidelity for many applications of interest. In this talk we present a non-overlapping Domain Decomposition Method for the simulation of such configurations, implemented with Impedance-Impedance Operators computed via a High-Order Perturbation of Surfaces algorithm. With an implementation of this approach we demonstrate the stable, robust, and highly accurate properties of our method. Time permitting, we also demonstrate how our formulation delivers a

straightforward proof of existence, uniqueness, and analyticity of solutions to this problem.

Xin Tong
University of Illinois at Chicago
xtong20@uic.edu

MS373

Radiation Hydrodynamics in GenASiS

GenASiS (General Astrophysical Simulation System) is a code under development initially and primarily for the core-collapse supernova problem, and as such neutrino radiation transport is a key part of the physics being developed within it. Classes representing phase space (position space plus momentum space) and particle distributions thereon are key pieces of infrastructure, and are being developed in a way that will be extensible from energy-dependent angular moment formalisms in the near term to full Boltzmann transport in the long term, including the potential for adaptive mesh refinement in both position space and momentum space. The spatial streaming operator is handled in conservative form in a time-explicit manner similar to fluid dynamics, while the collision operator is handled implicitly; these substeps are combined in an implicit-explicit Runge-Kutta scheme. Both finite-volume and discontinuous Galerkin versions are contemplated, as are gradual implementation of relativistic effects and increasingly sophisticated interactions with matter.

Christian Cardall
Oak Ridge National Laboratory
cardallcy@ornl.gov

MS373

Hybrid Defect Correction Methods for Time-dependent Radiation Transport Simulations

We present an arbitrarily high order solver for time-dependent radiation transport simulations that combines a hybrid splitting approach with a low-storage integral deferred correction (IDC) time-integration scheme. The hybrid splitting is based on a decomposition of the radiative flux into collided and uncollided components to which low- and high-resolution discrete ordinates approximations are applied, respectively. A novel method for mapping between arbitrary discrete ordinates quadratures is presented that is straightforward to implement and significantly more accurate than previous reconstruction approaches. An IDC time-integration scheme is used to attain high-order accuracy in time while simultaneously reducing the splitting error due to the hybrid approach. A test problem consisting of an inhomogeneous sphere radiating into a vacuum is used to demonstrate the effectiveness of the methods.

Michael Crockatt
Michigan State University
crockat1@msu.edu

Andrew J. Christlieb
Michigan State Univerity
Dept. of Comp. Math., Sci & Engr.
christli@msu.edu

Cory Hauck
Oak Ridge National Laboratory
hauckc@ornl.gov

MS373

Explosion Mechanisms in Multi-dimensional Stellar Core Collapse with Magnetic Fields

The collapse of the cores of massive stars allows for several distinct evolutionary paths such as supernova explosions driven by the energy deposition due to neutrinos behind the stalled shock wave, explosions driven by magneto-rotational stresses, as well as a failure of the shock revival. In the last case, the core will produce a black hole, while the compact remnant in the other cases may, depending on the dynamics of the explosion and the potentially ongoing accretion, be a neutron star or a black hole. Which of the evolutionary paths is taken depends on the interplay of various processes involving hydrodynamic flows, magnetic fields, neutrino transport, and the reactions between neutrinos and matter. Their relative importance is set by the conditions in the pre-collapse star. Using a two-moment scheme for spectral neutrino transport, we performed a series of multi-dimensional simulations of the collapse of stars of up to 35 solar masses, varying the rotational and magnetic energies. Among them, we find all of the aforementioned possible outcomes, which we attempt to describe in terms of important dynamical parameters such as the efficiency of neutrino heating and the strength of the magnetic field.

Martin Obergaulinger
University of Valencia
martin.obergaulinger@uv.es

MS373

FORNAX: A Flexible Code for Multiphysics Astrophysical Simulations

We describe the design and implementation of our new multi-group, multi-dimensional radiation hydrodynamics (RHD) code FORNAX, review code tests selected to validate its application in a wide range of physical regimes, and present the results of a recent successful 3D core-collapse supernova (CCSN) explosion model. Our code solves the comoving-frame radiation moment equations using the M1 algebraic closure, utilizes conservative high-order reconstruction, employs semi-explicit matter and multi-group radiation transport via a high-order time stepping scheme, and is suitable for application to a broad range of astrophysical problems in addition to CCSNe. We describe the philosophy, algorithms, and methodologies of FORNAX, and show the results of relevant code tests that demonstrate the numerical fidelity with which it captures the many physical effects of RHD and show excellent strong scaling well above 100k MPI tasks. Finally, we present results of a recent application of FORNAX to the simulation of a 16-$M_\odot$ CCSN progenitor that incorporates inelastic scattering, a many-body interaction correction to the neutrino-nucleon scattering rates, velocity perturbations in the progenitor core, and an approximate treatment of general relativity, including the effects of gravitational

redshifts.

Aaron Skinner
Lawrence Livermore National Laboratory
skinner15@llnl.gov

MS374

Splitting Up Method for Backward Doubly SDEs and Applications to Nonlinear Filtering Problems

Backward doubly SDEs (BDSDEs) are equivalent to a class of Zakai type stochastic partial differential equations. They also provide solutions for nonlinear filtering problems. In this talk we consider the splitting up method to solve BDSDEs numerically. In the splitting up scheme, a BSDE is first decomposed into a coupled BSDE and a SDE. Then standard numerical methods are applied to solve the BSDE and the SDE. It is proved the method is first order and feasible for scalable computing.

Yanzhao Cao
Department of Mathematics & Statistics
Auburn University
yzc0009@auburn.edu

Feng Bao
Florida State University
bao@math.fsu.edu

He Zhang
Auburn University
hzz0077@auburn.edu

MS374

A Multilevel Reduced-basis Method for Linear Parameterized PDEs

Surrogate modeling is the most common approach for reducing the computational burden of performing statistical analysis of complex models. We consider parametric partial differential equations (PDE) with operators that have affine dependence on the input parameters, in which case, the reduced basis (RB) method offers the theoretically fastest convergence rate in terms of accuracy per number of samples (model evaluations). However, theory assumes that the samples are drawn from an infinite dimensional space, while in practice, a finite element discretization is use with a fixed dense mesh. We present a method that incorporates multiple meshes with different density, which allow for further reduction of total computational cost in a manner similar to Multi-Level Monte Carlo (MLMC) and Multi-Level Stochastic Collocation (MLSC) methods. Unlike MLMC and MLSC methods, RB works with the PDE operators in addition to the model outputs which presents a set of unique challenges. Specifically, we derive a different PDE for each mesh density where we have to give considerationt o boundary conditions. Furthermor, MLMC and MLSC rely on sharp apriori error estimates to compute optimal distribution of the computation work, which is not available in our context. Thus, we employ a greedy seach similar to the knapsack multi-index problem. We present several examples where we demonstrate 60%-90% reduction in cost of the multi-level RB approach compared to the standard RB method.

Miroslav Stoyanov
Oak Ridge National lab
stoyanovmk@ornl.gov

MS374

Learning Deep Neural Network Surrogate Models for High Dimensional and Multi-fidelity Uncertainty Quantification Problems

State-of-the-art computer codes for simulating real physical phenomena are often characterized by vast numbers of input parameters. Often, these input parameters are uncertain, and one needs to rigorously assess the effect of these uncertainties on the outputs from the computer codes. Performing uncertainty quantification (UQ) tasks with Monte Carlo (MC) methods is almost always infeasible because of the need to perform hundreds of thousands or even millions of forward model evaluations in order to obtain convergent statistics. One, thus, tries to construct a cheap-to-evaluate surrogate model to replace the forward model solver. However, for systems with large numbers of input parameters, one has to deal with the curse of dimensionality - the exponential increase in the volume of the input space, as the number of parameters increases linearly. This necessitates the application of suitable dimensionality reduction techniques.

Rohit Tripathy, Ilias Bilionis
Purdue University
rtripath@purdue.edu, ebilionis@gmail.com

MS375

Towards the Distributed Burning Regime in Turbulent Premixed Flames

Abstract not available.

Andrew Aspden
Newcastle University
andrew.aspden@newcastle.ac.uk

MS375

An Efficient GPU Implementation for Time-implicit Integration of Arrhenius Combustion Kinetics

Chemistry processes in combustion applications are characterized by a broad range of timescales. Because of this inherent stiffness, inclusion of comprehensive chemistry in multi-physics CFD simulations is generally performed with stiff ODE solvers (such as CVODE) and often add-up to more than 50% of the total computational cost. Several strategies can be considered to minimize the cost of the kinetic evaluation. While the efficiency of strategies relying on reducing the size of the chemical mechanism (either a-priori or on-the-fly) is firmly established, considerable computational time savings can also be achieved by directly tackling the ODE evaluation and/or by taking advantage of GPU architectures. However, a clear quantification of the gain to be expected in simulations of realistic turbulent reactive flows, for example using fully stiff chemistry (e.g.

n-dodecane) has remained elusive. An attempt is proposed in this work, using the low Mach number combustion code PeleLM, where it is necessary to rely on an implicit solver for stability reasons. In a first step, the capabilities of several suitable state-of-the-art strategies for efficient kinetic evaluation designed for implementation on hybrid architectures is evaluated, on a 3-D turbulent flame. In a second step, we propose an alternative methodology based on the Spectral Deferred Correction, which in particular alleviates the start-up cost associated with traditional BDF methods.

Anne Felden, Marcus Day
Lawrence Berkeley National Laboratory
enna@lbl.gov, msday@lbl.gov

MS375

Direct Numerical Simulations of Compression Ignition Jet Flames

Abstract not available.

Evatt R. Hawkes
University of New South Wales
evatt.hawkes@unsw.edu.au

MS375

Direct Numerical Simulations of Buoyancy-driven Flows Using PeleLM

Computational simulations have the potential to provide new physical understanding of fire dynamics and evolution in both natural and built environments. However, such simulations are challenging due to the multi-physics, multi-scale nature of essentially all fires, as well as to the prevalence of compressible and buoyancy-driven flow effects. In this talk, we present new results from a computational effort focused on understanding and characterizing wildland fire spread at small scales (roughly 1m1mm) using direct numerical simulations (DNS). The cost of the simulations is reduced using adaptive mesh refinement (AMR) within PeleLM, where resolution is provided only when and where it is needed to resolve physically-relevant fine-scale features. Simulation results are shown for two canonical buoyancy-driven flows: helium plumes and pool fires. In the plume configuration, helium gas is released into ambient air from a 1m inlet. In the pool fire configuration, methane is released into ambient air from the same 1m inlet, before burning as a non-premixed diffusion flame. Comparisons are made between results from the PeleLM simulations with and without AMR, and between the simulations and available experimental data. Focus is placed, in particular, on the computational savings enabled by the use of AMR, in addition to simulation accuracy.

Nicholas Wimer
Colorado University, Boulder
niwi6248@colorado.edu

Marcus Day
Lawrence Berkeley National Laboratory
msday@lbl.gov

Amanda Makowiecki, Jeffrey Glusman, John Daily, Gregory Rieker, Peter Hamlington
University of Colorado, Boulder
amanda.mackowiecki@colorado.edu, jeffrey.glusman@colorado.edu, john.daily@colorado.edu, gregory.rieker@colorado.edu, peter.hamlington@colorado.edu

MS376

Data-driven Correction for Reduced Order Modeling of Nonlinear Systems

In this talk, we address the following question: Given a nonlinear equation $u' = f(u)$ and a basis of fixed dimension r, find the best Galerkin model of dimension r. We present the answer proposed by our group for reduced order models (ROMs), supporting numerical results, and open questions.
Specifically, we propose a data-driven correction ROM (DDC-ROM) framework, which can be formally written as DDC-ROM = Galerkin-ROM + Correction. To minimize the new DDC-ROM's noise sensitivity, we use the maximum amount of classical projection-based modeling and resort to data-driven modeling only when we cannot use the projection-based approach anymore (i.e., for the Correction term). The resulting minimalistic data-driven ROM (i.e., the DDC-ROM) is more robust to noise than standard data-driven ROMs, since the latter employ an inverse problem (which is sensitive to noise) to model all the ROM operators, whereas the former solves the inverse problem only for the Correction term.
We test the novel DDC-ROM in the numerical simulation of a 2D channel flow past a circular cylinder at Reynolds numbers $Re = 100$, $Re = 500$, and $Re = 1000$.

Traian Iliescu, Muhammad Mohebujjaman
Department of Mathematics
Virginia Tech
iliescu@vt.edu, jaman@vt.edu

Leo Rebholz
Clemson University
Department of Mathematical Sciences
rebholz@clemson.edu

Birgul Koc, Changhong Mou
Virginia Tech
birgul@vt.edu, cmou@math.vt.edu

MS376

Machine-learning Error Models for Rom Closure Via the ROMES Method

Characterizing the epistemic uncertainty introduced by substituting a full-order model with a reduced-order model (ROM) is crucial for developing rapid yet reliable uncertainty quantification procedures. We present a technique for constructing a statistical model of the state (i.e., solution field) error introduced by a ROM for problems characterized by parameterized systems of algebraic equations. The technique extends the ROM error surrogates (ROMES) method by applying Gaussian-process regression to construct a statistical model that maps inexpensive error indicators to a distribution over the generalized

coordinates of the state error. This approach enables error estimation for any quantity of interest a posteriori, as the state-error model can be propagated through the associated quantity-of-interest functional. Numerical experiments performed on both linear and nonlinear parameterized PDE problems illustrate the ability of the proposed method to significantly improve ROM prediction accuracy and to construct accurate statistical error models for arbitrary quantities of interest a posteriori.

Stefano Pagani
École polytechnique fédérale de Lausanne
stefano.pagani@polimi.it

Andrea Manzoni
Politecnico di Milano
andrea1.manzoni@polimi.it

Kevin T. Carlberg
Sandia National Laboratories
ktcarlb@sandia.gov

MS376

Data-informed Reduced-order Models with Memory Effects

Unresolved dynamics in projection-based reduced order models can lead to a degradation of accuracy and stability. This work investigates the Mori-Zwanzig formalism as a tool for quantifying and modeling unresolved dynamics in projection-based reduced-order models. The Mori-Zwanzig formalism is a tool that allows for a high dimensional Markovian dynamical system (i.e. a full-order model) to be recast as a lower dimensional non-Markovian system (i.e. a reduced-order model). In this lower dimensional system, the impact of the unresolved 'fine-scale' dynamics on the resolved coarse-scales appears as a 'memory' integral. In this work, features arising from expansions of the MZ memory integral are used in conjunction with data-informed modeling techniques to develop reduced-order models of non-linear dynamical systems. Examples are presented with a focus on fluid dynamic problems.

Eric Parish
Univ of Michigan
parish@umich.edu

MS376

A Hybrid Approach for Model Order Reduction of Barotropic Quasi-geostrophic Turbulence

We put forth a robust reduced-order modeling approach for near real-time prediction of mesoscale flows. In our hybrid-modeling framework, we combine physics-based projection methods with neural network closures to account for truncated modes. We introduce a weighting parameter between the physics-based and extreme learning machine models to explore its effectiveness, accuracy and generalizability. To illustrate the success of the proposed modeling paradigm, we predict both mean flow pattern and time series response of a single-layer quasi-geostrophic ocean model, which is a simplified prototype for the wind-driven general circulation models. We demonstrate that our approach yields significant improvements over the standard Galerkin projection method with a negligible computational overhead.

Sk. Mashfiqur Rahman
Oklahoma State University
skmashfiqur.rahman@okstate.edu

Omer San
Mechanical and Aerospace Engineering
Oklahoma State University
osan@okstate.edu

Adil Rasheed
SINTEF Digital
adil.rasheed@sintef.no

MS377

Structured Bayesian Gaussian Process Latent Variable Model: Applications to Data-driven Dimensionality Reduction and High-dimensional Inversion

We introduce a Bayesian Gaussian process latent variable model (GPLVM) that explicitly captures spatial correlations in data using a parameterized spatial kernel and leveraging structure-exploiting algebra on the model covariance matrices for computational tractability. Modeling high-dimensional time series is enabled through use of a dynamical GP latent variable prior. A methodology is developed for Bayesian inversion using the SGPLVM where the unknown quantity is a spatial field. An SGPLVM is used both to construct a low-dimensional generative model of the sample-based stochastic prior as well as a surrogate for the forward evaluation. Its Bayesian formulation captures epistemic uncertainty introduced by the limited number of input and output examples, automatically selects an appropriate dimensionality for the learned latent representation of the data, and rigorously propagates the uncertainty of the data-driven dimensionality reduction of the stochastic space through the forward model surrogate. Importantly, the Bayesian inversion is carried out by solving a variational optimization problem, replacing traditional computationally-expensive Monte Carlo sampling. Experiments including data imputation on partially-observed images and time series demonstrate the advantages of the model. Bayesian inversion of an elliptic PDE is shown to return well-calibrated posteriors and is tractable with latent spaces with over 100 dimensions.

Steven Atkinson
University of Notre Dame
steven@atkinson.mn

MS377

Randomized Methods for Regression-type Problems

Classically, highly accurate deterministic algorithms are used for the task of solving inverse problems. However, the emergence of large-scale datasets across the social, physical, engineering, biological, and ecological sciences has severely challenged our computational ability to analyze data. Over the last decade, the concept of randomness

has been demonstrated as an effective strategy to quickly produce approximate answers to familiar problems such as dimension reduction and regression. In this talk we discuss randomized methods for solving regularized inverse problems which appear in ridge, sparse and non-linear regression-type problems. We discuss different sketching techniques and give both an optimization and statistical perspective.

N. Benjamin Erichson
University of California, Berkeley
erichson@berkeley.edu

Michael Mahoney
UC Berkeley
mmahoney@stat.berkeley.edu

MS377

Pdenetwork - The PDE's that Govern Deep Neural Networks

Abstract not available.

Eldad Haber
Department of Mathematics
The University of British Columbia
haber@eoas.ubc.ca

MS377

Generalized Graph Based Probit in the Continuum Limit

Graphical semi-supervised learning (SSL) has attracted a lot of attention recently. SSL is inherently geometric and the success of this approach is tied to our ability to extract coarse scale structures in the data. A common approach is to construct a weighted graph that summarizes the similarities between pairs of points in the dataset. In this talk we study the graph Laplacian operator that is defined on such similarity graphs and study its spectral properties. In particular, we study the limit where the number of vertices (i.e., the size of the dataset) becomes large. It is known that in this limit the graph Laplacian converges to a differential operator that shares similar properties with the underlying discrete graph Laplacian. This suggests the definition of a continuum analog of SSL that can be analyzed to gain insight into the discrete problem. We will discuss the consistency of the continuum SSL and identify the SSL classifier in binary classification.

Bamdad Hosseini, Franca Hoffmann
California Institute of Technology
bamdadh@caltech.edu, fkoh@caltech.edu

MS378

Selecting Multiple Borehole Locations for Maximizing Bayesian Information Gain on Groundwater Transmissivity

Data are essential for calibration, prediction, and uncertainty quantification, but acquiring data for geophysical applications is often time-consuming, expensive, and even dangerous. When choosing borehole drilling locations in subsurface aquifers, one should then aim to maximize the value of data while minimizing costs and risks. In this work, we present optimal experimental design (OED) for selecting multiple borehole locations to infer groundwater transmissivity, while working under a fixed total budget and with geographic constraints. In the multi-borehole design scenario, the rigorous OED framework is able to discover non-obvious location combinations induced by aquifer spatial correlation. Adopting a Bayesian perspective, the amount of information gain is quantified by the Kullback-Leibler (KL) divergence from the posterior to the prior. Numerical evaluation of the KL divergence requires nested Monte Carlo, which is intractable especially with expensive forward models and when wrapped around by an optimization routine. To make this procedure computationally feasible, we introduce asymptotically exact local approximations targeting an inner component of the objective function. Infinitely refining the surrogate model ensures the asymptotic decay of the surrogate bias. Furthermore, balancing error contributions from different parts of the estimator implies an ideal and efficient refinement strategy.

Andrew Davis
Cold Regions Research and Engineering Laboratory
andrew.add@gmail.com

Xun Huan
Sandia National Laboratories
xhuan@umich.edu

Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu

MS378

Optimal Experimental Design under Model Uncertainty for Linear Inverse Problems

Inverse problems stemming from real-world applications often contain inherent uncertainties in the governing model in addition to the uncertainties in the inversion parameters. We present a method for computing optimal sensor placements for Bayesian linear inverse problems governed by PDEs with model uncertainties. More specifically, given a statistical distribution for the model uncertainties, we find optimal sensor placements that minimize the expected value of the posterior parameter uncertainty. That is, we follow an A-optimal design strategy, with the design criterion given by the expected value of the trace of the posterior covariance operator. As a model application, we consider inversion for the initial condition in a subsurface flow problem with inherent uncertainty in the permeability fields.

Karina Koval
Courant Institute of Mathematical Sciences
koval@cims.nyu.edu

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

Alen Alexanderian
NC State University
alexanderian@ncsu.edu

MS378

Goal Oriented Information-Theoretic Observation Selection Strategies for Linear Bayesian Inverse Problems

We study the Bayesian linearGaussian model with a large number of observations, and propose several algorithms for solving the combinatorial problem of observation selection/optimal experimental design in a goal-oriented setting. Here, the quantity of interest (QoI) is not the model parameters, but some (vector-valued) function of the parameters. We wish to select a subset of the candidate observations that is most informative for this QoI, in the sense of reducing its uncertainty. More precisely, we seek to maximize the mutual information between the selected observations and the QoI. Finding the true optimum is NP hard, and in this setting, the MI objective is in general not submodular. We thus introduce several algorithms that approximate the optimal solution, including a greedy approach, a minorizemaximize approach employing modular bounds, and certain score-based heuristics. We develop performance bounds for some of these algorithms and compare their performance numerically.

Fengyi Li
Massachusetts Institute of Technology
fengyil@mit.edu

Jayanth Jagalur Mohan
MIT
jagalur@mit.edu

Youssef M. Marzouk
Massachusetts Institute of Technology
ymarz@mit.edu

MS378

Sensitivity Analysis for Inverse Problems and Sensor Placement

We present an approach to determine the sensitivity of sensor locations in a PDE-constrained inverse problem. Our implementation computes a directional derivative through the optimality system and uses the Singular Value Decomposition (SVD) to determine dominate perturbation directions. Unlike standard sensitivity analysis which evaluates the sensitivity of a quantity of interest to uncertain variables, we compute the sensitivity of the inverse solution with respect to sensor weights and auxiliary uncertain variables. The methodology can make use of a full or reduced space approaches to solve the optimality system, and is coupled with a randomized eigenvalue solver for the SVD computation. We demonstrate a parallel implementation on a simple numerical prototype in which we evaluate the importance of sensor locations and auxiliary uncertain variables.

Bart G. Van Bloemen Waanders
Sandia National Laboratories
bartv@sandia.gov

Joseph L. Hart
North Carolina State University
jlhart3@ncsu.edu

MS379

Scaling Up and Randomizing Derivative-free Optimization for Machine Learning

In this talk, we will introduce several black-box optimization problems arising in machine learning and discuss what challenges and opportunities these applications present for current derivative-free optimization algorithms. In particular, for these applications derivative-free methods need to scale up to thousands of variables, and can utilize parallel computations effectively and can potentially benefit from randomized techniques. We will present several new and existing algorithms, specifically finite-difference quasi-Newton methods and model-based trust region methods, and compare their theoretical and numerical behavior.

Albert S. Berahas
Department of Engineering Sciences and Applied Mathematics
Northwestern University
albertberahas@u.northwestern.com

Katya Scheinberg
Lehigh University
kas410@lehigh.edu

Jorge Nocedal
Department of Electrical and Computer Engineering
Northwestern University
j-nocedal@northwestern.edu

Richard H. Byrd
University of Colorado
richard@cs.colorado.edu

MS379

Derivative-free Robust Optimization

The paradigm of robust optimization (RO) has attracted significant attention from the optimization community over the past decade as a means of modeling and obtaining solutions in the presence of uncertainty. In RO, a modeler specifies an uncertainty set in which the uncertain parameters involved in the optimization problem can lie, and the optimization model aims to inoculate a solution against the worst-case realization in the uncertainty set. However, much of the work in RO to date has been for derivative-based optimization and imposes strong assumptions on both the objective function and the uncertainty set for the sake of deriving tractable formulations of problems. In this talk, we will consider robust optimization of a black box function. We develop and analyze a derivative-free variant of an inexact outer approximations method for solving a general unconstrained minimax problem. We will compare with the stochastic approach of chance-constrained optimization, a class of problems in which a probability distribution for the uncertain parameters is inferred via sample

average approximation and only a certain percentage of constraints realized by the uncertain parameters must be satisfied. Using the framework of derivative-free manifold sampling, we develop and analyze a method for solving these problems.

Matt Menickelly
Lehigh University
mmenickelly@anl.gov

Stefan Wild
Argonne National Laboratory
wild@mcs.anl.gov

MS379

A Stochastic Levenberg-Marquardt Method using Random Models with Application to Data Assimilation

Globally convergent variants of the Gauss-Newton algorithm are often the preferred methods to tackle nonlinear least squares problems. Among such frameworks, Levenberg-Marquardt and trust-region methods are two well-established, similar paradigms, with analyses that are often close in spirit. Both schemes have been successfully studied when the Gauss-Newton model is replaced by a random model, only accurate with a given probability. Meanwhile, problems where even the objective value is subject to noise have gained interest, driven by the need for efficient methods in fields such as data assimilation. In this paper, we describe a stochastic Levenberg-Marquardt algorithm that handles noisy objective function values and random models, provided sufficient accuracy is achieved in probability. In particular, if the probability of accurate function estimates and models is sufficiently large, we establish that the proposed algorithm converges globally to a first-order stationary point of the problem with probability one. Furthermore, we bound the expected number of iterations needed to reach an approximate stationary point. Our method relies on a specific scaling of the regularization parameter, in order to leverage existing theory for trust-region algorithms. We finally present an application of our method to variational data assimilation, where stochasticity arises from the so-called ensemble methods.

El houcine Bergou
INRA, Jouy-en-Josas, France
elhoucine.bergou@inra.fr

Youssef Diouane
Institut Supérieur de lAéronautique et de lEspace (ISAE)
Toulouse, France
youssef.diouane@isae.fr

Vyacheslav Kungurtsev
Czech Technical University in Prague
Department of Computer Science
vyacheslav.kungurtsev@fel.cvut.cz

Clément W. Royer
Wisconsin Institute for Discovery
University of Wisconsin-Madison
croyer2@wisc.edu

MS379

Complexity of ASTRO-DF: Adaptive Sampling Trust-region Methods for Derivative-free Simulation Optimization

ASTRO-DF is a class of adaptive sampling algorithms for solving simulation optimization problems that are derivative-free in the sense that they do not rely on direct observations of the function derivatives. We devise an iterative trust-region algorithm whereby the salient feature is the incorporation of adaptive sampling and replication to ensure that the simulation effort is spent efficiently. ASTRO-DF has been demonstrated to generate iterates that globally converge to a first-order critical point with probability one. In this study, we identify the time complexity in two aspects: iteration complexity and work complexity that represents the total amount of sampling utilized by ASTRO-DF. We compare these complexities to those obtained from methods that use traditional random or fixed sampling rules.

Sara Shashaani
University of Michigan
sshasha2@ncsu.edu

Raghu Pasupathy, Daniel Carvajal
Purdue University
pasupath@purdue.edu, dvasque@purdue.edu

MS380

Global Sensitivity Analysis for PDE-constrained Optimization

Many problems in engineering and sciences require large scale optimization constrained by partial differential equations (PDEs). Though PDE-constrained optimization is itself challenging, most applications pose additional complexity, namely, uncertain parameters in the PDEs. This talk presents a computationally efficient and scalable approach to determine the sensitivity of the solution of a PDE-constrained optimization problem to changes in the uncertain parameters. In contrast to traditional sensitivity analysis, which focuses on the sensitivity of the PDE solution, the approach presented in this talk is oriented toward the ultimate goal, control or design of the system. A randomized eigenvalue solver is used to convert the computational bottleneck of the algorithm into an embarrassingly parallel loop. The method is demonstrated on a nonlinear multi-physics problem with 153 uncertain parameters.

Joseph L. Hart
North Carolina State University
jlhart3@ncsu.edu

Bart G. Van Bloemen Waanders
Sandia National Laboratories
bartv@sandia.gov

MS380

Sensitivity Indices for Computer Codes Taking Val-

ues in General Metric Spaces

We consider a black-box code f from $E := E1E2Ed$ valued in some separable metric space (X; d). The output is denoted by Z given by $Z = f(X(1); :::; X(d))$: Various authors perform a sensitivity analysis when X = Rk for example by considering the so-called Sobol indices. Nevertheless, for more general output spaces defining some sensitivity indices is not easy. In the first part of presentation, we explain how to define some nice sensitivity indices when X is a general metric spaces. Then we construct some nice estimators of these indices that converge almost surely and are asymptotically normal. This new approach for defining sensitivity indices allows to revover, when the output space is Rk the classical Sobol indices and the Cramér Von Mises indices. Moreover, the proposed estimation procedure improves the classical one, in the sense that we achieve the same quality of estimations with less call of the computer code. In the second part of the presentation, we present two applications, the first one when the space X is a compact Riemannian manifold and the second one when the space X is the Wasserstein space. This last case allows us to define and estimate well taillored indices for the so-called stochastic codes.

Ricardo Fraiman
Universidad de la Republica
Uruguay
rfraiman@cmat.edu.uy

Fabrice Gamboa
Institut de Mathématiques de Toulouse. AOC Project
University of Toulouse
fabrice.gamboa@math.univ-toulouse.fr

Thierry Klein
Institut de Mathématiques de Toulouse
UMR CNRS 5219 Université de Toulouse
thierry.klein@math.univ-toulouse.fr

Agnès Lagnoux
Institut de Mathématiques de Toulouse
Université Toulouse 2
lagnoux@univ-tlse2.fr

Leonardo Moreno
Universidad de la Republica
Uruguay
mrleomr@hotmail.com

MS380

Application of Global Sensitivity Analysis and Monte Carlo Filtering for Constrained Models

Global sensitivity analysis (GSA) is a widely used tool for identification of key parameters whose uncertainty most affects model output. Most existing techniques are designed for models with independent inputs. There is a wide class of models involving inequality constraints which impose structural dependences between model variables. Consider a model $f\left(\vec{x};\vec{\theta}\right)$, and the vector of constraint functions $\vec{g}\left(\vec{x};\vec{\theta}\right)$, where $\vec{x}$ is a vector of space variables, $\vec{\theta}$ is a random vector. Define probability $p(\vec{x})$ of occurrence of undesirable (reliability estimation analysis) or desirable (chemical engineering) event as

$$p(\vec{x}) = P_\theta\left[\vec{g}\left(\vec{x};\vec{\theta}\right) \geq \vec{g}^*\right].$$

Here $\vec{g}^*$ is a given threshold. Ranking of inputs of a random vector $\vec{\theta}$ in order of their importance belongs to a factor mapping setting of GSA. It typically requires rather complex techniques such as MC Filtering, constraint GSA, etc. We propose formulations via indicator functions which allow transformation of constrained problem into a class of problems with independent variables. Practical problems of identification domains such that $\Omega(\vec{x}; p(\vec{x}) \geq p^*)$, where p^* is a critical value of probability shows efficiency of the proposed method.

Sergei S. Kucherenko, Dimitris Giamalakis
Imperial College London
s.kucherenko@broda.co.uk, d.giamalakis@imperial.ac.uk

Nilay Shah
Imperial College London, London SW7 2AZ, UK
n.shah@imperial.ac.uk

MS380

Alternative to Sobol' Indices for the Multivariate Response Models

In uncertainty quantification, classical multivariate sensitivity analysis (MSA) extends variance-based sensitivity analysis to cope with the multivariate response models, and it aims to apportion the variability of the multivariate response into input factors and their interactions. The first-order and total-effect covariance matrices from MSA, which allow for assessing the effects of input factors, provide useful information about interactions among input factors, the order of interactions, and the magnitude of interactions over all model outputs. In this abstract, we propose and study new generalized sensitivity indices (GSIs) using the Frobenius norm when partial orders such as the Loewner ordering on covariance matrices is not possible. The new GSIs account for the off-diagonal components of the first-order and total-effect covariance matrices. We also derive the consistency, the asymptotic normality, and the asymptotic confidence regions of the GSIs estimators. The new GSIs come down to Sobol' indices for the single response models.

Matieyendou Lamboni
University of Guyane
matieyendou.lamboni@gmail.com

MS381

Optimizing Quantum Circuit Synthesis Using Numerical Methods

To have a deeper understanding of quantum algorithmic or to sketch a classical / quantum hybrid software, it is crucial to have compilation methods - synthesis or optimization of quantum circuits - that are efficient in classical time (time compilation) and in quantum time (execution time, fidelity of the circuit). To this end we present nonlinear optimization methods that can provide circuits of near optimal size in reasonable time and on various small-scale

problems. We present some results on specific unitaries e.g in quantum chemistry. Then we discuss the various implications of these algorithms. In particular, these methods highlight a tradeoff between classical and quantum time, but also the need to find efficient heuristics to make these methods scalable.

Timothée Goubault de Brugière, Marc Baboulin
University of Paris-Sud
timothee.goubault@lri.fr, marc.baboulin@lri.fr

Benot Valiron
CentraleSupélec
benoit.valiron@lri.fr

Cyril Allouche
ATOS/BULL
Cyril.Allouche@atos.net

MS381

Seeking Mid-term Quantum Computing Supremacy with Numerical Simulation

Quantum Computing promises to tackle computations which are intractable for todays supercomputers. Among the expected beneficiaries one can cite numerical analysis, AI, combinatorial optimization, chemistry... Though, most of the known quantum algorithms underlying these applications require a large number of *perfect* qubits. That is: unlimited coherence time, perfect fidelity, full connectivity. Designing such ideal qubits is a tremendous challenge for scientists and engineers, due to the very laws of quantum, which make quantum information very fragile and error correction extremely costly. Scaling them to thousands or millions is yet another challenge. This may take one decade or even more. A question is now emerging in the QC community: in the meantime can early devices with a limited number of *imperfect* qubits exhibit any kind of computing advantage? Brand new quantum algorithms have to be invented and optimized for these very contrained devices. Numerical techniques, in particular simulation and optimization, offer here an invaluable tooling. We will introduce some of these techniques through QLM, a quantum software framework developed at the Atos Research Labs. From a physical model of the quantum device, a quantum circuit can be simulated, then drastically optimized so as to maximize its fidelity. These techniques can then efficiently assist the algorithmics researchers, as well as the quantum hardware engineers, in their quest for quantum supremacy.

Cyril Allouche
ATOS/BULL
Cyril.Allouche@atos.net

Timothée Goubault de Brugière
University of Paris-Sud
timothee.goubault@lri.fr

MS381

Accelerating Scientific Applications with Quantum Computing

Quantum computing has the potential to accelerate innovation and discovery by enabling a new paradigm of scientific computing. However, there is the practical challenge of integrating these novel algorithms, devices, and execution models into state of the at scientific workflows. Following a brief motivation, we discuss how these scientific uses cases are realized by the development of domain-specific algorithms, software infrastructure, and hardware testbeds. We promote the use of metrics and benchmarks that enable comparisons between quantum computing and conventional high-performance computing methods, and we conclude by reporting on our recent demonstrations of these ideas for computational chemistry and nuclear physics.

Travis Humble, Alexander McCaskey
Oak Ridge National Laboratory
humblets@ornl.gov, mccaskeyaj@ornl.gov

MS381

Modern Methods in Quantum Compiling

Quantum compiling is concerned with the representation of general unitary operations by circuits built from some chosen set of quantum gates. The circuit representation of a unitary U is exact if the product of the gates composing the circuit is equal to U. The representation is approximate up to epsilon if this product is at distance epsilon of U in the operator norm. In the last few years, the field of quantum compiling was rejuvenated by the introduction of methods from algebraic number theory. In particular, such number-theoretic methods were used to provide an optimal solution to the problem of approximating single-qubit unitaries using Clifford+T circuits. In this talk, I will present an efficient algorithm for the optimal approximation of single-qubit unitaries using Clifford+T circuits and discuss recent advances in the multi-qubit decomposition of unitary operators.

Neil J. Ross
Dalhousie University
neil.jr.ross@dal.ca

MS382

On Some New Variants of the Multiscale Finite Element Method

The Multiscale Finite Element Method (MsFEM) is a Finite Element type approximation method for multiscale PDEs, where the basis functions used to generate the approximation space are precomputed and are specifically adapted to the problem at hand. Many ways to define these basis functions have been proposed in the literature over the past years. In this work, we introduce and analyze a specific MsFEM variant, the construction of which is inspired by component mode synthesis techniques [Hetmaniuk and Lehoucq, M2AN 2010; Hetmaniuk and Klawonn, ETNA 2014]. However, in contrast to these approaches, we do not solve eigenvalue problems but rather consider enrichment by Legendre polynomials of the boundary conditions imposed on the highly oscillatory basis functions. Motivation for this new approach, error estimates and numerical results will be discussed. Joint work with U. Hetmaniuk (University of Washington), C. Le Bris and P.-L.

Rothe (ENPC and Inria).

Frederic Legoll
Laboratoire Navier (ENPC) and Inria
frederic.legoll@enpc.fr

MS382

Constraint Energy Minimizing Generalized Multiscale Finite Element for Flows in Heterogeneous Media

Many practical applications contain multiple scales and high contrast. Due to the scale disparity and the contrast, we need to use some multiscale approaches to handle this type of problem. In this talk, we will introduce constraint Energy Minimizing Generalized Multiscale Finite Element Method (CEM-GMsFEM) for Darcy flow equation in heterogeneous media. This method gives a coarse mesh dependent convergence with a low dimension finite element space. In this approach, the multiscale basis functions are constructed in two steps. The first step is using a spectral problem to select some auxiliary basis functions which can capture the local feature of the solution. The next step is the constructing of the localized multiscale basis functions by solve an energy minimizing problem in an oversampling domain. We will present a convergence analysis of the method with numerical result to confirm our theoretical results. This method can be extended for solving vary problems, for example, nonlinear problem. We will give some numerical examples to demonstrate the performance of the method in different applications.

Tony Cheung
Texas A&M University
tonycsw2905@math.tamu.edu

Wing Tat Leung
University of Texas at Austin
sidnet123@gmail.com

Mary F. Wheeler
Center for Subsurface Modeling, ICES
University of Texas at Austin
mfw@ices.utexas.edu

Eric Chung
The Chinese University of Hong Kong
Department of Mathematics
tschung@math.cuhk.edu.hk

Yalchin Efendiev
Dept of Mathematics
Texas A&M University
efendiev@math.tamu.edu

MS382

On the Convergence Rates of GMsFEMs for Heterogeneous Elliptic Problems Without Oversampling Techniques

This work is concerned with the rigorous analysis on the Generalized Multiscale Finite Element Methods (GMsFEMs) for elliptic problems with high-contrast heterogeneous coefficients. GMsFEMs are popular numerical methods for solving flow problems with heterogeneous high-contrast coefficients, and it has demonstrated extremely promising numerical results for a wide range of applications. However, the mathematical justification of the efficiency of the method is still largely missing. In this work, we analyze two types of multiscale basis functions, i.e., local spectral basis functions and basis functions of local harmonic extension type, within the GMsFEM framework. These constructions have found many applications in the past few years. We establish their optimal convergence in the energy norm under a very mild assumption that the source term belongs to some weighted L^2 space, and without the help of any oversampling technique. Furthermore, we analyze the model order reduction of the local harmonic extension basis and prove its convergence in the energy norm. These theoretical findings shed insights into the mechanism behind the efficiency of the GMsFEMs.

Guanglian Li
Texas A&M University
lotusli0707@gmail.com

MS383

Energy-stable Open Boundary Conditions and Associated Algorithm for Multiphase Flows of N Immiscible Incompressible Fluids

We introduce a set of effective outflow/open boundary conditions, which is suitable for simulating N-phase flows (N¿=2) in domains involving outflows or open boundaries. These boundary conditions satisfy two properties: energy stability and reduction consistency. They are devised such that their contributions to the N-phase energy balance equation will not cause the total system energy to increase over time. Therefore, these outflow/open boundary conditions are very effective in overcoming the backflow instability. The reduction consistency ensures the inherent equivalence relations between N-phase system and the corresponding smaller system when some of the fluid components are absent from the N-phase system. We also present an efficient algorithm for numerically treating the proposed boundary conditions together with the N-phase governing equations. The proposed algorithm involves only solving a set of de-coupled individual Helmholtz equations in each time step with constant and time-independent coefficients. We present ample numerical examples to confirm that the proposed method produces physically accurate results.

Zhiguo Yang, Suchuan Dong
Purdue University
yang1508@purdue.edu, sdong@math.purdue.edu

MS383

Modeling and Computations for Actomyosin Systems in the Cell Cortex

Abstract not available.

Maosheng Jiang
Beijing Computational Science Research Center, China
msjiang@csrc.ac.cn

Qi Wang

University of South Carolina
Beijing Computational Science Research Center
qwang@math.sc.edu

MS383

Numerical Analysis of a Variable Step Bdf2 Scheme for the Cahn-Hilliard Equation

We propose and analyze a novel second order in time variable step backward difference scheme (BDF2) for the Cahn-Hilliard equation. The construction is based on convex splitting, variable step BDF2, and a viscous regularization at the discrete level. The scheme enjoys a discrete energy law and hence is long-time stable. It is also uniquely solvable. Moreover, we prove the second order accuracy utilizing a novel generalized discrete Gronwall type inequality. Numerical results will be presented as well .

Xiaoming Wang
Fudan University
Florida State University
wxm.math@outlook.com

Wenbin Chen
School of Mathematical Science,
Fudan University, China
wbchen@fudan.edu.cn

Zhuying Zhang
Fudan University
15210180019@fudan.edu.cn

MS384

Accelerating Training Phase in Time-dependent Nonlinear Model Order Reduction

Several reduced order models have been successfully developed for nonlinear dynamical systems. To achieve a considerable speedup, a hyper-reduction step is needed to reduce the computational complexity due to nonlinear terms. Many hyper-reduction techniques require the construction of nonlinear term basis, which introduces a computationally expensive offline phase. A novel way of constructing nonlinear term basis within the hyper-reduction process is introduced. In contrast to the traditional hyper-reduction techniques where the collection of nonlinear term snapshots is required, the SNS method completely avoids the use of the nonlinear term snapshots. Instead, it uses the solution snapshots that are used for building a solution basis. Furthermore, it avoids an extra data compression of nonlinear term snapshots. As a result, the SNS method provides a more efficient offline strategy than the traditional model order reduction techniques, such as the DEIM, GNAT, and ST-GNAT methods. The SNS method is justified by the conforming subspace condition and the subspace inclusion relation. Numerical results support that the accuracy of the solution from the SNS method is comparable to the traditional methods.

Youngsoo Choi
Lawrence Livermore National Laboratory
choi15@llnl.gov

Deshawn Coombs
Syracuse University
dmcoombs@syr.edu

Robert W. Anderson
Lawrence Livermore National Laboratory
anderson110@llnl.gov

MS384

Reduced-order Modeling of Coupled Flow-geomechanics Problems

A reduced-order modeling (ROM) procedure for coupled subsurface flow-geomechanics problems is presented. Coupled flow and geomechanics can be important in oil-gas production and in carbon storage operations. Earlier ROM-related developments in the context of subsurface flow have focused on the flow equations, while in this work the geomechanical stress equilibrium equations and associated displacement variables are also incorporated. The ROM methodology developed here entails the use of proper orthogonal decomposition (POD) combined with trajectory piecewise linearization (TPWL). Solutions with new sets of controls (in this case well settings) are represented via linearization around previously simulated solutions. The POD-TPWL ROM is tested on 2D and 3D coupled problems with two fluid phases. Examples involving time-varying injection and production rates, which differ from those used in training simulations, are presented. Runtime speedups, relative to high-fidelity simulations, of over a factor of 100 are achieved for the cases considered. These large speedups are attained because the model is (piecewise) linear and all computations are in reduced space. Output quantities of interest include phase production rates and the maximum and minimum principal stress fields. Accurate results for these quantities, relative to those from high-fidelity simulations, are achieved for a range of test cases.

Louis J. Durlofsky, Zhaoyang Larry Jin
Stanford University
Dept of Energy Resources Engineering
lou@stanford.edu, zjin@stanford.edu

Timur T. Garipov
Department of Energy Resources Engineering,
Stanford University
tgaripov@stanford.edu

Oleg Volkov
Stanford University
ovolkov@stanford.edu

MS384

Multiscale Two-stage Solver for Poroelastic Media

In this work we focus on numerical solution of Biots equations governing coupled flow and deformation in poroelastic media. Finite-element/finite volume spatial discretization combined with backward implicit time discretization results in block 2x2 linear systems to be solved at every step. Often the complexity and spatial resolution requirements lead to high computational cost of the full order model that

must be solved by a Krylov solver. We propose a preconditioning technique based on model reduction, in which a coarse problem is constructed using multiscale basis functions for pressure and displacement fields. An approximate fine-scale solution is obtained by interpolation of the coarse scale fields. In order to achieve robust convergence of the Krylov solver, the multiscale step is multiplicatively combined with a block-triangular smoother based on fixed-stress decoupling of unknowns that effectively dampens high-frequency error modes, resulting in a two-stage preconditioner. The resulting solver achieves accuracy of the full model at much lower computational cost. We demonstrate the efficiency and scalability of the proposed method using a number of synthetic benchmarks and realistic subsurface poromechanical systems.

Sergey Klevtsov
Stanford University
klevtsov@stanford.edu

Nicola Castelletto
Stanford University
Department of Energy Resources Engineering
castelletto1@llnl.gov

Hamdi Tchelepi
Stanford University
Energy Resources Engineering Department
tchelepi@stanford.edu

MS384

Lagrangian Data-driven Reduced Order Modeling of Finite Time Lyapunov Exponents.

We propose a Lagrangian data-driven reduced order model (ROM) for an efficient and relatively accurate numerical simulation of the finite time Lyapunov exponent (FTLE) field. To generate the basis, the new Lagrangian ROM explicitly uses Lagrangian data (i.e., the FTLE field) in addition to the Eulerian data (i.e., the velocity field); the Eulerian ROM, on the other hand, uses only Eulerian data. The Lagrangian ROM and the Eulerian ROM are compared in the numerical simulation of the quasi-geostrophic equations. It is shown that the new Lagrangian ROM outperforms the Eulerian ROM with respect to both Eulerian (velocity) and Lagrangian (FTLE field) accuracy criteria. Furthermore, the online CPU time of the Lagrangian ROM is dramatically lower than the CPU time of the corresponding fine resolution numerical simulation. Our numerical investigation suggests that Lagrangian ROMs could be used as an efficient and relative accurate alternative to fine resolution numerical simulations to generate the FTLE field.

Changhong Mou
Virginia Tech
cmou@math.vt.edu

Xuping Xie
Oak Ridge National Lab
xiex@ornl.gov

Peter Nolan
Virginia Polytechnic Institute and State University
pnolan86@vt.edu

Shane D. Ross
Virginia Tech
Engineering Mechanics program
sdross@vt.edu

Traian Iliescu
Department of Mathematics
Virginia Tech
iliescu@vt.edu

MS385

Dimensional Considerations in Probing the Upper Limits of Energy Storage in Nanostructured Electrochemical Devices

To enhance the capacity of a charge storage device, nanostructures such as one-dimensional carbon nanotubes (CNTs) have often been employed, much of the time with the effect of increasing the surface area exposed to electrolyte or enhancing the electrode inter-connections and reducing the electrical resistance. Here, for nanostructured battery electrodes, we propose a new electrical model based on 2D and 1D DOS (density of states) and compare with related experiments. The model is broadly based on Marcus-Hush-Chidsey (MHC) kinetics and is aimed to extend the utility of the MHC formulations, to a larger class of materials and situations. We employ an integral expression for current, which involves the DOS, the Fermi-Dirac function, and the Arrhenius rate law as the basis. An interesting feature of our model, is the observation in lower dimensions of electrical current oscillations as a function of the overpotential, corresponding to the gradual population (and de-population) of each successive sub-band. We posit that the consideration of a variable/non-constant DOS in MHC electrokinetics may yield tests of dimensional character and contribute to enhanced storage in electrochemical systems.

Hidenori Yamada, Rajaram Narayan,
Prabhakar R. Bandaru
UC San Diego
hyamada@ucsd.ed, ranaraya@ucsd.edu, pbandaru@ucsd.edu

MS385

A New Type of Flow Battery Based on Oxidation State Instability

We developed novel membraneless batteries based on two immiscible protic solvents (ethanol or methanol and saltwater). These batteries were also capable of producing hydrogen on demand, and hence allowed one to construct a system that both generated the power to compress hydrogen and simultaneously produce the hydrogen to be compressed. Several counter intuitive results from experiments related to the two-phase system led us to develop a totally new type of battery. This next generation battery involves a catholyte with a current collector and an anode imbedded in the catholyte with no membrane to separate them. A single cell generates around 2 volts and 1 watt at the batteries internal resistance. If one accounts for the energy

associated with hydrogen gas production, then the energy density exceeds 170Whr/kg. At the anode there is an oxidation reaction followed by proton reduction, resulting in hydrogen gas production. At the cathode current collector there is a separate reduction reaction. The question that should come to mind is, if there is oxidation-reduction at the anode, where do the electrons come from to create the reduction at the cathode current collector? We propose a novel theory to answer this question: There is an oxidation state instability at the anode, which to energetically stabilize itself, requires the release through the load of additional electrons. We will present videos to demonstrate the system and its workings.

John H. Cushman
Center for Applied Mathematics
Purdue University
jcushman@purdue.edu

MS385

Small-scale Acoustofluidics in Rechargeable Battery Technologies

Acoustics is undergoing a renaissance, finding new life at small scales in generating extraordinary forces upon fluids and particles in handheld devices for micro to nano-scale lab-on-a-chip applications. Helping to deliver on the many promises of micro to nano-scale fluidics in medicine, biology, chemistry, and forensics, the nonlinear aspects of transmitting sound through a medium at relatively high frequencies (10 MHz and more) has long been known but now take on entirely new importance as "acoustofluidics". In this talk, a few fundamental concepts of acoustofluidics are provided that shows why it is now so useful. We then show the inadequacy of classical acoustics to explain and understand small-scale acoustofluidics, producing an urgent need to recast past approaches in light of a myriad of recently observed and curious physical phenomena, particularly in the context of rechargeable batteries. Along the way, the fascinating underlying physics will be described and illustrated via specific examples, tying together the acoustics, fluid dynamics, and broader physical phenomena appearing in these systems.

James Friend
University of California, San Diego
jfriend@eng.ucsd.edu

MS385

A Diffuse Interface Method for Electrochemical Systems

Microstructures have significant effects on the performance of materials for electrochemical energy storage and conversion, in which thermodynamics, transport, and reaction play critical roles. The smoothed boundary method, wherein the microstructures are described by a field that changes rapidly over the length scale of interest but smoothly over the length scale of simulation resolution, have been developed to couple the governing partial differential equations in irregular-shaped domains. After an overview of the method, its application to transport in materials with multiple diffusion paths is discussed. The results are analyzed to yield a coarse-grained description of the transport properties that accurately accounts for the microstructures.

Katsuyo Thornton, W. Beck Andrews
Department of Materials Science and Engineering
University of Michigan
kthorn@umich.edu, wband@umich.edu

MS386

FFT Applications and Benchmarks

This talk will highlight the FFT project in the Exascale era followed by a description of an ongoing project that aims to speed up FFTs on the Cray XC40 machine by utilizing the full bandwidth offered by the cluster.

Samar A. Aseeri
Computational Scientist at KAUST
Extreme Computing Research Center (ECRC)
samar.aseeri@kaust.edu.sa

MS386

A Periodic Treecode Method for Electrostatics In Molecular Dynamics Simulations

Molecular dynamics (MD) simulations in periodic boundary conditions typically employ smooth particle mesh Ewald (SPME) to compute electrostatic interactions. The excellent serial scaling of SPME is primarily due to the use of FFT. However, this presents a problem in parallel since the communication requirement of the 3D-FFT leads to poor parallel scaling. In this work we develop an periodic tree code algorithm as an alternative to SPME for parallel MD simulations. We present simulation and parallel timing results compared to the smooth particle mesh Ewald method.

Henry A. Boateng
Department of Mathematics
Bates College
hboateng@bates.edu

MS386

High-dimensional Sparse FFT

We discuss the development of high-dimensional sparse FFT. One-dimensional sparse FFTs introduced in "Adaptive sublinear time Fourier algorithm" (2013) and "A multiscale sub-linear time Fourier algorithm for noisy data" (2016) by A. Christlieb, D. Lawlor and Y. Wang quickly approximate functions represented by only a few Fourier basis functions using a few samples (function evaluations) without noise and with noise respectively. The algorithms can be directly applied to compute the Fourier transform of the high-dimensional functions. However, it becomes hard to implement them if the dimension gets too large. In this talk, we introduce two new concepts: *partial unwrapping* and *tilting*. These two ideas allow us to efficiently compute the high-dimensional sparse FFT using the ideas in the two papers mentioned above whether the samples are noisy or not. Both numerics and average-case analysis will

be presented.

Bosu Choi
Michigan State University
choibosu@utexas.edu

Andrew J. Christlieb
Michigan State Univerity
Dept. of Comp. Math., Sci & Engr.
christli@msu.edu

Yang Wang
Hong Kong University of Science and Technology
yangwang@ust.hk

MS386

Rank-1 Lattice Based High-dimensional Approximation and FFT

We consider the approximate reconstruction of a high-dimensional (e.g. d=10) periodic function from samples using trigonometric polynomials. As sampling schemes, we use rank-1 lattices, which may be constructed by a component-by-component approach when the locations of the approximately largest Fourier coefficients are known. With the help of a single one-dimensional fast Fourier transform (FFT), we are able to compute the Fourier coefficients in the high-dimensional case. For functions from Sobolev Hilbert spaces of generalized mixed smoothness, error estimates are presented where the sampling rates are best possible up to logarithmic factors. We give numerical results which confirm our theoretical estimates. Additionally, we discuss an approach where we use multiple instances of rank-1 lattices. This allows for efficient construction algorithms and we obtain improved error estimates where the sampling rates are optimal up to small constant offset in the exponent. For the case where we do not know the locations of important Fourier coefficients, we present a method which approximately reconstructs high-dimensional sparse periodic signals in a dimension-incremental way based on adaptively chosen rank-1 lattices and 1-dimensional FFTs. This is based on joint work with Glenn Byrenheid, Lutz Kmmerer, Daniel Potts, and Tino Ullrich.

Toni Volkmer
Chemnitz University of Technology
toni.volkmer@mathematik.tu-chemnitz.de

MS387

Optimized Quadratures for the Sommerfeld Representations of Acoustic and Electromagnetic Fields

A classical representation for the solution of the Helmholtz or Maxwell equations is the Sommerfeld integral, which yields an expansion of acoustic or electromagnetic fields in the frequency domain in terms of propagating and evanescent plane waves. Using a change of variables and contour deformation, we show how to construct optimized quadrature schemes for specific ranges of the source and target over a wide range of frequencies. This is of particular importance in the implementation of "modern" fast multipole methods, where the number of quadrature nodes is a major determinant of algorithm performance. This is joint work with J. Sifuentes, L. Greengard, and M. Rachh.

Zydrunas Gimbutas
National Institute of Standards and Technology
zydrunas.gimbutas@nist.gov

MS387

Multiple Junctions and Transmission Problems in Elliptic PDEs

In this talk we describe the analytical behaviour of solutions to a boundary integral equation arising in the solution of the transmission problem for Laplace's equation in composite regions, ie. regions in which two or more interfaces meet at a point. Using these results we construct an efficient high-order Nystrom scheme for solving the boundary integral equation which requires relatively few discretization nodes. Finally, we illustrate the properties of the resulting algorithm with a few numerical examples.

Jeremy Hoskins
Yale University
jeremy.hoskins@yale.edu

MS387

Analytical Expressions for the Solutions of Elliptic PDEs Near Corners, Edges, and Conical Points

Solutions of elliptic partial differential equations (PDEs) are notoriously singular in the vicinity of geometric singularities such as corners, edges, and conical points. It was recently observed that, in many cases, when an elliptic PDE is formulated as a boundary integral equation, the solutions to the integral equation are representable as explicit series of elementary functions in the vicinities of the geometric singularities. The resulting expressions provide highly accurate representations of the solutions, and lend themselves to the construction of efficient numerical schemes. In this talk I will discuss recent results in this area, including the derivation of such explicit series representations for more general classes of curves and more general boundary conditions, as well as preliminary results in three dimensions.

Kirill Serkh
Yale University
Applied Mathematics Program
kirill.serkh@yale.edu

MS387

A Fast Boundary Integral Method for Generating High-order Surface Meshes

In order to develop truly high-order integral equation-based solvers for boundary value problems in three dimensions, all aspects of the solver must be high-order: discretization of the unknown, quadratures for singular Greens functions, and most importantly, the description of the underlying geometry (surface meshes in this case). During the last 30 years, there have been amazing advances in fast algorithms, discretization, and singular quadrature. Unfortunately, due to the lack of robust schemes for gen-

erating high-quality high-order triangulated surfaces, most integral equation-based simulations have been limited to simple analytically defined geometries, and real-world complex engineering geometries can only be described by flat-triangulations. In this talk we will describe a recently developed algorithm for transforming a flat triangulation (i.e. skeleton) of a smooth boundary of a domain into a high-order curvilinear triangulation that can then be coupled with high-order integral equation methods. The algorithm is based on the fact that convolution of the indicator function of the domain with a Gaussian results in an infinitely differentiable level-set function in the volume. Specific level-sets can then be meshed to high-order, and coincide closely with the original skeleton of the boundary. Various numerical examples will be shown, including high-order results from acoustic and electromagnetic scattering problems.

Felipe Vico
Departamento de Comunicaciones
Universidad Politecnica de Valencia
felipe.vico@gmail.com

MS388

Rethinking HPC Performance and the Role of Learning-based Analytics

Changes in the HPC landscape, such as increasing node count and increasing system complexity, have resulted in the need for automated problem detection and diagnosis. The available monitoring data that can be collected with low overhead has also increased to hundreds (or thousands) of numeric metrics from every node per second. Making the best use of this complex data requires advanced and scalable data analysis techniques. In our recent work, we have designed and applied supervised learning methods and machine learning frameworks that use numeric system monitoring data to diagnose performance variations or to detect illegal/malicious applications. Our frameworks use statistical feature extraction techniques for lowering the storage requirements and increasing performance. For performance variation diagnosis, we can detect over 98% of known performance variation causes. For application detection, we can detect known applications with an F-score over 0.95 and also correctly identify previously unseen applications as unknown. This talk will highlight our automated diagnostics framework and our most interesting results.

Emre Ates
Boston University
ates@bu.edu

Ayse K. Coskun
Boston University
Electrical and Computer Engineering Department
acoskun@bu.edu

MS388

Performance Modeling Experience with MILC on Blue Waters

Application performance modeling and system monitoring are two essential components of application development and compute system development as HPCD moves into the exascale. The ability to predict the impact of changing system characteristics on an application and to measure that impact allows for informed design and decision processes. In the case of the NCSAs Blue Waters system, located at the University of Illinois, application performance modeling was used to predict time to solution for a petascale benchmark of a Lattice Quantum Chromodynamics (LQCD) application prior to system deployment. The difficulties encountered in that modeling effort motivated NCSA to work on a semi-analytical performance-modeling method that combines a simple analytical model to capture the main characteristics of the code, such as numbers and sizes of passed messages and invocation counts of serial code blocks, with statistically sound curve fitting methods to characterize application performance. This is a non-trivial task due to variations in measurements caused by events such as network congestion or system noise effects. System monitoring of resources like the high-speed network allow for detection of abnormal behavior such as the fore mentioned congestion. In this session we will review the relationship between performance modeling and use of system monitoring on Blue Waters.

Greg Bauer
National Center for Supercomputing Applications (NCSA)
gbauer@ncsa.uiuc.edu

Celso Mendes
University of Illinois
cmendes@illinois.edu

William Kramer
NCSA / University of Illinois
wtkramer@illinois.edu

MS388

A Survey of Analysis Techniques Relevant to HPC Resource Modeling

HPC resource monitoring often yields large, high-dimensional datasets. While HPC systems have many broad commonalities, the differences in their details cause resource monitoring data to vary significantly in structure between deployments (even of the same hardware) based on administrative practices, expected workloads, platform size, and many other factors. Furthermore, the repetitive nature of much of the data can make it difficult to distinguish interesting variations in behavior from trivial variations. We discuss large-data analytical techniques that may be of interest to people maintaining current and planning next-generation HPC machines. These techniques may be categorized in several ways:

- based on the type of measurements available (e.g., free-form text logs to regularly-sampled numeric data);
- based on the results each technique provides or, alternatively, the hypothesis each technique tests (i.e, normality, change-detection, anomaly detection, classification);
- based on the decisions which the analysis will support (e.g., identifying and correcting operational problems

such as poor placements, evaluating next-generation HPC platforms).

We use these categorizations to motivate particular techniques.

David Thompson
Kitware, Inc
david.thompson@kitware.com

Thomas Hastings Greer
Kitware, Inc.
hastings.greer@kitware.com

MS388

Characterization of Contention and Failure in Large-scale Systems

The DEPEND group at the University of Illinois focuses on the research, design, and validation of highly available, reliable, and trustworthy computing systems and networks. As part of this work, we have been developing characterizations of contention and failure in large-scale HPC systems as they apply to a number of subsystems, such as network, filesystems, facilities, etc. We present methodologies in run time characterizations of state and regions of interest and how they can be classified and quantified to enable increased understanding and improved performance, throughput, and utilization.

Zbigniew Kalbarczyk
University of Illinois
kalbarcz@illinois.edu

MS389

The Chase Library for Large Hermitian Eigenvalue Problems

ChASE is a new library based on an optimized version of subspace iteration with polynomial acceleration. ChASE is written in C++ using the most current concepts in software engineering which favor a simple integration in application codes and effortless portability over heterogeneous platforms. When solving sequences of Hermitian eigenproblems for a portion of their exterior spectrum, ChASE experiences a considerable speedup and outperforms direct solvers in many scenarios. ChASE performance depends on a number of internal parameters whose choice has been optimized by the automatic estimation of the necessary spectral properties of the eigenproblem. Such feature greatly simplifies its use by the average user. The library is highly scalable, ships with two distinct parallelization schemes and is easily extensible to other computing architectures.

Edoardo A. Di Napoli
Juelich Supercomputing Centre
e.di.napoli@fz-juelich.de

MS389

Inexact Methods for Symmetric Stochastic Eigenvalue Problems

We study two inexact methods for solutions of random eigenvalue problems in the context of spectral stochastic finite elements. In particular, given a parameter-dependent, symmetric matrix operator, the methods solve for eigenvalues and eigenvectors represented using polynomial chaos expansions. Both methods are based on the stochastic Galerkin formulation of the eigenvalue problem and they exploit the Kronecker-product structure of the problem. The first method is an inexact variant of the stochastic inverse subspace iteration [B. Sousedík, H. C. Elman, *SIAM/ASA Journal on Uncertainty Quantification* 4(1), pp. 163–189, 2016]. The second method is based on an inexact variant of Newton iteration. In both cases, the problems are formulated so that the associated stochastic Galerkin matrices are symmetric, and the corresponding linear problems are solved using preconditioned Krylov subspace methods with several novel hierarchical preconditioners. The accuracy of the methods is compared with that of Monte Carlo and stochastic collocation, and the effectiveness of the methods is illustrated by numerical experiments.

Kookjin Lee
University of Maryland College Park
koolee@sandia.gov

Bedrich Sousedik
University of Maryland, Baltimore County
bedrich.sousedik@gmail.com

MS389

Computing Planetary Interior Normal Modes with a Highly Parallel Polynomial Filtering Eigensolver

A highly parallel algorithm has been developed and exploited to compute the planetary normal modes of the elastic-gravitational system, which is approximated via the mixed finite element method on unstructured tetrahedral meshes. The eigenmodes of the relevant generalized eigenvalue problem were extracted by a Lanczos approach combined with polynomial filtering. In contrast with the standard shift-and-invert and the full-mode coupling algorithms, the polynomial filtering technique is ideally suited for solving large-scale 3-D interior eigenvalue problems since it significantly enhances the memory and computational efficiency without loss of accuracy. The parallel efficiency and scalability of this approach are demonstrated on Stampede2 at the Texas Advanced Computing Center. To the best of our knowledge, this is the first time that the direct calculation of the normal modes of 3-D strongly heterogeneous planets, in particular, Earth, Moon, and Mars, is made feasible via a combination of multiple matrix-free methods and a separation of the essential spectra.

Jia Shi
Rice University
jia.shi@rice.edu

Ruipeng Li
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
li50@llnl.gov

Yuanzhe Xi, Yousef Saad
University of Minnesota

yxi26@emory.edu, saad@umn.edu

Maarten de Hoop
Rice University
mdehoop@rice.edu

MS389

Stopping Criteria for Computing Low Rank Approximations with Iterative Methods

Many applications in computational science as well as in machine learning are faced with the problem of calculating a rank-k approximation of a matrix. Although the optimal solution is simply the dominant rank-k singular vector space, if low accuracy is needed solutions far from the optimal can satisfy the requirement. This is one of the reasons why randomized SVD has proven effective on this problem. In this talk we present stopping criteria for traditional iterative methods that address the problem of low rank approximation, rather than computing the singular vectors in small angle. We show that with these criteria iterative methods can be equally effective and more efficient.

Andreas Stathopoulos
College of William & Mary
Department of Computer Science
andreas@cs.wm.edu

Eloy Romero
College of William and Mary
eloy@cs.wm.edu

MS390

Application of Multirate Methods to Multiphysics Problems

Operator splitting with sub cycling for time stepping. The approach is favored due to its simplicity, as it allows code modules to time step with minimal involvement of other modules. However, splitting methods are generally only viable to second order. Certain types of multirate methods may be good alternatives to splitting methods for multi physics applications. They allow for higher than second order time stepping, while maintaining most of the convenience of splitting. In this talk, we discuss some multirate methods that may be good alternatives to splitting and describe some early results in applying such methods to multi physics problems. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

John Loffeld
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
loffeld1@llnl.gov

Carol S. Woodward
Lawrence Livermore National Laboratory
woodward6@llnl.gov

Slaven Peles
Lawrence Livermore Natinoal Laboratory
peles2@llnl.gov

MS390

On the Derivation of a New Class of Multirate Methods Based on Exponential Integrators

In this talk, we develop a new class of multirate methods based on our newly derived exponential integrators applied to semilinear stiff differential equations with fast and slow dynamics. Using the idea of backward errors analysis, we constructed modified differential equations whose exact solutions coincide with the numerical solutions obtained by those exponential integrators. As a trade-off, this approach does not need to deal with a set of complicated coupling and decoupling order conditions like the previous approaches. Consequently, we are able to derive methods of order up to 5 in an elegant and systematic way. Our new methods allow for both fast and slow components to be integrated in explicit-explicit and implicit-explicit manners.

Vu Thai Luan
Department of Mathematics, SMU, Dallas, TX
vluan@smu.edu

MS390

Implicit Multirate Generalized Additive Runge-Kutta Methods

Multirate methods efficiently integrate ODEs that exhibit multiple time scales by taking small timesteps to accurately capture the fast dynamics and large timesteps to cheaply capture the slow dynamics. In this talk, we present new implicit multirate methods for solving systems that are stiff in all partitions. These methods are analyzed using the generalized additive Runge-Kutta (GARK) framework. Special attention is given to the linear stability analysis, which is vital to the development of practical and robust methods. In literature, there are many test problems used to access the stability of multirate methods. We compare them and discuss the implications on A- and L-stability. Finally, the efficiency of the new methods is compared to traditional Runge-Kutta methods.

Steven Roberts
Virginia Tech
steven94@vt.edu

Arash Sarshar
Virginia Polytechnic Institute and State University
USA
sarshar@vt.edu

Adrian Sandu
Virginia Polytechnic Institute
and State University
sandu@cs.vt.edu

MS390

Stable Time Integration for Coupled Ocean-atmosphere Models

Fully coupled ocean-atmosphere models are needed to rep-

resent and understand the complicated interactions, becoming increasingly important in climate change assessment in recent years. Numerical stability issues may arise because of cost-effective time integration. In particular, the contributing factors include using large time stepsize, lack of accurate interface flux, and singe-iteration coupling. We investigate the stability of the coupled ocean-atmosphere models for a variety of interface conditions such as DIrichlet-Neumman condition and bulk condition which is unique to climate modelling. We will also discuss the use of Schwarz-in-time iterative schemes that can add implicitness to the interface points for better stability. The efficiency can be achieved by using modified interface conditions that lead to faster convergence and by exploiting task-level parallelism.

Hong Zhang
Argonne National Laboratory
hongzhang@anl.gov

Paula Egging
UNIVERSITY OF NEBRASKALINCOLN
paula.egging@huskers.unl.edu

Emil M. Constantinescu
Argonne National Laboratory
Mathematics and Computer Science Division
emconsta@mcs.anl.gov

Robert Jacob
Argonne National Laboratory
jacob@mcs.anl.gov

MS391

Adaptive Covariance Tuning in Hybrid Ensemble-variational Data Assimilation

Hybrid ensemble-variational (EnVar) data assimilation systems (DAS) provide a dynamic representation of the statistical properties of the model forecast error through low-rank ensemble-based corrections of a static (climatological) component of the background error covariance matrix (**B**-model) The EnVar approach has shown the ability of reducing the analysis errors and improving the forecast skill and has been adopted for operational implementation at various numerical weather prediction (NWP) centers. Key ingredients of the hybrid EnVar assimilation are the ensemble size and the selection of the ensemble members, the design of the covariance localization operator, and specification of the hybridization coefficients used to weight the static and the ensemble-based components of the **B**-model. The ability of the ensemble perturbations to capture coherent 'errors of the day' is crucial in determining the performance of the hybrid EnVar DAS. This talk presents new developments in both theoretical formulation and practical implementation to optimize the flow-dependent specification of the **B**-model using a derivative-based adaptive covariance tuning procedures. A proof-of-concept and illustrative results are presented together with practical aspects of implementation for NWP applications.

Dacian N. Daescu
Portland State University
Department of Mathematics and Statistics
daescu@pdx.edu

MS391

Prior and Posterior Inflation for Ensemble Filters: Theoretical Formulation and Application to Community Atmosphere Model

In this talk, temporally and spatially varying adaptive inflation algorithms for ensemble filters are presented. The algorithms provide an efficient variance correction strategy for the prior and posterior ensemble statistics. The derivation of the adaptive prior scheme is first presented following Bayes Theorem. Then, I will introduce a new posterior inflation scheme featuring an observation-impact removal strategy as a way to sequentially compute the inflation distribution. The usefulness of posterior inflation is investigated and compared to prior inflation with an 80-member ensemble in the Community Atmosphere Model. 6-hour forecasts of the atmospheric state, in the Troposphere and lower Stratosphere, are generated over the month of September 2010. GPS Radio Occultation refractivity observations in addition to wind and temperature data from aircraft, ACARS and satellites are assimilated. The following questions are addressed: What inflation scheme is more effective at handling sampling errors? When model bias dominates other error sources, which inflation strategy yields a better fit to the observations? Does inflating both the prior and the posterior ensemble perturbations help mitigate for different error sources? The evaluation of the inflation algorithms is assessed using observation space diagnostics. The inflation patterns are studied and correlations to observation network densities are examined.

Mohamad El-Gharamti, Kevin Raeder
National Center for Atmospheric Research
gharamti@ucar.edu, raeder@ucar.edu

Jeffrey Anderson
National Center for Atmospheric Research
Institute for Math Applied to Geophysics
jla@ucar.edu

Xuguang Wang
The University of Oklahoma, Norman, OK
xuguang.wang@ou.edu

MS391

Variational Data Assimilation with Model Errors

Contrary to the reality, the traditional 4DVar approaches assume that the underlying computational models are perfect. However, the 'weak constraint 4DVar' relaxes this 'perfect model' assumption by enforcing the model constraints 'weakly'. Solving the 4DVar in a weak constrained framework requires the necessary statistics associated with the model errors. In the past, only simplistic approaches have been proposed and they do not entirely account for spatio-temporal correlations in model errors. We have developed a Gaussian multivariate space-time statistical method that combines historical observations and imperfect computational models to obtain the necessary statistical information that can be deployed in a weak con-

strained 4DVar framework. This method accounts for spatio-temporal correlations in model errors and the preliminary results with small models are promising (we see a 60% improvement in forecast with this framework when compared with the traditional 4DVar formulations). We demonstrate our results on the shallow water model.

Vishwas Rao
Argonne National Laboratory
vhebbur@anl.gov

Emil M. Constantinescu
Argonne National Laboratory
Mathematics and Computer Science Division
emconsta@mcs.anl.gov

Julie Bessac
Argonne National Laboratory
jbessac@anl.gov

Adrian Sandu
Virginia Polytechnic Institute
and State University
sandu@cs.vt.edu

MS392

Simulating the Swimming Motion of *C. elegans* and Amoeboids in Viscoelastic Fluids via the Immersed Boundary Method

Microorganisms often find themselves moving through viscoelastic environments, as biological fluids frequently have a rich microstructure owing to the presence of large polymeric molecules. Research on the effect fluid elasticity has on the swimming kinematics of these organisms has been focused on those that move via cilia or flagellum. Experimentally, it has been shown that the nematode C. elegans, a model organism used to study such undulatory motion, swims slower as the Deborah number of the fluid is increased. Yet, this phenomenon has not been thoroughly studied with a fully resolved 3D simulation; furthermore, the effect fluid elasticity has on the swimming speed of organisms that move via euglenoid movement, such as E. gracilis, is completely unknown. In this talk, we discuss the simulation of the arbitrary motion of an undulating or pulsating swimmer in 3D, with the ability to specify any viscoelastic model for the surrounding fluid. To accomplish this task, we use a modified version of the Immersed Finite Element Method presented by Zhang et al. (2007). After validating C. elegans simulations against experimental data provided by Shen and Arratia (2011), we compare and contrast the effect fluid elasticity has on the swimming speed of these two model swimmers. Finally, we provide a physical explanation for the mechanism by which viscoelasticity affects the swimming speed through an analysis of the surrounding flow field.

Jeremy Binagia, Christopher Guido
Stanford University
jbinagia@stanford.edu, cjguido@stanford.edu

Eric S. Shaqfeh
Professor, Chemical Engineering & Mechanical Engineering
Stanford University, Stanford, CA
esgs@stanford.edu

MS392

Flagellar Swimming in Complex Fluids

Many important biological functions depend on microorganisms' ability to move in viscoelastic fluids such as mucus and wet soil. The effects of fluid elasticity on motility remain poorly understood, partly because, the swimmer strokes depend on the properties of the fluid medium, which obfuscates the mechanisms responsible for observed behavioral changes. We use experimental data on the gaits of *Chlamydomonas reinhardtii* swimming in Newtonian and viscoelastic fluids as inputs to numerical simulations that decouple the swimmer gait and fluid type in order to isolate the effect of fluid elasticity on swimming. In viscoelastic fluids, cells employing the Newtonian gait swim faster but generate larger stresses and use more power, and as a result the viscoelastic gait is more efficient. Furthermore, we show that fundamental principles of swimming based on viscous fluid theory miss important flow dynamics: fluid elasticity provides an elastic memory effect which increases both the forward and backward speeds, and (unlike purely viscous fluids) larger fluid stress accumulates around flagella moving tangent to the swimming direction, compared to the normal direction.

Robert D. Guy
Mathematics Department
University of California Davis
guy@math.ucdavis.edu

MS392

A Regularized Stokeslet Approach to the Motion and Behavior of Microscopic Swimmers

In the world of bacterial locomotion, motile organisms generate wakes and eddies that affect their local fluid environment. When many of them are present around each other, they interact with the flows generated by other organisms and form active suspensions whose complex flow structures play a significant role in fluid transport and mixing. In this talk, we present a reduced model for bacterial locomotion that describes their self-propelled motion in a low Reynolds number viscous incompressible fluid. The model is based on a particular limit of regularized Stokeslets with built-in asymmetry in order to produce a swimming direction. The result is a single-particle model of a swimmer that does not require special treatment of the self velocity due to the regularization, while allowing us to efficiently study the collective motion of bacteria. With this model, we are able to model pusher and puller organisms in a straightforward manner for both free-space and periodic domains. We will characterize the particle dynamics and discuss the diffusion of these particles as a function of the concentration density. We will then take advantage of the regularized Stokeslets framework to understand how active suspensions interact with viscoelastic structures, such as biofilms.

Alexander Hoover
University of Akron
ahoover1@uakron.edu

Ricardo Cortez
Tulane University
Mathematics Department
rcortez@tulane.edu

MS392

Dynamics of Elastic Filaments in Viscous Fluids

The deformations of flagella are important in motility of single- and multi-flagellated bacteria. Flagella are nearly inextensible, and existing numerical models have treated them as extensible filaments with a large extensional modulus. To consider the flexibility of filaments, we introduce extensible and inextensible elastic rod models with surface distribution of regularized Stokeslets and compare results to the previously described model of centerline distribution of regularized Stokeslets and rotlets. The elasticity of the flexible filaments is described by the Kirchhoff rod model where intrinsic curvature and twist determines cross-sectional forces and torques in the extensible rod model. For the inextensible model, the corresponding cross-sectional torques are calculated from the configuration of space curved filaments and forces are calculated from the viscous flow. We compare the accuracy of the centerline distribution of the Stokeslets and rotlets to the more accurate surface distribution of Stokeslets in modeling fluid-structure interaction. We also compare computational time and memory usage to illustrate the efficiency of the different approaches. We show that our inextensible approach is markedly more efficient than extensible models considered here.

Mehdi Jabbarzadeh, Henry Fu
University of Utah
jabbarzadeh.mehdi@utah.edu, henry.fu@utah.edu

MS393

Improving Stability of Numerical Methods for Recovering Governing Equations from Noisy Data

In modern science and engineering, coupled systems of nonlinear ordinary differential equations are frequently employed to represent a dynamical system of interest. However, in many applications the underlying governing equations of a dynamical system are either unknown or not fully understood and must be informed by experimental data. Our work focuses on the discovery of governing equations for such systems when the amount of data is limited and the data is corrupted with noise. We explore improvements to the Sparse Identification of Non-linear Dynamics (SINDy) method. Approximating solutions via SINDy often involves solving a large over-determined linear system via sparse regression, where the data is evaluated using a candidate dictionary of potential basis functions and gradients are either measured or approximated from data. Problems may arise when the condition number of the dictionary grows due to a poor choice of basis functions or a large number of data points. Ill-conditioning limits the accuracy of approximations considerably given the presence of noise in the data. We explore methods from statistical regression to address this problem.

Paul Diaz
University of Colorado, Boulder
pdiaz303@gmail.com

Alireza Doostan
Department of Aerospace Engineering Sciences
University of Colorado, Boulder
Alireza.Doostan@Colorado.EDU

MS393

A Bayesian Framework for Robust Decisions in the Presence of Unobserved Heterogeneity

Measures of decision risk based on the distribution of losses over a posterior distribution obtained via standard Bayesian techniques may not be robust to model misspecification. Standard statistical models assume homogeneity: all individual observables are subject to the same latent variables. Here we consider model misspecification due to *unobserved* sources of heterogeneity that are not readily captured by known covariates. For example, many predictive physical models contain conservation laws that are globally consistent, but those physical models may be misspecified due to embedded phenomenological models that are only locally consistent within a narrow range of parameters. We propose an approach for parameter inference that is robust towards unobserved heterogeneity. We *localize* the likelihood for a given parameter value by discounting observations that are inconsistent with the predictive model. We also scale the resulting localized likelihood by the fraction of the observations that are consistent. The result is a "generalized' posterior distribution that captures additional uncertainty due to model discrepancy that is revealed by observed data. We argue that this generalized posterior distribution is more appropriate for decision making, especially in a risk-averse setting.

Chi Feng, Youssef M. Marzouk
Massachusetts Institute of Technology
chifeng@mit.edu, ymarz@mit.edu

MS393

Representing Model Inadequacy in Interacting Systems with Constrained Stochastic Operators

In mathematical models of interacting systems, typically a small subset of species is of interest. For example, in modeling chemical reactions, one may be interested in only the time behavior of the chemical reactants and products, but not the intermediate species. As another example, biological or epidemiological models often track a much smaller subset of species than what is active in the true system of interest. It is thus natural to use reduced models, but this introduces error and uncertainty with respect to the high-fidelity, or detailed, model. This work proposes a method to account for the discrepancy between the detailed and reduced models using stochastic operators, whose formulations are constrained by available physical information.

Rebecca E. Morrison, Youssef M. Marzouk
Massachusetts Institute of Technology

rebeccam@colorado.edu, ymarz@mit.edu

MS394

Identifying 2D SSS Representations

Abstract not available.

Shivkumar Chandrasekaran
University of California, Santa Barbara
Department of Electrical and Computer Engieering
shiv@ece.ucsb.edu

MS394

Robust and Efficient QR Decomposition of HODLR Matrices with Applications

The efficient and accurate QR decomposition for matrices with hierarchical low-rank structures, such as HODLR and hierarchical matrices, has been a challenge. Existing structure-exploiting algorithms are prone to numerical instability as they proceed indirectly, via Cholesky decompositions or a block Gram-Schmidt procedure. For a highly ill-conditioned matrix, they either break down in finite-precision arithmetic or result in significant loss of orthogonality. Although these issues can sometimes be addressed by regularization and iterative refinement, it would be more desirable to have an algorithm that avoids these detours and is numerically robust to ill-conditioning. In this work, we propose such an algorithm for HODLR matrices. It achieves accuracy by utilizing Householder reflectors. It achieves efficiency by utilizing fast operations in the HODLR format in combination with compact WY representations and the recursive QR decomposition by Elmroth and Gustavson. Numerical experiments demonstrate that our newly proposed algorithm is robust to ill-conditioning and capable of achieving numerical orthogonality down to the level of roundoff error.

Daniel Kressner
EPFL Lausanne
Mathicse
daniel.kressner@epfl.ch

Ana Susnjara
EPF Lausanne
ana.susnjara@epfl.ch

MS394

Generalizations to the Superfast Divide-and-Conquer Eigenvalue Algorithm

As generalizations and extensions of the superfast divide-and-conquer eigensolver for symmetric hierarchically semiseparable (HSS) matrices, we look at how techniques relying on highly structured eigenvector representations and the fast multipole method (FMM) can be applied to generalized, non-symmetric, and sparse eigenvalue problems. All results will be supported by theoretical complexity and analysis, as well as by numerical test results. We will discuss applications of these algorithms to problems from computational physics.

James Vogel
Purdue University
vogel13@purdue.edu

MS394

On the Accuracy and Stability of Hierarchical Structured Approximate Factorizations

Fast approximate hierarchical structured factorizations have been shown to be very effective in solving large linear systems. On the other hand, the approximation and numerical accuracies have rarely been rigorously studied. In this work, we investigate the error propagation and numerical stability of some fast hierarchical direct solvers, including some dense structured solvers and sparse structured multifrontal solvers. For some algorithms, we also give concrete backward error bounds. We show that, in general, the solvers are not only faster, but also have much better stability as compared with the corresponding standard matrix algorithms. The error propagation factors only involve low-degree powers of the relevant off-diagonal numerical rank bounds and the logarithm of the matrix size. We also show that factorization-based structured solvers are usually preferred, while inversion-based ones may suffer from numerical instability. The analysis builds a comprehensive framework for understanding the accuracy and stability control of hierarchical rank structured methods. The error propagation patterns also provide insights into the improvement of other types of structured solvers and the design of new stable hierarchical structured algorithms. This talk includes joint work with Yuanzhe Xi and Zixing Xin.

Jianlin Xia
Purdue University
xiaj@purdue.edu

Yuanzhe Xi
University of Minnesota
yxi26@emory.edu

Zixing Xin
Department of Mathematics
Purdue University
zxin@purdue.edu

MS395

Summation-by-Parts Methods for Inverse Problems in Exploration Seismology

We consider the inverse problem of estimating parameters of the elastic wave equation. Given seismograms at point locations, we set up a misfit functional (MF) that measures the difference between simulated and recorded data. The gradient of the MF with respect to the large number of parameters can be efficiently computed by solving the adjoint PDE, which involves the adjoint of the elastic differential operator. Further, the restriction operator that localizes the solution to the seismometer point in the MF gives rise to a singular point source term in the adjoint PDE. To mimic this property discretely, the discrete restriction operator must be carefully chosen so that it also functions as a high-order approximation of a singular point source. The summation-by-parts (SBP) framework provides a rig-

orous approach for source terms located at boundaries and material interfaces, which is a common situation in exploration seismology. It also provides a recipe for constructing a difference operator whose adjoint approximates the adjoint differential operator stably. Combining the adjoint-consistent SBP discretization with proper restriction operators in the discrete MF yields a discrete adjoint equation that approximates the adjoint PDE to high order. This consistency has several benefits. One is super-convergence for the MF. Another is that the computed gradient is both the exact gradient of the discrete MF and a high-order approximation of the gradient of the continuous MF.

Martin Almquist
Stanford
malmquist@stanford.edu

Eric M. Dunham, Joseph Jennings
Department of Geophysics
Stanford University
edunham@stanford.edu, joseph29@sep.stanford.edu

MS395

Optimization of Free Parameters in Cubature Rules and Accuracy Equations to Develop Optimal Multi-Dimensional Summation-by-Parts Operators

Summation-by-parts (SBP) finite-difference operators have primarily been studied as one-dimensional operators extended to two and three dimensions using tensor products. The application of SBP operators on simplex elements, so called multi-dimensional SBP operators, is recent and has opened new research areas. We focus on investigating symmetrical cubature rules for diagonal-norm SBP operators on simplex elements with free parameters that are used to optimize the operators. The free parameters include the nodal locations, the cubature weights and the parameters to construct the surface integral operator as well as the derivative operator. The free parameters are used to generate cubature rules of higher order or to minimize the truncation error, spectral radius, or the Lebesgue constant, or combinations of these objectives. Other considerations include having or not having nodes on the facets, using different cubature rules on the facet to construct the surface integral operator, and having a diagonal directional surface integral operator by having the volume cubature nodes collocated with the facet cubature nodes. SBP operators with more than the minimum number of nodes per element are also considered to investigate the trade-off between higher computational cost and additional optimization from the extra nodes. The stability and accuracy of the optimized SBP operators are investigated using linear convection problems.

Andre Marchildon, David Zingg
University of Toronto Institute for Aerospace Studies
andre.marchildon@mail.utoronto.ca,
dwz@utias.utoronto.ca

MS395

Entropy Stable Discontinuous Galerkin Schemes on Moving Meshes for Hyperbolic Conservation Laws

In this talk, a split form nodal discontinuous Galerkin spectral element method (DGSEM) on moving meshes for hyperbolic conservation will be presented. The discrete derivative approximations in space are summation-by-parts (SBP) operators. Further, in order to avoid aliasing errors of the interpolation operator in the volume integrals, the split form DG framework is used. The SBP property and suitable two-point flux functions in the split form DG discretization are used to mimic results from the continuous entropy analysis. In particular, the semi-discrete moving mesh DGSEM is entropy conservative (EC) without assuming any exactness of quadrature in the variational forms. On moving meshes, the condition for EC deviates slightly from Tadmor's well known entropy condition. Such EC fluxes on moving meshes will be presented for the homogeneous shallow water equations and the compressible Euler equations. To remain valid in the presence of discontinuities, appropriate numerical dissipation is added to the EC fluxes. Then, the method becomes entropy stable such that the discrete mathematical entropy is bounded at any time by its initial data. Besides the entropy stability, the time discretization of moving mesh DGSEM will be discussed, where special care is taken to guarantee the method satisfies the free stream preservation property for any explicit Runge-Kutta method. Finally, the theoretical properties of the moving mesh DGSEM will be validated by numerical experiments.

Gero Schnuecke, Andrew R. Winters, Gregor Gassner
Mathematical Institute
University of Cologne
gschnuec@math.uni-koeln.de, awinters@math.uni-koeln.de, ggassner@math.uni-koeln.de

MS396

Unstructured Computations on Emerging Architectures

We present a complete architecture- and algorithmic-specific re-implementation in C++ of the PETSc-FUN3D application code, which includes the call backs for PETSc's nonlinear solver as well as the PETSc core routines. We characterize the performance of PETSc-FUN3D on various emerging architectures, including Intel Xeon Phi Knights Landing, Intel Xeon Scalable Processor – Skylake, Nvidia GPU (Pascal and Volta), and ARM (ThunderX2), as representatives of the compute units in contemporary leading supercomputers, identifying and addressing performance challenges without compromising the floating point numerics of the original code. We employ low- and high-level code optimizations and tuning in light of thread- and data-level parallelism, with the focus on strong thread scaling. Our approach is based upon a novel multi-level hierarchical workload distribution of data across both the thread contexts and the vector units within every hardware core. We analyze the application and its key computational routines on specific computing architectures, by which we develop certain performance metrics and models to bespeak the upper and lower bound of the performance on various back-end hardware platforms. These optimizations are expected to be of value for many other unstructured mesh partial differential equation-based scientific applications as

multi- and many-core architecture evolves.

Mohammed A. Al Farhan
KAUST
mohammed.farhan@kaust.edu.sa

David E. Keyes
King Abdullah University of Science and Technology (KAUST)
david.keyes@kaust.edu.sa

MS396

Heterogeneous Computing in ORNL's Experimental Computing Laboratory (excl)

Abstract not available.

Jeffrey S. Vetter
Oak Ridge National Laboratory
University of Tennessee, Knoxville
vetter@computer.org

Mehmet E. Belviranli, Steve E. Moulton
Oak Ridge National Laboratory
belviranlime@ornl.gov, moultonsa@ornl.gov

MS396

A Variable Neighbourhood Descent Heuristic for Conformational Search using a Quantum Annealer

Discovering the low-energy conformations of a molecule is of great interest to computational chemists with potential applications in in-silico materials design and drug discovery. In this paper, we propose a variable neighbourhood search heuristic for the conformational search problem. Using the structure of the molecule, neighbourhoods are chosen to allow for efficient optimization, and also allowing the application of a quantum annealer for this step of the iteration. The proposed method can adapt to the size and topology of the available quantum annealer chip through careful definition of neighbourhoods, making it scalable with respect to future hardware specifications.

Brad Woods, Dominic Marchand, Gili Rosenberg, Moslem Noori
1QBit
brad.woods@1qbit.com, dominic.marchand@1qbit.com, gili.rosenberg@1qbit.com, moslem.noori@1qbit.com

MS396

Early Application Successes on D-Wave Quantum Computers

Abstract not available.

Steve P. Reinhardt
D-Wave Systems, Inc.
sreinhardt@dwavesys.com

MS397

Entropy Stable Gauss Collocation Discontinuous Galerkin Methods

High order discontinuous Galerkin (DG) methods offer several advantages in the approximation of solutions of nonlinear conservation laws. However, these methods also tend to suffer from instability in practice, requiring filtering, limiting, or artificial dissipation to prevent numerical solutions from blowing up. Entropy stable schemes address a primary cause of this instability by ensuring that the solution satisfies a semi-discrete entropy inequality, even in the presence of discrete effects such as quadrature error. Most entropy stable DG schemes collocate the solution at Gauss-Legendre-Lobatto (GLL) points. The resulting schemes can be interpreted as classical summation-by-parts (SBP) finite differences in order to prove a discrete entropy inequality. In this talk, we show that, for more general quadratures, high order DG methods can be reinterpreted within a modified SBP framework using discrete projection and decoupled SBP operators. We utilize this equivalence to construct new entropy stable collocation schemes using Gauss-Legendre points on tensor product elements. Numerical experiments suggest that Gauss collocation schemes are significantly more accurate than GLL collocation schemes on warped meshes.

Jesse Chan
Rice University
jesse.chan@rice.edu

MS397

Efficient High-order Discontinuous Galerkin Methods on Modern Architectures

High-order finite element simulations on modern supercomputer architectures forces us to move away from standard sparse linear algebra to achieve high performance. The approach considered in this talk uses a representation of the bilinear form as a tensor product decomposition instead of using a matrix. This decomposition only requires values at quadrature points to be assembled (stored or computed on the fly) in order to apply the bilinear operator. Using this so-called partial assembly formulation, instead of assembling and applying a sparse matrix, results both in less computation and less data movement, therefore increasing the performance and reducing the time to solution. Applying these ideas in the context of discontinuous Galerkin methods (DG) offers even more opportunities for efficient computing due to the locality of DG. This talk explores the advantages and challenges of such algorithms on modern architectures with the MFEM finite element library, http://mfem.org.

Yohann Dudouit
LLNL
dudouit1@llnl.gov

Johann Dahm
IBM
Johann.Dahm@ibm.com

Steve Rennich
NVIDIA Corporation
srennich@nvidia.com

Tzanio Kolev, Veselin Dobrev, Jean-Sylvain Camier
Lawrence Livermore National Laboratory
kolev1@llnl.gov, dobrev1@llnl.gov, camier1@llnl.gov

MS397

A Moving Discontinuous Galerkin Finite Element Method with Interface Condition Enforcement for Viscous Flows

The Moving Discontinuous Galerkin Method with Interface Condition Enforcement (MDG-ICE) [A. Corrigan, A. Kercher, and D. Kessler, *A moving discontinuous Galerkin finite element method for flows with interfaces*, Tech. Report NRL/MR/6040–17-9765, December 2017, Submitted to IJNMF], which can accurately and stably compute flows with interfaces without relying on interface or shock capturing, is extended to viscous flows. The proposed method is shown to reduce the need for h-refinement obtain accurate approximate solutions in regions with sharp, finite, gradients, by enforcing the conservation law and the corresponding interface condition simultaneously while treating the discrete domain geometry as a variable. The number of degrees of freedom required to obtain a desired accuracy for problems with boundary-layer type solutions is investigated and reported for various Pclet numbers. Time accurate solutions of unsteady problems are obtained via a spacetime formulation, in which the unsteady problem is formulated as a higher dimensional steady spacetime problem. The proposed method is shown to accurately transport of viscous structures, e.g., viscous shocks, throughout the spacetime domain. Since MDG-ICE does not rely on numerical dissipation for stabilization, the viscous structures remain coherent as they are transported by the flow field.

Andrew Kercher, Andrew Corrigan, David Kessler
Naval Research Laboratory
andrew.kercher@nrl.navy.mil, andrew.corrigan@nrl.navy.mil, david.kessler@nrl.navy.mil

MS397

A High-order Fully Implicit Incompressible Navier-Stokes DG Solver

An efficient Navier-Stokes solver holds the key to solving general multiphysics problems that have an incompressible fluid model as a component. The Navier-Stokes equations are decoupled by projection methods and the spatial discretization is based on discontinuous Galerkin; therefore efficient linear and nonlinear solvers are required. To fully utilize computational resources, the implementation exploits tensor product structure, uses sum-factorization technique and matrix-free methods. Time-stepping in the viscous substep is done with diagonal implicit Runge-Kutta methods. For the pressure Poisson equation, in particular a hybrid AMG-DG is employed where on the finest level matrix-free block smoothers are applied and the low order subspace correction is solved with AMG. Results on the performance of the implementation are presented. The discretization is applied to direct numerical simulation of turbulent flows as well as analyzed in underresolved turbulence.

Marian Piatkowski
IWR, Heidelberg University
marian.piatkowski@uni-heidelberg.de

MS398

On Directives-based Refactoring of Legacy CFD Codes for Accelerator-based Systems

High performance computing (HPC) graphics processing units (GPUs) have high computational power relative to CPUs. Offloading scientific and engineering simulations to such GPUs has the potential to reduce analysis time. To harness this power, applications must be refactored to (i) expose suitably high parallelism in order to utilize the many GPU cores and (ii) to manage the data transfers between the separate GPU and host CPU memories. Managing the data transfers between the CPU and GPU (and between GPUs on the same node) is critical for attaining high efficiency. Both of these refactoring efforts can be costly (developers time) and invasive (significant code restructuring). Recent advances in NVIDIA GPU technology (e.g., NVLink) have improved the data transfer rates and Unified Virtual Memory (UVM) automatically manages data clones between the CPU and GPU. The later could significantly reduce the development cost of refactoring applications to run on GPUs. We investigated the performance of the NAS BT-MZ benchmark when offloaded to GPUs using OpenACC directives with and without explicit data management. In the latter case, we rely upon UVM for data management. We report performance results from a GPU-accelerated HPC system with multiple GPUs per node to determine any performance loss associated with UVM. Furthermore, we report on the development cost savings, measured in lines-of-code (LOC), when relying on UVM compared to using OpenACC data management directives.

Christopher P. Stone
HPCMP PETTT ACE On-site (Contractor)
chris.stone.ctr@engility.com

Bracy H. Elton
HPCMP PETTT / Engility
bracy.elton@engility.com

MS398

Unstructured-grid CFD Algorithms for a Many-core Landscape

In the field of computational fluid dynamics (CFD), the Navier-Stokes equations are often solved using an unstructured-grid approach to accommodate geometric complexity. Furthermore, turbulent flows encountered in aerospace applications generally require highly anisotropic meshes, driving the need for implicit solution methodologies to efficiently solve the discrete equations. We explore the transition of an unstructured-grid CFD code from a dense MPI model to shared-memory models suitable for a many-core landscape. Node-level studies of computationally intensive CFD kernels on traditional x86 architectures are included as well as NVIDIA GPUs. Scaling studies are

also performed using resources at several large supercomputing facilities.

Eric Nielsen
NASA Langley
eric.j.nielsen@nasa.gov

Aaron Walden
NASA
aaron.walden@nasa.gov

M Zubair
Department of Computer Science
Old Dominion University
zubair@cs.odu.edu

Justin Luitjens
NVIDIA Corporation
juitjens@nvidia.com

MS398

Design of a Portable, Multithreaded, and Generic C/C++ Library for SIMD Processing

Vectorization is an optimization technique used to increase processing throughput in computing systems by generating single-instruction multiple-data (SIMD) instructions. Modern multiprocessors support vectorization on a per-core basis, with vector units widening as newer generations are designed, such that SIMD features have become imperative for high-performance computing. Compilers tend to automatically employ vector transformations when code blocks have regular access patterns and minimal data dependencies. For cases where hand-tuning vectorization is effective, programmers use explicit vectorization via either inline assembly or vendor/compiler-specific intrinsic functions, thus reducing software portability and maintenance. To mitigate these issues, a SIMD library is being developed to make vector features accessible from high-level languages. This work presents the design of a generic SIMD library for explicit vectorization in C and C++ programs with on-demand multithreading capabilities. Our vector library provides common function-based and object-oriented interfaces across compilers and microprocessor architectures. Hybrid parallel implementations of pseudo-random number generators are presented as case studies that benefit from explicit vectorization.

Eduardo Ponce, Gregory D. Peterson
University of Tennessee
eponcemo@utk.edu, gdp@utk.edu

MS398

Accelerating Dynamic Load Balancing with GPUs

High performance simulations running on distributed memory, parallel systems require even work distributions with minimal communications. To efficiently maintain these distributions on systems with accelerators, the balancing and partitioning procedures must utilize the accelerator. This work presents algorithms and speedup results using OpenCL and Kokkos to accelerate critical portions of the EnGPar (hyper)graph-based diffusive load balancer. Focus is first given to basic graph traversal and selection procedures and then turns to ongoing work on the more complex graph construction and migration procedures.

Cameron Smith
Scientific Computation Research Center
Rensselaer Polytechnic Institute
smithc11@rpi.edu

Gerrett Diamond
Rensselaer Polytechnic Institute
diamog@rpi.edu

Lucas Davis
RPI
davisl3@rpi.edu

Mark S. Shephard
Rensselaer Polytechnic Institute
Scientific Computation Research Center
shephard@rpi.edu

MS399

FFTX and SpectralPack: A First Look

We present a first look at FFTX and SpectralPack. These software packages are developed as part of the DOE ExaScale effort by LBL, Carnegie Mellon University, and SpiralGen, Inc. We aim at translating the LAPACK/BLAS approach from the numerical linear algebra world to the spectral algorithm domain. FFTX is extending and updating FFTW for the exascale era and beyond while providing backwards compatibility. SpectralPack captures higher level spectral algorithms and their variants, including convolutions, Poisson solvers, correlations, and numerical differentiation approaches that translate to FFT calls. The SPIRAL code generation and autotuning system–now available as open source under a BSD/Apache license–underpins the effort to provide performance portability.

Franz Franchetti
Department of Electrical and Computer Engineering
Carnegie Mellon University
franzf@ece.cmu.edu

MS399

High Performance Eigensolver Exploiting an Online Tuning Mechanism

Modern numerical libraries need to be embedded a lot of performance tuning parameters to adapt the target platforms. In particular for heterogeneous or multi-node parallel system, the developers of numerical libraries have to write several code segments and algorithms, then combine and find the best combination of them, which yields almost best performance. In this study, we introduced a Bayesian-based tuning mechanism onto the EigenExa, which is a high performance eigensolver, in order to determine the optimal parameters automatically with online samplings.

Toshiyuki Imamura
RIKEN Advance Institute for Computational Science

imamura.toshiyuki@riken.jp

MS399

Test Matrices for Numerical Computations of Linear Systems and Eigenvalue Problems

This talk concerns test matrices for numerical linear algebra. If the exact solution of a problem such as a linear system and an eigenvalue problem is known in advance, we can obtain the error of a numerical solution exactly. It is useful for checking reliability of numerical algorithms and their comparisons. We propose methods that generate linear systems with the exact solutions on the basis of error-free transformation of floating-point numbers. When a user inputs a matrix A and a vector x, our methods first perturb A to a new matrix A'. Here, if $a_{ij} = a_{k\ell}$, then $a'_{ij} = a'_{k\ell}$ is satisfied, i.e., some matrix structure such as symmetricity is preserved. Next, we compute a vector b without rounding errors such that $b = A'x$, which means x is the unique and exact solution of a linear system $A'x = b$, provided that A' is nonsingular. In addition, we introduce methods that generate matrices with exactly known eigenvalues and singular values. The computing time for generating a problem is much less than that for obtaining an approximate solution of a problem.

Katsuhisa Ozaki
Shibaura Institute of Technology / JST
ozaki@sic.shibaura-it.ac.jp

Takeshi Ogita
Tokyo Woman's Christian University
ogita@lab.twcu.ac.jp

MS399

A Contour Integral Eigensolver with Divide-and-Conquer and Locking Techniques for Organic Material Simulations

Numerical simulations of organic materials post an eigenvalue problem that a certain of interior eigenvalues within given regions are computed. While eigenvalue solvers based on contour integral, such as FEAST or CIRR, are capable of finding the desired eigenpairs, the convergence performance may be slow if there are too many eigenvalues in the selected regions. To improve the convergence performance, we develop a divide-and-conquer scheme on top of the eigenvalue solver FEAST. The scheme partitions the given regions into sub-regions based on the estimated number of eigenvalues or domain knowledge. Consequently, the partition can lead to faster convergence in the proposed eigenvalue solver. For those eigenpairs that remain slow convergence in these sub-regions, we apply the locking technique so that the solver only updates the slowly convergent eigenpairs. Numerical experiments suggest the proposed schemes are promising. The proposed methods can be further parallelized to reduce the computation time.

Yuhsiang M. Tsai
National Taiwan University
Institute of Applied Mathematical Sciences
yhmtsai@gmail.com

Takeo Hoshi
Tottori University
hoshi@damp.tottori-u.ac.jp

Weichung Wang
National Taiwan University
Institute of Applied Mathematical Sciences
wwang@ntu.edu.tw

MS400

An Adaptive Least Squares Radial Basis Function Partition of Unity Method for PDEs

A least squares (LS) radial basis function (RBF) partition of unity method (PUM) for solving PDE problems was recently developed. In RBF-PUM, a global approximation to the solution is constructed as a weighted sum of local RBF approximations on overlapping patches. The LS approach allows us to give each patch an identical set of degrees of freedom, which can then be determined through collocation at a denser global node set. The main advantages of the method are high-order accuracy, numerical robustness, and computational efficiency. An inherent property of PUM is that the local approximation in each patch can be chosen with respect to the local properties of the solution. We explore this to derive an adaptive LS-RBF-PUM. In this talk, we discuss some of the challenges in how to handle Dirichlet boundary conditions, how to choose a convergent refinement strategy, and how to manage the adaptive data structures. Numerical results that compare the adaptive method with a the non-adaptive case are shown for a Poisson test problem. The adaptive method is shown to perform better in some cases.

Elisabeth Larsson
Uppsala University, Sweden
Department of Information Technology
Elisabeth.Larsson@it.uu.se

Danilo Stocchino
University of Padova
Dept. of Mathematics
danilostocchino@hotmail.it

Stefano De Marchi
University of Padova (Italy)
demarchi@math.unipd.it

MS400

Meshfree Simulation of Human Respiratory Muscles

The main objective of our research is to understand the functionality of a human diaphragm, the main muscle of the respiratory system. Its action affects the volume of the thorax cavity such that the lungs can inflate and deflate, enabling a human to breathe. The aim is to enable medical researchers to perform studies on ventilator induced diaphragm disease (VIDD). The diaphragm models in the existing simulation tools are not advanced enough to capture the processes that lead to VIDD. Our model is based on the nonlinear elasticity equations. The goal is to solve the equations on a 3D diaphragm geometry. That can ex-

hibit some numerical difficulties related to the small thickness of the tissue, compared to the overall size of the muscle. Localized radial basis function methods are chosen to discretize the equations in space and the quasistatic approach is used to advance the movement of the diaphragm in time. In this talk we will present the results of our first steps. These include the solutions of the static linear elasticity equations on a simple thin-plate geometry and also the full diaphragm geometry.

Nicola Cacciani
Karolinska Institutet
Department of Physiology and Pharmacology
nicola.cacciani@ki.se

Elisabeth Larsson
Uppsala University, Sweden
Department of Information Technology
Elisabeth.Larsson@it.uu.se

Igor Tominec
Division of Scientific Computing
Uppsala University
igor.tominec@it.uu.se

Pierre-Frederic Villard
University of Lorraine
pierre-frederic.villard@univ-lorraine.fr

MS400

A High-order Meshfree Semi-Lagrangian Method for Advection on Manifolds: Mass-conservation

We present a new meshfree semi-Lagrangian method for simulating advection equations on two-dimensional manifolds embedded in three-dimensional space. The method is based on localized radial basis function (RBF) interpolation using polyharmonic splines with polynomials, formulated in the tangent space of the surface. The semi-Lagrangian framework allows the method to avoid the use of any stabilization terms (such as hyperviscosity) during time-integration, thus reducing the number of parameters that have to be tuned. Additionally, time-steps that exceed CFL condition can be used without suffering from temporal instabilities. Finally, the method is formulated in Cartesian coordinates, avoiding any artificial singularities that arise from surface-based coordinates. We present a new technique for making the scheme mass-conserving and illustrate its accuracy and stability for several example problems on various manifolds. We also discuss the extension of these methods to more general advection-reaction-diffusion equations.

Grady B. Wright
Department of Mathematics
Boise State University, Boise ID
gradywright@boisestate.edu

MS400

Stable Computations with Flat Anisotropic Gaussians Using Hermite Polynomials

In recent years, a series of stabilization algorithms for the interpolation with flat RBFs, so-called RBF-QR algorithms, has been developed. The core of RBF-QR methods is to find an equivalent basis that spans the same space as the RBF basis but is well-conditioned. The first RBF-QR method has been introduced in [Fornberg, Piret, A stable algorithm for flat radial basis functions on a sphere, 2007] and is based on an expansion of RBFs in terms of spherical harmonics. For Gaussian kernel, an alternative basis can be derived based on the Mercer expansion [Fasshauer, McCourt, Stable evaluation of Gaussian radial basis function interpolants, 2012]. In this talk, we derive a basis similar to the one by Fasshauer and McCourt but based on the Hermite generating function (HermiteGF) expansion [Yurova, Kormann, Stable evaluation of Gaussian radial basis functions using Hermite polynomials, 2017]. Moreover, we extend the stabilization method to the case of the interpolation via anisotropic Gaussians. We present a theoretical framework of the anisotropic stabilization method and evaluate approximation properties of the new basis.

Anna Yurova
Max Planck Institute of Plasma Physics
Numerical Methods in Plasma Physics
anna.yurova@ipp.mpg.de

Katharina Kormann
Max Planck Institute for Plasma Physics
katharina.kormann@ipp.mpg.de

Caroline Lasser
Technische Universität München
classer@ma.tum.de

MS401

Advances in Machine Learned Potentials for Molecular Dynamics Simulation

Recent machine learning techniques allow emulation of quantum physics with stunning fidelity. For example, deep neural networks can now predict molecular properties with accuracy comparable to density functional theory, and approaching that of coupled cluster theory, at a tiny fraction of the computational cost. We present methods for building machine learned potentials that will enable large-scale and highly accurate molecular dynamics simulations, e.g., for chemistry, materials science, and biophysics. Key ideas are: 1. Encoding known physical properties and symmetries in the neural network architecture. 2. Active learning to dynamically augment the training dataset with new quantum calculations in regions where the machine learned model is uncertain. 3. Transfer learning, such that we first train on a large quantity of relatively low-fidelity data, and then perform some final training iterations using a smaller quantity of high-fidelity data.

Kipton Barros, Nicholas Lubbers, Justin S. Smith
Los Alamos National Laboratory
kbarros@lanl.gov, nlubbers@lanl.gov, jsmith48@lanl.gov

MS401

Toward the Systematic Generation of Hypothetical Atomic Structures: Neural Networks and Geomet-

ric Motifs

Materials discovery, a multidisciplinary process, now increasingly relies on computational methods. We can now rapidly screen materials for desirable properties by searching materials databases and performing high-throughput first-principles calculations. However, high-throughput computational materials discovery pipelines are bottlenecked by our ability to hypothesize new structures, as these approaches to materials discovery often presuppose that a material already exists and is awaiting identification. In contrast to this assumption, synthesis efforts regularly yield materials that differ substantially from the structures in databases of previously known materials. In this talk, we discuss strategies for generating hypothetical atomic structures using the concepts of geometric motifs (the recurring patterns of atoms in materials) and neural networks that can manipulate discrete geometry. We demonstrate the use of toy models to test the expressiveness and accuracy of these neural network operations.

Tess Smidt
Lawrence Berkeley National Laboratory
tsmidt@lbl.gov

MS401

Predicting Molecular Energies Using Machine-learned Electron Densities

Density functional theory (DFT) is a computationally affordable electronic structure method widely used to calculate the energy of molecular systems. However, using Kohn-Sham DFT in molecular dynamics simulations where the energy must be recalculated at every step is still a costly approach, particularly for periodic systems. We show that machine-learned density functionals that map potentials to energies via electron densities allow us to bypass the Kohn-Sham equations in 1D model and molecular systems. Molecular dynamics based on machine-learned densities for isolated molecules and crystalline solids illustrate the versatility of this method.

Leslie Vogt-Maranto
New York University
leslie.vogt@nyu.edu

MS401

Accelerating Molecular Dynamics with Multi-scale Neural Networks

Neural networks have fueled several breakthroughs in machine learning tasks in the recent years; however, their application to traditional scientific computing problems have been somewhat limited. In this talk we aim to bridge the gap between both realms. We provide a new multi-scale framework to build approximation to non-linear maps using deep neural networks. We show how to approximate the Kohn-Sham map arising from Density Functional Theory using a multi-scale neural network, and how this network is able to accelerate molecular dynamics simulations in the Born-Oppenheimer approximation. Joint work with Yuwei Fan (Stanford), Lin Lin (UC Berkeley) and Lexing Ying (Stanford)

Leonardo Zepeda-Nunez
Lawrence Berkeley National Lab
lzepeda@math.mit.edu

Lin Lin
University of California, Berkeley
Lawrence Berkeley National Laboratory
linlin@math.berkeley.edu

Lexing Ying
Stanford University
Department of Mathematics
lexing@stanford.edu

Yuwei Fan
Stanford University
ywfan1989@gmail.com

PP1

Design Exploration of Fuel Injectors Based on Reduced Order Modeling

In the following work, we use model order reduction techniques in order to make a design exploration of the non-reacting turbulent fluid flow in a fuel injector of aircraft engines. This study is shown for the lean premixed burner PRECCINSTA test case that has been studied in the combustion community to validate large eddy simulation models. We are proposing an adaptative geometrical reduced order model of the unsteady and incompressible Navier-Stokes equations, that we solve efficiently with respect to a collection of a priori designs of the injector. This framework is based on a prediction step of the global aerodynamic field using the Gappy-POD on a local high-fidelity solution associated with a new design, and a correction step by extrapolation using the POD-Galerkin projection of the governing Navier-Stokes quations. We precise that the projection step is done thanks to a POD basis enriched with dissipative modes that enable physical stabilization of the ROM. This combination between data reconstruction techniques and physics based ROM, enables a good recognition of the geometrical aerodynamic field. Then, we compute a quantity of interest (for example the swirl number) associated with each reduced order prediction. In a final step, we run the high-fidelity large eddy simulation associated with the optimal design, to certify our novel geometrical reduced order strategy. Finally, we mention that the speed up associated with this reduced order strategy is equal to 55.

Nissrine Akkari
SAFRAN
nissrine.akkari@safrangroup.com

Fabien Casenave, Christian Rey
Safran
fabien.casenave@safrangroup.com,
christian.rey@safrangroup.com

Vincent R. Moureau
CORIA CNRS

vincent.moureau@coria.fr

PP1

Parallel Implementation of a Monolithic Li-ion Battery Model using Fenics

With the increase of battery-powered electronics (from cell phones to cars), the demand for higher performing Li-Ion batteries has sky-rocketed. Two ways to increase performance is to increase the batterys ability to charge faster or to increase its charge density. Either method can be explored with the aide of highly accurate simulations. A Li-ion battery contains three domains: two electrodes (anode and cathode) and the electrolyte that connects them. The mathematical model describes the concentration and potential in all three domains at the microscopic scale. The transport of Lithium ion at the interface of an electrode and the electrolyte are governed by the nonlinear Butler-Volmer equation. One solution method for this model is to separate the three physical domains into three computational domains creating a staggered scheme. This way, the concentration and potential is resolved in one domain assuming the other two domains are fixed. However, this requires an additional level of iteration, to reach consistency between all three domains. This research proposes a monolithic method that solves on all three domains simultaneously to minimize the number of iterations per time step. We compare the performance of this monolithic method with the staggered domain method. In addition, we explore the parallelism of both methods using the FEniCS and PETSc libraries.

Jeffery M. Allen
University of Colorado at Boulder
Jeffery.Allen@nrel.gov

Justin Chang
NREL
justin.chang@nrel.gov

Francois Usseglio-Viretta
National Renewable Energy Lab
francois.usseglioviretta@nrel.gov

Peter Graf
NREL
peter.graf@nrel.gov

PP1

An Edge-preserving Method for Joint Bayesian Inversion with Non-Gaussian Priors

In this work we develop a joint Bayesian inversion method with the goal of inferring for multiple parameter fields using noisy observations of systems governed by multiple PDE models. For instance, in geophysical exploration applications both acoustic and electromagnetic wave propagation models can be used for estimating the subsurface parameters. In this and many other cases, the unknowns can be viewed as different properties of a single medium and, therefore, can be spatially correlated. We incorporate such similarities between the unknown fields through a joint prior distribution in order to improve the parameter reconstruction quality as opposed to inverting for each field independently. The main challenge of this problem is to design a joint inversion method which is well-defined in the infinite-dimensional setting and at the same time is suitable for capturing edges, or discontinuities, in the material properties. These restrictions forbid the use of vectorial total variation (VTV) priors, as they are only well-defined in finite-dimensional case, and make Gaussian priors infeasible, as they fail to detect jumps in the parameters. We propose an alternative approach, which leads to a well-defined prior distribution and is suitable for edge-preserving joint inversion, as demonstrated through a series of numerical experiments.

Ilona Ambartsumyan
The University of Texas at Austin, ICES
ailona@austin.utexas.edu

Tan Bui-Thanh, Omar Ghattas
The University of Texas at Austin
tanbui@ices.utexas.edu, omar@ices.utexas.edu

Eldar Khattatov
ICES, The University of Texas at Austin
ekhattatov@austin.utexas.edu

PP1

Feasibility Study of Multilevel Schur Complement Methods for Sparse Linear Systems

Many sparse linear solver methods, such as Domain-Decomposition and Hybrid-Solvers use Schur Complements in order to simplify and speed up computation. These methods base this decision on the well-known property that the condition number of the Schur Complement is much smaller than that of the whole linear system. The better-behaved condition number allows for solutions using various iterative methods, e.g., Conjugate Gradient methods, that converge quickly. In computation, this section is traditionally solved on a single processor or in parallel over a small number of processors. Due to increasing thread and core counts on systems like GPUs and Intel KNL, more concurrency needs to be extracted to solve a Schur Complement system effectively. A natural approach would be to use ordering methods that decompose the Schur Complement into multiple levels (smaller systems) to be solved. These levels would be determined in likeness to Nested-Dissection or Hierarchical Interface orderings. The problem is that little to no information is understood on the conditioning of these smaller systems. Because of this problem our work focuses on performing extensive empirical tests to better understand the conditioning of these levels and how they affect multiple solving methods that use Schur Complements.

Gregory S. Bolet, Joshua Booth
Franklin & Marshall College
gbolet@fandm.edu, jbooth1@fandm.edu

PP1

Discontinuous-Galerkin Galerkin-differences for

the Wave Equation in Second-order Form

The construction of a high-order accurate upwind scheme for the second order wave equation is discussed. The scheme is an energy-stable upwind difference scheme constructed using the discontinuous Galerkin framework. Lagrange polynomial interpolation over a computational grid is used to construct basis functions specifically chosen to be discontinuous at the midpoints between grid points. The upwind numerical flux is determined by the exact solution to a Riemann problem at these midpoints. This framework allows for the construction of high-order accurate upwind schemes for the second order wave equation. The resulting one dimensional scheme is extended to higher dimensions by way of tensor product, which allows for the efficient inversion of the mass matrix, so that the computational cost of the scheme grows linearly in the number of grid points. The spectral radius of the discretization matrix is small, and grows slowly as the order increases, so that high-order schemes remain stable.

Benjamin B. Buckner, Jeff Banks
Rensselaer Polytechnic Institute
bucknb2@rpi.edu, banksj3@rpi.edu

T Hagstrom, K Juhnke
Southern Methodist University
thagstrom@mail.smu.edu, kjunkhe@mail.smu.edu

PP1

A Kronecker Product Implementation of Density Matrix Renormalization Group

Density Matrix Renormalization Group (DMRG) is the preferred method for nanoscale modeling of strongly correlated materials. One goal of DMRG is computation of the lowest eigen-vector of Hamiltonian (ground-state) defined on a N-site lattice. The full Hamiltonian can be written as Kronecker products of operators

$$H_{\text{full}} = H_L \otimes I_R + I_R \otimes H_R + \sum_{k=0}^{K} C_L^k \otimes C_R^k \ . \qquad (6)$$

One key computational kernel is matrix-vector multiplication of target sparse Hamiltonian matrix in Lanczos algorithm. By grouping the admissible states by their left and right quantum numbers, the Hamiltonian matrix can compactly expressed as sums of Kronecker products of smaller dense matrices. Our implementation exploits this property of Kronecker products to have an order of magnitude saving in memory and computation. By careful placement of these Kronecker matrices in memory, acceleration on GPU can be implemented using batched GEMM operations in MAGMA library. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing, Scientific Discovery through Advanced Computing (SciDAC) program, under contract number DE-AC05-00OR22725.

Eduardo F. D'Azevedo
Oak Ridge National Laboratory
Computer Science and Mathematics Division
dazevedoef@ornl.gov

Wael R. Elwasif, Arghya Chatterjee
Oak Ridge National Laboratory
elwasifwr@ornl.gov, chatterjeea@ornl.gov

Gonzalo Alvarez
Computational Sciences & Engineering Division
Oak Ridge National Laboratory
alvarezcampg@ornl.gov

PP1

Towards the Understanding of Multirate Schemes on Adaptive Mesh Refinement Grids

A wide range of scientific and engineering problems requires computer simulation of complex large-scale systems of differential equations that evolve over a wide range of spatial and temporal scales. Such so-called stiff systems are both ubiquitous in applications and also notoriously difficult to solve numerically. An important example of such systems is large-scale reactive flows encountered in a variety of industrial and scientific applications such as flame propagation or supernova. To alleviate numerical difficulties associated with the spatial multiscale nature of such flows, adaptive mesh refinement (AMR) is commonly employed to limit the use of high spatial resolution, and consequently cost, only to regions with fine spatial structure. On the other hand, for time integration, multirate methods have been developed to solve problems with multiple time scales. These schemes are able to advance each forcing term at its characteristic time scale and thus significantly improve efficiency. In this project, we investigate the possibility of combining these two methods (AMR and multirate time integration) in order to get more efficient and adaptive schemes. The results presented in this poster assess the feasibility of this approach and present an application using the AMReX framework.

Valentin Dallerit
University of California Merced
vdallerit@ucmerced.edu

John Loffeld
Lawrence Livermore National Laboratory
loffeld1@llnl.gov

PP1

Sensitivity-driven Dimension-adaptive Sparse Stochastic Approximations in Linear Gyrokinetics

The quantification of uncertainty in complex real-world world applications, such as plasma microturbulence simulation is very challenging, mainly due to the large number of stochastic parameters and associated computational cost. To overcome these challenges, we propose a novel structure-exploiting, adaptive sparse approximation methodology for quantifying uncertainty in these types of problems. We leverage local Sobol' decompositions and compute partial variances for all input directions as well as for their interaction. Afterward, we introduce a sensitivity scoring system that quantifies the local contribution of each subspace to the local variances, thus exploiting the local structure of the problem at hand. We remark that our approach is generic, i.e., it can be formulated in terms of arbitrary approximation operators and point distributions. In particu-

lar, in this work we consider sparse interpolation and sparse pseudo-spectral projection. We employ the proposed approach in two problems, (i) a modified benchmark, in which we consider eight stochastic inputs, and (ii) a more realistic test case with three or 12 uncertain input parameters. In addition, we compare the proposed method with a classical adaptive strategy. The results show that our strategy requires fewer grid points than the classical approach, while preserving the overall accuracy. Moreover, the underlying uncertainty analysis proves to be an useful tool to aid the physical interpretation of the results.

Ionut-Gabriel Farcas
Technical University of Munich
Chair of Scientific Computing in Computer Science
farcasi@in.tum.de

Tobias Goerler
Max Planck Institute for Plasma Physics
Tokamak Theory division
tobias.goerler@ipp.mpg.de

Hans-Joachim Bungartz
Technical University of Munich, Department of Informatics
Chair of Scientific Computing in Computer Science
bungartz@in.tum.de

Frank Jenko
Max Planck Institute for Plasma Physics
Tokamak Theory division
frank.jenko@ipp.mpg.de

Tobias Neckel
Technical University of Munich
neckel@in.tum.de

PP1

Tissue Geometry May Govern Lung Branching Mode Selection

Lung branching morphogenesis proceeds in three stereotyped modes (domain, planar, and orthogonal branching). Much is known about the molecular players, including growth factors such as fibroblast growth factor 10 but it is unknown how these signals could actuate the different branching patterns. With the aim of identifying mechanisms that may determine the different branching modes, we developed a computational model of the epithelial lung bud and its surrounding mesenchyme. We studied transport of morphogens and localization of morphogen flux at lobe surfaces and lobe edges. We find that a single simple mechanism is theoretically capable of directing an epithelial tubule to elongate, bend, flatten, or bifurcate, depending solely on geometric ratios of the tissues in the vicinity of a growing tubule tip. Furthermore, the same simple mechanism is capable of generating orthogonal or planar branching, depending only on the same geometric ratios.

Uduak Z. George
San Diego State University
ugeorge@sdsu.edu

Sharon Lubkin
North Carolina State University
lubkin@ncsu.edu

PP1

A Fixed Mesh Method With Immersed Finite Elements for Solving Interface Inverse Problems

Abstract: I will present a new fixed mesh algorithm for solving a class of interface inverse problems for the typical elliptic interface problems. These interface inverse problems are formulated as shape optimization problems whose objective functionals depend on the interface shape. Both the governing partial differential equations and objective functionals are discretized optimally by an immersed finite element (IFE) method on a fixed mesh independent of interface location. The formula for the shape sensitivities of the discretized objective functions is derived within the IFE framework that can be computed accurately and efficiently through the discretized adjoint method. We show its applications by a group of representative interface inverse/design problems.

Ruchi Guo
Virginia Polytechnic Institute and State University
The Department of Mathematics
ruchi91@vt.edu

Tao Lin
Virginia Tech
tlin@vt.edu

Yanping Lin
The Hong Kong Polytechnic University
Department of Applied Mathematics
yanping.lin@polyu.edu.hk

PP1

Performing Derivative-free Optimization using Active Subspaces and Random Walks

We investigate and propose several algorithms that optimize functions $f(x)$ where x is a high-dimension vector in $\mathbb{R}^n$, f is noisy, and we lack the analytical gradient, ∇f. In particular, we consider dimension reduction of f through active subspace analysis, which can be done without access to ∇f. If the action of f can be described in fewer variables than n (the dimension of x) perhaps we only need to optimize f over those variables. We also consider a derivative-free optimization technique based on taking random walks to points in the domain and forming finite-difference approximations with function evaluations at those points. Finally, we blend both techniques, analyzing the active subspace of f and performing derivative-free optimization of f together.

Jordan R. Hall
University of Colorado Denver
Department of Mathematical and Statistical Sciences
jordan.r.hall@ucdenver.edu

PP1

An Isofrequency Remapping Scheme for Harmonic

Balance Methods

In the discretization of periodic parabolic partial differential equations (PDEs), for example, in the modeling of transient semiconductor physics with periodic boundary conditions (BCs) arising from coupling to an electronic circuit, the harmonic balance (HB) method may be applied to analyze the equations in the frequency domain. A discrete Fourier transform is applied in order to obtain the HB residual equations. However, the Fourier transform is prohibitively expensive when too many or too great frequencies $\vec{\omega} = (\omega_1, \ldots, \omega_s)$ are applied at the BC. In the literature, frequency remapping schemes $\vec{\omega} \mapsto \vec{\eta}$ have been proposed for one- or two-tone BC problems. However, a recognized oversight in the existing methods lies in the lack of a means to handle degenerate frequencies. Degenerate frequencies are non-unique linear combinations of the fundamental frequencies which arise from the non-linear terms of the PDE. For instance, if there are two equal linear combinations $\vec{\alpha}\cdot\vec{\omega}$ and $\vec{\beta}\cdot\vec{\omega}$, then a robust frequency remapping must yield $\vec{\alpha}\cdot\vec{\eta} = \vec{\beta}\cdot\vec{\eta}$. In this work, we present a method to obtain an optimal frequency remapping scheme so that large Fourier transforms are avoided, and the degenerate frequencies are respected, even for many tones. Our approach casts the frequency remapping scheme as the solution to an integer linear programming problem.

Andy Huang
Sandia National Laboratories
ahuang@sandia.gov

PP1

Contour Integration and Moments for the Solution of Large Eigenproblems

BEAST-M and z-Pares are both scalable eigensolvers for the solution of generalized interior eigenproblems. In this work, we will compare and contrast the solvers in the solution of large-scale electronic structure calculations. Both solvers rely on Sakurai-Sugiura methods, using contour integration with moments to construct the approximate subspace for the desired eigenspace. They use distributed memory and utilize multiple levels of parallelism. The problems considered will be large and sparse, meaning the most expensive component of the solvers is the solution of the large, ill-conditioned linear systems arising at each quadrature point in the contour integral. To minimize computational expense, the BEAST-M solver relies on an iterative approach and utilizes various adaptive strategies, whereas the z-Pares solver is typically designed to converge within one iteration. We compare the solvers' scalability and performance as well as identifying common bottlenecks.

Sarah Huber
University of Wuppertal
shuber@math.uni-wuppertal.de

Yasunori Futamura
University of Tsukuba
futamura@cs.tsukuba.ac.jp

Martin Galgon
Yniversity of Wuppertal
galgon@math.uni-wuppertal.de

Bruno Lang
Bergische Universitat Wuppertal
lang@uni-wuppertal.de

Tetsuya Sakurai
Department of Computer Science
University of Tsukuba
sakurai@cs.tsukuba.ac.jp

PP1

3D Deep Learning in High Resolution Point Clouds using the Grid-Octree Data Structure

In recent years there has been an immense expansion of deep learning methods aiming to solve some of the common computer vision tasks on 2D images. However, there are not many deep convolutional network architectures capable of working with high-resolution 3D point clouds as input. This is due to the fact that in 3D the memory requirements and the number of computations increase cubically with respect to the input resolution. In order to construct 3D convolutional networks which have powerful expressiveness and at the same time can be applied directly to the high-resolution representation of the whole input 3D volume, we opt to use the recently published OctNet library. The OctNet library aims to exploit the fact that 3D data is sparse in its nature by hierarchically partitioning the 3D space into the grid of shallow octrees. In this work, we show that by using the proposed input representation we are indeed focusing memory and computation to the dense regions and are therefore able to train deep convolutional networks for tasks such as semantic segmentation of detailed 3D point clouds (i.e. comprising of up to million points). Using the standard GPUs we are able to feed deep network with the whole 3D scene discretized with high resolution of up to 256^3, what was considered not to be feasible up to recently.

Ivana Jovanovic
Technical University Munich
ivana.jovanovic@tum.de

PP1

A Two-stage Method for Spectral-Spatial Classification of Hyperspectral Images

This poster proposes a novel two-stage method for the classification of hyperspectral images. Pixel-wise classifiers, such as the classical support vector machine (SVM), consider spectral information only; therefore they would generate noisy classification results as spatial information is not utilized. Many existing methods, such as morphological profiles, superpixel segmentation, and composite kernels, exploit the spatial information too. We propose a two-stage approach to incorporate the spatial information. In the first stage, an SVM is used to estimate the class probability for each pixel. The resulting probability map for each class will be noisy. In the second stage, a variational denoising method is used to restore these noisy probability maps to get a good classification map. Our proposed method effectively utilizes both spectral and spatial infor-

mation of the hyperspectral data sets. Experimental results on three widely used real hyperspectral data sets indicate that our method is very competitive when compared with current state-of-the-art methods, especially when the inter-class spectra are similar or the percentage of the training pixels is high.

Raymond H. Chan
The Chinese Univ of Hong Kong
Department of Mathematics
rchan@math.cuhk.edu.hk

Kelvin Kan
Department of Mathematics, Emory University
kai.fung.kan@emory.edu

Mila Nikolova
CMLA, CNRS, ENS Cachan
nikolova@cmla.ens-cachan.fr

Robert Plemmons
Wake Forest University
plemmons@wfu.edu

PP1

Efficient Black Box Optimization for Thin Film Solar Cell Design using Transfer Learning

In this study a novel transfer learning aided design optimization framework is proposed which improves the accuracy and speed of surrogate based optimization. The proposed method is then used in a case study where thin film multilayer solar cells are optimized for maximum solar-to-electric performance. We assume that at least one optimization has taken place (base case). The aim is to repeat optimization for a structure with different material choices (transfer cases). Knowledge gained from the base case is transferred to the new problem by means of neural network layers. Improvement in the prediction performance due to transfer learning is studied using the out-sample mean squared error metric. The work has therefore two novelties: 1) proposing a neural network transfer learning based optimization framework for solving complex optimization problems and, 2) using the proposed method to design a multi-layer thin film solar cell structure. The results show that the mean squared error of the surrogate model is reduced by 2-3 times using the transfer learning approach, even though a smaller number of training data is used. Furthermore, the results obtained using transfer learning are in a good agreement with the direct optimization results and previous results.

Mine Kaya
Texas A&M University
minekaya@tamu.edu

Shima Hajimirza
Texas A&M University
Texas A&M University
shima.hm@tamu.edu

PP1

Vector Potentials for Incompressible Flows in Bounded Domains

Consider the following problem from fluid dynamics: given the vorticity $\vec{\omega}$, find the velocity field $\vec{u}$ that satisfies $\nabla \cdot \vec{u} = 0$ and $\nabla \times \vec{u} = \vec{\omega}$ in a given bounded domain Ω subject to boundary conditions $\vec{u} \cdot \vec{n} = g$ on $\partial\Omega$. It has long been known that the solution to this problem can be written in terms of a vector potential as $\vec{u} = \nabla \times \vec{\psi}$, where $\vec{\psi}$ is the solution to the Poisson equation $-\Delta\vec{\psi} = \vec{\omega}$. Solving this problem using a stream function has the great advantage that the resulting velocity field is *exactly divergence free*, even if one is only given an approximation of $\vec{\psi}$. This formulation is however not commonly used in practice: it has been considered too difficult to obtain boundary conditions for $\vec{\psi}$. In this poster we present a surprisingly simple numerical scheme and results to tackle this problem in general domains. Recently, the study of pressure–robust finite element formulations for the Navier–Stokes equations has received considerable interest. The vector potential can be shown to be the underlying reason for the intruiging performance of these elaborate schemes. We hope that our results using simple Lagrangian elements on the other hand, encourage more research in the velocity–vorticity formulation of the three-dimensional Navier–Stokes equations.

Matthias Kirchhart
MathCCES
RWTH Aachen University
kirchhart@mathcces.rwth-aachen.de

Yutaka Kobayashi
Department of Mechanical Engineering
Keio University
ykobayashiusa@gmail.com

PP1

NUMO A New Non-hydrostatic Ocean Model for Fjord Circulation and Ice-sheet/Ocean Interaction

Ice-sheet/ocean interaction in Greenland is one of the key outstanding challenges in climate modeling, yet present-day climate models are not able to resolve fine-scale processes in the fjords. This is due to orders of magnitude difference in spatial scales between the open ocean (1000km) and fjord (¡1km) as well as small-scale processes at the glacier terminus, complicated bathymetry and coastline. The NUMO project aims to address the issues arising in fjord modeling by developing a three-dimensional, unstructured mesh model using high-order Galerkin methods. Local non-conforming refinement allows for resolving particular areas of interests (i.e. ice-sheet / ocean interface), while maintaining coarser resolution in the rest of the domain. We will present progress of the NUMO project, with the goal to develop a non-hydrostatic ocean model able to resolve ice-sheet / ocean interactions, circulation within the fjord and exchanges with the ocean. Recent developments include a streamlined modeling process from GIS bathymetry information (i.e. BedMachine 3.0), through 3D geometry modeling, meshing, and simulations.

Michal Kopera
Boise State University
michalkopera@boisestate.edu

Wieslaw Maslowski, Frank Giraldo
Naval Postgraduate School
maslowsk@nps.edu, fxgirald@nps.edu

PP1

Linking Mechanistic Infectious Disease Models to Genomic Surveillance Data Reveals Transmission Properties

Genomic analysis is quickly becoming an indispensable tool in computational epidemiology. Genetic surveillance data can reveal how a disease evolves as it spreads and as it reacts to intervention efforts, which are important factors for eliminating a disease without introducing drug resistances. For example, phylogenetic studies of polio transmission have demonstrated how precise knowledge of the genetic variation in an endemic region can be used to deliver targeted responses to specific strains, increasing the effectiveness of the ongoing eradication campaign. As higher quality data become available, mechanistic genetic models coupled with infectious disease models have the ability to test hypotheses about specific transmission properties and sampling strategies. This is especially relevant for eliminating diseases like malaria. Understanding the interplay between parasite evolution, importation dynamics due to human movement, and complexity of infections will better enable operational decision-making. We present preliminary results from a new genetic modeling infrastructure that can trace the transmission pathways of diseases with higher granularity and produce synthetic genomic data. Specifically, we allow the underlying mechanistic model to vary, testing hypotheses about the transmission and evolution of the malaria parasite. We also discuss how such a tool has the potential to drive decision-making in elimination settings.

Albert Lee
Institute for Disease Modeling
allee@idmod.org

PP1

A High-order Cross-platform Incompressible Navier-Stokes Solver via Artificial Compressibility with Application to Submarine Hydrodynamics

Modern hardware architectures such as graphics processing units (GPU) and manycore processors are characterised by an abundance of compute capability relative to memory bandwidth. Efficiently utilising such architectures to solve projection-based incompressible Navier-Stokes formulations is known to be a challenge, due to the sparse linear problem arising from the Poisson equation. An alternative approach to solving the incompressible Navier-Stokes equations is the method of artificial compressibility (AC). When adopting explicit dual time stepping, and using high-order flux reconstruction (FR) in space, AC can be cast as combination of pointwise operations and matrix-matrix multiplications, which can have high arithmetic intensity. In this poster, a high-order cross-platform incompressible NavierStokes solver, via artificial compressibility and dual time stepping, is presented. The solver supports a range of hardware architectures via platform-unified templating that generates and compiles CUDA, OpenCL and C/OpenMP code at runtime. It also implements existing and novel explicit convergence acceleration techniques, including P-multigrid, locally adaptive pseudo-time-stepping, and FR-optimised maximum stability-schemes. To demonstrate the technology at scale, the solver is applied to a submarine aft-body simulation using over 1000 GPUs.

Niki A. Loppi
Imperial College London
n.loppi15@imperial.ac.uk

Freddie Witherden
Department of Aeronautics and Astronautics
Stanford University
freddie.witherden08@imperial.ac.uk

Peter E. Vincent
Department of Aeronautics
Imperial College London
p.vincent@imperial.ac.uk

PP1

Fast Manifold Updates for Non-stationary Data Streams

How can we maintain a low dimensional manifold approximation when faced with new data from a high dimensional stream? We revise local linear models with incremental SVD (singular value decomposition) and form an updateable partition definition with cover trees. High dimensional data may have low dimensional structure embedded in the coordinate space. Given a set of data $X \in \mathbb{R}^D$ and a metric δ, we model this assuming that data lie close to a d dimensional manifold, $d << D$. To construct a manifold we partition $\mathbb{R}^D$ and fit rank d principal component analysis (PCA) models of subsets. We maintain an active model of a data stream by updating with new data and downdating expired data. Our work updates and downdates many local data given only their principal components. We find the partition subset of new data (adding new partition centers as needed) then use incremental SVD methods locally to update new data, downdate old data, and recenter the PCA. Given a scale $i \in \mathbb{Z}$, we partition initial data X by exploiting cover tree invariants to locate $P \subset X$ st $\forall x \in X$ $\exists p \in P$ where $\delta(x,p) < 2^i$. To identify the partition membership of $y \in \mathbb{R}^D$, we find the nearest $p \in P$. If no p satisfies $\delta(p,z) \geq 2^i$ for a new datum z we append z to P as a center. Our poster presents incremental algorithms and illustrates their performance on a time-varying data set.

Anthony Marcich, Benjamin W. Ong
Michigan Technological University
ajmarcic@mtu.edu, ongbw@mtu.edu

PP1

Calibration, Propagation, and Validation of Model Discrepancy Across Experiments

Model inadequacy due to model form error remains a concern in all areas of mathematical modeling, despite continuing advances in statistical inversion and the availability of

cost-effective high-performance computing infrastructure. The Bayesian paradigm naturally integrates uncertainties from both experimental data and model formulation, including initial or boundary conditions, model form and parameters, and numerical approximation. However, model error is unavoidable in many situations due to incomplete understanding of the underlying physics, likely in addition to large and possibly poorly characterized uncertainties in calibration and validation data. Put simply, infinite amounts of data may still result in inadequate models. Model correction techniques have the potential to increase the range of applicability and enhance the predictive power of models that suffer from model-form error. Calibrating a discrepancy model requires careful consideration regarding problem-specific formulation, parameter estimation, and uncertainty quantification. Furthermore, the validity of the original physical model, the inadequacy model, and the combined model for the prediction of quantities of interest remains in question. A generalized approach and implementation of model discrepancy detection, construction, and propagation into a predictive setting is presented in the context of Bayesian model calibration.

Kathryn Maupin
Sandia National Laboratories
kmaupin@sandia.gov

Laura Swiler
Sandia National Laboratories
Albuquerque, New Mexico 87185
lpswile@sandia.gov

PP1

A Low-communication Method to Solve Poisson's Equation on Locally-structured Grids

This poster describes a new algorithm, Method of Local Corrections (MLC), and a high-performance implementation for solving Poisson's equation with infinite-domain boundary conditions, on locally-refined nested rectangular grids. MLC represents the potential as a linear superposition of small local discrete convolutions, with global coupling represented using a non-iterative form of geometric multigrid. Thus the data motion is comparable to that of only a single V-cycle of multigrid, and hence is an order of magnitude smaller than tradition multigrid iteration. The computational kernels where most of the time is spent are 3D FFTs on small domains. Our results show solution error that is fourth order in mesh spacing, down to a fixed barrier that can be made low using high-order Mehrstellen stencils. Strong scaling tests on 64 to 4096 cores on NERSC Cori I (Haswell) show over 60% efficiency, and weak scaling by replication tests over 64 to 32768 cores show 92% efficiency on the same platform. We find comparable solve times between HPGMG on a uniform grid with one billion grid points, and MLC on the same number of grid points adaptively distributed. Since MLC operates on a nested series of adaptive locally-refined grids, it is able to solve problems with much higher resolution at the finest level than an algorithm on a uniform grid.

Peter Mccorquodale, Phillip Colella
Lawrence Berkeley National Laboratory
PWMcCorquodale@lbl.gov, PColella@lbl.gov

Brian Van Straalen
Lawrence Berkeley National Laboratory
Compuational Research Division
bvstraalen@lbl.gov

Christos Kavouklis
Lawrence Livermore National Laboratory
kavouklis1@llnl.gov

PP1

Machine Learning of a Dynamical Systems Model of Aircraft Environmental Control System Pneumatics

As humans endeavor to engineer complex systems that push the boundary of historical performance envelopes, it becomes critical to compare our design assumptions with actual measured performance. Insight into where our assumptions fail is essential to troubleshooting unexpected system behavior and improving our predictions. Aircraft environmental control systems are pneumatic systems comprised of many separately designed complex subsystems, which must be integrated into a whole to control the temperature and air supply for both cooling avionics and providing flow for aircrew breathing. Predictive models of pneumatic systems often include nonlinear effects, which are sometimes not observable until flight test. Nonlinear differential equation models of several of these subsystems were discovered from flight data using a sparse regression algorithm by Brunton et al. [PNAS, 113, 5, (2016)]. By several different metrics, model predictions compared favorably with the flight data reserved for validation. These methods hold promise for combining different experimental data recordings to produce predictions of air pressure and oxygen content supplied to the aircrew. The integration of adsorption air separation units for aircrew breathing supply into larger more complex pneumatic systems in aircraft makes testing our design assumptions against flight data important for ensuring aircrew safety. DISTRIBUTION STATEMENT A. Approved for public release. PR#18-0186

Kirsten Meeker
Naval Air Warfare Center
kirsten.meeker@navy.mil

PP1

The Theoretical and Computational Analysis of the Non-conserved Interactive Driven Diffusive Systems

Stimulated by the biological transport phenomena and the vehicular traffic flow, we investigate the dynamics of interacting molecular motors or interacting vehicles that move along linear filaments or roads and can reversibly attach/detach from them. To analyze these processes, we introduced a new model adapting the interactions in a totally asymmetric simple exclusion process coupled with nonconserving Langmuir kinetics (TASEP-LK). The model was first analyzed using the continuum version of the simple mean-field approach that neglects the correlations between the particles. We compared the theoretical results with computer simulations of high fidelity and found them

to deviate significantly. To recover the problem, we developed a new numerical method that takes into account some correlations in the system. We explicitly evaluated the effect of interactions on stationary phase diagrams, particle currents, densities, and the two-point correlation function. We found the correlations to be stronger at the positions of the localized shocks. Moreover, all the theoretical calculations were in excellent union with Monte Carlo computer simulations.

Tripti Midha
Indian Institute of Technology Ropar
Indian Institute of Technology Ropar
tripti.midha@iitrpr.ac.in

Anatoly B. Kolomeisky
Rice University
tolya@rice.edu

Arvind K. Gupta
Indian Institute of Technology Ropar
India
akgupta@iitrpr.ac.in

PP1

The Impact of Noise on Krylov Methods

In this work, we document and analyze the impact of operating system noise and machine variance on Krylov and pipelined Krylov method execution on the Cray XC40 Theta supercomputer. We perform repeated runs and collect fine-grained data that shows the influence of system variability on Krylov and pipelined Krylov methods. To succinctly describe the execution performance, we employ stochastic models based on the distribution of iteration times of each method. We test the models with collected data and suggest ways to improve and expand them.

Hannah M. Morgan
Argonne National Laboratory
Mathematics and Computer Science
hmorgan@anl.gov

Patrick Sanan
USI - Università della Svizzera italiana
ICS
patrick.sanan@gmail.com

Matthew G. Knepley
University at Buffalo
Department of Computer Science and Engineering
knepley@gmail.com

PP1

Benchmarking as an Aid to Selecting the Right Hardware Architecture and Numerical Method

A number of recent benchmarks for high performance computers suggest the use of an iterative method for solving sparse linear systems of equations to rank computer performance. Particular methods used include multigrid and conjugate gradients, though other methods such as the fast Fourier transform are also applicable in some cases. Asymptotic operation counts are not however always illustrative of efficiency in practice. Similarly, the level of accuracy can also be important, with some numerical methods being efficient for low accuracy simulations, but others more efficient for high accuracy simulations. These tradeoffs are examined in the parallel numerical solution of semilinear partial differential equations, in particular for wave equations where in many cases, short computation times enable execution of the full application with a massively parallel, but limited total computational resource use.

Benson K. Muite
University of Tartu, Estonia
benson.muite@ut.ee

PP1

Refined Isogeometric Analysis (rIGA): A Fluid Flow Application

Refined Isogeometric Analysis (rIGA) is a discretization method used to solve numerical problems governed by partial differential equations (PDEs). Starting from a highly continuous Isogeometric Analysis (IGA) discretization, rIGA reduces the continuity over certain hyperplanes, which act as separators during the elimination of the degrees of freedom (DoF). This method dramatically reduces the computational time employed to approximate the PDEs solution using direct solvers while it simultaneously improves the best approximation error. When using rIGA to solve a Laplace problem in 2D and 3D, the total computational time decreases by a factor that is approximately p^2 (with p being the polynomial order) with respect to the IGA (with maximum continuity) counterpart. In this work, we apply rIGA method to solve an incompressible fluid flow problem on a bounded domain. To approximate the velocity field, we use a spline-based generalization of the Raviart-Thomas finite element spaces, while for the pressure field we use a spline-based L^2 space. rIGA delivers a reduction in the computational cost when solving the fluid flow problem that asymptotically reaches $\mathcal{O}(p^2)$. Moreover, it provides better accuracy than C^{p-1} IGA. However, to arrive at the asymptotic limit and reach the maximum possible savings, we require large problems. In 3D, we often observe gain factors of $\mathcal{O}(p)$ in the pre-asymptotic regime.

Daniel Garcia
Basque Center for Applied Mathematics (BCAM)
dagarcialhit@gmail.com

David Pardo
Department of Applied Mathematics, Statistics, and Operational Research, University of the Basque Country
dzubiaur@gmail.com

Victor Calo
KAUST
vmcalo@gmail.com

Judit Muñoz-Matute
University of the Basque Country (UPV/EHU), Leioa, Spain

judit.munoz@ehu.eus

PP1

Stability of Traveling Wave Solutions of Nonlinear Conservation Laws for Image Processing

In image processing, nonlinear partial differential equations have been applied to denoising images and edge detection. One of the most well-known methods is based on the idea of anisotropic diffusion, which was proposed by Perona and Malik (PM) to reduce image noise without losing important image contents. Kurganov, Levy and Rosenau introduced a convection term in the PM equation to investigate the behavior of solutions which are typical of edges in images. In this poster, we study the existence and asymptotic stability of smooth traveling wave solutions to the Cauchy problem which combines a more general anisotropic diffusion and a general convection term.

Jeungeun Park, Tong Li
The Department of Mathematics
The University of Iowa
jeungeun-park@uiowa.edu, tong-li@uiowa.edu

PP1

LFA Lab - Flexible Local Fourier Analysis Library

LFA Lab is a Python library for analyzing multigrid methods using local Fourier analysis (LFA). Multigrid methods efficiently solve many types of linear systems. They, however, require a proper choice of parameters. LFA can be used to determine the quality of a parameter configuration, and LFA Lab makes LFA analysis easy by providing a general set of components that can be combined in many ways. This poster presents the theoretical background needed to use the library, introductory examples and interesting applications.

Hannah Rittich
University of Wuppertal
h.rittich@fz-juelich.de

PP1

Improved Newton Linearization for L^1-Norm-Type Minimization with Application to Viscoplastic Fluid Solvers

We target highly nonlinear applications in CS&E that are modelled by optimization problems with Hessians exhibiting a problematic (near) null space upon linearization with Newton's method. The null space is caused by a projector-type coefficient in the Hessian, which is created by terms in the objective functional that resemble the L^1-norm. This occurs, e.g., in nonlinear incompressible Stokes flow in Earth's mantle with plastic yielding rheology, which effectively limits stresses in the mantle by weakening the viscosity depending on the strain rate. Another example is total variation regularization for inverse problems. Using a standard Newton linearization for such applications is known to produce severe Newton step length reductions due to backtracking line search and stagnating nonlinear convergence. Additionally, these effects become increasingly prevalent as the mesh is refined. We analyze issues with the standard Newton linearization in an abstract setting and propose an improved linearization, which can be applied straightforwardly to Stokes flow with yielding and total variation regularization. Finally, numerical experiments compare the standard and improved Newton linearizations in practice. When we employ our improved linearization within our inexact Newton-Krylov method, a fast and highly robust nonlinear solver is attained that exhibits mesh-independent convergence and scales to large numbers of cores with high parallel efficiency.

Johann Rudi
Argonne National Laboratory
The University of Texas at Austin
jrudi@anl.gov

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

Omar Ghattas
The University of Texas at Austin
omar@ices.utexas.edu

PP1

Adjoint Based H-(m)odel Adaptive Scheme for Linearised Boltzmann Equation

Interest in deterministic methods to solve the Boltzmann type equations has increased in the recent years. Owing to the large number of degrees of freedom, required to sufficiently resolve the kinetic solution, most of the well-known deterministic methods are computationally very expensive. An effective method to cope with high dimensionality of deterministic methods is to develop an adaptive algorithm which allocates degrees of freedom only where the solution is under-resolved. In the present work, we approximate the solution to the linearised Boltzmann equation with a Galerkin types approach. In the physical space we use a discontinuous galerkin scheme whereas in the velocity space we use the Hermite series expansion proposed by [Grad,1949]. We focus upon steady state problems and use the adjoint of the kinetic equation to provide us with an a-posteriori error indicator. By balancing the predicted error for the spatial discretization and the Hermite series expansion, we propose an adaptive algorithm which selectively refines the spatial grid and chooses the appropriate polynomial order in the Hermite expansion. Such an algorithm helps us reach the desired error tolerance with sufficiently less degrees of freedom thus reducing the overall computational cost.

Neeraj Sarna
RWTH
sarna@mathcces.rwth-aachen.de

Manuel Torrilhon
RWTH Aachen University
mt@mathcces.rwth-aachen.de

PP1

On Composable Block Solvers and Performance

Spectrum Model for the Four-Field Double Porosity/permeability Model

Flow in many geo-materials such as aggregated soils or fissured rocks, which contain complex pore-networks, can be correctly modeled under the double porosity/permeability (DPP) model. DPP is a four-field mathematical model that describes the flow of a single-phase incompressible fluid in a porous medium with two distinct pore-networks and with a possibility of mass transfer across them. To facilitate solving large-scale problems under the DPP model, we first propose two composable block solver methodologies to solve the discrete systems that arise from finite element (FE) discretizations of the DPP model. Using the composable solvers feature available in PETSc and the FE libraries under the Firedrake Project, we illustrate two approaches by which one can effectively precondition these large systems of equations. Second, we employ the recently developed performance model called the Time-Accuracy- Size (TAS) spectrum to demonstrate that the proposed solvers are scalable in both the parallel and algorithmic sense. Also, we utilize this spectrum analysis to compare the performance of three different FE discretizations (classical mixed formulation with H(div) elements, stabilized continuous Galerkin mixed formulation, and stabilized discontinuous Galerkin mixed formulation) for the DPP model. Our performance spectrum analysis demonstrates that the composable block solvers are fine choices for any of these three FE discretizations.

Mohammad Sarraf Joshaghani
University of Houston
University of Houston
m.sarraf.j@gmail.com

Justin Chang
Rice University
jychang48@gmail.com

Kalyana Nakshatrala
University of Houston
kalyanababu@gmail.com

Matthew G. Knepley
University at Buffalo
Department of Computer Science and Engineering
knepley@gmail.com

PP1

An Immersed Boundary, Fourier Pseudospectral Method for Advection-Diffusion-Reaction in Incompressible Flow

An immersed boundary method (IBM) for the advection, diffusion, and reaction (ADR) of chemicals or species in incompressible flow is proposed. Suitable for complex geometries, the IBM is used in this paper to apply Dirichlet boundary conditions to surfaces of systems modeled by the Navier-Stokes, continuity, and ADR equations. Embedded in a Fourier Pseudospectral scheme, the method lends itself to high accuracy and computational efficiency. The numerical algorithm is verified using analytical solutions given by 2-D counter-rotating Taylor-Green vorticies, implemented with and without immersed boundaries. High order convergence and machine precision is observed. A validation case is then run, simulating a 2-D Michaelis-Menten enzyme/catalyst mechanism in incompressible flow. Results achieved are in good agreement with available data in literature.

Joseph Schiff
California State, Long Beach
j.f.schiff@gmail.com

PP1

Communication-avoiding Mesh Multiplication for Large-scale Unstructured Meshes

Often unstructured grid methods, such as finite elements, are used in high fidelity simulations. In these methods, solution accuracy is associated with the interpolation order and to the grid size. Unstructured parallel mesh refinement is computationally expensive due to subdomains interface communication. In the present work, we develop a uniform edge-based parallel tetrahedral mesh multiplication scheme completely free of communication, fast, simple to implement and highly scalable. Avoiding communication is achieved by an index generation for subdomain interface grid points based on special pairing functions like the classical examples of Cantors and Szudziks pairing functions. Besides, exploring the mesh properties, we propose a method called economic pairing which is provable an $\mathcal{O}(n)$ function and has all the properties of the classical pairing functions. We also address the issue of repositioning the new nodes, using boundary smoothing where kD-trees or a perfect hash search for the nearest neighbors. We evaluate the parallel mesh multiplication strategy generating meshes of $\mathcal{O}(10^{12})$ tetrahedra.

Romulo M. Silva
Federal University of Rio de Janerio
romulo.silva@nacad.ufrj.br

Benaia Lima
Federal University of Rio de Janeiro
benaia@nacad.ufrj.br

José Camata
NACAD - High Performance Computing Center
COPPE/Federal University of Rio de Janeiro
camata@nacad.ufrj.br

Renato Elias
Federal University of Rio de Janeiro
renato@nacad.ufrj.br

Alvaro Coutinho
Dept. of Civil Engineering
Federal University of Rio de Janeiro
alvaro@nacad.ufrj.br

PP1

Performance Analytics for Computational Experiments (PACE)

The Energy Exascale Earth System Model (E3SM) is a high-resolution coupled Earth system model, designed to

address energy-related science challenges of national interest while effectively using Department of Energy (DOE) supercomputers. This work presents PACE (Performance Analytics for Computational Experiments), a framework to summarize performance data collected from E3SM experiments to derive insights and present them through a web portal. PACE is designed to help identify bottlenecks and targets for performance engineering and optimization. E3SM incorporates a default lightweight performance profiling capability that is based on the General Purpose Timing Library (GPTL). PACE ingests the performance data from a completed experiment to facilitate interactive performance exploration including deep-dive into performance of different parallel processes and threads. Furthermore, it enables multi-experiment comparisons including scalability analysis for well-defined problem configurations. PACE uses MariaDB database to store structured and unstructured experiments outputs; various tools in the Python ecosystem for the backend infrastructure and middleware; and JavaScript tools for frontend and visual analytics. PACE enabled climate scientists to view executive summary of E3SM experiments and interactively deep-dive as desired. PACE is designed to be generic with reusable components to facilitate performance data collection and analysis for diverse applications.

Sarat Sreepathi
Oak Ridge National Laboratory
sarat@ornl.gov

Zachary Mitchell
Pellissippi State Community College
zdmitchell1@pstcc.edu

Gaurab Kc
University of Tennessee Knoxville
gkc@vols.utk.edu

PP1

Physics Constrained Machine Learning

Our proposed methodology considers the creation of parametric surrogate models for applications in science and engineering where the goal is to predict high-dimensional output quantities of interest, such as pressure, temperature and strain fields. A low-dimensional parameterization of these quantities of interest is developed using the proper orthogonal decomposition (POD), and combines this parameterization with machine learning methods to learn the map between the input parameters and the POD expansion coefficients. The use of particular solutions in the POD expansion provides a way to embed physical constraints, such as boundary conditions and other features of the solution that must be preserved. The relative costs and effectiveness of four different machine learning techniques - neural networks, multivariate polynomial regression, k-nearest-neighbors and decision trees - are explored through multiple engineering examples.

Renee Swischuk
Texas A&M University
swischuk@mit.edu

Karen E. Willcox
Massachusetts Institute of Technology
kwillcox@ices.utexas.edu

PP1

A Computationally Inexpensive Approach to Set Partitioning with Applications to Optimal Planning for Large-scale Robot Swarms

This work introduces a method to reduce the computational cost of solving constrained optimization problems for high dimensional systems. We use a specific family of pre-fractal curves as a mapping under which high dimensional problems can be embedded in a real interval and solved in quasi-linear time complexity. For demonstration, we applied this method to address the optimal routing problem for energy-constrained swarms in rapidly-changing environments Using the proposed method, a swarm of 500 agents can find a near-optimal solutions to time-optimal and energy-optimal planning problems in about 0.1 seconds. We can compute the optimality bounds for the solution from the boundedness property of the introduced map. Results of this study suggests that the cost of our solution converges to the cost of optimal plan as the number of swarm agents or the number of covered zones increases, which makes this approach suitable for large-scale swarm systems. The effectiveness and efficiency of the method is evaluated through simulations with different number of swarm agents, different number of target zones and their spatial distribution.

Shadi Tasdighi Kalat
Worcester Polytechnic Institute
stasdighikalat@wpi.edu

PP1

Predicting Optimal Error Bounded Lossy Compression Configuration for Sampled Data

Traditional lossless compressors can compress information to a smaller size, but dont perform well on all data types. However, there are many domains do not require the precision of lossless compressors. It can be challenging and time consuming to determine what parameters of the compressor produce the best compression ratio or bandwidth. We propose to construct a model that predicts a high-quality configuration for the SZ error-bounded lossy compressor (EBLC). This model will aid designers in choosing an appropriate configuration given their data distribution. This poster presents and validates a modeling approach and model of the compression ratio, compression bandwidth, and decompression bandwidth for the SZ one of the leading ELBCs. We demonstrate that the configuration of the compressor has significant impacts on these performance parameters using synthetic datasets. We further show that these impacts are predictable considering only properties of the sampled distribution. Finally, we are in the process of validating our model of optimal by comparing our models suggested parameters for real datasets from high performance computing (HPC) and intelligent transportation systems (ITS) to an exhaustive parameter sweep. We present preliminary work in this area. The model produced by this poster has broad application to these domains. It will benefit systems designers who work with quantified

uncertainty who need higher performance out of existing system.

Robert R. Underwood
Clemson University
robertu@g.clemson.edu

Jon Calhoun
Clemson University
Holcombe Department of Electrical and Computing Engineering
jonccal@clemson.edu

Amy Apon
Clemson University
School of Computing
aapon@clemson.edu

PP1

A Dynamical Systems Approach to Transforming Disparate Timescales in Data Driven Equation-free Modeling of Disease Dynamics

Effective prediction of disease dynamics often requires knowledge of both endogenous and exogenous factors. Incorporating diverse types of information, however, introduces covariates at different observational timescales. Despite advances in time dependency and discretization in data driven equation-free modeling, dissimilarity in time intervals has remained a challenge. We present a method that addresses this limitation, judiciously unifying timescales of input data. To demonstrate, we develop a malaria transmission model. Relevant remote sensing data includes weekly land cover data alongside sub-daily precipitation and temperature. Plasmodium incidence is reported at a weekly to monthly frequency. Our approach rigorously connects these timescales using principles from non-autonomous dynamical systems and a recently developed method for sparse identification of nonlinear dynamics (SINDy). We employ this method in multiple layers: identifying the disease transmission model as well as optimizing the timescale transformation itself. This layered SINDy approach identifies the most appropriate functional form for converting inputs to a unified timescale. Incorporation of environmental and epidemiological data initially at different timescales expands the options for feature selection and increases the models predictive power. Applications of this approach are broadly relevant to disease modeling as well as any data driven equation-free modeling incorporating diverse input data.

Mollie Van Gordon, Joshua Proctor
Institute for Disease Modeling
mvangordon@idmod.org, jproctor@idmod.org

PP1

A Variational Approach to Design a Numerical Scheme for N-Fluid Flow

In some highly demanding fluid dynamics simulations, it appears necessary to simulate multiphase flows involving numerous constraints at the same time, such as: large number of fluids, both isentropic and strongly shocked compressible evolution, large deformations, transport over large distances, and highly variable equation of state stiffenesses. Fulfilling such a challenge in a robust and tractable way demands that thermodynamic consistency of the numerical scheme be carefully ensured. This is addressed here over an arbitrarily evolving computational grid (ALE or Arbitrary LagrangianEulerian approach) by a three-step mimicking derivation procedure named GEECS (Geometry, Energy, and Entropy Compatible). The resulting multiphase scheme named multiGEECS involves the following features: i) full conservation of mass, momentum, and total energy at discrete level; ii) direct ALE formalism where mass, momentum, and internal energy fluxes are taken into account directly into the discrete evolution equations; iii) thermodynamic consistency of the pressure work obtained by application of a variational principle; iv) pressure equilibrium through a simple and local (to the cells) procedure; and v) generic set of evolution equations written for an arbitrary number of fluids and derived without any constraint on structure or spatial dimension in order to simulate a broad category of multiphase flows.

Thibaud Vazquez-Gonzalez
CEA Bruyeres le Chatel
thibaud.vazquez-gonzalez@cea.fr

Antoine Llor
Commissariat à l'Energie Atomique, DAM-DIF
antoine.llor@cea.fr

PP1

Using Bayesian Inference to Evaluate Rare Event Probabilities

Inputs to computer models of physical reality are often uncertain. This uncertainty, which may arise due to measurement errors, lack of knowledge, or small fluctuations, gets propagated through the model, leading to uncertain outputs. In this work, we consider the following uncertainty quantification (UQ) task - given a map with uncertain inputs, compute the probability that the output evaluates inside a specific interval. Standard statistical approaches for this problem, such as the Monte Carlo method, become increasingly expensive as the event under consideration becomes rare. We propose an importance sampling (IS) algorithm to address this problem. Our algorithm leverages a statistical inverse problem to arrive at the importance sampling probability density. Hence, we describe how *inverse* UQ can be used to efficiently perform a goal-oriented *forward* UQ task. A well-known drawback of importance sampling is that is suffers from the curse-of-dimensionality. We take the connection between forward and inverse UQ further and describe how recent advances in dimension reduction for Bayesian inverse problems can be used to define the notion of an intrinsic dimension for the forward UQ task, as a step towards constructing an IS density that is insensitive to the ambient dimension of the problem.

Siddhant Wahal
The university of Texas at Austin
siddhant@ices.utexas.edu

Biros George
ICES, UT Austin

biros@ices.utexas.edu

PP1

Locally-implicit Discontinuous Galerkin Schemes with Limiters that Guarantee Moment-Invertibility for Hyperbolic Quadrature-Based Moment Closures

We consider the quadrature-based moment closure approach for approximating kinetic Boltzmann equations. The true distribution is replaced by Dirac deltas with variable weights and abscissas. We show how to construct these to obtain a set of conservation laws that are conditionally hyperbolic. We then develop a high-order numerical method with limiters that guarantee that the numerical solutions remain in the convex hyperbolic regions of solution space. Numerical examples in 1D and 2D will be presented.

Christine Wiersma, James A. Rossmanith
Iowa State University
Department of Mathematics
cwiersma@iastate.edu, rossmani@iastate.edu

PP1

Physics-informed Machine Learning for Data-driven Turbulence Modeling

Reynolds-averaged Navier-Stokes (RANS) models have large model-form uncertainties in engineering applications that involve turbulent flows. Recently, data-driven methods have been proposed as a promising alternative. In this poster, we present a physics-informed machine-learning-assisted framework for RANS modeling. This framework consists of four components: (1) extracting invariant inputs and outputs to ensure important invariances of the machine-learning-assisted model, (2) constructing a functional mapping between inputs and outputs via machine learning techniques, (3) assessing the prediction confidence a priori by evaluating the distance between different flows in the mean flow features space, and (4) propagating the predicted Reynolds stress field to mean velocity field by using physics-informed stabilization. With the machine-learning-assisted RANS model, significant improvements over the baseline RANS simulation are observed for both the Reynolds stress and the mean velocity fields. This work also has great potential as an alternative to the explicit modeling of closures for unresolved physics when simulating other complex systems such as Earths climate.

Jinlong Wu
Dept. of Aerospace and Ocean Engineering, Virginia Tech
jinlong@vt.edu

Carlos Michelen-Strofer
Virginia Tech
cmich@vt.edu

Jianxun Wang
University of Notre Dame
jwang33@nd.edu

Heng Xiao
Dept. of Aerospace and Ocean Engineering, Virginia Tech
hengxiao@vt.edu

PP1

Fast Multipole Method for Non-oscillating Kernel Based on Cartesian Tensor and Differential Algebra

The fast multipole method (FMM) scales linearly with the number of particles when calculating the non-oscillating pairwise interaction between the particles. In FMM the interaction is represented as multipole expansions around the sources and then converted into local expansions around the objectives. The Cartesian tensor-based FMM was originally derived for the Coulomb interaction between charged particles by representing the expansions as well as the locomotion and translation of the expansions using totally symmetric Cartesian tensors. But the Cartesian tensor-based FMM is known to be semi-kernel independent. The method can be applied on a kernel as long as the high order derivatives of the kernel can be calculated. The differential algebra (DA) technique allows to calculate the derivatives up to any predetermined order for differentiable kernel functions including those without explicit formulas, which can only be calculated numerically. Using DA, the Cartesian tensor FMM is easily extended to any non-oscillating kernel. We will present the algorithm, numerical experiments, and the performance analysis.

He Zhang
Jefferson Lab
hezhang@jlab.org

He Huang
Old Dominion University
huang@jlab.org

Li-Shi Luo
Dept of Math. & Stat.
Old Dominion University
lluo@odu.edu

PP2

Padé Time Stepping Method of Rational Form for Reaction Diffusion Equations

One of the difficulties of the numerical solution of time-dependent ordinary and partial differential equations is the severe restrictions on time-step sizes for stability in explicit methods. Otherwise, challenging, generally nonlinear systems of equations in implicit schemes would be imposed to solve such problem. A class of explicit methods based on use of Padé approximation is introduced. They are inexpensive per time-step and they possess stability restrictions similar to the one offered by implicit schemes. Due to the rationality form of PTS, some numerical error occurs and then some a kind of control is imposed. We find that the Padé time stepping (PTS) showed a preferable behaviour when reaction diffusion equations are considered. We also notice that the PTS schemes have less computational time than the compared ones. Finally, numerical runs are conducted to obtain the optimal local error control threshold.

Said Algarni
Assistant Professor - Department of Mathematics
King Fahd University of Petroleum and Minerals
garnis@kfupm.edu.sa

PP2

Sparse Approximate Matrix Multiplication in a Fully Recursive Distributed Task-based Parallel Framework

We consider a parallel implementation of approximate multiplication of large matrices with decay. Such matrices arise in computations related to electronic structure calculations and some other fields of science. The original algorithm for the sparse approximate multiplication was suggested by [N. Bock and M. Challacombe, An optimized sparse approximate matrix multiply for matrices with decay]. An implementation done using the Chunks and Tasks programming model and library [E.H. Rubensson and E. Rudberg, Chunks and Tasks: A programming model for parallelization of dynamic algorithms] is presented and discussed. We describe a two-level approach, where the outer one operates with tasks in parallel, while the inner one performs actual computations within tasks. The implementation of the algorithm is applied to large chemical systems with $> 10^6$ atoms. We found out that it is competitive to another popular approach, which performs truncation of small blocks before multiplication. A comparison between these two methods in terms of performance on a model problem is done. The method is then applied to real matrices arising in quantum chemistry.

Anton G. Artemov
Uppsala University
anton.artemov@it.uu.se

PP2

A Double Core Tensor Factorization for Heterogeneous Data

In tensor factorization, the core tensor and its entries provide the level of interaction between the factor matrices and their components. Standard tensor factorizations use a single and unsupervised core to achieve this objective. However, the presence of heterogeneous data may induce different structure for subsets of them. To address this issue, we propose a double core tensor factorization (DCOT), where the core tensor is given by the superposition of global and subset specific local cores. DCOT preserves structural properties of the heterogeneous datasets, such as joint structure within subsets and idiosyncratic structures across different subsets (e.g. such subsets may correspond to presepcified groups). Further, in order to characterize the underlying joint-manifold drawn from the model factors, we propose two new structural regularization schemes: (1) a Nystrom multilinear graph embedding regularization on the factor matrices to characterize their underlying manifold determined by them, and (2) an orthogonal total variation regularization, on individual factor matrices of the tensor to encourage uniqueness of the factorization. A fast and efficient sketch based alternating minimization algorithm is suggested for DCOT factorization. The performance of DCOT and its Nystrom version (called DCOT-N) are illustrated on clustering, tensor completion, and foreground-background separation in surveillance video applications.

Davoud Ataee Tarzanagh, George Michailidis
University of Florida
Tarzanagh@ufl.edu, gmichail@ufl.edu

PP2

Randomized Sub-sampled Methods for Matrix Approximation

This talk introduces a framework for randomized sub-sampling methods for matrix approximation. Modern computational problems require solving linear systems relating to increasingly larger data sets. In most cases it is impractical to use all of the data available and unnecessary for reasonable approximations. Current work uses sampling methods to iteratively construct an approximation to a matrix. Sub-sampling methods require less work per iteration and are more adaptable to different problem settings when compared to full sampling. Convergence of the sub-sampled iterates is proven in expectation. An accelerated sampling scheme is presented and experimentally shown to be comparable to current accelerated sampling methods.

Andrew T. Azzam, Benjamin W. Ong, Allan Struthers
Michigan Technological University
atazzam@mtu.edu, ongbw@mtu.edu, struther@mtu.edu

PP2

A Stabilized, Hypoelastic Constitutive Correspondence Framework for Peridynamics

Over the past decade, the peridynamic framework has attracted a broad interest from the computational mechanics community due to its natural capabilities in modeling crack nucleation and growth without any complicated numerical treatments at discontinuities. We present a recently developed constitutive correspondence theory for peridynamics that approximates the velocity gradient in the current configuration using an updated (semi-Lagrangian) kernel. We show this model to have improved stability over the standard correspondence theory of peridynamcis and believe it to be useful in applications of large plastic deformation and fragmentation. We demonstrate the relationship of this theory to classical hypoelastic constitutive theories under uniform deformations and show numerical examples of the new theory in applications of crack propagation with large plastic deformations.

Masoud Behzadinasab, John Foster
University of Texas at Austin
behzadi@utexas.edu, john.foster@utexas.edu

PP2

Anomalous Sub- and Super-diffusion in Image Pro-

cessing

Abstract not available

Toheeb A. Biala, Abdul Khaliq
Middle Tennessee State University
tab7v@mtmail.mtsu.edu, abbdul.khqliq@mtsu.edu

PP2

Blocked Multigrid Methods for Structured Matrices

Multigrid methods are well-known to be optimal solvers for different classes of matrices, including matrices arising from the discretization of partial differential equations (PDEs) and certain structured matrices, e.g., Toeplitz matrices. Usually, point-smoothers are used in multigrid methods. The use of larger blocks with or without overlap can speed up the convergence rate of multigrid methods. This becomes necessary when the convergence rate of the multigrid method is unsatisfactory, e.g., when larger coarsening ratios are used. Furthermore blocked methods can be implemented efficiently on modern computer architectures, effectively reducing the wall-clock time needed to obtain the solution of the systems. The analysis of blocked methods can be carried out by modifying the standard tools. We focus on matrices that arise in the discretization of partial elliptic differential equations with constant coefficients. This is independent from what discretization is used and from the dimension of the problem. The theoretical results that are obtained in these cases can be carried over to the non-constant coefficient case as usual. We will present analyses of blocked methods for structured matrices as well as numerical results of implementations thereof.

Matthias Bolten
Bergische Universität Wuppertal
Fakultät für Mathematik und Naturwissenschaften
bolten@math.uni-wuppertal.de

PP2

Exploring the Variability and Error Distributions of Lossy Compression Algorithms

Lossy compression algorithms are effective tools to reduce the size of HPC datasets. As established lossy compressors such as SZ and ZFP evolve they seek to improve the compression/decompression bandwidth and the compression ratio. As the the underlying algorithms of these compressors evolve, the spatial distribution of errors in the compressed data changes even for the same error bound type and error bound. Recent work has shown that properties of the simulation such as choice of boundary conditions, PDE properties, and domain geometry significantly impacts an application's ability to compute on state from lossy compressed data files. If HPC applications are to compute on data coming from compressed data files, we require an understanding of how the spatial distribution of error changes. This poster explores how spatial distributions of error, compression/decompression bandwidth, and compression ratio changes for HPC data sets from the applications PlasComCM and Nek5000 between various versions of SZ and ZFP. In addition, we explore how the spatial distribution of error impacts the correctness of the applications and the ability to create methodologies to recommend lossy compression error bounds.

Jon Calhoun
Clemson University
Holcombe Department of Electrical and Computing Engineering
jonccal@clemson.edu

Tasmia Reza, Pavlo Triantafyllides
Clemson University
treza@g.clemson.edu, ptriant@g.clemson.edu

PP2

Development of Efficient Preconditioners for Newton-Krylov Method in Spectral/*hp* Element, DG Compressible Flow Simulations

High-order spectral/*hp* element methods using a Discontinuous Galerkin (DG) formulation, are suitable for high fidelity compressible simulations around complex geometries because of their good resolution and compactness properties. However, explicit DG simulations of high Reynolds number flows are still very challenging because of overly restrictive time steps. To accelerate the DG simulations of compressible flows in the open source package Nektar++, implicit time integration methods based on Jacobian-Free Newton-Krylov (JFNK) method have been investigated. Although the JFNK method utilizes the relatively low memory footprint of the explicit formulation it also requires a preconditioner that lead to a much larger memory footprint. In this presentation we will discuss the effectiveness of different preconditioning methods, such as SOR iteration, Block ILU and p-multigrid method. In addition we are currently exploring eigenvalue analysis of the discrete system to understand the influences of different stiffness sources such as sound wave, viscosity and grid stretching on the eigenspectra. Based on this analysis, we will therefore discuss the potential to simplify the preconditioning Jacobian matrix in order to reduce the computation cost and memory consumption. To highlight the speed up of the implicit approach while maintaining a reasonable memory consumption, DG compressible flow simulations of vortex shedding and boundary layer flows will be presented.

Zhen-Guo Yan
Department of Aeronautics, Imperial College London, UK
China Aerodynamics Research and Development Center
yanzhg@mail.ustc.edu.cn

Yu Pan
Department of Aeronautics, Imperial College London, UK
yu.pan16@imperial.ac.uk

Joaquim Peiro
Dept of Aeronautics
Imperial College London, UK
j.peiro@imperial.ac.uk

Spencer Sherwin
Department of Aeronautics
Imperial College London
s.sherwin@imperial.ac.uk

Giacomo Castiglioni
Imperial College London, United Kingdom
g.castiglioni@imperial.ac.uk

PP2

Understanding the Structure of Germanium Sulfide (GeS)

Energy conversion in Photovoltaics (PVs) needs to improve fast, one way of achieving this goal is by focusing in how photons are absorbed using new thin film materials. GeS has arisen as a potential substitute due to its high abundance and low toxicity. The purpose of this project was to manipulate GeS Band Gap by applying bi-axial strain (5%, 1% step) in the 'a' and 'c' lattice directions with the goal of changing GeS' gap from 1.74 eV to 1.3 eV, which is the optimal light absorption gap. To complete this process, we used computational methods with density functional theory (DFT). Before applying strain, we obtained an indirect band gap of 0.78 eV, which turned out to be roughly half of the experimental value. After straining the structure, the gap remained indirect and increased to 0.87 eV for 5% tensile, while for 5% compressive, the gap closed completely. However, an interesting behavior was observed between 0% and 2% tensile, the conduction band minimum (CBM) changed from being at Γ, to a point in between the Γ and X path. This opens the possibility of a shift from indirect to a direct gap at some point in the same path.

Miguel A. Castro
Inter-American University of Puerto Rico, Bayamón Campus
miguelcr1997@gmail.com

Betul Pamuk, Darrell Schlom
Cornell University
Materials Science and Engineering
betul.pamuk@cornell.edu, schlom@cornell.edu

PP2

A Scalable Petsc Implementation of the 3-Phase Unbalanced Ac Power Flow Solver

A revolutionary transition of energy infrastructure, spanning across supply and demand sectors, is happening all around the world. Integration of renewable energy sources into our ever-complex power distribution networks is central to the infrastructure planning of the future. One step towards achieving this goal is to solve unbalanced transmission and distribution systems. Traditional approaches and software tools are not suitable because they assume a completely balanced 3-phase system, hence a single-phase equivalent can be solved. Renewable energy such as wind and solar power are intermittent and will be naturally unbalanced when integrated into the power grid. Furthermore, these power grids can be quite large at scale consisting of potentially thousands of buses, which can be computationally demanding. This poster presents a skeleton of a robust and scalable 3-phase unbalanced AC power flow solver, based on the Power Injection Method (PIM). The proposed power flow solver is built on top of the PETSc library and utilizes state-of-the-art parallel data structures like the DMNetwork. Key steps of the algorithm are outlined, and several network problems of varying size are solved across multiple compute nodes. The scalability in both the parallel and algorithmic sense using the proposed software framework is highlighted.

Justin Chang, Gord Stephen, Dheepak Krishnamurthy
NREL
justin.chang@nrel.gov, gord.stephen@nrel.gov, dheepak.krishnamurthy@nrel.gov

PP2

Density-equalizing Reference Map with Applications

Density-equalizing map is a technique originally developed for cartogram creation. In this work, we propose a new method for computing density-equalizing maps with the aid of the reference map technique, which is an Eulerian approach for large-strain solid mechanics. Our method can be effectively applied to flattening-based medical visualization, and be generalized to higher dimensions for shape modeling, morphing and data visualization. Experimental results will be presented to demonstrate the effectiveness of our proposed method.

Gary Choi, Chris H. Rycroft
Harvard University
pchoi@g.harvard.edu, chr@seas.harvard.edu

PP2

Sensitivity Analysis in Particle-in-Cell Methods

The ability of Particle-in-Cell (PIC) methods to efficiently represent distribution functions is attractive for simulating plasma. Additionally, computing derivatives using sensitivity analysis is a critical tool for uncertainty quantification and gradient-based optimization. However, computing sensitivities for PIC representations remains challenging due to the chaos inherent in particle dynamics. We examine this challenge from a dual particle–continuum perspective that motivates a new particle-based sensitivity discretization. In its development, two routes to sensitivity computation are presented and compared: a *particle-exact* approach provides sensitivities of each particle trajectory, and a new *sensitivity-pdf* discretization. This new approach is formulated as a continuum problem discretized efficiently using PIC. In this approach, the continuum perspective makes implicit the distinguishability of particles. Sensitivity particles are only indirectly linked to PIC particles, which circumvents the Lyapunov instability of the N-body PIC system. The advantage of the sensitivity-pdf approach is demonstrated and quantified for standard model configurations. Since the sensitivity particles are, in a sense, independent of the plasma particles, it allows adaptive methods to focus computation effort on the sensitivity calculation.

Seung Whan Chung
University of Illinois at Urbana-Champaign
schung58@illinois.edu

Stephen D. Bond
Sandia National Laboratories
sdbond@sandia.gov

Eric C. Cyr
Computational Mathematics Department
Sandia National Laboratotories
eccyr@sandia.gov

Jonathan B. Freund
University of Illinois at Urbana-Champaign
jbfreund@illinois.edu

PP2

Pipelined Krylov Subspace Methods with Improved Parallel Scalability

The need for solving algebraic problems of ever larger scales can be observed in virtually all academic and industrial applications, including the simulation of ocean circulation models, global climate prediction models, etc. These model calculations require the solution of a large-scale linear system that typically consists of millions of unknowns. Due to the huge size of these model calculations, computations are spread across parallel computer platforms to reduce the time-to-solution. Krylov subspace methods (KSMs) have been established as the benchmark iterative solvers for sparse linear algebra problems. However, classical formulations of these methods are not adapted to scale to future parallel hardware due to communication latencies, stemming mainly from the calculation of dot-products in the algorithms. By introducing auxiliary variables and rearranging operations the global synchronization phase can be 'hidden behind the computational work of the sparse matrix vector product (SPMV), resulting in so-called 'pipelined variants of KSMs. When dot-product latency is longer than the time to compute an SPMV, the overlap of a global reduction phase with a single SPMV will not be able to completely hide the latency. Recently deep pipelined KSMs have been developed that allow overlapping multiple SPMVs with the global synchronization phase. This poster presents recent advances in the development and implementation of pipelined KSMs with improved parallel scalability.

Jeffrey Cornelis
University of Antwerp
jeffrey.cornelis@uantwerp.be

Pieter Ghysels
Lawrence Berkeley National Laboratory
Computational Research Division
pghysels@lbl.gov

Siegfried Cools, Wim Vanroose
University of Antwerp
siegfried.cools@uantwerp.be, wim.vanroose@uantwerp.be

PP2

On the less than Reasonable Effectiveness of Floating Point Arithmetic in the Modelling of Physical Systems: the Case of the Generalised Bernoulli Map

We discuss the structure and late time asymptotic behaviour of the generalised Bernoulli map, also known as the -shift, and compare its known mathematical properties with those determined using the IEEE floating point numbers available on all modern digital computers. We show that much of the structure of the dynamical system is lost in the floating point representation with various exactly known properties, such as the number of unstable periodic oribits, their properties and ensemble averages being only poorly approximated. Our work demonstrates that digital computers can provide a severely depleted description of chaotic dynamical systems.

Peter Coveney
University College London, United Kingdom
p.v.coveney@ucl.ac.uk

PP2

Exponential Affine Solutions for Correlated Stochastic Process

We develop analytical solutions for the Laplace and Fourier transforms of a stochastically correlated integrated two weighted factors Feller diffusion process. We exhibit the exponential affine solution calculated via Riccati differential equations derived from the Feynman-Kac theorem. It is easy to see that the case of two uncorrelated CIR processes is a particularization of our model. Moreover, the simulations carried out show that the Laplace transform solution of our model tracks that of two CIR models with constant correlation inferred via Monte Carlo simulations. The result encounters applications in many real-world situations: via the Laplace transform we can price bonds via splitting the nominal interest rates as a combination of real interest rates and actual inflation, estimate the default probability of emerging market bonds sensitive to US interest rates and calculate the failure probability of a series system. Via the Fourier transform we can find the probability density function of the process.

Allan J. Da Silva
National Laboratory for Scientific Computing - Brazil
Federal Center of Technological Education - Brazil
allanjs@lncc.br

Jack Baczynski
National Laboratory for Scientific Computing - Brazil
jack@lncc.br

José Vicente
BCB/IBMEC
jose.valentim@bcb.gov.br

PP2

An EXCMG Accelerated Multiscale Multigrid Computation for 3D Poisson Equation

Mutiscale multigrid (MSMG) method is an effective computational framework for computing high accuracy and high efficiency solutions for elliptical partial differential equations. Although many efforts have been put to improve its performance by eliminating the iterative refinement procedure on the finest grid, it still has room to make it more efficient. In the current MSMG method, compared to the CPU cost on computing sixth-order solutions by applying extrapolation and other techniques on

two fourth-order solutions from different scales grids, much more CPU time is spent on computing fourth-order solutions themselves on coarse and fine grids, particularly for high-dimensional problems. Here we propose to embed extrapolation cascadic multigrid (EXCMG) method into the MSMG computation to accelerate the whole process. Numerical results on 3D Poisson equations show that the new MSMG method is much more efficient than the current MSMG method and the EXCMG method for sixth-order solution computation.

Ruxin Dai
University of Wisconsin River Falls
ruxin.dai@uwrf.edu

PP2

Accelerated Boundary Integral Treecodes

A fast-summation treecode algorithm using a two-parameter family of regularized kernels is implemented to calculate dynamic properties and render simulations of classical N-body systems. This two-parameter family of regularized kernels is able to control modelling error introduced by approximating the original dynamical system using regularized kernels while providing sufficient smoothness to capture the dynamics. A key advancement in this work is the development of a recurrence formula to efficiently compute derivatives of the regularized kernels used in the treecode approximations.

Satyen V. Dhamankar, Benjamin W. Ong
Michigan Technological University
sdhamank@mtu.edu, ongbw@mtu.edu

PP2

Registry Effects in Carbon Nanostructures

A graphene sheet is a single-atom thick macromolecule of carbon atoms arranged in a hexagonal lattice. Our research is motivated in large part by the exceptional physical properties of graphene and its potential applications in engineering and materials science. In particular, the phenomena motivating our work include pattern formation and localized wrinkling driven by lattice and orientation mismatches between a graphene sheet and its supporting substrate, moire patterns in bilayer graphene, and polygonization and faceting in multi-walled carbon nanotubes. Specifically, we derive continuum models that retain lattice registry effects to describe weak van der Waals interactions in carbon nanostructures. Atomistic simulations are needed to validate our models.

Malena I. Espanol
Department of Mathematics
The University of Akron
mespanol@uakron.edu

Dmitry Golovaty
The University of Akron
Department of Mathematics
dmitry@uakron.edu

J. Patrick Wilber
The University of Akron
jw50@uakron.edu

PP2

Towards Precision Medicine: Simulation Based Parameter Estimation for Drug Metabolism

Accurate estimation of physiological parameters is important for personalized drug dosing. Common practice is to use a probe drug and to measure the response of the body through blood samples for a population to establish bounds for safety and efficacy. Blood sample based adaptation for individuals is challenging due to low sampling rate given clinical constraints. We present iPBPK-R, our individualized physiologically based pharmacokinetic modelling approach for rate data to address this problem. A microdose of the radio-labeled drug erythromycin is given intravenously. The amount of radioactive CO_2 is measured in the subjects breath with 11 samples spaced over 2 hours. The resulting data is per subject fitted to model with 9 ODEs and 12 free parameters to estimate the physiological parameters of interest. Breath data is proportional to the first derivative of blood concentrations and higher sampling frequency leads to more accurate estimates. Compared to standard PBPK approaches, we utilize a reduced order model that facilitates consistent parameter estimation. Model fitting requires nested optimization of very sensitive ODE parameters and well-crafted objective functions. In an outer loop, initial parameter estimates are found in pharmacology literature and then adjusted for subjects physiology and drug properties. The approach is implemented in R and run on the PSC Bridges supercomputer under an NSF XSEDE grant, where individual estimations require 4 to 36 CPU hours.

Yoko Franchetti
University of Pittsburgh School of Pharmacy
yof8@pitt.edu

Thomas Nolin
University of Pittsburgh School of Pharmacy
Department of Pharmacy and Therapeutics
nolin@pitt.edu

Franz Franchetti
Department of Electrical and Computer Engineering
Carnegie Mellon University
franzf@ece.cmu.edu

PP2

Vesicle Adhesion in Constricted Geometries

Vesicle suspensions are used to model red blood cells, and adhesion can create vesicle aggregates. In many physiological geometries, the aggregate must be broken up so that vesicles can pass through a narrow constriction by forming a single file line. Therefore, the adhesive forces must be overcome by hydrodynamics. We examine the effect of the adhesion strength, adhesion range, and vesicles' reduced area on the formation and disunion of aggregates before and after constrictions.

Ashley R. Gannon
Department of Scientific Computing

Florida State University
ag12s@my.fsu.edu

PP2

High-order Accurate Vlasov Discretizations in 2D+2V

Kinetic simulation of multi-dimensional plasma waves through direct discretization of the Vlasov equation is a useful tool to study many physical interactions, and is particularly attractive for situations where minimal fluctuation levels are desired (for instance, when measuring growth rates of plasma wave instabilities). However, direct discretization of phase space can be computationally expensive, and as a result there are few examples of published results using Vlasov codes in more than a single configuration space dimension. In an effort to fill this gap we have developed the Eulerian-based kinetic code LOKI that evolves the Vlasov-Poisson system in 2+2-dimensional phase space. Here we discuss the the 4th- and 6th-order accurate finite difference algorithms used in the code, and present rigorous verification of convergence rates using both manufactured solutions and extraction of physically relevant quantities of interest such as Landau damping decay rates.

Andre Gianesini Odu, Jeffrey W. Banks
Rensselaer Polytechnic Institute
gianea@rpi.edu, banksj3@rpi.edu

PP2

Elasticity in Space and Time

Elastic materials occur in many constructions as, e.g. in the clutch of a car or within a bridge to prevent the risk of fracture under high stress. The Lagrangian equation models the dynamics of such elastic bodies caused by external forces. In this configuration, the Lagrangian is represented as the difference of the kinetic energy and the potential energy. Our aim is to provide stable space-time formulations, see e.g., [Urban, Patera, *An improved Error Bound for Reduced Basis Approximation of linear parabolic problems*, 2012], for these type of mechanical equations. We rephrase the Lagrangian in a space-time variational formulation and provide stable pairs for trail and test spaces in the case of linear materials (for which the Lagrangian results in an equation of *wave-type*). In particular, we consider two formulations, which are characterized as a system of first order in time and very weak formulation in time. We provide numerical investigations for both formulations with focus on the stability in the finite-dimensional spaces. This is joint work with Christian Hesch, University of Siegen.

Silke Glas
University of Ulm, Germany
silke.glas@uni-ulm.de

Karsten Urban
Institute of Numerical Mathematics, University of Ulm
karsten.urban@uni-ulm.de

PP2

Numerical Simulation of Basal Entrainment of a Viscoplastic Fluid

The basal entrainment is the increase in the volume of natural gravity-driven flows by eroding the bed on which they descend. Numerical simulations offer an interesting alternative compared to experimental models of such phenomena, even though the prediction of dense granular flows is still challenging. As the particle concentration raises above a certain amount, the flow has a viscoplastic behavior which can be described by a generalized Newtonian model. Here, a Herschel-Bulkley rheology model is applied, in which the viscosity is not a constant, but shear rate dependent. We study the basal entrainment reproducing a dam-break experiment composed by a fixed amount of a viscoplastic fluid on a sloping erodible bed. We implement this phenomenon as numerical models in libMesh, an open finite element library that provides a framework for multiphysics. The mathematical model results from the incompressible Navier-Stokes equations combined with an advection-diffusion transport equation for suspended sediments. Both the Navier-Stokes equations and the suspended load equations are treated with the residual-based variational multiscale finite element formulation. The interface between water and air is modeled using the volume of fluid (VOF) method, while the interface between sediment and water is tracked using an Arbitrary Eulerian-Lagrangian approach. Empirical models are used to represent the entrainment rate of sediments. The results are discussed and analyzed.

Malú Grave
Federal University of Rio de Janeiro
COPPE/UFRJ
malugrave@nacad.ufrj.br

Linda Gesenhues
Federal University of Rio de Janeiro
linda@nacad.ufrj.br

José Camata
NACAD - High Performance Computing Center
COPPE/Federal University of Rio de Janeiro
camata@nacad.ufrj.br

Alvaro Coutinho
Dept. of Civil Engineering
Federal University of Rio de Janeiro
alvaro@nacad.ufrj.br

PP2

Analysis of a Coupled Reaction-diffusion Model for Tumour Induced Angiogenesis in Breast Cancer Tissue

An important stage in the formation and development of solid tumours is tumour-induced angiogenesis. This process links the relatively benign avascular phase to the prospective lethal vascular stage of solid tumour growth. Capillary vessels from pre-existing vascular systems begin

to form. This formation is a response to secreted tumour angiogenic factor (TAF) proteins secreted by the avascular solid tumour when it reaches a hypoxic state. Once secreted, the endothelial cells in the nearby existing vascular host breast tissue, get activated and also secrete a chemical known as fibronectin to enable them to adhere to the extracellular matrix whilst migrating towards the TAF gradient to penetrate the tumour and become vascularized. By analyzing a computational model, the process from the avascular stage to the phase of tumour-induced angiogenesis to the vascularization stage which makes the tumour more lethal to invade and metastasize is understood. A continuum mathematical model with coupled non-linear partial differential equations is developed and analyzed to describe the initial migration process of the endothelial cells during tumour-induced angiogenesis. The numerical simulations of the system which uses parameter values based on experimental data are presented. The results indicate that both the processes of chemotaxis and haptotaxis have significant impacts on the movement of the endothelial cells to the solid tumour in density in order to permeate and it vascularized.

Phebe Mawuena A. Havor
National Institute for Mathematical Sciences, Ghana
Kwame Nkrumah University of Science and Technology, Kumasi
enawoode@gmail.com

PP2

Dynamics of Disease Models with Self-diffusion: A Study of Cholera in Ghana

The attention to cholera epidemiology has increased, as its epidemics have become a worldwide health problem most especially in Ghana. A deterministic compartmental model is proposed with stability analysis on the epidemic and endemic equilibrium. A reaction-diffusion SIR-B mathematical model of cholera epidemiology that incorporates an environmental reservoir of V. cholerae is formulated to capture the movement of human hosts and bacteria in a heterogeneous environment. Here our findings are supported by the results of numerical simulations. Based on these results, we present the evolutionary processes that involves organism distribution and their interaction of spatially distributed population with local diffusion, and find that the model dynamics exhibits a diffusion-controlled formation growth to certain hole-like pattern replication that indicates that diffusion has a great influence on the spread of the cholera epidemic.

Phebe Mawuena A. Havor
National Institute for Mathematical Sciences, Ghana
Kwame Nkrumah University of Science and Technology, Kumasi
enawoode@gmail.com

Anas Musah
Kwame Nkrumah University of Science and Technology, Kumasi
bunmalik11@gmail.com

PP2

A High-resolution Godunov Method with HLLC Riemann Solver for Two-phase Modeling of Practical Condensed-phase Explosives

We present a high-resolution Godunov method for a two-phase model of reactive flow which considers general equations of state of Mie-Gruneisen form relevant for simulations of the dynamical behavior of detonations in PBX-type granular explosives. The numerical approach employs a Riemann solver, and various options ranging from an exact solver to an approximate solver of HLLC type are described. We provide an assessment of the approximate Riemann solvers in terms of the accuracy and efficiency of numerical solutions of the two-phase model. In particular, the state-sensitive behavior in different regimes of the solutions in a reverse-impact configuration are considered, including compaction, run-to-detonation, and compaction-led and reaction-led steady detonations. For these regimes, the convergence properties of the numerical approach is examined, as well as the behavior of the different approximate solvers at lower grid resolution.

Michael P. Hennessey
Rensselaer Polytechnic Institute
hennem3@rpi.edu

Donald W. Schwendeman
Rensselaer Polytechnic Institute
Department of Mathematical Sciences
schwed@rpi.edu

Ashwani K. Kapila
Rensselaer Polytechnic Inst.
kapila@rpi.edu

PP2

High-order Runge-Kutta Discontinuous Galerkin Methods with Local Time-stepping for Conservation Laws

We present high-order, explicit, local time-stepping methods for hyperbolic conservation laws. Predictor-corrector type algorithms are proposed in which the predictor is obtained based on the Strong Stability Preserving Runge-Kutta (SSP-RK) time discretization and Taylor expansions, and the corrector is derived to conserve mass exactly and to preserve the accuracy in time. These algorithms with spatially variable time step sizes restricted by local CFL conditions are presented in a general setting of Runge-Kutta discontinuous Galerkin methods with the modified minmod limiter. Numerical results on various test cases confirm the stability and accuracy of the proposed methods.

Thi-Thao-Phuong Hoang
Department of Mathematics and Statistics, Auburn University

tzh0059@auburn.edu

PP2

Fully Kinetic PIC-DSMC Simulations to Study Backflow of Ion Thruster Plasma Plumes using Hybrid MPI-CUDA Paradigm

Ion thruster plasma plumes consist of ions and neutrals which undergo charge-exchange (CEX) reactions to produce slow-moving CEX ions. These CEX ions flow backward and impinge on the solar panels due to the electric field induced between the plume and the spacecraft surfaces. Electrons, emitted from an external hollow cathode to neutralize the plume, also affect the electric field significantly altering the energy of the backflow ions. Predicting the ion energy distribution in the backflow region is critical to estimate the damage to the solar panels and their lifetime. Typically, electrons are not modeled due to the computational cost in PIC simulations of thruster plumes. Instead, their distribution is obtained from the Boltzmann relation which assumes quasi-neutrality and a single electron temperature. However, from our recent neutralization studies, we have found that when the electron source is shifted from the ion source, the electron velocity distribution is non-Maxwellian and the plume is not quasi-neutral. Therefore, in this work, we perform fully kinetic PIC-DSMC simulations using a three-dimensional, octree-based multi-GPU code, called CHAOS, to predict the ion energy distribution in the backflow region for real ion thruster plasma densities of $10^{15}/\mathrm{m}^3$ densities, and also study the effect of hollow cathode location on the backflow contamination. These results will also be compared with the Boltzmann relation predictions.

Revathi Jambunathan
University Of Illinois, Urbana-Champaign
jambuna2@illinois.edu

Deborah Levin
University of Illinois, Urbana-Champaign
deblevin@illinois.edu

PP2

CFD and Heat Transport Study of Varying Pebble Diameters in Pebble Bed Nuclear Reactors

Pebble Bed Nuclear Reactors (PBR) are commonly designed as helium gas cooled reactors with fuel designs that feature thousands of nuclear fuel elements encased in thousands 6 cm graphite pebbles. The effectiveness of the heat and fluid transport in the reactor is ultimately determined by the fluid interaction with the pebbles and their respective diameters, so it leaves the open question on whether or not the 6 cm pebble diameter is most effective design for the PBR. In this work a Computational Fluid Dynamics (CFD) study of a PBR is undertaken using Open source Field Operation and Manipulation (OpenFOAM) in order to analyze the fluid and heat transport in an idealized packing configuration. The pebble domains are arranged in a 81 pebble cubic packing configuration and the respective sizes tested are 3 cm and 6 cm models. For our simulations we use physical parameters such mass flow and system pressure from the HTR-PM reactor design, with a fixed pebble surface temperature and helium gas acting as the coolant. Analysis for comparing the advantages of the heat and fluid transport capabilities of each pebble model will include outlet fluid velocity, pressure and temperature, and the pressure drop, fluid temperature and velocity over column length is also included in results.

Andrew Jones
Boise State University
andrewjones237@u.boisestate.edu

Tien Yee
Kennesaw State Unviersity
tyee@kennesaw.edu

Eduardo Farfan
Kennesaw State University
efarfan1@kennesaw.edu

PP2

Validation of Assumptions in MILP Scheme for Determining Optimal Orificing in Fast Reactors using SAM

This study investigates the errors introduced into optimal orificing calculations of sodium fast reactors determined using a recently proposed mixed-integer linear programming method. The finite element code SAM is used to simulate the thermal-hydraulic behavior of a small two-assembly cluster of prototypical sodium-cooled fast reactor assemblies to determine the impact of assembly-to-assembly heat transfer on the coolant outlet temperatures. The heat transfer is quantified for different assembly powers and flowrates, and the sensitivities to various geometric parameters are examined. It is found that assembly-to-assembly heat transfer has a roughly 5% impact on the outlet temperatures as predicted neglecting the heat transfer, serving to smear the outlet temperatures and leading to conservative results from the optimization procedure.

Chris Keckler
University of California, Berkeley
keckler@berkeley.edu

Yves Robert
INSA Lyon
yves.robert@insa-lyon.fr

Massimiliano Fratoni, Ehud Greenspan
University of California, Berkeley
maxfratoni@berkeley.edu, gehud@berkeley.edu

PP2

Bayesian Inversion of Fault Properties in Two-phase Flow in Fractured Media

We address the inverse problem of inferring fault properties from well observations of two-phase flow in fractured media. The two-phase flow model in its mixed form is modified to account for fractures in the medium, thus permitting pressure discontinuities across the faults. Adopting a Fully Implicit Method (FIM) for solving the variational problem that arises, provides a path to computing

adjoint-consistent derivatives of observables with respect to fault parameters. We formulate the inverse problem in the framework of Bayesian inference, and seek the solution as a posterior probability distribution characterizing uncertainty in the inferred fault parameters. We begin by constructing a Laplace approximation of the posterior with mean given by the MAP point, and covariance given by the inverse of the Hessian of the negative log posterior. The MAP point is found by a scalable deterministic adjoint-based inexact Newton-CG optimization algorithm, while the inverse Hessian is made tractable by a low-rank approximation of the log likelihood, constructed via a randomized eigensolver. Finally, the Laplace approximation is employed as a preconditioner for a Crank-Nicolson MCMC method to sample the posterior distribution. The resulting exploitation of the smoothness and sparsity of the posterior leads to scalable algorithms whose cost scales independent of the parameter and data dimensions, and depends only on the number of informed modes of the parameters.

Eldar Khattatov
ICES, The University of Texas at Austin
ekhattatov@austin.utexas.edu

Omar Ghattas, Tan Bui-Thanh
The University of Texas at Austin
omar@ices.utexas.edu, tanbui@ices.utexas.edu

Ilona Ambartsumyan
The University of Texas at Austin, ICES
ailona@austin.utexas.edu

Umberto Villa
University of Texas at Austin
uvilla@wustl.edu

PP2

Flexible Parallel Mesh Partitioning Strategies for Large-scale Multiphysics Simulations

Large-scale parallel numerical simulations are essential for a wide range of engineering problems that involve complex, coupled physical processes interacting across a broad range of spatial and temporal scales. The data structures (matrices, meshes, etc.) of the underlying multiphysics simulations are often represented by graphs that need to be "optimally" partitioned for having an efficient parallel calculation. Partitions are required to minimize the number of edge-cuts while maintaining a balance in the amount of computation work associated to each processor core. Most existing partitioning software work well for a few hundreds of processor cores, but they are far from ideal when the number of processor cores goes beyond 10,000. Further more, the existing partitioners also ignore the existence of multi-core per compute node that is a standard architecture of modern supercomputers. In this work, we propose novel partitioning algorithms that take into consideration the core count associated to each compute node to preserve the data locality and minimize the inter-node communication while maintaining the work balance for each processor core. We numerically demonstrate that our new algorithms work well with more than 30,000 processor cores for complex multiphysics problems with billion unknowns.

Fande Kong
Department of Modeling and Simulation
Idaho National Laboratory
fdkong.jd@gmail.com

Derek R. Gaston, John Peterson
Idaho National Laboratory
derek.gaston@inl.gov, jw.peterson@inl.gov

Cody J. Permann
Advanced Modeling and Simulation
Idaho National Laboratory
cody.permann@inl.gov

Richard Martineau
Idaho National Laboratory
richard.martineau@inl.gov

PP2

Machine Learning Algorithm to Support Interpretability and Accuracy Simultaneously

Increasing accuracy, tractability, and interpretability are main Machine Learning problems. While deep learning algorithms exceeded the 99% accuracy on multiple datasets, the most accurate models are often the least interpretable. In criminal justice, health insurance, credit markets, and others, the degree of a model's accuracy is moot, for any decision without explanation. The 'right to explanation' established by the EU effectively disables use of the most accurate black-box algorithms. We introduce a False Prediction Pattern Recognition (FPPR) algorithm based on our previous Dominance Classifier Predictor (DCP) algorithm. FPPR is interpretable, converges in less-than quadratic time, and has shown high accuracies on benchmark datasets. DCP finds intervals within training data where a single class dominates, then constructs interpretable vote-weighting classification models using them. While simple, interpretable, and more accurate than several interpretable algorithms, DCP is less accurate than non-interpretable algorithms in our experiments. FPPR bridges the accuracy gap between these interpretable and uninterpretable algorithms by learning patterns found in training cases where DCP predictions failed. FPPR discovers relations between attributes and then correct DCP predictions using these relations. FPPR having achieved accuracy greater than 99.8% on three benchmark datasets using 10-fold cross validation, has reached the accuracy of non-interpretable algorithms.

Boris Kovalerchuk
Dept. of Comp. Sci.
Central Washington Univ.
borisk@cwu.edu

Nathan Newhous
Central Washington University
nathan.neuhaus@cwu.edu

PP2

Algorithm Design at Scale: Porting Parallel FFT-

based Fortran Simulations to GPUs

Large-scale scientific simulations involving parallel Fast Fourier Transforms (FFTs) have extreme memory requirements and high communication overhead. Modern hardware accelerators such as Graphical Processing Units (GPUs) provide a lot of compute power, can provide significant performance advantage for numerical simulations and are relatively inexpensive compared to large clusters. However, it is difficult to use GPUs to accelerate legacy Fortran scientific codes because of memory constraints of GPUs. A co-design of algorithm and software is necessary to scale scientific simulations using accelerators. We describe such a co-design solution for high-performance simulation of a FFT-based Hookes law partial differential equation (PDE) iterative numerical solver for large-scale datasets. The dataset consists of a microstructure volume made up of irregular grains with smooth stress and strain fields in the grain interiors. Our method uses irregular domain decomposition on the grain volumes and compression on the domains to reduce data communication between CPU and GPU without destroying accuracy. We use domain-local FFTs on the GPU followed by adaptive downsampling to distribute and summarize the output over the full volume. In this work, we discuss proof-of-concept results for various aspects of algorithm development including compression technique, iteration cost and convergence.

Anuva Kulkarni
Carnegie Mellon University
anuvak@andrew.cmu.edu

Jelena Kovacevic
New York University
jelenak@nyu.edu

Franz Franchetti
Department of Electrical and Computer Engineering
Carnegie Mellon University
franzf@ece.cmu.edu

PP2

A Fourth-order Accurate Multigrid Solver for Overset Grids

We describe a fourth-order accurate parallel multigrid solver for overset grids. The coarse level overset grids are automatically generated. We describe how the coarse-grid operator can be automatically generated by Galerkin averaging from the fine-grid operator. We also show, by both local Fourier analysis and numerical results, that a coarse-grid operator of second-order accuracy can be used to achieve convergence rates that are as good as those from using a fourth-order accurate coarse grid solver, with significantly more efficiency. The approach to boundary conditions is discussed to ensure consistency with the eigenstructure of the problem so that the convergence rate is not degraded. Instead of optimizing the smoothing rates, we determine optimal values of an over-relaxation parameter for the Red-Black smoother based on minimizing the overall multigrid convergence rate; this results in improved performance.

Chang Liu
Dept. of Mathematical Sciences,
Rensselaer Polytechnic Institute
liuc10@rpi.edu

William Henshaw
Rensselaer Polytechnic Institute
henshw@rpi.edu

PP2

Experiences Porting a Large-scale Deterministic Neutron Transport Application to Heterogeneous Architectures

In the current high performance computing world, machines are typically hybrid architectures composed of traditional CPUs connected to "accelerators". No single accelerator architecture has yet emerged as a clear winner. The high degree of heterogeneity of today's machines means HPC applications are forced to be performant over a wide range of architectures, i.e. they must be "performance portable". Ardra is a large deterministic neutron transport application designed to be highly scalable over large node-count machines, such as the Blue Gene architectures of prior years. With the advent of accelerator-based architectures, the Ardra team undertook a major rewriting effort to make it portable across the various upcoming new architectures. It was imperative that the code be able to exploit the performance gains possible with the new accelerators, particularly Lawrence Livermore's own GPU-based machines, while maintaining acceptable performance on traditional CPU-based machines. In this poster, we describe the technical decisions made along the way in converting Ardra to a performance-portable structure, describe the speedups achieved on new architectures, and discuss the impact to performance of those design choices on CPU-based machines. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

John Loffeld
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
loffeld1@llnl.gov

Teresa S. Bailey
LLNL
bailey42@llnl.gov

Peter Brown, Adam Kunen, Bujar Tagani
Lawrence Livermore National Laboratory
brown42@llnl.gov, kunen1@llnl.gov, tagani1@llnl.gov

PP2

Meshfree Semi-Lagrangian Schemes for Advection on Surfaces: Polyharmonic Splines Augmented with Polynomials

We present a new meshfree semi-Lagrangian method for simulating advection on a smooth two-dimensional manifold embedded in three-dimensional space. The method is based on a local RBF stencil interpolation method that uses polyharmonic spline kernels augmented with polynomials. Two key features of the method are that it does not

require a surface-based coordinate system or any artificial stabilization terms. One key challenge with the method is that the polynomial space for the manifold can be degenerate. We discuss how to overcome this issue and illustrate the effectiveness of the method on several examples.

Daniel Malmuth
Boise State University
dan.malmuth@gmail.com

Grady B. Wright
Department of Mathematics
Boise State University, Boise ID
gradywright@boisestate.edu

Varun Shankar
University of Utah
varun.shankar2001@gmail.com

PP2

Coupling an Explicit Variable Density Projection Method to Finite Rate Kinetics

A variety of industrially-relevant flows, such as those occurring in energy production systems utilizing coal or natural gas combustion, occur under conditions in which acoustic perturbations have a negligible impact on the system. However, stability requirements imposed by acoustic timescales put severe limitations on time step size when solving the unmodified Navier-Stokes equations with an explicit integration method. Use of pressure projection methods circumvents time step restrictions imposed by acoustic waves. In this work, an explicit pressure projection method for use on uniform grids is coupled to a generalized finite-rate approach to chemical kinetics and species transport where scalars are transported in conservative form. The approach is verified against manufactured solutions of reacting variable density flows.

Josh T. Mcconnell
University of Utah
Dept. of Chemical Engineering
josh.mcconnell@chemeng.utah.edu

James C. Sutherland
Department of Chemical Engineering
The University of Utah
james.sutherland@chemeng.utah.edu

PP2

A Time Adaptive Multirate Dirichlet-Neumann Waveform Relaxation Method for Heterogeneous Coupled Heat Equations

The efficient simulation of thermal interaction between fluids and structures is crucial in the design of many industrial products, e.g. the cooling of rocket thrust chambers. Unsteady thermal fluid structure interaction is modeled using two partial differential equations describing a fluid and a structure which are coupled at an interface. The standard algorithm to find solutions of the coupled problem is the Dirichlet-Neumann iteration, where the PDEs are solved separately using Dirichlet-, respectively Neumann boundary with data given from the solution of the other problem. Previous numerical experiments show that this iteration is fast for the air-steel coupling. This method has the disadvantage that both fields are solved with a common time resolution. Using instead a time adaptive multirate scheme would be more efficient. In view of this, we present here a time adaptive, multirate numerical method for two heterogeneous coupled heat equations. We use the Dirichlet-Neumann waveform relaxation (DNWR) method. When choosing the relaxation parameter right, the iterative DNWR becomes a direct solver. We present an analysis of the DNWR algorithm that shows that the optimal relaxation parameter is highly dependent on the material coefficients. In order to get an adaptive multirate scheme, we use possibly different adaptive time integration methods on the two subdomains. Furthermore, two alternatives are presented, the implicit Euler method and SDIRK2.

Azahar Monge
Lund University
Lund University
azahar.monge@na.lu.se

Philipp Birken
Centre for the Mathematical Sciences
Lund University, Sweden
philipp.birken@na.lu.se

PP2

Sparse Fourier Interpolation in Computational Chemistry

Stable molecular geometries correspond to local minima on a potential energy surface (PES) $E_i(\mathbf{q})$, where $\mathbf{q}$ contains bond lengths, bond angles, and dihedral angles, and i is the energy level of the molecule's electronic state. To study how molecular geometry evolves in time after excitation and relaxation, we must integrate a system of equations involving the PES and its gradient. The PES is prohibitively expensive to evaluate directly at each step of the integration, so we construct a reduced-order model $\tilde{E}_i(\mathbf{q})$ with sparse interpolation for use within the time integration. With respect to the angular components of $\mathbf{q}$, the true PES $E_i(\mathbf{q})$ should be at least 2π-periodic in both the function values and gradient. However, polynomial interpolation of $E_i(\mathbf{q})$ does not guarantee the periodicity of $\nabla\tilde{E}_i(\mathbf{q})$ with respect to the angular components of $\mathbf{q}$, which leads to non-conservation of energy over the course of the time integration. This poster presents sparse interpolation using a Fourier basis, along with chemistry results showing energy conservation when we interpolate $E_i(\mathbf{q})$ with a Fourier basis in the angular dimensions of $\mathbf{q}$.

Zachary B. Morrow
Department of Mathematics
North Carolina State University
zbmorrow@ncsu.edu

ELENA Jakubikova
North Carolina State University
ejakubi@ncsu.edu

Tim Kelley
Department of Mathematics

North Carolina State University
tim_kelley@ncsu.edu

Chang Liu
Department of Chemistry
North Carolina State University
chang_liu@ncsu.edu

Miroslav Stoyanov
Oak Ridge National lab
stoyanovmk@ornl.gov

PP2

Forecasting Volcanic Eruption Time using a Stochastic Enhancement of the Failure Forecast Method

Abstract not available

Abani K. Patra
Department of Mechanical & Aerospace Engineering
University at Buffalo
abani@buffalo.edu

PP2

Randomized Methods for Recompression of Low-rank Tensors

Many basic linear algebra operations with low-rank tensors, like addition, matrix-vector multiplication or element-wise product, have a tendency to significantly increase the rank, even though the resulting tensors admit a good low-rank approximation. We use randomized algorithm to recompress these tensors when dealing with low-rank formats such as Tucker and Tensor Train, by employing random vectors with rank-1 structure which matches the structure of the tensors. In case of element-wise product of tensors, this has shown to significantly reduce the computational effort, while achieving a similar accuracy as the corresponding deterministic techniques.

Lana Perisa
EPFL
lana.perisa@epfl.ch

Daniel Kressner
EPFL Lausanne
Mathicse
daniel.kressner@epfl.ch

PP2

A Continuous Model of Discrete Scientific Data

We investigate modeling discrete scientific data by a functional basis representation of high-dimensional multivariate nonuniform rational B-spline functions (NURBS). The functional model is an approximation that is more compact than the original discrete data, while having geometric and analytic properties making it useful for further analysis without converting back to the original. Our research has several main themes, highlighted in this poster. The first is computing the model in high dimensions—not just curves or surfaces but hypervolumes—iterating over separable dimensions and resulting in lower computational complexity than modeling all the dimensions together. The second is separate models for the spatiotemporal domain and for each of the scientific attributes, allowing the domain geometry and science variables to each be stored optimally. The third is an adaptive algorithm that adds knots, control points, and weights until the maximum error everywhere is below a threshold set by the user, allowing different levels of refinement in local regions of the data. The fourth is parallelization on high-performance supercomputers, with user-specified levels of continuity across subdomain boundaries. The fifth is precise evaluation of values and high-order derivatives anywhere in the domain, not only at the input data points. We evaluate our model using both synthetic and actual scientific datasets.

Tom Peterka, Youssef Nashed
Mathematics and Computer Science Division
Argonne National Laboratory
tpeterka@mcs.anl.gov, ynashed@anl.gov

Iulian Grindeanu, Vijay Mahadevan
Argonne National Laboratory
iulian@mcs.anl.gov, mahadevan@anl.gov

Raine Yeh
Purdue University
yeh@gmail.com

Xavier M. Tricoche
Computer Science Department
Purdue University
xmt@purdue.edu

PP2

Using Push-forward Measures for Parameter Identification under Uncertainty

Here we present novel developments for the solution of stochastic inverse problems using push-forward and pull-back measures. We extend previous work that focused on quantifying parameter variability to perform parameter identification under uncertainty. The previous work focuses on the transformation of distributions using discrepancies in push-forward and observed measures in contrast to updating prior beliefs with likelihood functions. The novel developments presented here focus on the extension of this approach for identifying a single "true" parameter from the aggregation and direct use of noisy data.

Michael Pilosov
University of Colorado: Denver
michael.pilosov@ucdenver.edu

PP2

Multi-level and Multi-index Monte Carlo Discontinuous Galerkin Methods for Uncertainty Quantification of Nonlinear Hyperbolic Problems

Many physical phenomena are modelled by hyperbolic conservation laws. As a result of measurement noise and uncertainties in model-driven factors such as initial conditions, boundary conditions, domain geometry and other

model inputs, it is required to quantify uncertainties in the solution of hyperbolic problems. Towards this goal, the Monte Carlo method is one of the most popular approaches in handling high-dimensional and nonlinear uncertainties, but it comes with a high computational cost due to repeated model evaluations. The purpose of this work is to investigate and develop a suite of accelerated Monte Carlo methods for quantifying uncertainties in hyperbolic systems. In this work, we exploit combining the discontinuous Galerkin (DG) method with multi-level and multi-index Monte Carlo methods. We employ the DG method for the discretisation of the conservation laws. We then integrate DG with multi-level and multi-index Monte Carlo methods for accelerating the uncertainty quantification. Various numerical experiments, including for Burgers equation and the shallow water equations with uncertainties, are used to demonstrate the efficiency of our methods.

Stanislav Y. Polishchuk, Hans De Sterck, Tiangang Cui
Monash University
stanislav.polishchuk@monash.edu,
hans.desterck@monash.edu, tiangang.cui@monash.edu

PP2

A Highly Scalable, Fault-tolerant Implementation of the Sparse Grid Combination Technique

In this poster session we give an overview on the current advancements in the ExaHD project which uses the Sparse Grid Combination Technique to solve high-dimensional time-dependent PDEs. We focus on the implementational aspects of the Sparse Grid Combination Technique for massively parallel simulations and the integration into a fault-tolerant framework. The Combination Technique introduces a second level of parallelism, replacing a single simulation run by several independent runs. A crucial step of the method is the combination of the single grids which results in a sparse grid representation of the result. Due to its hierarchical nature, the Sparse Grid Combination Technique provides a unique approach to realize algorithm-based fault tolerance without the need of checkpoint-restart by calculating alternative combination schemes that exclude failed resources. A key aspect with respect to scalability is the efficient implementation of the combination step which is crucial to achieve good parallel performance for exascale simulations. Several optimizations based on a manager-worker approach and an efficient communication scheme demonstrate excellent scaling. Convergence and scaling results are shown for the application code GENE which solves the gyrokinetic equations to simulate hot plasma in a fusion reactor.

Michael Obersteiner
Technical University of Munich
oberstei@in.tum.de

Hans-Joachim Bungartz
Technical University of Munich, Department of Informatics
Chair of Scientific Computing in Computer Science
bungartz@in.tum.de

Dirk Pflüger
University of Stuttgart
Dirk.Pflueger@ipvs.uni-stuttgart.de

Theresa Pollinger
Universität Stuttgart
theresa.pollinger@ipvs.uni-stuttgart.de

PP2

Multigrid Preconditioning for Optimization-based Domain Decomposition of the Helmholtz Equation

Recent work by Draganescu et al. shows the low computational cost of using a preconditioned multigrid method to solve the control system for an optimization-based domain decomposition (OBDD) of elliptic PDEs. We present an adaptation of the OBDD method for solving the Helmholtz equation in a 2 dimensional bounded domain, again showing the efficiency of multigrid preconditioning. Numerical results show the effectiveness of our method for different wave numbers.

Michael A. Retzlaff
University of Maryland, Baltimore County
mretzla1@umbc.edu

Andrei Draganescu
Department of Mathematics and Statistics
University of Maryland, Baltimore County
draga@umbc.edu

PP2

Reconstruction of Sparse Sums with a Generalized Music Algorithm

Many inverse problems in science and engineering can be solved by reconstructing sparse sums of parameter-dependent functions. This is true, for example, in direction-of-arrival estimation for radio waves, in imaging of small scatterers with acoustic or electromagnetic waves, or in spectral density estimation. We introduce a non-iterative method for solving this class of problems which depends only on measured data and yields accurate solutions with low sensitivity to measurement noise. This technique is a generalization of the MUSIC (MUltiple SIgnal Classification) algorithm which was developed for direction-of-arrival estimation. Comparatively few measurements are required to apply this generalized MUSIC scheme and we discuss its flexibility with respect to sensor placement. We describe some relevant applications in signal processing and communications theory.

Jacob D. Rezac
Communications Technology Lab
National Institute of Standards and Technology
jacob.rezac@nist.gov

PP2

Different Measure Approximations for Efficient Constrained Multi-objective Optimization under Uncertainty

This work is devoted to tackle constrained multi-objective optimization under uncertainty problems. The SABBa framework has been developed in [1 - M. Rivier, P.M. Con-

gedo, *Surrogate-Assisted Bounding-Box Approach for Optimization Problems with Approximated objectives*, Technical Report RR-9155, 2018] and applied to uncertainty-based optimization in [2 - M. Rivier, P.M. Congedo, *Surrogate-Assisted Bounding-Box Approach Applied to Constrained Multi-Objective Optimization Under Uncertainty*, Submitted, 2018]. Through the use of Bounding-Box measure approximations coupled to a Surrogate-Assisting strategy, SABBa yields very robust and parsimonious results. A Bounding-Box (or conservative box) is defined as a multi-dimensional product of intervals centered on approximated objectives and containing the associated true values. This contribution aims at proposing other measure approximations such as sampling or gaussian approximation, that can yield independent or joint information, based on the underlying surrogate-model dimension. The use of surrogate model in coupled space or separated spaces has been investigated in [2] and shows best results when exploiting both coupled space correlations and the low-dimensional surrogate-assisting model simplicity. Hence, joint information for measure approximation should allow for more accurate Pareto Optimal Probability computations and thus better characterization of the Pareto-optimal area.

Mickaël Rivier
ArianeGroup Le Haillan
Inria Bordeaux Sud-Ouest
mickael.rivier@inria.fr

Pietro M. Congedo
INRIA Bordeaux Sud-Ouest (FRANCE)
pietro.congedo@inria.fr

PP2

Modeling of Low Salinity Water Flooding in Naturally Fractured Core Samples

In this work a methodology is presented for modeling and simulate oil recovery by low salinity water injection through a naturally fractured core sample. The fractured core sample is modeled with mixed dimensions elements, representing the fractures as elements of n-1 dimensions immersed in a porous matrix of n dimensions by isolated internal boundaries, taking into account interactions between the fractures and the surrounding porous media. The derived model is multiphase, composed of two phases and three components: oil, water and salt (NaCl). Salt is treated as an additional single component in the brine phase, based on the following physical considerations: salt is transported only within the brine phase by advection and diffusion. Relative permeability and residual oil saturation depend on salinity. The oil and water component, are assumed to be present only in their associated phases. The effect of the salinity reduction on the relative permeability curves is introduced as parametric changes in Modified Brooks Corey functions for each injection cycle. The model uses a interpolation between sets of parameters using an scaling factor that is equivalent to brine salinity fraction. For the numerical solution is applied a Control Volume Finite Element Method (Box-scheme) and its computational implementation is carried out in DuMuX. Finally, the model is numerically validated in a two-dimensional case study by experimental data on oil recovery and pressure drop.

Carlos A. Romano
Universidad Nacional Autónoma de México
Posgrado de Ingeniería
carlos_alberto@comunidad.unam.mx

Martin Diaz
Instituto Mexicano del Petróleo
mdiazv@imp.mx

PP2

Communication Lower Bounds for Computing a Matricized-tensor Times Khatri-Rao Product

The matricized-tensor times Khatri-Rao product (MTTKRP) computation is the typical bottleneck in algorithms for computing a CP decomposition of a tensor. In order to develop high performance sequential and parallel algorithms, we previously established communication lower bounds that identify how much data movement is required for this computation in the case of cubical dense tensors. We extend our previous work to communication lower bounds for this computation in the case of rectangular dense tensors. Additionally, we present a family of parallel algorithms that attain the lower bounds.

Kathryn Rouse, Grey Ballard
Wake Forest University
kathryn.rouse@inmar.com, ballard@wfu.edu

Nicholas Knight
UC Merced
nknight@ucmerced.edu

PP2

A Kernel-based Approach for Solving the Hamilton-Jacobi Formulation of the Vlasov-Poisson System

In this work, we propose a conservative, unconditionally stable, implicit framework for solving the Vlasov-Poisson system. By recasting the problem as a Hamilton-Jacobi (HJ) equation, see e.g., [P. J. Morrison, On the Hamilton-Jacobi Variational Formulation of the Vlasov Equation, Math-for-Industry Lecture Series 39, 6475 (2012)], we can utilize the newly developed techniques for HJ equations developed in [A. Christlieb, W. Guo, and Y. Jiang, A Kernel Based High Order "Explicit" Unconditionally Stable Scheme for Time Dependent Hamilton-Jacobi Equations, arXiv preprint arXiv:1802.00536, 2018]. We will discuss some of the details in the transformations used, as well as the consequences of this approach.

William A. Sands
University of Akron
sandswi3@msu.edu

Andrew J. Christlieb
Michigan State Univerity
Dept. of Comp. Math., Sci & Engr.
christli@msu.edu

Yan Jiang
Michigan State University
yanjiang@msu.edu

PP2

Sparse Grid Density Estimation Techniques for Clustering-based Collocation

This contribution investigates the application of clustering with adaptive sparse grids for stochastic collocation (SC). Recently, an alternative to Gaussian quadrature based on partitioning and clustering methods was proposed. The basic idea is to consider cluster centers as the nodes and the fraction of the data in each cluster as the corresponding weights, respectively. To this end, we extend the range of initially considered clustering methods by employing non-parametric clustering based on sparse grid density estimation (SGDE). One of the challenges in the context of SC is the efficient handling of both correlation and nonlinearity with the increase in input size. Adaptive SGDE grows only with the size of the grids employed, while being able to retrieve the characteristics of the underlying data distribution. Still, the regular threshold approach based on the obtained densities does not provide enough control on the number and sizes of the obtained clusters. Thus, we showcase several methods for extracting appropriate clusters using both use-case driven approaches, as well as known techniques for handling estimated densities. Our focus is two-fold. Firstly, we aim to implement an efficient SGDE clustering suitable for the application at hand. We apply our proposed SC algorithm on standard test functions and we compare our results with those yielded by existing methods. Secondly, we compare several techniques for partitioning and clustering using our estimated densities.

Paul Cristian Sarbu
Technical University Of Munich
Department of Informatics
sarbu@in.tum.de

Hans-Joachim Bungartz
Technical University of Munich, Department of Informatics
Chair of Scientific Computing in Computer Science
bungartz@in.tum.de

PP2

ETD Algorithm for Fractional Model of Pressure Distribution in Fractured Rocks

Fractured rocks are complex geological formations where flow paths are convoluted and location of discontinuities are unknown. Fractional diffusion models have proven to be efficient in describing the anomalous diffusion due to these fractures. We consider a fractional diffusion model for the pressure distribution in fractured rocks. The temporal derivative (sub-diffusion) accounts for the memory effect in the diffusion process which reflects trapping, existence of obstacles, cervices and cracks in the media. While spatial derivative (super-diffusion) accounts for the non-locality which incorporates long-range interactions depicting communications over large distances. We present an unconditionally stable exponential time differencing scheme for the model subject to suitable boundary conditions and initial data. Numerical experiments are provided to show the outcomes and advantages of our scheme compared to some existing ones. Effects of anomalous diffusion due to fractures will be identified through the numerical solutions obtained.

Ibrahim O. Sarumi
King Fahd University of Petroleum and Minerals
g201304910@kfupm.edu.sa

Khaled Furati
Department of Mathematics & Statistics
King Fahd University of Petroleum & Minerals
kmfurati@kfupm.edu.sa

Abdul Khaliq
Middle Tennessee State University
abdul.khaliq@mtsu.edu

PP2

Deep Learning - Neural Networks

In this project we applied various deep learning methods (convolutional neural networks) to identify the key seven human emotions: anger, disgust, fear, happiness, sadness, surprise and neutrality. We used the Kaggle (Facial Expression Recognition Challenge) and Karolinska Directed Emotional Faces datasets. The architectures we employed for our convolutional neural networks were VGG-16 and ResNet50. We used the support vector machine multiclass classifier as our baseline, which had an accuracy performance of 31.8%. To further improve our results, we leveraged ensemble and transfer learning techniques to achieve our best results. Thus, the accuracy using ensemble learning was 67.2% and with transfer learning was 78.3%, solid results given that the winner of the Kaggle Facial Expression Recognition Challenge had an accuracy of 71.2%, and those who ranked in the top 10 of the same competition only achieved accuracies starting at around 60%.

Alexandru Savoiu
Stanford University
Institute of Math and Computational Engineering
savoiu@stanford.edu

PP2

MPI-OpenMP Load Balanced Simulation of Inhomogeneous Particle Systems in ls1 Mardyn at Extreme Scale

N-body simulations are used in astrophysics for galaxy formation or in molecular dynamics to study droplet coalescence. These scenarios typically imply heterogeneous particle systems which have to be handled appropriately via load balancing. ls1 mardyn is a molecular dynamics program collaboratively developed by computer scientists and process engineers. Since a couple of years, a purely MPI-based load balancing approach exploiting k-d trees has been implemented in ls1 mardyn. In this contribution, we show how hybrid MPI+OpenMP load balancing techniques and especially performance measurements can increase the performance of ls1 mardyn and other n-body

simulation programs. We show scaling results of large-scale vapor-liquid equilibrium (VLE) simulations of the coalescence of two droplets within a vapor phase up to machine size of current HPC systems and compare the results produced by the new load balancing techniques of ls1 mardyn with results from the previous implementation. We further provide an outlook on performance enhancements, employing auto-tuning techniques.

Steffen Seckler, Nikola Tchipev
Technical University of Munich
Department of Informatics
steffen.seckler@tum.de, n.tchipev@tum.de

Matthias Heinen
Paderborn University
Faculty of Mechanical Engineering
mheinen@mail.uni-paderborn.de

Fabio A. Gratl
Technical University of Munich
seckler@in.tum.de

Hans-Joachim Bungartz
Technical University of Munich, Department of Informatics
Chair of Scientific Computing in Computer Science
bungartz@in.tum.de

Philipp Neumann
University of Hamburg, Germany
Scientific Computing
philipp.neumann@uni-hamburg.de

PP2

A Comparison of RBF-FD Methods for Solving Partial Differential Equations on Surfaces

Over the past five years a number of different radial basis function finite difference (RBF-FD) methods have been developed for solving partial differential equations on surfaces. In this poster, we make the first direct comparison of these methods on some standard test problems to highlight their potential strengths and weaknesses. We also develop a new method that avoids singularity issues when augmenting the kernel basis with low degree polynomials.

Sage B. Shaw
Boise State University
sageshaw@boisestate.edu

Grady B. Wright
Department of Mathematics
Boise State University, Boise ID
gradywright@boisestate.edu

Varun Shankar
University of Utah
varun.shankar2001@gmail.com

PP2

A Weighted Essentially Non-oscillatory Forward Semi-Lagrangian Scheme for Vlasov-Poisson Systems

The kinetic evolution of a collisionless electrostatic plasma is governed by the Vlasov-Poisson (VP) system and is described by the so-called distribution function, $f = f(t, x, v)$, a proxy containing averaged information from the exact inventory of discrete particles at every coordinate (x, v) in a continuous phase space. To date, most prominent numerical schemes for VP systems can achieve arbitrary high order accuracy on smooth data; however, these improvements have the adverse effect of manufacturing spurious oscillations near discontinuities. Since the VP system can drive even smooth initial conditions towards filamented structures with near-discontinuous transitions, their numerical resolution is crucial for simulating the correct physics. The weighted essentially non-oscillatory (WENO) idea has met this challenge by guiding the repurposing of several high order schemes to efficiently resolve discontinuities including the finite difference, finite volume, discontinuous Galerkin, backward semi-Lagrangian, and hybridized methods thereof. We add to this list the first forward semi-Lagrangian (FSL) WENO scheme by extending an efficient, arbitrarily high order, local FSL scheme known as the convected scheme by means of finite-difference-based WENO derivative calculations. We present a detailed study of WENO derivatives, and showcase 5th, 7th, and 9th order simulations on one- and two-dimensional hyperbolic problems including highly filamented solutions arising in VP systems.

David Sirajuddin, William Hitchon
University of Wisconsin Madison
Department of Electrical and Computer Engineering
david.sirajuddin@wisc.edu, whitchon@yahoo.com

PP2

Anisotropy in Two-dimensional and Plane Elasticity Bond-based Linear Peridynamics

Modeling material failure and damage is an essential consideration in engineering and materials science communities. In classical mechanics, the governing equations depend on spatial derivatives, which poses difficulties when discontinuities such as cracks develop. As a potential remedy, the nonlocal theory of peridynamics was proposed by Stewart Silling. Within the peridynamic framework we have developed linearized anisotropic models describing each of the eight symmetry classes in three-dimensional linear elasticity and the four symmetry classes in two-dimensional linear elasticity. From the three-dimensional anisotropic peridynamic model we further derived two-dimensional plane strain and stress models utilizing nonlocal analogues of the assumptions posed in classical plane strain and stress models. We conclude our work with numerical results of wave propagation in media with various symmetry types.

Jeremy Trageser, Pablo Seleson
Oak Ridge National Laboratory
trageserje@ornl.gov, selesonpd@ornl.gov

PP2

NFFT-based Fast Summation for the Graph Lapla-

cian of Fully Connected Networks

The graph Laplacian is a standard tool in data science, machine learning, and image processing. The corresponding matrix inherits the complex structure of the underlying network and is densely populated in certain applications. A typical task is the computation of a number of its eigenvalues and eigenvectors. Standard methods become infeasible as the number of nodes in the graph is too large. We propose to use Krylov subspace methods in combination with a fast summation approach based on the nonequispaced fast Fourier transform (NFFT) to perform dense matrix-vector products with the graph Laplacian in a fast way. The enormous flexibility of the NFFT algorithm allows us to embed the accelerated matrix-vector multiplication into Lanczos-based eigenvalues routines and iterative linear system solvers. We illustrate the feasibility and advantages of our approach on several numerical examples. In particular, we compare our approach with the Nyström method.

Toni Volkmer, Dominik Alfke, Daniel Potts, Martin Stoll
Chemnitz University of Technology
toni.volkmer@mathematik.tu-chemnitz.de,
dominik.alfke@mathematik.tu-chemnitz.de,
potts@mathematik.tu-chemnitz.de,
martin.stoll@mathematik.tu-chemnitz.de

PP2

Primal-dual Weak Galerkin Finite Element Methods for PDEs

Weak Galerkin (WG) finite element method is a numerical technique for PDEs where differential operators in the variational form are reconstructed by using a framework that mimics the theory of distributions for piecewise polynomials. The usual regularity of the approximating functions is compensated by carefully-designed stabilizers. The fundamental difference between WG and other existing methods is the use of weak derivatives and weak continuities in the design of numerical schemes based on conventional weak forms for the underlying PDE problems. Due to its great structural flexibility, WG is well suited to a wide class of PDEs by providing the needed stability and accuracy in approximations. This poster will outline a recent development of WG, called "Primal-Dual Weak Galerkin (PD-WG)". The essential idea of PD-WG is to interpret the numerical solutions as a constrained minimization of some functionals with constraints that mimic the weak formulation of the PDEs by using weak derivatives. The resulting Euler-Lagrange equation offers a symmetric scheme involving both the primal variable and the dual variable (Lagrange multiplier). PD-WG is applicable to several challenging problems for which existing methods may have difficulty in applying; these problems include the second order elliptic equations in nondivergence form, Fokker-Planck equation and elliptic Cauchy problems. An abstract framework for PD-WG will be presented for its potential in other scientific applications.

Chunmei Wang
Texas Tech University
chunmei.wang@ttu.edu

PP2

Sharp Uniform in Time Error Estimate on a Stochastic Structure-preserving Lagrangian Method and Computation of Effective Diffusivity in 3D Chaotic Flows

In this paper, we study the problem of computing the effective diffusivity for a particle moving in chaotic flows. Instead of solving a convection-diffusion type cell problem in the Eulerian formulation (arising from homogenization theory for the Fokker-Planck equation), we compute the motion of particles in the Lagrangian formulation, which is modeled by stochastic differential equations (SDEs). A robust numerical integrator based on a splitting method was proposed to solve the SDEs and a rigorous error analysis for the numerical integrator was provided using the backward error analysis (BEA) technique. However, the upper bound in the error estimate is not sharp. In this paper, we propose a completely new and sharp error analysis for the numerical integrator that allows us to get rid of the exponential growth factor in our previous error estimate. Our new error analysis is based on a probabilistic approach, which interprets the solution process generated by our numerical integrator as a Markov process. By exploring the ergodicity of the solution process, we prove the convergence analysis of our method in computing the effective diffusivity over infinite time. We present numerical results to demonstrate the accuracy and efficiency of the proposed method in computing effective diffusivity for several chaotic flows, especially the Arnold-Beltrami-Childress (ABC) flow and the Kolmogorov flow in three-dimensional space. arXiv: 1808.06309

Zhongjian Wang
Department of Mathematics
University of Hong Kong
ariswang@connect.hku.hk

Jack Xin
Department of Mathematics
UCI
jxin@math.uci.edu

Zhiwen Zhang
Department of Mathematics
University of Hong Kong
zhangzw@hku.hk

PP2

Spectral Velocity Discretization of the Vlasov Equation using Generalized Hermite Functions

The Vlasov equation models the evolution of a plasma in its external and self-consistent fields. It is an advection equation in phase space nonlinearly coupled to Maxwells equations. We consider a Galerkin discretization with Fourier basis in space and generalized Hermite functions in velocity. The proposed basis allows for exact integration and resembles the structure of the solution in velocity space. Two special cases of generalized Hermite functions, so-called symmetrically and asymmetrically weighted Hermite

bases, have been introduced in [Holloway, Spectral velocity discretizations for the Vlasov-Maxwell equations, 1996]. It has been shown that asymmetrically weighted basis allows for the exact conservation laws in discrete form [Delzanno, Multi-dimensional, fully-implicit, spectral method for the VlasovMaxwell equations with exact conservation laws in discrete form, 2015]. We introduce a theoretical framework for the generalized Hermite functions based on the recently developed HermiteGF theory [Yurova, Kormann, Stable evaluation of Gaussian radial basis functions using Hermite polynomials, 2017]. We investigate the influence of scaling and shifting the basis functions on the numerical solution of the Vlasov equation. We study different choices of the parameters in order to find a setup with a good representation of the initial distribution and, at the same time, long-term numerical stability.

Anna Yurova
Max Planck Institute of Plasma Physics
Numerical Methods in Plasma Physics
anna.yurova@ipp.mpg.de

Katharina Kormann
Max Planck Institute for Plasma Physics
katharina.kormann@ipp.mpg.de

Caroline Lasser
Technische Universität München
classer@ma.tum.de

PP2

Exploration of Numerical Precision in Deep Neural Networks

Reduced numerical precision is a common technique to reduce computational cost in deep neural network and to speed up data transfer for distributed algorithms. While it has been observed that deep neural networks are resilient to small errors and noises, not enough investigations had been done to evaluate the sensitivity and numerical stability of a general deep neural network architecture of reduced precision. In this paper, we emulated arbitrary bit-width using a specified floating-point representation via truncation, which was applied to the neural network weights after each batch. We explored the impact of several key model parameters on the networks training accuracy, including batch sizes, rounding schemes, number of layers, and number of dense units and showed results on MNIST and CIFAR10 datasets. We then present a preliminary theoretical investigation of the error scaling in both forward and backward propagations. Our results indicate that caution must be taken for future reduced precision applications. We end with a discussion of the implications of these results as well as the potential for generalization to other network architectures.

Yunkai Zhang
UC Santa Barbara
yunkai_zhang@ucsb.edu

Yu Ma
UC Berkeley
UC Berkeley
midsummer@berkeley.edu

Zhaoqi Li
University of Washington
zli@macalester.edu

Catalina Vajiac
University of Notre Dame
cvajiac@nd.edu

PP2

A Novel Method to Average Images by Averaging Diffeomorphisms

Applied on our method for Image Registration[New Development of Nonrigid Registration, Hsiao, etc, 2013], diffeomorphic transformations can be constructed between some similar given images. A geometrical connection in terms of the Jacobian determinant and curl-vector of the constructed deffeomorphic transformations had been realized in our method of averaging diffeomorphisms[New method of averaging diffeomorphisms based on Jacobian determinant and curl-vector, Chen, etc, 2016]. In this work, we combine our components and build a novel approach for averaging the similar given images. An algorithm is provided and numerical examples are analytically displayed.

Zicong Zhou
University of Texas at Arlington
Department of Mathematics
zicong.zhou@mavs.uta.edu

Guojun Liao
Department of Mathematics
University of Texas at Arlington
liao@uta.edu

PP2

High Degree Discontinuous Petrov-Galerkin Immersed Finite Element Methods using Fictitious Elements for Elliptic Interface Problems

We propose a new strategy for constructing the p-th degree immersed finite element (IFE) spaces by applying the least squares framework on fictitious elements. This new construction method significantly reduces the ill-conditioning, caused by the small sub-element issue in our previous work, of solving the local IFE shape function. The proposed IFE spaces are employed in a discontinuous Petrov-Galerkin (DPG) scheme to solve the second order elliptic interface problems. We present a group of numerical examples to show that the DPGIFE method with the new p-th degree IFE space as the trial function space have the optimal convergence rate, which improves the numerical results reported in our previous work.

Qiao Zhuang
Virginia Tech
qzhuang@vt.edu

Ruchi Guo
Virginia Polytechnic Institute and State University
The Department of Mathematics

ruchi91@vt.edu

PP2

Monolithic Simulation of Binary Alloy Phase-change Coupled to Thermal and Compositional Convection

Melting and solidification are often affected by liquid convection, posing a multi-physics problem involving fluid flow, convective and diffusive heat transfer, and latent heat reactions. We model this as a nonlinear system of PDE's on a single domain using a variable-viscosity enthalpy method, and solve the system monolithically using a mixed finite element method (FEM) with global Newton linearization. We implemented this into an open-source code, Phaseflow [Zimmerman, A. G., & Kowalski, J. (2018). Monolithic Simulation of Convection-Coupled Phase-Change: Verification and Reproducibility. *Springer LNCSE*] using a FEM software library based on UFL [Alns et al. (2014). Unified form language: A domain-specific language for weak formulations of partial differential equations. *ACM TOMS*]. In this contribution, first we present an extension of the above model, method, and code to simulate eutectic binary alloys (e.g. sea-water) via coupling a solute mass balance as written in [Worster, M. G. (2000). Solidification of fluids. *Perspectives in fluid dynamics*]. Next, we demonstrate a technique for maintaining optimal convergence of the nonlinear solver by successively reducing the phase interface regularization. Finally, we present a series of simulations showing effects of the solute, including the concentration dependent liquidus and compositional convection.

Alexander G. Zimmerman
RWTH Aachen University
zimmerman@aices.rwth-aachen.de

Julia Kowalski
AICES
RWTH Aachen University
kowalski@aices.rwth-aachen.de

PP101

Minisymposterium: Randomized Least Squares Regression: Combining Model and Algorithm-induced Uncertainties

We analyze the uncertainties in the minimum norm solution of full-rank regression problems, arising from Gaussian linear models, computed by randomized (row-wise sampling and, more generally, sketching) algorithms. From a deterministic perspective, our structural perturbation bounds imply that least squares problems are less sensitive to multiplicative perturbations than to additive perturbations. From a probabilistic perspective, our expressions for the total expectation and variance with regard to both model- and algorithm-induced uncertainties, are exact, hold for general sketching matrices, and make no assumptions on the rank of the sketched matrix. The relative differences between the total bias and variance on the one hand, and the model bias and variance on the other hand, are governed by two factors: (i) the expected rank deficiency of the sketched matrix, and (ii) the expected difference between projectors associated with the original and the sketched problems. A simple example, based on uniform sampling with replacement, illustrates the statistical quantities.

Jocelyn Chi
NC State
jtchi@ncsu.edu

Ilse Ipsen
North Carolina State University
Department of Mathematics
ipsen@ncsu.edu

PP101

Minisymposterium: Finite Difference Moving Mesh Methods for PDEs on Curved Domains

We present a Monge-Ampére method for grid generation and adaptation for solving partial differential equations on curved domains using finite difference methods. For mesh generation, we develop a mapping F between a fixed rectangular computational domain Ω_C and a curved physical domain Ω_P, using a Hamilton Jacobi equation to define the boundary mapping. For dynamic mesh adaptation, we solve a parabolic regularization of the Monge-Ampére equation which increases mesh density in areas of fine scale solution behavior based on a monitor function derived from the underlying solution of the PDE. This method can resolve dynamic fine scale behavior of PDEs on curved domains, while retaining the simplicity of performing all computations on a fixed rectangular grid. We display the efficacy of these new methods on a variety of linear and nonlinear PDE examples.

Kelsey Dipietro
University of Notre Dame
kdipiet1@nd.edu

Alan E. Lindsay
Applied Computational Mathematics and Statistics
University of Notre Dame
a.lindsay@nd.edu

PP101

Minisymposterium: Fast Algorithms for Cosmic Microwave Background Radiation Data on Healpix Points

Faintly glowing at the edge of the observable universe, the Cosmic Microwave Background Radiation (CMBR) represents the first light to travel during the early stages of the universe's development. After about 379,000 years, the forming universe cooled enough to allow photons to move freely through space, creating a sphere of relic radiation that gives the strongest evidence for the Big Bang theory to date. Further analysis of the CMBR data can lead to revolutionary developments in understanding the nature of dark matter and dark energy. Since the discovery of the CMBR in 1964, scientists have worked to measure it in full detail using a Hierarchical Equal Area isoLatitude Pixelization scheme on the sphere. While these "HEALPix" points allow for a quasiuniform discretization of the sphere,

they are not well suited for the fast algorithms necessary for mining the prodigious amounts of CMBR data. For this work, we apply the Double Fourier Sphere method and fast Fourier transforms for uniform and non-uniform point distributions in order to transform the HEALPix grid to a Cartesian grid. We then utilize a state-of-the-art method which computes spherical harmonic coefficients from those of a bivariate Fourier series. This algorithm enables us to quickly and more accurately compute the angular power spectrum of CMBR data sampled at the HEALPix points. Results illustrating the effectiveness and advantages of these algorithms over current methods are presented.

Kathryn P. Drake
Boise State University
kathryndrake@u.boisestate.edu

Grady B. Wright
Department of Mathematics
Boise State University, Boise ID
gradywright@boisestate.edu

PP101

Minisymposterium: Classification of Vascular Disease Based on the Persistence Diagram using Topological Data Analysis (TDA) of Vascular Data

Vascular disease is a leading cause of death worldwide. The disease is measured by the percent stenosis and fractional flow reserve. Recently, a new measure using TDA based on the fundamental projection of vascular data was proposed. We use persistence diagrams (PD) from the fundamental projection for classification. We apply the kernel method to the PD, obtain the 2D image, and classify using machine learning. We present numerical examples using both model and experimental data.

Megan Johnson
The University at Buffalo, State University of New York
meganjoh@buffalo.edu

PP101

Minisymposterium: A New Theory for Movement of Surface Meshes

Adaptive mesh refinement has been proven to be an extremely useful tool in the numerical solutions of partial differential equations (PDEs). This is because when solutions are calculated numerically, they are often limited to a predetermined mesh however, many problems do not require a uniform precision in the mesh. A more accurate solution may be obtained if specific regions of the mesh that require higher precision contain a higher concentration of mesh elements. Meshing functionals, and the theory within, can be used to adapt the mesh in this manner. I will present a new theory for variational mesh adaptation on surfaces for which, to our best knowledge, it is the only surface moving mesh method that can be directly applied to a general surface. A surface functional and a new surface moving mesh PDE method will be formulated. I will then discuss various theoretical results that ensure the mesh does not become tangled over all time. Finally, various numerical results will be presented in two and three-dimensions which include adapting the initial mesh to a uniform mesh as well as adapting the initial mesh to a curvature focused mesh.

Avary Kolasinski
The University of Kansas
avaryk@ku.edu

Weizhang Huang
University of Kansas
whuang@ku.edu

PP101

Minisymposterium: A Randomized, Inexact, Newton-based Approach for Quantitative Photoacoustic Tomography

Quantitative photo-acoustic tomography (QPAT) is an imaging modality that is used in breast and brain imaging. The goal of QPAT is to recover the spatial distribution of the absorption coefficient from photo-acoustic images which can be mathematically formulated as a non-linear inverse problem. We develop several preconditioned, inexact, Newton-based solvers for this inverse problem. We study various aspects of the solvers such as choice of preconditioner, choice of stopping criteria, the behavior with mesh refinement and increased number of sources. To address the computational cost of multiple sources, we employ a randomization-based approach. The performance of the solvers is demonstrated through a synthetic model problem from QPAT.

Katrina Petroske
North Carolina State University
kepetros@ncsu.edu

PP101

Minisymposterium: Troll Hunter: Understanding Swarm Behavior in Social Networks

Abusive speech online results in physical harm and manipulation of public discourse. Challenges in detection and tracking of these behaviors arise from the datas scale and linguistic variability. Utilizing multiple machine learning linguistic models and graph theoretical approaches to analyze sub-reddit threads on reddit, we inferred topic network structure and dynamic characteristics correlated with abuse. This will aid in identifying topics missed by our language models and near real time event detection.

Chrysm Watson Ross
University of New Mexico,
email not on file 3

PP101

Minisymposterium: Computationally Efficient Fast Optimization over Time-varying Directed Graphs

Distributed optimization has been recently an active area of research in multiple communities including control and systems theory and operations research, as more and more applications are arising in many areas of engineering, science, and economics. One of the fundamental problems of

interest in distributed optimization is the consensus optimization problem where a network of autonomous agents collaboratively optimizes a sum of private local objective functions while communicating according to a certain communication scheme, abstracted as a graph. In this work, we investigate the consensus optimization problem over multi-agent systems where the inter-agent communication topology is both directed and time-varying. In this context, we present a computationally efficient fast distributed optimization algorithm, TV-AB, over time-varying directed graphs. Our main contribution is that our approach, unlike most of the existing approaches over directed graphs, does not require the estimation of the Perron eigenvector of the underlying weight matrices, a technique referred to as push-sum, which imposes additional computation and communication burden on the agents and introduces further nonlinearity in the algorithm. We show that under the standard assumptions of strong-convexity of the global objective function and Lipschitz-continuity of local gradients, the proposed algorithm achieves a geometrical convergence to the global optimum provided a sufficiently small step-size.

Fakhteh Saadatniaki
Department of Electrical and Computer Engineering
Tufts University
Fakhteh.Saadatniaki@tufts.edu

PP101

Minisymposterium: Detecting Novel Structural Variants in Genomes by Leveraging Parent-child Relatedness

Genomic variation shared by members of the same species that are longer than a single nucleotide are commonly called structural variants (SVs). Though relatively rare, they represent an increasingly important class of variation as SVs have been associated with diseases and susceptibility to some types of cancer. Common approaches to SV detection require the sequencing and mapping of fragments from a test genome to a high-quality reference genome. Candidate SVs correspond to fragments with discordant mapped configurations, but because errors in the sequencing and mapping will also create discordant arrangements, many of these predictions will be false. When sequencing coverage is low, distinguishing true SVs from errors is even more complicated. In recent work, we have developed SV detection methods that simultaneously consider the genomes of closely related individuals – parents and children. Our approaches control false positive SVs by requiring children inherit all SVs in their genome from a parent. However, in doing so we may have missed true novel variants acquired by the child. In this work, we generalize our previous approaches to allow the child to carry novel variants but enforce sparsity through an ℓ_1 penalty (since novel SVs in the child should be rare). We present results on both simulated genomes as well as two-sequenced parent-child trios from the 1000 Genomes Project.

Melissa Spence
University of California, Merced
mspence5@ucmerced.edu

PP101

Minisymposterium: Modeling the Effects of Macrophages on Bone Fracture Healing

A new mathematical model is presented to study the effects of macrophages on the bone fracture healing process. The model consists of a system of nonlinear ordinary differential equations that represents the interactions among classically and alternatively activated macrophages, mesenchymal stem cells, osteoblasts, and pro- and anti-inflammatory cytokines. A qualitative analysis of the model is performed to determine the equilibria and their corresponding stability properties. Numerical simulations are also presented to support the theoretical results and to monitor the evolution of a broken bone for different types of fractures under various medical interventions. The model can be used to guide clinical experiments and to explore possible medical treatments that accelerate the bone fracture healing process either by surgical interventions or drug administrations.

Imelda Trejo
The University of Texas at Arlington
imelda.trejo@mavs.uta.edu

PP101

Minisymposterium: Fractional Derivatives and Laplacians in One and Two-sided Weighted Sobolev Spaces

The most general way to define Fractional derivatives and Laplacians is in the distributional sense. We will show that the pointwise formulas for these operators hold for functions in more general classes. The classes we consider are weighted Sobolev spaces with two-sided Muckenhoupt weights and one-sided Sawyer weights. The latter capture the one-sided nature of fractional derivatives. The pointwise and norm limits as the orders of the derivatives converge to an integer are also analyzed.

Pablo Stinga, Mary Vaughan
Iowa State University
stinga@iastate.edu, maryo@iastate.edu

PP101

Minisymposterium: AWM Selection of the Regularization Parameter in the Ambrosio - Tortorelli Approach to Image Segmentation

Image segmentation is a process to find the boundary sets of a given image, which can be approached by multiple methods, such as edge detection and dual clustering method. In this poster, we present a partial differential equation-based approach, which is based on the Ambrosio-Tortorelli approximation of the Mumford- Shah functional. Mumford and Shah (1989) propose a functional whose minimization leads to the optimal segmentation. However, the Mumford-Shah functional is inconvenient to carry out in practical computation due to its lack in regularity. Ambrosio and Tortorelli (1992) propose a phase-field regularization of the functional and show that it gamma-converges

to the original functional as the regularization parameter goes to zero. In actual computation, people find that the regularization term of the functional has physical dimension, and its choice can result in completely different results. Even worse, the functional is found not to gamma-converge to the Mumford-Shah functional in some cases. In this poster we present a strategy for choosing the regularization parameter for better segmentation effects, as well as numerical examples.

Yufei Yu
University of Kansas
no email on file 4

PP102

Minisymposterium: Modelling In-crib Drying of eAr Maize- A Case Study of Sunyani-West District

This study was aimed at improving the store-drying of maize in cribs, specifically in the Sunyani-West district in the Brong Ahafo Region of Ghana. A model containing three differ- ential equations developed from the drying rate equation, the mass balance equation and the energy balance equation was used to explain the drying process. These equations were solved numerically and simulated with data from the Sunyani-West district. The algorithm used was written in Python 2.7. The study provided information on the appropriate time for drying and duration for the drying periods. The work specified that maize haversted in the minor and major harvest season takes 2-3 months and 5 months to dry respectively. It also investigated the range of the drying rate constant for the Sunyani-West district as 0.0004-0.0008. This was achieved using the relative humidity data from the area of study.

Mercy Amankway
Montana State University
mleung@shinstitute.org

PP102

Minisymposterium: Joint Sequence Analysis Challenges: How to Handle Missing Values and Mixed Variable Types

This study focuses on developing methodologies to minimize the effects of incomplete data. Specifically, it hopes to reduced the noise and bias caused in categorical sequence data by data gaps. Some strategies investigated include choosing a substitution cost to replace missing values and deleting the missing values at the end of a sequence. Cluster validity metrics are used to determine the accuracy of the unsupervised clustering algorithms and t-SNE is employed to visualize clusters and age biases. It became clear that deleting missing values provided the best results, but all data sets are different. Thus this study recommends employing the studied procedures before conducting analysis on longitudinal sequence data to ensure the results are unbiased. After these tests optimize the data, clustering is conducted to understand the correlation between a persons state in life and their travel.

Alexandra Ballow
Youngstown University
alballow@student.ysu.edu

PP102

Minisymposterium: Deployment of Automatic Differentiation Package in R

ADOL-C (Automatic Differentiation by OverLoading in C++) is an useful tool for computing First and Higher order derivatives of vector functions that are defined by C and C++ source codes. This presentation will focus on the deployment of ADOLC in R environment in the form of a Package. Sri Hari Krishna Narayanan has already created a package which works on Mac system and other Unix variants of it. My project focused on creating the package which enables the package in Windows system. Previous approach of using the package requires prerequisite for Boost, ColPack and a software development tool Swig to be pre installed which we would like to automate in order to make the process simple for the end user. My discussion will highlight issues of building the package autodiffadolc for Windows system architecture, show some motivating optimization examples after successful installation of autodiffadolc and possibly discuss some scope to solve optimization problems involving matrix functions.

Chitrak Banerjee
Michigan State University
banerj25@stt.msu.edu

PP102

Minisymposterium: Capsule Networks for Protein Structure Classifciation and Prediction

Proteins are responsible for most functions in our body. Determining their 3D structure is key to understanding how they work, why they cause diseases, and how to design drugs to block or activate their functions. Convolutional Neural Networks (CNNs) have been recently applied to structural biology. Capsule Networks introduce capsules, a new building block that better models the hierarchical relationships in the internal knowledge representations of a neural network. In this research, the implementation of two Capsule Network architectures is discussed and their performance is compared to the ones of conventional CNN architectures. Although slower than conventional CNNs, Capsule Networks produce better testing accuracy values.

Julian Cuevas
LBNL & University of Puerto Rico at Mayaguez
julian.cuevas1@upr.edu

PP102

Minisymposterium: Finite Difference Moving Mesh for Nonlinear PDEs in Curved Domains

We develop numerical methods for solving nonlinear PDEs in two-dimensional curved domains using finite difference methods. The method utilizes the Monge-Ampére equation to create a mapping between a fixed, rectangular computational domain on which the PDE is discretized onto the curved physical domain where the PDE is posed. We illustrate the effectiveness of the method on several example problems, including blow-up, moving and splitting do-

mains, and interface dynamics.

Kelsey Dipietro
University of Notre Dame
kdipiet1@nd.edu

PP102

Minisymposterium: Gaussian Processing for Coarse-grained Potential Development

Molecular Dynamic (MD) simulations provide insight into the physical basis of the dynamical and structural information of the molecular system. However, if the system of interest has a large number of atoms, those systems are difficult to tackle by simple atomistic MD simulations to sample all possible phases-spaces in the system to calculate important thermodynamics properties. To efficiently simulate these systems, the coarse-grained (CG) approach has been widely used to reduce the number of degrees of freedom in the system. The challenge in developing the CG potentials is to correctly map atoms into the groups. Here we are presenting Gaussian Processing (GP) to simplify the CG potential mapping method. The GP method is a robust regression analysis tool to estimate a mathematical relationship within the observed data. In my research, trajectories from a short atomistic MD simulation was used as a basis for GP analysis. Here we are presenting a new CG development scheme and then test it to small and then large molecular systems.

Zachary Douglas
University of Saint Mary
zachary.douglas@my.stmary.edu

PP102

Minisymposterium: In Situ Analysis with Apache Spark

My project aims to research how an in situ framework like SENSEI can be connected to Apache Spark using the Spark streaming library. This will enable scientific codes to send their simulation data to Spark where analysis using different algorithms can be run easily and efficiently. My goal is to also stream in parallel on multiple nodes of a cluster like Cooley at Argonne National Laboratory.

Aparna S. Gollakota
Loyola University Chicago
gollakota.aparna@gmail.com

PP102

Minisymposterium: Utilizing Graphics Techniques as a Means of Calculating Light Absorption in Staple Crops

As populations continue to increase, finding ways to increase crop yield in order to feed them is becoming more and more prevalent. Because attempting to find the best combination of light, water, etc. to grow the most nutrient-rich crops is not feasible, it is necessary to turn to computationally driven simulations to solve this problem. Light intensity (i.e., how much light is absorbed by a leaf of a plant) can be calculated by utilizing a graphics technique known as ray tracing to trace the path of a single ray of sunlight and realistically determine how much of it is absorbed, reflected, etc. in order to create a real world simulation. This work aims to show the benefits of using ray tracing in the computational plant biology field and how it is used within the scope of the Crops in Silico project run by the University of Illinois at Urbana-Champaign.

Colleen Heinemann
University of Illinois at Urbana-Champaign
cheinemann@lbl.gov

PP102

Minisymposterium: Checkpoints Compression for Adjoint Computation

When we want to understand the sensitivity of a simulation model with respect to an input value or to optimize an objective function, the gradient usually provides a good hint. The adjoint state method is a widely used numerical method to compute the gradient of a function. It decomposes functions into a sequence of basic operations. It performs a forward sweep to evaluate the function, followed by a backward sweep to calculate the gradient using the chain rule iteratively. One limitation of the adjoint state method is that all intermediate values from the forward sweep are needed by the backward sweep. Usually, we keep only a portion of those values, called checkpoints, in the memory because of limited space. The remaining values are either stored on the hard disk or recomputed from the nearest checkpoint whenever needed. In this work, we seek to compress the intermediate values in order to better utilize limited space in the memory and to speed the I/O when checkpointing to the hard disk.

Kai-yuan Hou
Northwestern University
khou2020@outlook.com

PP102

Minisymposterium: CFD Study of Varying Pebble Diameters in Pebble Bed Nuclear Reactors

Commercial Pebble Bed Reactor designs have maintained a constant fuel element diameter of 6 cm for previous reactors and the new fourth generation nuclear reactors, there is physical justification that a pebbles at larger or smaller diameters could not improve the reactor performance. This work aims is to analyze fluid and heat transport if the pebble diameter adjusted, we use computational fluid dynamics techniques to determine the effectiveness of a 6 cm pebble model and an compare with pebble model half the standard diameter.

Andrew Jones
Boise State University
andrewjones237@u.boisestate.edu

PP102

Minisymposterium: Finite-difference Time-domain Method for AZO/ZnO Multilayered 1D Structures

There is an ever-increasing demand for devices that are

both smaller and faster. This work aims to develop unique materials that can process light at the subwavelength scale to answer this demand. Here we present FDTD simulations of ultrashort pulse propagation in Zinc Oxide and Aluminum-doped Zinc Oxide multilayers at the epsilon-near-zero point to understand material effects. We focus on the epsilon-near-zero frequency due to promising recent applications in ultrafast switching, deformation-resistant pulse propagation, and heat management.

Priscilla Kelly
San Diego State University
priscilla.noreen@gmail.com

PP102

Minisymposterium: Network Traffic Performance Prediction with Multivariate Clusters in Time Windows

This project analyzes Tstat logs collected from Lawrence Berkeley National Laboratory's ESNet data transfer nodes to obtain insights on data transfer characteristics and behavior. Detecting anomalies in network transfers at the package level will provide solutions for improving network transfer rate. Several feature subsets from the Tstat logs were identified as good predictors for low network throughput. Dimensionality reduction was used to reduce the number of features and to select several sets of prominent features. K-means clustering provided a way to group data transfers by transfer quality. T-SNE was used to visualize and verify multi-variate clustering results in two-dimensional scatter plots. The results indicate that there is high correlation between the percentage of the smallest cluster of transfers and average throughput per time window for low throughput.

Tyler Leibengood
Youngstown University
tjleibengood@student.ysu.edu

PP102

Minisymposterium: Modification and Application of a Method for Studying Stability of High-speed Boundary Layers

In fluid mechanics, boundary layers are thin regions on a scale typically much smaller than the entire flow field where viscous forces dominate near the surface of a solid boundary. Understanding the stability of these layers is an important step in understanding how turbulence begins to develop in these layers. Recently there has been a focus on understanding how incessant random motion of molecules impacts the stability of high-speed boundary layers. In this work we take a computational approach to study the receptivity and stability of high-speed boundary layers to kinetic fluctuation. The approach is verified through comparison with published results obtained using a different approach. Using our approach, significant and interesting results regarding how and where the fluctuations occur were found.

Kevin Luna
University of Arizona
kevinluna@email.arizona.edu

PP102

Minisymposterium: Computational Tools for the Reconstruction of Atmospheric Aerosols via Spectroscopy

The reconstruction of atmospheric aerosols from multi-axis differential optical absorption spectroscopy is an ill-posed inverse problem. By understanding the underlying mathematical model for the forward problem, we are able to develop computationally efficient techniques for stable and efficient reconstruction of atmospheric aerosol profiles. Here we have investigated hybrid iterative regularization methods for reconstruction and high-performance computing for efficient computations.

Aimee Maurais
Virginia Tech
amaurais@vt.edu

PP102

Minisymposterium: New Gene Editing Capabilities Against Vector-borne Infectious Diseases

Vector-borne diseases are an increasing concern for the world and the United States - including La Crosse encephalitis, West Nile Virus, and Zika. Evolving gene editing techniques represent a promising tool to disrupt mosquito transmission of these diseases. This research aims to support an agent-based model that will simulate gene editing outcomes; accounting for vector, intermediate host, human interactions, and environmental influences.

Pedro X. Medina
University of Puerto Rico-Arecibo
pedro.medina5@upr.edu

PP102

Minisymposterium: PDE-informed Covariance for Model-form Uncertainty

Techniques for quantifying or reducing model-form uncertainty for computational mechanics models often require statistical modeling of one or more random fields. An integral component of the statistical model is the spatial correlation within the field. Here we present a method for obtaining physically meaningful covariance fields using the governing PDEs for the unclosed terms in the model. A covariance structure for the terms of interests is obtained by assuming a covariance structure for the unclosed terms of the PDEs and propagating it through the PDEs. To demonstrate this we use the Reynolds-averaged Navier-Stokes equations, assume a Gaussian process for the unclosed terms of the turbulent kinetic energy (TKE) equation, and propagate this to covariance in the Reynolds stresses through the PDE governing TKE transport. The resulting covariance in the Reynold stresses is not Gaussian and shows physically intuitive results such as strong correlation along streamlines (convection) as well as diffusion.

Carlos Michelen-Strofer

Virginia Tech
cmich@vt.edu

PP102

A Low Order Mimetic Finite Difference Method in 2-D For Resistive MHD

2-D For Resistive MHD Abstract: We developed a low order Mimetic Finite Difference (MFD) discretization for the equations of magnetohydrodynamics (MHD) in two space dimensions. These equations describe the evolution of the electric and magnetic fields in the presence of prescribed velocity field. The method is designed to work on general polygonal meshes and preserves the divergence-free condition on the magnetic field. The electric field is discretized at the vertices/nodes and the magnetic field uses edge-based discretization. The method reconstructs the magnetic field to extract nodal values necessary to approximate some terms present in Ohm's law. We test the robustness of our numerical scheme on three different types of meshes: with triangular elements, quadrilateral and unstructured polyhedrons obtain from a Voronoi tesselation. Analysis of the convergence for each of the aforementioned mesh types is presented. We finish with a test problem that shows the method is capable of modelling magnetic reconnection.

Sebastian Naranjo
Oregon State University
mleung@shinstitute.org

PP102

Minisymposterium: Numerical Study of Impact of Interparticle Interactions on Particle Dynamics in Lab Generated Spark Discharge Plasma

Our goal is to numerically study the dynamics of the micron-sized particles following a spark discharge. To model the spark discharge, we use the non-reactive Navier-Stokes equations using U. Chicago code, FLASH. Particles are treated as spherical, dry, electrostatically neutral and almost rigid (small deformations are permitted), and are tracked in a Lagrangian fashion. To model the particle physics, an efficient collision model based on the hard sphere model has been implemented in Particles Unit of FLASH. Additionally, cohesion between the particles owing to attractive van der Waals forces has also been taken into account via extending the hard sphere model to include an energy based agglomeration criterion. By including such particle-interaction events, we can study the effect of collisions, agglomeration and possible deagglomeration on the particle dynamics. This can be a significant mechanism of particle lifting and clumping as observed in SEM images of some experiments. The particle volume fraction is varied to find its relationship with collision frequency.

Pooja Rao
University of Illinois at UrbanaChampaign
poojarao12@gmail.com

PP102

Minisymposterium: Capsule Networks for Protein Structure Classification

Proteins are responsible for most functions in our body. Determining their 3D structure is key to understanding how they work, why they cause diseases, and how to design drugs to block or activate their functions. Convolutional Neural Networks (CNNs) have been recently applied to structural biology. Capsule Networks introduce capsules, a new building block that better models the hierarchical relationships in the internal knowledge representations of a neural network. In this research, the implementation of two Capsule Network architectures is discussed and their performance is compared to the ones of conventional CNN architectures. Although slower than conventional CNNs, Capsule Networks produce better testing accuracy values.

Dan Rosa
University of Puerto Rico
dan.rosa@upr.edu

PP102

Minisymposterium: In Situ Performance Analysis of HPC Network Simulations

Optimistic parallel discrete event simulation (PDES), which allow events to be processed speculatively and provide mechanisms to recover from causality errors, can be highly scalable but difficult to tune so that time spent rolling back the simulation to fix causality errors is minimized. This work aims to make performance tuning of optimistic PDES easier by providing in-situ analysis and visualization of model and simulation performance data, which will help users to detect and understand the causes of performance bottlenecks. However, most of our analyses have been performed for relatively small-scale simulations, since collecting large-scale simulation performance data at a high-resolution result in more data output than is feasible to store on disk or transfer to other machines for visualization. We have developed our in-situ analysis system to perform feature extraction on the simulation engine data to help reduce the high-resolution data to a more manageable size that can be streamed to a visual analysis tool. It also detects potential performance problems in the simulation to determine some subsets of the higher resolution data that should be streamed for visualization, for the user to examine more detailed data for specific parts of the simulation where performance bottlenecks are found. We evaluate the usefulness of our in-situ analysis system by performing case studies with two high-performance computing (HPC) network simulations.

Caitlin Ross
Rensselaer Polytechnic Institute
rossc3@rpi.edu

PP102

Minisymposterium: Molecular Communication in Biological Cells: Foundational Study and Development of Computational Techniques

The focus of this work is to understand the limits in the controllability of a natural (non-engineered) biological cells behavior. We characterized the chemical activities that

take place inside the cell to provide energy and basic components, also known as cell metabolism, into two channel abstractions namely, the enzyme expression regulation and the metabolic reaction network. The enzyme expression regulation channel models the mechanisms underlying the regulation of cell metabolism as a function of the chemical compounds in the cells environment. The metabolic reaction network channel models the cell metabolism as a function of the exchange of metabolites and growth (biomass) of the cell. The composition of these two channels together quantifies the information flow in the end-to-end molecular communication system, where the exchange of metabolites and growth of the cell are modeled as a function of the chemical compounds in its environment. Based on these abstractions, the potential of cell metabolism is characterized and quantified in terms of the information-centric steady-state mutual information parameter. This upper bound quantifies how much information in bits can be transmitted from input to output in the case of both channels. We performed simulations and presented numerical results of our proposed models.

Zahmeeth Sayed Sakkaff
University of Nebraska, Lincoln
sszahmeeth@gmail.com

PP102

Minisymposterium: A Comparative Analysis of Parallel Louvain Algorithms for Community Detection

Community detection is of great importance in graph (network) mining. Louvain algorithm is an efficient algorithm for discovering communities. However, due to the emergence of large-scale network data, parallel algorithms leveraging high performance computing platforms are necessary. Although there are several shared memory-based parallel algorithms for Louvain method, those do not scale to a large number of cores and large networks. The existing only Message Passing Interface (MPI) based distributed-memory parallel implementation of Louvain algorithm has shown scalability to only 16 processors. In this work, first, we design a shared-memory based algorithm using Open MultiProcessing (OpenMP), which shows 4x speedup but is limited to the physical cores available to system. Our MPI-based distributed-memory parallel algorithm demonstrates scalability to a moderate number of processors. We then implement a hybrid algorithm combining both shared and distributed memory-based approaches. Our Hybrid Algorithm strikes a balance between both systems. Finally, our MPI-based parallel Louvain algorithm along with a load-balancing scheme shows around 12x speedup and scalable to increased processing power. Our load-balancing scheme minimizes the communication overhead to a large extent. We present a comparative analysis of these different implementations of parallel Louvain Algorithms for several real-world networks.

Naw Safrin Sattar
University of New Orleans
nsattar@uno.edu

PP102

Minisymposterium: Collective I/O using RAM Area Network (RAN)

Collective I/O operations allow parallel applications to significantly improve I/O performance by merging the requests of different processes. These collective operations are typically implemented using two-phase collective aggregation where a subset of ranks are chosen to serve as aggregators. Noncontiguous data from each rank is organized into contiguous portions of the file within these aggregators and written to the file system. Our research explores using a RAM area network as an aggregation layer to these collective operations. In this approach data is pooled into remote memory where aggregators can collect portions of the file without coordinating with other ranks. The goal of this approach is to drastically reduce the synchronization cost among the ranks.

Sergio Servantez
Illinois Institute of Technology
sergio.servantez.esq@gmail.com

PP102

Minisymposterium: Analyzing and Evaluating Resilience of Scheduling Scientific Applications

Large scale systems provide a powerful computing platform for solving large and complex scientific applications. However, the inherent wide distribution, heterogeneity, and dynamism of the computing environments along with frequent occurring resource failures can lead to performance degradation of the scientific applications executing on these computing systems. In general, various scheduling approaches are employed to maximize scientific throughput and to maximize reliability. A high level modeling methodology based on simulation for the analysis and evaluation of resilience of the scheduling approaches in presence of resource failures will be presented in this research poster. With the aid of this model, a wide class of dependencies existing between applications and computing system will be captured for quantifying the performance impact expected from changes in application and system characteristics.

Sonny R Sevin
Slippery Rock University of Pennsylvania
srs1030@sru.edu

PP102

Minisymposterium: Quantum Local Search for Graph Community Detection

We present Quantum Local Search (QLS) approach and demonstrate its efficacy by applying it to the problem of community detection in real-world networks. QLS is a hybrid algorithm that combines a classical machine with a small quantum device. QLS starts with an initial solution and searches its neighborhood, iteratively trying to find a better candidate solution. One of the main challenges of the quantum computing in NISQ era is the small number of available qubits. QLS addresses this challenge by us-

ing the quantum device only for the neighborhood search, which can be restricted to be small enough to fit on near-term quantum device. We implement QLS for modularity maximization graph clustering using QAOA on IBM Q Experience as a quantum local solver. We demonstrate the potential for quantum acceleration by showing that existing state-of-the-art optimization solvers cannot find a good solution to the local problems quickly and provide an estimate of how larger quantum devices can improve the performance of QLS. We apply QLS to the problem of clustering microbiome co-occurence networks and present the preliminary results.

Ruslan Shaydulin
Clemson University
rshaydu@g.clemson.edu

PP102

Minisymposterium: Blow-up Behavior of Conservation Law with Spatially Varying Flux

We consider a conservation law containing a multi-valued graph. Using the change of variables, we obtain the conservation law with a non-smooth space-dependent flux. We study the regularization of flux to understand the blow-up behavior of our problem. We show the convergence of the numerical solution using the Godunovs scheme and the local phase behavior solver. One application of this problem is the methane gas transport in subsea sediments.

Choah Shin
Oregon State University
Department of Mathematics
shinc@oregonstate.edu

PP102

Minisymposterium: Utilization of the Polymerase Chain Reaction and DNA Barcoding Method in Bioinformatics for the Identification of Unknown Fish Species

Food fraud from species substitution is an emerging risk given the increasingly global food supply chain and potential food safety issues. Economic food fraud is committed when food is deliberately placed on the market, for financial gain, with the intention of deceiving the consumer. DNA barcoding has received substantial attention as an accurate and comprehensively applicable device for animal species identifications. DNA barcoding is a taxonomic framework organized on succession data from a short extend of a center DNA arrangement. A region of approximately 648-pb of the mitochondrial gene COI was initially proposed as the barcode source to identify and delimit all animal species. The methodology involves the sequencing a portion of DNA from a species, followed by a comparison with other sequences previously deposited in a database. Species are identified by matching the obtained sequence with sequences of known identity already in the database. COI seems to have an incredible scope of phylogenetic flag, demonstrating quick rates of nucleotide substitutions that empower the separation of mysterious species, as well as uncover phylogeographic structures inside animal groups. The aim of this study was to investigate the utility of the DNA barcoding method identifying two different fish species. The names of the fish were not known at the time of this experiment; it was our onus to use the DNA barcoding technique to confirm the type of fish species it was.

Kaayla Tippins
Jarvis Christian College
kaayla.tippins@student.jarvis.edu

PP102

Minisymposterium: Using Deep Learning on Medical Data for Suicide Prevention

We present a suicide prediction scheme and predict the possibility of readmission for each day after last discharge based on our empirical investigation using the MIMIC-III intensive care database. Our feature mapping is working on determine the high-risk suicide group and generated structured data which contains Lab Work, Diagnosis, Medical, Procedures and unstructured data which selected Physician's Notes, Discharge Notes and all the notes information. The unstructured data (doctor notes) can be encoded into vectors using embedding techniques such as word2vec or document2vec. This allows us to represent textual data as numerical vectors. Combined structured data and notes, feed them into deep learning models such as RNN, LSTM and Feedforward Model with Time-Aware Attention to output the prediction of suicide and readmission. Techniques such as Principal Component Analysis (PCA) and T-Distributed Stochastic Neighbor Embedding (t-SNE) can be used to perform a dimensionality reduction on the vectors which we are able to use for groupings and visualizations.

Xiange "Shirley" Wang
Hood College
shirleywang5269@gmail.com

PP102

Minisymposterium: Locally-implicit Discontinuous Galerkin Schemes with Limiters that Guarantee Moment-invertibility for Hyperbolic Quadrature-based Moment Closures

We consider the quadrature-based moment closure approach for approximating kinetic Boltzmann equations. The true distribution is replaced by Dirac deltas with variable weights and abscissas. We show how to construct these to obtain a set of conservation laws that are conditionally hyperbolic. We then develop a high-order numerical method with limiters that guarantee that the numerical solutions remain in the convex hyperbolic regions of solution space. Numerical examples in 1D and 2D will be presented.

Christine Wiersma
Iowa State University
Department of Mathematics
cwiersma@iastate.edu

PP102

Minisymposterium: Application of the Ensemble

Kalman Filter in Tsunami Wavefield Reconstruction

Currently, the most popular method for forecasting the tsunami waveform is optimal interpolation, which uses a Kalman filter like approach, but holds the Kalman gain matrix fixed to reduce the computational cost. Here, we demonstrate that more accurate and stable forecasts can be obtained using the ensemble Kalman filter, a more computationally efficient variant of the Kalman filter, in which the gain matrix is updated according to the physical model and the evolution of the error covariance matrix. The ensemble representation is a form of dimensionality reduction, in that only a small ensemble is propagated, instead of the joint distribution including the full covariance matrix. We use a scenario tsunami in the Cascadia subduction zone, generated from a 2D fully-coupled dynamic rupture simulation (Lotto et al., submitted 2018). Randomly perturbed tsunami wave height data is used in the data assimilation process, as we propagate the wave using a 1D linear shallow water code on a staggered grid. Better waveform agreement is achieved even in the early stages of assimilation. Through the dynamic rupture simulation, we are also able to examine the seafloor pressure time series due to various factors that affect the pressure change: seismic waves, ocean acoustic waves, seafloor deformation and the change in the wave height. Our methods will significantly reduce the runtime, taking us a step closer to reliable real-time tsunami early warning.

Yuyun Yang
Stanford University
Institute for Computational and Mathematical Engineering
yyang85@stanford.edu

PP103

Minisymposterium: hIPPYlib: An Extensible Software Framework for Large-scale Inverse Problems

We present an Inverse Problem PYthon library (hIPPYlib) for solving large-scale deterministic and Bayesian inverse problems governed by partial differential equations (PDEs). hIPPYlib implements state-of-the-art scalable algorithms that exploit the structure of the problem, notably the Hessian of the log posterior. The key property of the algorithms implemented in hIPPYlib is that the solution is computed at a cost, measured in forward PDE solves, that is independent of the parameter dimension. The mean of the posterior is approximated by the MAP point, which is found by minimizing the negative log posterior. This deterministic nonlinear least squares optimization problem is solved with an inexact matrix-free Newton-CG method. The posterior covariance is approximated by the inverse of the Hessian of the negative log posterior evaluated at the MAP point. This Gaussian approximation is exact when the parameter-to-observable map is linear; otherwise it can serve as a proposal for Hessian-based MCMC methods. The construction of the posterior covariance is made tractable by invoking a low-rank approximation of the Hessian of the log likelihood. hIPPYlib makes these advanced algorithms easily accessible to domain scientists and provides an environment that expedites the development of new algorithms. hIPPYlib is also a teaching tool that can be used to educate students and practitioners who are new to inverse problems and to the Bayesian inference framework.

Olalekan Babaniyi
University of California, Merced
obabaniyi@ucmerced.edu

Omar Ghattas
The University of Texas at Austin
omar@ices.utexas.edu

Noemi Petra
University of California, Merced
npetra@ucmerced.edu

Umberto Villa
University of Texas at Austin
uvilla@wustl.edu

PP103

Minisymposterium: Shaping a Sustainable SUNDIALS: Applying Software Sustainability Practices to a CSE Library

The SUNDIALS Suite of Nonlinear Differential-Algebraic Solvers and integrators is an open source software library including highly robust and adaptive time integration methods for ODEs and DAEs as well as robust nonlinear solvers. SUNDIALS has a long history of user deployment with over 17,000 downloads worldwide in 2017. Therefore, the sustainability of SUNDIALS is an area of focus for the SUNDIALS team. In this poster, we will discuss specific strategies such as modernizing our CMake build system, increasing test coverage, and increasing modularity within the library. We will also discuss continuing issues and future plans for strengthening SUNDIALS sustainability. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC. LLNL-ABS-755764.

Cody J. Balos, David J. Gardner
Lawrence Livermore National Laboratory
balos1@llnl.gov, gardner48@llnl.gov

Daniel R. Reynolds
Southern Methodist University
Mathematics
reynolds@smu.edu

Carol S. Woodward
Lawrence Livermore National Laboratory
woodward6@llnl.gov

Slaven Peles
Lawrence Livermore Natinoal Laboratory
peles2@llnl.gov

Alan Hindmarsh
Lawrence Livermore National Laboratory

hindmarsh@llnl.gov

PP103

Minisymposterium: Outreach for Better Scientific Software

The primary mission of the IDEAS Productivity project is to help code teams supported by the U.S. Department of Energy to improve their software development practices in order to improve scientific productivity and produce software that is more sustainable. But were not a large project, and DOE comprises many code teams; moreover, these concerns resonate with projects throughout the broader computational science and engineering (CSE) community. Thus, we recognized from the start that we would benefit from extensive engagement with the CSE community to identify and tailor best practices for CSE software, including how best to assist CSE projects in implementing them. Consequently, the IDEAS project includes an extensive outreach component. We organize and deliver trainings and tutorials; we produce a monthly webinar series on Best Practices for HPC Software Developers; and we organize technical sessions at meetings like SIAM CSE19 to promote the discussion of experiences with software development and the exchange of ideas within the community. We have also recently spun out Better Scientific Software (https://BSSw.io) as a community-driven portal to share and find resources to improve CSE software development and experiences. This poster will provide more information about these activities and how they support our core mission while engaging with and giving back to the broader community.

David E. Bernholdt
Oak Ridge National Laboratory
Computer Science and Mathematics Division
bernholdtde@ornl.gov

PP103

Minisymposterium: Using Software Best Practices to Advance Research Across the Aviation Industry

Forward-looking aviation research typically employs advanced computational simulation techniques to mature the state of the industry. Technologies are matured though cross-aviation industry stakeholder interaction that seeks to build consensus via a research exploration activity that is constantly shaping both operational and performance requirements. Among the stakeholder participants are technology researchers, avionics manufacturers, airframe manufacturers, and other interested parties. This environment presents multiple challenges for the research software engineers that are shaping the technology, presenting results, and preparing test data. The simulation becomes a constantly evolving entity that responds to changing committee requirements, the simulation scenarios are regularly in flux (occasionally producing feedback to the simulation), and the certification test data that is to be produced must remain stable and represent the final state of committee decisions. MITRE has been at the center of this research software challenge and has been employing best practices from the software engineering discipline to meet the aviation industrys needs. These best-practices include: semantic versioning, library-focused software architecture and development techniques, unit testing of essential algorithms, regression testing to ensure stable statistical results, use of continuous integration, and a continuous open-source release of fundamental simulation elements.

Stuart Bowman
MITRE Corporation
sbowman@mitre.org

PP103

Minisymposterium: Ginkgo: Designing a Single-Node Linear Operator Framework for High Performance Computing

An important aspect in the design of a numerical software ecosystem is software sustainability and to ensure ease of use and contributing. With the primary intention of providing a numerical linear algebra library to the scientific computing community, we design the Ginkgo linear operator library based on the the guidelines and policies of the Extreme-scale Scientific Software Development Kit (xSDK) and the Better Scientific Software (BSSw) initiatives. We follow the open source strategy with a modified BSD license, which does not restrict commercial use of the software. The main repository is publicly available on github and only prototype implementations of ongoing research is kept in a private repository. The github repository is open to external contributions through a peer-review system and uses issues for bug tracking and to bolster development efforts. A Continuous Integration system (CI) realizes the automatic synchronization of repositories, and the compilation and testing of the distinct branches. Configuration and the compilation process are facilitated with CMake. The testing is realized using Google Test and comprises a comprehensive list of unit tests ensuring the library's functionality. Finally, a documentation is automatically kept up-to date with the software, and multiple wiki pages containing examples, tutorials, and contributor guidelines are available.

Hartwig Anzt
University of Tennessee, USA
hanzt@icl.utk.edu

Goran Flegar
Universitat Jaume I
flegar@uji.es

Terry Cojean, Pratik Nayak
Karlsruhe Institute of Technology
terry.cojean@kit.edu, pratik.nayak@kit.edu

Enrique Quintana-Orti
Universidad Jaume I
flegar@uji.es

PP103

Minisymposterium: Reproducible Computational Scientific Workflows with Signac

Researchers in computational science are regularly posed with the challenge of managing and analyzing large, het-

erogeneous, and highly dynamic data spaces. We present *signac*, an open-source Python framework that enables researchers to efficiently operate on primarily file-based data spaces while keeping track of all relevant metadata. The *signac* framework provides all components required to create a well-defined, collectively accessible data space and to implement reproducible workflows. The software is designed to be highly modular, decoupling its data and workflow management components so as to minimize the effort required for integration into existing workflows. The serverless data management and lightweight workflow model ensure that workflows are just as easily executed on laptops as in high-performance computing environments. Using *signac* not only increases research efficiency, it also improves reproducibility and lowers barriers for data sharing by transparently enabling the robust tracking, selection, and searching of data by its metadata.

Bradley Dice
Department of Physics
University of Michigan, Ann Arbor
bdice@umich.edu

Carl S. Adorf
Department of Chemical Engineering
University of Michigan, Ann Arbor
csadorf@umich.edu

Vyas Ramasubramani
University of Michigan
vramasub@umich.edu

Paul Dodd
Department of Chemical Engineering
University of Michigan, Ann Arbor
pdodd@umich.edu

Sharon C. Glotzer
Chemical Engineering
University of Michigan
sglotzer@umich.edu

PP103

Minisymposterium: Slate: Developing Sustainable Linear Algebra Software for Exascale

The SLATE software package will be a replacement for the LAPACK and ScaLAPACK numerical libraries, which over the last two decades set the bar for quality, sustainability, and community engagement. Improving upon the accomplishments of these legacy packages in those capacities is one of the major objectives of the SLATE project. SLATE will accomplish these improvements by applying contemporary software engineering techniques in hosting, development, documentation, automated testing, and continuous integration, creating communication channels with academic and commercial application teams to actively engage the larger HPC community, and using modern tools that enable these practices and techniques.

Jamie M. Finney
Innovative Computing Laboratory
jmfinney@icl.utk.edu

Jakub Kurzak
University of Tennessee, Knoxville
kurzak@icl.utk.edu

Gerald Ragghianti, Asim YarKhan
Innovative Computing Laboratory
University of Tennessee
ragghianti@icl.utk.edu, yarkhan@icl.utk.edu

Mark Gates
Computer Science Department
University of Tennessee
mgates3@icl.utk.edu

Piotr Luszczek
Innovative Computing Laboratory
University of Tennessee
luszczek@icl.utk.edu

Jack J. Dongarra
University of Tennessee, Oak Ridge National Laboratory, USA
dongarra@icl.utk.edu

PP103

Minisymposterium: A Survey of Development Practices in High Performance Computing Applications

In this survey, we seek to distill what is currently scattered knowledge about software practices that play a key role in success of scientific software projects. To do so, we use three methods in order to examine a set of well-managed applications that have been the richest source of engagement in productivity efforts. First, the published works that have resulted from these projects are investigated and provide information on the research areas encompassed and the citation counts that provide insight into the popularity and size of the user base. The second method is through interviews conducted with leading developers and group managers for a subset of projects. These interviews provide insights into the characteristics of the teams, their development tools and methods, and the challenges that they face. The third method is a quantitative analysis of the code base repositories that is capable of tracking the development tasks and communications performed by the principal developers.

Alexander Grannan
Argonne National Laboratory
agrannan@anl.gov

Kanika Sood
University Of Oregon
1991
kanikas@cs.uoregon.edu

Nisansa de Silva, Boyana Norris
University of Oregon
ndilushandesilva@anl.gov, norris@cs.uoregon.edu

Anshu Dubey
Argonne National Laboratory

adubey@anl.gov

PP103

Minisymposterium: A Look at PFLOTRAN's Cloud-based Continuous Integration

Continuous integration can play a valuable role in sustainable software development when properly utilized. Through continuous integration, source code and documentation can be automatically downloaded, compiled, and tested against a gold standard prior to deployment. One potential challenge with continuous integration is the ongoing maintenance of supporting software and hardware infrastructure, which may be routine for IT specialists, but somewhat bothersome for domain scientists. However, frameworks exist within cloud computing that greatly facilitate continuous integration, many of which are free to open source code development projects. The PFLOTRAN project has leveraged continuous integration since 2012. PFLOTRAN is a massively-parallel, subsurface reactive multiphase flow and transport simulator founded upon the PETSc framework. PFLOTRAN has been developed under open source GNU LGPL licensing for over ten years. According to OpenHub, PFLOTRAN is composed of over three hundred thousand lines of source code generated through an estimated 83 years of effort (COCOMO model) by an international group of 36 contributors. The project leverages several cloud resources (i.e. Bitbucket, Codeship, Github and Travis CI) to automatically build, test, and deploy PFLOTRAN source code and documentation. This presentation will detail the projects cloud-based implementation of continuous integration. SAND2018-8861 A.

Glenn Hammond
Sandia National Laboratories
gehammo@sandia.gov

Gautam Bisht
Lawrence Berkeley National Laboratory
gbisht@lbl.gov

PP103

Minisymposterium: LibEnsemble + PETSc/TAO Sustaining a Library for Dynamic Ensemble-based Computations

This poster presents the development and use cases for LibEnsemble, a new catalyst for massively parallel design, decision, and inference problems. We emphasize our strategy for productivity, sustainability, portability, functionality, and breadth of potential use cases while outlining our limitations and identifying collaboration and support needs. We outline our strategy for continuous integration and testing and avenues for efficient automation of incorporating user-contributed enhancements.

Stephen Hudson
Argonne National Laboratory
shudson@anl.gov

David S. Bindel
Cornell University
bindel@cs.cornell.edu

Jeffrey Larson, Barry F. Smith, Stefan Wild
Argonne National Laboratory
jmlarson@anl.gov, bsmith@mcs.anl.gov, wild@mcs.anl.gov

PP103

Minisymposterium: Sustaining Student Software

A university research group focused on CSE software development can generate an enormous amount of software very quickly. Not all of it can (or should) be sustained. A constantly increasing maintenance burden is unsustainable and the transient nature of student researchers further complicates software sustainability. However, computational research groups build their research capabilities by leveraging and sustaining software. This poster will summarize lessons learned and challenges encountered by an early career faculty member striking this balance. We will describe software life cycles in the context of the software development workflow, tool stack, and community strategies we use in the Advanced Reactors and Fuel Cycles group. This life cycle will touch on a number of topics, including: software birth/death/afterlife, the various roles of the PI, collaborative code review, skills training, best practices, responsible employment of undergraduates, contribution disambiguation, and contributor transitions.

Kathryn Huff
University of Illinois at Urbana–Champaign
kdhuff@illinois.edu

PP103

Minisymposterium: Productive and Sustainable Python Workflows in ParSL

Researchers frequently rely on large-scale and domain-specific applications that have been developed over long periods and with significant investments. Increasingly, however, the researchers do not simply run these codes once but instead run them multiple times with varying inputs to understand a range of outputs or to find a set of inputs that gives a particular desired or "best' output. Using a workflow system rather than ad hoc scripts offers better long-term productivity. In addition, similar patterns appear in multiple workflows, so that a script from one application can become a template for others, again improving productivity. And finally, using a workflow system increases sustainability since it can be supported by both its developers and a larger user community. Parsl (Parallel Scripting Library), a Python library for programming and executing data-oriented workflows in parallel, is one such workflow system that helps to increase productivity and sustainability for CSE application developers. They can write a Python script that wraps their applications and contains the higher level logic of their overall workflow, and then annotate the script with Parsl directives. Parsl permits execution of the script on clusters, clouds, grids, and other resources, orchestrates required data movement, and manages execution of Python functions and external applications in parallel. This poster describe Parsls architecture and highlights domains in which it has been used.

Daniel S. Katz
University of Illinois at Urbana-Champaign
d.katz@ieee.org

Yadu Babuji, Kyle Chard, Ben Clifford
University of Chicago
yadunand@uchicago.edu, chard@uchicago.edu,
b.clifford@gmail.com

Ian Foster
Argonne National Laboratory
Data Science and Learning Division
foster@anl.gov

Lukasz Lacinski
University of Chicago
lukasz@uchicago.edu

Connor Pigg
University of Illinois
cpigg2@illinois.edu

Michael Wilde
Parallel Works
wilde@parallelworks.com

Anna Woodard
University of Chicago
annawoodard@uchicago.edu

PP103

Minisymposterium: Improving the Development Workflow of the SETSM Photogrammetry Software

The Surface Extraction by TIN-based Search-space Minimization (SETSM) software is used to produce digital elevation maps (DEMs) from satellite imagery. SETSM has been used for both the ArcticDEM and REMA projects to produce high-resolution terrain maps for the polar regions. The goal of SETSM is to automatically extract a stereo-photogrammetric DEM from pairs of images without any user-defined or a-priori information and using only the sensor Rational Polynomial Coefficients for geometric constraints. The software is written entirely in C and is only dependent on the libtiff and libgeotiff libraries. The SETSM algorithm constructs a Triangular Irregular Network (TIN) in object-space domain to minimize the necessary search space and it employs the coarse-to-fine and vertical line locus strategies. In this poster we present some of the challenges and solutions involved in optimization of a code that is also concurrently under active development by domain scientists. We discuss the challenges of using continuous integration to move toward automated testing. We also present methods to improve the maintainability and usability of this code. One example, code modularization allows for the testing of different algorithms and makes concurrent development by computer scientists and photogrammetry experts more manageable.

Samuel Khuvis
University of Maryland, Baltimore County
Department of Mathematics and Statistics
skhuvis@osc.edu

Judith Gardiner
Ohio Supercomputer Center
judithg@osc.edu

Ian Howat
Ohio State University
School of Earth Sciences
ihowat@gmail.com

Caleb Lehman
Ohio Supercomputer Center
clehman@osc.edu

Myoung-Jong Noh
Ohio State University
School of Earth Sciences
ngnmj77@gmail.com

Karen Tomko
Ohio Supercomputer Center
ktomko@osc.edu

PP103

Minisymposterium: Automated Performance Analysis with PETSc

Comparing the performance of different discretizations or complex solvers is a challenging task, since pure flop rates or rooflines shed little light. Our recently developed TAS system for performance comparison incorporates error measures into the performance analysis, so that we can compare different solvers, discretizations, and models for the same problem. We have now extended this work to multiphysics and time-dependent systems.

Matthew G. Knepley
University at Buffalo
Department of Computer Science and Engineering
knepley@gmail.com

Justin Chang
NREL
justin.chang@nrel.gov

Albert Cowie
University at Buffalo
arcowie@buffalo.edu

PP103

Minisymposterium: Managing Software Development Requirements with user Stories

A user story is a simplified description of a requirement of a product or project that can assist in planning and discussion. Its basic elements are who, what, why, which are used to describe the type of user, what is required, and the reason for the requirement. User stories are typically used in Agile software development, where requirements and solutions evolve and involve a collaborative effort toward the

delivery of the finished product. This poster will cover the basics of user stories, will give examples of stories conceptualized by members of the IDEAS-Productivity component of DOEs Exascale Computing Project, and will evoke likely scenarios in scientific software development.

Osni A. Marques
Lawrence Berkeley National Laboratory
Berkeley, CA
oamarques@lbl.gov

PP103

Minismymposterium: Progress in CSE Software Ecosystems

Software—cross-cutting technology that connects advances in mathematics, computer science, and domain-specific science and engineering—is a cornerstone of long-term collaboration and progress in computational science and engineering (CSE). As we leverage unprecedented high-performance computing resources to work toward predictive science, software complexity is increasing due to multiphysics and multiscale modeling, the coupling of simulations and data analytics, and the demand for greater reproducibility and sustainability, all in the midst of disruptive architectural changes. Applications increasingly require the combined use of independent software packages, whose development teams have diverse sponsors, priorities, software engineering expertise, and processes for development and release. The developers of open-source scientific software are increasingly encouraging community contributions and considering more effective strategies for connections among complementary packages. This poster will discuss work toward broader software interoperability and scientific software ecosystems needed to support next-generation CSE.

Lois Curfman McInnes
Mathematics and Computer Science Division
Argonne National Laboratory
curfman@mcs.anl.gov

PP103

Minisymposterium: Software Engineering for Research Software

The increase in the importance of Research Software (i.e. software developed to support research) motivates the need to identify and understand which software engineering (SE) practices are appropriate. Because of the uniqueness of the research software domain, existing SE tools and techniques developed for the business/IT community are often not efficient or effective. Appropriate SE solutions must account for the salient characteristics of the research software development environment. To identify these solutions, members of the SE community must interact with members of the research software community. This poster will describe efforts to integrate appropriate software engineering practices into the development of research software. The poster will highlight: the findings from a series of case studies of research software projects, an ongoing workshop series, and the results of direct interactions between software engineers and research software projects.

Reed Milewicz
Sandia National Laboratory
rmilewi@sandia.gov

Jeffrey C. Carver
Department of Computer Science
University of Alabama
carver@cs.ua.edu

PP103

Minisymposterium: IDEAS PSIP in Practice: Adopting Continuous Integration for Exascale MD

As leadership-scale computing systems continue to grow in both size and complexity, so do the scientific software applications that are designed to use them. In order to leverage extreme scales efficiently, many modern applications are composed of a collection of independent packages and libraries. For these projects, the necessary implementation of sustainable software practices requires developers to navigate multiple code bases. One promising answer to this challenge is the Productivity and Sustainability Improvement Planning (PSIP) methodology being developed by the Interoperable Design of Extreme-scale Application Software (IDEAS) project. The PSIP methodology promotes the clear factoring of new software processes and capabilities into a manageable number of critical steps with simple completion criteria. The Exascale Atomistic Capability for Accuracy, Length and Time (EXAALT), currently developed under the US-DOE Exascale Computing Project (ECP), is an excellent example of a complex software framework in which PSIP has proven valuable. Specifically, EXAALT is comprised of three sub-projects (ParSplice, LAAMPS, and LATTE) with different physics, each with its own development processes and dependencies. In this work, we highlight a recent effort to implement an end-to-end continuous-integration pipeline within the EXAALT project repository to demonstrate the advantages of PSIP-based software development.

Richard J. Zamora
Argonne National Laboratory
rzamora@anl.gov

Christoph Junghans
Los Alamos National Laboratory
junghans@lanl.gov

David Moulton
Los Alamos National Laboratory
Applied Mathematics and Plasma Physics
moulton@lanl.gov

PP103

Minisymposterium: The Journal of Open Source Software

This poster describes the progress of the Journal of Open Source Software (JOSS), a free, open-access journal designed to publish brief papers about research software. The primary purpose of JOSS is to enable developers of re-

search software to receive citation credit equivalent to typical archival publications. JOSS papers are deliberately short and required to include a summary describing the purpose and high-level functionality of the software (written for a diverse, non-specialist audience), a statement of need, authors and their affiliations, and key references, as well as link to an archived version of the software (e.g., DOI obtained from Zenodo). Upon acceptance, papers receive a CrossRef DOI. Rather than a review of a lengthy software paper (including, e.g., methodology, validation, sample results), JOSS submissions undergo rigorous peer review of the article and software, including documentation, tests, continuous integration, and licensing. The JOSS review process is modeled on the established approach of the rOpenSci collaboration. The entire submission and review process occurs openly on GitHub; papers not yet accepted remain visible and under review until the authors make appropriate changes for acceptance—unlike other journals, papers requiring major revision are not rejected. JOSS was founded in May 2016, and it its first year published 111 articles in a variety of fields. Since then, JOSS has published over 350 articles, and submissions continue to grow.

Kyle E. Niemeyer
Oregon State University
Kyle.Niemeyer@oregonstate.edu

Arfon Smith
Space Telescope Science Institute
arfon@stsci.edu

Lorena Barba
George Washington University
labarba@email.gwu.edu

Jed Brown
University of Colorado Boulder
jed.brown@colorado.edu

Jason Clark
Montana State University
jaclark@montana.edu

George Githinji
KEMRI-Wellcome Trust Research Programme
ggithinji@kemri-wellcome.org

Melissa Gymrek
UC San Diego
mgymrek@ucsd.edu

Lindsey Heagy
University of British Columbia
lheagy@eos.ubc.ca

Kathryn Huff
University of Illinois at Urbana–Champaign
kdhuff@illinois.edu

PP103

Minisymposterium: How to Professionally Develop Reusable Scientific Software and When Not To

A critical challenge in scientific computing is balancing the production of high-quality, domain-specific, reusable software with the need to make immediate scientific progress. While a comprehensive design may provide poor returns on resources expended if many features see minimal use, prioritizing immediate research goals often produces software with limited reuse potential, thereby hindering future progress, impeding reproducibility, and hampering collaborative efforts. We present a flexible approach to scientific software development in which prototype code is made reusable only when required by imminent scientific applications. Our lazy refactoring technique emphasizes the importance of clearly defined interfaces and sharply delimited scopes for maximizing code reusability, and its effective application requires a systematic, thorough assessment of the existing software ecosystem to leverage existing tools wherever possible. We provide guidance on how this technique may be applied, and we develop concrete criteria by which the usability of existing tools may be evaluated within this framework. Finally, we showcase tools for particle simulations developed by the Glotzer Group at the University of Michigan, using their development progression to demonstrate the application of lazy refactoring. Although these tools are independent of one another, we focus on showing how they form a loosely integrated group of software tools that can be integrated into research

Vyas Ramasubramani
University of Michigan
vramasub@umich.edu

Carl S. Adorf
Department of Chemical Engineering
University of Michigan, Ann Arbor
csadorf@umich.edu

Joshua Anderson
University of Michigan
joaander@umich.edu

Sharon Glozter
Department of Chemical Engineering, University of Michigan
4Biointerfaces Institute, University of Michigan
sglotzer@umich.edu

PP103

Minisymposterium: Exahype: An Exascale Hyperbolic PDE Engine

ExaHyPE is a hyperbolic PDE engine capable of solving systems of first order hyperbolic PDEs. The engine provides a space-tree discretization of the computational domain, higher-order ADER DG schemes and a-posteriori subcell limiters. The two main applications currently tackled with this engine are long-range seismic risk assessment and the search for gravitational waves emitted by binary neutron stars. The primary goal of our project is to enable medium-sized interdisciplinary research teams to realise extreme-scale simulations quickly. Users of the engine write only their own application specific code and benefit immediately from the efficient adaptive mesh re-

finement algorithms and from the numerical schemes built into ExaHyPE. On our poster, we will provide several examples illustrating the effort necessary to add new hyperbolic PDE systems into the engine. This is joint work with groups from Frankfurt's FIAS, the University of Trento, Ludwig-Maximilians-University Munich and the University of Durham.

Anne Reinarz, Michael Bader
Technical University of Munich
reinarz@in.tum.de, bader@in.tum.de

PP103

Minisymposterium: The NWChemEX Simulation Development Environment - A General Computational Chemistry Software Framework

The range of scientific problems which can be addressed with computational quantum chemistry (QC) expands through advances in computational throughput, method development, and algorithm optimizations. But as we approach the era of exascale computing, very few QC codes are poised to take advantage of the most powerful supercomputers. Additionally, methods/algorithm implementations are rarely reusable across codes owing to the complexity and monolithic nature of many QC codes. A new QC software package, NWChemEx, is being developed to overcome these challenges. NWChemEx is designed for high performance and scalability on emerging exascale architectures. Software sustainability is addressed via a software framework called the Simulation Development Environment (SDE), by encapsulating scientific algorithms into units called modules which are decoupled from the rest of the program infrastructure. In this work, the capabilities and major design concepts of the SDE are presented. We show how developers can add and extend modules without intimate understanding of the software infrastructure outside of their module. We exhibit how the SDE infrastructure supports flexible data representations, simplifying the development of data interfaces with other software packages. The SDE approach to high-performance computing is briefly discussed. Finally, we provide an example implementation of the self-consistent field method which illustrates how developers interact with the SDE.

Ryan Richard
Ames National Laboratory
Iowa State University
rrichard@ameslab.gov

Kristopher Keipert
Argonne National Laboratory
kkeipert@anl.gov

Thom Dunning
Pacific Northwest National Laboratory
University of Washington
thdjr@uw.edu

Robert Harrison
Brookhaven National Laboratory
Stony Brook University
robert.harrison@stonybrook.edu

Theresa Windus
Iowa State University
twindus@iastate.edu

PP103

Minisymposterium: Changing Mindsets for Large-scale Modelling

Solving large-scale problems is substantially different from solving small-scale problems, and a change in approach is often required. Communicating this idea to end use practitioners can be difficult, yet important. We demonstrate with an example showing the need to move away from direct solvers in order to solve large solid mechanics problems in $O(n)$ time. Whilst this may be uncontroversial, it does result in new areas of focus. Rather than trying to reduce the number of cells in a mesh, there is now be a stronger emphasis on good cell quality. Additionally, iterative solvers are more difficult to drive than a black-box direct solver, so more experience is required to get good results. By demonstrating significant speed-ups on real world problems, we can make a good case for the investment of time and training which is required.

Chris Richardson
University of Cambridge
chris@bpi.cam.ac.uk

Garth Wells
Department of Engineering
University of Cambridge
gnw20@cam.ac.uk

Nathan Sime
University of Cambridge
njcs4@cam.ac.uk

PP103

Minisymposterium: Increasing Software Testing Coverage and Portability with Spack

Adequate testing coverage is critical for successful scientific software development. For numerically intensive codes, the test matrix should cover a variety of machines, compilers, third-party library dependencies, as well as possible other features such as memory testing, parallelization paradigms, and application options that may be enabled or disabled, as some examples. Such coverage can be a major burden on developers without powerful methods for managing the matrix. Spack is a great tool developed by Lawrence Livermore National Lab for managing software on a Unix-based system and in addition to this, we are leveraging Spack to simplify application testing development and management. Here we demonstrate exploiting many of Spack's features to orchestrate the test coverage for the Nalu-Wind application we are developing under the Exascale Computing Project. By using Spack underneath our testing infrastructure, we have increased our testing coverage, portability, robustness, and software quality. While we use our Nalu-Wind application as an example, the benefits of exploiting Spack for application testing can be had for other

applications as well.

Jon Rood, Shreyas Ananthan
National Renewable Energy Laboratory
jon.rood@nrel.gov, shreyas.ananthan@nrel.gov

PP103

Minisymposterium: Modernizing the Scientific Software Approach for the Fusion Analysis Code Transp

TRANSP is a time-dependent 1.5D MHD equilibrium and plasma transport solver for modeling tokamak fusion devices. This software is used by over 100 physicists at several research centers world-wide for interpreting experimental results, predictive analysis, experimental campaign planning, and rapid between shot analysis. The capabilities of the code has grown significantly over 30+ years of development, particularly with respect to the physics models. However, attention to best software practices has lagged the advancements in physics resulting in a non-modular code that is challenging the implementation of new models which require extensive HPC capabilities. This poster will describe the progress and plans for modernizing the development and maintenance cycles for TRANSP. Topics to be discussed will include dependency graph analysis and potential code refactoring, homegrown build system upgrades, approach to version control, containerization for portability and cloud computing, and code documentation.

Jai Sachdev
Princeton Plasma Physics Laboratory
Princeton Plasma Physics Laboratory
jsachdev@pppl.gov

Marina Gorelenkova, Xingqiu Yuan, Joshua Breslau, Francesca Poli
Princeton Plasma Physics Laboratory
mgorelen@pppl.gov, xyuan2@pppl.gov, jbreslau@pppl.gov, fpoli@pppl.gov

PP103

Minisymposterium: No Need for Excuses: Applying Software Engineering Principles to Facilitate Scientific Software Documentation

Software documents, such as requirements specification, design specification, verification reports and user manuals, improve such software qualities as maintainability, reusability, verifiability and usability. So why do scientific software developers underemphasize documentation? Typical reasons include the following: requirements are not known up-front (they emerge over time), change is frequent, rigid processes hamper creativity, the software is too complex and there is no test oracle. Although these are reasons that documentation is challenging, they are often used as excuses for avoiding it entirely. Complexity and frequent change are not unique to scientific software. So how do other domains deal with these challenges - by applying Software Engineering (SE) principles, techniques and tools. Specific SE ideas that should prove helpful for scientific software include: faking a rational design process, separation of concerns, abstraction, generalization, modularization, information hiding, documentation templates, unit testing, and adopting a program family perspective. Many developers are familiar with these ideas, but examples of adapting them to scientific software are rare. An example is shown here to illustrate SE ideas for software to analyze the blast resistance of a glass pane. Example ideas highlighted include abstraction to keep the requirements stable and using a design that hides the likely changes, such as the interpolation algorithm, behind a fixed interface.

Spencer Smith
McMaster University, Canada
smiths@mcmaster.ca

PP103

Minisymposterium: Analyzing Open-source Scientific Software Projects

The growing number and complexity of open-source numerical software projects offer an opportunity to study software development productivity in such projects. Unlike industry, research software development has more limited resources and is typically only a part of the each developers main occupation, making individual developers productivity critical for the success of the project. In this work, we analyze HPC software repositories and compute productivity metrics that can be used to better understand and potentially improve software development practices. We derive metrics that give insight into development practices and general software trends, such as monthly bug fix and feature request rates. We also perform fine-grain analysis for individual code files, patterns in software development, and individual developer efforts. We study the impact of code growth and changes in code complexity over project lifetimes. These repository-based measures are augmented by our analysis of natural language sources tied to the projects, such as mailing lists and issue discussions. We estimate the impact of these factors by analyzing several projects that have been in existence for at least three years. We demonstrate our approach on several HPC projects: SPACK, PETSc, MOOSE, ACME, and yt, among others.

Kanika Sood
University Of Oregon
1991
kanikas@cs.uoregon.edu

Nisansa Silva
University Of Oregon
nisansa@cs.uoregon.edu

Boyana Norris
University of Oregon
norris@cs.uoregon.edu

Anshu Dubey
Argonne National Laboratory
adubey@anl.gov

Lois Curfman McInnes
Mathematics and Computer Science Division
Argonne National Laboratory

curfman@mcs.anl.gov

PP103

Minisymposterium: A Software Productivity and Sustainability Case Study: Multithreaded Requests to Cloud Services for Intelligent Address Standardization

Software productivity and sustainability are two key concerns when developing or modifying scientific software. Here, we present a case study on incorporating multithreading and multiprocessing into the extension of a software package for speeding up large-scale address standardization. For long-term sustainability, we have developed the research software with substantial code reuse of existing software, Google-style Python documentation and doctests. Address standardization serves as an important preprocessing step to geocoding or a post-processing step to reverse geocoding. It is critical in various record linkage schemes of big data sources involving geographical fields. Numerous well-known address standardization software is available as cloud services, for example, usaddress, Data Science Toolkit, and Geocoder.us; their underlying models are based on intelligent machine learning methods such as neural networks trained from large samples of addresses in US or beyond. In this case study, we extend a Python open-source software package of Choi, Lin, and Mulrow (2017) that serially tested accuracies and response time of the aforementioned clouds services on parsing large samples of clean and noisy addresses. Specifically, we design and implement multithreading and multiprocessing wrappers for issuing RESTful APIs in parallel. Our parallelized testing approach achieves an average speed up factor of more than 25 in execution time on machines with multiple cores.

Aleksei Sorokin
Illinois Institute of Technology
asorokin@hawk.iit.edu

Andy Liu
Illinois Mathematics and Science Academy
aliu2@imsa.edu

Sou-Cheng T. Choi
Illinois Institute of Technology
schoi32@iit.edu

PP201

Minisymposterium: Modeling the Chemistry and Hydrodynamics of Micro-swimmers

Janus particles are inorganic micro-swimmers that have been extensively studied, both experimentally and numerically. The swimmer is partially coated with a catalyst that induces chemical reactions on one hemisphere. The resulting gradients in concentration drive a fluid flow in a process called self-diffusiophoresis. Simulating these microswimmers has high computational cost, caused by the complexity of the chemistry and modeling the fluid-structure interaction, so high performance computing is needed. For this problem we use AMReX, an open source software framework developed at Lawrence Berkeley National Laboratory used for solving partial differential equations. As an initial result we study the concentration profiles and flow generated by a single spherical swimmer.

Matea Alvarado
University of California Merced
malvarado27@ucmerced.edu

Johannes P. Blaschke
Max Planck Institute for Dynamics and Self-Organization
jpblaschke@lbl.gov

PP201

Minisymposterium: Modeling Electro-kinetic Flows with Fluctuating Hydrodynamics

We formulate a computational model of low Mach number fluctuating hydrodynamic (FHD) equations for electrolyte solutions. We are interested in studying transport in mixtures of charged species at the mesoscale, down to scales below the Debye length, where thermal fluctuations have a significant impact on the dynamics. We model charged species using a quasielectrostatic approximation, where localized charges create an electric field, which in turn provides additional forcing in the mass and momentum equations. We observe that thermal fluctuations induce large scale spatiotemporal oscillations in the electric potential, which in turn affect important quantities such as the mass flow rate and electric current.

Sean P. Carney
University of Michigan
scarney@math.utexas.edu

John B. Bell
CCSE
Lawrence Berkeley Laboratory
jbbell@lbl.gov

Andy J. Nonaka
Lawrence Berkeley National Laboratory
ajnonaka@lbl.gov

Alejandro Garcia
San Jose State University
alejgarcia987@gmail.com

PP201

Minisymposterium: An Overview of GPU Strategies for Porting Amrex-Based Applications to Next-generation HPC Systems

AMReX is a parallel computing framework for applying adaptive mesh refinement (AMR) to scientific applications. AMReX-based applications, including the astrophysics code Castro and the beam-plasma simulation code WarpX, have begun to implement AMReX's new GPU offloading paradigms to gain access to next generation HPC resources, including ORNL's Summit supercomputer. The AMReX library is exploring multiple paradigms using OpenMP, OpenACC, CUDA Fortran and CUDA to allow users to offload kernels in a manner that yields good speedups while maintaining readability for users. An

overview of the paradigms will be presented and compared on Summit, LBNL's Cori supercomputer and other applicable HPC platforms. Selected AMReX-based applications that have been ported to GPUs will be presented, focusing on paradigms implemented, the difficulty of the conversion, runtime improvement compared to modern CPU-based HPC systems, and where additional optimizations could be made.

Kevin N. Gott
LBNL
kngott@lbl.gov

PP201

Minisymposterium: Performance Study of GPU Offloading via CUDA, OpenACC, and OpenMP in AMReX

AMReX is a parallel computing framework for applying adaptive mesh refinement (AMR) to scientific applications. Many applications built with AMReX can benefit from offloading compute intensive routines to Graphics Processing Units (GPUs). In AMReX, an FArrayBox is the data structure that holds a box, which represents a patch in in the AMR hierarchy, as well as data pointers to values on the grid points within the patch. FArrayBox serves as a computational unit in an operation that loops over a set of grid points, i.e., CPU function calls and GPU kernels are applied to one FArrayBox at a time. Hence, the box size determines the amount of work the GPU has to do during each kernel launch. We investigate the effect of box sizes on GPU computation efficiency by adapting two existing CPU applications to run on Tesla P100 NVIDIA GPUs on Summitdev at Oak Ridge National Laboratory. One application solves the heat equation using data held on meshes, and the other solves the Maxwell equations using both mesh and particle data. Our study also compares three common GPU programming frameworks for offloading kernels: CUDA, OpenACC, and OpenMP, both in terms of performance, and in terms of ease of programming on the current Summitdev development environment.

Yuexia Lin
Harvard University
y_lin@g.harvard.edu

Ann S. Almgren
Lawrence Berkeley National Laboratory
asalmgren@lbl.gov

Brian Friesen, Andrew Myers
Lawrence Berkeley National Lab
bfriesen@lbl.gov, atmyers@lbl.gov

PP201

Minisymposterium: Application of Implicit Methods for Complex Fluid Flow

The overall goal of this project is to develop a code framework for modeling small-scale (sub-micron) flow, where thermal fluctuations play an important role. The applications include electrolyte solutions, electrochemical micropumps, and surface membranes. One such approach is Fluctuating Hydrodynamics (FHD) where the usual deterministic equations are augmented with random fluxes. The underlying numerical approach for FHD requires the use of a Stokes solver to solve the incompressible Navier-Stokes equations. We have developed a new incompressible Navier-Stokes solver using the massively parallel libraries provided by the AMReX framework. The development of this Stokes solver requires developing a multigrid module to discretize momentum diffusion implicitly, and a GMRES solver to simultaneously enforce the evolution of the velocity field subject to the divergence-free constraint. Results from the algorithm are verified using various stochastic and deterministic tests, such as static structure factor characterization for thermal equilibrium fluctuations.

Guy C. Moore, Andrew Nonaka
Lawrence Berkeley National Laboratory
g_moore5@u.pacific.edu, ajnonaka@lbl.gov

PP201

Minisymposterium: Overview of Amrex - a New Framework for Block-structured Adaptive Mesh Refinement Calculations

AMReX is a new software framework that supports the development of block-structured adaptive mesh refinement algorithms for solving systems of partial differential equations on emerging architectures. AMReX aims to provide all the tools necessary for performing complex multiphysics simulations on an adaptive hierarchy of meshes. We give an overview of the software components provided by AMReX, including support for cell, edge, face, and node-centered mesh data, particles, embedded boundary (cut cell) representations of complex geometries, linear solvers, profiling tools, and parallel load balancing. We describe the parallelization strategies supported, including straight MPI, hybrid MPI+OpenMP, and support for GPU systems. Finally, we also give an overview of the application codes built on top of AMReX, which span a wide range of scientific domains and include several ECP and SciDAC-supported projects.

Andrew Myers
Computational Research Division
Lawrence Berkeley National Laboratory
atmyers2@gmail.com

Ann S. Almgren
Lawrence Berkeley National Laboratory
asalmgren@lbl.gov

John B. Bell
CCSE
Lawrence Berkeley Laboratory
jbbell@lbl.gov

Marcus Day
Lawrence Berkeley National Laboratory
msday@lbl.gov

Brian Friesen
Lawrence Berkeley National Lab
bfriesen@lbl.gov

Kevin N. Gott
LBNL
kngott@lbl.gov

Andy J. Nonaka, Steven Reeves
Lawrence Berkeley National Laboratory
ajnonaka@lbl.gov, sireeves@lbl.gov

Weiqun Zhang
Lawrence Berkeley National Laboratory
Center for Computational Sciences and Engineering
weiqunzhang@lbl.gov

PP201

Minisymposterium: Building a Microphysics Cloud Model with AMReX and PADDI

Weather and climate prediction models remain inaccurate, largely due to the uncertainty of warm rain formation (which accounts for 30-70% precipitation budget on various regions of Earth). To understand this process, it is necessary to study how small-scale turbulence in a cloud affects droplet growth through condensation and droplet-droplet collisions. Current models are idealistic – either consisting of a mono/bi-disperse distribution of raindrops or only a subset of physical processes characteristic of a cloud. We propose a new HPC tool PADDI+AMReX to capture the full microphysics in a small volume of a cloud.

Sara Nasab
University of California, Santa Cruz
snasab@ucsc.edu

PP201

Minisymposterium: SedonaEx: A Monte Carlo Radiation Transfer Code for Astrophysical Events

We present SedonaEx, a multi-dimensional, time-dependent Monte Carlo radiation transfer code built on top of the AMReX framework for adaptive mesh refinement. This code is designed to address the three-dimensional, angle-dependent, multi-wavelength radiation transfer problem for astrophysical events such as supernovae. Using SedonaEx, one can calculate the broadband light curves, time-varying spectra, and polarization arising from thermonuclear supernovae explosions and other events where it is important to model the spatial asymmetries of the ejecta in a three-dimensional domain. SedonaEx is based on the Sedona code previously introduced for supernovae [Kasen et al., Time-Dependent Monte Carlo Radiative Transfer Calculations For Three-Dimensional Supernova Spectra, Light Curves, And Polarization, ApJ 651:366, 2006]. We describe the spatial grid layout and communication strategy for SedonaEx using AMReX along with scaling results on current supercomputing systems. We also discuss our treatment of opacities in local thermodynamic equilibrium (LTE) as well as non-LTE opacities in the context of Type Ia supernovae. Finally, we describe exploratory efforts to couple the Monte Carlo radiation transfer algorithm of SedonaEx to the AMReX-based hydrodynamics code Castro.

Donald E. Willcox, Ann S. Almgren
Lawrence Berkeley National Laboratory
dewillcox@lbl.gov, asalmgren@lbl.gov

Daniel Kasen
University of California, Berkeley
Lawrence Berkeley National Laboratory
kasen@berkeley.edu

Andrew Myers
Lawrence Berkeley National Lab
atmyers@lbl.gov

Weiqun Zhang
Lawrence Berkeley National Laboratory
Center for Computational Sciences and Engineering
weiqunzhang@lbl.gov

PP202

Minisymposterium: A Massively Parallel Solver for Poissons Equation on Block Structured Cartesian Grids

We are developing a massively parallel, open-source computational fluid dynamics (CFD) software package (GEM3D) to study microscale atmospheric flows over complex terrain with super-fine numerical resolutions. Our package uses block structured Cartesian adaptive mesh refinement (AMR) methods to provide a practical, accurate method for geometrically modeling complex terrain and large-eddy simulation techniques with dynamic subgrid scale models for physically modeling incompressible flows in complex terrain. In this poster, we report on one of the sub-projects of the main project: the development of a scalable solver for the Pressure Poisson equation in three dimensions on multi-block Cartesian grids. This method uses a fast adaptive composite-type (FAC) multigrid method as a preconditioner for BiCGSTAB. The FAC multigrid method employs a block Jacobi-smoother that uses Fourier transforms to solve Poissons equation on the Cartesian blocks. We compare this method to a Schur complement domain decomposition technique and a standard Krylov iterative method preconditioned with algebraic multigrid (AMG) implemented through the HYPRE library. Scalability results of the method are presented on both CPU and GPU architectures.

Scott Aiton, Donna Calhoun
Boise State University
scottaiton@u.boisestate.edu,
donnacalhoun@boisestate.edu

Grady B. Wright
Department of Mathematics
Boise State University, Boise ID
gradywright@boisestate.edu

PP202

Minisymposterium: Absorbing Outgoing Waves

with Coordinate Mappings in Clawpack

When a partial differential equation on a very large domain is solved numerically on a smaller domain, computational boundary conditions are needed to close the discrete system. For hyperbolic PDEs, these computational boundary conditions must avoid modifying the underlying physics of the problem or generating spurious reflections of waves from the non-physical boundary. There are many existing approaches, including extrapolation, exact boundary conditions, and perfectly matched layers. Many of these approaches are, however, limited to a particular class of hyperbolic PDE: one-dimensional, scalar, linear, constant coefficient, etc. As such, a coordinate mapping approach (also known as grid-stretching or super-grid) is selected for its robustness in problems that involve higher dimensions, systems of PDEs, and/or non-linear equations. Various mappings are investigated that map a bordering layer of the finite computational domain to the original infinite domain. An implementation of the approach is developed for high-resolution Godunov-type, finite-volume schemes that use solutions to Riemann problems at cell boundaries. The efficacy of each mapping with this implementation is presented using Clawpack (www.clawpack.org) for linear, nonlinear, and heterogenous problems. This work was partially performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Randall LeVeque
University of Washington
Applied Math
rjl@uw.edu

Christopher J. Vogl
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
vogl2@llnl.gov

PP202

Minisymposterium: Accelerating Wave-propagation Algorithms with Adaptive Mesh Refinement Using the Graphics Processing Unit (GPU)

Clawpack is a library for solving nonlinear hyperbolic partial differential equations using high-resolution finite volume methods based on Riemann solvers and limiters. It supports Adaptive Mesh Refinement (AMR), which is essential in solving multi-scale problems. Recently, we added capabilities to accelerate the code by using the Graphics Process Unit (GPU). Routines that manage CPU and GPU AMR data and facilitate the execution of GPU kernels are added. Customized and CPU thread-safe memory managers are designed to manage GPU and CPU memory pools, which is essential in eliminating the overhead of memory allocation and de-allocation. A global reduction is conducted every time step for dynamically adjusting the time step based on Courant number restrictions. Some small GPU kernels are merged into bigger kernels, which greatly reduces kernel launching overhead. A speed-up between 2 and 3 for the total running time is observed in an acoustics benchmark problem. The code and benchmark used in this study can be found at `https://github.com/xinshengqin/amrclaw/tree/gpu_amr_paper_benchm`

Xinsheng Qin
University of Washington, Seattle
xsqin@uw.edu

Randall LeVeque
University of Washington
Applied Math
rjl@uw.edu

Michael Motley
University of Washington
mrmotley@uw.edu

PP202

Minisymposterium: Evaluation of Shallow Water Models for Tsunami Prediction with a Near-field Seismic Source

The tsunami warning systems in operation today rely heavily on shallow water models to protect coastal communities. Once detected, the seismic signal is inverted to obtain fault-plane information, which is then used to determine seafloor deformation. Assuming the ocean is in a hydrostatic state, the seafloor deformation profile is used as an initial condition for the shallow water model. The solution to that model determines where initial warnings are issued. Warnings and forecasts are then updated by measuring the height of inbound waves, which are computed using ocean bottom pressure signals. Both the computation of ocean height and simulated tsunami propagation rely on a hydrostatic assumption that has been shown to be valid for far-field sources but is suspect for issuing warnings for near-field events such as a magnitude-9 earthquake in the Cascadia Subduction Zone. To evaluate the use of a hydrostatic assumption for near-field sources, a coupled seismic-acoustic-gravity model is used to capture seismic source, tsunami formation, and tsunami propagation. Two-dimensional results are obtained using Clawpack (www.clawpack.org). on shared-resource systems. Three-dimensional results are obtained from a recently developed code that scales across distributed-resource architectures. This work was partially performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Christopher J. Vogl
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
vogl2@llnl.gov

Randall LeVeque
University of Washington
Applied Math
rjl@uw.edu

PP203

Minisymposterium: Some Recent Advancements on Multifidelity Monte Carlo Techniques for Un-

certainty Quantification

Uncertainty Quantification (UQ) is critical in order to enable predictive numerical simulations for scientific discoveries and advanced engineering design. However, in the presence of complex high-fidelity simulations and a large number of uncertainty parameters the computational cost becomes prohibitive. In recent years, multifidelity (MF) UQ has been introduced in order to alleviate this issue and it is based on the aggregation of several lower accuracy models with an handful of higher fidelity computations. In this poster we focus on sampling based approaches and we will present an overview of our recent activities in this area. At least two different topics will be included in the poster. The first topic is an approximated control variate approach which enables to maximizing the information gain in the presence of multiple low-fidelity models when their expected values are not know. The second topic deals with the enhancement of correlation between models (potentially with dissimilar parametrizations) by leveraging their independent active subspaces structures. All the approaches will be demonstrated on a several test problems and some more realistic engineering applications.

Gianluca Geraci
Sandia National Laboratories, NM
ggeraci@sandia.gov

Alex A. Gorodetsky
University of Michigan
goroda@unimich.edu

Michael S. Eldred, John D. Jakeman
Sandia National Laboratories
mseldre@sandia.gov, jdjakem@sandia.gov

PP203

Minisymposterium: Markov Chain Monte Carlo Methods for Seismic Source Inversion

We focus on implementing sampling strategies for Bayesian inversion, using Markov Chain Monte Carlo (MCMC), to recover the probability distribution of the spatial location of the epicenter of an earthquake in a bounded domain. To this end, we first postulate a computational elasticity model (usually arising from the discretization of a partial differential equation) that describes the earthquake dynamics. In particular, this so-called forward problem is described by the elastodynamic wave equation and discretized using, e.g, the spectral element method. Despite having an efficient solver for the forward problem, efficiently sampling from the posterior distribution arising from the Bayesian paradigm is not a trivial task. In addition, material properties of the ground are typically considered to be uncertain and as such will be treated as random parameters in the source inversion problem. We discuss possible Markov Chain Monte Carlo strategies (and their multi-level extensions) suited for this task.

Juan Pablo Madrigal Cianci
University of New Mexico
juan.madrigalcianci@epfl.ch

Fabio Nobile
EPFL, Switzerland
fabio.nobile@epfl.ch

Raul F. Tempone
Mathematics, Computational Sciences & Engineering
King Abdullah University of Science and Technology
raul.tempone@kaust.edu.sa

PP203

Minisymposterium: Computation of Electromagnetic Fields Scattered From Objects of Uncertain Shapes Using Multilevel Monte Carlo

Computational tools for characterizing scattering from objects of uncertain shapes are highly useful in the fields of electromagnetics, optics, and photonics, where device performance oftentimes is subject to manufacturing tolerances. Often, such computational tools use the Monte Carlo (MC) method to sample a parametric space describing geometric uncertainties. For each sample, which corresponds to a realization of the geometry, a deterministic electromagnetic solver computes the scattered fields. However, for an accurate statistical characterization the number of MC samples has to be large. In this work, to address this challenge, the continuation multilevel Monte Carlo (CMLMC) method is used together with a surface integral equation solver. The CMLMC method optimally balances statistical errors due to sampling of the parametric space, and numerical errors due to the discretization of the geometry using a hierarchy of discretizations, from coarse to fine. The number of realizations of finer discretizations can be kept low, with most samples computed on coarser discretizations to minimize computational work. Consequently, the total execution time is significantly reduced, in comparison to the standard MC scheme.

Alexander Litvinenko
Stochastic Numerics and Bayesian Computational groups, KAUST
alitvinen@gmail.com

Raul F. Tempone
Mathematics, Computational Sciences & Engineering
King Abdullah University of Science and Technology
raul.tempone@kaust.edu.sa

Abdulkadir Yucel
School of Electrical and Electronics Engineering, Nanyang Te
abdulkadir.yucel@kaust.edu.sa

Hakan Bagci
King Abdullah University of Science and Technology
hakan.bagci@kaust.edu.sa

Jesper Oppelstrup
KTH Royal INstitute of Technology
jespero@kth.se

Erik Michielssen
Department of Electrical Engineering and Computer Science
University of Michigan

emichiel@umich.edu

PP203

Minisymposterium: Kullback-Leibler Algorithms for Bayesian Inverse Problems using Multi-Level Monte Carlo Methods

In this work we are attempting a multi-level Monte Carlo approach to efficiently carry out stochastic optimization procedures that arise in Bayesian inversion problems where the posterior distribution is approximated by a distribution chosen within a parametric family via minimization of the Kullback-Leibler divergence.

Panagiotis Tsilifis
University of Southern California
pantsili@gmail.com

Juan Pablo Madrigal Cianci
University of New Mexico
juan.madrigalcianci@epfl.ch

Fabio Nobile
EPFL, Switzerland
fabio.nobile@epfl.ch

PP204

Minisymposterium: Analysis of Equity Markets - A Graph Theory Approach

We develop a characterization of the structure of the US Stock Market by studying how correlations between the various stocks and sectors of the market fluctuate. Through this characterization, we hope to identify the "strongest' of stocks, and sectors and thus identify which investments are safest. This analysis will allow us to provide an alternate investment strategy for those wishing to avoid long term risk in equity markets. This is done using a correlation based graph representing the stock market. The central finding of this study is that transportation sector, a subset of the industrial sector, is the best indicator of the market's fluctuation.

Joshua R. Abrams
University of Arizona
rabrams12@math.arizona.edu

José Celaya-Alcalá
Brown University
Jose_Celaya-Alcala@Brown.edu

PP204

Minisymposterium: A Mathematical Model for Cost-Effectiveness Analysis and Early Detection of Leptospirosis in Humans

Leptospirosis is a zoonotic disease, which is endemic in many tropical regions. A recent estimate for Leptospirosis globally is 1.03 million cases and 58,900 deaths each year. In addition, since mild uncharacteristic signs and symptoms of fever, headache and muscle pain often are confused with other flu-like diseases, most mild Leptospirosis cases are misdiagnosed and not reported. Despite of this global burden of Leptospirosis and the greatest benefit of antibiotic therapy in the early stage of illness, early diagnosis of Leptospirosis still remains a challenge. In this study, we have developed a mathematical tool for analyzing cost effectiveness of various strategies for early detection of Leptospirosis in humans. This tool involves optimization and stochastic simulations.

Ibrahim H. Aslan
University of Tennessee, Knoxville
iaslan@vols.utk.edu

Suzanne Lenhart
National Institute for Mathematical and Biological Synthesis
and Department of Mathematics, University of Tennessee
slenhart@tennessee.edu

PP204

Minisymposterium: Comparison of Modern Langevin Integrators for Simulations of Coarse-Grained Polymer Melts

For a wide range of phenomena, current computational ability does not always allow for fully atomistic simulations of high-dimensional microscopic systems to reach time scales of interest. Coarse-graining is a popular approach to alleviate the impact of computational limits. Consequently, it is of importance to understand how Langevin integrators perform on non-trivial coarse-grained molecular systems, and in particular how large of an integration time step can be used without introducing unacceptable amounts of error into averaged quantities of interest. To investigate this, we examined three different Langevin integrators on a coarse-grained polymer melt: the recently developed Leimkuhler and Matthews method, called BAOAB, the Grnbech-Jensen and Farago method, denoted G-JF, and the frequently used Brnger-Brooks-Karplus integrator, also known as BBK. We compute and analyze key statistical properties for each. Our results indicate that the three integrators perform similarly when using a small friction parameter; however, outside of this regime the use of large integration steps produces significant deviations from the predicted diffusivity and steady-state distributions for all integration methods examined with the exception of G-JF.

Joshua Finkelstein, Giacomo Fiorin, Benjamin Seibold
Temple University
joshua.finkelstein@temple.edu, giacomo.fiorin@temple.edu, seibold@temple.edu

PP204

Minisymposterium: Weak Galerkin Method for Electrical Impedance Tomography Inverse Problem

In this work, we propose a weak Galerkin method for solving the electrical impedance tomography inverse problem based on bounded variation regularization and analyze its convergence. We use the complete electrode model as the forward model and its solution is approximated using the lowest order polynomial bases for weak Galerkin method.

The error estimates are studied for the forward problem. We establish the convergence of this weak Galerkin algorithm for the inverse problem, in the sense that the sequence of discrete solutions contains a convergent subsequence to a solution of the continuous bounded variation regularization formulation. We also present numerical results to verify the advantages and efficiency of the algorithm.

Ying Liang
The Chinese University of Hong Kong
yliang@math.cuhk.edu.hk

Jun Zou
The Chinese Univ. of Hong Kong
zou@math.cuhk.edu.hk

PP204

Minisymposterium: Conservative Higher-Order Method for the Solute Transport Equation in Unsaturated Porous Media

Dynamics of solutes in unsaturated porous media is not completely understood. Accurate numerical experiments are a key tool to find relations and parameterizations. The mathematical modeling of solute transport in unsaturated media is, however, quite challenging: the use of finite volume (FV) methods turns to be low order and not natural in unstructured grids while finite element methods need to be treated carefully in order to have local conservation of mass and to avoid spurious oscillations. In this poster, I show the application of the discontinuous Galerkin (DG) method to solve (i) the Richards equation, (ii) the transport equation for unsaturated porous media, and (iii) the use of a flux reconstruction technique with Raviart Thomas elements in order to obtain mass conservative and higher-order solute states. Comparisons show the advantage of using a DG scheme over a FV scheme by the rate of convergence of the effective hydrodynamic dispersion tensor in analytic and synthetic cases.

Santiago Ospina De Los Ríos
Heidelberg University
santiago.ospina@iup.uni-heidelberg.de

Peter Bastian
Interdisciplinary Center for Scientific Computing
University of Heidelberg
peter.bastian@iwr.uni-heidelberg.de

Kurt Roth, Lukas Riedel
Heidelberg University
kurt.roth@iup.uni-heidelberg.de, lriedel@iup.uni-heidelberg.de

PP204

Minisymposterium: Regularization of Rate-Independent Evolution of SMA Described by Non-Convex Energies

Regularization of Rate-Independent Evolution of SMA Described by Non-Convex Energies Abstract: The poster presents mathematical aspects of evolutionary material models for shape-memory alloys (SMAs) at finite-strains. The difficulty of related mathematical analysis consists in the non-linear and non-convex dependence of the energy on the deformation gradient. One-possible way, how maintain the analysis tractable, is to suppose that the energy depends also on the second deformation gradient and is convex in it. We relax this assumption by using the recently proposed concept of gradient-polyconvexity. Namely, we consider energies which are convex only in gradients of non-linear minors of the deformation gradient (i.e. cofactor and determinant in three dimensions). As a consequence, the whole second deformation gradient needs not to be integrable. Yet, at the same time, the obtained compactness is sufficient and, moreover, additional physically desirable properties (e.g. global invertibility) can be shown. We extend previous results for hyperelastic materials by incorporating a rate-independent dissipation to our model and by proving existence of an energetic solution to it. It is a joint work with Martin Kruik (Institute of Information Theory and Automation of the Czech Academy of Sciences) and Anja Schloemerkemper (University of Wuerzburg).

Petr Pelech
Charles University
pelech@karlin.mff.cuni.cz

Martin Kruzik
Academy of Sciences
Czech Technical University Prague
kruzik@utia.cas.cz

Anja Schloemerkemper
Institute for Mathematics
University of Wuerzburg
anja.schloemerkemper@mathematik.uni-wuerzburg.de

PP204

Minisymposterium: Computational Developments for the Bayesian Conjugate Gradient Method

Abstract not available

Timothy Reid
North Carolina State University
twreid@ncsu.edu

PP204

Minisymposterium: MORLAB - Model Order Reduction LABoratory

The modeling of real-world applications, like electrical circuits, computational fluid dynamics and mechanical systems results in linear dynamical systems. Often those systems are described by a large number of differential and algebraic equations, which make the evaluation in optimization and controller design difficult. The aim of model reduction is the construction of a surrogate model for the original one, that is much easier to evaluate. The MORLAB, Model Order Reduction LABoratory, toolbox is an efficient software solution in MATLAB and Octave for the model reduction of linear time-invariant continuous-time systems. It implements a wide range of techniques for medium-scale systems in standard, descriptor and second-

order form by using spectral projection methods.

Steffen W. R. Werner
Max Planck Institute for Dynamics of Complex Technical Syste
werner@mpi-magdeburg.mpg.de

Peter Benner
Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany
benner@mpi-magdeburg.mpg.de

PP204

Minisymposterium: Reduced Order Methods for Parametrized Optimal Flow Control Problems: Applications in Biomedical and Environmental Marine Sciences

In this work, we will present applications of reduced order methods for parametrized problems in computational fluid dynamics, with a special attention to inverse problems, such as optimal flow control problems and data assimilation. Our focus will be on applications in life sciences, specifically in biomedical and environmental sciences.
Among the former, we will show these methods applied to patient-specific coronary artery bypass grafts (CABGs), extracted from medical images provided by clinics/hospitals. The optimal control paradigm allows pursuing clinically provided physiological data through computational fluid dynamics, hence, bringing patient-specific hemodynamics modeling closer to real-life clinical cases.
Among the latter, we will discuss applications arising in environmental marine sciences and engineering, namely a pollutant control in the Gulf of Trieste, Italy and a solution tracking governed by quasigeostrophic equations describing North Atlantic Ocean dynamic. Also in this field, reduced optimal control framework is a useful approach to monitor, manage and predict (possibly dangerous) marine phenomena.
This work is in collaboration with L. Jiménez-Juan (Sunnybrook Health Sciences Centre, Toronto, Canada), P. Triverio (University of Toronto, Canada) and R. Mosetti (National Institute of Oceanography and Applied Geophysics, Trieste, Italy).

Zakia Zainib
SISSA International School for Advanced Studies
zzainib@sissa.it

Maria Strazzullo
mathLab, Mathematics Area, SISSA International School for Ad
mstrazzu@sissa.it

Francesco Ballarin
SISSA International School for Advanced Studies
fballarin@sissa.it

Gianluigi Rozza
SISSA University, Italy
Trieste, Italy
grozza@sissa.it

Organizer and Speaker Index

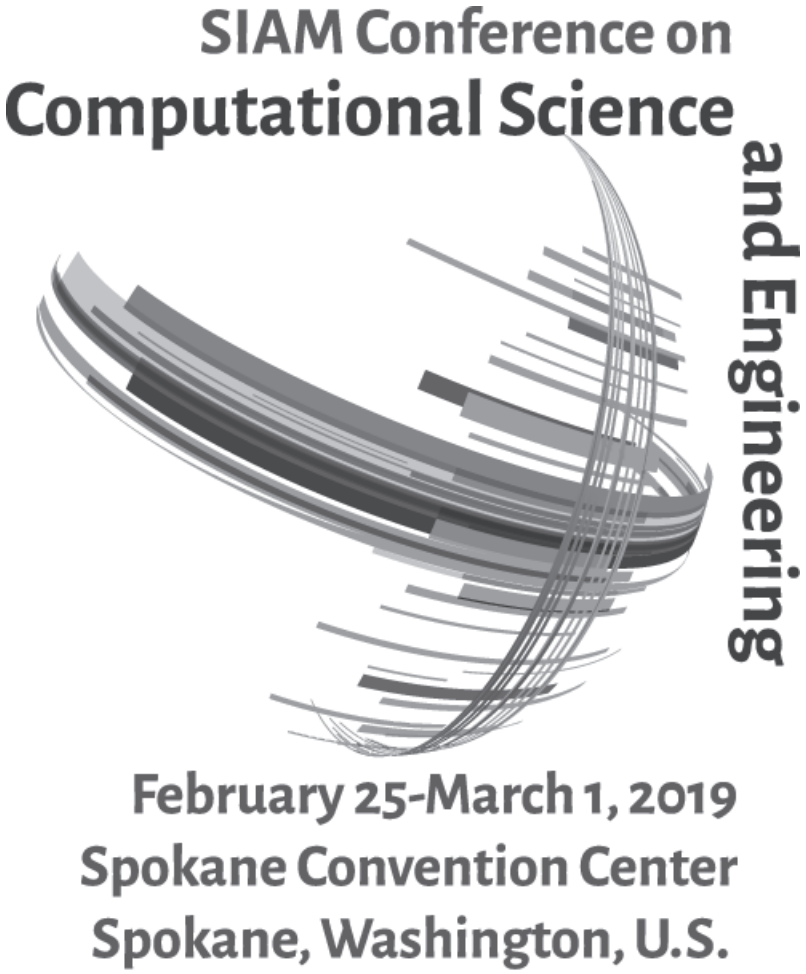

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

R

S

T

U

V

W

X

Y

Z

CSE19 Budget

Conference Budget
SIAM Conference on Computational Science & Engineering
February 25 - March 1, 2019
Spokane, WA

Expected Paid Attendance: 1620

Revenue		
Registration Income		$617,835
	Total	$617,835
Expenses		
Printing		$7,900
Organizing Committee		6,500
Invited Speakers		25,300
Food and Beverage		178,200
AV Equipment/ Room Rental and Telecommunication		200,500
Advertising		9,000
Conference Labor (including benefits)		215,097
Professional Services (Mobile App / Child Care)		6,200
Other (supplies, staff travel, freight, misc.)		24,900
Administrative		80,872
Accounting/Distribution & Shipping		45,941
Information Systems		79,298
Customer Service		27,866
Marketing		52,890
Office Space (Building)		31,907
Other SIAM Services		29,999
	Total	$1,022,370
Net Conference Expense		($404,535)
Support Provided by SIAM		$404,535
		$0
Estimated Support for Travel Awards not included above:		
Early Career and Students	131	$105,200

Spokane Convention Center

CAMPUS MAP

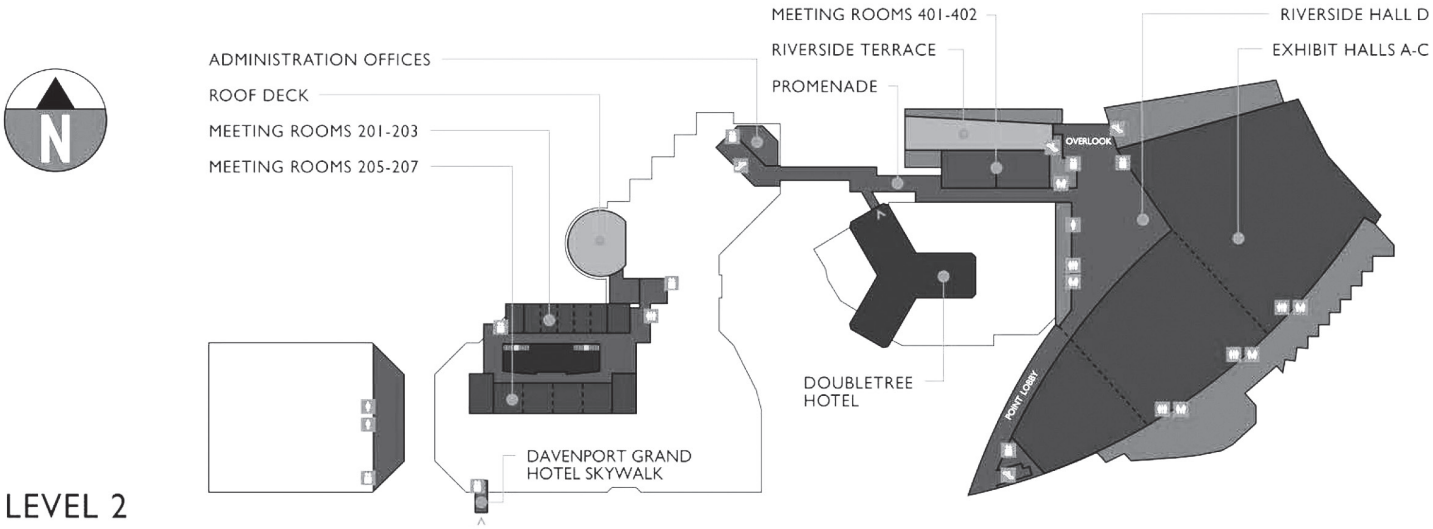

LEVEL 2

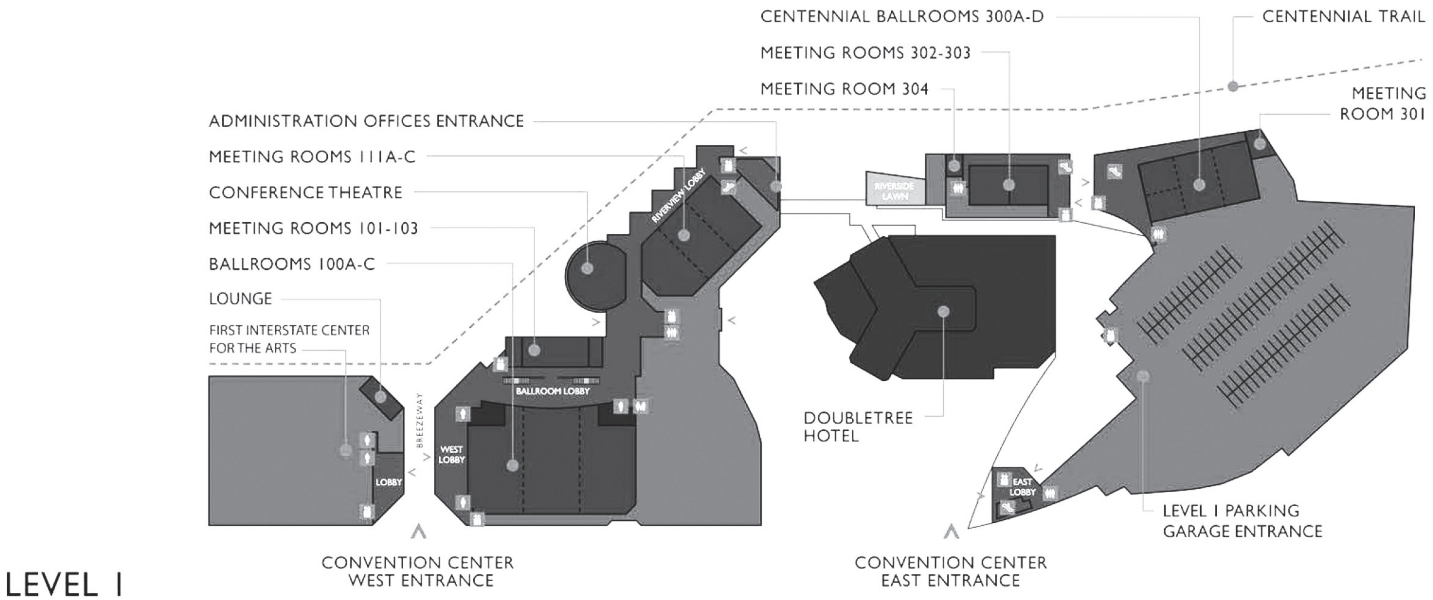

LEVEL 1