Final Program and Abstracts

SIAM Conference on
Parallel Processing for Scientific Computing

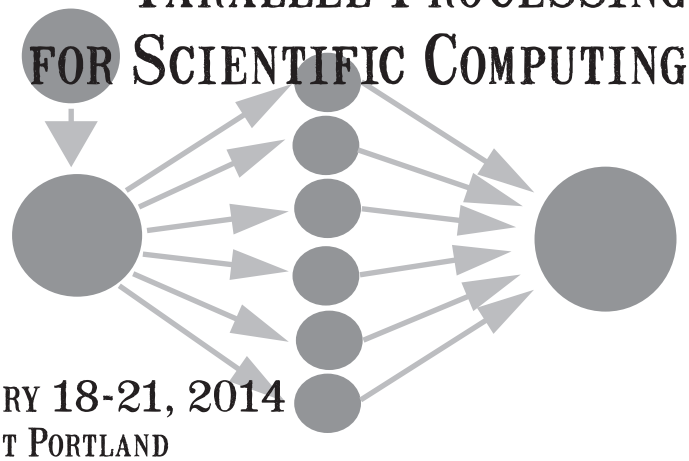

February 18-21, 2014
Marriott Portland
Downtown Waterfront
Portland, Oregon, USA

Sponsored by the SIAM Activity Group on Supercomputing

The SIAM Activity Group on Supercomputing provides a forum for computational mathematicians, computer scientists, computer architects, and computational scientists to exchange ideas on mathematical algorithms and computer architecture needed for high-performance computer systems. The activity group promotes the exchange of ideas by focusing on the interplay of analytical methods, numerical analysis, and efficient computation. The activity group organizes this SIAM Conference on Parallel Processing, awards the SIAG/Supercomputing Career Prize and the SIAG/Supercomputing Junior Scientist Prize, and maintains a member directory and an electronic mailing list.

Society for Industrial and Applied Mathematics
3600 Market Street, 6th Floor
Philadelphia, PA 19104-2688 USA
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 (US & Canada) or +1-215-382-9800 (worldwide)

www.siam.org/meetings/pp14

Table of Contents

Organizing Committee Co-Chairs

Edmond Chow
Georgia Institute of Technology, USA

Ali Pinar
Sandia National Laboratories, USA

Rob Schreiber
HP Labs, USA

Organizing Committee

Pavan Balaji
Argonne National Laboratory, USA

George Biros
The University of Texas, Austin, USA

Pradeep Dubey
Intel Corporation, USA

John Gilbert
University of California, Santa Barbara, USA

Laura Grigori
INRIA Paris, France

William Gropp
University of Illinois, Urbana-Champaign, USA

Torsten Hoefler
ETH Zürich, Switzerland

Jeff Hollingsworth
University of Maryland, College Park, USA

David Kahaner
Asian Technology Information Program, USA

Madhav Marathe
Virginia Tech, USA

Lenny Oliker
Lawrence Berkeley National Laboratory, USA

Stephen Wright
University of Wisconsin-Madison, USA

SIAM Registration Desk

The SIAM registration desk is located in the Ballroom Foyer. It is open during the following hours:

Tuesday, February 18
10:00 AM - 7:00 PM
Wednesday, February 19
7:45 AM - 5:15 PM
Thursday, February 20
7:45 AM - 5:15 PM
Friday, February 21
7:45 AM - 5:15 PM

Hotel Address

Portland Marriott Downtown Waterfront
1401 SW Naito Parkway
Portland, Oregon 97201
Phone Number: +1-503-226-7600
Fax: +1-503-221-1789
Toll Free (USA & Canada): +1-888-236-2427

Hotel Telephone Number

To reach an attendee or to leave a message, call 1-503-226-7600. The hotel operator can either connect you with the SIAM registration desk or to the attendee's room. Messages taken at the SIAM registration desk will be posted to the message board located in the registration area.

Hotel Check-in and Check-out Times

Check-in time is 4:00 PM and check-out time is 12:00 PM.

Child Care

Please contact the hotel concierge for local child care options, +1-503-226-7600.

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

Corporate/Institutional Members

The Aerospace Corporation

Air Force Office of Scientific Research

AT&T Laboratories - Research

Bechtel Marine Propulsion Laboratory

The Boeing Company

CEA/DAM

Department of National Defence (DND/CSEC)

DSTO- Defence Science and Technology Organisation

Hewlett-Packard

IBM Corporation

IDA Center for Communications Research, La Jolla

IDA Center for Communications Research, Princeton

Institute for Computational and Experimental Research in Mathematics (ICERM)

Institute for Defense Analyses, Center for Computing Sciences

Lawrence Berkeley National Laboratory

Lockheed Martin

Los Alamos National Laboratory

Mathematical Sciences Research Institute

Max-Planck-Institute for Dynamics of Complex Technical Systems

Mentor Graphics

National Institute of Standards and Technology (NIST)

National Security Agency (DIRNSA)

Oak Ridge National Laboratory, managed by UT-Battelle for the Department of Energy

Sandia National Laboratories

Schlumberger-Doll Research

Tech X Corporation

U.S. Army Corps of Engineers, Engineer Research and Development Center

United States Department of Energy

List current December 2013.

Funding Agency

SIAM and the Conference Organizing Committee wish to extend their thanks and appreciation to the Department of Energy (DOE) and the U.S. National Science Foundation for their support of this conference.

Leading the applied mathematics community . . .

Join SIAM and save!

SIAM members save up to $130 on full registration for the 2014 SIAM Conference on Parallel Processing for Scientific Computing! 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 *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 or Non-Member

Mini Speaker/Organizer rate to attend the conference, you can apply the difference between what you paid and what a member would have paid ($130 for a Non-Member and $65 for a Non-Member Mini Speaker/Organizer) 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 $10 to join the SIAM Activity Group on Supercomputing (SIAG/SC). As a SIAG/SC member, you are eligible for an additional $10 discount on this conference, so if you paid the SIAM member rate to attend the conference, you might be eligible for a free SIAG/SC membership. Check at the registration desk.

Free Student Memberships are available to students who attend an institution that is an Academic Member of SIAM, are members of Student Chapters of SIAM, or are nominated by a nonstudent member of SIAM.

Join onsite at the registration desk, go to *www.siam.org/joinsiam* 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

E-mail: *membership@siam.org*

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

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.

The Plenary Session Room will have two (2) screens, one (1) data projector and one (1) overhead projector. Cables or adaptors for Apple computers are not supplied, as they vary for each model. Please bring your own cable/adaptor if using an Apple computer.

All other concurrent/breakout rooms will have one (1) screen and one (1) data projector. Cables or adaptors for Apple computers are not supplied, as they vary for each model. Please bring your own cable/adaptor if using an Apple computer. Overhead projectors will be provided only if requested.

If you have questions regarding availability of equipment in the meeting room of your presentation, or to request an overhead projector for your session, please see a SIAM staff member at the registration desk.

E-mail Access

Attendees booked within the SIAM room block will have complimentary wireless Internet access in their guest rooms. All conference attendees will have complimentary wireless Internet access in the meeting space. In addition, a limited number computers with Internet access will be available for attendee use during registration hours.

Registration Fee Includes

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

Job Postings

Please check with the SIAM registration desk regarding the availability of job postings or visit *http://jobs.siam.org.*

Important Notice to Poster Presenters

The poster session is scheduled for Tuesday, February 18, 2014, 6:00 PM – 8:00 PM. Poster presenters are requested to set up their poster material on the provided poster boards in Salon E between the hours of 1:00 PM and 6:00 PM. All materials must be posted by 6:00 PM on Tuesday, February 18, 2014, official start time of the session. Posters will remain on display through Friday, February 21, 2014. Poster displays must be removed by 10:30 AM. Posters remaining after this time will be discarded. SIAM is not responsible for discarded posters.

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. If a SIAM books representative is not available, completed order forms and payment (credit cards are preferred) may be taken to the SIAM registration desk. The books table will close at 10:30 AM on Friday, February 21, 2014.

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:

Sven Leyffer, SIAM Vice President for Programs (*vpp@siam.org*)

Get-togethers

- Welcome Reception and Poster Session
  Tuesday, February 18,
  6:00 – 8:00 PM

- Business Meeting
  (open to SIAG/SC members)
  Thursday, February 29,
  6:30 – 7:15 PM

Complimentary beer and wine will be served.

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, pagers, 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, in order 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 #SIAMPP14.

SIAM Activity Group on Supercomputing (SIAG/SC)

www.siam.org/activity/supercomputing

A GREAT WAY TO GET INVOLVED!

Collaborate and interact with mathematicians and applied scientists whose work involves the applications of supercomputing.

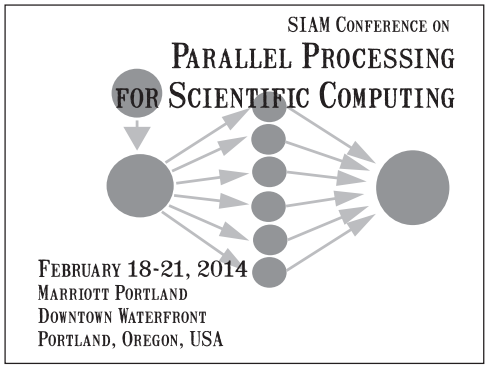

ACTIVITIES INCLUDE:

- Special sessions at SIAM Annual Meetings
- Biennial conference
- SIAG/Supercomputing Career Prize
- SIAG/Supercomputing Junior Scientist Prize
- Website

BENEFITS OF SIAG/DS MEMBERSHIP:

- Listing in the SIAG's online membership directory
- Additional $10 discount on registration at the SIAM Conference on Parallel Processing and Scientific Computing *(excludes students)*
- Electronic communications about recent developments in your specialty
- Eligibility for candidacy for SIAG/SC office
- Participation in the selection of SIAG/SC officers

ELIGIBILITY:

- Be a current SIAM member.

COST:

- $10 per year
- Student members can join 2 activity groups for free!

2014-15 SIAG/SC OFFICERS

Chair: Ali Pinar, Sandia National Laboratories
Vice Chair: Karen Devine, Sandia National Laboratories
Secretary: Bora Ucar, CNRS
Program Director: Laura Grigori, INRIA

TO JOIN:

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

Invited Plenary Speakers

** All Invited Plenary Presentations will take place in Salon F**

Tuesday, February 18

5:15 PM - 6:00 PM

IP1 The Fusion of Supercomputing with Large Scale Data Analytics

William Blake, *Cray, Inc., USA*

Wednesday, February 19

8:15 AM - 9:00 AM

IP2 On the Performance of Adaptive Mesh-Based Simulations on Modern HPC Architectures

Michael Bader, *Technische Universität München, Germany*

1:45 PM - 2:30 PM

IP3 How Much (Execution) Time, Energy, and Power Will my Algorithm Need?

Richard Vuduc, *Georgia Institute of Technology, USA*

Thursday, February 20

8:15 AM - 9:00 AM

IP4 Particles, Grids and HPC for Simulations in Fluid Dynamics

Petros Koumoutsakos, *ETH Zürich, Switzerland*

1:45 PM - 2:30 PM

IP5 Large-scale GPU Applications for Scientific Computing

Takayuki Aoki, *Tokyo Institute of Technology, Japan*

Invited Plenary Speakers

Friday, February 21

8:15 AM - 9:00 AM

IP6 Co-Design of Exascale Simulation of Combustion in Turbulence (ExaCT)

Jackie Chen, *Sandia National Laboratories, USA*

1:45 PM - 2:30 PM

IP7 Large-Scale Visual Data Analysis

Christopher Johnson, *University of Utah, USA*

Program Schedule

SIAM Conference on

Parallel Processing for Scientific Computing

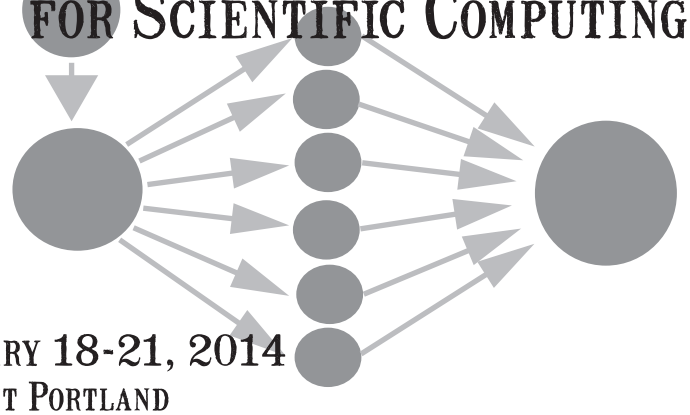

February 18-21, 2014
Marriott Portland
Downtown Waterfront
Portland, Oregon, USA

Tuesday, February 18

Registration

10:00 AM-7:00 PM

Room:Oregon Ballroom Foyer

Tuesday, February 18

MS1

Toward Resilient Applications for Extreme-Scale Systems Part I of IV

1:30 PM-3:10 PM

Room:Salon F

As leadership-class computing systems increase in complexity and transistor feature sizes decrease, application codes find themselves less and less able to treat a system as a reliable digital machine. In fact, the high performance computing community has grown increasingly concerned that applications will have to manage resilience issues beyond the current practice of global checkpoint restart. This is expensive at scale and not capable of fixing all types of errors. We discuss alternatives in software and numerical algorithms that can improve the resiliency of applications and manage a variety of faults anticipated in future extreme-scale computing systems.

Organizer: Keita Teranishi
Sandia National Laboratories, USA

Organizer: Mark Hoemmen
Sandia National Laboratories, USA

Organizer: Jaideep Ray
Sandia National Laboratories, USA

Organizer: Michael A. Heroux
Sandia National Laboratories, USA

1:30-1:50 Toward a Local Failure-Local Recovery Resiliency Model
Keita Teranishi and Michael A. Heroux, Sandia National Laboratories, USA

1:55-2:15 Reliability and Resiliency Trends: An Update
Larry Kaplan, Cray, Inc., USA

2:20-2:40 Fault Tolerance at Exascale: Is Checkpointing Dead?
Kathryn Mohror, Lawrence Livermore National Laboratory, USA

2:45-3:05 Models for Fault-Tolerance at Very Large Scale
Yves Robert, ENS, France

Tuesday, February 18

MS2

Hierarchical and Iteration Space Tiling - Part I of II

1:30 PM-3:10 PM

Room:Salon A

For Part 2 see MS10

The minisymposium discusses the revitalized importance of tiling optimization for effective use of memory hierarchy in the exascale era. The performance gap between the memory and processor is increasing and it will be unprecedentedly widened in the exascale systems. Consequently, cache optimizations such tiling will be not only relevant but indispensable to reduce adverse effects of this performance gap. Iteration space tiling, tuning of tile size, and representing tiles in a hierarchical organization of data are some of active research areas. The symposium will highlight the topics on programming model, compiler, auto-tuning, runtime and hardware supports for tiling optimization.

Organizer: Didem Unat
Lawrence Berkeley National Laboratory, USA

Organizer: John Shalf
Lawrence Berkeley National Laboratory, USA

1:30-1:50 Tiling Dense and Sparse Computations for Parallelism and the Memory Hierarchy of Future Architectures
Anand Venkat, Manu Shantharam, and *Mary Hall*, University of Utah, USA; Michelle Strout, Colorado State University, USA

1:55-2:15 Tiling for Parallel Execution of Stencil Computations
P. Sadayappan, Ohio State University, USA

2:20-2:40 Two Approaches for Scheduling and Tiling to Very Deep Hierarchies, Implemented and Available in R-Stream
Richard Lethin, Benoit Meister, Nicolas Vasilache, Muthu Baskaran, and David Wohlford, Reservoir Labs, USA

2:45-3:05 Performance Modeling for Tiling with PIN
John Bachan, Didem Unat, Cy Chan, and John Shalf, Lawrence Berkeley National Laboratory, USA

Tuesday, February 18

MS3

Parallel Kinetic Monte Carlo and Discrete-Event Simulation

1:30 PM-3:10 PM

Room:Salon B

Kinetic Monte Carlo methods are used to solve stochastic formulations in physics, biology, and in industrial applications and can be categorized as Discrete-Event simulations. Due to the sequential way of processing chronological events, parallelization of Discrete-Event simulations is known to be a non-trivial task. In this minisymposium we discuss algorithms, implementations, and applications where efficient solution to kinetic Monte Carlo problems on shared-memory and distributed systems are targeted.

Organizer: Stefan Engblom
Uppsala University, Sweden

Organizer: Pavol Bauer
Uppsala University, Sweden

1:30-1:50 Parallel Discrete Event Simulation of Infectious Disease Spread in Animal Populations
Pavol Bauer and Stefan Engblom, Uppsala University, Sweden

1:55-2:15 Parallelization of Kinetic Monte Carlo Using Synchronous Algorithms: Applications and Differences in Continuum and Discrete Systems
Jaime Marian, Lawrence Livermore National Laboratory, USA

2:20-2:40 Parallelization, Error Analysis and Partial Asynchrony in Lattice Kinetic Monte Carlo
Markos A. Katsoulakis, University of Massachusetts, Amherst, USA; Georgios Arampatzis, University of Crete, Greece; Petr Plechac, University of Delaware, USA

2:45-3:05 Performance of Time Warp on 1,966,080 Cores
Christopher Carothers, Rensselaer Polytechnic Institute, USA

Tuesday, February 18

MS4

Task Mapping: Parallel Resource Management after Scheduling, Allocation, and Partitioning

1:30 PM-3:10 PM

Room:Salon C

This minisymposium focuses on task mapping to improve computational performance, an important issue that will become crucial as system size continues to increase. Specifically we look at the assignment of tasks to allocated processors for an application. Speakers will present experimental results on algorithms for non-contiguous allocations, algorithms using geometric partitioners, a comparison of strategies on two generation of Cray systems, and algorithms for predicting performance using supervised learning.

Organizer: Vitus Leung
Sandia National Laboratories, USA

Organizer: Karen D. Devine
Sandia National Laboratories, USA

1:30-1:50 Task Mapping Stencil Computations for Non-Contiguous Allocations
Vitus Leung, Sandia National Laboratories, USA; David Bunde, Knox College, USA

1:55-2:15 Topology-Aware Task Mapping Using Geometric Partitioning
Mehmet Deveci, The Ohio State University, USA; Siva Rajamanickam, Karen D. Devine, and Vitus Leung, Sandia National Laboratories, USA

2:20-2:40 A Comparison of Task Mapping Strategies on Two Generations of Cray Systems
Kevin Pedretti, Sandia National Laboratories, USA; Torsten Hoefler, University of Illinois, USA

2:45-3:05 On Predicting Performance on Different Task Mappings Using Supervised Learning
Abhinav Bhatele, Lawrence Livermore National Laboratory, USA; Nikhil Jain, University of Illinois, USA; Todd Gamblin, Lawrence Livermore National Laboratory, USA; Laxmikant Kale, University of Illinois at Urbana-Champaign, USA

Tuesday, February 18

MS5

Integrating Communication Hiding and Avoiding in Exascale Preconditioned Krylov Solver

1:30 PM-3:10 PM

Room:Salon D

We report the progress of the Exascale Algorithms and Advanced Computational Techniques (Exa2CT) project funded by the European Union. In this project communication avoiding and hiding techniques are integrated in a preconditioned Krylov solver for exascale applications. We use the pipelining to hide the latencies of the global reductions, stencil compilers to increase the arithmetic intensity and communication avoiding at the level of the preconditioner. We report initial benchmarks of the components and the integrated prototypes on various proto applications.

Organizer: Wim I. Vanroose
University of Antwerp, Belgium

1:30-1:50 Latency Hiding of Global Reductions in Pipelined Krylov Methods
Wim I. Vanroose, University of Antwerp, Belgium

1:55-2:15 Acceleration of Domain Decomposition Based Algorithms by Communication Avoiding and Hiding Krylov Method
Tomas Kozubek, VSB-Technical University Ostrava, Czech Republic

2:20-2:40 Increasing the Arithmetic Intensity of Multigrid with Stencil Compilers
Pieter Ghysels, Lawrence Berkeley National Laboratory, USA

2:45-3:05 Recovery-Restart Techniques for Resilient Krylov Methods
Emmanuel Agullo, INRIA, France; Luc Giraud, INRIA Bordeaux Sud-Ouest, France; Salas Pablo, CERFACS, France; Zounon Mawussi, INRIA, France

Tuesday, February 18

MS6

Application and Impact of FASTMath (Frameworks, Algorithms and Scalable Technologies for Mathematics on Next-generation Computers) - Part I of II

1:30 PM-3:10 PM

Room:Salon G

For Part 2 see MS14

This minisymposium will describe the use and impact of FASTMath SciDAC software and technologies on application codes. The DOE FASTMath SciDAC project is focused on development of mathematical methods and software for DOE applications. Technologies included in this minisymposium include structured and unstructured mesh techniques, linear and nonlinear solvers, and time integration. Speakers will discuss the science application challenges, then how their technologies have impacted the simulation code technologies.

Organizer: Carol S. Woodward
Lawrence Livermore National Laboratory, USA

1:30-1:50 Nonlinear Solvers and Time Integrators for Dislocation Dynamics
Carol S. Woodward, Athanasio Arsenlis, and Sylvie Aubry, Lawrence Livermore National Laboratory, USA; David Gardner, Southern Methodist University, USA; Gregg Hommes and Kathryn Mohror, Lawrence Livermore National Laboratory, USA; Daniel R. Reynolds, Southern Methodist University, USA

1:55-2:15 Rapid Development of a New Ice Sheet Application Code Using Albany and Trilinos
Andrew Salinger, Irina Kalashnikova, and Mauro Perego, Sandia National Laboratories, USA; Stephen Price, Los Alamos National Laboratory, USA; Ray S. Tuminaro, Sandia National Laboratories, USA; Patrick H. Worley, Oak Ridge National Laboratory, USA

continued in next column

2:20-2:40 Application of Algebraic Multigrid (petsc) for Adaptive Mesh Refinement Applications (Chombo)
Mark Adams, Lawrence Berkeley National Laboratory, USA

2:45-3:05 Scalable Preconditioners for Atmospheric Climate Simulation
Aaron Lott, Lawrence Livermore National Laboratory, USA; Richard Archibald and Katherine J. Evans, Oak Ridge National Laboratory, USA; Carol S. Woodward, Lawrence Livermore National Laboratory, USA; Patrick H. Worley, Oak Ridge National Laboratory, USA

Tuesday, February 18

MS7

Matrix Computations

1:30 PM-3:10 PM

Room:Salon H

This minisymposium collects some recent developments in parallel methods for computing matrix functions, computing stochastic matrix-vector products for uncertainty quantification, solving dense complex systems with thousands of right-hand sides, and matrix bidiagonalization.

Organizer: James V. Lambers
University of Southern Mississippi, USA

1:30-1:50 A Parallel Approach to the Solution of Pde Through Componentwise Approximation of Matrix Functions
James V. Lambers, University of Southern Mississippi, USA

1:55-2:15 High Performance Solution of Dense Linear Systems with Application to Large 3D Electromagnetics Problems on a Petascale Computer and on a Cluster of Gpus
David Goudin, Cedric Augonnet, Agnes Pujols, and Muriel Sesques, CEA/CESTA, France

2:20-2:40 Exploring Emerging Manycore Architectures for Uncertainty Quantification Through Embedded Stochastic Galerkin Methods
Eric Phipps, H. Carter Edwards, Jonathan J. Hu, and Jakob Ostien, Sandia National Laboratories, USA

2:45-3:05 On Parallelization of MKL LAPACK SVD
Sergey V Kuznetsov, Intel Corporation, Russia

Tuesday, February 18

MS8

Creating Meaningful Network Models from Data

1:30 PM-3:10 PM

Room:Salon I

Network science has emerged as an important tool for studying complex systems of relational data. While much research focuses on developing algorithms for analyzing these networks, an equally important component consists of transforming the raw data into network models that accurately represent the underlying system. In this minisymposium, speakers from different domain areas will describe end-to-end case studies. Each will look at choices made as they went from raw data to network model to analysis. By bringing model creation to the forefront we hope to generate discussion in the network science community about the interplay between network modelling and algorithmic analysis.

Organizer: Tzu-Yi Chen
Pomona College, USA

Organizer: Sanjukta Bhowmick
University of Nebraska, Omaha, USA

1:30-1:50 Analyzing Shakespeare's Dramas Using Networks
Sanjukta Bhowmick, University of Nebraska, Omaha, USA

1:55-2:15 Stoichiometry of Steroidogenesis: Towards Understanding Optimal Design and Function
David Hala and Duane Huggett, University of North Texas, USA

2:20-2:40 Using Networks to Model Student Conceptual Understanding of Physics
Jesper Bruun, University of Copenhagen, Denmark

2:45-3:05 On Creating Networks from MRI Data
Tzu-Yi Chen, Pomona College, USA

Coffee Break

3:10 PM-3:40 PM

Room:Salon E

Tuesday, February 18

CP1

GPU Computing I

3:40 PM-5:00 PM

Room:Salon F

Chair: Jeffrey D. Blanchard, Grinnell College, USA

3:40-3:55 Parallel Rotor Wake Simulation on Multicore Architectures with Gpgpu Accelerators Using Openacc
Achim Basermann and Melven Zoellner, German Aerospace Center (DLR), Simulation and Software Technology (SISTEC), Germany; Johannes Hofmann, German Aerospace Center (DLR), Germany

4:00-4:15 Selecting Multiple Order Statistics with a Graphics Processing Unit
Jeffrey D. Blanchard, Erik Opavsky, and Emircan Uysaler, Grinnell College, USA

4:20-4:35 Cyclic Reduction Type Poisson and Helmholtz Solvers on a GPU
Mirko E. Myllykoski, Tuomo Rossi, and Jari Toivanen, University of Jyvaskyla, Finland

4:40-4:55 Accelerating Earthquake Simulations on General-Purpose Graphics Processors
Monish D. Tandale and Jason Kwan, Optimal Synthesis Inc., USA

Tuesday, February 18

CP2

Multigrid I

3:40 PM-5:00 PM

Room:Salon A

Chair: Paul Lin, Sandia National Laboratories, USA

3:40-3:55 Parallel Implementation of An Aggregation-Based Algebraic Multigrid Method
Meng-Huo Chen, University of Washington, USA

4:00-4:15 Efficient Parallel Adaptive Multi-Grid-Like Solver for CFD Applications
Jérôme Frisch and Ralf-Peter Mundani, Technische Universität München, Germany

4:20-4:35 Parallel Algebraic Multigrid and Linear Solvers Suitable for Implicit Particle in Cell Simulations
Pawan Kumar, Katholieke Universiteit Leuven, Belgium

4:40-4:55 Performance of Multigrid Smoothers for Large-Scale Finite Element Simulations
Paul Lin, Eric C. Cyr, and Sivasankaran Rajamanickam, Sandia National Laboratories, USA

Tuesday, February 18

CP3

Eigensolvers

3:40 PM-5:00 PM

Room:Salon B

Chair: Michael M. Wolf, Massachusetts Institute of Technology, USA

3:40-3:55 A Symmetry-Based Decomposition Approach to Eigenvalue Problems with Applications to Electronic Structure Calculations
Xingyu Gao, Beijing Institute of Applied Physics and Computational Mathematics, China; Jun Fang, Institute of Applied Physics and Computational Mathematics, China; Aihui Zhou, Chinese Academy of Sciences, China

4:00-4:15 CoarrayARPACK: Parallel ARPACK in Fortran 2008
Richard J. Hanson, Principal Consultant, USA; Richard B. Lehoucq, Sandia National Laboratories, USA

4:20-4:35 Parallel Strategy for Computing Eigenvalues of Non-Hermitian Matrices
James Kestyn and Eric Polizzi, University of Massachusetts, Amherst, USA; Peter Tang, Intel Corporation, USA

4:40-4:55 Effective Parallel Computation of Eigenpairs to Detect Anomalies in Very Large Graphs
Michael M. Wolf and Benjamin Miller, Massachusetts Institute of Technology, USA

Tuesday, February 18

CP4

Parallel Programming

3:40 PM-5:00 PM

Room:Salon C

Chair: Sunita Chandrasekaran, University of Houston, USA

3:40-3:55 ∇-Nabla: A Numerical-Analysis Specific Language for Exascale Scientific Applications
Jean-Sylvain Camier, CEA, France

4:00-4:15 Simplifying Heterogeneous Multicore Programming Using Industry Standards
Sunita Chandrasekaran and Barbara Chapman, University of Houston, USA

4:20-4:35 Refactorization of a Hybrid, Highly Efficient Parallel PDE Solver
Malin Källén, Uppsala University, Sweden

4:40-4:55 PGAS Programming Model for Manycore Computers
Yili Zheng, Lawrence Berkeley National Laboratory, USA

Tuesday, February 18

CP5

Algorithms I

3:40 PM-5:00 PM

Room:Salon D

Chair: David Starinshak, Lawrence Livermore National Laboratory, USA

3:40-3:55 Parallel Adaptive Cartesian Level-Set Methods
Mohammad Mirzadeh, University of California, Santa Barbara, USA

4:00-4:15 Polytope: A New Parallel Framework for Computing Voronoi Meshes on Complex Boundaries
David Starinshak and John Owen, Lawrence Livermore National Laboratory, USA; Jeffrey N. Johnson, Lawrence Berkeley National Laboratory, USA

4:20-4:35 Randomized Heuristics for the Bipartite Matching Problem on Shared Memory Parallel Computers
Fanny Dufosse, LAAS-CNRS, Toulouse, France; Kamer Kaya, The Ohio State University, USA; *Bora Ucar*, LIP-ENS Lyon, France

4:40-4:55 Partitioning and Parallel Computation of Electricity Production Cost Models
Clayton Barrows, Marissa Hummon, and Wesley Jones, National Renewable Energy Laboratory, USA

Tuesday, February 18

CP6

Applications

3:40 PM-5:00 PM

Room:Salon G

Chair: William F. Mitchell, National Institute of Standards and Technology, USA

3:40-3:55 Interactive Multi-Scale High-Performance Computing
Ralf-Peter Mundani and Jérôme Frisch, Technische Universität München, Germany

4:00-4:15 Recent Advances in PHAML
William F. Mitchell, National Institute of Standards and Technology, USA

4:20-4:35 A Parallel Scattered Node Finite Difference Scheme for the Shallow Water Equations on a Sphere
Martin Tillenius and Elisabeth Larsson, Uppsala University, Sweden; Erik Lehto and Natasha Flyer, National Center for Atmospheric Research, USA

4:40-4:55 Multi Component Polymer Flooding Two Dimensional Oil Reservoir Simulation
Sudarshan Kumar Kenettinkara, TIFR Centre, Bangalore, India

Tuesday, February 18

CP7

Planet Earth Applications

3:40 PM-5:00 PM

Room:Salon H

Chair: Helen Kershaw, National Center for Atmospheric Research, USA

3:40-3:55 Parallelization Challenges for Ensemble Data Assimilation
Helen Kershaw, Nancy Collins, and Jeffrey Anderson, National Center for Atmospheric Research, USA

4:00-4:15 Geophysical EULAG Model with Three-Directional MPI Parallelization
Zbigniew P. Piotrowski, Institute of Meteorology and Water Management, Poland; Piotr Smolarkiewicz, European Centre for Medium-Range Weather Forecasts, United Kingdom; Andrzej Wyszogrodzki, National Center for Atmospheric Research, USA

4:20-4:35 Massively Parallel CPU and GPU Implementation of Elliptic Solvers in Geophysical Modelling
Eike H. Mueller and Robert Scheichl, University of Bath, United Kingdom

4:40-4:55 Large-Scale Parallel Simulation of Urban Flooding Scenarios
Vasco Varduhn and Ralf-Peter Mundani, Technische Universität München, Germany

Tuesday, February 18

CP8

Performance Optimization for Stencils and Meshes

3:40 PM-5:00 PM

Room:Salon I

Chair: Charles R. Ferenbaugh, Los Alamos National Laboratory, USA

3:40-3:55 Modeling Stencil Code Optimizations
Mauricio Araya-Polo, Shell International Exploration and Production, USA; Raul de la Cruz, Barcelona Supercomputing Center, Spain

4:00-4:15 Unstructured Mesh Physics Performance on Current and Future Architectures
Charles R. Ferenbaugh, Los Alamos National Laboratory, USA

4:20-4:35 A Communication Algorithm for the Patch-Based Multiblock Structured Mesh Applications
Hong Guo, Zeyao Mo, and Aiqing Zhang, Institute of Applied Physics and Computational Mathematics, China

4:40-4:55 Numa-Aware Runtime Optimization of Structured Grid Numerical Simulations
Zhang Yang, Chinese Academy of Sciences, China; Aiqing Zhang and Zeyao Mo, Institute of Applied Physics and Computational Mathematics, China

Welcome Remarks

5:10 PM-5:15 PM

Room:Salon E

Tuesday, February 18

IP1

The Fusion of Supercomputing with Large Scale Data Analytics

5:15 PM-6:00 PM

Room:Salon F

Chair: John R. Gilbert, University of California, Santa Barbara, USA

Is highly scalable computing facing a branch in the road with one path leading to exascale supercomputers delivering billion-way parallel computing and another path leading to millions of servers and billions of cores in the cloud delivering results with advanced distributed computing? This paper will explore the technology and architectural trends facing system and application developers and speculate on whether the future will be an "either/or" or a "both/and" scenario.

William Blake
Cray, Inc., USA

Tuesday, February 18

Welcome Reception & Poster Session

6:00 PM-8:00 PM

Room:Salon E

Sparsity in Discontinuous Galerkin on Future Architectures
Zechariah J. Jibben, Arizona State University, USA

Large Scale Multi-Physics Simulation Using the Interoperable Exectuive Library
Kwai L. Wong, University of Tennessee and Oak Ridge National Laboratory, USA; Andrew Kail, University of Tennessee, USA; David White, Maryville College, USA

Scalable Parallel Simulated Annealing
Zhihao Lou, University of Chicago, USA; Rick Stevens, Argonne National Laboratory and University of Chicago, USA; John Reinitz, University of Chicago, USA

GPU Accelerated Shake and Rattle Algorithms for Systems with Holonomic Constraints
Georgiy Krylov and Nursultan Zarlyk, Nazarbayev University, Kazakhstan

Scaling of Entropy Based Moment Closures Vs Pn Moment Closures
Charles K. Garrett and Cory Hauck, Oak Ridge National Laboratory, USA

Investigation on Parallel Performance of Fluid Solvers Based on Different Numerical Methods
Xiaohe Zhufu, Xiaotong Dong, Weishan Deng, and Jin Xu, Chinese Academy of Sciences, China

Parallel Imperfection Study for Accelerating Cavities
Lixin Ge, Liling Xiao, Cho-Kuen Ng, and Zenghai Li, Stanford Linear Accelerator Center, USA

Adaptive Sampling for Large Scale Networks
Shehla Ahmed, University of Nebraska, Omaha, USA

Fast Sparse Direct Solvers for Systems from Elliptic Partial Differential Equations
Joshua D. Booth and Padma Raghavan, Pennsylvania State University, USA

continued in next column

Performance Study of Parallel Octree-based Conforming Tetrahedral Mesh Generation
Jose J. Camata, Federal University of Rio de Janerio, Brazil; Igor Ghisi, ESSS, Brazil; Alvaro Coutinho, COPPE/Universidade Federal do Rio e Janeiro, Brazil

What Can the Roofline Model of Energy Tell Us About How to Build the Next Supercomputer?
Jee W. Choi and Richard Vuduc, Georgia Institute of Technology, USA

HPCC and Undergraduate Applied Mathematics
Randall E. Cone, Virginia Military Institute, USA

A Simple, Parallel Scheme for Support Graph Preconditioning of Networks
Kevin Deweese, University of California, Santa Barbara, USA; Erik G. Boman, Sandia National Laboratories, USA

Evaluating the Impact of Silent Data Corruption in Numerical Methods
James Elliott, North Carolina State University, USA

Treatment of a Lane-Emden Type Equation Via Second Derivative Backward Differentiation Formula Using Boundary Value Technique
Ehigie Julius and Okunuga Solomon, University of Lagos, Nigeria

The Parallel Java 2 Library: Multicore and Cluster Parallel Programming in 100% Java
Alan Kaminsky, Rochester Institute of Technology, USA

Design and Implementation of a Distributed Spatial Data Structure
Sandeep Koranne, Mentor Graphics Corporation, USA

Paralution - a Library for Iterative Sparse Methods on Multi-Core CPU and GPU
Dimitar Lukarski and *Elisabeth Larsson*, Uppsala University, Sweden

A New Scalable Parallel Algorithm for Fock Matrix Construction
Xing Liu, Aftab Patel, and Edmond Chow, Georgia Institute of Technology, USA

Parallel Scaling of Camellia for {DPG} Using BlueGene/Q
Nathan Roberts and Ramesh Balakrishnan, Argonne National Laboratory, USA; Jesse Chan, Rice University, USA; Truman E. Ellis, University of Texas at Austin, USA

continued on next page

Feature Extraction for Hyperspectral Data Using Massive Parallel Processors
Stefan Robila, Montclair State University, USA

Applying Software Agents to Resource Management in Tsunami Modeling
Alexander P. Vazhenin, Kensaku Hayashi, and Yutaka Watanobe, University of Aizu, Japan; Michal Drozdowicz, Maria Ganzha, and Marcin Paprzycki, Polish Academy of Sciences, Poland

Performance Evaluation of the Most Program Package on Nvidia® Kepler GPUs
Alexander P. Vazhenin, Kazuya Matsumoto, Naohiro Nakasato, Stanislav Sedukhin, and Kensaku Hayashi, University of Aizu, Japan

Parallel Implementation of a Contour-Integral-Based Eigensolver for Dense Generalized Eigenvalue Problems on Gpu Clusters
Takahiro Yano, Yasunori Futamura, and Tetsuya Sakurai, University of Tsukuba, Japan

A Highly-Scalable Multi-Bsp Fast Fourier Transform
Albert-Jan N. Yzelman, K.U. Leuven, Belgium

Wednesday, February 19

Registration
7:45 AM-5:15 PM

Room:Oregon Ballroom Foyer

IP2
On the Performance of Adaptive Mesh-Based Simulations on Modern HPC Architectures
8:15 AM-9:00 AM

Room:Salon F

Chair: Laura Grigori, INRIA, France

For large mesh-based simulations, adaptive refinement is essential to limit the computational work, but also comes at an implementation and performance overhead. Depending on whether unstructured, block- or tree-structured approaches are followed, and whether dynamic adaptivity is desired, various challenges exist for multi-level parallelisation (incl. vectorisation) and memory-aware performance optimisation. I will present two respective case studies stemming from earthquake and tsunami simulation: For SeisSol, an ADER-DG code to simulate dynamic rupture and seismic wave propagation on unstructured tetrahedral meshes, I will report on a joint project to optimise SeisSol for the SuperMUC platform. For tsunami simulation, I will discuss parallel adaptive mesh refinement and respective performance optimisations based on space-filling curves.

Michael Bader
Technische Universität München, Germany

Intermission
9:00 AM-9:10 AM

Wednesday, February 19

CP9
GPU Computing II
9:10 AM-10:10 AM

Room:Salon F

Chair: Jhihming Huang, National Taiwan University, Taiwan

9:10-9:25 A Mixed Precision Eigenvalue Solver on GPU
Jhihming Huang, Tsung-Ming Huang, Pochuan Wang, and Weichung Wang, National Taiwan University, Taiwan

9:30-9:45 A Scalable and Fast Multi-Gpu Eigensolver for 3D Photonic Crystal Band Structure
Weichung Wang, Jiahong Chen, Jhihming Huang, and Tsung-Ming Huang, National Taiwan University, Taiwan

9:50-10:05 X-Ray Laser Imaging of Biomolecules Using Multiple Gpus
Jing Liu and Stefan Engblom, Uppsala University, Sweden

Wednesday, February 19

CP10

Multigrid II

9:10 AM-10:10 AM

Room:Salon A

Chair: Andrey Prokopenko, Sandia National Laboratories, USA

9:10-9:25 Reducing Coarse Grids Contention in a Parallel Algebraic Multigrid
Andrey Prokopenko, Sandia National Laboratories, USA

9:30-9:45 Parallel AMG Solver Based on Adaptive Setup Strategy with Applications in Large-Scale Radiation Hydrodynamics Simulations
Xiaowen Xu and Zeyao Mo, Institute of Applied Physics and Computational Mathematics, China

9:50-10:05 Compiler Based Generation and Autotuning of Communication-Avoiding Operators for Geometric Multigrid
Protonu Basu, University of Utah, USA

Wednesday, February 19

CP11

Domain Decomposition

9:10 AM-10:10 AM

Room:Salon B

Chair: Veronika S. Vasylkivska, Oregon State University, USA

9:10-9:25 Deflation Based Domain Decomposition Preconditioners
Pierre Jolivet and Frederic Nataf, Laboratoire Jacques-Louis Lions, France; Christophe Prud'homme, Institut de Recherche Mathématique Avancée de Strasbourg, France

9:30-9:45 A Domain Decomposition Method for Unsteady Flow Routing in Complex River Systems
Veronika S. Vasylkivska, Nathan L. Gibson, Arturo Leon, and Luis Gomez, Oregon State University, USA

9:50-10:05 Data Structures and Algorithms for High-Dimensional Structured Adaptive Mesh Refinement (SAMR)
Magnus Grandin, Uppsala University, Sweden

Wednesday, February 19

CP12

Computational Physics

9:10 AM-10:10 AM

Room:Salon C

Chair: Ying Wai, Oak Ridge National Laboratory, USA

9:10-9:25 A Multi-Physics Approach for Time-Parallel Plasma Physics Applications
Mathias Winkel and Paul Gibbon, Jülich Supercomputing Centre, Germany; Michael Minion, Stanford University, USA; Matthew Emmett, Lawrence Berkeley National Laboratory, USA

9:30-9:45 Replica-Exchange Wang-Landau Sampling - a Highly Scalable Framework for Stochastic Optimization
Thomas Wuest, Swiss Federal Research Institute WSL, Switzerland; Ying Wai Li, Oak Ridge National Laboratory, USA; Thomas Vogel, Los Alamos National Laboratory, USA; David Landau, University of Georgia, USA

9:50-10:05 Scalability and Performance Analysis for Replica-Exchange Wang-Landau Sampling
Ying Wai Li, Oak Ridge National Laboratory, USA; Thomas Vogel, Los Alamos National Laboratory, USA; Thomas Wuest, Swiss Federal Research Institute WSL, Switzerland; Markus Eisenbach, Oak Ridge National Laboratory, USA; David Landau, University of Georgia, USA

Wednesday, February 19

CP13

Algorithms II

9:10 AM-10:10 AM

Room:Salon D

Chair: Robert Robey, Los Alamos National Laboratory, USA

9:10-9:25 Hashing in the Scientific World

Rebecka Tumblin and *Robert Robey*, Los Alamos National Laboratory, USA; Peter Ahrens, University of California, Berkeley, USA; Sara Hartse, Brown University, USA

9:30-9:45 Co-Design of Extremely Scalable Algorithms/Architecture for 3-Dimensional Linear Transforms

Stanislav Sedukhin, University of Aizu, Japan

9:50-10:05 Parallel Locking Sweeping for Static Hamilton-Jacobi Problems of Many Dimensions

Miles L. Detrixhe, University of California, Santa Barbara, USA

Wednesday, February 19

CP14

Fast Fourier Transforms

9:10 AM-10:10 AM

Room:Salon G

Chair: Andrew M. Canning, Lawrence Berkeley National Laboratory, USA

9:10-9:25 A Hybrid MPI/openmp 3D FFT Implementation for Parallel Plane Wave First-Principles Materials Science Codes

Andrew M. Canning, Lawrence Berkeley National Laboratory, USA

9:30-9:45 Exploiting Data Reuse for Reduction of Communication Volume in Parallelization of Multi-Dimensional FFTs

Truong Vinh Truong Duy, University of Tokyo, Japan; Taisuke Ozaki, Japan Advance Institute of Science and Technology, Japan

9:50-10:05 Multiple GPU FFT Algorithms

Kevin Wadleigh, Alex Fit-Florea, Slawomir Kierat, Lukasz Ligowski, and Paul Sidenblad, NVIDIA, USA

Wednesday, February 19

CP15

Checkpointing and Scheduling

9:10 AM-10:10 AM

Room:Salon H

Chair: Zachary B. Spears, Naval Research Laboratory, USA

9:10-9:25 A Comparison of Checkpoint Techniques for HPC Applications

Yves Robert, ENS, France

9:30-9:45 Checkpointing Schemes for Adjoint-Based Optimization of Jet Engine Noise

Zachary B. Spears and Andrew Corrigan, Naval Research Laboratory, USA

9:50-10:05 Resource-Aware Scheduling in Task Parallel Frameworks

Elisabeth Larsson and Martin Tillenius, Uppsala University, Sweden; Rosa M. Badia and Xavier Martorell, Barcelona Supercomputing Center, Spain

Wednesday, February 19

CP16

Performance Optimization

9:10 AM-10:10 AM

Room:Salon I

Chair: Robert L. Kelsey, Los Alamos National Laboratory, USA

9:10-9:25 Timing Performance Surrogates in Auto-Tuning for Qualitative and Quantitative Factors
Jiahong Chen, National Taiwan University, Taiwan; Ray-Bing Chen, National Cheng Kung University, Taiwan; Akihiro Fujii, Kogakuin University, Japan; Reiji Suda, University of Tokyo, Japan; Weichung Wang, National Taiwan University, Taiwan

9:30-9:45 Towards Parallel Library Generation for Polynomial Arithmetic
Lingchuan Meng, Drexel University, USA

9:50-10:05 Performance Analysis of the Pagosa Application
Robert L. Kelsey, Los Alamos National Laboratory, USA

Coffee Break

10:10 AM-10:35 AM

Room:Salon E

Wednesday, February 19

MS9

Toward Resilient Applications for Extreme-Scale Systems Part II of IV

10:35 AM-12:15 PM

Room:Salon F

As leadership-class computing systems increase in complexity and transistor feature sizes decrease, application codes find themselves less and less able to treat a system as a reliable digital machine. In fact, the high performance computing community has grown increasingly concerned that applications will have to manage resilience issues beyond the current practice of global checkpoint restart. This is expensive at scale and not capable of fixing all types of errors. We discuss alternatives in software and numerical algorithms that can improve the resiliency of applications and manage a variety of faults anticipated in future extreme-scale computing systems.

Organizer: Keita Teranishi
Sandia National Laboratories, USA

Organizer: Mark Hoemmen
Sandia National Laboratories, USA

Organizer: Jaideep Ray
Sandia National Laboratories, USA

Organizer: Michael A. Heroux
Sandia National Laboratories, USA

10:35-10:55 Fault-tolerant Iterative Linear Solvers with Multilevel Fault Detection
Mark Hoemmen, Sandia National Laboratories, USA; James Elliott, North Carolina State University, USA

11:00-11:20 Self-Stabilizing Iterative Solvers
Piyush Sao and Richard Vuduc, Georgia Institute of Technology, USA

11:25-11:45 A Fault Tolerant Implementation of Multi-Level Monte Carlo Methods
Stefan Pauli, Manuel Kohler, and Peter Arbenz, ETH Zürich, Switzerland

11:50-12:10 Towards an Unified ABFT Approach for Resilient Dense Linear Algebra
Piotr Luszczek, University of Tennessee, Knoxville, USA

Wednesday, February 19

MS10

Hierarchical and Iteration Space Tiling - Part II of II

10:35 AM-12:15 PM

Room:Salon A

For Part 1 see MS2

The minisymposium discusses the revitalized importance of tiling optimization for effective use of memory hierarchy in the exascale era. The performance gap between the memory and processor is increasing and it will be unprecedentedly widened in the exascale systems. Consequently, cache optimizations such tiling will be not only relevant but indispensable to reduce adverse effects of this performance gap. Iteration space tiling, tuning of tile size, and representing tiles in a hierarchical organization of data are some of active research areas. The symposium will highlight the topics on programming model, compiler, auto-tuning, runtime and hardware supports for tiling optimization.

Organizer: Didem Unat
Lawrence Berkeley National Laboratory, USA

Organizer: John Shalf
Lawrence Berkeley National Laboratory, USA

10:35-10:55 Programming with Tiles
David Padua, University of Illinois at Urbana-Champaign, USA

11:00-11:20 Chapel Language Features for Hierarchical Tiling and Exascale Architectures
Bradford L. Chamberlain, Cray, Inc., USA

11:25-11:45 Hardware Support for Collective Data Transfers
Georgios Michelogiannakis, Lawrence Berkeley National Laboratory, USA

11:50-12:10 Programming Model Support for Tiling
Didem Unat, Lawrence Berkeley National Laboratory, USA

Wednesday, February 19

MS11

Large-scale Multiphysics Simulation for Nuclear Reactor Analysis and Design

10:35 AM-12:15 PM

Room:Salon B

The next generation of simulation tools for nuclear reactor design are focused not on a single physics, but on multiple coupled disciplines that can span a wide range of time and length scales. This session is focused on advances in multiphysics integration applied to large-scale simulation for nuclear reactor design. Areas of interest include solution algorithms, discretization algorithms, scalability, conservation enforcement, data transfer strategies and the design and APIs of supporting software frameworks/ libraries.

Organizer: Roger Pawlowski
Sandia National Laboratories, USA

Organizer: Eric C. Cyr
Sandia National Laboratories, USA

Organizer: Eric Phipps
Sandia National Laboratories, USA

10:35-10:55 Physics Based Coupling for Multiscale Full Core Nuclear Reactor Simulation
Derek R. Gaston, Cody Permann, John Peterson, Andrew Slaughter, David Andrs, and Yaqi Wang, Idaho National Laboratory, USA; Michael Short, Massachusetts Institute of Technology, USA; Michael Tonks and Richard Martineau, Idaho National Laboratory, USA

11:00-11:20 Tools Supporting the Assembly of Multiphysics Simulation Codes from Standalone Parallel Physics Codes
Timothy J. Tautges, Vijay Mahadevan, and Rajeev Jain, Argonne National Laboratory, USA

continued in next column

11:25-11:45 Code Integration Strategies for Large-Scale Reactor Simulation
Roger Pawlowski, Sandia National Laboratories, USA; Roscoe Bartlett, Oak Ridge National Laboratory, USA; Kenneth Belcourt, Sandia National Laboratories, USA; Kevin Clarno, Oak Ridge National Laboratory, USA; Derek R. Gaston, Idaho National Laboratory, USA; Russell W. Hooper, Sandia National Laboratories, USA; Robert Salko, Oak Ridge National Laboratory, USA

11:50-12:10 Advanced Coupling Explorations for Parallel Coupled Neutronic and Thermal-Hydraulic Simulation
Kevin Clarno, Mark Berrill, and Steven Hamilton, Oak Ridge National Laboratory, USA; Roger Pawlowski, Sandia National Laboratories, USA; John Turner, Oak Ridge National Laboratory, USA

Wednesday, February 19

MS12

Density Functional Theory: Large-scale Algorithms and O(N) Methods - Part I of II

10:35 AM-12:15 PM

Room:Salon C

For Part 2 see MS20
Density Functional Theory (DFT) is a fundamental predictive model to study matter at atomistic scale. It is widely used to help understand and design advanced materials. DFT simulations, however, are limited in size due to their high computational cost and the cubic complexity of typical solvers. In addition, traditional DFT implementations typically involve global operations which prevent efficient use of a large number of processors. We will discuss recent developments to address these issues and enable scaling to modern multi-petaflop supercomputers, including advanced $O(N^3)$ algorithms and their implementations, O(N) algorithm, and efforts to reduce global communications and time to solution.

Organizer: Jean-Luc Fattebert
Lawrence Livermore National Laboratory, USA

Organizer: Erik W. Draeger
Lawrence Livermore National Laboratory, USA

10:35-10:55 Scalable Density Functional Theory on Blue Gene/Q
Erik W. Draeger, Lawrence Livermore National Laboratory, USA

11:00-11:20 Algorithmic Rethinking and Code Reengineering for Truly Massively Parallel Ab Initio Molecular Dynamics Simulations
Costas Bekas, IBM Research, USA

11:25-11:45 Decay Properties of Density Matrices: Rigorous Results
Michele Benzi, Emory University, USA

11:50-12:10 A Scalable and Accurate O(N) Parallel Algorithm for Large-Scale First-Principles Molecular Dynamics Simulations
Daniel Osei-Kuffuor and Jean-Luc Fattebert, Lawrence Livermore National Laboratory, USA

Wednesday, February 19

MS13

Minimizing Communication in Linear Algebra - Part I of II

10:35 AM-12:15 PM

Room:Salon D

For Part 2 see MS21

The relative costs of communication compared to computation continue to increase, and historically computation-bound algorithms in linear algebra are becoming communication bound. In order to make efficient use of today's and future hardware, algorithms must be designed in a way that reduces the amount of communication they perform, both across the network and within the memory hierarchy. This minisymposium discusses recent progress in both the practice of designing and implementing dense and sparse linear algebra algorithms and in the theory of analyzing lower bounds on their communication costs.

Organizer: Oded Schwartz
University of California, Berkeley, USA

Organizer: Grey Ballard
Sandia National Laboratories, USA

10:35-10:55 Reconstructing Householder Vectors from TSQR
Grey Ballard, Sandia National Laboratories, USA; James Demmel, University of California, Berkeley, USA; Laura Grigori, INRIA, France; Mathias Jacquelin, Lawrence Berkeley National Laboratory, USA; Hong Diep Nguyen and Edgar Solomonik, University of California, Berkeley, USA

11:00-11:20 Avoiding Synchronization in Geometric Multigrid
Erin C. Carson and Nicholas Knight, University of California, Berkeley, USA; Samuel Williams, Michael Lijewski, and Ann S. Almgren, Lawrence Berkeley National Laboratory, USA; James Demmel, University of California, Berkeley, USA

continued in next column

11:25-11:45 Recent Progress in Nested and Communication Avoiding Incomplete Factorization Preconditioners
Laura Grigori, INRIA, France; Frederic Nataf, Laboratoire Jacques-Louis Lions, France; Sophie Moufawad and Long Qu, INRIA, France

11:50-12:10 High-Bandwidth Communication Avoidance: Oxymoron Or Recipe?
James A. Edwards and Uzi Vishkin, University of Maryland, USA

Wednesday, February 19

MS14

Application and Impact of FASTMath (Frameworks, Algorithms and Scalable Technologies for Mathematics on Next-generation Computers) - Part II of II

10:35 AM-12:15 PM

Room:Salon G

For Part 1 see MS6

This minisymposium will describe the use and impact of FASTMath SciDAC software and technologies on application codes. The DOE FASTMath SciDAC project is focused on development of mathematical methods and software for DOE applications. Technologies included in this minisymposium include structured and unstructured mesh techniques, linear and nonlinear solvers, and time integration. Speakers will discuss the science application challenges, then how their technologies have impacted the simulation code technologies.

Organizer: Carol S. Woodward
Lawrence Livermore National Laboratory, USA

10:35-10:55 Performance Optimization of Block Eigensolvers for Nuclear Structure Computations
H. Metin Aktulga, Sam Williams, Aydin Buluc, Chao Yang, and Esmond G. Ng, Lawrence Berkeley National Laboratory, USA

11:00-11:20 Progress in Adaptive Computational Mechanics Applications Using the Albany Agile Components Framework
Jakob T. Ostien, Glen Hansen, James W. Foulk III, Alejandro Mota, and Mike Veilleux, Sandia National Laboratories, USA

11:25-11:45 Automated Massively Parallel Simulations Starting with Circuit Design Data
Max Bloomfield, Brian Granzow, and Mark S. Shephard, Rensselaer Polytechnic Institute, USA

11:50-12:10 Parallel Infrastructure for Multiscale Simulations
William R. Tobin, Daniel Fovargue, and Mark S. Shephard, Rensselaer Polytechnic Institute, USA

Wednesday, February 19

MS15

Challenges in Parallel Adaptive Mesh Refinement. Part I: Scalable Algorithms

10:35 AM-12:15 PM

Room:Salon H

For Part 2 see MS23

Parallel adaptive mesh refinement (AMR) is a key technique when simulations are required to capture multiscale features. Frequent re-adaptation and repartitioning of the mesh during the simulation can impose significant overhead, particularly in large-scale parallel environments. Further challenges arise due to the availability of accelerated or special-purpose hardware, and the trend toward hierarchical and hybrid compute architectures. Our minisymposium addresses algorithms, scalability, and software issues of parallel AMR on HPC and multi-/manycore platforms. It will discuss novel techniques and applications that demonstrate particular use cases for AMR.

Organizer: Michael Bader
Technische Universität München, Germany

Organizer: Martin Berzins
University of Utah, USA

Organizer: Carsten Burstedde
Universität Bonn, Germany

10:35-10:55 MPI/OpenMP Parallelization of Sam(oa)2 Using Sierpinski Curves
Oliver Meister, Michael Bader, and Kaveh Rahnema, Technische Universität München, Germany

11:00-11:20 Hybrid Octree/Quadtree AMR for Anisotropic Domains
Tobin Isaac, University of Texas at Austin, USA; Carsten Burstedde, Universität Bonn, Germany; Omar Ghattas, University of Texas at Austin, USA

11:25-11:45 A Nested Partitioning Scheme for Adaptive Meshes on Parallel Heterogeneous Clusters
Hari Sundar, George Biros, and Omar Ghattas, University of Texas at Austin, USA

11:50-12:10 Parallel Computing Research Topics in the Context of the Chombo AMR Code
Brian Van Straalen, Lawrence Berkeley National Laboratory, USA

Wednesday, February 19

MS16

Parallel Algorithms for MapReduce-Based Scientific Computing - Part I of II

10:35 AM-12:15 PM

Room:Salon I

For Part 2 see MS24

The goal of our minisymposium is to give an overview of the rapidly emerging use of new parallel processing techniques for scientific computing problems that make use of the Hadoop and MapReduce frameworks and related technologies. We are interested in novel developments on numerical or combinatorial algorithms for Hadoop/MapReduce applied to scientific computing problems in the broad sense, including numerical linear algebra, graph processing, and machine learning. We are also interested in developments on new scalable data analytics environments that go beyond Hadoop/MapReduce, and in success stories about applications using Hadoop/MapReduce.

Organizer: David F. Gleich
Purdue University, USA

Organizer: Paul Constantine
Colorado School of Mines, USA

Organizer: Hans De Sterck
University of Waterloo, Canada

10:35-10:55 A First and Second Introduction to MapReduce
David F. Gleich, Purdue University, USA; *Hans De Sterck*, University of Waterloo, Canada

11:00-11:20 Matrix Factorizations in MapReduce with Applications to Model Reduction
Paul Constantine, Colorado School of Mines, USA; *Austin Benson*, University of California, Berkeley, USA

11:25-11:45 Scaling Up Tensor Decompositions with MapReduce
Evangelos Papalexakis, Carnegie Mellon University, USA

11:50-12:10 Generating Large Graphs with Desired Community Structure
Todd Plantenga, Sandia National Laboratories, USA

Wednesday, February 19

Lunch Break

12:15 PM-1:45 PM

Attendees on their own

IP3

How Much (Execution) Time, Energy, and Power Will my Algorithm Need?

1:45 PM-2:30 PM

Room:Salon F

Chair: Lenny Oliker, Lawrence Berkeley National Laboratory, USA

Do we need to design algorithms differently if our goal is to save energy and power, rather than (or in addition to) time or space? This talk presents a simple thought exercise and a collection of actual experiments on real systems that suggest when and why the answer could be "yes." Importantly, this talk is about speculative ideas more than it is a set of well-developed results. As such, your questions, healthy skepticism, (constructive!) feedback, and offers of collaboration may be even more welcome than usual!

Richard Vuduc
Georgia Institute of Technology, USA

Wednesday, February 19

MS17

Toward Resilient Applications for Extreme-Scale Systems Part III of IV

2:40 PM-4:20 PM

Room:Salon F

As leadership-class computing systems increase in complexity and transistor feature sizes decrease, application codes find themselves less and less able to treat a system as a reliable digital machine. In fact, the high performance computing community has grown increasingly concerned that applications will have to manage resilience issues beyond the current practice of global checkpoint restart. This is expensive at scale and not capable of fixing all types of errors. We discuss alternatives in software and numerical algorithms that can improve the resiliency of applications and manage a variety of faults anticipated in future extreme-scale computing systems.

Organizer: Keita Teranishi
Sandia National Laboratories, USA

Organizer: Mark Hoemmen
Sandia National Laboratories, USA

Organizer: Jaideep Ray
Sandia National Laboratories, USA

Organizer: Michael A. Heroux
Sandia National Laboratories, USA

2:40-3:00 Evaluating the Impact of Faults and Recovery Mechanisms in Exascale Applications
Martin Schulz and Ignacio Laguna, Lawrence Livermore National Laboratory, USA

3:05-3:25 Managing Resilience in Exascale Computing
Saurabh Hukerikar, Pedro Diniz, and Bob Lucas, University of Southern California, USA

3:30-3:50 Tolerable Fault Tolerance
Jeff Keasler, Lawrence Livermore National Laboratory, USA

3:55-4:15 GVR-Enabled Trilinos: An Outside-In Approach for Resilient Computing
Ziming Zheng and Zachary Rubenstein, University of Chicago, USA; Andrew Chien, University of Chicago and Argonne National Laboratory, USA

Wednesday, February 19

MS18

Frontiers of Performance Analysis and Performance Tools - Part I of II

2:40 PM-4:20 PM

Room:Salon A

For Part 2 see MS26
This minisymposium will explore the "frontiers" of understanding the performance and scalability of parallel applications as represented by the latest research efforts in field. The increasing complexity of parallel systems and variety of programming methodologies redefines the challenges for performance analysis techniques and the tools that support them. Speakers will give their impression of the new territory to trek in areas such as modeling and measurement of energy and power; new and enhanced tools for measuring, finding, and explaining bottlenecks; and applications of models and tools for autotuning. The minisymposium aims to address these themes at all scales.

Organizer: Richard Vuduc
Georgia Institute of Technology, USA

Organizer: Allen Malony
Paratools Inc., USA

2:40-3:00 Holistic Performance Measurement and Analysis for High End Applications
Karen Karavanic, Portland State University, USA

3:05-3:25 Pattern-Driven Node-Level Performance Engineering
Jan Treibig, Erlangen Regional Computing Center, Germany

3:30-3:50 Application of the Pi Theorem from Dimensional Analysis to Computer Performance Modeling
Bob Numrich, City University of New York, USA

3:55-4:15 A Methodology for Characterizing the Opportunity and Feasibility of Reconfigurable Memory Hierarchies for Improved Energy Efficiency
Laura Carrington, San Diego Supercomputer Center, USA

Wednesday, February 19

MS19

Implementation Aspects of Parallel-in-time Methods on HPC Systems

2:40 PM-4:20 PM

Room:Salon B

Parallel-in-time integration methods for time-dependent PDEs are a promising approach to augment existing parallelization strategies. Adding concurrency in the temporal direction can help extend prevailing strong-scaling limits. Several theoretical results exist as well as small-scale application examples, illustrating convergence properties, possible application areas and bounds on parallel efficiency of some of these methods. However, their implementation, especially in combination with existing spatial decomposition techniques in large- and extreme-scale applications, poses new and often unforeseen challenges. A key goal of the minisymposium is the discussion of implementation strategies for different time-parallel methods and recent results on HPC platforms.

Organizer: Daniel Ruprecht
University of Lugano, Switzerland

Organizer: Robert Speck
Jülich Supercomputing Centre, Germany

Organizer: Matthew Emmett
Lawrence Berkeley National Laboratory, USA

2:40-3:00 Implementation Strategies for Parallel-in-Time Methods
Daniel Ruprecht, University of Lugano, Switzerland; Robert Speck, Jülich Supercomputing Centre, Germany; Rolf Krause, University of Lugano, Switzerland

continued on next page

3:05-3:25 The Parallel Full Approximation Scheme in Space and Time (PFASST) on Extreme Scales
Matthew Emmett, Lawrence Berkeley National Laboratory, USA

3:30-3:50 Hybrid Mpi-OpenMP Implementation of Waveform Relaxation
Benjamin Ong and *Scott High*, Michigan State University, USA; Felix Kwok, University of Geneva, Switzerland

3:55-4:15 The Parareal Algorithm - Applications to Fusion Plasma Physics
Debasmita Samaddar, UK Atomic Energy Authority, United Kingdom

Wednesday, February 19

MS20

Density Functional Theory: Large-scale Algorithms and O(N) Methods - Part II of II

2:40 PM-4:20 PM

Room:Salon C

For Part 1 see MS12

Density Functional Theory (DFT) is a fundamental predictive model to study matter at atomistic scale. It is widely used to help understand and design advanced materials. DFT simulations, however, are limited in size due to their high computational cost and the cubic complexity of typical solvers. In addition, traditional DFT implementations typically involve global operations which prevent efficient use of a large number of processors. We will discuss recent developments to address these issues and enable scaling to modern multi-petaflop supercomputers, including advanced $O(N^3)$ algorithms and their implementations, O(N) algorithm, and efforts to reduce global communications and time to solution.

Organizer: Jean-Luc Fattebert
Lawrence Livermore National Laboratory, USA

Organizer: Erik W. Draeger
Lawrence Livermore National Laboratory, USA

2:40-3:00 High-Order and Enriched Finite Element Methods for Electronic Structure
John Pask, Lawrence Livermore National Laboratory, USA

3:05-3:25 Recent Progress on the Pole Expansion and Selected Inversion Method for Solving Kohn-Sham Density Functional Theory
Lin Lin, Lawrence Berkeley National Laboratory, USA

3:30-3:50 Divide-Conquer-Recombine Algorithms for Metascalable Quantum Molecular Dynamics Simulations
Aiichiro Nakano, University of Southern California, USA

3:55-4:15 Multiresolution DFT and O(N) Methods
George Fann, Oak Ridge National Laboratory, USA

Wednesday, February 19

MS21

Minimizing Communication in Linear Algebra - Part II of II

2:40 PM-4:20 PM

Room:Salon D

For Part 1 see MS13

The relative costs of communication compared to computation continue to increase, and historically computation-bound algorithms in linear algebra are becoming communication bound. In order to make efficient use of today's and future hardware, algorithms must be designed in a way that reduces the amount of communication they perform, both across the network and within the memory hierarchy. This minisymposium discusses recent progress in both the practice of designing and implementing dense and sparse linear algebra algorithms and in the theory of analyzing lower bounds on their communication costs.

Organizer: Oded Schwartz
University of California, Berkeley, USA

Organizer: Grey Ballard
Sandia National Laboratories, USA

2:40-3:00 A New Sparse Inertia-Revealing Factorization
Sivan A. Toledo, Tel Aviv University, Israel

3:05-3:25 A Lower Bound Technique for Communication on BSP with Application to the FFT
Gianfranco Bilardi, Michele Scquizzato, and Francesco Silvestri, University of Padova, Italy

3:30-3:50 Tradeoffs between Synchronization, Communication, and Work in Parallel Linear Algebra Computations
Edgar Solomonik, Erin C. Carson, Nicholas Knight, and James Demmel, University of California, Berkeley, USA

3:55-4:15 Communication Optimal Loop Nests
Michael Christ, James W. Demmel, and *Nicholas Knight*, University of California, Berkeley, USA; Thomas Scanlon, University of California, USA; Katherine Yelick, University of California, Berkeley, USA

Wednesday, February 19

MS22

Parallel Processing with Python - Part I of II

2:40 PM-4:20 PM

Room:Salon G

For Part 2 see MS30

Python is establishing itself as an effective language for high-performance computing. Its expressiveness makes it a natural language for interactive and exploratory parallel computation; it easily interfaces with HPC libraries and frameworks to provide high-level interfaces; and its dynamism and flexibility allows for domain-specific languages and compilers. These two minisymposia present advances in Python for HPC focused on two areas: HPC data structures and libraries for easy application development, and a selection of HPC Python packages that span the foundational (MPI-centric) to modern (interactive parallelism with IPython, and python-based domain-specific languages for HPC).

Organizer: Kurt W. Smith
Enthought, Inc., USA

Organizer: William F. Spotz
Sandia National Laboratories, USA

2:40-3:00 Distributing NumPy Using Global Arrays
Jeff Daily, Pacific Northwest National Laboratory, USA

3:05-3:25 PyTrilinos: Parallel Solvers and Simulation Tools for Python
William F. Spotz, Sandia National Laboratories, USA

3:30-3:50 ODIN: Bringing NumPy's Strengths to Distributed Computing
Kurt W. Smith, Enthought, Inc., USA

3:55-4:15 Bohrium: Unmodified NumPy Code on CPU, GPU, and Cluster
Mads Kristensen, *Simon Lund*, Troels Blum, Kenneth Skovhede, and Brian Vinter, University of Copenhagen, Denmark

Wednesday, February 19

MS23

Challenges in Parallel Adaptive Mesh Refinement. Part II: Advanced Applications

2:40 PM-4:20 PM

Room:Salon H

For Part 1 see MS15
For Part 2 see MS31

Parallel adaptive mesh refinement (AMR) is a key technique when simulations are required to capture multiscale features. Frequent re-adaptation and repartitioning of the mesh during the simulation can impose significant overhead, particularly in large-scale parallel environments. Further challenges arise due to the availability of accelerated or special-purpose hardware, and the trend toward hierarchical and hybrid compute architectures. Our minisymposium addresses algorithms, scalability, and software issues of parallel AMR on HPC and multi-/manycore platforms. It will discuss novel techniques and applications that demonstrate particular use cases for AMR.

Organizer: Michael Bader
Technische Universität München, Germany

Organizer: Martin Berzins
University of Utah, USA

Organizer: Carsten Burstedde
Universität Bonn, Germany

2:40-3:00 Parallel Lattice Boltzmann Methods with Static Adaptivity
Florian Schornbaum, Ulrich J. Ruede, Harald Koestler, and David Staubach, University of Erlangen-Nuremberg, Germany

3:05-3:25 A Block-Structured Parallel Adaptive Lattice-Boltzmann Method for Rotating Geometries
Ralf Deiterding, German Aerospace Center (DLR), Germany; Stephen L. Wood, University of Tennessee, Knoxville, USA

3:30-3:50 AMR for Fluid-Structure and Radiation Problems on Recent and Novel Architectures
Qingyu Meng, University of Utah, USA

3:55-4:15 A Volume Integral Equation Solver for Boundary Value Problems with Highly Heterogeneous Coefficients
George Biros, University of Texas at Austin, USA

Wednesday, February 19

MS24

Parallel Algorithms for MapReduce-Based Scientific Computing - Part II of II

2:40 PM-4:20 PM

Room:Salon I

For Part 1 see MS16

The goal of our minisymposium is to give an overview of the rapidly emerging use of new parallel processing techniques for scientific computing problems that make use of the Hadoop and MapReduce frameworks and related technologies. We are interested in novel developments on numerical or combinatorial algorithms for Hadoop/MapReduce applied to scientific computing problems in the broad sense, including numerical linear algebra, graph processing, and machine learning. We are also interested in developments on new scalable data analytics environments that go beyond Hadoop/MapReduce, and in success stories about applications using Hadoop/MapReduce.

Organizer: David F. Gleich
Purdue University, USA

Organizer: Paul Constantine
Colorado School of Mines, USA

Organizer: Hans De Sterck
University of Waterloo, Canada

2:40-3:00 Apache Giraph: Large-Scale Graph Processing Infrastructure on Hadoop
Avery Ching, Facebook, USA

3:05-3:25 Large-Scale Numerical Computation Using a Data Flow Engine
Matei Zaharia, Massachusetts Institute of Technology, USA

3:30-3:50 REEF - Beyond MapReduce by Re-Layering the Big Data Stack
Markus Weimer, Microsoft Research, USA

3:55-4:15 Traditional and Streaming MapReduce via MPI for Graph Analytics
Steve Plimpton, Karen D. Devine, and Timothy Shead, Sandia National Laboratories, USA

Wednesday, February 19

Coffee Break

4:20 PM-4:50 PM

Room:Salon E

MS25

Scalable Algorithms for Computational Statistical Inference - Part I of II

4:50 PM-6:05 PM

Room:Salon I

For Part 2 see MS40

High-performance algorithms for statistical inference are a required technology for enabling analysis of large datasets. By "large", we refer to datasets that are in the same order of magnitude as the total RAM of current leadership architectures, which currently can reach and exceed one petabyte. Despite the many mathematical breakthroughs in algorithms for statistical inference, there remain several challenges in scaling this algorithms to large datasets. In this minisymposium, we will discuss recent advances in scalable methods for statistical inference algorithms.

Organizer: George Biros
University of Texas at Austin, USA

4:50-5:10 Parallel Methods for Bayesian Network Structure Learning
Srinivas Aluru, Georgia Institute of Technology, USA

5:15-5:35 Parallel Algorithms for Point-Correlation Functions
William March, University of Texas at Austin, USA

5:40-6:00 Parallel Algorithms Fo Nearest Neighbor Searches
Bo Xiao, Georgia Institute of Technology, USA

Wednesday, February 19

MS26

Frontiers of Performance Analysis and Performance Tools - Part II of II

4:50 PM-6:30 PM

Room:Salon A

For Part 1 see MS18

This minisymposium will explore the "frontiers" of understanding the performance and scalability of parallel applications as represented by the latest research efforts in field. The increasing complexity of parallel systems and variety of programming methodologies redefines the challenges for performance analysis techniques and the tools that support them. Speakers will give their impression of the new territory to trek in areas such as modeling and measurement of energy and power; new and enhanced tools for measuring, finding, and explaining bottlenecks; and applications of models and tools for autotuning. The minisymposium aims to address these themes at all scales.

Organizer: Richard Vuduc
Georgia Institute of Technology, USA

Organizer: Allen Malony
Paratools Inc., USA

4:50-5:10 Mummi: A Modeling Infrastructure for Exploring Power and Execution Time Tradeoffs
Valerie Taylor, Texas A&M University, USA

5:15-5:35 Automating Performance Optimization for Compute Nodes of Hpc Systems
James C. Browne and Leonardo Fialho, University of Texas at Austin, USA

5:40-6:00 Performance Analysis of Mpi+openmp Programs on Scalable Parallel Systems
John Mellor-Crummey, Rice University, USA

6:05-6:25 Customizing Libraries with Dsls and Autotuning
Boyana Norris, Argonne National Laboratory, USA

Wednesday, February 19

MS27

Abstractions for Finite Element Assembly on Multi-Core Architectures

4:50 PM-6:30 PM

Room:Salon B

With the emergence of high-concurrency architectures, many high-performance finite element codes face a significant refactoring challenge to take advantage of this new capability. While linear algebra kernels have been considered, the process of assembling the global vectors and sparse matrices does not generalize easily across applications. This minisymposium explores software and algorithmic abstractions that enable efficient implementation of parallel finite element assembly on a myriad of multi-core architectures. Topics considered will include architecture independent abstractions for: dense computation of local stiffness matrices and vectors, scatter of local dense operators to global sparse data structures, and management of multiphysics systems.

Organizer: Eric Phipps
Sandia National Laboratories, USA

Organizer: Eric C. Cyr
Sandia National Laboratories, USA

Organizer: Roger P. Pawlow
Sandia National Laboratories, USA

4:50-5:10 Design Paradigms to Accommodate Architectural Uncertainty in Multiphysics Applications
James C. Sutherland, Christopher Earl, Tony Saad, and Matthew Might, University of Utah, USA

5:15-5:35 Open-Source, Multi-Physics Finite Element Analysis Using LibMesh and Grins
Roy Stogner and Paul T. Bauman, University of Texas at Austin, USA

5:40-6:00 Portable Manycore Sparse Linear System Assembly Algorithms and Performance Tradeoffs
H. Carter Edwards, Sandia National Laboratories, USA

6:05-6:25 Development of a Mini-Driver Application to Test Fem Assembly on Modern Architectures
Matthew Bettencourt and Eric C. Cyr, Sandia National Laboratories, USA

Wednesday, February 19

MS28

Parallel Quantum Chemistry

4:50 PM-6:30 PM

Room:Salon C

Recent developments in parallel quantum chemistry on large clusters and accelerators are presented. This minisymposium seeks to engage computer scientists and others outside the domain area. Topics include acceleration and offload of computationally intensive components (integral calculations), distributed memory tensor contractions exploiting symmetries, and portable parallel infrastructures. The speakers will also outline their ideas for future directions for the parallel quantum chemistry community.

Organizer: Edmond Chow
Georgia Institute of Technology, USA

Organizer: Robert Harrison
Brookhaven National Laboratory, USA

Organizer: Daniel Crawford
Virginia Tech, USA

4:50-5:10 Losing the Barriers: Increased Performance and Parallelism through Control and Data Flow
Theresa Windus, Iowa State University, USA

5:15-5:35 Tensor Hypercontraction and Graphical Processing Units for Electronic Structure and Ab Initio Molecular Dynamics
Todd Martinez, Stanford University, USA

5:40-6:00 Qc4
J Ramanujam, Louisiana State University, USA

6:05-6:25 Distributed Contraction of Symmetric Tensors
P. Sadayappan, Samyam Rajbhandari, Akshay Nikam, Kevin Stock, and Pai-Wei Lai, Ohio State University, USA; Sriram Krishnamoorthy, Pacific Northwest National Laboratory, USA

Wednesday, February 19

MS29

Solving Large-scale Linear Systems of Equations on Heterogeneous Distributed-memory Computers

4:50 PM-6:30 PM

Room:Salon D

Due to their computational capabilities and power efficiencies, compute nodes on high performance distributed-memory computers are commonly equipped with either accelerators or coprocessors. However, it is still challenging to effectively utilize these devices due to their limited amount of the device memory and the high cost of data transfer. This minisymposium highlights recent efforts to overcome these challenges in order to solve large-scale linear systems of equations on such heterogeneous distributed-memory computers. Our topic includes out-of-core algorithms and dynamic or static scheduling schemes to obtain scalable performance on such architectures.

Organizer: Ed D'Azevedo
Oak Ridge National Laboratory, USA

Organizer: Ichitaro Yamazaki
University of Tennessee, Knoxville, USA

4:50-5:10 Parallel LU Factorizations on Intel MIC in AORSA
Ed D'Azevedo and *Judith Hill*, Oak Ridge National Laboratory, USA

5:15-5:35 A Performance Study of Solving a Large Dense Matrix for Radiation Heat Transfer Using Intel Xeon Phi Coprocessors
Kwai L. Wong, University of Tennessee and Oak Ridge National Laboratory, USA; Ed D'Azevedo, Oak Ridge National Laboratory, USA; Shiquan Su, University of Tennessee, USA; Ki Sing Chan and Yue Qian, Chinese University of Hong Kong, Hong Kong

5:40-6:00 Optimized GPU Kernels for Sparse Factorization
Ichitaro Yamazaki, University of Tennessee, Knoxville, USA; Stan Tomov, University of Tennessee, USA

6:05-6:25 Overview of Distributed Dense Linear Algebra over StarPU Runtime
Mathieu Faverge, Bordeaux Institute of Technology, France; Emmanuel Agullo, INRIA, France

Wednesday, February 19

MS30

Parallel Processing with Python - Part II of II

4:50 PM-6:30 PM

Room:Salon G

For Part 1 see MS22

Python is establishing itself as an effective language for high-performance computing. Its expressiveness makes it a natural language for interactive and exploratory parallel computation; it easily interfaces with HPC libraries and frameworks to provide high-level interfaces; and its dynamism and flexibility allows for domain-specific languages and compilers. These two minisymposia present advances in Python for HPC focused on two areas: HPC data structures and libraries for easy application development, and a selection of HPC Python packages that span the foundational (MPI-centric) to modern (interactive parallelism with IPython, and python-based domain-specific languages for HPC).

Organizer: Kurt W. Smith
Enthought, Inc., USA

Organizer: William F. Spotz
Sandia National Laboratories, USA

4:50-5:10 Scaling Up Python with mpi4py
Aron Ahmadia, US Army Engineer Research and Development Center, USA; Lisandro Dalcin, Centro Int. de Métodos Computacionales en Ingeniería, Argentina

5:15-5:35 Using the IPython Notebook for Reproducible Parallel Computing
Brian E. Granger, California Polytechnic State University, San Luis Obispo, USA; Matthias Bussonnier, Institut Curie, France; Jonathan Frederic, California Polytechnic State University, San Luis Obispo, USA; Bradley Froehle, Paul Ivanov, and Thomas Kluyver, University of California, Berkeley, USA; Evan Patterson, Enthought, Inc., USA; Fernando Perez and Min Ragan-Kelley, University of California, Berkeley, USA; Zachary Sailer, California Polytechnic State University, San Luis Obispo, USA

continued on next page

5:40-6:00 IPYTHON.PARALLEL, Tools for Interactive Parallel Computing
Min Ragan-Kelley, University of California, Berkeley, USA; Brian E. Granger, California Polytechnic State University, San Luis Obispo, USA; Fernando Perez, University of California, Berkeley, USA

6:05-6:25 Going from a Python Embedded DSL to a Massively Parallel Heterogeneous AMR CFD Code
Andy R. Terrel and Kyle T. Mandli, University of Texas at Austin, USA; Donna Calhoun, Boise State University, USA; Carsten Burstedde, Universität Bonn, Germany

Wednesday, February 19

MS31

Challenges in Parallel Adaptive Mesh Refinement. Part III: Advanced Applications

4:50 PM-6:30 PM

Room:Salon H

For Part 2 see MS23

Parallel adaptive mesh refinement (AMR) is a key technique when simulations are required to capture multiscale features. Frequent re-adaptation and repartitioning of the mesh during the simulation can impose significant overhead, particularly in large-scale parallel environments. Further challenges arise due to the availability of accelerated or special-purpose hardware, and the trend toward hierarchical and hybrid compute architectures. Our minisymposium addresses algorithms, scalability, and software issues of parallel AMR on HPC and multi-/manycore platforms. It will discuss novel techniques and applications that demonstrate particular use cases for AMR.

Organizer: Michael Bader
Technische Universität München, Germany

Organizer: Martin Berzins
University of Utah, USA

Organizer: Carsten Burstedde
Universität Bonn, Germany

4:50-5:10 Parallel, Adaptive, Multilevel Solution of Nonlinear Systems Arising in Phase Field Problems
Peter Bollada, *Peter K. Jimack*, and Andrew Mullis, University of Leeds, United Kingdom

5:15-5:35 Parallel, Adaptive Finite Volume Method for Solving Conservation Laws on Mapped, Multiblock Domains
Donna Calhoun, Boise State University, USA; Carsten Burstedde, Universität Bonn, Germany

continued in next column

5:40-6:00 Parallel Strategies for Modeling Storm Surge With Adaptive Mesh Refinement
Kyle T. Mandli, University of Texas at Austin, USA; Marsha Berger, Courant Institute of Mathematical Sciences, New York University, USA; Clint Dawson, University of Texas at Austin, USA

6:05-6:25 Adaptive Discontinuous Galerkin Methods for Compressible Flow Applications
Robert Kloefkorn, National Center for Atmospheric Research, USA

Wednesday, February 19

MS32

Toward Resilient Applications for Extreme-Scale Systems Part IV of IV

4:50 PM-6:55 PM

Room:Salon F

As leadership-class computing systems increase in complexity and transistor feature sizes decrease, application codes find themselves less and less able to treat a system as a reliable digital machine. In fact, the high performance computing community has grown increasingly concerned that applications will have to manage resilience issues beyond the current practice of global checkpoint restart. This is expensive at scale and not capable of fixing all types of errors. We discuss alternatives in software and numerical algorithms that can improve the resiliency of applications and manage a variety of faults anticipated in future extreme-scale computing systems.

Organizer: Keita Teranishi
Sandia National Laboratories, USA

Organizer: Mark Hoemmen
Sandia National Laboratories, USA

Organizer: Jaideep Ray
Sandia National Laboratories, USA

Organizer: Michael A. Heroux
Sandia National Laboratories, USA

4:50-5:10 Finite Difference Stencils Robust to Silent Data Corruption
Jaideep Ray, *Jackson Mayo*, and Rob Armstrong, Sandia National Laboratories, USA

5:15-5:35 Reexamining Algorithm-Based Fault Tolerance for Exascale Architectures
Jeff Vetter and Dong Li, Oak Ridge National Laboratory, USA

5:40-6:00 Low-Communication Multigrid, with Applications to Time-Dependent Adjoints, in-Situ Visualization, and Resilience
Jed Brown, Argonne National Laboratory, USA; Mark Adams, Lawrence Berkeley National Laboratory, USA

continued in next column

6:05-6:25 Experimentally Tuned Algorithm-Based Fault Tolerance Techniques for Matrix Multiplication and FFT on GPUs
Paolo Rech and Luigi Carro, Federal University of Rio Grande do Sul, Brazil

6:30-6:50 An Algorithmic Approach to Silent Error Resilience
Rob Schreiber, Hewlett Packard Corporation, USA; Austin Benson, University of California, Berkeley, USA

Wednesday, February 19

PD1

High-Performance Computing on Graphs

8:00 PM-9:00 PM

Room:Salon A

Chair: John R. Gilbert, University of California, Santa Barbara, USA

Computations on large graphs are becoming ubiquitous, in such diverse areas as genomics, finance, social networks, commercial analytics, and many others. A new ecosystem of algorithms and systems for graph analysis is evolving. What will form the foundations of this ecosystem, and what role will parallel and high-performance computing play? The panel will focus on five questions about the future evolution of high-performance computing on graphs:

1. What are the key directions for algorithms research?

2. What are the key directions for software tools and libraries?

3. What are the key directions for programming languages and runtimes?

4. What are the key directions for processor and network architectures?

5. What are the key directions for education?

We expect a stimulating discussion among the attendees and panelists.

Panelists:

To Be Announced

Wednesday, February 19

PD2

Resilience at Exascale: Should it Worry Application Developers?

8:00 PM-9:00 PM

Room:Salon F

Chair: Marc Snir, Argonne National Laboratory, USA

Many fear that supercomputers are becoming less reliable, so that software will need to detect and correct errors, or else be immune to them. This has engendered work in algorithm-based fault tolerance (ABFT) techniques that "ruggedize" algorithms. But fault-tolerant kernels don't protect whole applications, We will discuss this situation and consider several questions, including: Must exascale hardware be less reliable than today? Will it be mandatory to expose failures to the application? Do we have general ABFT techniques, or must we develop a new technique for each important algorithm? How will the possibility of errors and failures affect application development?

Panelists:

Shekhar Borkar
Intel Corporation, USA

Michael Heroux
Sandia National Laboratories, USA

Rob Schreiber
Hewlett Packard Corporation, USA

Wednesday, February 19

PD3

Parallel Computing for Quantum Models in Material Sciences, Chemistry and Biology: Current and Future Directions

8:00 PM-9:00 PM

Room:Salon I

We have recently witnessed significant progress in improving the efficiency and throughput of electronic structure calculations through algorithmic advances and effective parallelization on distributed multi-core systems. Moderately-sized calculations are now routine, but these calculations remain challenging for large and complex systems. New ideas are needed to address complexity and scalability issues in computationally demanding excited state calculations. We will lead a discussion on how we, as a community, should go forward in addressing current and future computational challenges in materials science and chemistry simulations: what are the viable paths, and what lessons have we learned in the last several decades?

Panelists:

Chair: Daniel Crawford
Virginia Tech, USA

Chair: Jean-Luc Fattebert
Lawrence Livermore National Laboratory, USA

Chair: Robert Harrison
Brookhaven National Laboratory and Stony Brook University, USA

Chair: Eric Polizzi
University of Massachusetts, Amherst, USA

Chair: Chao Yang
Lawrence Berkeley National Laboratory, USA

Thursday, February 20

Registration

7:45 AM-5:15 PM

Room:Oregon Ballroom Foyer

IP4

Particles, Grids and HPC for Simulations in Fluid Dynamics

8:15 AM-9:00 AM

Room:Salon F

Chair: George Biros, University of Texas at Austin, USA

Simulation techniques in fluid dynamics can be broadly distinguished as grid based (e.g finite volume, finite differences, finite elements) and particle/meshless methods (including SPH and vortex methods). In this talk I present a methodology to transition between these two classes of discretisations and discuss in particular issues pertaining to their accuracy and their HPC implementation. I will present results from large scale flow simulations of compressible and incompressible flows, including recent simulations of two-phase flows reaching 14.5 PFLops.

Petros Koumoutsakos
ETH Zürich, Switzerland

Intermission

9:00 AM-9:10 AM

Forward Looking Panel

9:10 AM-10:10 AM

Room:Salon F

Thursday, February 20

Coffee Break

10:10 AM-10:35 AM

Room:Salon E

MS33

Parallel Methods and Algorithms for Extreme Computing - Part I of III

10:35 AM-12:15 PM

Room:Salon F

For Part 2 see MS41

Multicore processors and accelerators are universally available as both collections of homogeneous standard microprocessors and as attached heterogeneous co-processors. Application and library software developers may often effectively use these processors and some general approaches have emerged. It is widely recognized that careful design of software and data structures, with effective memory management, are the most critical issues to obtain scalable optimized performance on those systems. In these minisymposia we discuss current experiences and development of applications, libraries and frameworks using a variety of hardware. Speakers will address performance results and software design.

Organizer: Serge G. Petiton
CNRS/LIFL and INRIA, France

Organizer: Michael A. Heroux
Sandia National Laboratories, USA

Organizer: Kengo Nakajima
University of Tokyo, Japan

10:35-10:55 Toward a Portable, Resilient Application Design for Scalable Manycore Computing Systems
Michael A. Heroux, Sandia National Laboratories, USA

continued in next column

11:00-11:20 Preconditioned Iterative Solvers on Manycore Architectures
Satoshi Ohshima, Takahiro Katagiri, and Kengo Nakajima, University of Tokyo, Japan

11:25-11:45 Programming Paradigms for Emerging Architectures Applied to Asynchronous Krylov Eigensolvers
Christophe Calvin, CEA Saclay, France; France Boillod-Cerneux, CNRS/LIFL, France; Fan Ye, CEA, France; Hervé Galicher, CEA Saclay, France; Serge Petiton, CNRS, France

11:50-12:10 Achieving Many-core Performance Portability with Kokkos
Christian Trott and H. Carter Edwards, Sandia National Laboratories, USA

Thursday, February 20

MS34

Auto-tuning Technologies for Extreme-Scale Solvers - Part I of III

10:35 AM-12:15 PM

Room:Salon A

For Part 2 see MS42

Numerical solvers that accommodate extreme levels of parallelism will be essential for realizing exa-flops simulations. The development of such solvers will require not only algorithmic innovations but also significant effort for their implementations. Noticeably, auto-tuning (AT) technology has provided sustained high performance on various computer architectures, and it is expected to also play a role in the exa-flops era, especially in the context of solvers. In this minisymposium, we will discuss algorithms and implementations for extreme-scale numerical solvers with regards to AT. We target eigenvalue solvers, iterative linear equations solvers, and fundamental matrix decompositions such as the QR decomposition.

Organizer: Takahiro Katagiri
University of Tokyo, Japan

Organizer: Osni A. Marques
Lawrence Berkeley National Laboratory, USA

Organizer: Toshiyuki Imamura
RIKEN, Japan

Organizer: Leroy A. Drummond
Lawrence Berkeley National Laboratory, USA

10:35-10:55 Extreme-Scale Parallel Symmetric Eigensolver for Very Small-Size Matrices Using A Communication-Avoiding for Pivot Vectors
Takahiro Katagiri, Junichi Iwata, and Kazuyuki Uchida, University of Tokyo, Japan

continued on next page

11:00-11:20 A Parallel Two-grid Polynomial Jacobi-Davidson Algorithm for Large Sparse PDE Eigenvalue Problems
Feng-Nan Hwang, National Central University, Taiwan; Tsung-Ming Huang, National Taiwan Normal University, Taiwan; Weichung Wang, National Taiwan University, Taiwan; Yu-Fen Cheng, National Center for High-performance Computing, Taiwan

11:25-11:45 A Hierarchical Parallel Software Package of a Complex Moment Based Eigensolver
Yasunori Futamura and Tetsuya Sakurai, University of Tsukuba, Japan

11:50-12:10 Xabclib: Parallel Iterative Linear Solver with Run-Time Auto-Tuning
Takao Sakurai, Hitachi Ltd., Japan; Takahiro Katagiri, University of Tokyo, Japan; Mitsuyoshi Igai, Hitachi ULSI Systems Corporation, Japan; Satoshi Ohshima and Hisayasu Kuroda, University of Tokyo, Japan; Ken Naono, Hitachi Asia Malaysia, Malaysia

Thursday, February 20

MS35

Efficient Approaches to Architectural Simulation for Exascale System Evaluation

10:35 AM-12:15 PM

Room:Salon B

With the number of processing elements predicted to soar in future generations of high performance computers, brute force application of cycle accurate system simulation will not be a practical means for the full system performance predictions that will be required for system codesign and forward looking application development. Appropriate simulation approaches must adopt high level abstractions for computational efficiency, yet retain enough accuracy and detail in order to properly predict the broad metrics of interest. This minisymposium surveys recent progress in the development and application of efficient approaches for full scale system simulation.

Organizer: Joseph Kenny
Sandia National Laboratories, USA

Organizer: Jeremiah Wilke
Sandia National Laboratories, USA

10:35-10:55 Title Not Available at Time of Publication
Joseph Kenny, Sandia National Laboratories, USA

11:00-11:20 Modeling the Performance Repurcussions of Hpc Applications on Bandwidth Limited Exascale Systems
Laura Carrington, San Diego Supercomputer Center, USA

11:25-11:45 Title Not Available at Time of Publication
Adolfy Hoisie, Los Alamos National Laboratory, USA

11:50-12:10 Constructing High-Level Application Models for Exascale Co-Design Simulations
Sudhakar Yalamanchili, Georgia Institute of Technology, USA

Thursday, February 20

MS36

Electronic Structure Calculations: Parallel Algorithms and Applications Part I

10:35 AM-12:15 PM

Room:Salon C

For Part 2 see MS44

Electronic structure calculations and their applications are among the most challenging and computationally demanding science and engineering problems. This minisymposium aims at presenting and discussing new numerical and parallel processing avenues that are suitable for modern computing architectures, for achieving ever higher level of accuracy and scalability in DFT, TDDFT and other types of ground and excited states simulations. We propose to bring together physicists/chemists who are involved in improving the numerical development of widely known quantum chemistry and solid-state physics application software packages, with mathematicians/computer scientists who are focusing on advancing the required state-of-the-art mathematical algorithms and parallel implementation.

Organizer: Chao Yang
Lawrence Berkeley National Laboratory, USA

Organizer: Eric Polizzi
University of Massachusetts, Amherst, USA

10:35-10:55 Berkeleygw for Excited States Calculations
Jack Deslippe, National Energy Research Scientific Computing Center, USA

11:00-11:20 New Development in Nwchem
Bert de Jong, Lawrence Berkeley National Laboratory, USA

11:25-11:45 Using Siesta to Solve Large-Scale Electronic Structure Problem
Georg Huh, Barcelona Supercomputing Center, Spain

continued on next page

Thursday, February 20

MS36

Electronic Structure Calculations: Parallel Algorithms and Applications Part I

10:35 AM-12:15 PM

continued

11:50-12:10 Electronic Structure Calculation based on Daubechies Wavelets: BigDFT

Laura Ratcliff, CEA, France; Stephan Mohr, University of Basel, Switzerland; Paul Boulanger, CNRS, France; Luigi Genovese and Damien Caliste, CEA, France; Stefan Goedecker, University of Basel, Switzerland; *Thierry Deutsch*, CEA, France

Thursday, February 20

MS37

Toward Multilevel Solvers for Exascale - Part I of II

10:35 AM-12:15 PM

Room:Salon D

For Part 2 see MS45

The discretization of PDEs arising from applications such as flow problems, chemistry, or material science demands the fast solution of large systems of equations. The minisymposium will highlight the algorithmic design of robust highly parallel multilevel solvers. Here, the nature of the considered equations enforces a global transport of information which is getting increasingly challenging with increasing parallelism. Further, non-linearities have to be properly treated and time-dependencies may be exploited for parallelism. The talks will present optimization and parallelization approaches in order to construct the next generation of efficient solvers.

Organizer: Matthias Bolten
University of Wuppertal, Germany

Organizer: Björn Gmeiner
Universität Erlangen, Germany

10:35-10:55 Scalable Multilevel Stokes Solver for Mantle Convection Problems

Björn Gmeiner, Universität Erlangen, Germany

11:00-11:20 Upscaling Multigrid Towards Exascale Computing

Gabriel Wittum, Goethe University Frankfurt, Germany

11:25-11:45 Improving the Performance of Algebraic Multigrid Using Structured Coarse Grids

Ray S. Tuminaro, Jonathan J. Hu, and Andrey Prokopenko, Sandia National Laboratories, USA

11:50-12:10 Scalable Solvers for Multi-Phase Flow: Algebraic Multigrid for Discontinuous Galerkin and Accelerator Integration

Steffen Müthing, University of Stuttgart, Germany

Thursday, February 20

MS38

Software Productivity for the Next Generation of Scientific Applications - Part I of II

10:35 AM-12:15 PM

Room:Salon G

For Part 2 see MS46

While extreme-scale architectures provide unprecedented resources for scientific discovery, the community faces daunting productivity challenges for parallel application development. Difficulties include increasing complexity of algorithms and computer science techniques required in multiscale and multiphysics applications, the imperative of portable performance in the midst of dramatic and disruptive architectural changes, the realities of large legacy code bases, and human factors arising in distributed multidisciplinary research teams pursuing extreme parallel performance. Speakers in this minisymposium will discuss promising approaches underway in software development tools, libraries, and methodologies to address these productivity challenges in extreme-scale parallel scientific software.

Organizer: Lois C. McInnes
Argonne National Laboratory, USA

Organizer: Michael A. Heroux
Sandia National Laboratories, USA

Organizer: Hans Johansen
Lawrence Berkeley National Laboratory, USA

10:35-10:55 Challenges and Opportunities in Extreme-Scale Application Software Productivity

Hans Johansen and Jeffrey N. Johnson, Lawrence Berkeley National Laboratory, USA; Lois C. McInnes, Argonne National Laboratory, USA

11:00-11:20 NSF Activities to Support Software for Next Generation Systems

Daniel Katz, National Science Foundation, USA

continued on next page

11:25-11:45 Software/Science Co-Development: Software Engineering for Large-Scale in Silico Neuroscience Research
Felix Schürmann, Nenad Buncic, Fabien Delalondre, Stefan Eilemann, and Jeffrey Muller, École Polytechnique Fédérale de Lausanne, Switzerland

11:50-12:10 Runtime Configurability in PETSc
Matthew G. Knepley, University of Chicago, USA; Barry F. Smith, Lois C. McInnes, Jed Brown, and Peter R. Brune, Argonne National Laboratory, USA

Thursday, February 20

MS39

Recent Advances in Parallel Meshing Algorithms - Part I of II

10:35 AM-12:15 PM

Room:Salon H

For Part 2 see MS47

Meshes with millions to billions of elements are now commonly used in today's large-scale scientific codes which simulate various types of physical and biological phenomena. This is mainly due to the increase in computing power with the advent of multicore and many core machines, larger supercomputers, and graphics processing units (GPUs). Thus, parallel meshing techniques are needed in order to generate and manipulate such large meshes. This minisymposium will focus on recent advances in parallel meshing techniques on various types of problems, e.g., mesh generation, mesh quality improvement, adaptive mesh refinement, mesh motion, mesh untangling, and visualization of meshes.

Organizer: Suzanne M. Shontz
Mississippi State University, USA

Organizer: Yongjie Zhang
Carnegie Mellon University, USA

10:35-10:55 Hybrid MPI/openmp Anisotropic Mesh Generation
Gerard J Gorman, Imperial College London, United Kingdom

11:00-11:20 A Parallel Log-Barrier Algorithm for Untangling and Mesh Quality Improvement
Suzanne M. Shontz, Mississippi State University, USA; Shankar Prasad Sastry, University of Utah, USA

11:25-11:45 Multicore CPU or GPU Accelerated Geometry Modeling for Proteins
Yongjie Zhang, Carnegie Mellon University, USA

11:50-12:10 Overview of Parallel Mesh Generation and Optimizations Methods
Andrey Chernikov, Old Dominion University, USA; Suzanne M. Shontz, Mississippi State University, USA; Nikos P. Chrisochoides, College of William & Mary, USA

Thursday, February 20

MS40

Scalable Algorithms for Computational Statistical Inference - Part II of II

10:35 AM-12:15 PM

Room:Salon I

For Part 1 see MS25

High-performance algorithms for statistical inference are a required technology for enabling analysis of large datasets. By "large," we refer to datasets that are in the same order of magnitude as the total RAM of current leadership architectures, which currently can reach and exceed one petabyte. Despite the many mathematical breakthroughs in algorithms for statistical inference, there remain several challenges in scaling this algorithms to large datasets. In this minisymposium, we will discuss recent advances in scalable methods for statistical inference algorithms.

Organizer: George Biros
University of Texas at Austin, USA

10:35-10:55 Scalable Algorithms for Non-Negative Matrix Factorization
Haesun Park, Georgia Institute of Technology, USA

11:00-11:20 Sparse Inverse Covariance Estimation for a Million Variables
Inderjit S. Dhillon, University of Texas at Austin, USA

11:25-11:45 Parallel Algorithms for Sparse Grids
Christoph Kowitz, Technische Universität München, Germany

11:50-12:10 Parallel Algorithms for Prior Functions in Bayesian Inference
George Biros, University of Texas at Austin, USA

Lunch Break

12:15 PM-1:45 PM

Attendees on their own

Thursday, February 20

IP5

Large-scale GPU Applications for Scientific Computing

1:45 PM-2:30 PM

Room:Salon F

Chair: Rob Schreiber, Hewlett Packard Corporation, USA

GPU (Graphics Processing Unit) has been widely used in science and engineering and it has both high computational performance and wide memory bandwidth. On the whole TSUBAME system equipped with 4,224 GPUs and 5.7 PFLOPS of the peak performance at the Tokyo Institute of Technology, we carried out a meso-scale weather model with 500-m horizontal resolution, an air flow simulation of a central part of metropolitan Tokyo for 10 km x 10 km area with 1-m resolution, a phase-field simulation for the dendritic solidification of a binary alloy with 0.3 trillion cells and a granular simulation using 0.1 billion particles.

Takayuki Aoki
Tokyo Institute of Technology, Japan

Thursday, February 20

MS41

Parallel Methods and Algorithms for Extreme Computing - Part II of III

2:40 PM-4:20 PM

Room:Salon F

For Part 1 see MS33
For Part 3 see MS49

Multicore processors and accelerators are universally available as both collections of homogeneous standard microprocessors and as attached heterogeneous co-processors. Application and library software developers may often effectively use these processors and some general approaches have emerged. It is widely recognized that careful design of software and data structures, with effective memory management, are the most critical issues to obtain scalable optimized performance on those systems. In these minisymposia we discuss current experiences and development of applications, libraries and frameworks using a variety of hardware. Speakers will address performance results and software design.

Organizer: Serge G. Petiton
CNRS/LIFL and INRIA, France

Organizer: Michael A. Heroux
Sandia National Laboratories, USA

Organizer: Kengo Nakajima
University of Tokyo, Japan

2:40-3:00 Optimization of Communications/Synchronizations for Preconditioned Iterative Linear Solvers
Kengo Nakajima, University of Tokyo, Japan

continued in next column

3:05-3:25 Probabilistic Approaches for Fault-Tolerance and Scalability in Extreme-Scale Computing
Bert J. Debusschere and Khachik Sargsyan, Sandia National Laboratories, USA; Francesco Rizzi, Duke University, USA; Cosmin Safta and Karla Morris, Sandia National Laboratories, USA; Omar M. Knio, Duke University, USA; Habib N. Najm, Sandia National Laboratories, USA

3:30-3:50 Parallel H-Matrices with Adaptive Cross Approximation for Large-Scale Simulation
Akihiro Ida and Takeshi Iwashita, Kyoto University, Japan

3:55-4:15 What Krylov Basis Computation for GMRES on Cluster of Accelerators
Langshi Chen, CNRS, France; Serge G. Petiton, CNRS/LIFL and INRIA, France; Leroy A. Drummond, Lawrence Berkeley National Laboratory, USA; Maxime Hugues, LIFL, France

Thursday, February 20

MS42

Auto-tuning Technologies for Extreme-Scale Solvers - Part II of III

2:40 PM-4:20 PM

Room:Salon A

For Part 1 see MS34
For Part 3 see MS50

Numerical solvers that accommodate extreme levels of parallelism will be essential for realizing exa-flops simulations. The development of such solvers will require not only algorithmic innovations but also significant effort for their implementations. Noticeably, auto-tuning (AT) technology has provided sustained high performance on various computer architectures, and it is expected to also play a role in the exa-flops era, especially in the context of solvers. In this minisymposium, we will discuss algorithms and implementations for extreme-scale numerical solvers with regards to AT. We target eigenvalue solvers, iterative linear equations solvers, and fundamental matrix decompositions such as the QR decomposition.

Organizer: Toshiyuki Imamura
RIKEN, Japan

Organizer: Takahiro Katagiri
University of Tokyo, Japan

Organizer: Osni A. Marques
Lawrence Berkeley National Laboratory, USA

Organizer: Leroy A. Drummond
Lawrence Berkeley National Laboratory, USA

2:40-3:00 Communication Avoiding-hiding and Auto-tuning for Exteme-scale Eigensolver
Toshiyuki Imamura and Yusuke Hirota, RIKEN, Japan

continued in next column

3:05-3:25 A Performance Model based Approach to Auto-tuning Tall and Skinny QR Factorizations
Takeshi Fukaya and Yusaku Yamamoto, Kobe University, Japan

3:30-3:50 Spiral on (K)
Franz Frachetti, Carnegie Mellon University, USA

3:55-4:15 BCBCG: Iterative Solver with Less Number of Global Communications
Cong Li, Reiji Suda, Kohei Shimane, and Hongzhi Chen, University of Tokyo, Japan

Thursday, February 20

MS43

Hardware/Software Know-how for Effective Parallel Computing

2:40 PM-4:20 PM

Room:Salon B

There are some persistent challenges/ questions in HPC. How to analyze performance through memory for a given parallel architecture in order to use up the memory efficiently for a given application? How to design and program application-specific multicore systems? How to create and design a software infrastructure with power and energy models integrated? Last but not the least, how to solve computational issues related to applications involving depth imaging and visualization? The speakers will throw light upon these challenges and discuss some of the suitable solutions along with their recent research results.

Organizer: Sunita Chandrasekaran
University of Houston, USA

Organizer: Barbara Chapman
University of Houston, USA

2:40-3:00 Architectural Performance Analysis of Emerging Workloads
Matthew E. Tolentino, Intel Corporation, USA

3:05-3:25 Design and Programming of Application-Specific Multi-Core Architectures
Diana Goehringer, Ruhr-Universitat Bochum, Germany

3:30-3:50 Integrating Power and Energy Models in Optimization Tools
Kirk Cameron, Virginia Tech, USA

3:55-4:15 HPC from Cluster to Accelerator-level
Mauricio Araya-Polo, Amik St-Cyr, Ligang Lu, and Detlef Hohl, Shell International Exploration and Production, USA

Thursday, February 20

MS44

Electronic Structure Calculations: Parallel Algorithms and Applications Part II

2:40 PM-4:20 PM

Room:Salon C

For Part 1 see MS36

Electronic structure calculations and their applications are among the most challenging and computationally demanding science and engineering problems. This minisymposium aims at presenting and discussing new numerical and parallel processing avenues that are suitable for modern computing architectures, for achieving ever higher level of accuracy and scalability in DFT, TDDFT and other types of ground and excited states simulations. We propose to bring together physicists/chemists who are involved in improving the numerical development of widely known quantum chemistry and solid-state physics application software packages, with mathematicians/computer scientists who are focusing on advancing the required state-of-the-art mathematical algorithms and parallel implementation.

Organizer: Chao Yang
Lawrence Berkeley National Laboratory, USA

Organizer: Eric Polizzi
University of Massachusetts, Amherst, USA

2:40-3:00 A Parallel Orbital-Updating Approach for Electronic Structure Calculations
Aihui Zhou, Chinese Academy of Sciences, China

3:05-3:25 Parallel Pexsi for Electronic Structure Calculations
Mathias Jacquelin, Lawrence Berkeley National Laboratory, USA

3:30-3:50 Density of States and Eigenvalue Counts via Approximation Theory Methods
Yousef Saad, University of Minnesota, USA

3:55-4:15 Parallel Electron Transport Calculation Using Green's Function Methodology
Olaf Schenk, Universita' della Svizzera Italiana, Italy; Eric Polizzi, University of Massachusetts, Amherst, USA

Thursday, February 20

MS45

Toward Multilevel Solvers for Exascale - Part II of II

2:40 PM-4:20 PM

Room:Salon D

For Part 1 see MS37

The discretization of PDEs arising from applications such as flow problems, chemistry, or material science demands the fast solution of large systems of equations. The minisymposium will highlight the algorithmic design of robust highly parallel multilevel solvers. Here, the nature of the considered equations enforces a global transport of information which is getting increasingly challenging with increasing parallelism. Further, non-linearities have to be properly treated and time-dependencies may be exploited for parallelism. The talks will present optimization and parallelization approaches in order to construct the next generation of efficient solvers.

Organizer: Matthias Bolten
University of Wuppertal, Germany

Organizer: Björn Gmeiner
Universität Erlangen, Germany

2:40-3:00 Multigrid for Structured Grids on 100.000s of Cores
Matthias Bolten, University of Wuppertal, Germany

3:05-3:25 Extending Strong-Scaling Limits with Parallel Integration in Time
Robert Speck, Jülich Supercomputing Centre, Germany; Daniel Ruprecht, University of Lugano, Switzerland

3:30-3:50 Parallel Time Integration with Multigrid
Robert Falgout, Lawrence Livermore National Laboratory, USA

3:55-4:15 Inherently Nonlinear Domain Decomposition and Multigrid Methods for Strongly Nonlinear Problems
Rolf Krause, University of Lugano, Switzerland

Thursday, February 20

MS46

Software Productivity for the Next Generation of Scientific Applications - Part II of II

2:40 PM-4:20 PM

Room:Salon G

For Part 1 see MS38

While extreme-scale architectures provide unprecedented resources for scientific discovery, the community faces daunting productivity challenges for parallel application development. Difficulties include increasing complexity of algorithms and computer science techniques required in multiscale and multiphysics applications, the imperative of portable performance in the midst of dramatic and disruptive architectural changes, the realities of large legacy code bases, and human factors arising in distributed multidisciplinary research teams pursuing extreme parallel performance. Speakers in this minisymposium will discuss promising approaches underway in software development tools, libraries, and methodologies to address these productivity challenges in extreme-scale parallel scientific software.

Organizer: Lois C. McInnes
Argonne National Laboratory, USA

Organizer: Michael A. Heroux
Sandia National Laboratories, USA

Organizer: Hans Johansen
Lawrence Berkeley National Laboratory, USA

2:40-3:00 Software Lifecycle Models in Trilinos
Michael A. Heroux, Sandia National Laboratories, USA

3:05-3:25 What We Have Learned About Using Software Engineering Practices in Computational Science and Engineering
Jeffrey C. Carver, University of Alabama, USA

3:30-3:50 Expressive Environments and Code Generation for High Performance Computing
Garth Wells, University of Cambridge, United Kingdom

3:55-4:15 Tools for Change
David E. Bernholdt, Oak Ridge National Laboratory, USA

Thursday, February 20

MS47

Recent Advances in Parallel Meshing Algorithms - Part II of II

2:40 PM-4:20 PM

Room:Salon H

For Part 1 see MS39

Meshes with millions to billions of elements are now commonly used in today's large-scale scientific codes which simulate various types of physical and biological phenomena. This is mainly due to the increase in computing power with the advent of multicore and many core machines, larger supercomputers, and graphics processing units (GPUs). Thus, parallel meshing techniques are needed in order to generate and manipulate such large meshes. This minisymposium will focus on recent advances in parallel meshing techniques on various types of problems, e.g., mesh generation, mesh quality improvement, adaptive mesh refinement, mesh motion, mesh untangling, and visualization of meshes.

Organizer: Suzanne M. Shontz
Mississippi State University, USA

Organizer: Yongjie Zhang
Carnegie Mellon University, USA

2:40-3:00 Scalable Lattice Cleaving
Jonathan Bronson, University of Utah, USA; Jonathon Nooner and *Joshua Levine*, Clemson University, USA; Tom Peterka, Argonne National Laboratory, USA; Ross Whitaker, University of Utah, USA

3:05-3:25 Parallel Algorithms for Overlay Grid Methods
Steve J. Owen, Sandia National Laboratories, USA

3:30-3:50 Parallel Mesh Generation and Adaptation with Distributed Geometry
Mark Beall and Saurabh Tendulkar, Simmetrix, Inc., USA; *Mark S. Shephard*, Rensselaer Polytechnic Institute, USA

3:55-4:15 Advances in Parallel Unstructured Mesh Adaptation
Onkar Sahni and Daniel Ibanez, Rensselaer Polytechnic Institute, USA; Kedar Chitale, University of Colorado Boulder, USA; Mark S. Shephard, Rensselaer Polytechnic Institute, USA

Thursday, February 20

MS48

Graph Analysis for Scientific Discovery - Part I of III

2:40 PM-4:20 PM

Room:Salon I

For Part 2 see MS56

As data sets continue to explode exponentially, they urgently require sophisticated analysis methods to extract key information in a computationally tractable way. Graph theory is a powerful abstraction that is increasingly driving big-data analysis and is making an impact on a broad range of fields from social media, to cyber-security, to traditional scientific computing realms. In this minisymposium, we will explore the latest algorithmic approaches, optimization techniques, and computational domains that will drive the next generation of graph-based analysis of scientific dataset.

Organizer: Leonid Oliker
Lawrence Berkeley National Laboratory, USA

Organizer: Aydin Buluc
Lawrence Berkeley National Laboratory, USA

Organizer: John R. Gilbert
University of California, Santa Barbara, USA

2:40-3:00 High-performance and High-productivity Semantic Graph Analysis
Aydin Buluc, Lawrence Berkeley National Laboratory, USA; John R. Gilbert and Adam Lugowski, University of California, Santa Barbara, USA; Shoaib Kamil, Massachusetts Institute of Technology, USA; Leonid Oliker and Samuel Williams, Lawrence Berkeley National Laboratory, USA

continued in next column

3:05-3:25 Algorithms for Aligning Massive Networks
Alex Pothen, Arif Khan, and David F. Gleich, Purdue University, USA; Alessandro Morari, Mahantesh Halappanavar, Daniel Chavarria, and Oreste Villa, Pacific Northwest National Laboratory, USA

3:30-3:50 Managing Centrality in Large Scale Graphs
Ahmet Erdem Sariyuce, Erik Saule, Kamer Kaya, and *Umit V. Catalyurek*, The Ohio State University, USA

3:55-4:15 A Scalable Querying System for Large-Scale Graph Pattern Enumeration
Steve Harenberg, Rob Seay, Sriram Lakshminarasimhan, David Boyuka Ii, Gonzalo bello, Rada chirkova, and Nagiza Samatova, North Carolina State University, USA

Coffee Break

4:20 PM-4:50 PM

Room:Salon E

Thursday, February 20

MS49

Parallel Methods and Algorithms for Extreme Computing - Part III of III

4:50 PM-6:30 PM

Room:Salon F

For Part 2 see MS41

Multicore processors and accelerators are universally available as both collections of homogeneous standard microprocessors and as attached heterogeneous co-processors. Application and library software developers may often effectively use these processors and some general approaches have emerged. It is widely recognized that careful design of software and data structures, with effective memory management, are the most critical issues to obtain scalable optimized performance on those systems. In these minisymposia we discuss current experiences and development of applications, libraries and frameworks using a variety of hardware. Speakers will address performance results and software design.

Organizer: Serge G. Petiton
CNRS/LIFL and INRIA, France

Organizer: Michael A. Heroux
Sandia National Laboratories, USA

Organizer: Kengo Nakajima
University of Tokyo, Japan

4:50-5:10 Toward Intelligent Krylov-Based Linear Algebra Methods for Future Extreme Programming and Computing
Serge G. Petiton, CNRS/LIFL and INRIA, France

5:15-5:35 High-Performance Gpu Kernels for Multifrontal Sparse Factorization
Timothy A. Davis, Sanjay Ranka, and Nuri Yeralan, University of Florida, USA

5:40-6:00 Accelerating CFD Applications Without Accelerator-Specific Programming
Naoya Maruyama, RIKEN, Japan

6:05-6:25 Re-Architecting DFT Kernels for Sustainable Performance
Osni A. Marques, Andrew M. Canning, and David Prendergast, Lawrence Berkeley National Laboratory, USA; Jack Deslippe, National Energy Research Scientific Computing Center, USA

Thursday, February 20

MS50

Auto-tuning Technologies for Extreme-Scale Solvers - Part III of III

4:50 PM-6:30 PM

Room:Salon A

For Part 2 see MS42

Numerical solvers that accommodate extreme levels of parallelism will be essential for realizing exa-flops simulations. The development of such solvers will require not only algorithmic innovations but also significant effort for their implementations. Noticeably, auto-tuning (AT) technology has provided sustained high performance on various computer architectures, and it is expected to also play a role in the exa-flops era, especially in the context of solvers. In this minisymposium, we will discuss algorithms and implementations for extreme-scale numerical solvers with regards to AT. We target eigenvalue solvers, iterative linear equations solvers, and fundamental matrix decompositions such as the QR decomposition.

Organizer: Osni A. Marques
Lawrence Berkeley National Laboratory, USA

Organizer: Leroy A. Drummond
Lawrence Berkeley National Laboratory, USA

Organizer: Takahiro Katagiri
University of Tokyo, Japan

Organizer: Toshiyuki Imamura
RIKEN, Japan

4:50-5:10 The Impact of Workload Strategies on a UPC-based CG
Jorge González-Domínguez, University of A Coruña, Spain; Osni A. Marques, Lawrence Berkeley National Laboratory, USA; María Martín and Juan Touriño, University of A Coruña, Spain

continued in next column

5:15-5:35 Dynamic Parallel Algebraic Multigrid Coarsening for Strong Scaling
Akihiro Fujii, Takuya Nomura, and Teruo Tanaka, Kogakuin University, Japan; Osni A. Marques, Lawrence Berkeley National Laboratory, USA

5:40-6:00 Parameter Selection and Prediction to Tune the Performance of Krylov Subspace Methods
France Boillod-Cerneux, CNRS/LIFL, France; *Leroy A. Drummond*, Lawrence Berkeley National Laboratory, USA; Herve Galicher, CEA, France; Christophe Calvin, CEA Saclay, France; Serge G. Petiton, CNRS/LIFL and INRIA, France

6:05-6:25 Tuning Asynchronous Co-Methods for Large-scale Eigenvalue Calculations
Nahid Emad, University of Versailles, France; *Leroy A. Drummond*, Lawrence Berkeley National Laboratory, USA; Miwako Tsuji and Makarem Dandouna, University of Versailles, France

Thursday, February 20

MS51

DOE Computational Science Graduate Fellowship Program Showcase: Design and Application of Massively Parallel Codes - Part I of II

4:50 PM-6:30 PM

Room:Salon B

For Part 2 see MS59

Current fellows and alumni of the DOE-CSGF program will present their work in scaling complex application codes to massively parallel computers, including at least three codes that run on more than 1M cores. Best practices for end-to-end optimization of computational science workflows is demonstrate in multiple disciplines.

Organizer: Jeff R. Hammond
Argonne National Laboratory, USA

Organizer: Jeffrey A. Hittinger
Lawrence Livermore National Laboratory, USA

4:50-5:10 Design of HACC for Extreme-Scale Simulation

Hal Finkel, Salman Habib, Vitali Morozov, Adrian Pope, Katrin Heitmann, Kalyan Kumaran, Tom Peterka, and Joe Insley, Argonne National Laboratory, USA; David Daniel and Patricia Fasel, Los Alamos National Laboratory, USA; Nicholas Frontiere, Argonne National Laboratory and University of Chicago, USA; Zarija Lukic, Lawrence Berkeley National Laboratory, USA

5:15-5:35 Scalability and Performance of a Legacy Coastal Ocean Model

Jesse Lopez, Oregon Health & Science University, USA; Jed Brown, Argonne National Laboratory, USA

5:40-6:00 Scalability of Sweep Algorithms in Discrete Ordinates Transport

Teresa S. Bailey, Lawrence Livermore National Laboratory, USA

6:05-6:25 Scalable Adaptive Mesh Refinement

Brian Gunney, Lawrence Livermore National Laboratory, USA

Thursday, February 20

MS52

Electronic Structure Calculations: Parallel Algorithms and Applications Part III

4:50 PM-6:30 PM

Room:Salon C

Electronic structure calculations and their applications are among the most challenging and computationally demanding science and engineering problems. This minisymposium aims at presenting and discussing new numerical and parallel processing avenues that are suitable for modern computing architectures, for achieving ever higher level of accuracy and scalability in DFT, TDDFT and other types of ground and excited states simulations. We propose to bring together physicists/chemists who are involved in improving the numerical development of widely known quantum chemistry and solid-state physics application software packages, with mathematicians/computer scientists who are focusing on advancing the required state-of-the-art mathematical algorithms and parallel implementation.

Organizer: Chao Yang
Lawrence Berkeley National Laboratory, USA

Organizer: Eric Polizzi
University of Massachusetts, Amherst, USA

4:50-5:10 Real-Space DFT for Plane-Wave GW/BSE Calculations

David A. Strubbe, Massachusetts Institute of Technology, USA

5:15-5:35 Real-space Electronic Structure on Graphics Processing Units

Xavier Andrade, Harvard University, USA

5:40-6:00 Improved Iterative Subspace Methods for Large-Scale Response Calculations

Filipp Furche, Brandon Krull, and Jake Kwon, University of California, Irvine, USA

6:05-6:25 FEAST Applied to DFT and Real-time TDDFT Calculations

Eric Polizzi, University of Massachusetts, Amherst, USA

Thursday, February 20

MS53

Sparse Computations on Accelerators

4:50 PM-6:30 PM

Room:Salon D

Sparse matrix and other irregular computations are challenging to implement on hardware such as GPUs and Intel MIC that demand high levels of both thread and data parallelism. In this minisymposium, the speakers will present novel solutions for sparse computations on accelerators, including solutions applicable to both GPUs and MIC. Kernels explored include sparse matrix- vector multiplication, preconditioning operations including sparse triangular solves, and irregular computations arising from particle simulations.

Organizer: Edmond Chow
Georgia Institute of Technology, USA

4:50-5:10 Fine-Grained Parallel Preconditioning

Edmond Chow, Georgia Institute of Technology, USA

5:15-5:35 Sparse Matrix-Vector Multiplication with Wide SIMD Units: Performance Models and a Unified Storage Format

Moritz Kreutzer, *Georg Hager*, and Gerhard Wellein, Erlangen Regional Computing Center, Germany

5:40-6:00 Rethinking Multigrid and Sparse Matrix Computations in High Throughput Environments

Luke Olson, University of Illinois at Urbana-Champaign, USA; Steven Dalton, University of Illinois at Urbana-Champaign, USA

6:05-6:25 Accelerating Brownian Dynamics Simulations on Intel MIC

Xing Liu and Edmond Chow, Georgia Institute of Technology, USA

Thursday, February 20

MS54

Frameworks, Algorithms and Scalable Technologies for Mathematics on Next-generation Computers - Part I of II

4:50 PM-6:30 PM

Room:Salon G

For Part 2 see MS62

This minisymposium series focuses on algorithms and software developed by the FASTMath SciDAC team to improve the reliability and robustness of application codes. We describe advances in the scalable implementation of structured and unstructured mesh techniques, as well as linear and nonlinear solvers. A pervasive theme in our work is understanding the most effective ways to implement our algorithms efficiently and at scale on many-core architectures with million-way parallelism. This minisymposium discusses lessons learned in using hybrid parallelism models, the impact of new algorithms on communication costs and scalability, and scalability up to 750K cores.

Organizer: Lori A. Diachin
Lawrence Livermore National Laboratory, USA

4:50-5:10 An Overview of FASTMath Technology Developments
Lori A. Diachin, Lawrence Livermore National Laboratory, USA

5:15-5:35 Chombo-Crunch: High Performance Simulation of Pore Scale Reactive Transport Processes Associated with Carbon Sequestration
David Trebotich, Mark Adams, and Brian Van Straalen, Lawrence Berkeley National Laboratory, USA; Anshu Dubey, University of Chicago, USA; Terry J. Ligocki and Dan Graves, Lawrence Berkeley National Laboratory, USA

continued in next column

5:40-6:00 New Developments in Parallel Unstructured Mesh Methods
Daniel Ibanez, Seegyoung Seol, Cameron Smith, Ian Dunn, and Mark S. Shephard, Rensselaer Polytechnic Institute, USA

6:05-6:25 PHASTA Unstructured Mesh Scalability to 3.1M Processes
Michel Rasquin, Argonne National Laboratory, USA; Cameron Smith, Rensselaer Polytechnic Institute, USA; Kedar Chitale, University of Colorado Boulder, USA; Seegyoung Seol and Benjamin Matthews, Rensselaer Polytechnic Institute, USA; Jeffrey Martin, University of Colorado Boulder, USA; Onkar Sahni and Mark S. Shephard, Rensselaer Polytechnic Institute, USA; Ray Loy, Argonne National Laboratory, USA; *Kenneth Jansen*, University of Colorado Boulder, USA

Thursday, February 20

MS55

GPU and Manycore Heterogeneous Computing for Stencil Applications

4:50 PM-6:30 PM

Room:Salon H

Stencil applications such as Computational Fluid Dynamics, structural analysis and seismic wave propagation are successful applications on GPU and manycore heterogeneous computing supplying high computational performance and high memory bandwidth in HPC area. The application performances in large-scale computation depend on compute intensity, memory access and internode communication. The improvement of computational performance is much faster than memory and interconnection. From the viewpoint of future supercomputing, we would like to discuss the approach to increase compute intensity such as temporal cache blocking, overlapping communication with computation and switching to new numerical algorithms suitable for future processor architectures.

Organizer: Takayuki Aoki
Tokyo Institute of Technology, Japan

4:50-5:10 Parallel Computing of Finite-Volume Solver for Euler Equation using a Ghost-cell Immersed Boundary Method using Multiple Graphics Processor Units
Chongsin Gou and *Jong-Shinn Wu*, National Chiao Tung University, Taiwan

5:15-5:35 High Scalability of Lattice Boltzmann Simulations with Turbulence Models Using Heterogeneous Clusters Equipped with GPUs
Christoph Riesinger, Technische Universität München, Germany

5:40-6:00 Physics-based Seismic Hazard Analysis on Petascale Heterogeneous Supercomputers
Heming Xu, San Diego Supercomputer Center, USA

6:05-6:25 The waLBerla Framework: Multi-physics Simulations on Heterogeneous Platforms
Ulrich J. Ruede and Harald Koestler, University of Erlangen-Nuremberg, Germany

Thursday, February 20

MS56

Graph Analysis for Scientific Discovery - Part II of III

4:50 PM-6:30 PM

Room:Salon I

For Part 1 see MS48
For Part 3 see MS64

As data sets continue to explode exponentially, they urgently require sophisticated analysis methods to extract key information in a computationally tractable way. Graph theory is a powerful abstraction that is increasingly driving big-data analysis and is making an impact on a broad range of fields from social media, to cyber-security, to traditional scientific computing realms. In this minisymposium, we will explore the latest algorithmic approaches, optimization techniques, and computational domains that will drive the next generation of graph-based analysis of scientific dataset.

Organizer: Leonid Oliker
Lawrence Berkeley National Laboratory, USA

Organizer: Aydin Buluc
Lawrence Berkeley National Laboratory, USA

Organizer: John R. Gilbert
University of California, Santa Barbara, USA

4:50-5:10 Characterizing Biological Networks Using Subgraph Counting and Enumeration
George Slota and *Kamesh Madduri*, Pennsylvania State University, USA

5:15-5:35 Connectome Coding Via Latent Position Estimation
Carey Priebe, Johns Hopkins University, USA; Joshua Vogelstein, Duke University, USA; Jacob Vogelstein, Johns Hopkins University, USA

continued in next column

5:40-6:00 Fast Clustering Methods for Genetic Mapping in Plants
Veronika Strnadova, University of California, Santa Barbara, USA; Aydin Buluc and Leonid Oliker, Lawrence Berkeley National Laboratory, USA; Joseph Gonzalez and Stefanie Jegelka, University of California, Berkeley, USA; Jarrod Chapman, DOE Joint Genome Institute, USA; John R. Gilbert, University of California, Santa Barbara, USA

6:05-6:25 Large-Scale Metagenomic Sequence Clustering Via Maximal Quasi-Clique Enumeration
Srinivas Aluru, Georgia Institute of Technology, USA

SIAG/SC Business Meeting

6:30 PM-7:15 PM

Room:Salon F

Complimentary wine and beer will be served.

Friday, February 21

Registration

7:45 AM-5:15 PM

Room:Oregon Ballroom Foyer

IP6

Co-Design of Exascale Simulation of Combustion in Turbulence (ExaCT)

8:15 AM-9:00 AM

Room:Salon F

Chair: Ali Pinar, Sandia National Laboratories, USA

Exascale computing will enable direct numerical simulation (DNS) of turbulent combustion science at engine relevant thermo-chemical conditions. These DNS will be used to develop predictive models that couple chemistry with turbulent transport ultimately used to design fuel efficient, clean engines and gas turbines utilizing alternative fuels including biofuels. The mission of co-design within ExaCT is to have technology capabilities and constraints inform the formulation of new algorithms and software, and for combustion requirements to guide computer architecture and systems software design. ExaCT iteratively co-designs all aspects of the stack affecting combustion simulation including algorithms, domain-specific programming environments, scientific data management and analytics for in situ uncertainty quantification and architectural modeling and simulation to explore hardware tradeoffs with combustion proxy applications.

Jackie Chen
Sandia National Laboratories, USA

Intermission

9:00 AM-9:10 AM

Friday, February 21

SIAG Prize Lectures (Junior and Career Prizes)

9:10 AM-10:10 AM

Room:Salon F

To Be Announced

Coffee Break

10:10 AM-10:35 AM

Room:Salon E

Friday, February 21

MS57

Asynchronous Adaptive Runtimes for Exascale HPC

10:35 AM-12:15 PM

Room:Salon F

For Part 2 see MS65

Exascale computing will require managing increased complexity in both hardware and software systems. Asynchronous runtimes allow programs to deal with dynamic execution environments that arise under these conditions. Sources of variability include algorithmic sources such as dynamic data dependencies (e.g. AMR), and hardware sources such as thermal throttling. Responding effectively to this variability requires increased flexibility in when and where computations and data are executed, stored, or moved. Topics addressed in this symposium will include: applications of asynchronous runtimes; programming and execution models; task scheduling and data management; trade-offs in runtime overhead, task granularity, and data locality; fault-tolerance; and introspection.

Organizer: Cy Chan
Lawrence Berkeley National Laboratory, USA

10:35-10:55 Runtime Requirements for Scalable Semantic Data Analysis
John T. Feo, Pacific Northwest National Laboratory, USA

11:00-11:20 The Charm++ Applications Experience: Production Use of an Asynchronous, Adaptive Runtime
Laxmikant Kale, University of Illinois at Urbana-Champaign, USA

11:25-11:45 An Open Community Runtime (OCR) for Exascale Systems
Wilfred R. Pinfold, Intel Labs, USA

11:50-12:10 Habanero Execution Model and the Open Community Runtime Project
Vivek Sarkar, Rice University, USA

Friday, February 21

MS58

Optimizing Stencil-based Algorithms - Part I of II

10:35 AM-12:15 PM

Room:Salon A

For Part 2 see MS66

Stencil or stencil-like algorithms are the core of many numerical solvers and simulation codes. There is vast literature on parallelizing and optimizing stencil codes on modern computer architectures, and work is ongoing in many directions. Hardware features like wide SIMD parallelism, (massive) threading, multi-level caches, and increasing core counts complicate matters and fuel the trend towards software abstractions and automatic tuning frameworks. We bring together experts who provide a comprehensive overview of the state of the art and ongoing work. Various approaches, from domain-specific languages to performance models, and from auto-tuning to hardware-specific optimizations, will be covered.

Organizer: David E. Keyes
King Abdullah University of Science & Technology (KAUST), Saudi Arabia

Organizer: Jan Treibig
Erlangen Regional Computing Center, Germany

Organizer: Georg Hager
Erlangen Regional Computing Center, Germany

Organizer: Gerhard Wellein
Erlangen Regional Computing Center, Germany

10:35-10:55 Relevant Stencil Structures for Modern Numerics
David E. Keyes, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

continued on next page

11:00-11:20 Performance Engineering for Stencil Updates on Modern Processors
Jan Treibig, Georg Hager, and *Gerhard Wellein*, Erlangen Regional Computing Center, Germany

11:25-11:45 Compiler-Automated Communication-Avoiding Optimization of Geometric Multigrid
Protonu Basu, University of Utah, USA; Samuel Williams and Brian Van Straalen, Lawrence Berkeley National Laboratory, USA; Anand Venkat, University of Utah, USA; Leonid Oliker, Lawrence Berkeley National Laboratory, USA; Mary Hall, University of Utah, USA

11:50-12:10 Automatic Generation of Algorithms and Data Structures for Geometric Multigrid
Harald Koestler and *Sebastian Kuckuk*, Universität Erlangen-Nürnberg, Germany

Friday, February 21

MS59

DOE Computational Science Graduate Fellowship Program Showcase: Design and Application of Massively Parallel Codes - Part II of II

10:35 AM-12:15 PM

Room:Salon B

For Part 1 see MS51
Current fellows and alumni of the DOE-CSGF program will present their work in scaling complex application codes to massively parallel computers, including at least three codes that run on more than 1M cores. Best practices for end-to-end optimization of computational science workflows is demonstrate in multiple disciplines.

Organizer: Jeff R. Hammond
Argonne National Laboratory, USA

Organizer: Jeffrey A. Hittinger
Lawrence Livermore National Laboratory, USA

10:35-10:55 Multi-Scale Parallelism in Yt: Lessons from a Community Driven Analysis Package
Samuel W. Skillman, Stanford University, USA

11:00-11:20 Scaling the Lattice Boltzmann Method on Massively Parallel Systems
Amanda Randles and Erik W. Draeger, Lawrence Livermore National Laboratory, USA; Franziska Michor, Harvard University, USA

11:25-11:45 Evolution and Revolution in Massively Parallel Quantum Chemistry Codes
Jeff R. Hammond, Argonne National Laboratory, USA

11:50-12:10 Parallel Algorithms for Quantum Many-body Methods: Lessons Learned from NWChem
David Ozog, University of Oregon, USA; Jeff R. Hammond and Pavan Balaji, Argonne National Laboratory, USA; James Dinan, Ohio State University, USA; Allen Malony and Sameer Shende, University of Oregon, USA

Friday, February 21

MS60

Towards Scalable Dependency-aware Task-based Programming Models for Distributed Memory Systems

10:35 AM-12:15 PM

Room:Salon C

Dependency-aware task-based programming models have been successful for extracting performance from multicore-based shared memory computer systems. Basically, the computational problem is divided into tasks, dependencies are inferred through high-level, user supplied annotations, and a run-time system schedules the tasks onto available cores. High-performance computer systems typically have shared memory nodes, which may be equipped with accelerators, combined into an overall distributed memory system. This minisymposium asks the question if task-based programming models can be efficient also in this case. Aspects such as distributed dependency tracking, scheduling, and load balancing in current solutions will be discussed in relation to scalability requirements.

Organizer: Elisabeth Larsson
Uppsala University, Sweden

Organizer: Afshin Zafari
Uppsala University, Sweden

Organizer: Martin Tillenius
Uppsala University, Sweden

10:35-10:55 DuctTeip: A Task-Based Parallel Programming Framework with Modularity, Scalability and Adaptability Features
Afshin Zafari, Martin Tillenius, and Elisabeth Larsson, Uppsala University, Sweden

11:00-11:20 A Flexible Programming Environment for Distributed Heterogeneous Resources
George Bosilca, University of Tennessee, Knoxville, USA

11:25-11:45 OmpSs for Modern Cluster Architectures
Javier Bueno, Barcelona Supercomputing Center, Spain

11:50-12:10 Starpu-MPI: Extending Task Graphs from Heterogeneous Platforms to Clusters Thereof
Samuel Thibault, University of Bordeaux, France

Friday, February 21

MS61

Preconditioning Techniques for Sparse Linear Systems on GPUs

10:35 AM-12:15 PM

Room:Salon D

Graphical Processing Units represent a new frontier for scientific computing and their use as accelerators is becoming increasingly popular. However, GPU hardware model is quite different from classical CPUs, and may require an ad hoc algorithmic re-formulation for better exploiting the GPU potential. This includes in particular existing preconditioning techniques for sparse linear systems. Most existing algorithms can be hardly implemented efficiently on GPUs in their native formulations. This minisymposium aims at bringing together experts working on the parallel implementation of modern preconditioning techniques on GPU systems, with the presentation of the most recent algorithmic developments and numerical experiments.

Organizer: Carlo Janna
University of Padova, Italy

Organizer: Massimiliano Ferronato
University of Padova, Italy

10:35-10:55 Factorized Sparse Approximate Inverses on GPUs
Massimo Bernaschi, Istituto per le Applicazioni del Calcolo "Mauro Picone", Italy; Carlo Fantozzi, Massimiliano Ferronato, Giuseppe Gambolati, and *Carlo Janna*, University of Padova, Italy

11:00-11:20 Bone Structure Analysis with Multiple GPGPUs
Peter Arbenz, Cyril Flaig, and Daniel Kellenberger, ETH Zürich, Switzerland

11:25-11:45 Deflated Preconditioned Conjugate Gradient for Bubbly Flows: Multi-GPU/CPU Implementations
Rohit Gupta, Martin B. van Gijzen, and Kees Vuik, Delft University of Technology, Netherlands

11:50-12:10 Preconditioning Techniques for GPU-Accelerated Environments
Ruipeng Li and Yousef Saad, University of Minnesota, USA

Friday, February 21

MS62

Frameworks, Algorithms and Scalable Technologies for Mathematics on Next-generation Computers - Part II of II

10:35 AM-12:15 PM

Room:Salon G

For Part 1 see MS54

This minisymposium series focuses on algorithms and software developed by the FASTMath SciDAC team to improve the reliability and robustness of application codes. We describe advances in the scalable implementation of structured and unstructured mesh techniques, as well as linear and nonlinear solvers. A pervasive theme in our work is understanding the most effective ways to implement our algorithms efficiently and at scale on many-core architectures with million-way parallelism. This minisymposium discusses lessons learned in using hybrid parallelism models, the impact of new algorithms on communication costs and scalability, and scalability up to 750K cores.

Organizer: Lori A. Diachin
Lawrence Livermore National Laboratory, USA

10:35-10:55 Strategies for Reducing Setup Costs in Parallel Algebraic Multigrid
Jonathan J. Hu, Andrey Prokopenko, Ray S. Tuminaro, and Christopher Siefert, Sandia National Laboratories, USA

11:00-11:20 Algorithmic Advances for Algebraic Multigrid with Reduced Communication
Ulrike Meier Yang, Robert Falgout, Jacob Schroder, and Panayot Vassilevski, Lawrence Livermore National Laboratory, USA

continued in next column

11:25-11:45 ARKode: A Library of High Order Implicit/explicit Methods for Multi-rate Problems
Daniel R. Reynolds, Southern Methodist University, USA; Carol S. Woodward, Lawrence Livermore National Laboratory, USA; David Gardner, Southern Methodist University, USA; Alan Hindmarsh, Lawrence Livermore National Laboratory, USA

11:50-12:10 Integration of Albany and Mesh Adaptation for Parallel Applications
Glen Hansen, Sandia National Laboratories, USA; Brian Granzow, E. Seegyoung Seol, and Mark S. Shephard, Rensselaer Polytechnic Institute, USA

Friday, February 21

MS63

Recent Advances in Particle-in-Cell Method on Multicore and GPU Systems

10:35 AM-12:15 PM

Room:Salon H

Particle-in-cell (PIC) method are used in important applications such as modeling hot plasma in a fusion device, in understanding plasma in space physics, and in the design of high-energy accelerators. PIC methods have been shown to scale well on most distributed-memory architectures. We will present and discuss recent efforts to improve the performance and scalability of PIC methods on multicore and GPU systems.

Organizer: Eduardo F. D'Azevedo
Oak Ridge National Laboratory, USA

Organizer: Jan Hesthaven
Brown University, USA

10:35-10:55 Hybrid MPI/OpenMP/ GPU Parallelization of XGC1 Fusion Simulation Code
Patrick H. Worley and Ed D'Azevedo, Oak Ridge National Laboratory, USA; Jianying Lang, Seung-Hoe Ku, and Choong-Seock Chang, Princeton Plasma Physics Laboratory, USA

11:00-11:20 Migrating the GTC-P Gyrokinetic Particle-in-Cell Code to Multicore and Manycore Systems
Bei Wang, Princeton University, USA; Stephane Ethier, Princeton Plasma Physics Laboratory, USA; William Tang, Princeton University, USA; Khaled Z. Ibrahim, Lawrence Berkeley National Laboratory, USA; Kamesh Madduri, Pennsylvania State University, USA; Samuel Williams and Leonid Oliker, Lawrence Berkeley National Laboratory, USA

11:25-11:45 Accelerator Beam Dynamics on Multicore and Gpu Systems
James Amundson and Qiming Lu, Fermi National Accelerator Laboratory, USA

11:50-12:10 Exploiting the Power of Heterogeneous Computing for Kinetic Simulations of Plasmas
Kai Germaschewski and *Narges Ahmadi*, University of New Hampshire, USA; Homayoun Karimabadi, University of California, San Diego, USA

Friday, February 21

MS64

Graph Analysis for Scientific Discovery - Part III of III

10:35 AM-12:15 PM

Room:Salon I

For Part 2 see MS56

As data sets continue to explode exponentially, they urgently require sophisticated analysis methods to extract key information in a computationally tractable way. Graph theory is a powerful abstraction that is increasingly driving big-data analysis and is making an impact on a broad range of fields from social media, to cyber-security, to traditional scientific computing realms. In this minisymposium, we will explore the latest algorithmic approaches, optimization techniques, and computational domains that will drive the next generation of graph-based analysis of scientific dataset.

Organizer: Leonid Oliker
Lawrence Berkeley National Laboratory, USA

Organizer: Aydin Buluc
Lawrence Berkeley National Laboratory, USA

Organizer: John R. Gilbert
University of California, Santa Barbara, USA

10:35-10:55 Streaming Graph Analytics for Real-World Problems
David A. Bader and *James Fairbanks*, Georgia Institute of Technology, USA

11:00-11:20 Tracking Combustion Features Using Topological Methods Involving Large Graphs
Valerio Pascucci, University of Utah, USA; Peer-Timo Bremer, Lawrence Livermore National Laboratory and University of Utah, USA; Attila Gyulassy and Aaditya Landge, University of Utah, USA; Janine Bennet and Jackie Chen, Sandia National Laboratories, USA

11:25-11:45 Scalable Graph Methods for Functional Characterization of Environmental Microbial Communities
Ananth Kalyanaraman, Washington State University, USA

11:50-12:10 Generating Large Graphs for Benchmarking
Ali Pinar, C. Seshadhri, Tamara G. Kolda, and Todd Plantenga, Sandia National Laboratories, USA

Friday, February 21

Lunch Break

12:15 PM-1:45 PM

Attendees on their own

IP7

Large-Scale Visual Data Analysis

1:45 PM-2:30 PM

Room:Salon F

Chair: William D. Gropp, University of Illinois at Urbana-Champaign, USA

Modern high performance computers have speeds measured in petaflops and handle data set sizes measured in terabytes and petabytes. Although these machines offer enormous potential for solving very large-scale realistic computational problems, their effectiveness will hinge upon the ability of human experts to interact with their simulation results and extract useful information. One of the greatest scientific challenges of the 21st century is to effectively understand and make use of the vast amount of information being produced. Visual data analysis will be among our most important tools to understand such large-scale information. In this talk, I will present state-of-the-art visualization techniques, including scalable visualization algorithms and cluster-based methods applied to problems in science, engineering, and medicine.

Christopher Johnson
University of Utah, USA

Friday, February 21

MS65

Asynchronous Adaptive Runtimes for Exascale HPC, Part II

2:40 PM-4:20 PM

Room:Salon F

For Part 1 see MS57

Exascale computing will require managing increased complexity in both hardware and software systems. Asynchronous runtimes allow programs to deal with dynamic execution environments that arise under these conditions. Sources of variability include algorithmic sources such as dynamic data dependencies (e.g. AMR), and hardware sources such as thermal throttling. Responding effectively to this variability requires increased flexibility in when and where computations and data are executed, stored, or moved. Topics addressed in this symposium will include: applications of asynchronous runtimes; programming and execution models; task scheduling and data management; trade-offs in runtime overhead, task granularity, and data locality; fault-tolerance; and introspection.

Organizer: Cy Chan
Lawrence Berkeley National Laboratory, USA

2:40-3:00 ParalleX: Defining the Computer of the Future
Thomas Sterling, Indiana University, USA

3:05-3:25 Distributed Memory Runtimes for Dense Linear Algebra Software
Jakub Kurzak, George Bosilca, and Jack J. Dongarra, University of Tennessee, Knoxville, USA

3:30-3:50 SWARM: Maximizing Compute Efficiency and Scalability Today While Designing for Reliability and Power Efficiency for Tomorrow
Rishi Khan, ETI International, USA

3:55-4:15 Dynamic Runtime Systems for Scalable Graph Processing
Andrew Lumsdaine, Indiana University, USA

Friday, February 21

MS66

Optimizing Stencil-based Algorithms - Part II of II

2:40 PM-4:20 PM

Room:Salon A

For Part 1 see MS58

Stencil or stencil-like algorithms are the core of many numerical solvers and simulation codes. There is vast literature on parallelizing and optimizing stencil codes on modern computer architectures, and work is ongoing in many directions. Hardware features like wide SIMD parallelism, (massive) threading, multi-level caches, and increasing core counts complicate matters and fuel the trend towards software abstractions and automatic tuning frameworks. We bring together experts who provide a comprehensive overview of the state of the art and ongoing work. Various approaches, from domain-specific languages to performance models, and from auto-tuning to hardware-specific optimizations, will be covered.

Organizer: David E. Keyes
King Abdullah University of Science & Technology (KAUST), Saudi Arabia

Organizer: Jan Treibig
Erlangen Regional Computing Center, Germany

Organizer: Georg Hager
Erlangen Regional Computing Center, Germany

Organizer: Gerhard Wellein
Erlangen Regional Computing Center, Germany

2:40-3:00 Stencil Computations: from Academia to Industry
Raul de la Cruz, Mauricio Hanzich, and Jose Maria Cela, Barcelona Supercomputing Center, Spain

3:05-3:25 Evaluating Compiler-driven Parallelization of Stencil Micro-applications on a GPU-enabled Cluster
Dmitry Mikushin and Olaf Schenk, Universita' della Svizzera Italiana, Italy

continued in next column

3:30-3:50 Firedrake: a Multilevel Domain Specific Language Approach to Unstructured Mesh Stencil Computations
Gheorghe-Teodor Bercea, *David Ham*, Paul Kelly, Nicolas Loriant, Fabio Luporini, Lawrence Mitchell, and Florian Rathgeber, Imperial College London, United Kingdom

3:55-4:15 Tuning Sparse and Dense Matrix Operators in SeisSol
Alexander Breuer, Sebastian Rettenberger, and Alexander Heinecke, Technische Universität München, Germany; Christian Pelties, Ludwig-Maximilians-Universität München, Germany; Michael Bader, Technische Universität München, Germany

Friday, February 21

MS67

Applications and Algorithms for Industrial Numerical Reservoir Simulation on Massively Parallel Computers

2:40 PM-4:20 PM

Room:Salon B

The solution of coupled PDEs describing the flow of hydrocarbons and water through structurally complex porous media provides major industrial oil producing entities estimates of oil recovery for different production strategies. The recent acquisition of massively parallel computers by the petroleum industry is driving scientific parallel computing into algorithms that were historically serial. New linear solver technologies, domain partitioning algorithms, and parallel software architectures are being developed for deployment on large-scale parallel clusters with heterogeneous computing elements. This minisymposium presents recent advances, industrial applications and continuing issues in the use of massively-parallel machines for reservoir simulation.

Organizer: Ali H. Dogru
Saudi Aramco Oil Company, Saudi Arabia

Organizer: Bret Beckner
ExxonMobil Upstream Research Company, USA

2:40-3:00 Strong Scalability of Reservoir Simulation on Massively Parallel Computers: Issues and Results
Serguei Maliassov, Bret Beckner, and Vadim Dyadechko, ExxonMobil Upstream Research Company, USA

3:05-3:25 A New Linear Solver Algorithm for Solving Difficult Large Scale Thermal Models
Gary Li, Chevron Corporation, USA; John Wallis, Wallis Consulting, USA; Gareth Shaw, Schlumberger-Doll Research, USA

continued in next column

3:30-3:50 Challenges and Opportunities in Incorporating State-of-the Art Parallel Computation Capabilities in Modern End-User Based Dynamic Reservoir Simulation Systems
Sander Belien, Shell Global Solutions International B.V., Rijswijk, Netherlands

3:55-4:15 Addressing Domain Decomposition, Load Balancing and Coupled Solution in An Unstructured Parallel Reservoir Simulator
Larry Fung, Saudi Aramco Oil Company, Saudi Arabia

Friday, February 21

MS68

High Performance Symbolic Computation - Part I of II

2:40 PM-4:20 PM

Room:Salon C

For Part 2 see MS76

Attaining high performance in Computer Algebra applications is often a challenging task because of the underlying complexity of computations and higher-order data types. There continues to be a growing need to develop computer algebra systems that are able to extend the range of feasible computations and to do so in a resourceful manner. This minisymposium is devoted to high performance Symbolic Computing, where "high performance" encompasses both scalable parallel algorithms as well serial algorithms that address memory bound problems in the context of the memory hierarchy, thus paving the way for parallel algorithms that reduce communication and resource contention. Talks involving algorithm design, complexity analysis, and algorithm engineering are welcome.

Organizer: Fatima K. Abu Salem
American University of Beirut, Lebanon

Organizer: Marc Moreno Maza
University of Western Ontario, Canada

2:40-3:00 Space, Time, and I/O Efficient Polynomial Factorisation Via the Funnel Heap
Fatima K. Abu Salem and Khalil El-Harake, American University of Beirut, Lebanon; Karl Gemayel, University of Oxford, United Kingdom

3:05-3:25 CUMODP: A CUDA Library for Symbolic Computation with Polynomials over Finite Fields
Marc Moreno Maza, University of Western Ontario, Canada; Wei Pan, Intel Corporation, China; Sardar Anisul Haque, Farnam Mansouri, and Ning Xie, University of Western Ontario, Canada

3:30-3:50 Toward High Performance Matrix Multiplication for Exact Computation
Pascal Giorgi, Universite de Montpellier II, France

3:55-4:15 Accelerating Path Tracking for Polynomial Homotopies
Jan Verschelde and Xiangcheng Yu, University of Illinois, Chicago, USA

Friday, February 21

MS69

Sparse Matrix Computations using Low-rank Representations

2:40 PM-4:20 PM

Room:Salon D

Low-rank representations are an increasingly popular way of speeding-up sparse matrix algorithms. They can be used for designing fast matrix-vector products, direct solvers with linear or near-linear complexity, and robust preconditioners. Many different approaches, such as H-matrices, HSS representations, or the BLR format, are currently under study by different research groups. The speakers will present their recent results in using these techniques.

Organizer: Francois-Henry Rouet
Lawrence Berkeley National Laboratory, USA

Organizer: Xiaoye S. Li
Lawrence Berkeley National Laboratory, USA

2:40-3:00 Designing Multifrontal Solvers Using Hierarchically Semiseparable Structures
François-Henry Rouet and Xiaoye S. Li, Lawrence Berkeley National Laboratory, USA; Artem Napov, Université Libre de Bruxelles, Belgium

3:05-3:25 Fast Direct Solvers Using Sparse Rocket Matrices
Eric F. Darve, Stanford University, USA

3:30-3:50 Parallelization and Pivoting in a Block-Low Rank Multifrontal Solver
Patrick Amestoy, ENSEEIHT-IRIT, France; Cleve Ashcraft, Livermore Software Technology Corporation, USA; Olivier Boiteau, EDF, France; Alfredo Buttari, CNRS, France; Jean-Yves L'Excellent, INRIA-LIP-ENS Lyon, France; *Clément Weisbecker*, ENSEEIHT-IRIT, France

3:55-4:15 Multilevel Low-Rank Approximation Preconditioners
Yousef Saad and Ruipeng Li, University of Minnesota, USA

Friday, February 21

MS70

Co-Design with Proxy Applications and Programming Abstractions: Results and Experiences - Part I of II

2:40 PM-4:20 PM

Room:Salon G

For Part 2 see MS78

Effective use of computing environments for scientific and engineering applications is determined by a combination of issues throughout a co-design optimization space: hardware, runtime environment, programming models, languages and compilers, algorithm choice and implementation, and more. The US DoE has embarked upon multiple co-design efforts that use open source proxy applications to explore this large optimization space. Presentations in this minisymposium will introduce these proxy applications, describe them in some detail so that others may use them for their own research, and report recent co-design results in DoE projects using these proxies.

Organizer: Allen McPherson
Los Alamos National Laboratory, USA

Organizer: Richard Barrett
Sandia National Laboratories, USA

Organizer: Charles (Bert) H. Still
Lawrence Livermore National Laboratory, USA

2:40-3:00 Exploring Workloads of Adaptive Mesh Refinement
Courtenay Vaughan and Richard Barrett, Sandia National Laboratories, USA

3:05-3:25 Lulesh 2.0 Changes and Initial Experiences
Ian Karlin, Lawrence Livermore National Laboratory, USA

3:30-3:50 Algorithmic Acceleration of Ocean Models
Chris Newman, Geoff Womeldorff, and Dana Knoll, Los Alamos National Laboratory, USA

3:55-4:15 Unified Task-Data-Vector Parallelism on Manycore Architectures
H. Carter Edwards and Stephen Olivier, Sandia National Laboratories, USA

Friday, February 21

MS71

Scale-bridging Algorithmic Acceleration Methods

2:40 PM-4:20 PM

Room:Salon H

Scale-bridging methods are used to target work more efficiently in space or time and to couple scale-dependent physical models. One approach is that of moment-based methods, where original problem is coupled to moment equations thereof, obtained by integration of energy, velocity, or physical space. Another example is that of parallel-in-time methods which use coarse and fine temporal scales to inform each other. In both moment-based and parallel-in-time, coarse descriptions of a problem are used to accelerate a full version. These hierarchical problem descriptions, in space or time, also naturally support heterogeneous computing; a critical path to exascale.

Organizer: Geoff Womeldorff
Los Alamos National Laboratory, USA

Organizer: Chris Newman
Los Alamos National Laboratory, USA

2:40-3:00 An Algorithmically Accelerated Ocean Model Framework: Methods and Communication Strategies
Geoff Womeldorff, Chris Newman, Dana A. Knoll, and Luis Chacon, Los Alamos National Laboratory, USA

3:05-3:25 An Asymptotic Parallel-in-Time Method for Highly Oscillatory PDEs
Terry Haut, Los Alamos National Laboratory, USA

3:30-3:50 Multi-rate and Multi-level Spectral Deferred Correction Methods: High-order Integrators with Strong Coupling for Models with Disparate Time Scales and/or Multi-physics
Matthew Emmett, Lawrence Berkeley National Laboratory, USA; Michael Minion, Stanford University, USA; Max Duarte, Lawrence Berkeley National Laboratory, USA

3:55-4:15 A Fully Implicit, Moment Accelerated, Electromagnetic Particle-in-Cell Algorithm
Guangye Chen and Luis Chacon, Los Alamos National Laboratory, USA

Friday, February 21

MS72

In-situ Data Analysis for Extreme-scale Scientific Discovery

2:40 PM-4:20 PM

Room:Salon I

The move to exascale computing is creating a disruptive shift in simulation workflows as saving the entirety of raw data to persistent storage becomes increasingly expensive. Consequently, we are moving away from a post-process centric data analysis paradigm towards a concurrent analysis framework in which raw simulation data is processed as it is computed. In this minisymposium we will explore the latest algorithmic advances, software solutions, and scientific insights made possible by this shift to in-situ processing. We will also discuss challenges, including those posed by anticipated future system architectures and the need to maintain exploratory post-processing capabilities.

Organizer: Janine C. Bennett
Sandia National Laboratories, USA

2:40-3:00 In-Situ Processing with EAVL, the Extreme-Scale Analysis and Visualization Library
Jeremy Meredith, Oak Ridge National Laboratory, USA

3:05-3:25 In-Situ Visualization with Explorable Images
Kwan-Liu Ma, Robert Miller, and Yucong Ye, University of California, Davis, USA

3:30-3:50 Approaching Production In-Situ Visualization for Extreme Scale Analysis
Kenneth Moreland, Ron A. Oldfield, and Fabian Fabian, Sandia National Laboratories, USA; Andrew Bauer, Berk Geveci, and David Lonie, Kitware, Incorporated, USA

3:55-4:15 Computing, Exploring, and Tracking Features in Massive Scientific Data
Peer-Timo Bremer, Lawrence Livermore National Laboratory and University of Utah, USA; Valerio Pascucci and Attila Gyulassy, University of Utah, USA; Janine C. Bennett, Sandia National Laboratories, USA

Friday, February 21

Coffee Break

4:20 PM-4:50 PM

Room:Salon E

MS73

Productivity, Portability, and Scalability through Domain-Specific Languages

4:50 PM-6:30 PM

Room:Salon F

Complexity in software for solving multiphysics applications is compounded by a proliferation of current and forthcoming hardware architectures. Domain-specific languages (DSLs) are emerging as an effective way of reducing the burden on application programmers by allowing abstraction of algorithmic details (e.g., stencil computations, discretization, etc.) from the deployment of these algorithms on various architectures (e.g., GPU, multicore/manycore, etc.). This minisymposium explores several DSLs being actively developed and discusses key issues such as the role of domain experts in DSL abstraction, and the balance of functionality, scalability and performance in a DSL.

Organizer: James C. Sutherland
University of Utah, USA

4:50-5:10 Nebo: A DSL for Numerical Solution of PDEs on Current and Emerging Architectures
Christopher Earl, James C. Sutherland, and Matthew Might, University of Utah, USA

5:15-5:35 What Should Yafcc Look Like?
Robert C. Kirby, Baylor University, USA

5:40-6:00 Layered DSLs for Portable Manycore Scalability
Daniel Sunderland, Patrick G. Xavier, and H. Carter Edwards, Sandia National Laboratories, USA

6:05-6:25 Harlan: High-Level Programming for GPU Systems
Eric Holk and Andrew Lumsdaine, Indiana University, USA; Matthew Might, University of Utah, USA

Friday, February 21

MS74

Visualization of Performance Data on Large Scale Systems and Applications

4:50 PM-6:30 PM

Room:Salon A

The hardware complexity of HPC systems has increased in parallel with the complexity of modern HPC applications, which has made writing efficient software difficult. Understanding the interactions between hardware and software and their impacts at large scale is essential for optimizing HPC systems, but results from classic performance tools are often too low-level and difficult to comprehend. To overcome this challenge and to gain real insight into an application's performance, we need both novel techniques in performance analysis as well as a close collaboration between the fields of performance analysis and data analytics/visualization, which this minisymposium aims at fostering.

Organizer: Martin Schulz
Lawrence Livermore National Laboratory, USA

Organizer: Joshua Levine
Clemson University, USA

Organizer: Peer-Timo Bremer
Lawrence Livermore National Laboratory and University of Utah, USA

Organizer: Paul Rosen
University of Utah, USA

4:50-5:10 Visions of Tau Dancing in Your Head - Ruminations on Performance Visualization
Allen Malony, University of Oregon, USA

5:15-5:35 Visual Characterization of High-End Computing
Kwan-Liu Ma, University of California, Davis, USA

5:40-6:00 Visual Performance Analysis for the Exascale Era
Peer-Timo Bremer, Lawrence Livermore National Laboratory and University of Utah, USA

6:05-6:25 Effectively Presenting Application Performance Using Simple Techniques
Nathan Tallent, Pacific Northwest National Laboratory, USA

Friday, February 21

MS75

CS&E Applications on Novel Architectures

4:50 PM-6:30 PM

Room:Salon B

In this minisymposium, we will present real Computational Science and Engineering applications on GPUs and heterogeneous architectures and discuss scalability issues and other practical challenges.

Organizer: Christopher Johnson
University of Utah, USA

4:50-5:10 Towards Parallel Multi-Scale Materials by Design
Alan Humphrey and Martin Berzins, University of Utah, USA

5:15-5:35 A Scalable, Efficient Scheme for Evaluation of Stencil Computations over Unstructured Meshes
Mike Kirby and James King, University of Utah, USA

5:40-6:00 Stream Architectures for PDEs on Unstructured Grids
Ross Whitaker, University of Utah, USA

6:05-6:25 Title Not Available at Time of Publication
Christopher Johnson, University of Utah, USA

Friday, February 21

MS76

High Performance Symbolic Computation - Part II of II

4:50 PM-6:30 PM

Room:Salon C

For Part 1 see MS68
Attaining high performance in Computer Algebra applications is often a challenging task because of the underlying complexity of computations and higher-order data types. There continues to be a growing need to develop computer algebra systems that are able to extend the range of feasible computations and to do so in a resourceful manner. This minisymposium is devoted to high performance Symbolic Computing, where "high performance" encompasses both scalable parallel algorithms as well serial algorithms that address memory bound problems in the context of the memory hierarchy, thus paving the way for parallel algorithms that reduce communication and resource contention. Talks involving algorithm design, complexity analysis, and algorithm engineering are welcome.

Organizer: Fatima K. Abu Salem
American University of Beirut, Lebanon

Organizer: Marc Moreno Maza
University of Western Ontario, Canada

4:50-5:10 BPAS: A CilkPlus Library for Basic Polynomial Algebra Subroutines
Yuzhen Xie, Changbo Chen, Farnam Mansouri, Marc Moreno Maza, and Ning Xie, University of Western Ontario, Canada

5:15-5:35 Symbolic Homotopies
Eric Schost, University of Western Ontario, Canada

5:40-6:00 Parallel Computation of Multivariate Polynomial Greatest Common Divisors
Michael B. Monagan, Simon Fraser University, Canada

6:05-6:25 Parallel Computation of Echelon Forms and Rank Profiles
Clément Pernet, CNRS/INRIA/LIP-ARIC ENS Lyon, France; Jean-Guillaume L. Dumas, MNC-IMAG, UJF Grenoble, France; Ziad Sultan, Grenoble University, France

Friday, February 21

MS77

Randomized Algorithms in Parallel Matrix Computations

4:50 PM-6:30 PM

Room:Salon D

Randomized algorithms are becoming very attractive in high-performance computing applications due to the significant acceleration they can provide for large size simulations while still giving very accurate results. This minisymposium will present innovative randomized algorithms and solvers in parallel matrix computations. The speakers will describe several applications of randomized algorithms in parallel computing including direct solvers for dense or sparse matrices, iterative linear solver for SPD matrices and how randomized algorithms can be implemented in parallel and distributed environments.

Organizer: Marc Baboulin
INRIA/University of Paris-Sud, France

Organizer: Xiaoye S. Li
Lawrence Berkeley National Laboratory, USA

4:50-5:10 Randomized Algorithms for Dense Linear Algebra
Marc Baboulin, INRIA/University of Paris-Sud, France

5:15-5:35 Employing Random Butterfly Transformation in Sparse Direct Solvers
Xiaoye S. Li and Francois-Henry Rouet, Lawrence Berkeley National Laboratory, USA; Marc Baboulin, INRIA/University of Paris-Sud, France

5:40-6:00 Implementing Randomized Matrix Algorithms in Parallel and Distributed Environments
Michael Mahoney, Stanford University, USA

6:05-6:25 Randomized Asynchronous Iterative Linear Solver for Spd Matrices
Haim Avron, IBM T.J. Watson Research Center, USA

Friday, February 21

MS78

Co-Design with Proxy Applications and Programming Abstractions: Results and Experiences - Part II of II

4:50 PM-6:30 PM

Room:Salon G

For Part 1 see MS70

Effective use of computing environments for scientific and engineering applications is determined by a combination of issues throughout a co-design optimization space: hardware, runtime environment, programming models, languages and compilers, algorithm choice and implementation, and more. The US DoE has embarked upon multiple co-design efforts that use open source proxy applications to explore this large optimization space. Presentations in this minisymposium will introduce these proxy applications, describe them in some detail so that others may use them for their own research, and report recent co-design results in DoE projects using these proxies.

Organizer: Allen McPherson
Los Alamos National Laboratory, USA

Organizer: Richard Barrett
Sandia National Laboratories, USA

Organizer: Charles (Bert) H. Still
Lawrence Livermore National Laboratory, USA

4:50-5:10 Co-Design Via Proxy Applications: MiniMD in Chapel
Bradford L. Chamberlain, Cray, Inc., USA; Ben Harshbarger, University of Waterloo, Canada

5:15-5:35 C++ Compiler Analysis and RAJA Programming Techniques for Performance Portability
Richard Hornung and Jeff Keasler, Lawrence Livermore National Laboratory, USA

continued in next column

5:40-6:00 Performance and Energy Comparisons between FPGA and GPU Implementations of Data Assembly
Li Tang and Xiaobo Sharon Hu, University of Notre Dame, USA; Richard Barrett, Sandia National Laboratories, USA

6:05-6:25 Non-Traditional Approaches to Development of Multi-Scale Simulation Codes
Allen McPherson, Christopher Mitchell, and Kipton Barros, Los Alamos National Laboratory, USA

Friday, February 21

MS79

Application Experiences with the Intel® Xeon Phi™ Coprocessor

4:50 PM-6:30 PM

Room:Salon H

With physical limitations imposing increasingly significant performance limitations on future generations of computing hardware, computer architects are turning to increased parallelism and specialized hardware to accelerate key applications and workloads. This minisymposium examines the impact of the Intel® Xeon Phi™ coprocessor on applications in computational science and engineering over the first year of its general availability. Each talk presents current results, provides an overview of the efforts required to achieve the results, and summarizes any lessons learned.

Organizer: R. Glenn Brook
University of Tennessee, Knoxville, USA

4:50-5:10 The Beacon Project: Application Experiences
R. Glenn Brook, University of Tennessee, Knoxville, USA

5:15-5:35 Experiences Moving Gpu-Accelerated Applications to the Intel® Xeon Phi™ Coprocessor
Bronson Messer, Oak Ridge National Laboratory & University of Tennessee, USA

5:40-6:00 Plasma Simulations on the Intel® Xeon Phi™ Coprocessor
Kai Germaschewski and *Narges Ahmadi*, University of New Hampshire, USA; Homayoun Karimabadi, University of California, San Diego, USA

6:05-6:25 Lattice Quantum Chromodynamics on the Intel® Xeon Phi™ Coprocessor
Balint Joo, Jefferson National Accelerator Laboratory, USA

Friday, February 21

MS80

Partitioning and Load-balancing for Networks and Scale-free Graphs

4:50 PM-6:30 PM

Room:Salon I

An emerging area for parallel computing is the analysis of large networks. Real-life networks (e.g., social networks, web graphs, communication networks) are often scale-free with a power-law degree distribution. How to partition (load-balance) such graphs on distributed-memory systems is a challenge, especially for large numbers of processors. Traditional graph partitioners were developed for meshes, which have small separators. This is typically not the case for scale-free graphs, where some vertices have very high degree. In this minisymposium, researchers will present recent research in this area that aims to improve scalable computations on large parallel systems.

Organizer: Erik G. Boman
Sandia National Laboratories, USA

Organizer: Karen D. Devine
Sandia National Laboratories, USA

Organizer: Siva Rajamanickam
Sandia National Laboratories, USA

4:50-5:10 Scalable Matrix Computations on Large Scale-Free Graphs Using 2D Graph Partitioning
Erik G. Boman, Karen D. Devine, and *Siva Rajamanickam*, Sandia National Laboratories, USA

5:15-5:35 Scaling Techniques for Massive Scale-Free Graphs in Distributed (External) Memory
Roger Pearce, Texas A&M University, USA

5:40-6:00 Multilevel Graph Partitioning: a Combinatorial and QP-based Method
Nuri Yeralan and Timothy A. Davis, University of Florida, USA

6:05-6:25 Panel Discussion: Parallel Computing for Big Data
Erik G. Boman, Sandia National Laboratories, USA

PP14 Abstracts

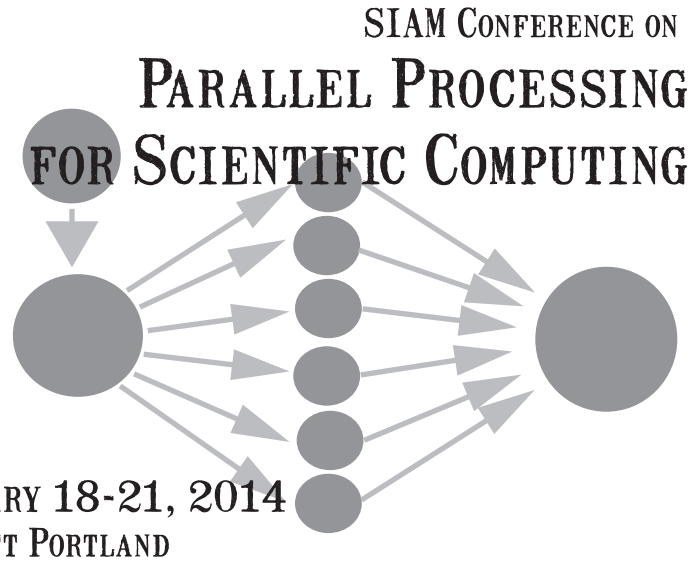

February 18-21, 2014
Marriott Portland
Downtown Waterfront
Portland, Oregon, USA

Abstracts are printed as submitted.

PP14 Speaker and Organizer Index

SIAM Conference on
Parallel Processing
for Scientific Computing

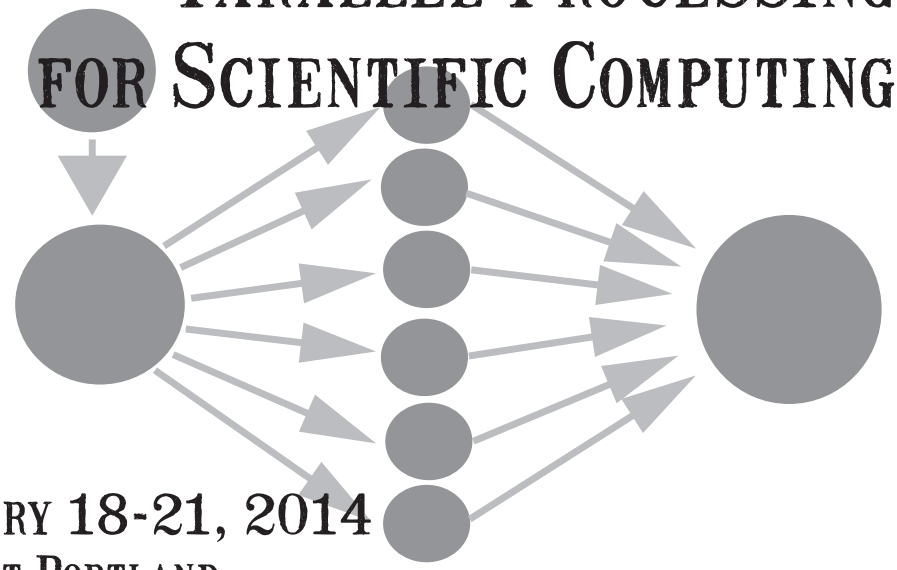

February 18-21, 2014
Marriott Portland
Downtown Waterfront
Portland, Oregon, USA

Italicized names indicate session organizers.

continued on next page

D

E

F

G

H

continued on next page

I

J

K

L

M

continued on next page

N

O

P

R

S

continued on next page

T

U

V

W

X

Y

Z

PP14 Budget

Conference Budget
SIAM Conference on Parallel Processing
February 18-21, 2014
Portland, OR

Expected Paid Attendance	430	
Revenue		
Registration Income		$152,175
Total		$152,175
Expenses		
Printing		$3,100
Organizing Committee		$4,000
Invited Speakers		$8,000
Food and Beverage		$25,670
AV Equipment and Telecommunication		$17,000
Advertising		$6,800
Conference Labor (including benefits)		$49,948
Other (supplies, staff travel, freight, misc.)		$7,919
Administrative		$14,655
Accounting/Distribution & Shipping		$7,195
Information Systems		$12,867
Customer Service		$4,766
Marketing		$7,400
Office Space (Building)		$4,041
Other SIAM Services		$4,581
Total		$177,942
Net Conference Expense		($25,767)
Support Provided by SIAM		$25,767
		$0
Estimated Support for Travel Awards not included above:		
Post Docs and Students	20	$14,300

portland marriott downtown waterfront

lower level 1

Marriott
PORTLAND DOWNTOWN
WATERFRONT
1401 SW Naito Parkway • Portland, Oregon 97201
Hotel (503) 226-7600
Sales Facsimile (503) 226-1209
portlandmarriott.com

BALLROOM LEVEL

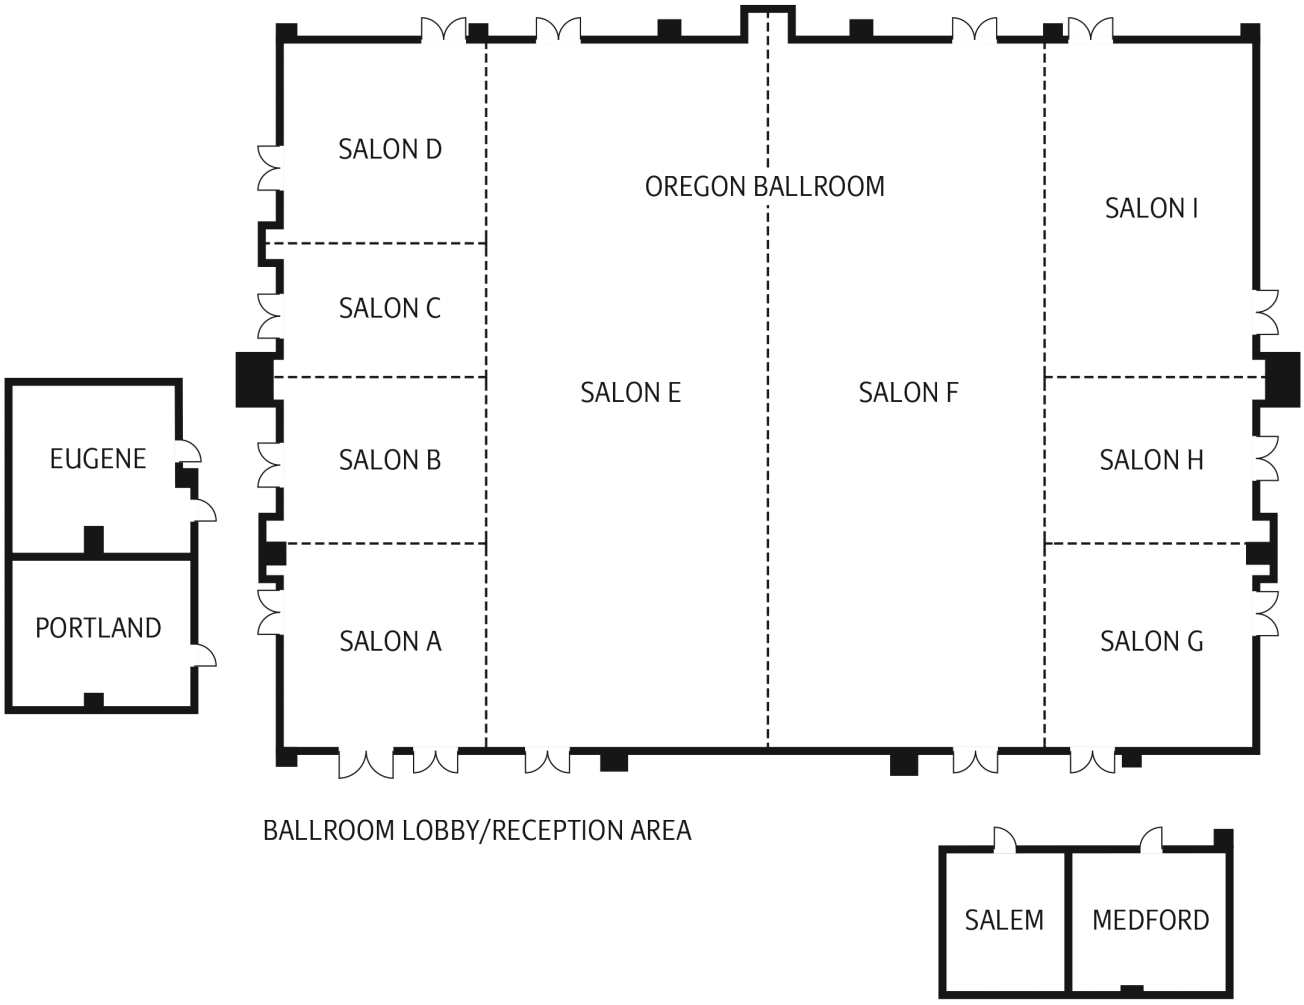

BALLROOM LOBBY/RECEPTION AREA

FSC logo text box indicating size & layout of logo. Conlins to insert logo.