Final Program and Abstracts

SIAM Conference on
Parallel Processing for Scientific Computing

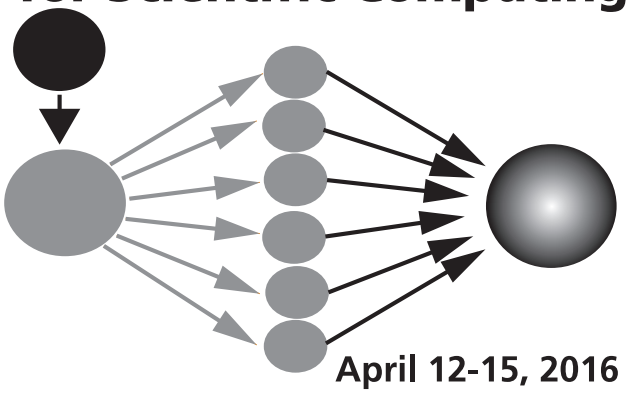

April 12-15, 2016
Université Pierre et Marie Curie,
Cordeliers Campus
Paris, France

This conference is sponsored by the SIAM Activity Group on Supercomputing, and is co-sponsored by Inria and Université Pierre et Marie Curie.

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 a biennial SIAM Conference on Parallel Processing, awards the SIAG/Supercomputing Career Prize, the SIAG/Supercomputing Early Career Prize, and the SIAG/Supercomputing Best Paper 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 Email: 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/pp16

Table of Contents

Organizing Committee Co-Chairs

Laura Grigori
Inria and Université Pierre et Marie Curie, France

Rich Vuduc
Georgia Institute of Technology, USA

Organizing Committee

Rosa M. Badia
Barcelona Supercomputing Center, Spain

Costas Bekas
IBM Research - Zurich, Switzerland

Brad Chamberlain
Cray Inc., USA

Ron Dror
Stanford University, USA

Mary Hall
University of Utah, USA

Bruce Hendrickson
Sandia National Laboratories, USA

Naoya Maruyama
RIKEN, Japan

Philippe Ricoux
Total, France

Olaf Schenk
Università della Svizzera italiana (USI), Switzerland

John Shalf
Lawrence Berkeley National Laboratory, USA

Peter Tang
Intel Corporation, USA

Sivan Toledo
Tel Aviv University, Israel

Weichung Wang
National Taiwan University, Taiwan

Barbara Wohlmuth
Technische Universität München, Germany

Local Organizing Committee

Xavier Claeys
Université Pierre et Marie Curie and Inria, France

Pierre Fortin
Université Pierre et Marie Curie, France

Jean-Frederic Gerbeau
Inria and Université Pierre et Marie Curie, France

Chantal Girodon
Inria, France

Frederic Hecht
Université Pierre et Marie Curie and Inria, France

Yvon Maday
Université Pierre et Marie Curie and IUF, France and DAM, Brown University, USA

Frederic Nataf
CNRS, Université Pierre et Marie Curie and Inria, France

SIAM Registration Desk

Registration and all sessions will be located at Université Pierre et Marie Curie, Cordeliers Campus.

The SIAM registration desk is located in the Entrance Hall. It is open during the following hours:

Tuesday, April 12
10:00 AM – 7:00 PM

Wednesday, April 13
7:45 AM- 5:00 PM

Thursday, April 14
7:45 AM - 5:15 PM

Friday, April 15
8:15 AM - 5:15 PM

Conference Location

All sessions and on-site registration will take place at the Université Pierre et Marie Curie, Cordeliers Campus. The entrance is 21 rue de l'Ecole de Médecine.

Université Pierre et Marie Curie, Cordeliers Campus

21 rue de l'Ecole de Médecine

75006 Paris

France

Child Care

Please visit *http://paris.angloinfo.com/af/705/baby-sitting-nanny-and-au-pair-services.html* for information about Baby Sitting services in Paris and Ile-de-France.

Corporate Members and Affiliates

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

Corporate Institutional Members

The Aerospace Corporation

Air Force Office of Scientific Research

Aramco Services Company

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

ExxonMobil Upstream Research

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

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 March 2016.

Funding Agencies

SIAM and the Conference Organizing Committee wish to extend their thanks and appreciation to the U.S. National Science Foundation and DOE Office of Advanced Scientific Computing Research for their support of this conference.

Leading the applied mathematics community

Join SIAM and save!

SIAM Members save up to 100 € on full registration for the 2016 SIAM Conference on Parallel Processing for Scientific Computing (PP16)! Join your peers in supporting the premier professional society for applied mathematicians and computational scientists. SIAM members receive subscriptions to *SIAM Review*, *SIAM News*, and *SIAM Unwrapped*, and enjoy substantial discounts on SIAM books, journal subscriptions, and conference registrations.

For a limited time, nonmember registrants to PP16 receive a 25% discount off of their annual membership dues.* Please complete and return the application you received with your registration materials or sign up online at *https://my.siam.org* using code MBPP16 when prompted.

***this discount is available to new SIAM members only.**

Internet Access

Wireless internet details will be distributed to participants onsite.

Complimentary wireless Internet access will be available in the meeting rooms.

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.

Each meeting room will have one (1) screen and one (1) data projector. The data projectors support VGA connections only. Presenters requiring alternate connection must provide their own adaptor.

Registration Fee Includes

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

Optional Conference Dinner Banquet

A Conference Dinner Banquet will be held on Thursday, April 14, from 7:30 PM - 10:30 PM at Salons Hoche (*http://www.salons-hoche.fr/)*, 9 Avenue Hoche, 75008, Paris.

A separate fee of 82.50 Euros applies. Dinner includes appetizer, main dish, dessert, and coffee.

Salons Hoche is 1/2 mile from Metro Charles de Gaulle Etoile (Arc de Triomphe)

From Les Cordeliers:

Metro: line 4 from station Odeon (Direction Porte de Clignancourt) to station Chatelet

then, change to line 1 (Direction LA DEFENSE) to station Charles de Gaulle-Etoile

and then 10 minutes walk on Avenue Hoche

ETA : 35 minutes

Visit the registration desk for more information.

Job Postings

Visit *http://jobs.siam.org.*

Important Notice to Poster Presenters

The poster session is scheduled for Wednesday, April 13 at 7:15 PM - 9:15 PM. Poster presenters are requested to set up their poster material on the provided 157 cm x 115 cm poster boards in the Exhibition Space. All materials must be posted by Wednesday, April 13 at 7:15 PM, the official start time of the session. Poster displays must be removed by the published time. 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, orders may be placed according to the instructions on the Titles on Display form.

Get-togethers

Welcome Reception
Tuesday, April 12, 2016
7:00 PM - 9:00 PM

Poster Session
Wednesday, April 13, 2016
7:15 PM - 9:15 PM

Business Meeting
(open to SIAG/SC members)
Wednesday, April 13, 2016
6:30 PM - 7:15 PM

Dinner Banquet (separate fees apply)
Thursday, April 14,
7:30 PM - 10:30 PM

Salons Hoche
(http://www.salons-hoche.fr/)

Conference Sponsors

Comments?

Comments about SIAM meetings are encouraged! Please send to:

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

Please Note

The local organizers are 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.

Statement on Inclusiveness

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

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.

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 mathematical algorithms and computer architecture.

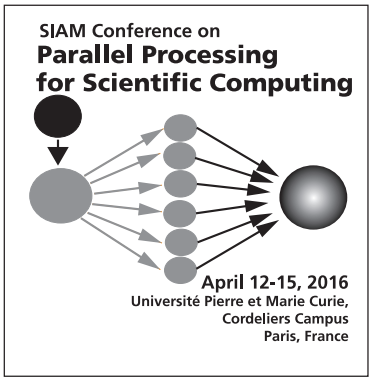

ACTIVITIES INCLUDE:

- Special sessions at SIAM Annual Meetings
- Biennial conference
- SIAG/SC Career Prize
- SIAG/SC Early Career Prize
- SIAG/SC Best Paper Prize

BENEFITS OF SIAG/SC MEMBERSHIP:

- Listing in the SIAG's online-only membership directory
- 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 two activity groups for free!

2016-17 SIAG/SC OFFICERS

- Chair: Laura Grigori, INRIA
- Vice Chair: Richard Vuduc, Georgia Institute of Technology
- Program Director: Olaf Schenk, Universita Della Svizzera italiana
- Secretary, Mark Hoemmen, Sandia National Laboratories

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 Farabeuf .

Tuesday, April 12

5:15 PM - 6:00 PM

IP1 Scalability of Sparse Direct Codes

Iain Duff, Science & Technology Facilities Council, United Kingdom and CERFACS, Toulouse, France

Wednesday, April 13

8:15 AM - 9:00 AM

IP2 Towards Solving Coupled Flow Problems at the Extreme Scale

Ulrich J. Ruede, University of Erlangen-Nuremberg, Germany

1:45 PM - 2:30 PM

IP3 Why Future Systems should be Data Centric

Paul Coteus, IBM Research, USA

Invited Plenary Speakers

All Invited Plenary Presentations will take place in Farabeuf .

Thursday, April 14

8:15 AM - 9:00 AM

IP4 Next-Generation AMR

Ann S. Almgren, Lawrence Berkeley National Laboratory, USA

1:45 PM - 2:30 PM

IP5 Beating Heart Simulation Based on a Stochastic Biomolecular Model

Takumi Washio, University of Tokyo, Japan

Friday, April 15

8:45 AM - 9:30 AM

IP6 Conquering Big Data with Spark

Ion Stoica, University of California, Berkeley, USA

1:45 PM - 2:30 PM

IP7 A Systems View of High Performance Computing

Simon Kahan, Institute for Systems Biology and University of Washington, USA

SIAG Supercomputing Prizes

Thursday, April 14

9:25 AM - 9:50 AM

SP1 SIAG/Supercomputing Career Prize: Supercomputers and Superintelligence

Horst D. Simon, *Lawrence Berkeley National Laboratory, USA*

9:55 AM - 10:10 AM

SIAG/Supercomputing Early Career Prize

Tobin Isaac, *University of Chicago, USA*

Friday, April 15

SP2 SIAG/Supercomputing Best Paper Prize:

Communication-Optimal Parallel and Sequential QR and LU Factorizations

James W. Demmel, *University of California, Berkeley, USA*

Laura Grigori, *Inria, France*

Mark Hoemmen, *Sandia National Laboratories, USA*

Julien Langou, *University of Colorado, Denver, USA*

PP16 Program

SIAM Conference on
Parallel Processing for Scientific Computing

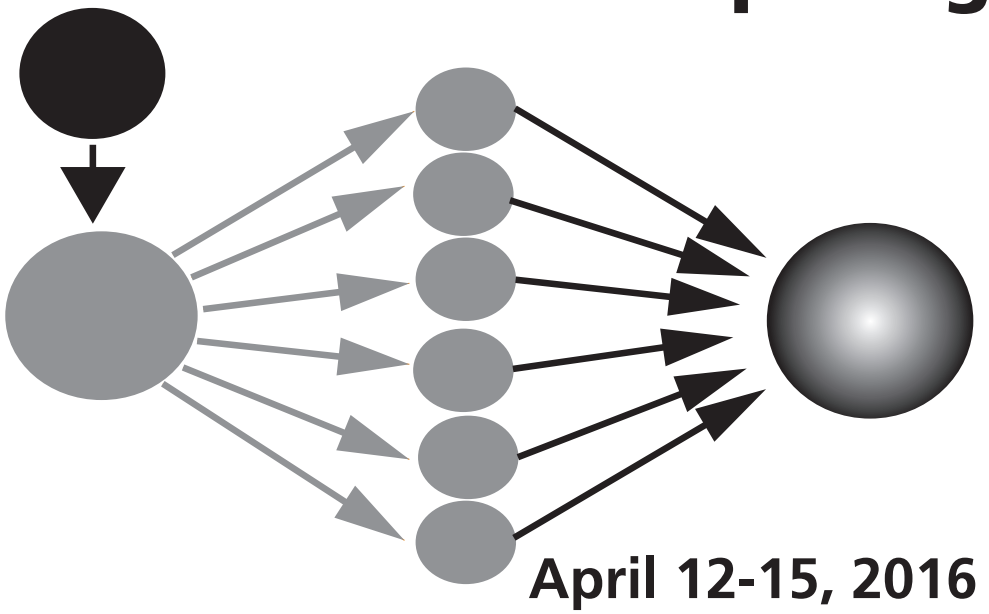

April 12-15, 2016
Université Pierre et Marie Curie,
Cordeliers Campus
Paris, France

Tuesday, April 12

Registration

10:00 AM-7:00 PM

Room:Entrance Hall

Tuesday, April 12

MS1

Parallel Programming Models, Algorithms and Frameworks for Extreme Computing - Part I of III

1:10 PM-2:50 PM

Room:Farabeuf

For Part 2 see MS9

Multicore/manycore 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: Kengo Nakajima
University of Tokyo, Japan

Organizer: Michael Heroux
Sandia National Laboratories, USA

Organizer: Serge G. Petiton
Université Lille 1 and CNRS, France

1:10-1:30 Opportunities and Challenges in Developing and Using Scientific Libraries on Emerging Architectures
Michael Heroux, Sandia National Laboratories, USA

1:35-1:55 Linear Algebra for Data Science Through Examples
Nahid Emad, University of Versailles, France

2:00-2:20 Coarse Grid Aggregation for Sa-Amg Method with Multiple Kernel Vectors
Akihiro Fujii, Naoya Nomura, and Teruo Tanaka, Kogakuin University, Japan; Kengo Nakajima, University of Tokyo, Japan; Osni A. Marques, Lawrence Berkeley National Laboratory, USA

2:25-2:45 Evaluation of Energy Profile for Krylov Iterative Methods on Cluster of Accelerators
Langshi Chen, CNRS, France; Serge G. Petiton, Université Lille 1 and CNRS, France

continued in next column

Tuesday, April 12

MS2

Improving Performance, Throughput, and Efficiency of HPC Centers through Full System Data Analytics - Part I of II

1:10 PM-2:50 PM

Room:Pasquier

For Part 2 see MS10

As the complexity of HPC centers increases, understanding how to optimize full system performance will require big data analytics. HPC centers comprise parallel supercomputing systems, application codes, and facility systems such as power and cooling, and applications compete for shared compute, storage, networking, and power resources. Techniques to monitor, manage, and analyze these resources to optimize HPC center efficiency are in their infancy. This minisymposium will bring together experts on:
1) measurement and monitoring;
2) data integration and storage; and
3) visualization and analytics, to develop parallel analysis techniques for optimizing the performance, throughput, and efficiency of HPC centers. LLNL-ABS-677080.

Organizer: Abhinav Bhatele
Lawrence Livermore National Laboratory, USA

Organizer: Todd Gamblin
Lawrence Livermore National Laboratory, USA

Organizer: James Brandt
Sandia National Laboratories, USA

Organizer: Ann Gentile
Sandia National Laboratories, USA

continued in next column

1:10-1:30 Correlating Facilities, System and Application Data Via Scalable Analytics and Visualization
Abhinav Bhatele and Todd Gamblin, Lawrence Livermore National Laboratory, USA; Alfredo Gimenez and Yin Yee Ng, University of California, Davis, USA

1:35-1:55 A Monitoring Tool for HPC Operations Support
R. Todd Evans, James C. Browne, John West, and Bill Barth, University of Texas at Austin, USA

2:00-2:20 Developing a Holistic Understanding of I/O Workloads on Future Architectures
Glenn Lockwood and Nicholas Wright, Lawrence Berkeley National Laboratory, USA

2:25-2:45 MuMMI: A Modeling Infrastructure for Exploring Power and Execution Time Tradeoffs in Parallel Applications
Valerie Taylor and Xingfu Wu, Texas A&M University, USA

Tuesday, April 12

MS3

Combinatorial Scientific Computing - Part I of III

1:10 PM-2:50 PM

Room:Roussy

For Part 2 see MS11

Combinatorial algorithms and tools are used in enabling parallel scientific computing applications. The general approach is to identify performance issues in an application and design, analyze, implement combinatorial algorithms to tackle the identified issues. The proposed minisymposium gathers 12 talks, covering applications in bioinformatics, solvers of linear systems, and data analysis; and graph algorithms for those applications. The objective is to summarize the latest combinatorial algorithmic developments and the needs of the applications. The goal is to cross-fertilize the both domains: the applications will raise new challenges to the combinatorial algorithms, and the combinatorial algorithms will address some of the existing problems of the applications.

Organizer: Aydin Buluc
Lawrence Berkeley National Laboratory, USA

Organizer: Alex Pothen
Purdue University, USA

1:10-1:30 Parallel Algorithms for Automatic Differentiation
Alex Pothen and Mu Wang, Purdue University, USA

1:35-1:55 Parallel Combinatorial Algorithms in Sparse Matrix Computation?
Mathias Jacquelin and *Esmond G. Ng*, Lawrence Berkeley National Laboratory, USA; Barry Peyton, Dalton State College, USA; Yili Zhrng and Kathy Yelick, Lawrence Berkeley National Laboratory, USA

2:00-2:20 Parallel Graph Matching Algorithms Using Matrix Algebra
Ariful Azad and Aydin Buluc, Lawrence Berkeley National Laboratory, USA

2:25-2:45 Randomized Linear Algebra in Large-Scale Computational Environments
Michael Mahoney, Alex Gittens, and Farbod Roosta-Khorasani, University of California, Berkeley, USA

Tuesday, April 12

MS4

Minimizing Communication in Numerical Algorithms - Part I of II

1:10 PM-2:50 PM

Room:Salle des theses

For Part 2 see MS12

The growing gap between the costs of computation and communication in terms of both time and energy necessitates the development of new algorithms which minimize data movement, both across the network and within the memory hierarchy, in order to make efficient use of today's and future hardware. 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

Hebrew University of Jerusalem, Israel

Organizer: Erin C. Carson

New York University, USA

1:10-1:30 Communication-Avoiding Krylov Subspace Methods in Theory and Practice

Erin C. Carson, New York University, USA

1:35-1:55 Enlarged GMRES for Reducing Communication When Solving Reservoir Simulation Problems

Hussam Al Daas, UPMC-Inria-TOTAL, France; Laura Grigori, Inria, France; Pascal Henon, Total E&P, France; Philippe Ricoux, TOTAL SA, France

2:00-2:20 CA-SVM: Communication-Avoiding Support Vector Machines on Distributed Systems

Yang You, University of California, Berkeley, USA

2:25-2:45 Outline of a New Roadmap to Permissive Communication and Applications That Can Benefit

James A. Edwards and Uzi Vishkin, University of Maryland, USA

Tuesday, April 12

MS5

Parallel-in-Time Integration Methods - Part I of III

1:10 PM-2:50 PM

Room:Marie Curie

For Part 2 see MS13

With the expected further increase in available computing power, new parallelisation strategies have to be developed to tackle the scalability challenges arising from the trend towards massively parallel architectures. In association with upcoming Exascale computing, simulations with real-time requirements face particular challenges in HPC. These challenges led to a growing interest in parallel-in-time methods over the last decade. Adding temporal parallelisation has been shown to be able to improve or extend the scalability of programs relying on traditional parallelisation-in-space methods. We invite speakers to present their current research progress related to parallel-in-time methods from all areas.

Organizer: Martin Schreiber

University of Exeter, United Kingdom

Organizer: Tobias Neckel

Technische Universität München, Germany

Organizer: Daniel Ruprecht

University of Leeds, United Kingdom

1:10-1:30 Time Parallel and Space Time Solution Methods - From Heat Equation to Fluid Flow

Rolf Krause, Università della Svizzera Italiana, Italy; *Pietro Benedusi*, USI, Switzerland; Peter Arbenz and Daniel Hupp, ETH Zürich, Switzerland

1:35-1:55 Time Parallelizing Edge Plasma Simulations Using the Parareal Algorithm

Debasmita Samaddar, UK Atomic Energy Authority, United Kingdom; David Coster, Max-Planck-Institut für Plasmaphysik, Germany; Xavier Bonnin, ITER Organization, France; Eva Havlickova, UK Atomic Energy Authority, United Kingdom; Wael R. Elwasif, Donald B. Batchelor, and Lee A. Berry, Oak Ridge National Laboratory, USA; Christoph Bergmeister, University of Innsbruck, Austria

2:00-2:20 Micro-Macro Parareal Method for Stochastic Slow-Fast Systems with Applications in Molecular Dynamics

Giovanni Samaey and Keith Myerscough, KU Leuven, Belgium; Frederic Legoll, Ecole Nationale des Ponts et Chaussées, France; Tony Lelievre, École des Ponts ParisTech, France

2:25-2:45 A Parallel-in-Time Solver for Time-Periodic Navier-Stokes Problems

Daniel Hupp and Peter Arbenz, ETH Zürich, Switzerland; Dominik Obrist, University of Bern, Switzerland

continued in next column

Tuesday, April 12

MS6

The Strong Scaling Limit: From Industrial Applications to New Algorithms - - Part I of III

1:10 PM-2:50 PM

Room:Leroux

For Part 2 see MS14

There is a strong interest in running parallel applications in the strong scaling limit. There, however, there is an increasing communication cost and a decreasing compute cost. In this mini symposium, we analyse how various applications behave in this limit and how novel communication hiding/avoiding algorithms can help. We start with industrial and scientific applications and discuss both the software and algorithmic challenges. This symposium reports work of the Exa2ct project, a EU supported project.

Organizer: Wim I. Vanroose

University of Antwerp, Belgium

1:10-1:30 Toward Large-Eddy Simulation of Complex Burners with Exascale Super-Computers: A Few Challenges and Solutions

Vincent R. Moureau, CORIA, France; Pierre Bénard, Université de Rouen, France

1:35-1:55 Extraction the Compute and Communications Patterns from An Industrial Application

Thomas Guillet, Intel Corporation, France

2:00-2:20 A Tiny Step Toward a Unified Task Scheduler for Large Scale Heterogeneous Architectures

Clement Fontenaille, PRiSM - UVSQ, France

2:25-2:45 Parallel GMRES Using A Variable S-Step Krylov Basis

David Imberti and *Jocelyne Erhel*, Inria-Rennes, France

Tuesday, April 12

MS7

Task-based Scientific, High Performance Computing on Top of Runtime Systems - Part I of II

1:10 PM-2:50 PM

Room:Dejerine

For Part 2 see MS15

The extreme complexity of hardware platforms makes them harder and harder to program. To fully exploit such machines, the High Performance Computing community often uses a MPI + X (X being pthreads, OpenMP, Cuda ...) programming models. In this minisymposium, we overview an alternative solution consisting of programming at a higher level of abstractions by descrbing a scientific, high performance computing application as a sequence of tasks whose execution is delegated to a runtime system. We discuss the potential of numerical libraries that have followed that design as well as recent advances made by fully-featured runtime systems to support them.

Organizer: Emmanuel Agullo

Inria, France

Organizer: Alfredo Buttari

CNRS, France

Organizer: Abdou Guermouche

LaBRI, France

1:10-1:30 Programming Linear Algebra with Python: A Task-Based Approach

Rosa M. Badia, Barcelona Supercomputing Center, Spain

1:35-1:55 Parsec: A Distributed Runtime System for Task-Based Heterogeneous Computing

Aurelien Bouteiller and George Bosilca, University of Tennessee, Knoxville, USA

2:00-2:20 Optimizing Numerical Simulations of Elastodynamic Wave Propagation Thanks to Task-Based Parallel Programming

Lionel Boillot, Inria Bordeaux Sud-Ouest, France; Corentin Rossignon, Total, France; George Bosilca, University of Tennessee, Knoxville, USA; Emmanuel Agullo, Inria, France; Henri Calandra, Total, France

2:25-2:45 Controlling the Memory Subscription of Applications with a Task-Based Runtime System

Marc Sergent, Inria, France; David Goudin, CEA/CESTA, France; Samuel Thibault, University of Bordeaux, France; Olivier Aumage, Inria, France

continued in next column

Tuesday, April 12

MS8

Extreme Scale Solvers for Coupled Systems

1:10 PM-2:50 PM

Room:Delarue

Exascale computers will exhibit billion-way parallelism. Computing on such extreme scale needs methods scaling perfectly with optimal complexity. This minisymposium combines talks on crucial aspects of extreme scale solving. The solver must be of optimal complexity, which is more an more severe with increasing problem size, and scale efficiently on extreme scales of parallelism. To that end, the minisymposium brings together talks on parallel adaptive multigrid methods in space and time, as well as optimziation and uncertainty quantification techniques. Also reducing power consumption will be a topic of the minisymposium.

Organizer: Gabriel Wittum
University of Frankfurt, Germany

1:10-1:30 Uncertainty Quantification Using Tensor Methods and Large Scale Hpc
Lars Grasedyck and Christian Löbbert, RWTH Aachen University, Germany

1:35-1:55 Parallel Methods in Space and Time and Their Application in Computational Medicine
Rolf Krause, Università della Svizzera Italiana, Italy; Andreas Kreienbühl, USI, Switzerland; Dorian Krause, University of Lugano, Switzerland; Gabriel Wittum, University of Frankfurt, Germany

2:00-2:20 Scalable Shape Optimization Methods for Structured Inverse Modeling Using Large Scale Hpc
Volker H. Schulz and *Martin Siebenborn*, University of Trier, Germany

2:25-2:45 Parallel Adaptive and Robust Multigrid
Gabriel Wittum, University of Frankfurt, Germany

Tuesday, April 12

CP1

Advances in Material Science

1:10 PM-2:10 PM

Room:Danton

Chair: Masha Sosonkina, Old Dominion University, USA

1:10-1:25 Parallel Recursive Density Matrix Expansion in Electronic Structure Calculations
Anastasia Kruchinina, Emanuel H. Rubensson, and Elias Rudberg, Uppsala University, Sweden

1:30-1:45 Comparison of Optimization Techniques in Parallel Ab Initio Nuclear Structure Calculations
Masha Sosonkina, Old Dominion University, USA

1:50-2:05 Scale-Bridging Modeling of Material Dynamics: Petascale Assessments on the Road to Exascale
Timothy C. Germann, Los Alamos National Laboratory, USA

Coffee Break

2:50 PM-3:20 PM

Room:Exhibition Center

Tuesday, April 12

MS9

Parallel Programming Models, Algorithms and Frameworks for Extreme Computing - Part II of III

3:20 PM-5:00 PM

Room:Farabeuf

For Part 1 see MS1
For Part 3 see MS17
Multicore/manycore 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: Kengo Nakajima
University of Tokyo, Japan

Organizer: Michael Heroux
Sandia National Laboratories, USA

Organizer: Serge G. Petiton
Université Lille 1 and CNRS, France

3:20-3:40 Thread-Scalable FEM
Kengo Nakajima, University of Tokyo, Japan

3:45-4:05 Iterative ILU Preconditioning
Hartwig Anzt, University of Tennessee, USA; Edmond Chow, Georgia Institute of Technology, USA

4:10-4:30 ICB: IC Preconditioning with a Fill-in Strategy Considering Simd Instructions
Takeshi Iwashita, Hokkaido University, Japan; Naokazu Takemura, Akihiro Ida and Hiroshi Nakashima, Kyoto University, Japan

4:35-4:55 Kokkos: Manycore Programmability and Performance Portability
H. Carter Edwards and Christian Trott, Sandia National Laboratories, USA

Tuesday, April 12

MS10

Improving Performance, Throughput, and Efficiency of HPC Centers through Full System Data Analytics - Part II of II

3:20 PM-5:00 PM

Room:Pasquier

For Part 1 see MS2

As the complexity of HPC centers increases, understanding how to optimize full system performance will require big data analytics. HPC centers comprise parallel supercomputing systems, application codes, and facility systems such as power and cooling, and applications compete for shared compute, storage, networking, and power resources. Techniques to monitor, manage, and analyze these resources to optimize HPC center efficiency are in their infancy. This mini-symposium will bring together experts on: 1) measurement and monitoring; 2) data integration and storage; and 3) visualization and analytics, to develop parallel analysis techniques for optimizing the performance, throughput, and efficiency of HPC centers. LLNL-ABS-677080.

Organizer: Abhinav Bhatele
Lawrence Livermore National Laboratory, USA

Organizer: Todd Gamblin
Lawrence Livermore National Laboratory, USA

Organizer: James Brandt
Sandia National Laboratories, USA

Organizer: Ann Gentile
Sandia National Laboratories, USA

3:20-3:40 Smart HPC Centers: Data, Analysis, Feedback, and Response
James Brandt and Ann Gentile, Sandia National Laboratories, USA; Cindy Martin, Los Alamos National Laboratory, USA; Benjamin Allan and Karen D. Devine, Sandia National Laboratories, USA

3:45-4:05 Low-Overhead Monitoring of Parallel Applications
John Mellor-Crummey and Mark Krentel, Rice University, USA

4:10-4:30 Data Management and Analysis for An Energy Efficient HPC Center
Ghaleb Abdulla, Anna Maria Bailey, and John Weaver, Lawrence Livermore National Laboratory, USA

4:35-4:55 Discovering Opportunities for Data Analytics Through Visualization of Large System Monitoring Datasets
Michael Showerman, University of Illinois at Urbana-Champaign, USA

continued in next column

Tuesday, April 12

MS11

Combinatorial Scientific Computing - Part II of III

3:20 PM-5:00 PM

Room:Roussy

For Part 1 see MS3
For Part 3 see MS19

Combinatorial algorithms and tools are used in enabling parallel scientific computing applications. The general approach is to identify performance issues in an application and design, analyze, implement combinatorial algorithms to tackle the identified issues. The proposed minisymposium gathers 12 talks, covering applications in bioinformatics, solvers of linear systems, and data analysis; and graph algorithms for those applications. The objective is to summarize the latest combinatorial algorithmic developments and the needs of the applications. The goal is to cross-fertilize the both domains: the applications will raise new challenges to the combinatorial algorithms, and the combinatorial algorithms will address some of the existing problems of the applications.

Organizer: Aydin Buluc
Lawrence Berkeley National Laboratory, USA

Organizer: Alex Pothen
Purdue University, USA

3:20-3:40 A Partitioning Problem for Load Balancing and Reducing Communication from the Field of Quantum Chemistry
Edmond Chow, Georgia Institute of Technology, USA

3:45-4:05 Community Detection on GPU
Md Naim and Fredrik Manne, University of Bergen, Norway

4:10-4:30 Scalable Parallel Algorithms for De Novo Assembly of Complex Genomes
Evangelos Georganas, University of California, Berkeley, USA

4:35-4:55 Talk Title To Be Announced
Aydin Buluc, Lawrence Berkeley National Laboratory, USA

Tuesday, April 12

MS12

Minimizing Communication in Numerical Algorithms - Part II of II

3:20 PM-5:00 PM

Room:Salle des theses

For Part 1 see MS4

The growing gap between the costs of computation and communication in terms of both time and energy necessitates the development of new algorithms which minimize data movement, both across the network and within the memory hierarchy, in order to make efficient use of today's and future hardware. 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
Hebrew University of Jerusalem, Israel

Organizer: Erin C. Carson
New York University, USA

3:20-3:40 Communication-Efficient Evaluation of Matrix Polynomials
Sivan A. Toledo, Tel Aviv University, Israel

3:45-4:05 Communication-Optimal Loop Nests
Nicholas Knight, New York University, USA

4:10-4:30 Write-Avoiding Algorithms
Harsha Vardhan Simhadri, Lawrence Berkeley National Laboratory, USA

4:35-4:55 Lower Bound Techniques for Communication in the Memory Hierarchy
Gianfranco Bilardi, University of Padova, Italy

Tuesday, April 12

MS13

Parallel-in-Time Integration Methods - Part II of III

3:20 PM-5:00 PM

Room:Marie Curie

For Part 1 see MS5
For Part 3 see MS21

With the expected further increase in available computing power, new parallelisation strategies have to be developed to tackle the scalability challenges arising from the trend towards massively parallel architectures. In association with upcoming Exascale computing, simulations with real-time requirements face particular challenges in HPC. These challenges led to a growing interest in parallel-in-time methods over the last decade. Adding temporal parallelisation has been shown to be able to improve or extend the scalability of programs relying on traditional parallelisation-in-space methods. We invite speakers to present their current research progress related to parallel-in-time methods from all areas.

Organizer: Martin Schreiber
University of Exeter, United Kingdom

Organizer: Tobias Neckel
Technische Universität München, Germany

Organizer: Daniel Ruprecht
University of Leeds, United Kingdom

3:20-3:40 Implementating Parareal - OpenMP Or MPI?
Daniel Ruprecht, University of Leeds, United Kingdom

3:45-4:05 Fault-Tolerance of the Parallel-in-Time Integration Method PFASST
Robert Speck and *Torbjörn Klatt*, Jülich Supercomputing Centre, Germany; Daniel Ruprecht, University of Leeds, United Kingdom

4:10-4:30 Parallel-in-Time Coupling of Models and Numerical Methods
Franz Chouly, Université de Franche-Comté & CNRS, France; Matteo Astorino, École Polytechnique Fédérale de Lausanne, Switzerland; Alexei Lozinski, CNRS, France; Alfio Quarteroni, École Polytechnique Fédérale de Lausanne, Switzerland

4:35-4:55 Parareal in Time and Fixed-Point Schemes
Olga Mula, CEREMADE Universite Paris 9 Dauphine, France

continued in next column

Tuesday, April 12

MS14

The Strong Scaling Limit: From Industrial Applications to New Algorithms - Part II of III

3:20 PM-5:00 PM

Room:Leroux

For Part 1 see MS6
For Part 3 see MS22

There is a strong interest in running parallel applications in the strong scaling limit. There, however, there is an increasing communication cost and a decreasing compute cost. In this minisymposium, we analyse how various applications behave in this limit and how novel communication hiding/avoiding algorithms can help. We start with industrial and scientific applications and discuss both the software and algorithmic challenges. This symposium reports work of the Exa2ct project, a EU supported project.

Organizer: Wim I. Vanroose
University of Antwerp, Belgium

3:20-3:40 Rounding Errors in Pipelined Krylov Solvers
Siegfried Cools, University of Antwerp, Belgium

3:45-4:05 Combining Software Pipelining with Numerical Pipelining in the Conjugate Gradient Algorithm
Emmanuel Agullo, Luc Giraud, and Stojce Nakov, Inria, France

4:10-4:30 Notification Driven Collectives in Gaspi
Christian Simmendinger, T-Systems International, Germany

4:35-4:55 Espreso: ExaScale PaRallel Feti Solver
Lubomir Riha, Tomas Brzobohaty, and Alexandros Markopoulos, VSB-Technical University Ostrava, Czech Republic

Tuesday, April 12

MS15

Task-based Scientific, High Performance Computing on Top of Runtime Systems - Part II of II

3:20 PM-5:00 PM

Room:Dejerine

For Part 1 see MS7

The extreme complexity of hardware platforms makes them harder and harder to program. To fully exploit such machines, the High Performance Computing community often uses a MPI + X (X being pthreads, OpenMP, Cuda ...) programming models. In this minisymposium, we overview an alternative solution consisting of programming at a higher level of abstractions by descrbing a scientific, high performance computing application as a sequence of tasks whose execution is delegated to a runtime system. We discuss the potential of numerical libraries that have followed that design as well as recent advances made by fully-featured runtime systems to support them.

Organizer: Emmanuel Agullo
Inria, France

Organizer: Alfredo Buttari
CNRS, France

Organizer: Abdou Guermouche
LaBRI, France

3:20-3:40 Task-Based Multifrontal QR Solver for GPU-Accelerated Multicore Architectures
Florent Lopez, Universite Paul Sabatier, France

3:45-4:05 An Effective Methodology for Reproducible Research on Dynamic Task-Based Runtime Systems
Arnaud Legrand, Inria Grenoble, France; Luka Stanisic, Inria Grenoble Rhône-Alpes, France

4:10-4:30 High Performance Task-Based QDWH-eig on Manycore Systems
Dalal Sukkari, Hatem Ltaief, and David E. Keyes, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

4:35-4:55 Exploiting Two-Level Parallelism by Aggregating Computing Resources in Task-Based Applications Over Accelerator-Based Machines
Terry Cojean, Inria, France

continued in next column

Tuesday, April 12

MS16

Numerical Reproducibility for High-Performance Computing

3:20 PM-5:00 PM

Room:Delarue

Numerical reproducibility is the property to get a bitwise identical floating-point result from multiple runs of the same code. Floating-point operations are not associative which hinder numerical reproducibility especially on large-scale heterogeneous platform. This property of floating-point calculation causes validations and debugging issues and may lead to deadlock. Those problems occur especially with parallel implementation embedding exchange of floating-point results in linear algebra libraries. This minisymposium will address some of those problems and speakers will present both theoretical and practical solutions.

Organizer: David Defour
University of Perpignan, France

Organizer: Stef Graillat
University Pierre and Marie Curie (UPMC), France

3:20-3:40 Reproducible and Accurate Algorithms for Numerical Linear Algebra
Roman Iakymchuk, KTH Royal Institute of Technology, Sweden; David Defour, University of Perpignan, France; Sylvain Collange, Inria Rennes, France; Stef Graillat, University Pierre and Marie Curie (UPMC), France

3:45-4:05 Summation and Dot Products with Reproducible Results
Siegfried M. Rump, Hamburg University of Technology, Germany

4:10-4:30 Bit-Reproducible Climate and Weather Simulations: From Theory to Practice
Andrea Arteaga, ETH Zürich, Switzerland; Christophe Charpilloz and Olive Fuhrer, MeteoSwiss, Switzerland; Torsten Hoefler, ETH Zürich, Switzerland

4:35-4:55 Parallel Interval Algorithms, Numerical Guarantees and Numerical Reproducibility
Nathalie Revol, Inria - LIP, Ecole Normale Superieure de Lyon, France

continued in next column

Tuesday, April 12

CP2

Eigensolvers

3:20 PM-5:00 PM

Room:Danton

Chair: Aaron Melman, Santa Clara University, USA

3:20-3:35 Parallel Computation of Recursive Bounds for Eigenvalues of Matrix Polynomials
Aaron Melman, Santa Clara University, USA

3:40-3:55 Exploring Openmp Task Priorities on the MR3 Eigensolver
Jan Winkelmann and Paolo Bientinesi, RWTH Aachen University, Germany

4:00-4:15 Increasing the Arithmetic Intensity of a Sparse Eigensolver for Better Performance on Heterogenous Clusters
Jonas Thies, German Aerospace Center (DLR), Germany

4:20-4:35 Feast Eigenvalue Solver Using Three Levels of MPI Parallelism
James Kestyn and Eric Polizzi, University of Massachusetts, Amherst, USA; Peter Tang, Intel Corporation, USA

4:40-4:55 Highly Scalable Selective Inversion Using Sparse and Dense Matrix Algebra
Fabio Verbosio and Drosos Kourounis, Università della Svizzera italiana, Switzerland; Olaf Schenk, Università della Svizzera Italiana, Italy

Welcome Remarks

5:05 PM-5:15 PM

Room:Farabeuf

Tuesday, April 12

IP1

Scalability of Sparse Direct Codes

5:15 PM-6:00 PM

Room:Farabeuf

Chair: Sivan A. Toledo, Tel Aviv University, Israel

As part of the H2020 FET-HPC Project NLAFET, we are studying the scalability of algorithms and software for using direct methods for solving large sparse equations. We examine how techniques that work well for dense system solution, for example communication avoiding strategies, can be adapted for the sparse case. We also study the benefits of using standard run time systems to assist us in developing codes for extreme scale computers. We will look at algorithms for solving both sparse symmetric indefinite systems and unsymmetric systems.

Iain Duff
Science & Technology Facilities Council, United Kingdom and CERFACS, Toulouse, France

Concurrent Panels

6:00 PM-7:00 PM

PD1 Towards Exascale Memory Systems

Room: Farabeuf

PD2 What Do You Need to Know About Task-Based Programming for Future Systems?

Room: Pasquier

Welcome Reception

7:00 PM-9:00 PM

Room: Exhibition Space

Wednesday, April 13

Registration

7:45 AM-5:00 PM

Room:Entrance Hall

IP2

Towards Solving Coupled Flow Problems at the Extreme Scale

8:15 AM-9:00 AM

Room:Farabeuf

Chair: Philippe Ricoux, TOTAL SA, France

Scalability is only a necessary prerequisite for extreme scale computing. Beyond scalability, we propose here an architecture aware co-design of the models, algorithms, and data structures. The benefit will be demonstrated for the Lattice Boltzmann method as example of an explicit time stepping scheme and for a finite element multigrid method as an implicit solver. In both cases, systems with more than ten trillion (10^{13}) degrees of freedom can be solved already on the current peta-scale machines.

Ulrich J. Ruede
University of Erlangen-Nuremberg, Germany

Wednesday, April 13

CP3

Data Analysis

9:10 AM-9:50 AM

Room:Pasquier

Chair: Robert Robey, Los Alamos National Laboratory, USA

9:10-9:25 Fast Mesh Operations Using Spatial Hashing: Remaps and Neighbor Finding

Robert Robey, David Nicholaeff, and Gerald Collom, Los Alamos National Laboratory, USA; Redman Colin, Arizona State University and Los Alamos National Laboratory, USA

9:30-9:45 Traversing Time-Dependent Networks in Parallel

Weijian Zhang, University of Manchester, United Kingdom

Wednesday, April 13

CP4

Uncertainty Quantification

9:10 AM-10:10 AM

Room:Roussy

Chair: Michael McKerns, California Institute of Technology, USA

9:10-9:25 "Mystic": Highly-Constrained Non-Convex Optimization and Uncertainty Quantification
Michael McKerns, California Institute of Technology, USA

9:30-9:45 A Resilient Domain Decomposition Approach for Extreme Scale UQ Computations
Paul Mycek and Andres Contreras, Duke University, USA; Olivier P. Le Maître, LIMSI-CNRS, France; Khachik Sargsyan, Francesco Rizzi, Karla Morris, Cosmin Safta, and Bert J. Debusschere, Sandia National Laboratories, USA; Omar M. Knio, Duke University, USA

9:50-10:05 Scalability Studies of Two-Level Domain Decomposition Algorithms for Stochastic Pdes Having Large Random Dimensions
Abhijit Sarkar, *Ajit Desai*, and Mohammad Khalil, Carleton University, Canada; Chris Pettit, United States Naval Academy, USA; Dominique Poirel, Royal Military College, Canada

Wednesday, April 13

CP5

Parallel Software Tools

9:10 AM-10:10 AM

Room:Salle des theses

Chair: Elmar Peise, RWTH Aachen University, Germany

9:10-9:25 The Elaps Framework: Experimental Linear Algebra Performance Studies
Elmar Peise and Paolo Bientinesi, RWTH Aachen University, Germany

9:30-9:45 HiPro: An Interactive Programming Tool for Parallel Programming
Li Liao and Aiqing Zhang, Institute of Applied Physics and Computational Mathematics, China; Zeyao Mo, CAEP Software Center for High Performance Numerical Simulations, China

9:50-10:05 Reliable Bag-of-Tasks Dispatch System
Tariq L. Alturkestani, King Abdullah University of Science and Technology (KAUST), Saudia Arabia and University of Pittsburgh, USA; Esteban Meneses, Costa Rica Institute of Technology, Costa Rica; Rami Melhem, University of Pittsburgh, USA

Wednesday, April 13

PD3

Forward-Looking Panel: Supercomputers Beyond the Exascale

9:10 AM-10:10 AM

Room:Farabeuf

Chair: Bruce Hendrickson, Sandia National Laboratories, USA

Chair: Sivan A. Toledo, Tel Aviv University, Israel

One-hundred petaflops machines are being built today. Exascale machines are clearly feasible and the design of their hardware and software is emerging from the haze of uncertainty. But although current technologies look likely to be able to get us to the exascale, further performance improvements will require solutions to enormous challenges in power, resiliency and other areas. Now is the time to ask what the first 10 exaflops machine might look like. Will it require post-CMOS technologies? Will it exploit reconfigurable hardware (FPGAs)? Perhaps most importantly, how will we program it? Our expert panelists will lead a discussion that will aim to help attendees formulate high-impact long-term research agendas in high-performance scientific computing.

Coffee Break

10:10 AM-10:35 AM

Room:Exhibition Space

Wednesday, April 13

MS17

Parallel Programming Models, Algorithms and Frameworks for Extreme Computing - Part III of III

10:35 AM-12:15 PM

Room:Farabeuf

For Part 2 see MS9

Multicore/manycore 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: Kengo Nakajima
University of Tokyo, Japan

Organizer: Michael Heroux
Sandia National Laboratories, USA

Organizer: Serge G. Petiton
Université Lille 1 and CNRS, France

10:35-10:55 Hierarchical Distributed and Parallel Programming for Auto-Tunned Extreme Computing
Serge G. Petiton, Université Lille 1 and CNRS, France

11:00-11:20 Scalable Sparse Solvers on GPUS
Timothy A. Davis and *Wissam M. Sid-Lakhdar*, Texas A&M University, USA

continued in next column

11:25-11:45 Parallel Generator of Non-Hermitian Matrices with Prescribed Spectrum and Shape Control for Extreme Computing
Herve Galicher, King Abdullah University of Science & Technology (KAUST), Saudi Arabia; France Boillod, CNRS, France; Serge G. Petiton, Université Lille 1 and CNRS, France

11:50-12:10 Monte Carlo Algorithms and Stochastic Process on Manycore Architectures
Christophe Calvin, CEA Saclay, France

Wednesday, April 13

MS18

Challenges in Parallel Adaptive Mesh Refinement - Part I of III

10:35 AM-12:15 PM

Room:Pasquier

For Part 2 see MS26

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 Nonconforming Meshes and Mesh Adaptivity with Petsc
Toby Isaac, University of Texas at Austin, USA; Matthew G. Knepley, Rice University, USA

11:00-11:20 Space-Filling Curves and Adaptive Meshes for Oceanic and Other Applications
Michael Bader, Technische Universität München, Germany

11:25-11:45 A Tetrahedral Space-Filling Curve Via Bitwise Interleaving
Johannes Holke and Carsten Burstedde, Universität Bonn, Germany

11:50-12:10 ADER-DG on Spacetrees in the ExaHyPE Project
Dominic E. Charrier, Durham University, United Kingdom; Angelika Schwarz, Technische Universität München, Germany; Tobias Weinzierl, Durham University, United Kingdom

Wednesday, April 13

MS19

Combinatorial Scientific Computing - Part III of III

10:35 AM-12:15 PM

Room:Roussy

For Part 2 see MS11

Combinatorial algorithms and tools are used in enabling parallel scientific computing applications. The general approach is to identify performance issues in an application and design, analyze, implement combinatorial algorithms to tackle the identified issues. The proposed minisymposium gathers 12 talks, covering applications in bioinformatics, solvers of linear systems, and data analysis; and graph algorithms for those applications. The objective is to summarize the latest combinatorial algorithmic developments and the needs of the applications. The goal is to cross-fertilize the both domains: the applications will raise new challenges to the combinatorial algorithms, and the combinatorial algorithms will address some of the existing problems of the applications.

Organizer: Aydin Buluc
Lawrence Berkeley National Laboratory, USA

Organizer: Alex Pothen
Purdue University, USA

10:35-10:55 Directed Graph Partitioning
Umit V. Catalyurek, The Ohio State University, USA; Kamer Kaya, Sabanci University, Turkey; Bora Ucar, LIP-ENS Lyon, France

11:00-11:20 Parallel Approximation Algorithms for b-Edge Covers and Data Privacy
Arif Khan and Alex Pothen, Purdue University, USA

continued in next column

11:25-11:45 Clustering Sparse Matrices with Information from Both Numerical Values and Pattern
Iain Duff, Science & Technology Facilities Council, United Kingdom and CERFACS, Toulouse, France; Philip Knight, University of Strathclyde, United Kingdom; Sandrine Mouysset, Universite de Toulouse, France; Daniel Ruiz, ENSEEIHT, Toulouse, France; Bora Ucar, LIP-ENS Lyon, France

11:50-12:10 Parallel Graph Coloring on Manycore Architectures
Mehmet Deveci, Erik G. Boman, and Siva Rajamanickam, Sandia National Laboratories, USA

Wednesday, April 13

MS20

Algorithm Comparison, Selection, and Tuning for Large Scale Network and Graph Problems

10:35 AM-12:15 PM

Room:Salle des theses

This minisymposium focuses on using multiple algorithms for analyzing large-scale networks. The topics include methods for selecting among the rapidly growing set of available algorithms and their implementation, comparing the accuracy, performance, and power of different algorithms for static as well as dynamic networks.

Organizer: Boyana Norris
University of Oregon, USA

Organizer: Sanjukta Bhowmick
University of Nebraska, Omaha, USA

10:35-10:55 Parallel Algorithms for Analyzing Dynamic Networks
Sanjukta Bhowmick and *Sriram Srinivasan*, University of Nebraska, Omaha, USA; Sajal Das, Missouri University of Science and Technology, USA

11:00-11:20 The Structure of Communities in Social Networks
Sucheta Soundarajan, Syracuse University, USA

11:25-11:45 Generating Random Hyperbolic Graphs in Subquadratic Time
Moritz von Looz, Henning Meyerhenke, and Roman Prutkin, Karlsruhe Institute of Technology, Germany

11:50-12:10 Performance and Power Characterization of Network Algorithms
Boyana Norris, University of Oregon, USA; Sanjukta Bhowmick, University of Nebraska, Omaha, USA

Wednesday, April 13

MS21

Parallel-in-Time Integration Methods - Part III of III

10:35 AM-12:15 PM

Room:Marie Curie

For Part 2 see MS13

With the expected further increase in available computing power, new parallelisation strategies have to be developed to tackle the scalability challenges arising from the trend towards massively parallel architectures. In association with upcoming Exascale computing, simulations with real-time requirements face particular challenges in HPC. These challenges led to a growing interest in parallel-in-time methods over the last decade. Adding temporal parallelisation has been shown to be able to improve or extend the scalability of programs relying on traditional parallelisation-in-space methods. We invite speakers to present their current research progress related to parallel-in-time methods from all areas.

Organizer: Martin Schreiber
University of Exeter, United Kingdom

Organizer: Tobias Neckel
Technische Universität München, Germany

Organizer: Daniel Ruprecht
University of Leeds, United Kingdom

10:35-10:55 Improving Applicability and Performance for the Multigrid Reduction in Time (mgrit) Algorithm
Robert D. Falgout, Veselin Dobrev, and Tzanio V. Kolev, Lawrence Livermore National Laboratory, USA; Ben O'Neill, University of Colorado Boulder, USA; Jacob B. Schroder, Lawrence Livermore National Laboratory, USA; Ben Southworth, University of Colorado Boulder, USA; Ulrike Meier Yang, Lawrence Livermore National Laboratory, USA

11:00-11:20 Multigrid Methods with Space-Time Concurrency
Stephanie Friedhoff, University of Cologne, Germany; Robert D. Falgout and Tzanio V. Kolev, Lawrence Livermore National Laboratory, USA; Scott Maclachlan, Memorial University, Newfoundland, Canada; Jacob B. Schroder, Lawrence Livermore National Laboratory, USA; Stefan Vandewalle, KU Leuven, Belgium

11:25-11:45 PinT for Climate and Weather Simulation
Martin Schreiber, University of Exeter, United Kingdom; Pedro Peixoto, University of Sao Paulo, Brazil; Terry Haut, Los Alamos National Laboratory, USA; Beth Wingate, University of Exeter, United Kingdom

11:50-12:10 Convergence Rate of Parareal Method with Modified Newmark-Beta Algorithm for 2nd-Order Ode
Mikio Iizuka and Kenji Ono, RIKEN, Japan

continued in next column

Wednesday, April 13

MS22

The Strong Scaling Limit: From Industrial Applications to New Algorithms - Part III of III

10:35 AM-12:15 PM

Room:Leroux

For Part 2 see MS14

There is a strong interest in running parallel applications in the strong scaling limit. There, however, there is an increasing communication cost and a decreasing compute cost. In this minisymposium, we analyse how various applications behave in this limit and how novel communication hiding/avoiding algorithms can help. We start with industrial and scientific applications and discuss both the software and algorithmic challenges. This symposium reports work of the Exa2ct project, a EU supported project.

Organizer: Wim I. Vanroose
University of Antwerp, Belgium

10:35-10:55 Increasing Arithmetic Intensity Using Stencil Compilers.
Simplice Donfack and Patrick Sanan, Università della Svizzera italiana, Switzerland; Olaf Schenk, Università della Svizzera Italiana, Italy; Reps Bram and Wim Vanroose, Antwerp University, Belgium

11:00-11:20 Toward Exascale Qp Solvers
D. Horak, Technical University Ostrava, Czech Republic

11:25-11:45 Many Core Acceleration of the Boundary Element Library BEM4I
Michal Merta, Technical University Ostrava, Czech Republic; Jan Zapletal, VSB-Technical University Ostrava, Czech Republic

11:50-12:10 HPGMG-FV Benchmark with GASPI/GPI-2: First Steps
Dimitar M. Stoyanov, Fraunhofer Institute for Industrial Mathematics, Germany

Wednesday, April 13

MS23

Extreme Scale Implicit PDE Solvers: Parallel Algorithms and Applications

10:35 AM-12:15 PM

Room:Dejerine

Many PDE problems in computational science and engineering are characterized by highly nonlinear phenomena, a wide range of length and time scales, and extreme variability in material properties. These in turn often necessitate aggressive adaptivity and advanced discretizations. Achieving scalability to the extreme core counts that are necessary for problems with such complexity often requires novel implementation strategies as well as new algorithmic developments. This minisymposium bring together researchers who have developed solvers capable of scaling to $O(10^5)$ cores and have demonstrated them on challenging problems in mantle convection, subsurface flows, ice sheet dynamics, materials processing, and others.

Organizer: Toby Isaac
University of Texas at Austin, USA

Organizer: Omar Ghattas
University of Texas at Austin, USA

Organizer: Georg Stadler
Courant Institute of Mathematical Sciences, New York University, USA

Organizer: Johann Rudi
University of Texas at Austin, USA

10:35-10:55 Extreme-Scale Solver for Earth's Mantle -- Spectral/Geometric/Algebraic Multigrid Methods for Nonlinear, Heterogeneous Stokes Flow
Johann Rudi, University of Texas at Austin, USA; A. Cristiano I. Malossi, IBM Research-Zurich, Switzerland; Tobin Isaac, University of Chicago, USA; Georg Stadler, Courant Institute of Mathematical Sciences, New York University, USA; Michael Gurnis, California Institute of Technology, USA; Peter W. J. Staar, Yves Ineichen, Costas Bekas, and Alessandro Curioni, IBM Research-Zurich, Switzerland; Omar Ghattas, University of Texas at Austin, USA

11:00-11:20 FE2TI - Computational Scale Bridging for Dual-Phase Steels
Martin Lanser, University of Cologne, Germany; Axel Klawonn, Universitaet zu Koeln, Germany; Oliver Rheinbach, Technische Universitaet Bergakademie Freiberg, Germany

11:25-11:45 3D Parallel Direct Elliptic Solvers Exploiting Hierarchical Low Rank Structure
Gustavo Chavez, Hatem Ltaief, *George M Turkiyyah*, George M Turkiyyah, and David E. Keyes, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

11:50-12:10 Title Not Available
Steffen Muething, Universität Heidelberg, Germany

continued in next column

Wednesday, April 13

MS24

Massively Parallel Phase-field Simulations of Phase Transformation Processes

10:35 AM-12:15 PM

Room:Delarue

The minisymposium focuses on massively parallel phase-field simulations to study pattern formations during solidification or solid-solid phase transformations in alloys. The various material and process parameters influence the evolving microstructure and hence the macroscopic material behavior. To minimize boundary effects and obtain statistically reliable results, simulations on large domain sizes are necessary. The computational study of phase transformation processes in real systems under real processing conditions requires the implementation of highly parallel algorithms. This minisymposium shows the usage of highly optimized and parallel methods for engineering applications related to material development and microstructure design.

Organizer: Harald Köstler
Universität Erlangen-Nürnberg, Germany

Organizer: Martin Bauer
University of Erlangen-Nuremberg, Germany

10:35-10:55 Performance Optimization of a Massively Parallel Phase-Field Method Using the Hpc Framework Walberla
Martin Bauer, University of Erlangen-Nuremberg, Germany; Harald Köstler, Universität Erlangen-Nürnberg, Germany; Ulrich Rüde, University of Erlangen-Nuremberg, Germany

11:00-11:20 Large Scale phase-field simulations of directional ternary eutectic solidification
Johannes Hötzer, Hochschule Karlsruhe Technik und Wirtschaft, Germany; Philipp Steinmetz and Britta Nestler, Karlsruhe Institute of Technology, Germany

continued on next page

11:25-11:45 Building High Performance Mesoscale Simulations using a Modular Multiphysics Framework
Cody Permann, Idaho National Laboratory, USA

11:50-12:10 Equilibrium and Dynamics in Multiphase-Field Theories: A Comparative Study and a Consistent Formulation
Gyula I. Toth, University of Bergen, Norway; Laszlo Granasy and Tamas Pusztai, Wignes Research Centre for Physics, Hungary; Bjorn Kvamme, University of Bergen, Norway

Wednesday, April 13

CP6

Matrix Factorization

10:35 AM-12:15 PM

Room:Danton

Chair: F. Sukru Torun, Bilkent University, Turkey

10:35-10:50 Solving Sparse Underdetermined Linear Least Squares Problems on Parallel Computing Platforms
F. Sukru Torun, Bilkent University, Turkey; Murat Manguoglu, Middle East Technical University, Turkey; Cevdet Aykanat, Bilkent University, Turkey

10:55-11:10 A New Approximate Inverse Technique Based on the GLT Theory
Ali Dorostkar, Uppsala University, Sweden

11:15-11:30 Sparse Matrix Factorization on GPUs
Mohamed Gadou, University of Florida, USA; Tim Davis, Texas A&M University, USA; Sanjay Ranka, University of Florida, USA

11:35-11:50 Locality-Aware Parallel Sparse Matrix-Matrix Multiplication
Emanuel H. Rubensson and Elias Rudberg, Uppsala University, Sweden

11:55-12:10 Designing An Efficient and Scalable Block Low-Rank Direct Solver for Large Scale Clusters
Xavier Lacoste, Cédric Augonnet, and David Goudin, CEA/DAM, France

Lunch Break

12:15 PM-1:45 PM

Attendees on their own

Wednesday, April 13

IP3

Why Future Systems should be Data Centric

1:45 PM-2:30 PM

Room:Farabeuf

Chair: John Shalf, Lawrence Berkeley National Laboratory, USA

We are immersed in data. Classical high performance computing often generates enormous structured data through simulation; but, a vast amount of unstructured data is also being generated and is growing at exponential rates. To meet the data challenge, we expect computing to evolve in two fundamental ways: first through a focus on workflows, and second to explicitly accommodate the impact of big data and complex analytics. As an example, HPC systems should be optimized to perform well on modeling and simulation; but, should also focus on other important elements of the overall workflow which include data management and manipulation coupled with associated cognitive analytics. At a macro level, workflows will take advantage of different elements of the systems hierarchy, at dynamically varying times, in different ways and with different data distributions throughout the hierarchy. This leads us to a data centric design point that has the flexibility to handle these data-driven demands. I will explain the attributes of a data centric system and discuss challenges we see for future computing systems. It is highly desirable that the data centric architecture and implementation lead to commercially viable Exascale-class systems on premise and in the cloud.

Paul Coteus
IBM Research, USA

Intermission

2:30 PM-2:40 PM

Wednesday, April 13

MS25

Middleware in the Era of Extreme Scale Computing - Part I of II

2:40 PM-4:20 PM

Room:Farabeuf

For Part 2 see MS33

The unprecedented complexity of future exascale systems will make the role of resource management increasingly important for allocating system resources to users at policy-driven levels of quality of service while meeting power and reliability requirements of the entire system. The key challenge lies in maintaining scalability while managing and monitoring allocations targeting multiple, often conflicting optimization criteria. In this minisymposium, we will cover several R&D activities on HPC resource management systems both in academia and industry. In particular, we will discuss their design and implementation to identify how to improve the sustained throughput, system availability and power efficiency at scale.

Organizer: Keita Teranishi
Sandia National Laboratories, USA

Organizer: Martin Schulz
Lawrence Livermore National Laboratory, USA

2:40-3:00 Middleware for Efficient Performance in HPC
Larry Kaplan, Cray, Inc., USA

3:05-3:25 {PMI}-Exascale: Enabling Applications at Exascale
Ralph Castain, Intel Corporation, USA

3:30-3:50 Co-Scheduling: Increasing Efficency for HPC
Josef Weidendorfer, Technische Universität München, Germany

3:55-4:15 Power Aware Resource Management
Tapasya Patki, Lawrence Livermore National Laboratory, USA

Wednesday, April 13

MS26

Challenges in Parallel Adaptive Mesh Refinement - Part II of III

2:40 PM-4:20 PM

Room:Pasquier

For Part 1 see MS18
For Part 3 see MS34

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 Recent Parallel Results Using Dynamic Quad-Tree Refinement for Solving Hyperbolic Conservation Laws on 2d Cartesian Meshes
Donna Calhoun, Boise State University, USA

continued in next column

3:05-3:25 Basilisk: Simple Abstractions for Octree-Adaptive Schemes
Stephane Popinet, CNRS / UPMC, France

3:30-3:50 Fast Solvers for Complex Fluids
George Biros and Dhairya Malhotra, University of Texas at Austin, USA

3:55-4:15 Simulation of Incompressible Flows on Parallel Octree Grids
Arthur Guittet and Frederic G. Gibou, University of California, Santa Barbara, USA

Wednesday, April 13

MS27

To Thread or Not To Thread - Part I of II

2:40 PM-3:55 PM

Room:Roussy

For Part 2 see MS35

Emerging architectures, in particular the large, hybrid machines now being built by the Department of Energy, will have smaller local memory/core, more cores/socket, multiple types of user-addressable memory, increased memory latencies -- especially to global memory -- and expensive synchronization. Two popular strategies, MPI+threads and flat MPI, need to be contrasted in terms of performance metrics, code complexity and maintenance burdens, and interoperability between libraries, applications, and operating systems. This session will present talks aimed at clarifying the tradeoffs involved for developers and users.

Organizer: Matthew G. Knepley
Rice University, USA

Organizer: Barry F. Smith
Argonne National Laboratory, USA

Organizer: Richard T. Mills
Intel Corporation, USA

Organizer: Jed Brown
Argonne National Laboratory and University of Colorado Boulder, USA

2:40-3:00 Tradeoffs in Domain Decomposition: Halos Vs Shared Contiguous Arrays
Jed Brown, Argonne National Laboratory and University of Colorado Boulder, USA

3:05-3:25 Firedrake: Burning the Thread at Both Ends
Gerard J Gorman, Imperial College London, United Kingdom

3:30-3:50 Threads: the Wrong Abstraction and the Wrong Semantic
Jeff R. Hammond and Timothy Mattson, Intel Corporation, USA

Wednesday, April 13

MS28

Extreme-Scale Linear Solvers - Part I of II

2:40 PM-4:20 PM

Room:Salle des theses

For Part 2 see MS36

This minisymposium highlights the challenges of solving sparse linear systems at extreme scale. Both scalable algorithms and approaches that can exploit manycore and massively parallel architectures will be discussed. Topics of interest include (but is not limited to) hierarchical solvers, low-rank compression, fine-grain parallelism, asynchronous iterative methods, random sampling, and communication-avoiding methods.

Organizer: Erik G. Boman
Sandia National Laboratories, USA

Organizer: Siva Rajamanickam
Sandia National Laboratories, USA

2:40-3:00 Recent Directions in Extreme-Scale Solvers
Erik G. Boman, Sandia National Laboratories, USA

3:05-3:25 Preconditioning Using Hierarchically Semi-Separable Matrices and Randomized Sampling
Pieter Ghysels, Francois-Henry Rouet, and Xiaoye Sherry Li, Lawrence Berkeley National Laboratory, USA

3:30-3:50 Preconditioning Communication-Avoiding Krylov Methods
Ichitaro Yamazaki, University of Tennessee, Knoxville, USA; Siva Rajamanickam, Andrey Prokopenko, Erik G. Boman, Mark Hoemmen, and Michael Heroux, Sandia National Laboratories, USA; Stan Tomov, University of Tennessee, USA; Jack J. Dongarra, University of Tennessee and Oak Ridge National Laboratory, USA

3:55-4:15 A New Parallel Hierarchical Preconditioner for Sparse Matrices
Chao Chen and Eric F. Darve, Stanford University, USA; Siva Rajamanickam and Erik G. Boman, Sandia National Laboratories, USA

Wednesday, April 13

MS29

The Present and Future Bulk Synchronous Parallel Computing - Part I of III

2:40 PM-4:20 PM

Room:Marie Curie

For Part 2 see MS37

The Bulk Synchronous Parallel (BSP) bridging model, proposed by Valiant in 1990, has had great influence on algorithm and software design over the past decades. It has inspired programming frameworks and structured parallel programming of communication minimising algorithms. In the current era of post-petascale HPC and of big data computing, BSP has again led to successful programming frameworks such as MapReduce, Pregel, and Giraph. This minisymposium explores the use of BSP in contemporary high performance and big data computing, highlighting new BSP algorithms and frameworks, new BSP-like models, and BSP-inspired auto-generation and verification tools.

Organizer: Albert-Jan N. Yzelman
Huawei Technologies, France

Organizer: Rob H. Bisseling
Utrecht University, The Netherlands

2:40-3:00 Towards Next-Generation High-Performance BSP
Albert-Jan N. Yzelman, Huawei Technologies, France

3:05-3:25 Computing Almost Asynchronously: What Problems Can We Solve in *O(log P)* Supersteps?
Alexander Tiskin, University of Warwick, United Kingdom

3:30-3:50 Deductive Verification of BSP Algorithms with Subgroups
Frédéric Gava, Universite de Paris-Est, France

3:55-4:15 A BSP Scalability Analysis of 3D LU Factorization with Row Partial Pivoting
Antoine Petitet, Huawei Technologies, France

Wednesday, April 13

MS30

Scientific Computing towards Ultrascale Systems - Part I of II

2:40 PM-4:20 PM

Room:Leroux

For Part 2 see MS38

Ultrascale computing systems are large-scale complex computing systems combining technologies from high performance computing, distributed systems, and cloud computing. In this minisymposium we discuss challenges and experiences with the development of applications for Ultrascale systems. The speakers are involved in the EU COST action 'Network for Sustainable Ultrascale Computing' (NESUS) and will address design and performance aspects of algorithms and applications from scientific computing for Ultrascale systems, including numerical and computational efficiency, energy aspects as well as implementations for heterogeneous architectures. The minisymposium is planned to consist of two parts.

Organizer: Svetozar Margenov
Bulgarian Academy of Science, Bulgaria

Organizer: Maya Neytcheva
Uppsala University, Sweden

Organizer: Thomas Rauber
Universität Bayreuth, Germany

Organizer: Gudula Rünger
Chemnitz University of Technology, Germany

2:40-3:00 Simulations for Optimizing Lightweight Structures on Heterogeneous Compute Resources

Robert Dietze, Michael Hofmann, and *Gudula Rünger*, Chemnitz University of Technology, Germany

continued in next column

3:05-3:25 Balanced Dense-Sparse Parallel Solvers for Nonlinear Vibration Analysis of Elastic Structures

Svetozar Margenov, Bulgarian Academy of Science, Bulgaria

3:30-3:50 High Performance Computing in Micromechanics

Radim Blaheta, O. Jakl, and J. Stary, Academy of Sciences of the Czech Republic, Ostrava, Czech Republic; I. Georgiev, Institute of Information Technologies, Bulgaria

3:55-4:15 Parallel Block Preconditioning Framework and Efficient Iterative Solution of Discrete Multi-Physics Problems

Milan D. Mihajlovic, University of Manchester, United Kingdom

Wednesday, April 13

MS31

Novel Educational Programs for CSE and Big Data

2:40 PM-4:20 PM

Room:Dejerine

People skilled in the principles of Computational Science and Engineering (CSE) and Big Data are the most important asset in bringing these technologies to bear on science and industry. We showcase four different educational programs that in various ways are all different from standard faculty-oriented university degree programs in order to address the cross-disciplinary and practical requirements of computational and data-driven science. Novelties include a blend of Germany's vocational training concept with academic instruction (Aachen), a strong project orientation with business participation (Maastricht), the tie-in of non-university research institutions (Aachen), or a strong algorithmic orientation (Jena).

Organizer: Christian H. Bischof
Technische Universität Darmstadt, Germany

2:40-3:00 Exploiting Synergies with Non-University Partners in CSE Education

Marek Behr, RWTH Aachen University, Germany

3:05-3:25 Supporting Data-Driven Businesses Through Talented Students in KnowledgeEngineering@Work

Carlo Galuzzi, Maastricht University, The Netherlands

3:30-3:50 Blending the Vocational Training and University Education Concepts in the Mathematical Technical Software Developer (matse) Program

Benno Willemsen, RWTH Aachen University, Germany

3:55-4:15 Designing a New Master's Degree Program in Computational and Data Science

Martin Buecker, Friedrich Schiller Universität Jena, Germany

Wednesday, April 13

MS32

High Performance Computing in Optimization

2:40 PM-4:20 PM

Room:Delarue

High performance computing (HPC) becomes increasingly relevant in solving large-scale optimization problems in many applications areas such as operations research, energy systems, industrial engineering, advanced manufacturing, and others. Despite being in its infancy, HPC optimization was successfully used to solve problems of unprecedented sizes. The mainstream computational approaches are parallelization of linear algebra, optimization decomposition algorithms and branch-and-bound searches. This minisymposium on HPC optimization will host a series of presentations on parallel methods and software implementations and is aimed at fostering interdisciplinary collaborations between the optimization and parallel processing communities.

Organizer: Kibaek Kim
Argonne National Laboratory, USA

Organizer: Cosmin G. Petra
Argonne National Laboratory, USA

2:40-3:00 Parallel Decomposition Methods for Structured MIP
Kibaek Kim, Argonne National Laboratory, USA

3:05-3:25 Parallel Optimization of Complex Energy Systems on High-Performance Computing Platforms
Cosmin G. Petra, Argonne National Laboratory, USA

continued in next column

3:30-3:50 PIPS-SBB: A Parallel Distributed Memory Solver for Stochastic Mixed-Integer Optimization
Deepak Rajan, Lawrence Livermore National Laboratory, USA; Lluis Munguia, Georgia Institute of Technology, USA; Geoffrey M. Oxberry, Lawrence Livermore National Laboratory, USA; Cosmin G. Petra, Argonne National Laboratory, USA; Pedro Sotorrio and Thomas Edmunds, Lawrence Livermore National Laboratory, USA

3:55-4:15 How to Solve Open MIP Instances by Using Supercomputers with Over 80,000 Cores
Yuji Shinano, Zuse Institute Berlin, Germany; Tobias Achterberg, Gurobi GmbH, Germany; Timo Berthold and Stefan Heinz, Fair Issac Europe Ltd, Germany; Thorsten Koch, Zuse Institute Berlin, Germany; Stefan Vigerske, GAMS Software GmbH, Germany; Michael Winkler, Gurobi GmbH, Germany

Wednesday, April 13

CP7

Preconditioning and Multigrid

2:40 PM-4:20 PM

Room:Danton

Chair: Paul Lin, Sandia National Laboratories, USA

2:40-2:55 Towards Exascalable Algebraic Multigrid
Amani Alonazi and David E. Keyes, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

3:00-3:15 A Matrix-Free Preconditioner for Elliptic Solvers Based on the Fast Multipole Method
Huda Ibeid, King Abdullah University of Science & Technology (KAUST), Saudi Arabia; Rio Yokota, Tokyo Institute of Technology, Japan; David E. Keyes, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

3:20-3:35 Efficiency Through Reuse in Algebraic Multigrid
Andrey Prokopenko and Jonathan J. Hu, Sandia National Laboratories, USA

3:40-3:55 A Dynamical Preconditioning Strategy and Its Applications in Large-Scale Multi-Physics Simulations
Xiaowen Xu, Institute of Applied Physics and Computational Mathematics, China

4:00-4:15 Performance of Smoothers for Algebraic Multigrid Preconditioners for Finite Element Variational Multiscale Incompressible Magnetohydrodynamics
Paul Lin and John Shadid, Sandia National Laboratories, USA; Paul Tsuji, Lawrence Livermore National Laboratory, USA; Jonathan J. Hu, Sandia National Laboratories, USA

Coffee Break

4:20 PM-4:50 PM

Room:Exhibition Space

Wednesday, April 13

MS33

Middleware in the Era of Extreme Scale Computing - Part II of II

4:50 PM-6:05 PM

Room:Farabeuf

For Part 1 see MS25

The unprecedented complexity of future exascale systems will make the role of resource management increasingly important for allocating system resources to users at policy-driven levels of quality of service while meeting power and reliability requirements of the entire system. The key challenge lies in maintaining scalability while managing and monitoring allocations targeting multiple, often conflicting optimization criteria. In this minisymposium, we will cover several R&D activities on HPC resource management systems both in academia and industry. In particular, we will discuss their design and implementation to identify how to improve the sustained throughput, system availability and power efficiency at scale.

Organizer: Keita Teranishi
Sandia National Laboratories, USA

Organizer: Martin Schulz
Lawrence Livermore National Laboratory, USA

4:50-5:10 Operation Experiences of the K Computer
Atsuya Uno, RIKEN Advanced Institute for Computational Science, Japan

5:15-5:35 Improving Resiliency Through Dynamic Resource Management
Suraj Prabhakaran and Felix Wolf, Technische Universität Darmstadt, Germany

5:40-6:00 The Influence of Data Center Power and Energy Constraints on Future HPC Scheduling Systems
Torsten Wilde, LRZ: Leibniz Supercomputing Centre, Germany

Wednesday, April 13

MS34

Challenges in Parallel Adaptive Mesh Refinement - Part III of III

4:50 PM-6:30 PM

Room:Pasquier

For Part 2 see MS26

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 A Simple Robust and Accurate a Posteriori Subcell Finite Volume Limiter for the Discontinuous Galerkin Method on Space-Time Adaptive Meshes
Olindo Zanotti and Michael Dumbser, University of Trento, Italy

5:15-5:35 Adaptive Mesh Refinement in Radiation Problems
Alan Humphrey and Todd Harman, University of Utah, USA; Daniel Sunderland, Sandia National Laboratories, USA; Martin Berzins, University of Utah, USA

5:40-6:00 Resiliency in AMR Using Rollback and Redundancy
Anshu Dubey, Argonne National Laboratory, USA

6:05-6:25 Resolving a Vast Range of Spatial and Dynamical Scales in Cosmological Simulations
Kevin Schaal and Volker Springel, Heidelberg Institute for Theoretical Studies and Karlsruhe Institute of Technology, Germany

continued in next column

Wednesday, April 13

MS35

To Thread or Not To Thread - Part II of II

4:50 PM-6:30 PM

Room:Roussy

For Part 1 see MS27

Emerging architectures, in particular the large, hybrid machines now being built by the Department of Energy, will have smaller local memory/core, more cores/socket, multiple types of user-addressable memory, increased memory latencies -- especially to global memory -- and expensive synchronization. Two popular strategies, MPI+threads and flat MPI, need to be contrasted in terms of performance metrics, code complexity and maintenance burdens, and interoperability between libraries, applications, and operating systems. This session will present talks aimed at clarifying the tradeoffs involved for developers and users.

Organizer: Matthew G. Knepley
Rice University, USA

Organizer: Barry F. Smith
Argonne National Laboratory, USA

Organizer: Richard T. Mills
Intel Corporation, USA

Organizer: Jed Brown
Argonne National Laboratory and University of Colorado Boulder, USA

4:50-5:10 MPI+MPI: Using MPI-3 Shared Memory As a MultiCore Programming System
William D. Gropp, University of Illinois at Urbana-Champaign, USA

5:15-5:35 MPI+X: Opportunities and Limitations for Heterogeneous Systems
Karl Rupp, Josef Weinbub, Florian Rudolf, Andreas Morhammer, Tibor Grasser, and Ansgar Juengel, Vienna University of Technology, Austria

5:40-6:00 How I Learned to Stop Worrying and Love {MPI}: Prospering with an {MPI}-Only Model in {PFLOTRAN} + {PETSc} in the Manycore Era
Richard T. Mills, Intel Corporation, USA

6:05-6:25 The Occa Abstract Threading Model: Implementation and Performance for High-Order Finite-Element Computations
Tim Warburton, Virginia Tech, USA; David Medina, Rice University, USA

continued in next column

Wednesday, April 13

MS36

Extreme-Scale Linear Solvers - Part II of II

4:50 PM-6:30 PM

Room:Salle des theses

For Part 1 see MS28

This minisymposium centers on the challenges of solving linear systems at extreme scale. Both scalable algorithms and approaches that can exploit manycore and massively parallel architectures will be discussed. Topics of interest include (but is not limited to) hierarchical solvers, low-rank compression, fine-grain parallelism, asynchronous iterative methods, random sampling, and communication-avoiding methods.

Organizer: Erik G. Boman
Sandia National Laboratories, USA

Organizer: Siva Rajamanickam
Sandia National Laboratories, USA

4:50-5:10 Task and Data Parallelism Based Direct Solvers and Preconditioners in Manycore Architectures
Siva Rajamanickam, Erik G. Boman, *Joshua D. Booth*, and Kyungjoo Kim, Sandia National Laboratories, USA

5:15-5:35 Parallel Ilu and Iterative Sparse Triangular Solves
Thomas K. Huckle and Juergen Braeckle, Technische Universität München, Germany

5:40-6:00 Aggregation-Based Algebraic Multigrid on Massively Parallel Computers
Yvan Notay, Université Libre de Bruxelles, Belgium

6:05-6:25 Revisiting Asynchronous Linear Solvers: Provable Convergence Rate Through Randomization
Haim Avron, IBM T.J. Watson Research Center, USA; Alex Druinsky, Lawrence Berkeley National Laboratory, USA; Anshul Gupta, IBM T.J. Watson Research Center, USA

Wednesday, April 13

MS37

The Present and Future Bulk Synchronous Parallel Computing - Part II of III

4:50 PM-6:30 PM

Room:Marie Curie

For Part 1 see MS29
For Part 3 see MS45

The Bulk Synchronous Parallel (BSP) bridging model, proposed by Valiant in 1990, has had great influence on algorithm and software design over the past decades. It has inspired programming frameworks and structured parallel programming of communication minimising algorithms. In the current era of post-petascale HPC and of big data computing, BSP has again led to successful programming frameworks such as MapReduce, Pregel, and Giraph. This minisymposium explores the use of BSP in contemporary high performance and big data computing, highlighting new BSP algorithms and frameworks, new BSP-like models, and BSP-inspired auto-generation and verification tools.

Organizer: Albert-Jan N. Yzelman
Huawei Technologies, France

Organizer: Rob H. Bisseling
Utrecht University, The Netherlands

4:50-5:10 From Cilk to MPI to BSP: The Value of Structured Parallel Programming
Gaetan Hains, Huawei Technologies, France

5:15-5:35 Systematic Development of Functional Bulk Synchronous Parallel Programs
Frédéric Loulergue, Universite d'Orleans, France

5:40-6:00 Scalable Parallel Graph Matching for Big Data
Rob H. Bisseling, Utrecht University, The Netherlands

6:05-6:25 Parallel Data Distribution for a Sparse Eigenvalue Problem Based on the BSP Model
Dirk Roose, Karl Meerbergen, and Joris Tavernier, KU Leuven, Belgium

Wednesday, April 13

MS38

Scientific Computing towards Ultrascale Systems - Part II of II

4:50 PM-6:30 PM

Room:Leroux

For Part 1 see MS30

Ultrascale computing systems are large-scale complex computing systems combining technologies from high performance computing, distributed systems, and cloud computing. In this minisymposium we discuss challenges and experiences with the development of applications for Ultrascale systems. The speakers are involved in the EU COST action 'Network for Sustainable Ultrascale Computing' (NESUS) and will address design and performance aspects of algorithms and applications from scientific computing for Ultrascale systems, including numerical and computational efficiency, energy aspects as well as implementations for heterogeneous architectures. The minisymposium is planned to consist of two parts.

Organizer: Svetozar Margenov
Bulgarian Academy of Science, Bulgaria

Organizer: Maya Neytcheva
Uppsala University, Sweden

Organizer: Thomas Rauber
Universität Bayreuth, Germany

Organizer: Gudula Rünger
Chemnitz University of Technology, Germany

4:50-5:10 Parallelization of Numerical Methods for Time Evolution Processes, Described by Integro-Differential Equations
Ali Dorostkar and *Maya Neytcheva*, Uppsala University, Sweden

continued in next column

5:15-5:35 Optimizing the Use of GPU and Intel Xeon Phi Clusters to Accelerate Stencil Computations
Roman Wyrzykowski, Krzysztof Rojek, and Lukasz Szustak, Czestochowa University of Technology, Poland

5:40-6:00 Efficient Parallel Algorithms for New Unidirectional Models of Nonlinear Optics
Raimondas Ciegas, University Vilnius, Lithuania

6:05-6:25 Online Auto-Tuning for Parallel Solution Methods for ODEs
Thomas Rauber, Universität Bayreuth, Germany; Natalia Kalinnik and Matthias Korch, University Bayreuth, Germany

Wednesday, April 13

MS39

Parallel Methodologies for Large Data Assimilation and Inverse Problems - Part I of II

4:50 PM-6:30 PM

Room:Dejerine

Data assimilation is the process of fusing information from priors, models, and observations, in order to obtain best estimates of the state of a physical system such as the atmosphere. Data assimilation relies on comprehensive physical models with large numbers of variables, and on complex observational data sets. The resulting large scale inverse problems need to be solved in faster than real time in order to, for example, meet the demands of weather forecasting. This minisymposium discusses the latest developments in parallel methodologies for solving large data assimilation problems.

Organizer: Adrian Sandu

Virginia Polytechnic Institute & State University, USA

Organizer: Marc Bocquet

Ecole Nationale des Ponts et Chaussées, France

4:50-5:10 Parallel Implementations of Ensemble Kalman Filters for Huge Geophysical Models

Jeffrey Anderson, Helen Kershaw, Jonathan Hendricks, and Nancy Collins, National Center for Atmospheric Research, USA

5:15-5:35 Parallel 4D-Var and Particle Smoother Data Assimilation for Atmospheric Chemistry and Meteorology

Hendrik Elbern, University of Cologne, Germany; Jonas Berndt and Philipp Franke, Forschungszentrum Jülich, Germany

5:40-6:00 Parallel in Time Strong-Constraint 4D-Var

Adrian Sandu, Virginia Polytechnic Institute & State University, USA

6:05-6:25 Development of 4D Ensemble Variational Assimilation at Météo-France

Gerald Desroziers, Meteo, France

Wednesday, April 13

MS40

High Performance Computing for Astronomical Instrumentation

4:50 PM-6:30 PM

Room:Delarue

Due to the always increasing complexity of astronomical instruments, research in high performance computing has become a major vector of scientific discovery. Indeed, the design of instruments for large telescopes requires efficient numerical simulations of observing conditions and instrument components. The latter can include large computing facilities for operations and data storage and the extraction of observables from these data also requires significant computation capabilities. We will cover most of these aspects through examples of use of parallel processing to design instruments, operate them and process the obtained data in the context of the largest projects in astronomy, namely the European Extremely Large Telescope and the Square Kilometer Array.

Organizer: Damien Gratadour

Observatoire de Paris, Meudon, France

4:50-5:10 Adaptive Optics Simulation for the European Extremely Large Telescope on Multicore Architectures with Multiple Gpus

Hatem Ltaief, King Abdullah University of Science & Technology (KAUST), Saudi Arabia; Damien Gratadour, Observatoire de Paris, Meudon, France; Ali Charara, King Abdullah University of Science & Technology (KAUST), Saudi Arabia; Eric Gendron, Observatoire de Paris, Meudon, France; David E. Keyes, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

5:15-5:35 High Performance Computing for Astronomical Instrumentation on Large Telescopes

Nigel Dipper, Durham University, United Kingdom

5:40-6:00 Parallel Processing for the Square Kilometer Array Telescope

David De Roure, University of Oxford, United Kingdom

6:05-6:25 The SKA (Square Kilometre Array) and Its Computational Challenges

Stefano Salvini, University of Oxford, United Kingdom

continued in next column

Wednesday, April 13

CP8

Numerical Optimization

4:50 PM-5:50 PM

Room:Danton

Chair: Scott P. Kolodziej, Texas A&M University, USA

4:50-5:05 Vertex Separators with Mixed-Integer Linear Optimization
Scott P. Kolodziej and Tim Davis, Texas A&M University, USA

5:10-5:25 Large-Scale Parallel Combinatorial Optimization -- the Role of Estimation and Ordering the Sub-Problems
Bogdan Zavalnij and Sandor Szabo, University of Pécs, Hungary

5:30-5:45 Parallel Linear Programming Using Zonotopes
Max Demenkov, Institute of Control Sciences, Russia

Business Meeting

6:30 PM-7:15 PM

Room:Farabeuf

Wednesday, April 13

PP1

Poster Session and Reception

7:15 PM-9:15 PM

Room:Exhibition Space

Improving Scalability of 2D Sparse Matrix Partitioning Models Via Reducing Latency Cost
Cevdet Aykanat and Oguz Selvitopi, Bilkent University, Turkey

Hybrid Parallel Simulation of Helicopter Rotor Dynamics
Melven Roehrig-Zoellner, *Achim Basermann*, and Johannes Hofmann, German Aerospace Center (DLR), Germany

Poisson Solver for large-scale Simulations with one Fourier Diagonalizable Direction
Ricard Borrell, Guillermo Oyarzun, Jorje Chiva, Oriol Lehmkuhl, Ivette Rodríguez, and Assensi Oliva, Universitat Politecnica de Catalunya, Spain

High Performance Algorithms and Software for Seismic Simulations
Alexander Breuer, University of California, San Diego, USA; Alexander Heinecke, Intel Corporation, USA; Michael Bader, Technische Universität München, Germany

Large Scale Parallel Computations in R Through Elemental
Rodrigo Canales, Elmar Peise, and Paolo Bientinesi, RWTH Aachen University, Germany

Parallel Programming Using Hpcmatlab
Mukul Dave, Arizona State University, USA; Xinchen Guo, Purdue University, USA; Mohamed Sayeed, Arizona State University, USA

Using the Verificarlo Tool to Assess the Floating Point Accuracy of Large Scale Digital Simulations
Christophe Denis, CMLA, ENS de Cachan, France; Pablo De Oliveira Castro and Eric Petit, PRiSM - UVSQ, France

Chase: A Chebyshev Accellerated Subspace Iteration Eigensolver
Edoardo A. Di Napoli, Jülich Supercomputing Centre, Germany; Mario Berljafa, University of Manchester, United Kingdom; Paul Springer and Jan Winkelmann, RWTH Aachen University, Germany

A Parallel Fast Sweeping Method on Nonuniform Grids
Miles L. Detrixhe, University of California, Santa Barbara, USA

Evaluation of the Intel Xeon Phi and Nvidia K80 As Accelerators for Two-Dimensional Panel Codes
Lukas Einkemmer, University of Innsbruck, Austria

An Impact of Tuning the Kernel of the Structured QR Factorization in the TSQR
Takeshi Fukaya, Hokkaido University, Japan; Toshiyuki Imamura, RIKEN, Japan

Filters of the Improvement of Multiscale Data from Atomistic Simulations
David J. Gardner, Lawrence Livermore National Laboratory, USA; Daniel R. Reynolds, Southern Methodist University, USA

Parallel Multigrid and FFT-Type Algorithms for High Order Compact Discretization of the Helmholtz Equation.
Yury A. Gryazin, Idaho State University, USA

Computational Modeling and Aerodynamic Analysis of Simplified Automobile Body Using Open Source Cfd (OpenFoam)
Javeria Hashmi, National University of Science and Technology, Zimbabwe

Performance Evaluation of a Numerical Quadrature Based Eigensolver Using a Real Pole Filter
Yuto Inoue, Yasunori Futamura, and Tetsuya Sakurai, University of Tsukuba, Japan

Synthesis and Verification of BSP Programs
Arvid Jakobsson and Thibaut Tachon, Huawei Technologies, France

continued in next column

continued on next page

A Reconstruction Algorithm for Dual Energy Computed Tomography and Its Parallelization
Kiwan Jeon, National Institute for Mathematical Sciences, Korea; Sungwhan Kim, Hanbat National University, Korea; Chi Young Ahn and Taeyoung Ha, National Institute for Mathematical Sciences, Korea

Nlafet - Parallel Numerical Linear Algebra for Future Extreme-Scale Systems
Bo T. Kågström, Lennart Edblom, and Lars Karlsson, Umeå University, Sweden; Dongarra Jack, University of Manchester, United Kingdom; Iain Duff, Science & Technology Facilities Council, United Kingdom and CERFACS, Toulouse, France ; Jonathan Hogg, Rutherford Appleton Laboratory, United Kingdom; Laura Grigori, Inria Paris-Rocquencourt, France

An Interior-Point Stochastic Approximation Method on Massively Parallel Architectures
Juraj Kardoš and Drosos Kourounis, Università della Svizzera italiana, Switzerland; Olaf Schenk, Università della Svizzera Italiana, Italy

Seigen: Seismic Modelling Through Code Generation
Michael Lange, Christian Jacobs, Fabio Luporini, and Gerard J. Gorman, Imperial College London, United Kingdom

High-Performance Tensor Algebra Package for GPU
Ian Masliah, University of Paris-Sud, France; Ahmad Abdelfattah, University of Tennessee, Knoxville, USA; Marc Baboulin, University of Paris-Sud, France; Jack J. Dongarra, University of Tennessee and Oak Ridge National Laboratory, USA; Joel Falcou, University of Paris-Sud, France; Azzam Haidar and Stanimire Tomov, University of Tennessee, Knoxville, USA

Multilevel Approaches for FSAI Preconditioning in the Iterative Solution of Linear Systems of Equations
Victor A. Paludetto Magri, Andrea Franceschini, Carlo Janna, and Massimiliano Ferronato, University of Padova, Italy

Large Scale Computation of Long Range Interactions in Molecular Dynamics Simulations
Pyeong Jun Park, Korea National University of Transportation, Korea

Accelerated Subdomain Smoothing for Geophysical Stokes Flow
Patrick Sanan, Università della Svizzera italiana, Switzerland

ExaHyPE - Towards an Exascale Hyperbolic PDE Engine
Angelika Schwarz, Vasco Varduhn, and Michael Bader, Technische Universität München, Germany; Michael Dumbser, University of Trento, Italy; Dominic E. Charrier and Tobias Weinzierl, Durham University, United Kingdom; Luciano Rezzolla, Frankfurt Institute for Advanced Studies, Germany; Alice Gabriel and Heiner Igel, Ludwig-Maximilians-Universität München, Germany

Malleable Task-Graph Scheduling with a Practical Speed-Up Model
Bertrand Simon, ENS Lyon, France; Loris Marchal, CNRS, France; Oliver Sinnen, University of Auckland, New Zealand; Frédéric Vivien, Inria and ENS Lyon, France

Overview of Intel® Math Kernel Library (Intel® MKL) Solutions for Eigenvalue Problems
Irina Sokolova, Peter Tang and Alexander Kalinkin, Intel Corporation, USA

Achieving Robustness in Domain Decomposition Methods: Adpative Spectral Coarse Spaces
Nicole Spillane, Ecole Polytechnique, France; Victorita Dolean, University of Nice, France; Patrice Hauret, Ecole Polytechnique, France; Frédéric Nataf, CNRS and UPMC, France; Clemens Pechstein, University of Linz, Austria; Daniel J. Rixen, Technische Universität München, Germany; Robert Scheichl, University of Bath, United Kingdom

A Proposal for a Batched Set of Blas.
Pedro Valero-Lara, University of Manchester, United Kingdom; Jack J. Dongarra, University of Tennessee and Oak Ridge National Laboratory, USA; Azzam Haidar, University of Tennessee, Knoxville, USA; Samuel Relton, University of Manchester, United Kingdom; Stanimire Tomov, University of Tennessee, Knoxville, USA; Mawussi Zounon, University of Manchester, United Kingdom

Spectrum Slicing Solvers for Extreme-Scale Eigenvalue Computations
Eugene Vecharynski, Lawrence Berkeley National Laboratory, USA; Ruipeng Li, Lawrence Livermore National Laboratory, USA; Yuanzhe Xi, University of Minnesota, USA; Chao Yang, Lawrence Berkeley National Laboratory, USA; Yousef Saad, University of Minnesota, USA

Compressed Hierarchical Schur Linear System Solver for 3D Photonic Device Analysis with Gpu Accelerations
Cheng-Han Du and *Weichung Wang*, National Taiwan University, Taiwan

Parallel Adaptive and Robust Multigrid
Gabriel Wittum, University of Frankfurt, Germany

High performance eigenvalue solver for Hubbard model on CPU-GPU hybrid platform
Susumu Yamada, Japan Atomic Energy Agency, Japan; Toshiyuki Imamura, RIKEN, Japan; Masahiko Machida, Japan Atomic Energy Agency, Japan

continued in next column

continued in next column

Thursday, April 14

Registration
7:45 AM-5:15 PM

Room:Entrance Hall

IP4
Next-Generation AMR
8:30 AM-9:15 AM

Room:Farabeuf

Chair: Costas Bekas, IBM Research-Zurich, Switzerland

Block-structured adaptive mesh refinement (AMR) is a powerful tool for improving the computational efficiency and reducing the memory footprint of structured-grid numerical simulations. AMR techniques have been used for over 25 years to solve increasingly complex problems. I will give an overview of what we are doing in Berkeley Lab's AMR framework, BoxLib, to address the challenges of next-generation multicore architectures and the complexity of multiscale, multiphysics problems, including new ways of thinking about multilevel algorithms and new approaches to data layout and load balancing, in situ and in transit visualization and analytics, and run-time performance modeling and control.

Ann S. Almgren
Lawrence Berkeley National Laboratory, USA

Intermission
9:15 AM-9:25 AM

Thursday, April 14

SP1
SIAG/Supercomputing Career Prize: Supercomputers and Superintelligence
9:25 AM-9:50 AM

Room:Farabeuf

Chair: Laura Grigori, Inria, France

In recent years the idea of emerging superintelligence has been discussed widely by popular media, and many experts voiced grave warnings about its possible consequences. This talk will use an analysis of progress in supercomputer performance to examine the gap between current technology and reaching the capabilities of the human brain. In spite of good progress in high performance computing (HPC) and techniques such as machine learning, this gap is still very large. I will then explore two related topics through a discussion of recent examples: what can we learn from the brain and apply to HPC, e.g., through recent efforts in neuromorphic computing? And how much progress have we made in modeling brain function? The talk will be concluded with my perspective on the true dangers of superintelligence, and on our ability to ever build self-aware or sentient computers.

Horst D. Simon, Lawrence Berkeley National Laboratory, USA

SIAG/Supercomputing Early Career Prize
9:55 AM-10:00 AM

Room:Farabeuf

Coffee Break
10:10 AM-10:35 AM

Room:Exhibition Space

Thursday, April 14

MS41
Resilience Toward Exascale Computing - Part I of VI
10:35 AM-12:15 PM

Room:Farabeuf

For Part 2 see MS49

Future extreme scale systems are expected to suffer more frequent system faults resulting from the unprecedented scale of parallelism. Combined with the technology trend such as the imbalance between computing and I/O throughput and a tight power budget in the system operations, the traditional hardware-level redundancy and checkpoint restart may not be a feasible solution. Instead, every layer of the systems should address these faults to mitigate their impact in holistic manners. In this minisymposium, we will address the deficiencies of the current practice from a spectrum of high performance computing research from hardware, systems, runtime, algorithm and applications.

Organizer: Keita Teranishi
Sandia National Laboratories, USA

Organizer: Luc Giraud
Inria, France

Organizer: Emmanuel Agullo
Inria, France

Organizer: Francesco Rizzi
Sandia National Laboratories, USA

Organizer: Michael Heroux
Sandia National Laboratories, USA

10:35-10:55 Resilience Research is Essential but Failure is Unlikely
Michael Heroux, Sandia National Laboratories, USA

11:00-11:20 Memory Errors in Modern Systems
Steven Raasch and Vilas Sridharan, AMD, USA

11:25-11:45 Characterizing Faults on Production Systems
Martin Schulz, Lawrence Livermore National Laboratory, USA

11:50-12:10 The Missing High-Performance Computing Fault Model
Christian Engelmann, Oak Ridge National Laboratory, USA

Thursday, April 14

MS42

High-End Computing for Next-Generation Scientific Discovery - Part I of III

10:35 AM-12:15 PM

Room:Pasquier

For Part 2 see MS50

Computational modeling and simulation at large scale has become indispensable for science and engineering. However, the recent end of Dennard scaling is causing significant disruption to computing systems, resulting in increasingly complex and heterogeneous HPC platforms. In this minisymposium, we discuss the critical components for effectively building and leveraging HPC systems. We begin by examining the world's most powerful supercomputers, and discussing their architectural tradeoffs. Next we present key application drivers, whose science depends on the successful deployment of future exascale systems. Finally, we explore emerging technology solutions that have the potential to enable unprecedented next- generation computing capability.

Organizer: Leonid Oliker
Lawrence Berkeley National Laboratory, USA

Organizer: Rupak Biswas
NASA Ames Research Center, USA

Organizer: David Donofrio
Lawrence Berkeley National Laboratory, USA

10:35-10:55 From Titan to Summit: Hybrid Multicore Computing at ORNL
Arthur S. Bland, Oak Ridge National Laboratory, USA

11:00-11:20 Sequoia to Sierra: The LLNL Strategy
Bronis R. de Supinski, Lawrence Livermore National Laboratory, USA

11:25-11:45 Hazelhen – Leading HPC Technology and its Impact on Science in Germany and Europe
Thomas Bönisch and Michael Resch, University of Stuttgart, Germany

11:50-12:10 Lessons Learned from Development and Operation of the K Computer
Fumiyoshi Shoji, RIKEN, Japan

Thursday, April 14

MS43

DOE FASTMath Numerical Software on Next-generation Computers - Part I of II

10:35 AM-12:15 PM

Room:Roussy

For Part 2 see MS51

This two part minisymposium series focuses on algorithms and software developed by the DOE FASTMath SciDAC team to improve the reliability and robustness of application codes. We describe advances in the scalable implementation of unstructured mesh techniques, partitioning libraries, as well as linear and nonlinear solvers. We describe our experiences implementing advanced numerical software on the latest computer platforms including techniques for achieving million-parallelism, the use of heterogeneous nodes, and dragonfly networks. We give specific examples of the use of these libraries in large-scale scientific computing applications.

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

10:35-10:55 Lessons Learned in the Development of FASTMath Numerical Software for Current and Next Generation HPC
Lori A. Diachin, Lawrence Livermore National Laboratory, USA

11:00-11:20 Parallel Unstructured Mesh Generation/Adaptation Tools Used in Automated Simulation Workflows
Mark S. Shephard, Cameron Smith, Dan A. Ibanez, Brian Granzow, and Onkar Sahni, Rensselaer Polytechnic Institute, USA; Kenneth Jansen, University of Colorado Boulder, USA

11:25-11:45 Partitioning and Task Placement with Zoltan2
Mehmet Deveci, Karen D. Devine, Erik G. Boman, Vitus Leung, Siva Rajamanickam, and Mark A. Taylor, Sandia National Laboratories, USA

11:50-12:10 Scalable Nonlinear Solvers and Preconditioners in Climate Applications
David J. Gardner, Lawrence Livermore National Laboratory, USA; Richard Archibald and Katherine J. Evans, Oak Ridge National Laboratory, USA; Paul Ullrich, University of California, Davis, USA; Carol S. Woodward, Lawrence Livermore National Laboratory, USA; Patrick H. Worley, Oak Ridge National Laboratory, USA

continued in next column

Thursday, April 14

MS44

Parallel Algorithms for Tensor Computations - Part I of II

10:35 AM-12:15 PM

Room:Salle des theses

For Part 2 see MS52

Tensors, or multidimensional arrays, are a natural way to represent high-dimensional data arising in a multitude of applications. Tensor decompositions, such as the CANDECOMP/PARAFAC and Tucker models, help to identify latent structure, achieve data compression, and enable other tools of data analysis. This minisymposium identifies the key operations in related algorithms for tensor decomposition. The focus of the minisymposium is on recent developments in efficient, parallel algorithms and implementations for handing large tensors. The eight talks cover three key operations in different shared memory and distributed memory parallel computing systems environments and investigates applications in scientific computing and recommender systems.

Organizer: Bora Ucar
LIP-ENS Lyon, France

Organizer: Grey Ballard
Sandia National Laboratories, USA

Organizer: Tamara G. Kolda
Sandia National Laboratories, USA

10:35-10:55 Parallel Tensor Compression for Large-Scale Scientific Data
Woody N. Austin, University of Texas at Austin, USA; *Grey Ballard* and Tamara G. Kolda, Sandia National Laboratories, USA

11:00-11:20 Parallelizing and Scaling Tensor Computations
Muthu M. Baskaran, Benoit Meister, and Richard Lethin, Reservoir Labs, USA

11:25-11:45 Coupled Matrix and Tensor Factorizations: Models, Algorithms, and Computational Issues
Evrim Acar, University of Copenhagen, Denmark

11:50-12:10 Parallel Algorithms for Tensor Completion and Recommender Systems
Lars Karlsson, Umeå University, Sweden; *Daniel Kressner*, EPFL, Switzerland; Andre Uschmajew, University of Bonn, Germany

continued in next column

Thursday, April 14

MS45

The Present and Future Bulk Synchronous Parallel Computing - Part III of III

10:35 AM-12:15 PM

Room:Marie Curie

For Part 2 see MS37

The Bulk Synchronous Parallel (BSP) bridging model, proposed by Valiant in 1990, has had great influence on algorithm and software design over the past decades. It has inspired programming frameworks and structured parallel programming of communication minimising algorithms. In the current era of post-petascale HPC and of big data computing, BSP has again led to successful programming frameworks such as MapReduce, Pregel, and Giraph. This minisymposium explores the use of BSP in contemporary high performance and big data computing, highlighting new BSP algorithms and frameworks, new BSP-like models, and BSP-inspired auto-generation and verification tools.

Organizer: Albert-Jan N. Yzelman
Huawei Technologies, France

Organizer: Rob H. Bisseling
Utrecht University, The Netherlands

10:35-10:55 New Directions in BSP Computing
Bill McColl, Huawei Technologies, France

11:00-11:20 Multi-ML: Functional Programming of Multi-BSP Algorithms
Victor Allombert, Université Paris-est Créteil, France; Julien Tesson, Université Paris-Est, France; Frédéric Gava, Universite de Paris-Est, France

11:25-11:45 Parallelization Strategies for Distributed Machine Learning
Eric Xing, Carnegie Mellon University, USA

11:50-12:10 Panel: The Future of BSP Computing
Albert-Jan N. Yzelman, Huawei Technologies, France

Thursday, April 14

MS46

Parallel Solvers for High Performance Computing - Part I of II

10:35 AM-12:15 PM

Room:Leroux

For Part 2 see MS54

Large scale scientific computing is demanded by the most complex applications, generally multi scale, multi physics, non-linear, and/ or transient in nature. In the past, improving performance was only a matter of waiting for the next generation processors. As a consequence, research on solvers would take second place to the search for new discretisation schemes. But since approximately year 2005 the increase in performance is almost entirely due to the increase in the number of cores per processor. To keep doubling performance parallelism must double. Tailored algorithms for numerical simulations are of paramount importance.

Organizer: Victorita Dolean
University of Nice, France

10:35-10:55 Linear and Non-Linear Domain Decomposition Methods for Coupled Problems
Rolf Krause, Università della Svizzera Italiana, Italy; *Erich Foster*, USI, Switzerland; Marco Favino, University of Lugano, Switzerland; Patrick Zulian, USI, Switzerland

11:00-11:20 Improving Strong Scalability of Linear Solver for Reservoir Simulation
Tom B. Jonsthovel, Schlumberger, United Kingdom

11:25-11:45 PCBDDC: Robust BDDC Preconditioners in PETSc
Stefano Zampini, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

11:50-12:10 Solving Large Systems for the Elastodynamics Equation to Perform Acoustic Imaging
Dimitri Komatitsch, CNRS & Universite de Marseille, France; Paul Cristini, CNRS, France

Thursday, April 14

MS47

Parallel Methodologies for Large Data Assimilation and Inverse Problems Part II of II

10:35 AM-12:15 PM

Room:Dejerine

Data assimilation is the process of fusing information from priors, models, and observations, in order to obtain best estimates of the state of a physical system such as the atmosphere. Data assimilation relies on comprehensive physical models with large numbers of variables, and on complex observational data sets. The resulting large scale inverse problems need to be solved in faster than real time in order to, for example, meet the demands of weather forecasting. This minisymposium discusses the latest developments in parallel methodologies for solving large data assimilation problems.

Organizer: Adrian Sandu
Virginia Polytechnic Institute & State University, USA

Organizer: Marc Bocquet
Ecole Nationale des Ponts et Chaussées, France

10:35-10:55 From Data to Prediction with Quantified Uncertainties: Scalable Parallel Algorithms and Applications to the Flow of Ice Sheets
Omar Ghattas, University of Texas at Austin, USA; Tobin Isaac, University of Chicago, USA; Noemi Petra, University of California, Merced, USA; Georg Stadler, Courant Institute of Mathematical Sciences, New York University, USA

11:00-11:20 Parallel Aspects of Diffusion-Based Correlation Operators in Ensemble-Variational Data Assimilation
Anthony Weaver and Jean Tshimanga, CERFACS, France

11:25-11:45 Title Not Available
Mike Fisher, European Center for Medium Range Weather Forecasts, United Kingdom

11:50-12:10 Using Adjoint Based Numerical Error Estimates and Optimized Data Flows to Construct Surrogate Models
Abani Patra, State University of New York, Buffalo, USA

Thursday, April 14

MS48

Particle Method Based Applications and their Parallel Optimizations towards Exascale - Part I of II

10:35 AM-12:15 PM

Room:Delarue

For Part 2 see MS56

A large number of industrial problems can be modeled using particle-based methods, The recent published DOE exascale report identified the particles based methods as “well-suited for exascale computing”. The reason is that particles based methods provide extremely fine-grained parallelism and allow the exploitation of asynchrony. Efficient parallel particle method applications require the efficient implementation: domain decomposition, dynamic load balancing, optimized structured and unstructured communication, parallel I/O, neighbour lists searching, particle-to-mesh, mesh-to- particle interpolation, sparse linear solvers. The present proposal aims to identify the common efficient implementation of the above kernels and their optimizations among different applications for the future exascale systems

Organizer: Xiaohu Guo
Science and Technology Facilities Council, United Kingdom

Organizer: Tao Cui
Chinese Academy of Sciences, China

10:35-10:55 Applications of Smoothed Particle Hydrodynamics (SPH) for Real Engineering Problems: Particle Parallelisation Requirements and a Vision for the Future
Benedict Rogers, University of Manchester, United Kingdom

continued on next page

Thursday, April 14

MS48

Particle Method Based Applications and their Parallel Optimizations towards Exascale - Part I of II

10:35 AM-12:15 PM

Room:Delarue

continued

11:00-11:20 Parallel Application of Monte Carlo Particle Simulation Code JMCT in Nuclear Reator Analysis
Gang Li, Institute of Applied Physics and Computational Mathematics, Beijing, P.R. China

11:25-11:45 Parallel Pic Simulation Based on Fast Electromagnetic Field Solver
Tao Cui, Chinese Academy of Sciences, China

11:50-12:10 An Hybrid OpenMP-MPI Parallelization for Variable Resolution Adaptive SPH with Fully Object-Oriented SPHysics
Ranato Vacondio, University of Parma, Italy

Thursday, April 14

CP9

System-level Scalability

10:35 AM-12:15 PM

Room:Danton

Chair: Courtenay T. Vaughan, Sandia National Laboratories, USA

10:35-10:50 On the Path to Trinity - Experiences Bringing Code to the Next Generation ASC Capability Platform
Courtenay T. Vaughan and Simon D. Hammond, Sandia National Laboratories, USA

10:55-11:10 Implicitly Parallel Interfaces to Operational Quantum Computation
Steve P. Reinhardt, D-Wave Systems, Inc., USA

11:15-11:30 Investigating IBM Power8 System Configuration To Minimize Energy Consumption During the Execution of CoreNeuron Scientific Application
Fabien Delalondre and Timothee Ewart, École Polytechnique Fédérale de Lausanne, Switzerland; Pramod Shivaj and Felix Schuermann, EPFL, France

11:35-11:50 Scalable Parallel I/O Using HDF5 for HPC Fluid Flow Simulations
Christoph M. Ertl and Ralf-Peter Mundani, Technische Universität München, Germany; Jérôme Frisch, RWTH Aachen University, Germany; Ernst Rank, Technische Universität München, Germany

11:55-12:10 Quantifying Performance of a Resilient Elliptic PDE Solver on Uncertain Architectures Using SST/Macro
Karla Morris, Francesco Rizzi, Khachik Sargsyan, and Kathryn Dahlgren, Sandia National Laboratories, USA; Paul Mycek, Duke University, USA; Cosmin Safta, Sandia National Laboratories, USA; Olivier P. Le Maître, LIMSI-CNRS, France; Omar M. Knio, Duke University, USA; Bert J. Debusschere, Sandia National Laboratories, USA

Thursday, April 14

Lunch Break

12:15 PM-1:45 PM

Attendees on their own

IP5

Beating Heart Simulation Based on a Stochastic Biomolecular Model

1:45 PM-2:30 PM

Room:Farabeuf

Chair: Naoya Maruyama, RIKEN, Japan

In my talk, I will introduce a new parallel computational approach indispensable to simulate a beating human heart driven by stochastic biomolecular dynamics. In this approach, we start from a molecular level modeling, and we directly simulate individual molecules in sarcomere by the Monte Carlo (MC) method, thereby naturally handle the cooperative and stochastic molecular behaviors. Then we imbed these sarcomere MC samples into contractile elements in continuum material model discretized by finite element method. The clinical applications of our heart simulator are also introduced.

Takumi Washio

University of Tokyo, Japan

Intermission

2:30 PM-2:40 PM

Thursday, April 14

MS49

Resilience Toward Exascale Computing - Part II of VI

2:40 PM-4:20 PM

Room:Farabeuf

For Part 1 see MS41
For Part 3 see MS57

Future extreme scale systems are expected to suffer more frequent system faults resulting from the unprecedented scale of parallelism. Combined with the technology trend such as the imbalance between computing and I/O throughput and a tight power budget in the system operations, the traditional hardware-level redundancy and checkpoint restart may not be a feasible solution. Instead, every layer of the systems should address these faults to mitigate their impact in holistic manners. In this minisymposium, we will address the deficiencies of the current practice from a spectrum of high performance computing research from hardware, systems, runtime, algorithm and applications.

Organizer: Keita Teranishi
Sandia National Laboratories, USA

Organizer: Luc Giraud
Inria, France

Organizer: Emmanuel Agullo
Inria, France

Organizer: Francesco Rizzi
Sandia National Laboratories, USA

Organizer: Michael Heroux
Sandia National Laboratories, USA

continued in next column

2:40-3:00 A Soft and Hard Faults Resilient Solver for 2D Elliptic PDEs Via Server-Client-Based Implementation

Francesco Rizzi and Karla Morris, Sandia National Laboratories, USA; Paul Mycek, Duke University, USA; Olivier P. Le Maître, LIMSI-CNRS, France; Andres Contreras, Duke University, USA; Khachik Sargsyan and Cosmin Safta, Sandia National Laboratories, USA; Omar M. Knio, Duke University, USA; Bert J. Debusschere, Sandia National Laboratories, USA

3:05-3:25 Fault-Tolerant Multigrid Algorithms Based on Error Confinement and Full Approximation

Mirco Altenbernd, Universität Stuttgart, Germany

3:30-3:50 Soft Error Detection for Finite Difference Codes on Regular Grids

Peter Arbenz and Manuel Kohler, ETH Zürich, Switzerland

3:55-4:15 Resilience for Multigrid at the Extreme Scale

Ulrich J. Ruede, University of Erlangen-Nuremberg, Germany; *Markus Huber*, Technische Universität München, Germany; Bjoern Gmeiner, University Erlangen-Nürnberg, Germany; Barbara Wohlmuth, Technische Universität München, Germany

Thursday, April 14

MS50

High-End Computing for Next-Generation Scientific Discovery - Part II of III

2:40 PM-4:20 PM

Room:Pasquier

For Part 1 see MS42
For Part 3 see MS58

Computational modeling and simulation at large scale has become indispensable for science and engineering. However, the recent end of Dennard scaling is causing significant disruption to computing systems, resulting in increasingly complex and heterogeneous HPC platforms. In this minisymposium, we discuss the critical components for effectively building and leveraging HPC systems. We begin by examining the world's most powerful supercomputers, and discussing their architectural tradeoffs. Next we present key application drivers, whose science depends on the successful deployment of future exascale systems. Finally, we explore emerging technology solutions that have the potential to enable unprecedented next-generation computing capability.

Organizer: Leonid Oliker
Lawrence Berkeley National Laboratory, USA

Organizer: Rupak Biswas
NASA Ames Research Center, USA

2:40-3:00 Extreme-Scale Genome Assembly

Leonid Oliker, Aydin Buluc, and Rob Egan, Lawrence Berkeley National Laboratory, USA; Evangelos Georganas, University of California, Berkeley, USA; Steven Hofmeyr, Lawrence Berkeley National Laboratory, USA; Daniel Rokhsar and Katherine Yelick, Lawrence Berkeley National Laboratory and University of California Berkeley, USA

continued on next page

Thursday, April 14

MS50

High-End Computing for Next-Generation Scientific Discovery - Part II of III

2:40 PM-4:20 PM

Room:Pasquier

continued

3:05-3:25 Research Applications and Analysis of a Fully Eddying Virtual Ocean.
Chris Hill, Massachusetts Institute of Technology, USA

3:30-3:50 Progress in Performance Portability of the XGC Plasma Microturbulence Codes: Heterogeneous and Manycore Node Architectures
Eduardo F. D'Azevedo, Oak Ridge National Laboratory, USA; *Mark Adams*, Lawrence Berkeley National Laboratory, USA; Choong-Seock Chang, Stephane Ethier, Robert Hager, Seung-Hoe Ku, and Jianying Lang, Princeton Plasma Physics Laboratory, USA; Mark Shepard, Rensselaer Polytechnic Institute, USA; Patrick H. Worley, Oak Ridge National Laboratory, USA; Eisung Yoon, Rensselaer Polytechnic Institute, USA

3:55-4:15 Challenges in Adjoint-Based Aerodynamic Design for Unsteady Flows
Eric Nielsen, NASA Langley Research Center, USA

Thursday, April 14

MS51

DOE FASTMath Numerical Software on Next-generation Computers - Part II of II

2:40 PM-4:20 PM

Room:Roussy

For Part 1 see MS43

This two part minisymposium series focuses on algorithms and software developed by the DOE FASTMath SciDAC team to improve the reliability and robustness of application codes. We describe advances in the scalable implementation of unstructured mesh techniques, partitioning libraries, as well as linear and nonlinear solvers. We describe our experiences implementing advanced numerical software on the latest computer platforms including techniques for achieving million-parallelism, the use of heterogeneous nodes, and dragonfly networks. We give specific examples of the use of these libraries in large-scale scientific computing applications.

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

2:40-3:00 An Extension of Hypre's Structured and Semi-Structured Matrix Classes
Ulrike M. Yang and Robert D. Falgout, Lawrence Livermore National Laboratory, USA

3:05-3:25 Unstructured Mesh Adaptation for MPI and Accelerators
Dan A. Ibanez, Rensselaer Polytechnic Institute, USA

3:30-3:50 Comparing Global Link Arrangements for Dragonfly Networks
Vitus Leung, Sandia National Laboratories, USA; David Bunde, Knox College, USA

3:55-4:15 Scaling Unstructured Mesh Computations
Kenneth Jansen, University of Colorado Boulder, USA; Jed Brown, Argonne National Laboratory and University of Colorado Boulder, USA; Michel Rasquin and Benjamin Matthews, University of Colorado Boulder, USA; Cameron Smith, Dan A. Ibanez, and Mark S. Shephard, Rensselaer Polytechnic Institute, USA

continued in next column

Thursday, April 14

MS52

Parallel Algorithms for Tensor Computations - Part II of II

2:40 PM-4:20 PM

Room:Salle des theses

For Part 1 see MS44

Tensors, or multidimensional arrays, are a natural way to represent high-dimensional data arising in a multitude of applications. Tensor decompositions, such as the CANDECOMP/PARAFAC and Tucker models, help to identify latent structure, achieve data compression, and enable other tools of data analysis. This minisymposium identifies the key operations in related algorithms for tensor decomposition. The focus of the minisymposium is on recent developments in efficient, parallel algorithms and implementations for handing large tensors. The eight talks cover three key operations in different shared memory and distributed memory parallel computing systems environments and investigates applications in scientific computing and recommender systems.

Organizer: Grey Ballard
Sandia National Laboratories, USA

Organizer: Tamara G. Kolda
Sandia National Laboratories, USA

Organizer: Bora Ucar
LIP-ENS Lyon, France

2:40-3:00 High Performance Parallel Sparse Tensor Decompositions
Oguz Kaya, Inria and ENS Lyon, France; Bora Ucar, LIP-ENS Lyon, France

3:05-3:25 Efficient Factorization with Compressed Sparse Tensors
Shaden Smith, University of Minnesota, USA; George Karypis, University of Minnesota and Army HPC Research Center, USA

continued in next column

3:30-3:50 Low Rank Bilinear Algorithms for Symmetric Tensor Contractions
Edgar Solomonik, ETH Zürich, Switzerland

3:55-4:15 An Input-Adaptive and In-Place Approach to Dense Tensor-Times-Matrix Multiply
Jiajia Li, Casey Battaglino, Ioakeim Perros, Jimeng Sun, and Richard Vuduc, Georgia Institute of Technology, USA

Thursday, April 14

MS53

Large-scale Electronic Structure Calculation: Algorithms, Performance and Applications - Part I of III

2:40 PM-4:20 PM

Room:Marie Curie

For Part 2 see MS61

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 DGDFT: A Massively Parallel Electronic Structure Calculation Tool
Chao Yang, Lawrence Berkeley National Laboratory, USA

continued on next page

Thursday, April 14

MS53

Large-scale Electronic Structure Calculation: Algorithms, Performance and Applications - Part I of III

2:40 PM-4:20 PM

Room:Marie Curie

continued

3:05-3:25 Design and Implementation of the PEXSI Solver Interface in SIESTA
Alberto Garcia, Institut de Ciencia de Materials de Barcelona (ICMAB-CSIC), Spain; Lin Lin, University of California, Berkeley and Lawrence Berkeley National Laboratory, USA; Georg Huhs, Barcelona Supercomputing Center, Spain; Chao Yang, Lawrence Berkeley National Laboratory, USA

3:30-3:50 Wavelets as a Basis Set for Electronic Structure Calculations: The BigDFT Project
Stefan Goedecker, University of Basel, Switzerland

3:55-4:15 BigDFT: Flexible DFT Approach to Large Systems Using Adaptive and Localized Basis Functions
Thierry Deutsch, CEA, France

Thursday, April 14

MS54

Parallel Solvers for High Performance Computing - Part II of II

2:40 PM-4:20 PM

Room:Leroux

For Part 1 see MS46

Large scale scientific computing is demanded by the most complex applications, generally multi scale, multi physics, non-linear, and/or transient in nature. In the past, improving performance was only a matter of waiting for the next generation processors. As a consequence, research on solvers would take second place to the search for new discretisation schemes. But since approximately year 2005 the increase in performance is almost entirely due to the increase in the number of cores per processor. To keep doubling performance parallelism must double. Tailored algorithms for numerical simulations are of paramount importance.

Organizer: Victorita Dolean
University of Nice, France

2:40-3:00 Preconditioning of High Order Edge Elements Type Discretizations for the Time-Harmonic Maxwell's Equations
Victorita Dolean, University of Nice, France

3:05-3:25 Parallel Preconditioners for Problems Arising from Whole-Microwave System Modelling for Brain Imaging
Pierre-Henri Tournier, Université Pierre et Marie Curie and Inria, France; Pierre Jolivet, CNRS, France; Frédéric Nataf, CNRS and UPMC, France

3:30-3:50 Block Iterative Methods and Recycling for Improved Scalability of Linear Solvers.
Pierre Jolivet, CNRS, France; Pierre-Henri Tournier, Université Pierre et Marie Curie and Inria, France

3:55-4:15 Reliable Domain Decomposition Methods with Multiple Search
Nicole Spillane, Ecole Polytechnique, France; Christophe Bovet, CNRS, Université Paris-Saclay, France; Pierre Gosselet, LMT-Cachan, France; Daniel J. Rixen, Technische Universität München, Germany; François-Xavier Roux, Université Pierre et Marie Curie, France

continued in next column

Thursday, April 14

MS55

Auto-Tuning for the Post Moore's Era - Part I of II

2:40 PM-4:20 PM

Room:Dejerine

For Part 2 see MS63

It is expected that the progress of semiconductor technology will be saturated by the end of 2020 due to physical limitations in device miniaturization. As a result, no gains in FLOPS computations will be achieved with a constant power budget. This change is referred to as "Post Moore's Era". We anticipate that auto-tuning (AT) technologies will still have the ability of providing optimized, high performance optimization for architectures of the Post Moore's Era. This minisymposium will discuss AT frameworks and technologies that target the core of most large applications, including linear solvers, eigensolvers and stencil computations for new algorithm towards the Post Moore's Era.

Organizer: Takahiro Katagiri
University of Tokyo, Japan

Organizer: Toshiyuki Imamura
RIKEN, Japan

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

2:40-3:00 Auto-Tuning of Hierarchical Computations with ppOpen-AT
Takahiro Katagiri, Masaharu Matsumoto, and Satoshi Ohshima, University of Tokyo, Japan

3:05-3:25 Scaling Up Autotuning
Jeffrey Hollingsworth, University of Maryland, USA

3:30-3:50 User-Defined Code Transformation for High Performance Portability
Hiroyuki Takizawa, Shoichi Hirasawa, and Hiroaki Kobayashi, Tohoku University, Japan

3:55-4:15 Auto-Tuning for the SVD of Bidiagonals
Osni A. Marques, Lawrence Berkeley National Laboratory, USA

Thursday, April 14

MS56

Particle Method Based Applications and their Parallel Optimizations towards Exascale - Part II of II

2:40 PM-4:20 PM

Room:Delarue

For Part 1 see MS48

A large number of industrial problems can be modeled using particle-based methods, The recent published DOE exascale report identified the particles based methods as "well-suited for exascale computing". The reason is that particles based methods provide extremely fine-grained parallelism and allow the exploitation of asynchrony. Efficient parallel particle method applications require the efficient implementation: domain decomposition, dynamic load balancing, optimized structured and unstructured communication, parallel I/O, neighbour lists searching, particle-to-mesh, mesh-to-particle interpolation, sparse linear solvers. The present proposal aims to identify the common efficient implementation of the above kernels and their optimizations among different applications for the future exascale systems

Organizer: Xiaohu Guo
Science and Technology Facilities Council, United Kingdom

Organizer: Tao Cui
Chinese Academy of Sciences, China

2:40-3:00 H5hut: a High Performance Hdf5 Toolkit for Particle Simulations
Achim Gsell, PSI, USA

3:05-3:25 Weakly Compressible Smooth Particle Hydrodynamics Performance on the Intel Xeon Phi
Sergi-Enric Siso, Science and Technology Facilities Council, United Kingdom

3:30-3:50 Parallelization of Smoothed Particle Hydrodynamics for Engineering Applications with the Heterogeneous Architectures
Jose M. Domínguez Alonso, Universidade de Vigo, Spain

3:55-4:15 Towards a Highly Scalable Particles Method Based Toolkit: Optimization for Real Applications
Xiaohu Guo, Science and Technology Facilities Council, United Kingdom

continued in next column

Thursday, April 14

CP10

Scalable Applications

2:40 PM-4:20 PM

Room:Danton

Chair: Axel Modave, Virginia Tech, USA

2:40-2:55 Efficient Solver for High-Resolution Numerical Weather Prediction Over the Alps
Zbigniew P. Piotrowski, Damian Wójcik, Andrzej Wyszogrodzki, and Milosz Ciznicki, Poznan Supercomputing and Networking Center, Poland

3:00-3:15 An Adaptive Multi-Resolution Data Model for Feature-Rich Environmental Simulations on Multiple Scales
Vasco Varduhn, Technische Universität München, Germany

3:20-3:35 An Overlapping Domain Decomposition Preconditioner for the Helmholtz Equation
Wei Leng, Chinese Academy of Sciences, China; Lili Ju, University of South Carolina, USA

3:40-3:55 Wave Propagation in Highly Heterogeneous Media: Scalability of the Mesh and Random Properties Generator
Regis Cottereau, CNRS, CentraleSupelec, Université Paris-Saclay, France; José Camata, Federal University of Rio de Janerio, Brazil; Lucio de Abreu Corrêa and Luciano de Carvalho Paludo, CNRS, CentraleSupelec, Université Paris-Saclay, France; Alvaro Coutinho, Federal University of Rio de Janerio, Brazil

4:00-4:15 GPU Performance Analysis of Discontinuous Galerkin Implementations for Time-Domain Wave Simulations
Axel Modave, Rice University, USA; Jesse Chan and Tim Warburton, Virginia Tech, USA

Coffee Break

4:20 PM-4:50 PM

Room:Exhibition Space

Thursday, April 14

MS57

Resilience Toward Exascale Computing - Part III of VI

4:50 PM-6:30 PM

Room:Farabeuf

For Part 2 see MS49
For Part 4 see MS65

Future extreme scale systems are expected to suffer more frequent system faults resulting from the unprecedented scale of parallelism. Combined with the technology trend such as the imbalance between computing and I/O throughput and a tight power budget in the system operations, the traditional hardware-level redundancy and checkpoint restart may not be a feasible solution. Instead, every layer of the systems should address these faults to mitigate their impact in holistic manners. In this minisymposium, we will address the deficiencies of the current practice from a spectrum of high performance computing research from hardware, systems, runtime, algorithm and applications.

Organizer: Keita Teranishi
Sandia National Laboratories, USA

Organizer: Luc Giraud
Inria, France

Organizer: Emmanuel Agullo
Inria, France

Organizer: Francesco Rizzi
Sandia National Laboratories, USA

Organizer: Michael Heroux
Sandia National Laboratories, USA

4:50-5:10 Resilience for Asynchronous Runtime Systems and Algorithms
Marc Casas, Barcelona Supercomputing Center, Spain

continued in next column

5:15-5:35 Silent Data Corruption and Error Propagation in Sparse Matrix Computations
Luke Olson, Jon Calhoun, and Marc Snir, University of Illinois at Urbana-Champaign, USA

5:40-6:00 Experiments with FSEFI: The Fine-Grained Soft Error Fault Injector
Nathan A. DeBardeleben and Qiang Guan, Los Alamos National Laboratory, USA

6:05-6:25 What Error to Expect When You Are Expecting a Bit Flip.
James Elliott, North Carolina State University, USA; Mark Hoemmen, Sandia National Laboratories, USA; Frank Mueller, North Carolina State University, USA

Thursday, April 14

MS58

High-End Computing for Next-Generation Scientific Discovery - Part III of III

4:50 PM-6:30 PM

Room:Pasquier

For Part 2 see MS50

Computational modeling and simulation at large scale has become indispensable for science and engineering. However, the recent end of Dennard scaling is causing significant disruption to computing systems, resulting in increasingly complex and heterogeneous HPC platforms. In this minisymposium, we discuss the critical components for effectively building and leveraging HPC systems. We begin by examining the world's most powerful supercomputers, and discussing their architectural tradeoffs. Next we present key application drivers, whose science depends on the successful deployment of future exascale systems. Finally, we explore emerging technology solutions that have the potential to enable unprecedented next- generation computing capability.

Organizer: Leonid Oliker
Lawrence Berkeley National Laboratory, USA

Organizer: Rupak Biswas
NASA Ames Research Center, USA

Organizer: David Donofrio
Lawrence Berkeley National Laboratory, USA

4:50-5:10 Quantum Computing: Opportunities and Challenges
Rupak Biswas, NASA Ames Research Center, USA

5:15-5:35 Emerging Hardware Technologies for Exascale Systems
David Donofrio, Lawrence Berkeley National Laboratory, USA

5:40-6:00 Silicon Photonics for Extreme Scale Computing
Keren Bergman and Sebastien Rumley, Columbia University, USA

6:05-6:25 Task-Based Execution Models and Co-Design
Romain Cledat and Joshua Fryman, Intel Corporation, USA

Thursday, April 14

MS59

Nonlinear Preconditioning

4:50 PM-6:30 PM

Room:Roussy

Increasing local computational work and reducing communication are key ingredients for efficiently using future exascale machines. In Newton-Krylov methods using domain decomposition as preconditioner, these aspects can be mainly treated at the level of solving the linear systems. Computational work can be localized and communication reduced by a reordering of operations: first decomposing the nonlinear problem, then linearizing it, leading to nonlinear domain decomposition. Nonlinear preconditioning can also be seen as a globalization strategy. Recent approaches to nonlinear domain decomposition methods and the composition of algebraic nonlinear solvers as well as results on several hundred thousand cores will be presented.

Organizer: Axel Klawonn
Universitaet zu Koeln, Germany

Organizer: Oliver Rheinbach
Technische Universitaet Bergakademie Freiberg, Germany

4:50-5:10 A Framework for Nonlinear Feti-Dp and Bddc Domain Decomposition Methods
Axel Klawonn, Universitaet zu Koeln, Germany; Martin Lanser and Martin Lanser, University of Cologne, Germany; Oliver Rheinbach, Technische Universitaet Bergakademie Freiberg, Germany

5:15-5:35 Scaling Nonlinear FETI-DP Domain Decomposition Methods to 786 452 Cores
Axel Klawonn, Universitaet zu Koeln, Germany; Martin Lanser and Martin Lanser, University of Cologne, Germany; *Oliver Rheinbach*, Technische Universitaet Bergakademie Freiberg, Germany

5:40-6:00 Convergence Estimates for Composed Iterations
Matthew G. Knepley, Rice University, USA

6:05-6:25 Additive and Multiplicative Nonlinear Schwarz
David E. Keyes and Lulu Liu, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

Thursday, April 14

MS60

Advances in Parallel Dense and Sparse Linear Algebra

4:50 PM-6:30 PM

Room:Salle des theses

Dense and sparse linear algebra are important kernels for many scientific problems. Over the past decade important developments in the use of runtime systems and other techniques have been developed for basic dense linear algebra problems that aim to manage the increasing complexity of HPC resources. This minisymposium addresses recent progress on dense and sparse linear algebra issues that must be solved as we move towards ever more extreme- scale systems.

Organizer: Jonathan Hogg
Rutherford Appleton Laboratory, United Kingdom

Organizer: Lars Karlsson
Umeå University, Sweden

4:50-5:10 The Challenge of Strong Scaling for Direct Methods
Jonathan Hogg, Rutherford Appleton Laboratory, United Kingdom

5:15-5:35 NUMA-Aware Hessenberg Reduction in the Context of the QR Algorithm
Mahmoud Eljammaly and Lars Karlsson, Umeå University, Sweden

5:40-6:00 Low Rank Approximation of Sparse Matrices for Lu Factorization Based on Row and Column Permutations
Laura Grigori and Sebastien Cayrols, Inria, France; James W. Demmel, University of California, Berkeley, USA

6:05-6:25 Assessing Recent Sparse Matrix Storage Formats for Parallel Sparse Matrix-Vector Multiplication
Weifeng Liu, University of Copenhagen, Denmark

Thursday, April 14

MS61

Large-scale Electronic Structure Calculation: Algorithms, Performance and Applications - Part II of III

4:50 PM-6:30 PM

Room:Marie Curie

For Part 1 see MS53
For Part 3 see MS69

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 High-Performance Real-Time TDDFT Excited-States Calculations
Eric Polizzi, University of Massachusetts, Amherst, USA

continued in next column

5:15-5:35 MOLGW: A Parallel Many-Body Perturbation Theory Code for Finite Systems
Fabien Bruneval, CEA, DEN, SRMP, France; Samia M. Hamed, Tonatiuh Rangel-Gordillo, and Jeffrey B. Neaton, University of California, Berkeley, USA

5:40-6:00 New Building Blocks for Scalable Electronic Structure Theory
Volker Blum, Duke University, USA

6:05-6:25 GPU-Accelerated Simulation of Nano-Devices from First-Principles
Mathieu Luisier, Mauro Calderara, Sascha Brueck, Hossein Bani-Hashemian, and Joost VandeVondele, ETH Zürich, Switzerland

Thursday, April 14

MS62

Scalable Task-based Programming Models - Part I of II

4:50 PM-6:30 PM

Room:Leroux

For Part 2 see MS70

Programming models that enable users to write sequential code that can be efficiently executed in parallel, for example by a dedicated run-time system, on heterogeneous architecture, are becoming ubiquitous. Such programming models provide an easy interface to developers while exploiting the performance of the underlying infrastructure. In this minisymposium, we explore the current state-of-the-art of task-based programming models, the potential for unification, user interfaces, and special issues such as energy or resource awareness and resilience. The minisymposium is organized in two parts, with contributions from key players in task-based programming.

Organizer: Elisabeth Larsson
Uppsala University, Sweden

Organizer: Rosa M. Badia
Barcelona Supercomputing Center, Spain

4:50-5:10 A Single Interface for Multiple Task-Based Parallel Programing Paradigms/Schemes
Afshin Zafari and Elisabeth Larsson, Uppsala University, Sweden

5:15-5:35 Composition and Compartmentalisation As Enabling Features for Data-Centric, Extreme Scale Applications: An MPI-X Approach
Marco Aldinucci, Maurizio Drocco, and Claudia Misale, University of Torino, Italy; Massimo Torquati, University of Pisa, Italy

5:40-6:00 Task Dataflow Scheduling and the Quest for Extremely Fine-Grain Parallelism
Hans Vandierendonck and Mahwish Arif, Queen's University, Belfast, United Kingdom

6:05-6:25 Recent Advances in the Task-Based StarPU Runtime Ecosystem
Samuel Thibault, University of Bordeaux, France; Marc Sergent and Suraj Kumar, Inria, France; Luka Stanisic, INRIA Grenoble Rhône-Alpes, France

Thursday, April 14

MS63

Auto-Tuning for the Post Moore's Era - Part II of II

4:50 PM-6:30 PM

Room:Dejerine

For Part 1 see MS55

It is expected that the progress of semiconductor technology will be saturated by the end of 2020 due to physical limitations in device miniaturization. As a result, no gains in FLOPS computations will be achieved with a constant power budget. This change is referred to as "Post Moore's Era". We anticipate that auto-tuning (AT) technologies will still have the ability of providing optimized, high performance optimization for architectures of the Post Moore's Era. This minisymposium will discuss AT frameworks and technologies that target the core of most large applications, including linear solvers, eigensolvers and stencil computations for new algorithm towards the Post Moore's Era.

Organizer: Toshiyuki Imamura
RIKEN, Japan

Organizer: Takahiro Katagiri
University of Tokyo, Japan

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

4:50-5:10 Auto-Tuning for Eigenvalue Solver on the Post Moore's Era
Toshiyuki Imamura, RIKEN, Japan

5:15-5:35 Hardware/Algorithm Co-Optimization and Co-Tuning in the Post Moore Era
Franz Franchetti, Carnegie Mellon University, USA

5:40-6:00 Automatic Tuning for Parallel FFTs on Intel Xeon Phi Clusters
Daisuke Takahashi, University of Tsukuba, Japan

6:05-6:25 On Guided Installation of Basic Linear Algebra Routines in Nodes with Manycore Components
Javier Cuenca, University of Murcia, Spain; Luis-Pedro García, Technical University of Cartagena, Spain; *Domingo Giménez*, University of Murcia, Spain

Thursday, April 14

MS64

Multi-Resolution Models for Environmental Simulations

4:50 PM-6:30 PM

Room:Delarue

Many important scientific and technological developments in the area of environmental simulations have resulted in an enormous increase in availability of models on different scales. While such tremendous amounts of information become available, the pure visual data exploration is not satisfying at all in order to gain sufficient insight. Hence, coupling multi-resolution models with numerical simulations is an indispensable step towards meaningful predictions of physical phenomena within modern engineering tasks. Important questions to be discussed in this Minisymposium will include efficient coupling of numerical models on different scales and how a predictive quality of such coupled simulations can be obtained.

Organizer: Jérôme Frisch
RWTH Aachen University, Germany

Organizer: Ralf-Peter Mundani
Technische Universität München, Germany

4:50-5:10 Thermal Comfort Evaluation for Indoor Airflows Using a Parallel CFD Approach
Jérôme Frisch, RWTH Aachen University, Germany; Ralf-Peter Mundani, Technische Universität München, Germany

continued in next column

5:15-5:35 High-Performance Computing in Free Surface Fluid Flow Simulations
Nevena Perovic, Technische Universität München, Germany; Jérôme Frisch, RWTH Aachen University, Germany; Ralf-Peter Mundani and Ernst Rank, Technische Universität München, Germany

5:40-6:00 Fast Contact Detection for Granulates
Tomasz Koziara, *Konstantinos Krestenitis*, Jon Trevelyan, and Tobias Weinzierl, Durham University, United Kingdom

6:05-6:25 LB Particle Coupling on Dynamically-Adaptive Cartesian Grids
Michael Lahnert and Miriam Mehl, Universität Stuttgart, Germany

Thursday, April 14

CP11

Domain-specific Programming Tools

4:50 PM-5:50 PM

Room:Danton

Chair: Paul Springer, RWTH Aachen University, Germany

4:50-5:05 Maximizing CoreNeuron Application Performance on All CPU Architectures Using Cyme Library
Timothee Ewart, École Polytechnique Fédérale de Lausanne, Switzerland; Kai Langen, Pramod Kumbhar, and Felix Schuermann, EPFL, France; Fabien Delalondre, École Polytechnique Fédérale de Lausanne, Switzerland

5:10-5:25 Fresh: A Framework for Real-World Structured Grid Stencil Computations on Heterogeneous Platforms
Yang Yang, Aiqing Zhang, and Zhang Yang, Institute of Applied Physics and Computational Mathematics, China

5:30-5:45 Ttc: A Compiler for Tensor Transpositions
Paul Springer and Paolo Bientinesi, RWTH Aachen University, Germany

Conference Dinner (*separate fees apply*)

7:30 PM-10:30 PM

Room:Offsite (Salons Hoche)

Friday, April 15

Registration

8:15 AM-5:15 PM

Room:Entrance Hall

IP6

Conquering Big Data with Spark

8:45 AM-9:30 AM

Room:Farabeuf

Chair: Rosa M. Badia, Barcelona Supercomputing Center, Spain

Today, big and small organizations alike collect huge amounts of data, and they do so with one goal in mind: extract "value" through sophisticated exploratory analysis, and use it as the basis to make decisions as varied as personalized treatment and ad targeting. To address this challenge, we have developed Berkeley Data Analytics Stack (BDAS), an open source data analytics stack for big data processing.

Ion Stoica
University of California, Berkeley, USA

Intermission

9:30 AM-9:40 AM

Friday, April 15

SP2

SIAG/Supercomputing Best Paper Prize: Communication-Optimal Parallel and Sequential QR and LU Factorizations

9:40 AM-10:05 AM

Room:Farabeuf

Chair: Olaf Schenk, Università della Svizzera Italiana, Italy

We present parallel and sequential dense QR factorization algorithms that are both optimal (up to polylogarithmic factors) in the amount of communication they perform and just as stable as Householder QR. We prove optimality by deriving new lower bounds for the number of multiplications done by "non-Strassen-like" QR, and using these in known communication lower bounds that are proportional to the number of multiplications. We not only show that our QR algorithms attain these lower bounds (up to polylogarithmic factors), but that existing LAPACK and ScaLAPACK algorithms perform asymptotically more communication. We derive analogous communication lower bounds for LU factorization and point out recent LU algorithms in the literature that attain at least some of these lower bounds. The sequential and parallel QR algorithms for tall and skinny matrices lead to significant speedups in practice over some of the existing algorithms, including LAPACK and ScaLAPACK, for example, up to 6.7 times over ScaLAPACK. A performance model for the parallel algorithm for general rectangular matrices predicts significant speedups over ScaLAPACK.

Joint work with James W. Demmel, University of California, Berkeley, USA; Mark Hoemmen, Sandia National Laboratories, US; Julien Langou, University of Colorado, Denver, USA

Laura Grigori
Inria, France

Coffee Break

10:10 AM-10:35 AM

Room: Exhibition Space

Friday, April 15

MS65

Resilience Toward Exascale Computing - Part IV of VI

10:35 AM-12:15 PM

Room:Farabeuf

For Part 3 see MS57
For Part 5 see MS79

Future extreme scale systems are expected to suffer more frequent system faults resulting from the unprecedented scale of parallelism. Combined with the technology trend such as the imbalance between computing and I/O throughput and a tight power budget in the system operations, the traditional hardware-level redundancy and checkpoint restart may not be a feasible solution. Instead, every layer of the systems should address these faults to mitigate their impact in holistic manners. In this minisymposium, we will address the deficiencies of the current practice from a spectrum of high performance computing research from hardware, systems, runtime, algorithm and applications.

Organizer: Keita Teranishi
Sandia National Laboratories, USA

Organizer: Luc Giraud
Inria, France

Organizer: Emmanuel Agullo
Inria, France

Organizer: Francesco Rizzi
Sandia National Laboratories, USA

Organizer: Michael Heroux
Sandia National Laboratories, USA

10:35-10:55 Will Burst Buffers Save Checkpoint/Restart?
Kathryn Mohror, Lawrence Livermore National Laboratory, USA

11:00-11:20 From System-Level Checkpointing to User-Level Failure Mitigation in MPI-3
Camille Coti, Université Paris 13, France

11:25-11:45 A Flexible "Area-Time" Model for Analyzing and Managing Application Resilience
Andrew A. Chien, University of Chicago and Argonne National Laboratory, USA; Aiman Fang, University of Chicago, USA

11:50-12:10 Local Recovery at Extreme Scales with FenixLR: Implementation and Evaluation
Marc Gamell, Rutgers University, USA; Keita Teranishi and Michael Heroux, Sandia National Laboratories, USA; Manish Parashar, Rutgers University, USA

continued in next column

Friday, April 15

MS66

In Situ Methods and Infrastructure: Answers Without All the I/O

10:35 AM-12:15 PM

Room:Pasquier

Due to the widening gap between the FLOP and I/O capacity of HPC platforms, it is increasingly impractical for computer simulations to save full-resolution computations to disk for subsequent analysis. "In situ" methods offer hope for managing this increasingly acute problem by performing as much analysis, visualization, and related processing while data is still resident in memory. While in situ methods are not new, they are presently the subject of much active R&D, which aims to produce novel algorithms and production-quality infrastructure deployed at HPC facilities for use by the scientific community.

Organizer: Wes Bethel
Lawrence Berkeley National Laboratory, USA

10:35-10:55 In Situ Processing Overview and Relevance to the Hpc Community
E. Wes Bethel, Lawrence Berkeley National Laboratory, USA

11:00-11:20 Overview of Contemporary In Situ Infrastructure Tools and Architectures
Julien Jomier and Patrick O'Leary, Kitware, Inc., USA

11:25-11:45 Complex In Situ and In Transit Workflows
Matthew Wolf, Georgia Institute of Technology, USA

11:50-12:10 In Situ Algorithms and Their Application to Science Challenges
Venkatram Vishwanath, Argonne National Laboratory, USA

Friday, April 15

MS67

Parallel Programming Frameworks: Technologies, Performance and Applications - Part I of III

10:35 AM-12:15 PM

Room:Roussy

For Part 2 see MS75

Parallel programming frameworks play an important role in computational sciences. This minisymposium will focus on aspects of parallel programming frameworks, including key technologies, software design and applications. Major topics includes programming models, data structures, load balancing algorithms, performance optimization techniques, and application developments. Both structured and unstructured grid frameworks are covered, with an emphasis on AMR capabilities. It brings together researchers from Europe, US and China, to share experiences and exchange ideas on designing, developing, optimizing and the applications of parallel programming frameworks.

Organizer: Ulrich J. Ruede
University of Erlangen-Nuremberg, Germany

Organizer: Zeyao Mo
CAEP Software Center for High Performance Numerical Simulations, China

Organizer: Zhang Yang
Institute of Applied Physics and Computational Mathematics, China

10:35-10:55 Exascale Software
Hans-Joachim Bungartz, Technische Universität München, Germany; Ulrich J. Ruede, University of Erlangen-Nuremberg, Germany

11:00-11:20 Progress on High Performance Programming Frameworks for Numerical Simulations
Zeyao Mo, CAEP Software Center for High Performance Numerical Simulations, China

11:25-11:45 Addressing the Future Needs of Uintah on Post-Petascale Systems
Martin Berzins, University of Utah, USA

11:50-12:10 Can We Really Taskify the World? Challenges for Task-Based Runtime Systems Designers
Raymond Namyst, University of Bordeaux, France

continued in next column

Friday, April 15

MS68

Scalable Network Analysis: Tools, Algorithms, Applications - Part I of II

10:35 AM-12:15 PM

Room:Salle des theses

For Part 2 see MS76

Graph analysis provides tools for analyzing the irregular data sets common in health informatics, computational biology, climate science, sociology, security, finance, and many other fields. These graphs possess different structures than typical finite element meshes. Scaling graph analysis to the scales of data being gathered and created has spawned many directions of exciting new research. This minisymposium includes talks on massive graph generation for testing and evaluating parallel algorithms, novel streaming techniques, and parallel graph algorithms for new and existing problems. It also covers existing parallel frameworks and interdisciplinary applications, e.g. the analysis of climate networks.

Organizer: Henning Meyerhenke
Karlsruhe Institute of Technology, Germany

Organizer: Jason Riedy
Georgia Institute of Technology, USA

Organizer: David A. Bader
Georgia Institute of Technology, USA

10:35-10:55 (Generating And) Analyzing Large Networks with NetworKit
Henning Meyerhenke, Elisabetta Bergamini, Moritz von Looz, Roman Prutkin, and Christian L. Staudt, Karlsruhe Institute of Technology, Germany

11:00-11:20 Generating and Traversing Large Graphs in External-Memory.
Ulrich Meyer, Goethe University Frankfurt am Main, Germany

continued on next page

11:25-11:45 The GraphBLAS Effort: Kernels, API, and Parallel Implementations
Aydin Buluc, Lawrence Berkeley National Laboratory, USA

11:50-12:10 Parallel Tools to Reconstruct and Analyze Large Climate Networks
Hisham Ihshaish, University of West England, Bristol, United Kingdom; Alexis Tantet, University of Utrecht, Netherlands; Johan Dijkzeul, Delft University of Technology, Netherlands; Henk A Dijkstra, Utrecht University, The Netherlands

Friday, April 15

MS69

Large-scale Electronic Structure Calculation: Algorithms, Performance and Applications - Part III of III

10:35 AM-12:15 PM

Room:Marie Curie

For Part 2 see MS61

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 Advancing Algorithms to Increase Performance of Electronic Structure Simulations on Many-Core Architectures
Wibe A. De Jong, Lawrence Berkeley National Laboratory, USA

11:00-11:20 A Parallel Algorithm for Approximating the Single Particle Density Matrix for O(N) Electronics Structure Calculations
Daniel Osei-Kuffuor and Jean-Luc Fattebert, Lawrence Livermore National Laboratory, USA

11:25-11:45 Parallel Eigensolvers for Density Functional Theory
Antoine Levitt, Inria and Ecole des Ponts ParisTech, France

11:50-12:10 Real-Valued Technique for Solving Linear Systems Arising in Contour Integral Eigensolver
Yasunori Futamura and Tetsuya Sakurai, University of Tsukuba, Japan

Friday, April 15

MS70

Scalable Task-based Programming Models - Part II of II

10:35 AM-12:15 PM

Room:Leroux

For Part 1 see MS62

Programming models that enable users to write sequential code that can be efficiently executed in parallel, for example by a dedicated run-time system, on heterogeneous architecture, are becoming ubiquitous. Such programming models provide an easy interface to developers while exploiting the performance of the underlying infrastructure. In this minisymposium, we explore the current state-of-the-art of task-based programming models, the potential for unification, user interfaces, and special issues such as energy or resource awareness and resilience. The minisymposium is organized in two parts, with contributions from key players in task-based programming.

Organizer: Elisabeth Larsson
Uppsala University, Sweden

Organizer: Rosa M. Badia
Barcelona Supercomputing Center, Spain

10:35-10:55 Developing Task-Based Applications with Pycompss
Javier Conejero, Barcelona Supercomputing Center, Spain

11:00-11:20 Resilience in Distributed Task-Based Runtimes (2)
George Bosilca, University of Tennessee, Knoxville, USA

11:25-11:45 DuctTeip: A Framework for Distributed Hierarchical Task-Based Parallel Computing
Elisabeth Larsson and Afshin Zafari, Uppsala University, Sweden

11:50-12:10 Heterogeneous Computing with GPUs and FPGAs Through the OpenMP Task-Parallel Model
Artur Podobas, KTH Royal Institute of Technology, Sweden

Friday, April 15

MS71

High Performance Simulations for Subsurface Applications

10:35 AM-12:15 PM

Room:Dejerine

Simulation of underground flow and transport has several important applications such as oil reservoirs, deep geological storage of nuclear waste, or carbon dioxide or geothermal energy. Common to these areas is the complexity of the phenomena involved (multi-phase flow, chemistry, geomechanics, thermics, surface flow) that are coupled, and the heterogeneity of the geological medium over several scales. High performance computing has thus been recognized as one important way for increasing the realism of these simulations. This minisymposium will gather researchers whose collective experience covers a large part of the above domains, or who have developed tools enabling better, faster, or more detailed numerical simulations.

Organizer: Michel Kern
Inria & Maison de la Simulation, France

Organizer: Roland Masson
University of Nice, France

Organizer: Feng Xing
University of Nice Sophia-Antipolis and Inria, France

10:35-10:55 Hybrid Dimensional Multiphase Compositional Darcy Flows in Fractured Porous Media and Parallel Implementation in Code ComPASS
Feng Xing, University of Nice Sophia-Antipolis and Inria, France; Roland Masson, University of Nice, France

continued in next column

11:00-11:20 PFLOTRAN: Practical Application of High Performance Computing to Subsurface Simulation
Glenn Hammond, Sandia National Laboratories, USA

11:25-11:45 Parallel Performance of a Reservoir Simulator in the Context of Unstable Waterflooding
Romain De Loubens, Total, France; *Pascal Henon*, Total E&P, France; Pavel Jiranek, CERFACS, France

11:50-12:10 Ongoing Development and HPSC Applications with the ParFlow Hydrologic Model
Stefan J. Kollet, Agrosphere Institute, Germany; Klaus Goergen, Jülich Supercomputing Centre, Germany; Carsten Burstedde and Jose Fonseca, Universität Bonn, Germany; Fabian Gasper, Agrosphere Research Center Jülich, Germany; Reed M. Maxwell, Colorado School of Mines, USA; Prabhakar Shresta and Mauro Sulis, Universität Bonn, Germany

Friday, April 15

MS72

Scientific Computing with Parallel Adaptive Mesh Refinement Methods

10:35 AM-12:15 PM

Room:Delarue

In the past years the popularization of important parallel computational resources have helped spreading parallel Adaptive Mesh Refinement (AMR) across a large range of applications that use scientific computing. We propose in this minisymposium to present a sample of these applications. We shall focus on the specific AMR method that was chosen, its implementation in a parallel framework and how these choices impact the discretization strategy.

Organizer: Samuel Kokh
Maison de la Simulation, France

10:35-10:55 Bb-Amr and Numerical Entropy Production
Frédéric Golay, Universite de Toulon, France

11:00-11:20 Title Not Available
Pierre Kestener, Maison de la Simulation, France

11:25-11:45 Title Not Available
Donna Calhoun, Boise State University, USA

11:50-12:10 Title Not Available
Thierry Coupez, Ecole Centrale de Nantes, France

Friday, April 15

CP12

Advanced Simulations - I of III

10:35 AM-12:15 PM

Room:Danton

Chair: Sever Hirstoaga, Inria, France

10:35-10:50 Parallel in Time Solvers for DAEs
Matthieu Lecouvez, Robert D. Falgout, and Carol S. Woodward, Lawrence Livermore National Laboratory, USA

10:55-11:10 Implicitly Dealiased Convolutions for Distributed Memory Architectures
Malcolm Roberts, University of Strasbourg, France; John C. Bowman, University of Alberta, Canada

11:15-11:30 A Massively Parallel Simulation of the Coupled Arterial Cells in Branching Network
Mohsin A. Shaikh, Pawsey Supercomputing Center, Australia; Constantine Zakkaroff, University of Canterbury, New Zealand; Stephen Moore, University of Melbourne, Australia; Tim David, University of Canterbury, New Zealand

11:35-11:50 Load Balancing Issues in Particle In Cell Simulations
Arnaud Beck, Ecole Polytechnique, France; Julien Derouillat, Maison de la Simulation, France

11:55-12:10 Particle-in-Cell Simulations for Nonlinear Vlasov-Like Models
Sever Hirstoaga, Inria, France; Edwin Chacon-Golcher, Institute of Physics of the ASCR, Prague, Czech Republic

Lunch Break

12:15 PM-1:45 PM

Attendees on their own

Friday, April 15

IP7

A Systems View of High Performance Computing

1:45 PM-2:30 PM

Room:Farabeuf

Chair: Bruce Hendrickson, Sandia National Laboratories, USA

HPC systems comprise increasingly many fast components, all working together on the same problem concurrently. Yet, HPC efficient-programming dogma dictates that communicating is too expensive and should be avoided. Program ever more components to work together with ever less frequent communication?!? Seriously? I object! In other disciplines, deriving efficiency from large Systems requires an approach qualitatively different from that used to derive efficiency from small Systems. Should we not expect the same in HPC? Using Grappa, a latency-tolerant runtime for HPC systems, I will illustrate through examples that such a qualitatively different approach can indeed yield efficiency from many components without stifling communication.

Simon Kahan
Institute for Systems Biology and University of Washington, USA

Intermission

2:30 PM-2:40 PM

Friday, April 15

MS73

How Will Next-generation Architectures Change How We Reason About and Write Applications?

2:40 PM-4:20 PM

Room:Farabeuf

Both European and US research programs have ambitious energy goals for extreme-scale computing - 200 petaFLOPS at 10 MW to exaFLOPS at 20 MW. Achieving performance within energy constraints requires fundamentally new architectures, potentially including heterogeneous CPUs with network-on-chip, multi-level memories, and optical interconnects. While accelerators, e.g, have started to influence programming models, many applications are agnostic to energy and hardware trends. To what extent will achieving energy-efficiency and utilizing new architectures be an algorithms problem? What can be achieved automatically at the runtime-level versus application-level?Leading architecture researchers demystify emerging hardware trends into specific consequences for algorithm design.

Organizer: Jeremiah Wilke
Sandia National Laboratories, USA

Organizer: Arun Rodrigues
Sandia National Laboratories, USA

2:40-3:00 Helping Applications Reason About Next-Generation Architectures Through Abstract Machine Models
John Shalf and David Donofrio, Lawrence Berkeley National Laboratory, USA

3:05-3:25 How Will Application Developers Navigate the Cost/ Peformance Trade-Offs of New Multi-Level Memory Architectures?
Arun Rodrigues, Sandia National Laboratories, USA

3:30-3:50 When Is Energy Not Equal to Time? Understanding Application Energy Scaling with EAudit
Sudhakar Yalamanchili and Eric Anger, Georgia Institute of Technology, USA

3:55-4:15 Performance-Energy Trade-Offs in Reconfigurable Optical Data-Center Networks Using Silicon Photonics
Sebastien Rumley and Keren Bergman, Columbia University, USA

Friday, April 15

MS74

Parallel Tree-Code Algorithms for Large Scale and Multiphysics Simulations - Part I of II

2:40 PM-4:20 PM

Room:Pasquier

For Part 2 see MS82

This minisymposium aims to bring together computational scientists that work on large- scale tree-codes and their applications. Tree-codes are hierarchical data structures that are used for nonuniform discretization of partial differential equations. Tree-codes are also used for integral-equation formulations and N-body methods. Designing efficient parallel tree-code algorithms that scale to thousands of processors is challenging. It is even harder to account for multiple scales and multiple physics due to the different resolution requirements of the underlying spatiotemporal fields. As a result load balancing, minimization of communication, and efficient memory utilization become even harder than the single field case.

Organizer: Arash Bakhtiari
Technische Universitaet Muenchen, Germany

Organizer: George Biros
University of Texas at Austin, USA

2:40-3:00 A Distributed-Memory Advection-Diffusion Solver with Dynamic Adaptive Trees
Arash Bakhtiari, Technische Universitaet Muenchen, Germany; Dhairya Malhotra and George Biros, University of Texas at Austin, USA

3:05-3:25 A Volume Integral Equation Solver for Elliptic PDEs in Complex Geometries
Dhairya Malhotra and George Biros, University of Texas at Austin, USA

3:30-3:50 Efficient Parallel Simulation of Flows, Particles, and Structures
Miriam Mehl and Michael Lahnert, Universität Stuttgart, Germany; Benjamin Uekermann, Technische Universität München, Germany

3:55-4:15 Communication-Reducing and Communication-Avoiding Techniques for the Particle-in-Dual-Tree Storage Scheme
Tobias Weinzierl, Durham University, United Kingdom

Friday, April 15

MS75

Parallel Programming Frameworks: Technologies, Performance and Applications - Part II of III

2:40 PM-4:20 PM

Room:Roussy

For Part 1 see MS67
For Part 3 see MS83

Parallel programming frameworks play an important role in computational sciences. This minisymposium will focus on aspects of parallel programming frameworks, including key technologies, software design and applications. Major topics includes programming models, data structures, load balancing algorithms, performance optimization techniques, and application developments. Both structured and unstructured grid frameworks are covered, with an emphasis on AMR capabilities. It brings together researchers from Europe, US and China, to share experiences and exchange ideas on designing, developing, optimizing and the applications of parallel programming frameworks.

Organizer: Ulrich J. Ruede
University of Erlangen-Nuremberg, Germany

Organizer: Zeyao Mo
CAEP Software Center for High Performance Numerical Simulations, China

Organizer: Zhang Yang
Institute of Applied Physics and Computational Mathematics, China

2:40-3:00 Cello: An Extreme Scale Amr Framework, and Enzo-P, An Application for Astrophysics and Cosmology Built on Cello
Michael L. Norman, University of California, San Diego, USA

continued in next column

3:05-3:25 Dune - A Parallel Pde Framework
Christian Engwer, University of Münster, Germany; Steffen Müthing and Peter Bastian, Universität Heidelberg, Germany

3:30-3:50 Peta-scale Simulations with the HPC Software Framework Walberla
Florian Schornbaum and Martin Bauer, University of Erlangen-Nuremberg, Germany; Christian Godenschwager, University of Erlangen-Nuernberg, Germany; Ulrich R\"ude, University of Erlangen-Nuremberg, Germany

3:55-4:15 PHG: A Framework for Parallel Adaptive Finite Element Method
Deng Lin, Wei Leng, and Tao Cui, Chinese Academy of Sciences, China

Friday, April 15

MS76

Scalable Network Analysis: Tools, Algorithms, Applications - Part II of II

2:40 PM-4:20 PM

Room:Salle des theses

For Part 1 see MS68

Graph analysis provides tools for analyzing the irregular data sets common in health informatics, computational biology, climate science, sociology, security, finance, and many other fields. These graphs possess different structures than typical finite element meshes. Scaling graph analysis to the scales of data being gathered and created has spawned many directions of exciting new research. This minisymposium includes talks on massive graph generation for testing and evaluating parallel algorithms, novel streaming techniques, and parallel graph algorithms for new and existing problems. It also covers existing parallel frameworks and interdisciplinary applications, e.g. the analysis of climate networks.

Organizer: Henning Meyerhenke
Karlsruhe Institute of Technology, Germany

Organizer: Jason Riedy
Georgia Institute of Technology, USA

Organizer: David A. Bader
Georgia Institute of Technology, USA

2:40-3:00 Scalable and Efficient Algorithms for Analysis of Massive, Streaming Graphs

Jason Riedy and David A. Bader, Georgia Institute of Technology, USA

3:05-3:25 Communication Avoiding Ilu Preconditioner

Sebastien Cayrols and Laura Grigori, Inria, France

3:30-3:50 Parallel Approximate Graph Matching

Fanny Dufossé, Inria Lille, France; *Kamer Kaya*, Sabanci University, Turkey; Bora Ucar, LIP-ENS Lyon, France

3:55-4:15 Diamond Sampling for Approximate Maximum All-pairs Dot-product (MAD) Search

Grey Ballard, C. Seshadhri, Tamara G. Kolda, and *Ali Pinar*, Sandia National Laboratories, USA

Friday, April 15

MS77

Low-Rank Approximations for Fast Linear Algebra Computations - Part I of II

2:40 PM-4:20 PM

Room:Marie Curie

For Part 2 see MS85

Low-rank representations are a popular way of speeding-up matrix algorithms. They can be used for designing fast matrix-vector products, direct solvers with linear or near-linear complexity, and robust preconditioners, both in the dense and the sparse case. Many different approaches, such as H-matrices, HSS representations, or the BLR format, are currently under study by different research groups. The minisymposium features researchers working on all these techniques. They will present recent algorithmic ideas as well as software packages.

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

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

2:40-3:00 A Comparison of Parallel Rank-Structured Solvers

Francois-Henry Rouet, Lawrence Berkeley National Laboratory, USA; Patrick Amestoy, Université of Toulouse, France; Cleve Ashcraft, Livermore Software Technology Corporation, USA; Alfredo Buttari, CNRS, France; Pieter Ghysels, Lawrence Berkeley National Laboratory, USA; Jean-Yves L'Excellent, Inria-LIP-ENS Lyon, France; Xiaoye Sherry Li, Lawrence Berkeley National Laboratory, USA; Theo Mary, Universite de Toulouse, France; Clément Weisbecker and Clément Weisbecker, Livermore Software Technology Corporation, USA

continued in next column

3:05-3:25 Complexity and Performance of Algorithmic Variants of Block Low-Rank Multifrontal Solvers

Patrick Amestoy, ENSEEIHT-IRIT, France; Alfredo Buttari, CNRS, France; Jean-Yves L'Excellent, Inria-LIP-ENS Lyon, France; *Theo Mary*, Universite de Toulouse, France

3:30-3:50 Hierarchical Matrix Arithmetics on Gpus

Wajih Halim Boukaram, Hatem Ltaief, *George M Turkiyyah*, and David E. Keyes, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

3:55-4:15 A Block Low-Rank Multithreaded Factorization for Dense BEM Operators.

Julie Anton, Cleve Ashcraft, and Clément Weisbecker, Livermore Software Technology Corporation, USA

Friday, April 15

MS78

Scheduling Techniques in Parallel Sparse Matrix Solvers

2:40 PM-4:20 PM

Room:Leroux

In this minisymposium, we focus on the scheduling problem in parallel sparse matrix solvers. Scheduling is likely to play a crucial role towards solvers that are scalable on extremely large platforms. On such platforms, communications will be relatively expensive and subject to performance variation. Overlapping them with computations becomes therefore of prime importance. Efficient scheduling frameworks are likely to be extremely valuable with respect to this objective, as well as tolerating with performance variability. This minisymposium will focus on both static and dynamic scheduling techniques in the context of sparse direct solvers.

Organizer: Mathias Jacquelin
Lawrence Berkeley National Laboratory, USA

Organizer: Esmond G. Ng
Lawrence Berkeley National Laboratory, USA

2:40-3:00 Optimizing Memory Usage When Scheduling Tree-Shaped Task Graphs of Sparse Multifrontal Solvers
Guillaume Aupy, ENS Lyon, France; Lionel Eyraud-Dubois, Inria, France; *Loris Marchal*, Centre National de la Recherche Scientifique, France; Frédéric Vivien, ENS Lyon, France; Oliver Sinnen, University of Auckland, New Zealand

3:05-3:25 Impact of Blocking Strategies for Sparse Direct Solvers on Top of Generic Runtimes
Gregoire Pichon, Inria, France; Mathieu Faverge, Bordeaux INP, Inria, LaBRI, France; *Pierre Ramet*, Université de Bordeaux, Inria, LaBRI, France; Jean Roman, Inria, France

continued in next column

3:30-3:50 Traversing a BLR Factorization Task Dag Based on a Fan-All Wraparound Map
Julie Anton, *Cleve Ashcraft*, and Clément Weisbecker, Livermore Software Technology Corporation, USA

3:55-4:15 Scheduling Sparse Symmetric Fan-Both Cholesky Factorization
Mathias Jacquelin, Esmond G. Ng, and Yili Zheng, Lawrence Berkeley National Laboratory, USA; Katherine Yelick, Lawrence Berkeley National Laboratory and University of California Berkeley, USA

Friday, April 15

MS79

Resilience Toward Exascale Computing - Part V of VI

2:40 PM-4:20 PM

Room:Dejerine

For Part 4 see MS65
For Part 6 see MS87
Future extreme scale systems are expected to suffer more frequent system faults resulting from the unprecedented scale of parallelism. Combined with the technology trend such as the imbalance between computing and I/O throughput and a tight power budget in the system operations, the traditional hardware-level redundancy and checkpoint restart may not be a feasible solution. Instead, every layer of the systems should address these faults to mitigate their impact in holistic manners. In this minisymposium, we will address the deficiencies of the current practice from a spectrum of high performance computing research from hardware, systems, runtime, algorithm and applications.

Organizer: Keita Teranishi
Sandia National Laboratories, USA

Organizer: Luc Giraud
Inria, France

Organizer: Emmanuel Agullo
Inria, France

Organizer: Francesco Rizzi
Sandia National Laboratories, USA

Organizer: Michael Heroux
Sandia National Laboratories, USA

2:40-3:00 Soft Errors in PCG: Detection and Correction
Emmanuel Agullo, Inria, France; Siegfried Cools, Antwerp University, Belgium; Luc Giraud, Inria, France; Wim Vanroose, Antwerp University, Belgium; *Fatih Yetkin*, Inria, France

3:05-3:25 Fault Tolerant Sparse Iterative Solvers
Simon McIntosh-Smith, University of Bristol, United Kingdom

3:30-3:50 Silent Data Corruption in BLAS Kernels
Wilfried N. Gansterer, University of Vienna, Austria

3:55-4:15 Algorithm-Based Fault Tolerance of the Sakurai-Sugiura Method
Tetsuya Sakurai, Akira Imakura, and Yasunori Futamura, University of Tsukuba, Japan

Friday, April 15

MS80

Towards Reliable and Efficient Scientific Computing at Exascale - Part I of II

2:40 PM-4:20 PM

Room:Delarue

For Part 2 see MS88

Exascale supercomputers will bring huge computational power. Unfortunately, their architectures have to be complex and heterogeneous to counteract the memory and power walls. Consequently, scientific codes designed for homogeneous supercomputers will have a low performance on Exascale supercomputers. Besides, intensive floating point computation and massive concurrency can degrade accuracy and reproducibility. Therefore arises a strong need for tools for efficient and reliable port of scientific codes. The aim of this minisymposium is to present jointly some research works around performance characterization and verification of floating point accuracy.

Organizer: Christophe Denis
CMLA, ENS de Cachan, France

2:40-3:00 Checking the Floating Point Accuracy of Scientific Code Through the Monte Carlo Arithmetic
Christophe Denis, CMLA, ENS de Cachan, France

3:05-3:25 Code Verification in Cfd
Daniel Chauveheid, Bruyères-le-Châtel, France

3:30-3:50 Performance Study of Industrial Hydrocodes: Modeling, Analysis and Aftermaths on Innovative Computational Methods
Florian De Vuyst, Ecole Normale Superieur de Cachan, France

3:55-4:15 Code Modernization by Mean of Proto-Applications
Eric Petit, PRiSM - UVSQ, France

Friday, April 15

CP13

Advanced Simulations - II of III

2:40 PM-3:40 PM

Room:Danton

Chair: Francois-Xavier Roux, ONERA, France

2:40-2:55 Efficient Two-Level Parallel Algorithm Based on Domain Decomposition Method for Antenna Simulation
Francois-Xavier Roux, ONERA, France

3:00-3:15 {GPU} Acceleration of Polynomial Homotopy Continuation
Jan Verschelde and Xiangcheng Yu, University of Illinois, Chicago, USA

3:20-3:35 Adaptive Transpose Algorithms for Distributed Multicore Processors
John C. Bowman, University of Alberta, Canada; *Malcolm Roberts*, University of Strasbourg, France

Coffee Break

4:20 PM-4:50 PM

Room:Exhibition Space

Friday, April 15

MS81

Evolution of Co-Design and Scientific HPC Applications

4:50 PM-6:30 PM

Room:Farabeuf

Modernizing scientific codes to run performantly on multiple emerging architectures is a task of great complexity and challenge for computational scientists today. It is typical the set of performance necessary optimizations for a code on a given platform will be distinct from those for another. To ameliorate this difficulty, we turn to abstraction layers and alternative programming models. This minisymposium will present both experiments with proxy applications and full applications towards the goal of developing strategies for performance portability on current and future HPC platforms.

Organizer: Geoff Womeldorff
Los Alamos National Laboratory, USA

Organizer: Ian Karlin
Lawrence Livermore National Laboratory, USA

4:50-5:10 On the Performance and Productivity of Abstract C++ Programming Models?
Simon D. Hammond, Christian Trott, *H. Carter Edwards*, Jeanine Cook, Dennis Dinge, Mike Glass, Robert J. Hoekstra, Paul Lin, Mahesh Rajan, and Courtenay T. Vaughan, Sandia National Laboratories, USA

5:15-5:35 Demonstrating Advances in Proxy Applications Through Performance Gains and/or Performance Portable Abstractions: CoMD and Kripke with RAJA
Richard Hornung, Olga Pearce, Adam Kunen, Jeff Keasler, Holger Jones, Rob Neely, and *Todd Gamblin*, Lawrence Livermore National Laboratory, USA

continued on next page

Friday, April 15

MS81

Evolution of Co-Design and Scientific HPC Applications

4:50 PM-6:30 PM

Room:Farabeuf

continued

5:40-6:00 Kokkos and Legion Implementations of the SNAP Proxy Application
Geoff Womeldorff, Ben Bergen, and Joshua Payne, Los Alamos National Laboratory, USA

6:05-6:25 Future Architecture Influenced Algorithm Changes and Portability Evaluation Using BLAST
Christopher Earl, Barna Bihari, Veselin Dobrev, *Ian Karlin*, Tzanio V. Kolev, Robert Rieben, and Vladimir Tomov, Lawrence Livermore National Laboratory, USA

Friday, April 15

MS82

Parallel Tree-Code Algorithms for Large Scale and Multiphysics Simulations - Part II of II

4:50 PM-6:30 PM

Room:Pasquier

For Part 1 see MS74

This minisymposium aims to bring together computational scientists that work on large- scale tree-codes and their applications. Tree-codes are hierarchical data structures that are used for nonuniform discretization of partial differential equations. Tree- codes are also used for integral-equation formulations and N-body methods. Designing efficient parallel tree-code algorithms that scale to thousands of processors is challenging. It is even harder to account for multiple scales and multiple physics due to the different resolution requirements of the underlying spatiotemporal fields. As a result load balancing, minimization of communication, and efficient memory utilization become even harder than the single field case.

Organizer: Arash Bakhtiari
Technische Universitaet Muenchen, Germany

Organizer: George Biros
University of Texas at Austin, USA

4:50-5:10 Recent Developments on Forests of Linear Trees
Carsten Burstedde and Johannes Holke, Universität Bonn, Germany

5:15-5:35 Exploiting Octree Meshes with a High Order Discontinuous Galerkin Method
Sabine P. Roller, Harald Klimach, and Jens Zudrop, Universitaet Siegen, Germany

5:40-6:00 Parallel Level-Set Algorithms on Tree Grids
Mohammad Mirzadeh, Massachusetts Institute of Technology, USA; Arthur Guittet, University of California, Santa Barbara, USA; Carsten Burstedde, Universität Bonn, Germany; Frederic G. Gibou, University of California, Santa Barbara, USA

6:05-6:25 A Sharp Numerical Method for the Simulation of Diphasic Flows
Maxime Theillard, University of California, San Diego, USA; Frederic Gibou, University of California, Santa Barbara, USA; David Saintillan, University of California, San Diego, USA; Arthur Guittet, University of California, Santa Barbara, USA

continued in next column

Friday, April 15

MS83

Parallel Programming Frameworks: Technologies, Performance and Applications - Part III of III

4:50 PM-6:05 PM

Room:Roussy

For Part 2 see MS75

Parallel programming frameworks play an important role in computational sciences. This minisymposium will focus on aspects of parallel programming frameworks, including key technologies, software design and applications. Major topics includes programming models, data structures, load balancing algorithms, performance optimization techniques, and application developments. Both structured and unstructured grid frameworks are covered, with an emphasis on AMR capabilities. It brings together researchers from Europe, US and China, to share experiences and exchange ideas on designing, developing, optimizing and the applications of parallel programming frameworks.

Organizer: Ulrich J. Ruede
University of Erlangen-Nuremberg, Germany

Organizer: Zeyao Mo
CAEP Software Center for High Performance Numerical Simulations, China

Organizer: Zhang Yang
Institute of Applied Physics and Computational Mathematics, China

4:50-5:10 Design and Implementation of Adaptive Spmv Library for Multicore and Manycore Architecture
Guangming Tan, Chinese Academy of Sciences, China

5:15-5:35 Towards Efficient Data Communications for Frameworks on Post-petascale Numa-Multicore Systems
Zhang Yang and Aiqing Zhang, Institute of Applied Physics and Computational Mathematics, China

5:40-6:00 Optimal Domain Decomposition Strategies for Parallel Geometric Multigrid
Gaurav Saxena, Peter K. Jimack, and Mark A. Walkley, University of Leeds, United Kingdom

Friday, April 15

MS84

Acclerating HPC and Datacenter Workloads with FPGAs

4:50 PM-6:30 PM

Room:Salle des theses

The importance of power efficiency has been changing the architectural landscape of high performance computing systems. Heterogeneous computing with accelerators such as GPUs gained popularity thanks to its superior power efficiency for typical HPC workloads. Recent technological developments such as hardened floating point units as well as the end of Moore's Law, however, are making reconfigurable chips such as FPGAs increasingly promising in a wide range of workloads in HPC and datacenter systems. In this minisymposium, we will discuss performance and programmability of a single FPGA to large-scale multi-FPGA systems.

Organizer: Naoya Maruyama
RIKEN, Japan

4:50-5:10 Benchmarking FPGAs with Opencl
Naoya Maruyama, RIKEN, Japan; Hamid Zohouri, Tokyo Institute of Technology, Japan; Aaron Smith, Microsoft Research, USA; Motohiko Matsuda, RIKEN Advanced Institute for Computational Science, Japan; Satoshi Matsuoka, Tokyo Institute of Technology, Japan

5:15-5:35 Architectural Possibilities of Fpga-Based High-Performance Custom Computing
Kentaro Sano, Tohoku University, Japan

5:40-6:00 Novo-G#: Strong Scaling Using Fpga-Centric Clusters with Direct and Programmable Interconnects
Martin Herbordt, Boston University, USA

6:05-6:25 Are Datacenters Ready for Reconfigurable Logic?
Aaron Smith, Microsoft Research, USA

Friday, April 15

MS85

Low-Rank Approximations for Fast Linear Algebra Computations - Part II of II

4:50 PM-6:30 PM

Room:Marie Curie

For Part 1 see MS77

Low-rank representations are a popular way of speeding-up matrix algorithms. They can be used for designing fast matrix-vector products, direct solvers with linear or near-linear complexity, and robust preconditioners, both in the dense and the sparse case. Many different approaches, such as H-matrices, HSS representations, or the BLR format, are currently under study by different research groups. The minisymposium features researchers working on all these techniques. They will present recent algorithmic ideas as well as software packages.

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

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

4:50-5:10 Exploiting H-Matrices in Sparse Direct Solvers
Gregoire Pichon, Inria, France; Eric F. Darve, Stanford University, USA; Mathieu Faverge, Bordeaux INP, Inria, LaBRI, France; Pierre Ramet, Université de Bordeaux, Inria, LaBRI, France; Jean Roman, Inria, France

5:15-5:35 Fast Updating of Skeletonization-Based Hierarchical Factorizations
Victor Minden, Anil Damle, Ken Ho, and Lexing Ying, Stanford University, USA

continued on next page

Friday, April 15

MS85

Low-Rank Approximations for Fast Linear Algebra Computations - Part II of II

4:50 PM-6:30 PM

Room:Marie Curie

continued

5:40-6:00 Algebraic and Geometric Low-Rank Approximations
Rio Yokota, Tokyo Institute of Technology, Japan; François-Henry Rouet and Xiaoye S. Li, Lawrence Berkeley National Laboratory, USA; David E. Keyes, King Abdullah University of Science & Technology (KAUST), Saudi Arabia

6:05-6:25 Parallel Performance of the Inverse FMM and Application to Mesh Deformation for Fluid-Structure Problems
Pieter Coulier, Chao Chen, and Eric F. Darve, Stanford University, USA

Friday, April 15

MS86

Code Generation for High Performance Geoscientific Simulation

4:50 PM-6:30 PM

Room:Leroux

The demands of ever-more realstic simulations, more sophisticated numerical techniques, and parallelism at multiple granularities leads to an explosion in the complexity of simulation software. Code generation approaches enable the complexity of these three layers to be decoupled, so that sophisticated techiques developed by experts in domain science, numerical analysis and computer science can be composed together to create ever-more capabile simulation toolchains. This minisymposium focusses on the development and use of these code generation techniques to achieve high performance in simulation, with a particular focus on the computational issues created by geoscience applications.

Organizer: David Ham
Imperial College London, United Kingdom

Organizer: William Sawyer
Swiss Centre of Scientific Computing, Switzerland

4:50-5:10 Generating High Performance Finite Element Kernels Using Optimality Criteria.
Fabio Luporini, David Ham, and Paul Kelly, Imperial College London, United Kingdom

5:15-5:35 Architecture-Agnostic Numerics in GridTools
Paolo Crosetto, École Polytechnique Fédérale de Lausanne, Switzerland

5:40-6:00 Kernel Optimization for High-Order Dynamic Rupture Simulation - Flexibility Vs. Performance
Carsten Uphoff and Michael Bader, Technische Universität München, Germany

6:05-6:25 Source-to-Source Translation for Code Validation and Acceleration of the Icon Atmospheric General Circulation Model
William Sawyer, Swiss Centre of Scientific Computing, Switzerland

Friday, April 15

MS87

Resilience Toward Exascale Computing - Part VI of VI

4:50 PM-6:30 PM

Room:Dejerine

For Part 5 see MS79

Future extreme scale systems are expected to suffer more frequent system faults resulting from the unprecedented scale of parallelism. Combined with the technology trend such as the imbalance between computing and I/O throughput and a tight power budget in the system operations, the traditional hardware-level redundancy and checkpoint restart may not be a feasible solution. Instead, every layer of the systems should address these faults to mitigate their impact in holistic manners. In this minisymposium, we will address the deficiencies of the current practice from a spectrum of high performance computing research from hardware, systems, runtime, algorithm and applications.

Organizer: Keita Teranishi
Sandia National Laboratories, USA

Organizer: Luc Giraud
Inria, France

Organizer: Emmanuel Agullo
Inria, France

Organizer: Francesco Rizzi
Sandia National Laboratories, USA

Organizer: Michael Heroux
Sandia National Laboratories, USA

4:50-5:10 {FA-MPI}: Fault-Aware and QoS-Oriented Scalable Message Passing
Anthony Skjellum, Auburn University, USA

5:15-5:35 Spare Node Substitution of Failed Node
Atsushi Hori, Kazumi Yoshinaga, and Yutaka Ishikawa, RIKEN Advanced Institute for Computational Science, Japan

5:40-6:00 Reinit: A Simple and Scalable Fault-Tolerance Model for {MPI} Applications
Ignacio Laguna, Lawrence Livermore National Laboratory, USA

6:05-6:25 Title Not Available
George Bosilca, University of Tennessee, Knoxville, USA

Friday, April 15

MS88

Towards Reliable and Efficient Scientific Computing at Exascale - Part II of II

4:50 PM-6:30 PM

Room:Delarue

For Part 1 see MS80

Exascale supercomputers will bring huge computational power. Unfortunately, their architectures have to be complex and heterogeneous to counteract the memory and power walls. Consequently, scientific codes designed for homogeneous supercomputers will have a low performance on Exascale supercomputers. Besides, intensive floating point computation and massive concurrency can degrade accuracy and reproducibility. Therefore arises a strong need for tools for efficient and reliable port of scientific codes. The aim of this minisymposium is to present jointly some research works around performance characterization and verification of floating point accuracy.

Organizer: Christophe Denis
CMLA, ENS de Cachan, France

4:50-5:10 Cere: Llvm Based Codelet Extractor and Replayer for Piecewise Benchmarking and Optimization
Pablo De Oliveira Castro, PRiSM - UVSQ, France

5:15-5:35 Using Uncertainties in High Level Performance Prediction
Philippe Thierry, Intel Corporation, France

5:40-6:00 Performance Modeling of a Compressible Hydrodynamics Solver on Multicore Cpus
Mathieu Peybernes, CEA Saclay, France

6:05-6:25 Topology-Aware Performance Modeling of Parallel SpMVs
Amanda Bienz, Luke Olson, and William D. Gropp, University of Illinois at Urbana-Champaign, USA

Friday, April 15

CP14

Advanced Simulations - III of III

4:50 PM-6:10 PM

Room:Danton

Chair: Ioan C. Hadade, Imperial College London, United Kingdom

4:50-5:05 Robust Domain Decomposition Methods for Structural Mechanics Computations
Christophe Bovet, CNRS, Université Paris-Saclay, France; Augustin Parret-Fréaud, Safran Tech, France; Pierre Gosselet, LMT-Cachan, France; Nicole Spillane, Ecole Polytechnique, France

5:10-5:25 A Parallel Contact-Impact Simulation Framework
Jie Cheng and Qingkai Liu, Institute of Applied Physics and Computational Mathematics, China

5:30-5:45 Optimization of the Photon's Simulation on GPU
Clement P. Rey, Benjamin Auer, Emmanuel Medernach, and Ziad El Bitar, CNRS, Université de Strasbourg, France

5:50-6:05 Turbomachinery CFD on Modern Multicore and Manycore Architectures
Ioan C. Hadade and Luca Di Mare, Imperial College London, United Kingdom

Conference Adjourns

6:30 PM

Notes

Abstracts

SIAM Conference on Parallel Processing for Scientific Computing

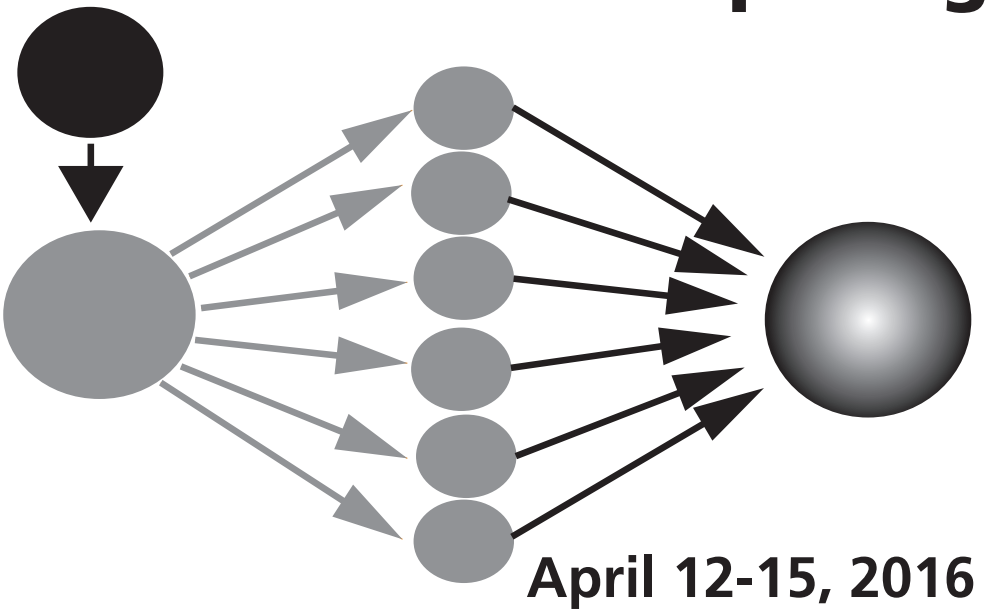

April 12-15, 2016
Université Pierre et Marie Curie, Cordeliers Campus
Paris, France

Abstracts are printed as submitted by the authors.

Notes

Speaker and Organizer Index

SIAM Conference on
Parallel Processing for Scientific Computing

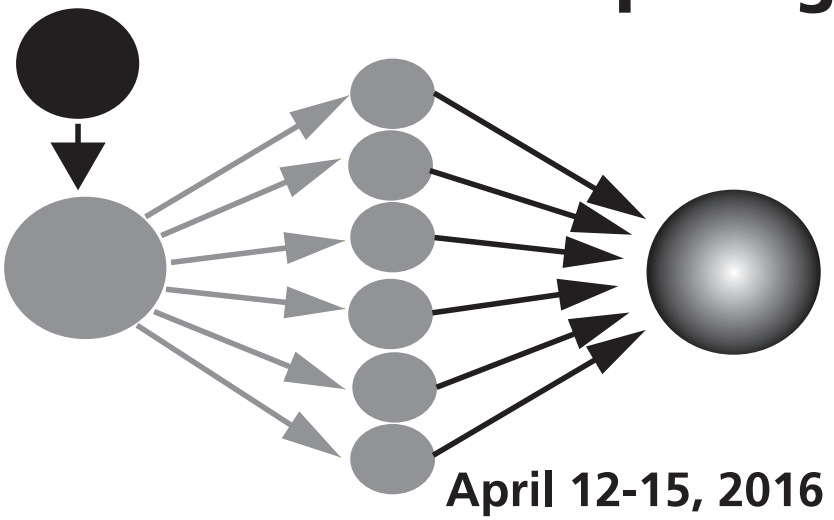

April 12-15, 2016
Université Pierre et Marie Curie, Cordeliers Campus
Paris, France

D

E

F

G

H

I

J

K

L

M

N

O

P

R

S

T

U

V

W

X

Y

Z

Notes

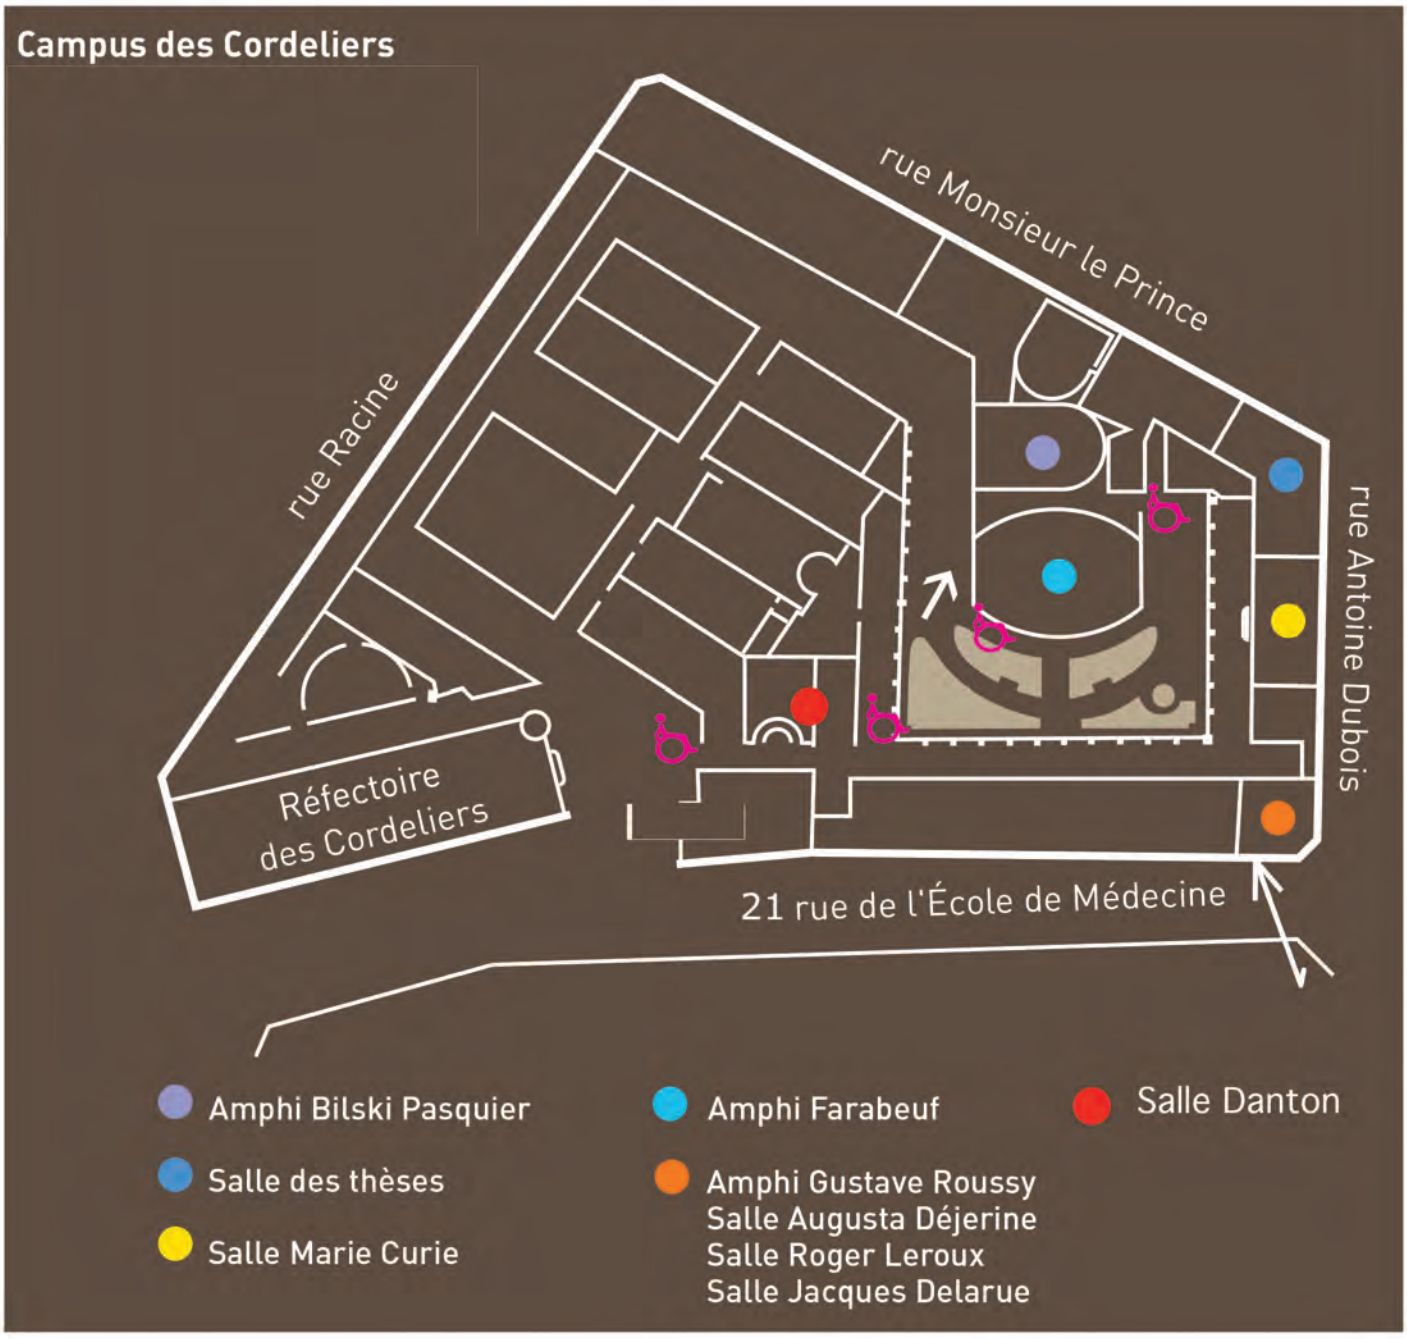
Campus des Cordeliers
rue Monsieur le Prince
rue Racine
rue Antoine Dubois
Réfectoire
des Cordeliers
21 rue de l'École de Médecine
Amphi Bilski Pasquier
Salle des thèses
Salle Marie Curie
Amphi Farabeuf
Amphi Gustave Roussy
Salle Augusta Déjerine
Salle Roger Leroux
Salle Jacques Delarue
Salle Danton