IP1

Scalability of Sparse Direct Codes

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
Rutherford Appleton Laboratory, Oxfordshire, UK OX11 0QX
and CERFACS, Toulouse, France
iain.duff@stfc.ac.uk

IP2

Towards Solving Coupled Flow Problems at the Extreme Scale

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
Department of Computer Science (Simulation)
ulrich.ruede@fau.de

IP3

Why Future Systems should be Data Centric

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
coteus@us.ibm.com

IP4

Next-Generation AMR

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 Labs 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
asalmgren@lbl.gov

IP5

Beating Heart Simulation Based on a Stochastic Biomolecular Model

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
washio@sml.k.u-tokyo.ac.jp

IP6

Conquering Big Data with Spark

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
istoica@cs.berkeley.edu

IP7

A Systems View of High Performance Computing

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
simon.kahan@gmail.com

SP0

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

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.

James W. Demmel
University of California
Division of Computer Science
demmel@berkeley.edu

Laura Grigori
INRIA
France
Laura.Grigori@inria.fr

Mark Hoemmen
Sandia National Laboratories
mhoemme@sandia.gov

Julien Langou
University of Colorado Denver
julien.langou@ucdenver.edu

SP1

SIAG/Supercomputing Career Prize: Supercomputers and Superintelligence

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 Laboratory
hdsimon@lbl.gov

CP1

Scale-Bridging Modeling of Material Dynamics: Petascale Assessments on the Road to Exascale

Within the multi-institutional, multi-disciplinary Exascale Co-design Center for Materials in Extreme Environments (ExMatEx), we are developing adaptive physics refinement methods that exploit hierarchical, heterogeneous architectures to achieve more realistic large-scale simulations. To evaluate and exercise the task-based programming models, databases, and runtime systems required to perform many-task computation workflows that accumulate a distributed response database from fine-scale calculations, we are preforming petascale demonstrations on the Trinity supercomputer, combining Haswell, Knights Landing, and burst buffer nodes.

Timothy C. Germann
Los Alamos National Laboratory
tcg@lanl.gov

CP1

Parallel Recursive Density Matrix Expansion in Electronic Structure Calculations

We propose a parallel recursive polynomial expansion algorithm for construction of the density matrix in linear scaling electronic structure theory. The expansion contains operations on matrices with irregular sparsity patterns, in particular matrix-matrix multiplication. By using the Chunks and Tasks programming model[Parallel Comput. 40, 328 (2014)] beforehand unknown sparsity pattern is exploited while preventing load imbalance. We evaluate the performance in electronic structure calculations, investigating linear scaling with system size and parallel strong and weak scaling.

Anastasia Kruchinina
Uppsala universitet
anastasia.kruchinina@it.uu.se

Emanuel H. Rubensson
Uppsala University
emanuel.rubensson@it.uu.se

Elias Rudberg
Division of Scientific Computing, Department of I.T.
Uppsala University
elias.rudberg@it.uu.se

CP1

Comparison of Optimization Techniques in Parallel

Ab Initio Nuclear Structure Calculations

Many tasks in ab initio nuclear structure calculations may rely on search techniques, such as tuning nucleon-nucleon and three-nucleon interactions to fit light nuclei. This field is still under-served, however, by the existing optimization software, mainly due to the complexity of the underlying function and derivative evaluations. Previously, we have developed a flexible framework to interface the MFDn package for large-scale ab initio nuclear structure calculations with derivative-free optimization packages. Recently, we have obtained results with the derivative-based optimization techniques. In this talk we compare these two broad optimization approaches as applied to parallel ab initio nuclear structure calculations and note on efficient function evaluations when several nuclei are fitted simultaneously.

Masha Sosonkina
Old Dominion University
Ames Laboratory/DOE
masha@scl.ameslab.gov

CP2

Feast Eigenvalue Solver Using Three Levels of MPI Parallelism

We present an additional level of MPI parallelism for the FEAST Eigenvalue solver that can augment its numerical scalability performances and data processing capabilities for challenging eigenvalue problems at extreme scale. Any MPI solver including domain-decomposition methods, can be interfaced with the software kernel through an RCI mechanism. The algorithm contains three levels of parallelism and is ideally suited to run on parallel architectures. Performance results from a real-world engineering application are discussed.

James Kestyn
University of Massachusetts, Amherst, MA
jkestyn@umass.edu

Eric Polizzi
University of Massachusetts, Amherst, USA
polizzi@ecs.umass.edu

Peter Tang
Intel Corporation
2200 Mission College Blvd, Santa Clara, CA
peter.tang@intel.com

CP2

Parallel Computation of Recursive Bounds for Eigenvalues of Matrix Polynomials

We derive bounds for eigenvalues of polynomial matrices by generalizing to matrix polynomials a recent result for scalar polynomials that can be recursively refined, increasing the computational cost and decreasing sparsity, making it too costly to carry out sequentially for large matrices. Instead, we use a parallel architecture (GPU) and achieve better bounds in the same amount of time necessary to obtain a single unrefined bound on a sequential machine. We illustrate with engineering examples.

Aaron Melman
Santa Clara University
Department of Applied Mathematics
amelman@scu.edu

CP2

Increasing the Arithmetic Intensity of a Sparse Eigensolver for Better Performance on Heterogenous Clusters

The performance of sparse eigensolvers is typically limited by the network and/or main memory bandwidth. In this 'bandwidth bounded' regime it is possible to perform additional 'free flops' to improve the robustness and/or convergence rate of a method. We present block orthogonalization techniques in simulated quadruple precision that exploit specific hardware features. As an example of use, we investigate the performance of a block Jacobi-Davidson method on compute clusters comprising multi/manycore CPUs and GPGPUs.

Jonas Thies
German Aerospace Center (DLR)
jonas.thies@dlr.de

CP2

Highly Scalable Selective Inversion Using Sparse and Dense Matrix Algebra

The computation of a particular subset of entries of inverse matrices is required in several different disciplines, such as genomic prediction, reservoir characterization and finance. Methods based on the matrix factorization can be applied to obtain a significant reduction of the complexity in a parallel and distributed-memory environment. We introduce an algorithm that implements an efficient, parallel, and scalable selective inversion method based on LU and Schur-complement factorization of sparse and dense matrices.

Fabio Verbosio
USI Università della Svizzera Italiana
fabio.verbosio@usi.ch

Drosos Kourounis
USI - Universita della Svizzera italiana
Institute of Computational Science
drosos.kourounis@usi.ch

Olaf Schenk
USI - Università della Svizzera italiana
Institute of Computational Science
olaf.schenk@usi.ch

CP2

Exploring Openmp Task Priorities on the MR3 Eigensolver

As part of the OpenMP 4.1 draft, the runtime incorporates task priorities. We use the Method of Multiple Relatively Robust Representations (MR3), for which a pthreads-based task parallel version already exists (MR3SMP), to analyze and compare the performance of MR3SMP with three different OpenMP runtimes, with and without the support of priorities. From a dataset consisting of application matrices, it appears that OpenMP is always on par or better than the pthreads implementation.

Jan Winkelmann
AICES
RWTH Aachen

winkelmann@aices.rwth-aachen.de

Paolo Bientinesi
AICES, RWTH Aachen
pauldj@aices.rwth-aachen.de

CP3

Fast Mesh Operations Using Spatial Hashing: Remaps and Neighbor Finding

Mesh operations based on hashing techniques are fast and scale linearly with problem-size. Hashing is also comparison-free which is better suited for the fine grained parallelism of Exascale hardware. We present how to leverage hashing for fast remaps and neighbor finding and compare performance with standard tree-based approaches on GPUs, MICs and multicore systems.

Robert Robey
Los Alamos National Lab
XCP-2 Eulerian Applications
brobey@lanl.gov

David Nicholaeff
XCP-2 Eulerian Applications
Los Alamos National Lab
dnic@lanl.gov

Gerald Collom
University of Washington
Los Alamos National Lab
gcollom@gmail.com

Redman Colin
Arizona State University
Los Alamos National Laboratory
13eakers@gmail.com

CP3

Traversing Time-Dependent Networks in Parallel

Large-scale time-dependent networks arise in a range of digital technologies, such as citation networks, schedule networks and networks of social users interacting through messaging. In this talk, we consider ways of mining such complex networks, including computing in parallel the shortest temporal path and the connected components. We generalize breadth first search and depth first search algorithms for time-dependent networks. All algorithms are implemented in Julia, a new dynamic programming language for scientific computing.

Weijian Zhang
School of Mathematics
The University of Manchester
weijian.zhang@manchester.ac.uk

CP4

Scalability Studies of Two-Level Domain Decomposition Algorithms for Stochastic Pdes Having Large Random Dimensions

The scalability of two-level domain decomposition algorithms developed by Subber (PhD thesis, Carleton University, 2012) for uncertainty quantification of large-scale stochastic PDEs is enhanced to handle high-dimensional stochastic problems. Parallel sparse matrix-vector operations are used to cut floating-point operations and memory requirements. Both numerical and parallel scalabilities of the algorithm are presented for a diffusion equation having spatially varying diffusion coefficient modeled by a non-Gaussian stochastic process.

Abhijit Sarkar
Associate Professor
carleton University, Ottawa, Canada
abhijit_sarkar@carleton.ca

Ajit Desai
Carleton University, Canada
ajit.desai@carleton.ca

Mohammad Khalil
Carleton University, Canada
(Currently at Sandia National Laboratories, USA)
mohammad.khalil@carleton.ca

Chris Pettit
United States Naval Academy, USA
pettitcl@usna.edu

Dominique Poirel
Royal Military College of Canada, Canada
poirel-d@rmc.ca

CP4

"Mystic": Highly-Constrained Non-Convex Optimization and Uncertainty Quantification

Highly-constrained, large-dimensional, and nonlinear optimizations are at the root of many difficult problems in uncertainty quantification (UQ), risk, operations research, and other predictive sciences. The prevalence of parallel computing has stimulated a shift from reduced-dimensional models toward more brute-force methods for solving high-dimensional nonlinear optimization problems. The 'mystic' software enables the solving of difficult UQ problems as embarrassingly parallel non-convex optimizations; and with the OUQ algorithm, can provide validation and certification of standard UQ approaches.

Michael McKerns
California Institute of Technology
mmckerns@caltech.edu

CP4

A Resilient Domain Decomposition Approach for Extreme Scale UQ Computations

One challenging aspect of extreme scale computing concerns combining uncertainty quantification (UQ) methods with resilient PDE solvers. We present a resilient polynomial chaos solver for uncertain elliptic PDEs. In a domain decomposition framework, our solver constructs the Dirichlet-to-Dirichlet operator by means of sampling and robust regression techniques. This leads to a resilient UQ solver that requires only a limited amount of global communication, thus benefiting the scalability of the computations.

Paul Mycek, Andres Contreras
Duke University
paul.mycek@duke.edu, andres.contreras@duke.edu

Olivier P. Le Matre
LIMSI-CNRS

olm@limsi.fr

Khachik Sargsyan, Francesco Rizzi, Karla Morris, Cosmin Safta
Sandia National Laboratories
ksargsy@sandia.gov, fnrizzi@sandia.gov, knmorri@sandia.gov, csafta@sandia.gov

Bert J. Debusschere
Energy Transportation Center
Sandia National Laboratories, Livermore CA
bjdebus@sandia.gov

Omar M. Knio
Duke University
omar.knio@duke.edu

CP5

Reliable Bag-of-Tasks Dispatch System

We introduce a reliable bag-of-tasks dispatch system that will allow those bag-of-tasks applications run efficiently on non-reliable hardware at excascale. The system is built on top of MPI to guarantee general applicability across different architectures. It utilizes a mathematical model that, based on the execution time of some of the tasks and their sizes, determines, at runtime, if a task failure is recovered via task re-execution or whether the task requires a replica.

Tariq L. Alturkestani
King Abdullah University of Science and Technology
University of Pittsburgh
tariq.alturkestani@kaust.edu.sa

Esteban Meneses
Costa Rica Institute of Technology
emeneses@ic-itcr.ac.cr

Rami Melhem
University of Pittsburgh, USA
melhem@cs.pitt.edu

CP5

HiPro: An Interactive Programming Tool for Parallel Programming

The scientific computing community is suffering from a lack of good development tool that can handle well the unique problems of coding for high performance computing.HiPro is an interactive graphical tool designed for rapid development of large-scale numerical simulations. The programming tool is based on the domain specific framework JASMIN and is designed to ease parallel programming for domain experts. Real applications demonstrate that it is helpful on enabling increased software productivity.

Li Liao, Aiqing Zhang
Institute of Applied Physics and Computational Mathematics
liliao@iapcm.ac.cn, zhang_aiqing@iapcm.ac.cn

Zeyao Mo
Laboratory of Computational Physics, IAPCM
P.O. Box 8009, Beijing 100088, P.R. China
zeyao_mo@iapcm.ac.cn

CP5

The Elaps Framework: Experimental Linear Algebra Performance Studies

The multi-platform open source framework ELAPS facilitates easy and fast, yet powerful performance experimentation and prototyping of dense linear algebra algorithms. In contrast to most existing performance analysis tools, it targets the initial stages of the development process and assists developers in both algorithmic and optimization decisions. Users construct experiments to investigate how performance and efficiency vary from one algorithm to another, depending on factors such as caching, algorithmic parameters, problem size, and parallelism.

Elmar Peise, Paolo Bientinesi
AICES, RWTH Aachen
peise@aices.rwth-aachen.de, pauldj@aices.rwth-aachen.de

CP6

Sparse Matrix Factorization on GPUs

We present a sparse multifrontal QR Factorization method on a eterogeneous platform multiple GPUs. Our method is extremely efficient for different sparse matrices. Our method benefits from both the highly parallel general-purpose computing cores available on a graphics processing unit (GPU) and from multiple GPUs provided on a single platform. It exploits two types of parallelism: the first over the available GPUs on the platform and the second over the Streaming Multiprocessor (SMPs)within a GPU. We are also developing a technique for factorizing a dense frontal matrix on multiple GPUs using our GPU engine.

Mohamed Gadou
University of Florida
mgadou@cise.ufl.edu

Tim Davis
Texas A&M University
davis@tamu.edu

Sanjay Ranka
CISE, University of Florida, USA
ranka@cise.ufl.edu

CP6

A New Approximate Inverse Technique Based on the GLT Theory

Sparse approximate inverse is a well-known preconditioning technique. Given a matrix A, the approximation M of the A^{-1}, in general, is based on Frobenius norm minimization $|I - MA|$. The sparsity pattern of the inverse matrix is prescribed a priori or adjusted dynamically during its construction. In this presentation, a new method for computing an approximate inverse of a matrix is presented. We utilize the available knowledge of the matrix eigenvalues attained by applying the Generalized Locally Toeplitz (GLT) theory. The sparsity pattern of the resulting approximation can be controlled dynamically. Moreover, the construction of the approximate inverse is parallel and generally cheap since the function describing the matrix A is known a priori.

Ali Dorostkar

Uppsala University
Department of Information Technology
ali.dorostkar@it.uu.se

CP6

Designing An Efficient and Scalable Block Low-Rank Direct Solver for Large Scale Clusters

Electromagnetic problems with an integral equation formulation lead to solve large dense linear systems. To address large problems, we enhanced our Full-MPI direct solver with a block low-rank compression method. We show how we derived a block low-rank version from our existing solver, and detail the optimizations we used to increase its scalability on thousands of cores of the TERA supercomputer at CEA/DAM, for example by introducing hybrid programming models, or with algorithmic improvements.

Xavier Lacoste
CEA - CESTA
xavier.lacoste@cea.fr

Cédric Augonnet, David Goudin
CEA/DAM
cedric.augonnet@cea.fr, david.goudin@cea.fr

CP6

Locality-Aware Parallel Sparse Matrix-Matrix Multiplication

We propose a parallel sparse matrix-matrix multiplication for distributed memory clusters. Our method exploits locality of the nonzero matrix elements without beforehand knowledge of the matrix sparsity pattern. This is achieved by using a distributed hierarchical matrix representation, straightforward to implement using the Chunks and Tasks programming model [Parallel Comput. 40, 328 (2014)]. We demonstrate favorable weak and strong scaling of the communication cost with the number of processes, both theoretically and in numerical experiments.

Emanuel H. Rubensson
Uppsala University
emanuel.rubensson@it.uu.se

Elias Rudberg
Division of Scientific Computing, Department of I.T.
Uppsala University
elias.rudberg@it.uu.se

CP6

Solving Sparse Underdetermined Linear Least Squares Problems on Parallel Computing Platforms

Computing the minimum 2-norm solution is essential for many areas such as geophysics, signal processing and biomedical engineering. In this work, we present a new parallel algorithm for solving sparse underdetermined linear systems where the coefficient matrix is block diagonal with overlapping columns. The proposed parallel approach handles the diagonal blocks independently and the reduced system involving the shared unknowns. Experimental results show the effectiveness of the proposed scheme on various parallel computing platforms.

F. Sukru Torun
Bilkent University
sukruf@cs.bilkent.edu.tr

Murat Manguoglu
Middle East Technical University
manguoglu@ceng.metu.edu.tr

Cevdet Aykanat
Bilkent University
aykanat@cs.bilkent.edu.tr

CP7

Towards Exascalable Algebraic Multigrid

Algebraic multigrid is the solver of choice in many and growing applications in today's petascale environments. Exascale architectures exacerbate the challenges of increased SIMT concurrency, reduced memory capacity and bandwidth per core, and vulnerability to global synchronization. The recent Vassilevski-Yang (2014) mult-additive AMG takes important steps in reducing communication and synchronization frequency while controlling memory growth, but remains bulk-synchronous. We extend the algorithm further towards exascale environments with a task-based directed acyclic graph implementation.

Amani Alonazi, David E. Keyes
KAUST
amani.alonazi@kaust.edu.sa, david.keyes@kaust.edu.sa

CP7

A Matrix-Free Preconditioner for Elliptic Solvers Based on the Fast Multipole Method

We combine the Fast Multipole Method, originally developed as a free-standing solver, with Krylov iteration as a scalable and performant preconditioner for traditional low-order finite discretizations of elliptic boundary value problems. FMM possesses advantages for emergent architectures including: low synchrony, high arithmetic intensity, and exploitable SIMT concurrency. We compare FMM-preconditioned Krylov solvers with multilevel and other solvers on a variety of elliptic problems and illustrate their effectiveness in a computational fluid dynamics application.

Huda Ibeid
King Abdullah University of Science and Technology
huda.ibeid@kaust.edu.sa

Rio Yokota
Tokyo Institute of Technology
rioyokota@gsic.titech.ac.jp

David E. Keyes
KAUST
david.keyes@kaust.edu.sa

CP7

Performance of Smoothers for Algebraic Multigrid Preconditioners for Finite Element Variational Multiscale Incompressible Magnetohydrodynamics

We examine the performance of various smoothers for a fully-coupled algebraic multigrid preconditioned Newton-Krylov solution approach for a finite element variational multiscale turbulence model for incompressible magneto-hydrodynamics (MHD). Our focus will be on large-scale,

high fidelity transient MHD simulations on unstructured meshes. We present scaling results for resistive MHD test cases, including results for over 500,000 cores on an IBM Blue Gene/Q platform.

Paul Lin
Sandia National Laboratories
ptlin@sandia.gov

John Shadid
Sandia National Laboratories
Albuquerque, NM
jnshadi@sandia.gov

Paul Tsuji
Lawrence Livermore National Laboratory
tsuji1@llnl.gov

Jonathan J. Hu
Sandia National Laboratories
Livermore, CA 94551
jhu@sandia.gov

CP7

Efficiency Through Reuse in Algebraic Multigrid

Algebraic multigrid methods (AMG) are utilized by many large-scale scientific applications. Even well developed AMG libraries may have few scalability areas of concern, be it an incomplete factorization smoother or coarse solver, or a graph traversal algorithms required in smoothed aggregation. We will discuss strategies of mitigating the effect of such kernels through reuse of multigrid hierarchy information in multiple setups during the simulation, and demonstrate the effectiveness of such strategies in parallel applications.

Andrey Prokopenko
Sandia National Laboratories
aprokop@sandia.gov

Jonathan J. Hu
Sandia National Laboratories
Livermore, CA 94551
jhu@sandia.gov

CP7

A Dynamical Preconditioning Strategy and Its Applications in Large-Scale Multi-Physics Simulations

In the implicit time-stepping based simulations, a fixed setup-based preconditioner, e.g. AMG, ILU etc, is usually used due to its robustness, which we called the static preconditioning strategy. Setup-based preconditioners, however, lead to poor parallel scalability due to its setup phase. In this work, a dynamical preconditioning strategy is introduced in order to reduce the setup overhead while maintain the robustness. Results for a practical multi-physics simulation on $O(10^4)$ cores show its efficiency and improvement.

Xiaowen Xu
Institute of Applied Physics and Computational Mathematics
xwxu@iapcm.ac.cn

CP8

Parallel Linear Programming Using Zonotopes

Zonotope is an affine projection of n-dimensional cube. The solution of a linear programming problem on a zonotope is given by a closed-form expression. We consider linear programming problem $\min c^T x \; s.t. \; Ax = b, ||x||_\infty \leq 1$. Following work of Fujishige et al., the problem with the help of lifting method can be reduced to finding an intersection between a zonotope and a line. We consider parallelization of this approach using multiple projections of line points on the zonotope via first-order methods.

Max Demenkov
Institute of Control Sciences
max.demenkov@gmail.com

CP8

Vertex Separators with Mixed-Integer Linear Optimization

We demonstrate a method of computing vertex separators in arbitrary graphs using a multi-level coarsening and refinement strategy coupled with a mixed-integer linear programming solver. This approach is easily parallelizable due to the use of branch-and-bound trees, leading to significant speedups on many-core systems. Our method contrasts with other parallelized graph partitioning schemes in that the addition of processors does not decrease the separator quality.

Scott P. Kolodziej, Tim Davis
Texas A&M University
scottk@tamu.edu, davis@tamu.edu

CP8

Large-Scale Parallel Combinatorial Optimization – the Role of Estimation and Ordering the Sub-Problems

One major problem in parallel algorithms for a combinatorial optimization problems is that the search space for the sub-problems differ in several magnitudes. This makes the scheduling hard, and the resulting algorithm often inefficient. Measuring the search space of these sub-problems helps to solve this, and have other usage as well. We will show actual running results using several hundred processors and show that our algorithm can solve previously unsolved maximum clique search problems.

Bogdan Zavalnij, Sandor Szabo
Institute of Mathematics and Informatics
University of Pecs
bogdan@ttk.pte.hu, sszabo7@hotmail.com

CP9

Investigating IBM Power8 System Configuration To Minimize Energy Consumption During the Execution of CoreNeuron Scientific Application

Leveraging from AMESTER library [P. Klavik et al. Changing Computing Paradigms Towards Power Efficiency, Klavk P, Malossi ACI, Bekas C, Curioni A. 2014 Changing computing paradigms towards power efficiency. Phil. Trans. R. Soc. A 372: 20130278.

http://dx.doi.org/10.1098/rsta.2013.0278], the power consumption of the POWER8 system will be investigated when executing CoreNeuron scientific application. Extending the study presented in [T. Ewart et al., Performance Evaluation of the IBM POWER8 Architecture to Support Computational Neuroscientific Application Using Morphologically Detailed Neurons, International Benchmarking and Optimization Workshop, 2015 Supercomputing Conference, Austin, TX, USA] both time to solution and energy consumption of several configurations of the POWER8 system will be analyzed. The optimal design point which allows maximizing application performance while lowering energy consumption will be suggested, offering useful insights for future system designs.

Fabien Delalondre
Ecole Polytechnique Federale de Lausanne
fabien.delalondre@epfl.ch

Timothee Ewart
EPFL
timothee.ewart@epfl.ch

Pramod Shivaj
EPFL - Blue Brain Project
pramod.kumbhar@epfl.ch

Felix Schuermann
EPFL
felix.schuermann@epfl.ch

CP9

Scalable Parallel I/O Using HDF5 for HPC Fluid Flow Simulations

According to the latest US Presidents Council of Advisors on Science and Technology *"high-performance computing must now assume the ability to efficiently manipulate vast and rapidly increasing quantities of data"*. Therefore, we present a framework for scientific I/O using Hierarchical Data Format 5 (HDF5), focusing on highly scalable parallel read and write routines. Furthermore, this framework allows to visualise subsets of the computed data already during runtime of a CFD simulation on massive parallel systems.

Christoph M. Ertl
Chair of Computation in Engineering, Prof. Rank
Technische Universität München
christoph.ertl@tum.de

Ralf-Peter Mundani
TUM, Faculty of Civil, Geo and Environmental Engineering
Chair for Computation in Engineering
mundani@tum.de

Jérôme Frisch
Institute of Energy Efficiency and Sustainable Building E3D
RWTH Aachen University
frisch@e3d.rwth-aachen.de

Ernst Rank
Computation in Engineering
Technische Universität München
ernst.rank@tum.de

CP9

Quantifying Performance of a Resilient Elliptic PDE Solver on Uncertain Architectures Using SST/Macro

Architectural uncertainty and resiliency are two key challenges of exascale computing. In this talk, we present a performance study of a resilient PDE solver on uncertain architectures using the architecture simulator software package SST/macro. We adopt a task-based server-client model to enable resiliency to hard faults. Using this framework, we conduct an uncertainty quantification study allowing us to perform a sensitivity analysis to identify the main parameters affecting performance and their correlations.

Karla Morris, Francesco Rizzi, Khachik Sargsyan
Sandia National Laboratories
knmorri@sandia.gov, fnrizzi@sandia.gov,
ksargsy@sandia.gov

Kathryn Dahlgren
Sandia National Labs
kmdahlg@sandia.gov

Paul Mycek
Duke University
paul.mycek@duke.edu

Cosmin Safta
Sandia National Laboratories
csafta@sandia.gov

Olivier P. Le Matre
LIMSI-CNRS
olm@limsi.fr

Omar M. Knio
Duke University
omar.knio@duke.edu

Bert J. Debusschere
Energy Transportation Center
Sandia National Laboratories, Livermore CA
bjdebus@sandia.gov

CP9

Implicitly Parallel Interfaces to Operational Quantum Computation

Many high-value computational problems involve search spaces that grow exponentially with problem size. Applying traditional parallelism to such problems has had limited success. An alternative strategy exploits *quantum* parallelism occurring within operational quantum computers. Building on recently demonstrated benefits of quantum effects in such devices, we show how subject matter experts using familiar interfaces can seamlessly tap into quantum superposition, entanglement, and tunneling to perform their computations in entirely new ways,without learning quantum mechanics.

Steve P. Reinhardt
Cray, Inc.

sreinhardt@dwavesys.com

CP9

On the Path to Trinity - Experiences Bringing Code to the Next Generation ASC Capability Platform

Our next system, Trinity, a Cray XC40, will be a mixture of Intel Haswell and Intel Xeon Phi Knights Landing processors. We are currently engaged in porting applications to these new processor designs where vectorization, threading and use multi-level memories are required for high performance. We will present initial results from these application ports and comment on the level of work required to get codes running on each partition of the machine.

Courtenay T. Vaughan
Sandia National Laboratories
ctvaugh@sandia.gov

Simon D. Hammond
Scalable Computer Architectures
Sandia National Laboratories
sdhammo@sandia.gov

CP10

Wave Propagation in Highly Heterogeneous Media: Scalability of the Mesh and Random Properties Generator

We present in this talk elastic wave propagation simulations in highly randomly heterogeneous geophysical media using a spectral element solver. When using thousand of cores, meshing and generation of properties become issues. We propose in this talk: a scalable meshing algorithm for hexahedra that considers topography; and a scalable random field generator that constructs elastic parameter fields with chosen statistics. Large scale simulations and scalability analyses illustrate the effectiveness of the approach.

Regis Cottereau
CNRS, CentraleSupelec, Université Paris-Saclay
regis.cottereau@centralesupelec.fr

José Camata
NACAD - High Performance Computing Center
COPPE/Federal University of Rio de Janeiro
camata@nacad.ufrj.br

Lucio de Abreu Corra, Luciano de Carvalho Paludo
CNRS, CentraleSupelec, Université Paris-Saclay
lucio.de-abreu-correa@ecp.fr, luciano.de-carvalho@ecp.fr

Alvaro Coutinho
The Federal University of Rio de Janeiro
Brazil
alvaro@coc.ufrj.br

CP10

An Overlapping Domain Decomposition Preconditioner for the Helmholtz Equation

An overlapping domain decomposition preconditioner is proposed to solve the high frequency Helmholtz equation. The computation domain is decomposed in d directions for problem in R^d, and the local solution on each subdomain is updated simultaneously in one iteration, thus there is no sweeping along certain directions. In these ways, the overlapping DDM is similar to the popular DDM for Poisson problem. The complexity of the overlapping DDM is $O(Nn_{\text{iter}})$, where n_{iter} is the number of iteration of the method, and it is shown numerically that n_{iter} is proportional to the number of subdomains in one direction. 2D Helmholtz problem with more than a billion unknowns and 3D Helmholtz problem are solved efficiently with the preconditioner on massively parallel machines.

Wei Leng
Laboratory of Scientific and Engineering Computing
Chinese Academy of Sciences, Beijing, China
lengweee@gmail.com

Lili Ju
University of South Carolina
Department of Mathematics
ju@math.sc.edu

CP10

GPU Performance Analysis of Discontinuous Galerkin Implementations for Time-Domain Wave Simulations

Finite element schemes based on discontinuous Galerkin methods exhibit interesting features for massively parallel computing and GPGPU. Nowadays, they represent a credible alternative for large-scale industrial applications. Computational performance of such schemes however strongly depends on their implementations, and several strategies have been proposed or discarded in the past. In this talk, we present and compare up-to-date performance results of different implementations tested in a unified framework.

Axel Modave, Jesse Chan, Tim Warburton
Virginia Tech
modave@vt.edu, jlchan@vt.edu, tcew@math.vt.edu

CP10

Efficient Solver for High-Resolution Numerical Weather Prediction Over the Alps

Prototype implementation of new "EULAG" dynamical core, based on nonoscillatory forward in time methods, into the operational COSMO framework for numerical weather prediction is discussed. The core aims at the next-generation very high regional weather prediction over the Alps. Currently, its numerical formulation based on MPDATA advection and iterative Krylov solver offers competitive performance to the operational dynamical core based on Runge-Kutta methods. Moreover, it offers an unprecedented choice of the soundproof/compressible equation systems, offering similar quality of the solution but different ratio of global and local communication cost.

Zbigniew P. Piotrowski
Institute of Meteorology and Water Management
zbigniew.piotrowski@imgw.pl

Damian Wójcik
Poznan Supercomputing and Networking Center
damian.wojcik@imgw.pl

Andrzej Wyszogrodzki
Institute of Meteorology and Water Management, Poland
andrzej.wyszogrodzki@imgw.pl

Milosz Ciznicki
Institute of Meteorology and Water Management

miloszc@man.poznan.pl

CP10

An Adaptive Multi-Resolution Data Model for Feature-Rich Environmental Simulations on Multiple Scales

The accuracy and predictive value of environmental simulations depend largely on the underlying geometric and semantic data. Therefore, a data-model suitable for driving such simulations with special emphasis on multi-scale urban flooding scenarios is presented. This includes fusion and orchestration of heterogeneous and ubiquitous data ranging from GIS to BIM sources, efficient hybrid parallelization of storage, access, and processing of large datasets, and the full integration and interaction with the pipeline for parallel numerical simulations.

Vasco Varduhn
University of Minnesota
varduhn@tum.de

CP11

Maximizing CoreNeuron Application Performance on All CPU Architectures Using Cyme Library

The Cyme library first presented in [T. Ewart et al., A Library Maximizing SIMD Computation on User-Defined Containers. In J.M. Kunkel, T. Ludwig, and H.W. Meuer (Eds.): ISC 2014, LNCS 8488, pp. 440–449. Springer International Publishing Switzerland] supports the autogeneration of optimized code for specific targeted system architectures. Its last extension allows supporting both IBM POWER8 and ARM processors in addition to Intel x86 and MIC as well as IBM Blue Gene/Q. In this presentation we present the integration of the Cyme library into CoreNeuron scientific application which proved to offer close optimal performance on IBM Blue Gene/Q systems except in the case of data hazard. Thanks to the integration, we demonstrate that better performance could ben obtained, overcoming the shortcomings of the XLC compiler.

Timothee Ewart, Kai Langen, Pramod Kumbhar, Felix Schuermann
EPFL
timothee.ewart@epfl.ch, kai.j.langen@gmail.com, pramod.kumbhar@epfl.ch, felix.schuermann@epfl.ch

Fabien Delalondre
Ecole Polytechnique Federale de Lausanne
fabien.delalondre@epfl.ch

CP11

Ttc: A Compiler for Tensor Transpositions

We present TTC, an open-source compiler for multidimensional tensor transpositions. Thanks to a range of optimizations (software prefetching, blocking, loop-reordering, explicit vectorization), TCC generates high-performance parallel C/C++ code. We use generic heuristics and autotuning to manage the huge search space. Performance results show that TTC achieves close to peak memory bandwidth on both the Intel Haswell and the AMD Steamroller architectures, yielding performance gains of up to 15× over modern compilers.

Paul Springer
AICES
RWTH Aachen University
paul.springer@rwth-aachen.de

Paolo Bientinesi
AICES, RWTH Aachen
pauldj@aices.rwth-aachen.de

CP11

Fresh: A Framework for Real-World Structured Grid Stencil Computations on Heterogeneous Platforms

To tackle the heterogeneous programming challenge, for an important computation pattern structured-grid stencil computations, we propose FRESH (Framework for REal-world structured-grid Stencil computations on Heterogeneous platforms), which consists of a Fortran-like domain-specific language and a source-to-source compiler. The DSL contains domain specific features that ease programming. The compiler generates code for multiple platforms. A real-world computational electromagnetic application was ported , and speedups of 1.47, 16.68 and 15.32 were achieved on CPU, GPU and MIC.

Yang Yang, Aiqing Zhang
Institute of Applied Physics and Computational Mathematics
yang_yang@iapcm.ac.cn, zhang_aiqing@iapcm.ac.cn

Zhang Yang
Institute of Applied Physics and Computational Mathematics
China Academy of Engineering Physics
yang_zhang@iapcm.ac.cn

CP12

Load Balancing Issues in Particle In Cell Simulations

Recent simulation campaigns of laser wakefield electrons acceleration, showed that, performance-wise, the most urgent concern is to find a way to face the strong load imbalance that arises on very large full 3D simulations. The hybrid MPI-openMP typical implementation performs quite well on systems with a couple hundreds of cores. But the accessible number of openMP threads is limited and, as the number of MPI processes increases, this relatively small number of threads is not able to balance the load any more. Imbalance issues come back, wasting our precious computation time again. Hence the need for an efficient dynamic load balancing algorithm. The algorithm we present is based on the division of each MPI domain into many smaller, so called, patches. These patches are used as data sorting structures and can be exchanged between MPI processes in order to balance the computational load. It has been implemented in the SMILEI code, a new open source Particle In Cell (PIC) code, developed jointly by physicists and HPC experts in order to make sure that it performs well on the newest supercomputers architectures.

Arnaud Beck
K.U.Leuven
beck@llr.in2p3.fr

Julien Derouillat
CEA - Maison De La Simulation

julien.derouillat@cea.fr

CP12

Particle-in-Cell Simulations for Nonlinear Vlasov-Like Models

We report on the progress in the development of a particle-in-cell code for high-performance computing. We focus on nonlinear Vlasov-like models as applications and on absolute performance per core on the implementation side. Results on classical physics test cases are presented. We report on the code's organization and the approach to multiprocess and multithread issues.

Sever Hirstoaga
Inria
hirstoaga@math.unistra.fr

Edwin Chacon-Golcher
Institute of Physics of the ASCR, ELI-Beamlines,18221 Prague
edwin.chacon-golcher@eli-beams.eu

CP12

Parallel in Time Solvers for DAEs

Time integration of dynamic systems is a sequential process by nature. With the increase of the number of cores and processors in modern computers, making this process parallel becomes an important issue. In this context, Braid software has been developed at LLNL as a non-intrusive parallel-in-time black box overlay. In this talk, we will focus on Differential Algebraic Equations and how parallel-in-time algorithms perform for such problems. These DAEs arise in many fields including the modelling of power grid systems.

Matthieu Lecouvez
Lawrence Livermore National Laboratory
lecouvez1@llnl.gov

Robert D. Falgout
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
rfalgout@llnl.gov

Carol S. Woodward
Lawrence Livermore Nat'l Lab
woodward6@llnl.gov

CP12

Implicitly Dealiased Convolutions for Distributed Memory Architectures

Convolutions are a fundamental tool of scientific computing. For multi-dimensional data, implicitly dealiased convolutions [Bowman and Roberts, SIAM J. Sci. Comput. 2011] are faster and require less memory than conventional FFT-based methods. We present a hybrid `OpenMP/MPI` implementation in the open-source software library `FFTW++`. The reduced memory footprint translates to a reduced communication cost, and the separation of input and work arrays allows one to overlap computation and communication.

Malcolm Roberts
Université de Strasbourg
malcolm.i.w.roberts@gmail.com

John C. Bowman
Dept. of Mathematical and Statistical Sciences
University of Alberta
bowman@ualberta.ca

CP12

A Massively Parallel Simulation of the Coupled Arterial Cells in Branching Network

Arterial wall is composed of a network of multiple types of cells connected via intercellular gap junctions. Under the environmental stimuli, e.g. the blood borne agonists and flow, a network of these cells elicits vasomotion. We present a massively parallel computational model simulating the coupled arterial cells in branching network, to study the mechanistic basis of the coordinated vascular response upstream and downstream of a localized stimulus, mediated by intracellular calcium signaling. The goal is to emulate and validate wet lab experiments.

Mohsin A. Shaikh
Pawsey Supercomputing Center
mohsin.shaikh@pawsey.org.au

Constantine Zakkaroff
UC HPC
University of Canterbury
constantine.zakkaroff@canterbury.ac.nz

Stephen Moore
IBM Research Collaboratory for Life Sciences-Melbourne
The University of Melbourne
stevemoore@au1.ibm.com

Tim David
University of Canterbury
tim.david@cantebrury.ac.nz

CP13

Adaptive Transpose Algorithms for Distributed Multicore Processors

An adaptive parallel block transpose algorithm optimized for distributed multicore hybrid OpenMP/MPI architectures is presented. A hybrid configuration can decrease communication costs by reducing the number of MPI nodes, thereby increasing message sizes. This allows for a more slab-like than pencil-like domain decomposition for multidimensional transforms (such as the FFT), reducing the cost of, or even eliminating the need for, a second distributed transpose. Nonblocking all-to-all transfers enable user computation and communication to be overlapped.

John C. Bowman
Dept. of Mathematical and Statistical Sciences
University of Alberta
bowman@ualberta.ca

Malcolm Roberts
Université de Strasbourg
malcolm.i.w.roberts@gmail.com

CP13

Efficient Two-Level Parallel Algorithm Based on Domain Decomposition Method for Antenna Simulation

This paper presents a scalable method for solving complex

antenna problems. The method is based on FETI-2LM domain decomposition method plus a block ORTHODIR approach for multiple right hand sides. The block approach allows efficient multi-threading parallelization on multi-core nodes for simultaneous local backward substitutions within the subdomains and decreases the overhead due to message passing between subdomains for the orthogonalization of search directions. Performance results for large antenna arrays on systems with up to tens of thousands of cores are analyzed.

Francois-Xavier Roux
ONERA
Francois-Xavier.Roux@onera.fr

CP13

GPU Acceleration of Polynomial Homotopy Continuation

Homotopy continuation methods apply predictor-corrector algorithms to solve polynomial systems. Our implementation on a Graphics Processing Unit (GPU) combines algorithmic differentiation with double double and quad double arithmetic. Our experiments with GPU acceleration show that we can offset the cost overhead of double double precision arithmetic. We report on our implementation on the NVIDIA K20C and the integration into the software PHCpack and phcpy.

Jan Verschelde
Department of Mathematics, Statistics and Computer Science
University of Illinois at Chicago
jan@math.uic.edu

Xiangcheng Yu
Department of Mathematics, Statistics, and Computer Science
University of Illinois at Chicago
xiangchengyu@outlook.com

CP14

Robust Domain Decomposition Methods for Structural Mechanics Computations

Classical domain decomposition strategies (FETI, BDD) are known to slowly converge when applied to real engineering problems. We present more robust block-based domain decomposition methods which can be viewed as multiprecondioned CG algorithms, tailored for medium-sized computer cluster. The talk focuses on numerical implementation of this solver. Some numerical results demonstrating its robustness regarding classical hard points (heterogeneity, incompressibility) will be shown. Slight modifications allowing to deal with large displacement finite element problem will be introduced.

Christophe Bovet
LMT, ENS Cachan, CNRS, Université Paris-Saclay, 94235 Cachan
christophe.bovet@lmt.ens-cachan.fr

Augustin Parret-Fréaud
Safran Tech, Rue des Jeunes Bois, Chteaufort
CS 80112 78772 Magny-Les-Hameaux
augustin.parret-freaud@safran.fr

Pierre Gosselet
LMT, ENS Cachan
gosselet@lmt.ens-cachan.fr

Nicole Spillane
Centre de Mathématiques Appliquées
Ecole Polytechnique
nicole.spillane@cmap.polytechnique.fr

CP14

A Parallel Contact-Impact Simulation Framework

Parallel contact-impact simulation poses great challenges on both software design and performance. We present our framework approach targeting at these challenges. Parallel details such as global contact search, dynamic load balancing among processes and threads, data exchange among sub-domains, etc., are encapsulated into our framework. Application developers are then able to focus only on the physics and numerical schemes within each sub-domain. Algorithms and numerical experiments over 10K cores are presented.

Jie Cheng, Qingkai Liu
Institute of Applied Physics and Computational Mathematics
cheng_jie@iapcm.ac.cn, liuqk@iapcm.ac.cn

CP14

Turbomachinery CFD on Modern Multicore and Manycore Architectures

The exploitation of parallelism across all granularities (instruction, data, core) is mandatory for achieving a high degree of performance on current and future architectures. We present the performance tuning of multiblock and unstructured CFD solvers used for aero-engine design on leading multicore (SandyBridge, Haswell) and manycore (Xeon Phi) platforms. Optimisations include efficient SIMD computations of stencil operations derived from PDE discretizations and domain decomposition for thread parallelism for both structured and unstructured grids.

Ioan C. Hadade, Luca Di Mare
Imperial College London
i.hadade@imperial.ac.uk, l.di.mare@imperial.ac.uk

CP14

Optimization of the Photon's Simulation on GPU

Fast simulation of photon's propagation has applications in medical imaging and dosimetry. Our aim is to implement a CUDA-based GPU method for the simulation of photon's propagation in materials. Using profiling tools, we present how the algorithm has been modified in order to improve thread level parallelism and memory access pattern design and so the execution time. A large-scale test will be implemented to evaluate the total execution time.

Clement P. Rey, Benjamin Auer, Emmanuel Medernach, Ziad El Bitar
CNRS,UMR7178,Université de Strasbourg
clement.rey@iphc.cnrs.fr, benjamin.auer@iphc.cnrs.fr, emmanuel.medernach@iphc.cnrs.fr, ziad.elbitar@iphc.cnrs.fr

MS1

Evaluation of Energy Profile for Krylov Iterative Methods on Cluster of Accelerators

Krylov iterative methods are used to solve large sparse lin-

ear problems, which require optimization work to address the communication bottleneck. With respect to the energy profile, we shall investigate whether optimizations in communication are also energy efficient and scalable. In this presentation, we will reveal the relation between time and energy profile of our optimized kernels for GMRES through an extensive test on GPU clusters, which demands a smart-tuning scheme for optimal performances.

Langshi Chen
CNRS/Maison de la Simulation
chenlangshi@gmail.com

Serge G. Petiton
University Lille 1, Science and Technologies - CNRS
serge.petiton@lifl.fr

MS1

Linear Algebra for Data Science Through Examples

This talk presents some linear algebra problems behind the techniques for analyzing data. Through examples from disciplines as astronomy and health, which generate data masses that must to be properly analyzed, we show the role of certain methods of linear algebra. Then we highlight the importance of high performance computing in the analysis of these data.

Nahid Emad
PRiSM laboratory and Maison de la Simulation,
University of
Nahid.Emad@prism.uvsq.fr

MS1

Coarse Grid Aggregation for Sa-Amg Method with Multiple Kernel Vectors

The Smoothed Agregation Algebraic Multigrid (SA-AMG) method is a multi-level linear solver that can be used in many kinds of applications. Its coarse grid, however, does not have enough number of unknowns for parallelism the computing environment may offer, so special care has to be taken at that stage. In this contribution we discuss the coarse grid aggregation approach. Convergence of the SA-AMG solver can be much improved by specifying multiple near-kernel vectors, and our study investigates the effectiveness of the coarse grid aggregation for SA-AMG solver with multiple kernel vectors.

Akihiro Fujii, Naoya Nomura, Teruo Tanaka
Kogakuin University
fujii@cc.kogakuin.ac.jp, em15016@ns.kogakuin.ac.jp,
teru@cc.kogakuin.ac.jp

Kengo Nakajima
The University of Tokyo
Information Technology Center
nakajima@cc.u-tokyo.ac.jp

Osni A. Marques
Lawrence Berkeley National Laboratory
Berkeley, CA
oamarques@lbl.gov

MS1

Opportunities and Challenges in Developing and Using Scientific Libraries on Emerging Architectures

Abstract not available

Michael Heroux
Sandia National Laboratories
maherou@sandia.gov

MS2

Correlating Facilities, System and Application Data Via Scalable Analytics and Visualization

Understanding application, system and facilities performance and throughput over time necessitates monitoring and recording massive amounts of data. The next step is storing and performing scalable analysis over this data. We will present efforts at LLNL to create an ecosystem for managing historical performance data and using data analytics and visualization to find root causes and insights into performance issues. We will also present some preliminary results on correlating data from various sources.

Abhinav Bhatele, Todd Gamblin
Lawrence Livermore National Laboratory
bhatele@llnl.gov, tgamblin@llnl.gov

Alfredo Gimenez, Yin Yee Ng
University of California, Davis
aagimenez@ucdavis.edu, karenyng@ucdavis.edu

MS2

A Monitoring Tool for HPC Operations Support

TACC Stats is a performance and resource usage monitoring tool that has been running on multiple HPC systems for the past five years. It collects data at regular intervals for every job from hardware counters, sysfs, and procfs. TACC uses this data in a variety of system management and user support efforts. Analyses characterizing common performance issues on TACC systems will be presented using components of the tool that facilitate data exploration and visualization.

R. Todd Evans
University of Texas
rtevans@tacc.utexas.edu

James C. Browne
Dept. of Computer Sciences
University of Texas-Austin
browne@cs.utexas.edu

John West
University of Texas at Austin
Texas Advanced Computing Center
tba

Bill Barth
Texas Advanced Computing Center
University of Texas at Austin
bbarth@tacc.utexas.edu

MS2

Developing a Holistic Understanding of I/O Workloads on Future Architectures

The I/O subsystems of extreme-scale computing systems

are becoming more complex as they stratify into tiers optimized for different balances of performance and capacity. Consequentially, obtaining a coherent picture of the workload that reflects everything from application-level I/O to back-end file system performance will be significantly harder. To this end, we are developing a framework for collecting and normalizing I/O workload data from disparate sources throughout the I/O subsystem to characterize the holistic utilization of future storage platforms.

Glenn Lockwood
NERSC
Lawrence Berkeley National Laboratory
glock@lbl.gov

Nicholas Wright
Lawrence Berkeley National Laboratory
njwright@lbl.gov

MS2

MuMMI: A Modeling Infrastructure for Exploring Power and Execution Time Tradeoffs in Parallel Applications

MuMMI (Multiple Metrics Modeling Infrastructure) facilitates the measurement, modeling and prediction of performance and power consumptions for parallel applications. The MuMMI framework develops models of runtime and power based upon performance counters; these models are used to explore tradeoffs and identify methods for improving energy efficiency. In this talk we will discuss the MuMMI framework and its recent development, and present examples of the use of MuMMI to improve the energy efficiency of parallel applications.

Valerie Taylor, Xingfu Wu
Texas A&M University
taylor@cs.tamu.edu, wuxf@tamu.edu

MS3

Parallel Graph Matching Algorithms Using Matrix Algebra

We present distributed-memory parallel algorithms for computing matchings in bipartite graphs. We consider both exact and approximate algorithms for cardinality and weighted matching problems. We substitute the asynchronous data access patterns of traditional matching algorithms by a small subset of more structured, bulk-synchronous functions based on matrix algebra. Relying on communication-avoiding algorithms for the underlying matrix-algebric modules, different matching algorithms achieve good speedups on tens of thousands of cores on current supercomputers.

Ariful Azad, Aydin Buluc
Lawrence Berkeley National Laboratory
azad@lbl.gov, abuluc@lbl.gov

MS3

Randomized Linear Algebra in Large-Scale Computational Environments

Randomized linear algebra (RLA) is an area that exploits randomness for the development of improved algorithms for fundamental matrix problems. It has led to the best worst-cast algorithms for problems such as over-determined least squares approximation, and high-quality implementations have beaten Lapack as well as specialized code in scientific applications. Here, we review recent RLA results in implementing in parallel and distributed environments, e.g., in high performance computing applications and in distributed data centers.

Michael Mahoney, Alex Gittens
UC Berkeley
mmahoney@stat.berkeley.edu, gittens@icsi.berkeley.edu

Farbod Roosta-Khorasani
Computer Science
UBC
farbod@stat.berkeley.edu

MS3

Parallel Combinatorial Algorithms in Sparse Matrix Computation?

Combinatorial techniques are used in several phases of sparse matrix computation. For large-scale problems, while numerical phases are often executed in parallel, most of these combinatorial techniques are serial and can become bottlenecks. We are investigating the extent to which some of the combinatorial techniques can be performed in parallel.

Mathias Jacquelin
Lawrence Berkeley National Lab
mjacquelin@lbl.gov

Esmond G. Ng
Lawrence Berkeley National Laboratory
egng@lbl.gov

Barry Peyton
Dalton State College
bpeyton@daltonstate.edu

Yili Zhrng, Kathy Yelick
Lawrence Berkeley National Laboratory
yzheng@lbl.gov, kayelick@lbl.gov

MS3

Parallel Algorithms for Automatic Differentiation

Despite three decades of effort, parallel algorithms for Automatic Differentiation have been challenging to design. We describe a new approach to parallelization for second and higher order derivatives using the Reverse mode of AD implemented with the MPI library. By interpreting a send message as creating a dependent variable, and a receive message as creating an independent variable, we perform all communications in the Forward mode, thus avoiding the reversal of communication operations.

Alex Pothen
Purdue University
Department of Computer Science
apothen@purdue.edu

Mu Wang
Purdue University
wangmu0701@gmail.com

MS4

Enlarged GMRES for Reducing Communication

When Solving Reservoir Simulation Problems

In this work we discuss enlarged Krylov subspaces, [Grigori et al., Enlarged Krylov subspace conjugate gradient methods for reducing communication], to solve linear systems with one right hand side, as arising from reservoir simulations. Our work focuses on GMRES since the considered linear systems are neither positive definite nor symmetric. We discuss the usage of deflation, inexact breakdowns, [Robbé and Sadkane, Exact and inexact breakdowns in the block GMRES method], in each iteration and deflation of approximated eigenvalues after each cycle when restarting GMRES.

Hussam Al Daas
Université Paris-Sud 11
hussam.al-daas@inria.fr

Laura Grigori
INRIA
France
Laura.Grigori@inria.fr

Pascal Henon
TOTAL E&P
pascal.henon@total.com

Philippe Ricoux
TOTAL SA
philippe.ricoux@total.com

MS4

Communication-Avoiding Krylov Subspace Methods in Theory and Practice

Solvers for sparse linear algebra problems, ubiquitous throughout scientific codes, often limit application performance due to a low computation/communication ratio. In this talk, we discuss the development of parallel communication-avoiding Krylov subspace methods, which can achieve performance improvements in a wide variety of scientific applications. We present results of a parallel scaling study and discuss complex tradeoffs that arise due to both machine parameters and the numerical properties of the problem. We conclude with recommendations on when we expect communication-avoiding variants to be beneficial in practice.

Erin C. Carson
New York University
erin.carson@nyu.edu

MS4

Outline of a New Roadmap to Permissive Communication and Applications That Can Benefit

FFT has been a classic computation engine for numerous applications but is bandwidth-intensive, limiting its performance on off-the-shelf parallel machines. Using the FFT and other applications as examples, we examine the impact that adoption of some enabling technologies would have on the performance of a manycore architecture. These technologies include 3D VLSI and silicon photonics, with a special midwifing role for microfluidic cooling, in a roadmap towards permissive communication.

James A. Edwards, Uzi Vishkin
University of Maryland
jedward5@umd.edu, vishkin@umd.edu

MS4

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

We consider the design and implementation of communication-efficient parallel support vector machines, widely used in statistical machine learning, for distributed memory clusters and supercomputers. Prior to our study, the parallel isoefficiency of a state-of-the-art implementation scaled as $W = \Omega(P^3)$, where W is the problem size and P the number of processors. Our Communication-Avoiding SVM method improves the isoefficiency to nearly $W = \Omega(P)$. We evaluate these methods on 96 to 1536 processors, and show speedups of up to 16× over Dis-SMO, and a 95% weak-scaling efficiency on six real-world datasets with only modest losses in overall classification accuracy.

Yang You
University of California, Berkeley
youyang@eecs.berkeley.edu

MS5

Time Parallel and Space Time Solution Methods - From Heat Equation to Fluid Flow

In this talk, we discuss the parallel solution in space and time of the Navier Stokes equation using space-time multigrid methods. The choice of the smoother is discussed as well as the proper interwtwining of the Picard iteration, which is used for resolving the nonlinear terms in the Navier Stokes equations, with the linear multigrid method used.

Rolf Krause
Institute of Computational Science
University of Lugano
rolf.krause@usi.ch

Pietro Benedusi
Institute of Computational Science
USI Lugano
pietro,benedusi@usi.ch

Peter Arbenz
ETH Zurich
Computer Science Department
arbenz@inf.ethz.ch

Daniel Hupp
ETH Zürich
huppd@inf.ethz.ch

MS5

A Parallel-in-Time Solver for Time-Periodic Navier-Stokes Problems

In science and engineering, many problems are driven by time-periodic forcing. In fluid dynamics, this occurs for example in turbines or in human blood flows. In the absence of turbulence, if a fluid is excited long enough, a time-periodic steady state establishes. In numerical models of such flows, the steady state is computed by simulating this transient phase with a time-stepping method until the periodic state is reached. In this work, we present an alternative approach by solving directly for the steady-state

solution, introducing periodic boundary conditions in time. This allows us to use a multi-harmonic ansatz in time for the solution of the Navier-Stokes equations. The solution of the different Fourier modes is distributed to different groups of processing units, which then have to solve a spatial problem. Therefore, the time axis is parallelized such that the number of parallel threads scales with the number of resolved Fourier modes.

Daniel Hupp
ETH Zürich
huppd@inf.ethz.ch

Peter Arbenz
ETH Zurich
Computer Science Department
arbenz@inf.ethz.ch

Dominik Obrist
University of Bern
dominik.obrist@artorg.unibe.ch

MS5

Time Parallelizing Edge Plasma Simulations Using the Parareal Algorithm

This work explores the exploitation of the parareal algorithm to achieve temporal parallelization of simulations of magnetically confined, fusion plasma at the edge of the device. The physics is extremely complex due to the presence of neutrals as well as the interaction of the plasma with the wall. The SOLPS code package, widely adapted by the fusion community, is used here to scale the performance of the parareal algorithm in various ITER relevant plasmas.

Debasmita Samaddar
CCFE
UK Atomic Energy Authority
dsamaddar@alaska.edu

David Coster
Max-Planck-Institut für Plasmaphysik, Germany
david.coster@ipp.mpg.de

Xavier Bonnin
ITER Organization, France
xavier.bonnin@iter.org

Eva Havlickova
UK Atomic Energy Authority
evikuv@gmail.com

Wael R. Elwasif
ORNL, USA
elwasifwr@ornl.gov

Donald B. Batchelor
Oak Ridge National Laboratory
batchelordb@ornl.gov

Lee A. Berry
ORNL, USA
leeaberry223@gmail.com

Christoph Bergmeister
University of Innsbruck, Austria
christoph.bergmeister@student.uibk.ac.at

MS5

Micro-Macro Parareal Method for Stochastic Slow-Fast Systems with Applications in Molecular Dynamics

We introduce a micro-macro parareal algorithm for the time-parallel integration of multiscale-in-time systems. The algorithm first computes a cheap, but inaccurate, solution using a coarse propagator (simulating an approximate slow macroscopic model), which is iteratively corrected using a fine-scale propagator (accurately simulating the full microscopic dynamics). This correction is done in parallel over many subintervals, thereby reducing the wall-clock time needed to obtain the solution, compared to the integration of the full microscopic model. We provide a numerical analysis of the algorithm for a prototypical example of a micro-macro model, namely singularly perturbed ordinary differential equations. We show that the computed solution converges to the full microscopic solution (when the parareal iterations proceed) only if special care is taken during the coupling of the microscopic and macroscopic levels of description. The convergence rate depends on the modeling error of the approximate macroscopic model. We illustrate these results with numerical experiments, including a non-trivial model inspired by molecular dynamics.

Giovanni Samaey
Department of Computer Science, K. U. Leuven
giovanni.samaey@cs.kuleuven.be

Keith Myerscough
Department of Computer Science
KU Leuven
keith.myerscough@cs.kuleuven.be

Frederic Legoll
LAMI, ENPC
legoll@lami.enpc.fr

Tony Lelievre
Ecole des Ponts ParisTech
lelievre@cermics.enpc.fr

MS6

Parallel GMRES Using A Variable S-Step Krylov Basis

Sparse linear systems arise in computational science and engineering. The goal is to reduce the memory requirements and the computational cost, by means of high performance computing algorithms. We introduce a new variation on s-step GMRES in order to improve its stability, reduce the number of iterations necessary to ensure convergence, and thereby improve parallel performance. In doing so, we develop a new block variant that allows us to express the stability difficulties in s-step GMRES more fully.

David Imberti
Inria-Rennes
dmimb@hushmail.com

Jocelyne Erhel
INRIA-Rennes, France
Campus de Beaulieu

Jocelyne.Erhel@inria.fr

MS6

A Tiny Step Toward a Unified Task Scheduler for Large Scale Heterogeneous Architectures

Abstract not available

Clement Fontenaille
UVSQ
clement.fontenaille@prism.uvsq.fr

MS6

Extraction the Compute and Communications Patterns from An Industrial Application

Abstract not available

Thomas Guillet
Intel
thomas.guillet@intel.com

MS6

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

The modeling of turbulent combustion in realistic devices is a multi-scale problem, which can be tackled with future exascale super-computers. However, the multi-scale nature of the flow often requires h- or p-adaptivity and the use of splitting approaches to decouple the different time-scales. These methods introduce a large load unbalance that needs to be addressed when thousands of cores are used.

Vincent R. Moureau
CORIA Laboratory
Université et INSA de Rouen
vincent.moureau@coria.fr

Pierre Bénard
CORIA - CNRS UMR6614, Université et INSA de Rouen
pierre.benard@coria.fr

Ghislain Lartigue
CORIA - CNRS UMR6614
ghislain.lartigue@coria.fr

MS7

Programming Linear Algebra with Python: A Task-Based Approach

Task-based programming has been extensively used to program linear algebra applications. While most of times the chosen programming language falls in the category of traditional ones, there is a trend in the HPC community in developing applications in scripting languages. COMPSs programming model has been recently extended to support a syntax in Python (PyCOMPSs). The talk will present the features of PyCOMPSs and how linear algebra Python codes can be run in parallel in clusters.

Rosa M. Badia
Barcelona Supercomputing Center / CSIC
rosa.m.badia@bsc.es

MS7

Optimizing Numerical Simulations of Elastodynamic Wave Propagation Thanks to Task-Based Parallel Programming

Elastodynamics equations can be discretized into an intrinsic parallelizable algorithm. The programming paradigm choice is critical as its capabilities to take advantage of the exposed degree of parallelism will vary. MPI may probe to be too static, especially concerning load balancing issues and ease of granularity adaptation. Instead of MPI, we used a task-based programming over runtime systems, PaRSEC. We analyze this novel implementation and compare the performances on both shared and distributed memory machines.

Lionel Boillot
Team-Project Magique-3D
INRIA Bordeaux Sud-Ouest
lionel.boillot@inria.fr

Corentin Rossignon
TOTAL
corentin.rossignon@total.com

George Bosilca
University of Tennessee - Knoxville
bosilca@icl.utk.edu

Emmanuel Agullo
INRIA
emmanuel.agullo@inria.fr

Henri Calandra
Total
CSTJF, Geophysical Operations & Technology, R&D Team
henri.calandra@total.com

MS7

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

Abstract not available

Aurelien Bouteiller
University of Tennessee
bouteill@eecs.utk.edu

George Bosilca
University of Tennessee - Knoxville
bosilca@icl.utk.edu

MS7

Controlling the Memory Subscription of Applications with a Task-Based Runtime System

This talk will discuss the cooperation between a task-based *compressed* linear algebra code and a runtime to control the memory subscription levels throughout the execution. The task paradigm enables controlling the memory footprint of the application by throttling the task submission rate, striking a compromise between the performance benefits of anticipative task submission and the resulting memory consumption. We illustrate our contribution on an ACA-

based *compressed* dense linear algebra solver.

Marc Sergent
INRIA, France
marc.sergent@inria.fr

David Goudin
CEA/CESTA
david.goudin@gmail.com

Samuel Thibault
University of Bordeaux 1, Inria
samuel.thibault@labri.fr

Olivier Aumage
INRIA, France
olivier.aumage@inria.fr

MS8

Uncertainty Quantification Using Tensor Methods and Large Scale Hpc

We consider the problem of uncertainty quantification for extreme scale parameter dependent problems where an underlying low rank property of the parameter dependency is assumed. For this type of dependency the hierarchical Tucker format offers a suitable framework to approximate a given output function of the solutions of the parameter dependent problem from a number of samples that is linear in the number of parameters. In particular we can a posteriori compute the mean, variance or other interesting statistical quantities of interest. In the extreme scale setting it is already assumed that the underlying fixed-parameter problem is distributed and solved for in parallel. We provide in addition a parallel evaluation scheme for the sampling phase that allows us on the one hand to combine several solves and on the other hand parallelise the sampling. Finally, we address the problem of estimating the accuracy of the computations in low rank tensor format.

Lars Grasedyck, Christian Löbbert
RWTH Aachen
lgr@igpm.rwth-aachen.de, loebbert@igpm.rwth-aachen.de

MS8

Parallel Methods in Space and Time and Their Application in Computational Medicine

In this talk, we consider parallel discretization methods in space and time for the efficient solution of problems arising from medicine. We start from permeation through the human skin, were we combine parallelization in time through the parareal method with multigrid methods in space. Then, we discuss the full space time discretization of reaction diffusion problems as occuring in electrophysiological simulations of cardiac activity.

Rolf Krause
Institute of Computational Science
University of Lugano
rolf.krause@usi.ch

Andreas Kreienbühl
Institute of Computational Science
USI Lugano
andreas.kreienbuehl@usi.ch

Dorian Krause
Institute of Computational Science
University of Lugano
dorian.krause@usi.ch

Gabriel Wittum
G-CSC
University of Frankfurt, Germany
wittum@techsim.org

MS8

Scalable Shape Optimization Methods for Structured Inverse Modeling Using Large Scale Hpc

In this talk, we present a method combining HPC techniques and large scale applications with an optimized inverse shape identification procedure. It is demonstrated, how existing multigrid infrastructure can be incorporated into a shape optimization framework with a focus on scalability of the algorithm. These techniques are utilized in order to fit a model of the human skin to data measurements and not only estimate the permeability coefficients but also the shape of the cells.

Volker H. Schulz
University of Trier
Department of Mathematics
Volker.Schulz@uni-trier.de

Martin Siebenborn
University of Trier
siebenborn@uni-trier.de

MS8

Parallel Adaptive and Robust Multigrid

Abstract not available

Gabriel Wittum
G-CSC
University of Frankfurt, Germany
wittum@techsim.org

MS9

Iterative ILU Preconditioning

Iterative and asynchronous algorithms, with their ability to tolerate memory latency, form an important class of algorithms for modern computer architectures. In this talk we present computational and numerical aspects of replacing traditional ILU preconditioners with iterative and potentially asynchronous variants.

Hartwig Anzt
Innovate Computing Lab
University of Tennessee
hanzt@icl.utk.edu

Edmond Chow
School of Computational Science and Engineering
Georgia Institute of Technology
echow@cc.gatech.edu

MS9

Kokkos: Manycore Programmability and Performance Portability

Kokkos enables development of HPC applications that are

performance portable across disparate manycore devices such as NVIDIA Kepler, Intel Xeon Phi, and AMD Fusion. Kokkos abstractions leverage C++11 features to enable developer productivity in implementing thread parallel algorithms; especially hierarchical parallel algorithms and architecture-polymorphic data structures required for high degree of parallelism and performant memory access patterns. In contrast, language extensions such as OpenMP, OpenACC, or OpenCL fail to address memory access patterns and are unwieldy to implement non-trivial parallel algorithms. Kokkos abstractions, highlights of its standard C++11 library interface, and sampling of performance results will be presented.

H. Carter Edwards, Christian Trott
Sandia National Laboratories
hcedwar@sandia.gov, crtrott@sandia.gov

MS9

ICB: IC Preconditioning with a Fill-in Strategy Considering Simd Instructions

Most of current processors are equipped with SIMD instructions that are used to increase the performance of application programs. We analyze the effective use of SIMD instructions in the Incomplete Cholesky Conjugate Gradient (ICCG) solver. A new fill-in strategy in the IC factorization is proposed for the SIMD vectorization of the preconditioning step and to increase the convergence rate. Our numerical results confirm that the proposed method has better performance than the conventional IC(0)-CG method.

Takeshi Iwashita
Hokkaido University
iwashita@iic.hokudai.ac.jp

Naokazu Takemura
Kyoto University
na.takemura@sys.i.kyoto-u.ac.jp

Akihiro Ida
Academic Center for Computing and Media Studies
Kyoto University
ida@media.kyoto-u.ac.jp

Hiroshi Nakashima
ACCMS
Kyoto University
h.nakashima@media.kyoto-u.ac.jp

MS9

Thread-Scalable FEM

Finite Element Method (FEM) is widely used for solving Partial Differential Equations (PDE) in various types of applications of computational science and engineering. Matrix assembly and sparse matrix solver are the most expensive processes in FEM procedures. In the present work, these processes are optimized on Intel Xeon Phi and NVIDIA Tesla K40 based on features of each architecture. This talk describes details of optimization and results of performance evaluation.

Kengo Nakajima
The University of Tokyo
Information Technology Center
nakajima@cc.u-tokyo.ac.jp

MS10

Data Management and Analysis for An Energy Efficient HPC Center

The HPC center at LLNL is concerned about the quality, cost, and environmental impact; in addition, we share with our providers the interest to reduce energy costs and improve electrical grid reliability. We implemented an extensive monitoring and data collection system to collect and analyze data from power meters, PMU's, and computer environmental and power sensors. We will describe our work on managing (storage and integration), analysis and visualization of the collected data.

Ghaleb Abdulla, Anna Maria Bailey, John Weaver
Lawrence Livermore National Laboratory
abdulla1@llnl.gov, bailey31@llnl.gov, weaver11@llnl.gov

MS10

Smart HPC Centers: Data, Analysis, Feedback, and Response

More efficient HPC operation can be enabled by analysis of full system data, feedback of analysis results to applications and subsystems, and dynamic adaptive response. However, significant gaps for support of such operations still exist. We present progress on an integrated full system approach to HPC operations addressing our work on a) application and global platform analysis b) architecture for data integration, analysis, and feedback and c) system (e.g., resource manager, facilities, partitioners) responses.

James Brandt, Ann Gentile
Sandia National Laboratories
brandt@sandia.gov, gentile@sandia.gov

Cindy Martin
Los Alamos National Laboratory
c_martin@lanl.gov

Benjamin Allan
Sandia National Laboratories
Livermore, CA 94550
baallan@sandia.gov

Karen D. Devine
Sandia National Laboratories
kddevin@sandia.gov

MS10

Low-Overhead Monitoring of Parallel Applications

Next-generation parallel systems will have hundreds of hardware threads per node and in some cases attached accelerators. Growth in node-level parallelism is a challenge for monitoring. This talk will lay out a vision for integrated, low-overhead, sampling based measurement of computation, communication, and I/O of hybrid programs on next-generation supercomputers. Using sampling, one can pinpoint inefficiencies and scaling bottlenecks at all levels within and across nodes of parallel systems.

John Mellor-Crummey
Computer Science Dept.
Rice University.
johnmc@rice.edu

Mark Krentel
Rice University
krentel@rice.edu

MS10

Discovering Opportunities for Data Analytics Through Visualization of Large System Monitoring Datasets

Blue Waters is a Cray supercomputer comprised of over 27,000 nodes and has been instrumented with detailed monitoring across all of its subsystems. NCSA has been focused on integrating the data and producing insights into application and systems behavior. While realtime event processing and visualization of the data has proved successful in demonstrating some clear data patterns, we look forward to the opportunities in numerical analysis and data analytics on this long term dataset.

Michael Showerman
NCSA
University of Illinois at Urbana-Champaign
mung@illinois.edu

MS11

A Partitioning Problem for Load Balancing and Reducing Communication from the Field of Quantum Chemistry

We present a combinatorial problem and potential solutions arising in parallel computational chemistry. The Hartree-Fock (HF) method has a very complex data access pattern. Much research has been devoted over 20 years for parallelizing this important method, based primarily on intuition and experience. A formal approach for parallelizing HF while reducing communication may come from graph and hypergraph partitioning. Besides providing a potential solution, this approach may also shed light on the optimality of existing approaches.

Edmond Chow
School of Computational Science and Engineering
Georgia Institute of Technology
echow@cc.gatech.edu

MS11

Scalable Parallel Algorithms for De Novo Assembly of Complex Genomes

A critical problem for computational genomics is the problem of *de novo* genome assembly: the development of robust scalable methods for transforming short randomly sampled sequences into the contiguous and accurate reconstruction of complex genomes. While advanced methods exist for assembling the small and haploid genomes of prokaryotes, the genomes of eukaryotes are more complex. We address this challenge head on by developing HipMer, an end-to-end high performance de novo assembler designed to scale to massive concurrencies. HipMer employs an efficient Unified Parallel C (UPC) implementation and computes the assembly of the human genome in only 8.4 minutes using 15,360 cores of a Cray XC30 system.

Evangelos Georganas
EECS
UC Berkeley
egeor@eecs.berkeley.edu

MS11

Community Detection on GPU

There has been considerable interest in community detection for finding the modularity structure in real world data. Such data sets can arise from social networks as well as various scientific domains. The Louvain method is one popular method for this problem as it is simple and fast. It can also be used to detect hierarchical structures in the data. However, its inherently sequential nature and cache unfriendly workloads makes it difficult to parallelize. This is particularly true for co-processor architectures. In this work we show how these obstacles can be overcome and present results from implementing the algorithm on a GPU.

Md Naim
University of Bergen
Norway
md.naim@uib.no

Fredrik Manne
University of Bergen, Norway
Fredrik.Manne@ii.uib.no

MS12

Lower Bound Techniques for Communication in the Memory Hierarchy

Communication requirements between different levels of the memory hierarchy of a DAG computation are strictly connected to the space requirements of the subset in suitable partitions of the computation DAG. We discuss a versatile technique based on a generalized notion of visit of a DAG to obtain space lower bounds. We illustrate the technique by deriving space lower bounds for specific DAGs. We also establish some limitations to the power of the proposed technique.

Gianfranco Bilardi
University of Padova
bilardi@dei.unipd.it

MS12

Communication-Optimal Loop Nests

Communication (data movement) often dominates a computation's runtime and energy costs, motivating organizing an algorithm's operations to minimize communication. We study communication costs of a class of algorithms including many-body and matrix/tensor computations and, more generally, loop nests operating on array variables subscripted by linear functions of the loop iteration vector. We use this algebraic relationship between variables and operations to derive communication lower bounds for these algorithms. We also discuss communication-optimal implementations that attain these bounds.

Nicholas Knight
New York University
nknight@nyu.edu

MS12

Write-Avoiding Algorithms

Communication complexity of algorithms on parallel machines and memory hierarchies is well understand. Most of

this work analyzes the amount of communication without separately counting induced reads and writes. However, in upcoming memory technologies, such as non-volatile memories, there is a significant asymmetry in the cost of reads and writes. This motivates us to study lower bounds on the number of writes, which are relatively expensive. While we prove that for many algorithms there can not be asymptotically fewer writes than reads, in the case of direct linear algebra and N-body methods the opposite is true. We present algorithms that match these lower bounds.

Harsha Vardhan Simhadri
Lawrence Berkeley National Laboratory
harshas@lbl.gov

MS12

Communication-Efficient Evaluation of Matrix Polynomials

Abstract not available

Sivan A. Toledo
Tel Aviv University
stoledo@tau.ac.il

MS13

Parallel-in-Time Coupling of Models and Numerical Methods

We present a new numerical procedure for the simulation of time-dependent problems based on the coupling between the finite element method and the lattice Boltzmann method. The procedure is based on the Parareal paradigm and allows to couple two numerical methods having optimal efficiency at different space and time scales. The main motivation, among other, is that one technique may be more efficient, or physically more appropriate or less memory consuming than the other depending on the target of the simulation and/or on the sub-region of the computational domain. We detail the theoretical and numerical framework for linear parabolic equations even though its potential applicability should be wider. We show various numerical examples on the heat equation that will validate the proposed procedure and illustrate its advantages. Finally some recent extensions and/or variants of our original algorithm will be discussed.

Franz Chouly
Besancon
franz.chouly@univ-fcomte.fr

Matteo Astorino
Ecole Pol. Fed. de Lausanne
matteo.astorino@gmail.com

Alexei Lozinski
Laboratoire de Mathématiques de Besançon
CNRS et Université de Franche-Comté
alexei.lozinski@univ-fcomte.fr

Alfio Quarteroni
Ecole Pol. Fed. de Lausanne
Alfio.Quarteroni@epfl.ch

MS13

Fault-Tolerance of the Parallel-in-Time Integration Method PFASST

For time-dependent PDEs, parallel-in-time integration with e.g. the "parallel full approximation scheme in space and time" (PFASST) is a promising way to accelerate existing space-parallel approaches beyond their scaling limits. In addition, the iterative, multi-level nature of PFASST can be exploited to recover from hardware failures, preventing a breakdown of the time integration process. In this talk, we present various recovery strategies, show their impact and discuss challenges for the implementation on HPC architectures.

Robert Speck, Torbjörn Klatt
Juelich Supercomputing Centre
Forschungszentrum Juelich GmbH
r.speck@fz-juelich.de, t.klatt@fz-juelich.de

Daniel Ruprecht
School of Mechanical Engineering
University of Leeds
d.ruprecht@leeds.ac.uk

MS13

Parareal in Time and Fixed-Point Schemes

In this talk, we will discuss how the parallel efficiency of the paralleal in time algorithm can be improved by coupling it with fixed-point schemes.

Olga Mula
CEREMADE (Paris Dauphine University)
mula@ceremade.dauphine.fr

MS13

Implementating Parareal - OpenMP Or MPI?

I will introduce an OpenMP implementation of Parareal with pipelining and compare it to a standard MPI implementation with respect to runtime, memory and energy. The used benchmark code is PararealF90, which provides three different implementations of Parareal with and without pipelining using OpenMP or MPI to parallelise in time. Benchmarks from a Linux cluster and a Cray XC40 show that while both versions differ only very little in runtime, using OpenMP can save memory.

Daniel Ruprecht
School of Mechanical Engineering
University of Leeds
d.ruprecht@leeds.ac.uk

MS14

Combining Software Pipelining with Numerical Pipelining in the Conjugate Gradient Algorithm

Modern hardware architectures provide an extremely high level of concurrency. One the one hand, new programming models such as task-based methods are being investigated by the HPC community in order to pipeline tasks at a fine-grain level. On the other hand, new numerical methods are being studied in order to relieve numerical synchronization points of most inner-most kernels used in simulations. In this talk, we propose a formulation of the Conjugate Gradient algorithm that combines software pipelining through task-based programming with numerical pipelining through the use of a recently introduce latency-hiding algorithm. We study the potential of

the method on modern multi-GPU, multicore and heterogeneous arcthitectures.

Emmanuel Agullo
INRIA
emmanuel.agullo@inria.fr

Luc Giraud
Inria
Joint Inria-CERFACS lab on HPC
luc.giraud@inria.fr

Stojce Nakov
INRIA
stojce.nakov@inria.fr

MS14

Rounding Errors in Pipelined Krylov Solvers

Abstract not available

Siegfried Cools
University of Antwerpen
Belgium
siegfried.cools@uantwerpen.be

MS14

Espreso: ExaScale PaRallel Feti Solver

This talk will present a new approach for efficient implementation of the FETI and Hybrid FETI solver for manycore accelerators using Schur complement (SC). By using the SC solver can avoid working with sparse Cholesky factors of the stiffness matrices. Instead a dense structure is computed and used by the iterative solver. The main bottleneck of this method is the computation of the SC. This bottleneck has been significantly reduced by using new PARDISO-SC sparse direct solver.

Lubomir Riha
VSB Technical University of Ostrava
lubomir.riha@vsb.cz

Tomas Brzobohaty, Alexandros Markopoulos
IT4Innovations National Supercomputing Centre
VSB Technical University of Ostrava
tomas.brzobohaty@vsb.cz, alexandros.markopoulos@vsb.cz

MS14

Notification Driven Collectives in Gaspi

Abstract not available

Christian Simmendinger
T-Systems
christian.simmendinger@t-systems-sfr.com

MS15

Exploiting Two-Level Parallelism by Aggregating Computing Resources in Task-Based Applications Over Accelerator-Based Machines

Computing platforms are now extremely complex providing an increasing number of CPUs and accelerators. This trend makes balancing computations between these heterogeneous resources performance critical. In this paper we tackle the task granularity problem and we propose aggregating several CPUs in order to execute larger parallel tasks and thus find a better equilibrium between the workload assigned to the CPUs and the one assigned to the GPUs. To this end, we rely on the notion of scheduling contexts in order to isolate the parallel tasks and thus delegate the management of the task parallelism to the inner scheduling strategy. We demonstrate the relevance of our approach through the dense Cholesky factorization kernel implemented on top of the StarPU task-based runtime system. We allow having parallel elementary tasks and using Intel MKL parallel implementation optimized through the use of the OpenMP runtime system. We show how our approach handles the interaction between the StarPU and the OpenMP runtime systems and how it exploits the parallelism of modern accelerator-based machines. We present experimental results showing that our solution outperforms state of the art implementations to reach a peak performance of 4.5 TFlop/s on a platform equipped with 20 CPU cores and 4 GPU devices.

Terry Cojean
INRIA
terry.cojean@inria.fr

MS15

Task-Based Multifrontal QR Solver for GPU-Accelerated Multicore Architectures

In this talk we present the design of task-based sparse direct solvers on top of runtime systems. In the context of the qr_mumps solver, we prove the usability and effectiveness of our approach with the implementation of a sparse matrix multifrontal factorization based on a Sequential Task flow (STF) parallel programming model. Using this programming model, we developed features such as the integration of dense 2D Communication Avoiding algorithms in the multifrontal method allowing for better scalability compared to the original approach used in qr_mumps. Furthermore, following this approach, we show that we are capable of exploiting heterogeneous architectures with the design of data partitioning and a scheduling algorithms capable of handling the heterogeneity of resources.

Florent Lopez
Université Paul Sabatier - IRIT (ENSEEIHT)
florent.lopez@enseeiht.fr

MS15

An Effective Methodology for Reproducible Research on Dynamic Task-Based Runtime Systems

Reproducibility of experiments and analysis by others is one of the pillars of science. Yet, the description of experimental protocols (particularly in computer science) is often incomplete and rarely allows for reproducing a study. In this talk, we introduce the audience to the notions of replicability, repeatability and reproducibility and we discuss their relevance in the context of HPC research. Then, we present an effective approach tailored to the specificities of studying HPC systems.

Arnaud Legrand, Luka Stanisic
INRIA Grenoble Rhône-Alpes
arnaud.legrand@imag.fr, luka.stanisic@inria.fr

MS15

High Performance Task-Based QDWH-eig on

Manycore Systems

Current dense symmetric eigensolver implementations suffer from the lack of concurrency during the reduction step toward a condensed tridiagonal matrix and typically achieve a small fraction of theoretical peak performance. Indeed, whether the reduction to tridiagonal form employs a single or multi-stage approach as in the state-of-the-art numerical libraries, the memory-bound nature of some of the computational kernels is still an inherent problem. A novel dense symmetric eigensolver, QDWH-eig, has been presented by Y. Nakatsukasa and N. J. Higham (SIAM SISC, 2012), which addresses these drawbacks with a highly parallel and compute-bound numerical algorithm. However, the arithmetic complexity is an order of magnitude higher than the standard algorithm, which makes it unattractive, at first glance, even with todays relatively low cost of flops. By (1) implementing QDWH-eig framework by the DAG-based tile algorithms (fine-grained) to allow pipelining the various computational stages and to increase computational usage of all processing cores, (2) leveraging the execution of based tile implementation of QDWH-eig with GPUs, the QDWH-eig eigensolver becomes a competitive alternative.

Dalal Sukkari
KAUST
dalal.sukkari@kaust.edu.sa

Hatem Ltaief
King Abdullah University of Science & Technology (KAUST)
hatem.Ltaief@kaust.edu.sa

David E. Keyes
KAUST
david.keyes@kaust.edu.sa

MS16

Bit-Reproducible Climate and Weather Simulations: From Theory to Practice

The problem of reproducibility in floating-point numerical simulations is an active research topic due to the many advantages of reproducible computations, ranging from the simplification of testing and debugging to the deployment of applications on heterogeneous systems maintaining numerical consistency. A reproducible simulation always produces the same result for the same settings regardless of the system it is run on and the configuration of the software. We show principles for achieving bit reproducibility with a case study of the COSMO climate and weather code in the context of the recently started crClim project. We use bit reproducibility to develop a fine-resolution climate model for next-generation supercomputers which necessitates high-performance bit-reproducibility. We discuss sources of non-reproducibility and our strategies for solving them. We will also show how these techniques affect the performance of the application. Overall, we aim at demystifying the belief that determinism is too hard to obtain in high-performance codes and show how reproducibility can be a desirable feature for a numerical model.

Andrea Arteaga
ETH Zurich
andrea.arteaga@env.ethz.ch

Christophe Charpilloz, Olive Fuhrer
MeteoSwiss
christophe.charpilloz@meteoswiss.ch,
oliver.fuhrer@meteoswiss.ch

Torsten Hoefler
ETH Zürich
htor@inf.ethz.ch

MS16

Reproducible and Accurate Algorithms for Numerical Linear Algebra

On modern multi-core, many-core, and heterogeneous architectures, floating-point computations, especially reductions, may become non-deterministic and, therefore, non-reproducible mainly due to the non-associativity of floating-point operations and the dynamic scheduling. We address the problem of reproducibility in the context of fundamental linear algebra operations – like the ones included in the Basic Linear Algebra Subprograms (BLAS) library – and propose algorithms that yields both reproducible and accurate results. We extend this approach to the higher level linear algebra algorithms, e.g. the LU factorization, that are built on top of these BLAS kernels. We present these reproducible and accurate algorithms for the BLAS routines and the LU factorization as well as their implementations in parallel environments such as Intel server CPUs, Intel Xeon Phi, and both NVIDIA and AMD GPUs. We show that the performance of the proposed implementations is comparable to the standard ones.

Roman Iakymchuk
UPMC - LIP6
riakymch@kth.se

David Defour
DALI/LIRMM
University of Perpignan
david.defour@univ-perp.fr

Sylvain Collange
INRIA – Rennes
sylvain.collange@inria.fr

Stef Graillat
University Pierre and Marie Curie (UPMC)
LIP6
stef.graillat@lip6.fr

MS16

Parallel Interval Algorithms, Numerical Guarantees and Numerical Reproducibility

This talk focuses on the parallel implementation of interval algorithms. A first issue is to preserve the inclusion property, that is, the guarantee that a result computed using interval arithmetic contains the exact result. We will illustrate the fact that the inclusion property is at risk with parallel implementations and give some hints to preserve the inclusion property even for parallel interval algorithms. A second issue concerns the ir-reproducibility of numerical computations, i.e. the fact that the same numerical computations deliver different results depending on the environment. We will detail the impact of numerical ir-reproducibility on interval computations and point to several developments that remedy this problem.

Nathalie Revol
INRIA - LIP, Ecole Normale Superieure de Lyon

Nathalie.Revol@ens-lyon.fr

MS16

Summation and Dot Products with Reproducible Results

It can be shown that under very general conditions floating-point summation cannot be associative. Basically, a fixed point system is the only exception. Nevertheless summation and dot product algorithms with reproducible results can be derived, for example using so-called error-free transformations (EFT). We discuss several possibilities of algorithms with identical results regardless of the order of summation, blocking or other techniques to speed up summation.

Siegfried M. Rump
TU Hamburg-Harburg
rump@tuhh.de

MS17

Parallel Generator of Non-Hermitian Matrices with Prescribed Spectrum and Shape Control for Extreme Computing

When designing eigensolvers it is necessary to provide backtesting sample matrices whose spectrum is known beforehand even with constraints on the matrices. We derive a method that produces rigorously non-hermitian matrixes with an exact control of its eigenvalues, shape, sparcity. The method is parallel and scalable in very high dimension. Moreover we show that it is possible to reach absolute arithmetic precision by avoiding all potential troncatures coming from number representations in the computer.

Herve Galicher
KAUST
herve.galicher@kaust.edu.sa

France Boillod
CNRS/LIFL
france.boillod-cerneux@cea.fr

Serge G. Petiton
University Lille 1, Science and Technologies - CNRS
serge.petiton@lifl.fr

Christophe Calvin
CEA-Saclay/DEN/DANS/DM2S
christophe.calvin@cea.fr

MS17

Monte Carlo Algorithms and Stochastic Process on Manycore Architectures

Monte Carlo algorithms are well-known candidates for massively parallel architectures. Nevertheless, from theory to implementation on many core architectures, there always some gaps. We will illustrate in this talk two examples of stochastic algorithms and their practical implementation onto many core architecture. The 1st one comes from Monte Carlo neutron transport and the 2nd is the investagation on potential benefits of stochastic algorithms for solving linear systems.

Christophe Calvin
CEA-Saclay/DEN/DANS/DM2S
christophe.calvin@cea.fr

MS17

Hierarchical Distributed and Parallel Programming for Auto-Tunned Extreme Computing

Exascale hypercomputers will have highly hierarchical architectures with nodes composed by lot-of-core processors and, often, accelerators ; mixing both distributed and parallel computing paradigms. Methods have to be redesigned and some will be rehabilitated or hybridized. The different programming levels will generate new difficult algorithm issues and numerical methods would have to be auto-tuned at runtime to accelerate convergence, minimize energy consumption and scalabilities. In this talk, we present a short survey of auto-tuning of such methods and show that we may use high level graph of component language, associated with multilevel programming paradigm, to help scientists to develop efficient software.

Serge G. Petiton
University Lille 1, Science and Technologies - CNRS
serge.petiton@lifl.fr

MS17

Scalable Sparse Solvers on GPUS

The sparse direct method we present is a multifrontal QR factorization intended specifically for GPU accelerators. Our approach relies on the use of a bucket scheduler that exploits an irregular parallelism on both a coarse grain, among a set of fronts with different characteristics, and on a fine grain, through the exploitation of the staircase shape of these fronts. The scheduler then relies on dense GPU kernels which design and implementation target recent GPU architectures.

Timothy A. Davis, Wissam M. Sid-Lakhdar
Texas A&M
Computer Science and Engineering
davis@tamu.edu, wissam@tamu.edu

MS18

Space-Filling Curves and Adaptive Meshes for Oceanic and Other Applications

sam(oa)^2 (space-filling curves and adaptive meshes for oceanic and other applications) is a software package for parallel adaptive mesh refinement on structured triangular grids created via newest vertex bisection. Exploiting Sierpinski order of the grid cells, it features high memory efficiency and throughput as well as efficient dynamic load balancing on shared and distributed memory. We will report on recent extensions of sam(oa)^2, such as implementation of 3D problems (with 2D adaptivity) and optimization for latest hardware arcitectures, and show application scenarios from tsunami simulation and porous media flow.

Michael Bader
Technische Universität München
bader@in.tum.de

MS18

ADER-DG on Spacetrees in the ExaHyPE Project

We introduce a first version of ExaHyPE, a spacetree-based AMR engine for hyperbolic equation system solvers. It

combines multilevel, higher-order ADER-DG with patch-based finite volumes. The spacetree either acts as compute or as organisational data structure - either to host the higher order shape functions or to organise the patches. Our presentation introduces a reordering of the classic ADER-DG solver phases that allows us to realise the algorithm with a pipelined single-touch policy, to process cells concurrently and to overlap MPI data exchange with the actual computation.

Dominic E. Charrier
Durham University
dominic.e.charrier@durham.ac.uk

Angelika Schwarz
Technische Universität München
angelika.schwarz@in.tum.de

Tobias Weinzierl
Durham University
tobias.weinzierl@durham.ac.uk

MS18

A Tetrahedral Space-Filling Curve Via Bitwise Interleaving

We introduce a space-filling curve for tetrahedral red-refinement that can be computed using bitwise interleaving operations similar to the well-known Morton curve for cubical meshes. We present algorithms that compute the parent, children and face-neighbors of a given mesh element in constant time, as well as the next and previous element in the space-filling curve and whether a given element is on the domain boundary. We close with numerical results obtained with our implementation.

Johannes Holke, Carsten Burstedde
Universität Bonn
holke@ins.uni-bonn.de, burstedde@ins.uni-bonn.de

MS18

Nonconforming Meshes and Mesh Adaptivity with Petsc

Efficient mesh adaptivity requires a sensible combination of data structures and numerical algorithms. Unlike AMR frameworks that try to control both components (structure on the outside, numerics on the inside), this talk will discuss work on using PETSc to drive AMR libraries that provide only data structures (numerics on the outside, structure on the inside). An implementation of this interface using the p4est library will be presented.

Toby Isaac
ICES, University of Texas at Austin
tisaac@ices.utexas.edu

Matthew G. Knepley
Rice University
knepley@gmail.com *Preferred

MS19

Directed Graph Partitioning

In scientific computing directed graphs are commonly used for modeling dependencies among entities. However, while modeling some of the problems as graph partitioning problems, directionality is generally ignored. Accurate modeling of some of the problems necessitates to take the directionally into account, which adds additional constraints that cannot be easily addressed in the current state-of-the-art partitioning methods and tools. In this talk, we will discuss some example problems, models and potential solution approaches for them.

Umit V. Catalyurek
The Ohio State University
Department of Biomedical Informatics
umit@bmi.osu.edu

Kamer Kaya
Sabanci University,
Faculty of Engineering and Natural Sciences
kaya@sabanciuniv.edu

Bora Ucar
LIP, ENS-Lyon, France.
Bora.Ucar@ens-lyon.fr

MS19

Parallel Graph Coloring on Manycore Architectures

In scientific computing, the problem of finding sets of independent tasks is usually addressed with graph coloring. We study performance portable graph coloring algorithms for many-core architectures. We propose a novel edge-based algorithm and enhancements of the speculative Gebremedhin-Manne algorithm that exploit architectures. We show superior quality and execution time of the proposed algorithms on GPUs and Xeon Phi compared to previous work. We present effects of coloring on applications such as Gauss-Seidel preconditioned solvers.

Mehmet Deveci
The Ohio State University
mndevec@sandia.gov

Erik G. Boman
Center for Computing Research
Sandia National Labs
egboman@sandia.gov

Siva Rajamanickam
Sandia National Laboratories
srajama@sandia.gov

MS19

Parallel Approximation Algorithms for b-Edge Covers and Data Privacy

We propose a new 3/2-approximation algorithm, called *LSE* for computing *b-Edge Cover* and it's application to a data privacy problem called adaptive k-Anonymity. *b-Edge Cover* is a special case of the well-known *Set Multicover* problem and also a generalization of *Edge Cover* problem in graphs. The objective is to choose a subset of C edges in the graph with weights on the edges, such that at least a specified number $b(v)$ of edges in C are incident on each vertex v and the sum of edge weights is minimized. We implement the algorithm on serial and shared-memory parallel processors and compare it's performance against a collection of inherently sequential approximation algorithms that have been proposed for the *Set Multicover* problem. With *LSE*, i) we propose the first shared-memory parallel algorithm for the adaptive k-Anonymity problem and

ii) give new theoretical results regarding privacy guarantees which are significantly stronger than the best known previous results.

Arif Khan
Purdue University
khan58@purdue.edu

Alex Pothen
Purdue University
Department of Computer Science
apothen@purdue.edu

MS19

Clustering Sparse Matrices with Information from Both Numerical Values and Pattern

Considering any square fully indecomposable matrix A, we can apply a two-sided diagonal scaling to $|A|$ to render it into doubly stochastic form. The Perron-Frobenius theorem is a key tool to exploit and we aim to use spectral properties of doubly stochastic matrices to reveal hidden block structure in matrices. We also combine this with classical graph analysis techniques to design partitioning algorithms for large sparse matrices based on both numerical values and pattern information.

Iain Duff
Rutherford Appleton Laboratory, Oxfordshire, UK OX11 0QX
and CERFACS, Toulouse, France
iain.duff@stfc.ac.uk

Philip Knight
University of Strathclyde, Glasgow, Scotland
p.a.knight@strath.ac.uk

Sandrine Mouysset
Université deToulouse, UPS-IRIT, France
sandrine.mouysset@irit.fr

Daniel Ruiz
INPT-IRIT-Université deToulouse, France
daniel.ruiz@enseeiht.fr

Bora Ucar
LIP, ENS-Lyon, France.
Bora.Ucar@ens-lyon.fr

MS20

Performance and Power Characterization of Network Algorithms

We present a comparative analysis of the performance and power/energy attributes of several community detection algorithms. We consider the Louvain, WalkTrap, and InfoMap algorithms on several synthetic and real-world networks. We introduce a new fine-grained power/performance measurement infrastructure that is used to estimate tradeoffs among performance, power, and accuracy of community detection algorithms.

Boyana Norris
Argonne National Laboratory
norris@cs.uoregon.edu

Sanjukta Bhowmick
Department of Computer Science
University of Nebraska, Omaha
sbhowmick@unomaha.edu

MS20

The Structure of Communities in Social Networks

Within the area of social network analysis, community detection and understanding community structure are critical to a variety of fields. However, there is little consensus as to what a community actually is, and characterizing the structure of real communities is an important research challenge. I present a machine-learning framework to better understand the structure of real-world communities. Application of this framework to real networks gives us surprising insight into the structure of communities.

Sucheta Soundarajan
Department of Computer Science
Rutgers University
ssoundarajan@gmail.com

MS20

Parallel Algorithms for Analyzing Dynamic Networks

We present, a parallel graph sparsification-based framework for updating the topologies of dynamic large-scaled networks. The challenge in computing dynamic network properties is to update them without recomputing the entire network from scratch. This process becomes more challenging in the parallel domain, because the concurrency in the updates is difficult to determine. To address this issue, we use a technique known as graph sparsification, that divides the graph into subgraphs. The updates can be performed concurrently on each graph and then aggregated over in a reduction-like format over the entire path. We show that this process can provide a general framework for developing scalable algorithms for updating different topologies, including MST and graph connectivity.

Sanjukta Bhowmick
Department of Computer Science
University of Nebraska, Omaha
sbhowmick@unomaha.edu

Sriram Srinivasan
University of Nebraska at Omaha
sriramsrinivasan@unomaha.edu

Sajal Das
Missouri Univeristy of Science and Technology
sdas@mst.edu

MS20

Generating Random Hyperbolic Graphs in Subquadratic Time

NetworKit is our open-source toolkit for large-scale network analysis. Besides numerous analytics kernels it contains various graph generators. *Random hyperbolic graphs* are a promising family of scale-free geometric graphs using hyperbolic geometry. We provide the first generation algorithm for random hyperbolic graphs with subquadratic running time. Using the code available in NetworKit, networks with billions of edges can be generated in a few minutes on a 16-core workstation.

Moritz von Looz

Karlsruhe Institute of Technology (KIT)
Institute of Theoretical Informatics
moritz.looz-corswarem@kit.edu

Henning Meyerhenke
Karlsruhe Institute of Technology
Institute for Theoretical Informatics, Parallel Computing
meyerhenke@kit.edu

Roman Prutkin
Karlsruhe Institute of Technology (KIT)
Institute of Theoretical Informatics
roman.prutkin@kit.edu

MS21

Improving Applicability and Performance for the Multigrid Reduction in Time (mgrit) Algorithm

Since clock speeds are no longer increasing, time integration is becoming a sequential bottleneck. In this talk, we discuss new advances in the MGRIT algorithm and corresponding code XBraid. MGRIT can be easily integrated into existing codes, but more intrusive approaches allow for better performance. We present developments in XBraid to broaden its applicability (e.g., support for temporal adaptivity) and improve efficiency by allowing for varying levels of intrusiveness, dictated by the user's code.

Robert D. Falgout
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
rfalgout@llnl.gov

Veselin Dobrev
Lawrence Livermore National Laboratory
dobrev1@llnl.gov

Tzanio V. Kolev
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
tzanio@llnl.gov

Ben O'Neill
University of Colorado, Boulder
Ben.Oneill@Colorado.edu

Jacob B. Schroder
Lawrence Livermore National Laboratory
schroder2@llnl.gov

Ben Southworth
University of Colorado, Boulder
ben.southworth@colorado.edu

Ulrike Meier Yang
Lawrence Livermore National Laboratory
yang11@llnl.gov

MS21

Multigrid Methods with Space-Time Concurrency

One feasible approach for doing effective parallel-in-time integration is with multigrid. In this talk, we consider the comparison of the three basic options of multigrid algorithms on space-time grids that allow space-time concurrency: coarsening in space and time, semicoarsening in the spatial dimensions, and semicoarsening in the temporal dimension. We discuss advantages and disadvantages of the different approaches and their benefit compared to traditional space-parallel algorithms with sequential time stepping in a parallel computational environment.

Stephanie Friedhoff
University of Leuven
stephanie.friedhoff@uni-koeln.de

Robert D. Falgout, Tzanio V. Kolev
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
rfalgout@llnl.gov, tzanio@llnl.gov

Scott Maclachlan
Department of Mathematics and Statistics
Memorial University of Newfoundland
smaclachlan@mun.ca

Jacob B. Schroder
Lawrence Livermore National Laboratory
schroder2@llnl.gov

Stefan Vandewalle
Department of Computer Science
Katholieke Universiteit Leuven
stefan.vandewalle@cs.kuleuven.be

MS21

Convergence Rate of Parareal Method with Modified Newmark-Beta Algorithm for 2nd-Order Ode

In this study, we introduce a modified Newmark-beta algorithm (MNBA) as a time integrator in the parareal method for 2nd-order ODEs. The MNBA allows us to calculate a highly accurate amplitude and phase regardless of time step width. Therefore, the MNBA is effective for coarse-solver in the parareal algorithm. The accurate calculation is accomplished by a phase correction coefficient. The computed results showed excellent convergence rate for simple harmonic motion with the MNBA.

Mikio Iizuka
RIKEN
mikio.iizuka@kiken.jp

Kenji Ono
VCAD System Research Program, RIKEN
Japan
keno@riken.jp

MS21

PinT for Climate and Weather Simulation

Over the last decade, the key component for increasing the compute performance for HPC systems was an increase in data parallelism. However, for simulations with a fixed problem size, this increase in data parallelism clearly leads to circumstances with the communication latency dominating the overall compute time which again results in a stagnation or even decline of scalability. For weather simulations, further requirements on wallclock time restrictions are given and exceeding these restrictions would make these simulation results less beneficial. For such circumstances, the so-called parallelization-in-time method provide a promising approach. We present our recent results and methods on applying the Parareal approach as one of the parallelization-in-time method to a prototype of nu-

merical weather/climate simulations.

Martin Schreiber
University of Exeter
College of Engineering, Mathematics and Physical Sciences
schreiberx@gmail.com

Pedro S. Peixoto
University of São Paulo, Brazil
pedrosp@ime.usp.br

Terry Haut
Los Alamos National Laboratory
terryhaut@lanl.gov

Beth Wingate
University of Exeter
B.Wingate@exeter.ac.uk

MS22

Increasing Arithmetic Intensity Using Stencil Compilers.

Stencil calculations and matrix-free Krylov sub-space solvers represent an important class of kernels in many scientific computing applications. In such types of solvers, applications of stencil kernels are often the dominant part of the computation, and an efficient parallel implementation of the kernel is therefore crucial in order to reduce the time to solution. Inspired by polynomial preconditioning, we increase the arithmetic intensity of this Krylov subspace building block by replacing the matrix with a higher-degree matrix polynomial. This allows for better use of SIMD vectorization and as a consequence shows better speed-ups on the latest hardware. Portability is made possible with the use of state- of-the-art stencil compiler programs, and an extension to a distributed-memory environment with local shared-memory parallelism is demonstrated. Thus, we demonstrate an effective, portable approach which can easily be adapted to extract high performance from newly developed hardware.

Simplice Donfack
USI
sidonfack@gmail.com

Patrick Sanan
USI - Università della Svizzera italiana
ICS
patrick.sanan@gmail.com

Olaf Schenk
USI - Università della Svizzera italiana
Institute of Computational Science
olaf.schenk@usi.ch

Reps Bram
Universiteit Antwerpen
bram.reps@ua.ac.be

Wim Vanroose
Antwerp University
wim.vanroose@uantwerpen.be

MS22

Many Core Acceleration of the Boundary Element Library BEM4I

We present a library of parallel boundary element based solvers for solution of engineering problems. The library supports solution of problems modelled by the Laplace, Helmholtz, Lame, and time-dependent wave equation. In our contribution we focus on the acceleration of the system matrices assembly using the Intel Xeon Phi coprocessors and parallel solution of linear elasticity problems using the boundary element tearing and interconnecting method (BETI) in cooperation with the ESPRESO library.

Michal Merta
VSB CZ
michal.merta@vsb.cz

Jan Zapletal
IT4Innovations
VSB-Technical University of Ostrava
jan.zapletal@vsb.cz

MS22

HPGMG-FV Benchmark with GASPI/GPI-2: First Steps

The communication library GASPI/GPI is thread based and offers a full flexibility to follow the task-model parallelization approach, being free from the restrictions of the standard hybridization with MPI and threads/OpenMP. Here, after replacing the MPI-communication by a thread-parallel GPI-communication in the prolongation module of HPGMG, we compare the timings of the original implementation and the new one. Our results show that GASPI/GPI might be an appropriate choice for the hierarchically parallel architectures nowadays.

Dimitar M. Stoyanov
ITWM Fraunhofer
dimitar.stoyanov@itwm.fraunhofer.de

MS23

FE2TI - Computational Scale Bridging for Dual-Phase Steels

The computational scale bridging approach (FE2TI) combines the well-known FE2 method with the FETI-DP domain decomposition method. This approach incorporates phenomena on the microscale into a macroscopic problem. FE2TI is used in the project "EXASTEEL - Bridging Scales for Multiphase Steels" (part of the German priority program SPPEXA) for simulations of dual-phase steels. Scalability results for three-dimensional nonlinear and micro-heterogeneous hyperelasticity problems are presented, filling the complete Mira at Argonne National Laboratory (786,432 BlueGene/Q cores).

Martin Lanser
Mathematical Institute
Universit
mlanser@math.uni-koeln.de

Martin Lanser
Mathematical Institute
Universität zu Köln, Germany
mlanser@math.uni-koeln.de

Axel Klawonn
Universität zu Köln
axel.klawonn@uni-koeln.de

Oliver Rheinbach
Technische Universität Bergakademie Freiberg
oliver.rheinbach@math.tu-freiberg.de

MS23

Title Not Available

Abstract not available

Steffen Muething
Universitaet Heidelberg
steffen.muething@iwr.uni-heidelberg.de

MS23

Extreme-Scale Solver for Earth's Mantle – Spectral/Geometric/Algebraic Multigrid Methods for Nonlinear, Heterogeneous Stokes Flow

We present an implicit solver that exhibits optimal algorithmic performance and is capable of extreme scaling for hard PDE problems, such as mantle convection. To minimize runtime the solver advances include: aggressive mesh adaptivity, mixed continuous-discontinuous discretization with high-order accuracy, hybrid spectral/geometric/algebraic multigrid, and novel Schur-complement preconditioning. We demonstrate such algorithmically optimal implicit solvers that scale out to 1.5 million cores for severely nonlinear, ill-conditioned, heterogeneous, and anisotropic PDEs.

Johann Rudi
The University of Texas at Austin
The University of Texas at Austin
johann@ices.utexas.edu

A. Cristiano I. Malossi
IBM Research - Zurich
acm@zurich.ibm.com

Tobin Isaac
The University of Chicago
tisaac@ices.utexas.edu

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

Michael Gurnis
Caltech
gurnis@caltech.edu

Peter W. J. Staar, Yves Ineichen, Costas Bekas
IBM Research - Zurich
taa@zurich.ibm.com, zrlyin@ch.ibm.com,
BEK@zurich.ibm.com

Alessandro Curioni
IBM Research-Zurich
cur@zurich.ibm.com

Omar Ghattas
The University of Texas at Austin
omar@ices.utexas.edu

MS23

3D Parallel Direct Elliptic Solvers Exploiting Hierarchical Low Rank Structure

We describe a parallel direct solver based on a octree-nested condensation strategy and hierarchical compressed matrix representations for the linear algebra operations involved. The algorithm is asymptotically work-optimal and its overall memory requirements grow only linearly with problem size assuming bounded ranks of off-diagonal matrix blocks. The algorithm exposes substantial concurrency in the octree traversals. We show its parallel efficiency with an implementation running on multi-thousand cores on the KAUST Shaheen supercomputer.

Gustavo Chavez
KAUST
gustavo.chavezchavez@kaust.edu.sa

Hatem Ltaief
King Abdullah University of Science & Technology (KAUST)
hatem.Ltaief@kaust.edu.sa

George M Turkiyyah, George M Turkiyyah
American University of Beirut
george.turkiyyah@kaust.edu.sa,
george.turkiyyah@kaust.edu.sa

David E. Keyes
KAUST
david.keyes@kaust.edu.sa

MS24

Performance Optimization of a Massively Parallel Phase-Field Method Using the Hpc Framework Walberla

We present a massively parallel phasefield code, based on the HPC framework waLBerla, for simulating solidification processes of ternary eutectic systems. Different patterns are forming during directional solidification of ternary eutectics, depending on the physical parameters, with significant influence on the mechanical properties of the material. We show simulations of large structure formations, which give rise to spiral growth. Our implementation of this method in the HPC framework waLBerla uses explicit vectorization, asynchronous communication and I/O bandwith reduction techniques. Optimization and scaling results are presented for simulations run on three German supercomputers: SuperMUC, Hornet and JUQUEEN.

Martin Bauer
University of Erlangen-Nuremberg
Department of Computer Science
martin.bauer@fau.de

Harald Köstler
Universität Erlangen-Nürnberg
Department of Computer Science 10 (System Simulation)
harald.koestler@fau.de

Ulrich Rüde
Friedrich-Alexander-University Erlangen-Nuremberg, Germany
ulrich.ruede@fau.de

MS24

Large Scale phase-field simulations of directional

ternary eutectic solidification

Materials with defined properties are crucial for the development and optimization of existing applications. During the directional solidification of ternary eutectic alloys, a wide range of patterns evolve. This patterns are related to the mechanical properties of the material. Large scale phase-field simulation, based on a grand-potential approach, allow to study the underlying physical processes of microstructure evolution. Based on the massive parallel waLBerla framework, a highly optimized solver was developed. Large phase-field simulations of the ternary system Al-Ag-Cu and the three-dimensional structure are presented and visually and quantitatively compared to experiments.

Johannes Hötzer
Hochschule Karlsruhe Technik und Wirtschaft, Germany
tba

Philipp Steinmetz, Britta Nestler
Karlsruhe Institute of Technology (KIT)
Institute of Applied Materials (IAM-ZBS)
philipp.steinmetz@kit.edu, britta.nestler@kit.edu

MS24

Building High Performance Mesoscale Simulations using a Modular Multiphysics Framework

The phasefield module included with MOOSE (the open-source Multiphysics Object Oriented Simulation Environment) allows for the rapid creation and evolution of a wide variety of multiphysics mesoscale simulations. Researchers utilizing this tool have access to powerful finite element capabilities while enjoying highly scalable runs using the built-in hybrid parallelism through the underlying framework. An overview of the advanced capabilities of the phasefield module will be provided including our parsed kernel and material system, fully-coupled linkage to other physics modules, and advances in algorithmic development with the grain tracking capability.

Cody Permann
Center for Advanced Modeling and Simulation
Idaho National Laboratory
cody.permann@inl.gov

MS24

Equilibrium and Dynamics in Multiphase-Field Theories: A Comparative Study and a Consistent Formulation

We present a new multiphase-field theory for describing interface driven phase transitions in multi-domain and/or multi-component systems. We first analyize the most widespread multiphase theories presently in use, to identify their advantages/disadvantages. Then, based on the results of the analysis, a new way of constructing the free energy surface is introduced, and a generalized multiphase description is derived for arbitrary number of phases (or domains). The construction of the free energy functional and the dynamic equations is based on criteria that ensure mathematical and physical consistency. The presented approach retains the variational formalism; reduces (or extends) naturally to lower (or higher) number of fields on the level of both the free energy functional and the dynamic equations; enables the use of arbitrary pairwise equilibrium interfacial properties; penalizes multiple junctions increasingly with the number of phases; ensures non-negative entropy production, and the convergence of the dynamic solutions to the equilibrium solutions; and avoids the appearance of spurious phases on binary interfaces. The new theory is tested for multi-component liquid phase separation (with/without hydrodynamics) and grain boundary roughening.

Gyula I. Toth
Department of Physics and Technology
University of Bergen, NORWAY
Gyula.Toth@ift.uib.no

Laszlo Granasy
Wignes Research Centre for Physics
Budapest, Hungary
granasy.laszlo@wigner.mta.hu

Tamas Pusztai
Wigner Research Centre for Phyics
Budapest, Hungary
pusztai.tamas@wigner.mta.hu

Bjorn Kvamme
Department of Physics and Technology
University of Bergen
bjorn.kvamme@ift.uib.no

MS25

PMI-Exascale: Enabling Applications at Exascale

The PMI-Exascale (PMIx) community is chartered with providing an open source, standalone library that supports a wide range of programming models at exascale and beyond. Critical to the success of these large systems is the ability to rapidly start applications, and the emerging partnership between resource managers and applications for managing workflows. This talk will provide an overview of PMIx objectives, and a status report on the community's progress.

Ralph Castain
Intel
rhc@open-mpi.org

MS25

Middleware for Efficient Performance in HPC

Two middleware components are attracting significant attention today given their notable contributions to the efficiency of HPC systems in terms of energy and throughput, while also improving performance. The first is the use of "local" storage hierarchy, sometimes referred to as "Burst Buffers", though the functionality goes beyond the task of buffering checkpoint files. The second is the deployment of advance power management functionality, from the system level, down through jobs, and into applications themselves.

Larry Kaplan
Cray
lkaplan@cray.com

MS25

Power Aware Resource Management

Maximizing application performance under a power constraint is a key research area for exascale supercomputing. Current HPC systems are designed to be worst-case

power provisioned, leading to under-utilized power and limited application performance. In this talk I will present RMAP, a low-overhead, power-aware resource manager implemented in SLURM that addresses the aforementioned issues. RMAP uses hardware overprovisioning and power-aware backfilling to improve job turnaround times by up to 31% when compared to traditional resource managers.

Tapasya Patki
Lawrence Livemore National Laboratory
tapasya.patki@gmail.com

MS25

Co-Scheduling: Increasing Efficency for HPC

An approach to improve efficiency on HPC systems without changing applications is co-scheduling, i.e., more than one job is executed simultaneously on the same nodes of a system. If co-scheduled applications use different resource types, improved efficiency can be obtained. However, applications may also slow down each other when sharing resources. This talk presents our research activities and results within FAST, a three year project funded by the German ministry of research and education.

Josef Weidendorfer
LRR-TUM
TU München
josef.weidendorfer@tum.de

MS26

Fast Solvers for Complex Fluids

Complex fluids are challenging to simulate as they are characterized by stationary or evolving microstructure. We will discuss the discretization of integral equation formulations for complex fluids. In particular, we will discuss the formulation, numerical challenges and scalability of algorithms for volume integral equations and we will present a new open-source library for such problems. We will compare their performance to other state-of-the art codes.

George Biros
University of Texas at Austin
biros@ices.utexas.edu

Dhairya Malhotra
Institute of Computational Engineering and Sciences
The University of Texas at Austin
dmalhotra@ices.utexas.edu

MS26

Recent Parallel Results Using Dynamic Quad-Tree Refinement for Solving Hyperbolic Conservation Laws on 2d Cartesian Meshes

We discuss recent parallel results using ForestClaw, an adaptive code based on the Berger-Oliger-Colella patch-based scheme for solving hyperbolic conservation laws coupled with the quad/octree mesh refinement available in the p4est library. We show how parallel communication costs vary across different refinement strategies and patch sizes and detail the strategies we use to minimize these costs. Results will be shown for several hyperbolic problems, including shallow water wave equations on the cubed sphere.

Donna Calhoun
Boise State University
donnacalhoun@boisestate.edu

MS26

Simulation of Incompressible Flows on Parallel Octree Grids

We present a second-order accurate solver for the incompressible Navier-Stokes equations on parallel Octree grids. The viscosity is treated implicitly and we use a stable projection method. We rely on the p4est library to handle the distributed Octree structure and capture the presence of irregular interfaces with the level-set method. This leads to a scalable, stable and accurate solver for incompressible flows past complex geometries.

Arthur Guittet
UCSB
arthur.guittet@gmail.com

Frederic G. Gibou
UC Santa Barbara
fgibou@engineering.ucsb.edu

MS26

Basilisk: Simple Abstractions for Octree-Adaptive Schemes

I will present the basic design and algorithms of a new octree-adaptive free software library: Basilisk (http://basilisk.fr). I will show how using simple stencil-based abstractions, one can very simply construct a wide range of octree-adaptive numerical schemes including: multigrid solvers for elliptic problems, volume-of-fluid advection schemes for interfacial flows and high-order flux-based schemes for systems of conservation laws. The automatic and transparent generalisation of these abstractions to large-scale distributed parallel systems will also be discussed.

Stephane Popinet
National Institute of Water and Atmospheric Research
NIWA
stephane.popinet@upmc.fr

MS27

Tradeoffs in Domain Decomposition: Halos Vs Shared Contiguous Arrays

Abstract not available

Jed Brown
Argonne National Laboratory; University of Colorado Boulder
jedbrown@mcs.anl.gov

MS27

Firedrake: Burning the Thread at Both Ends

Abstract not available

Gerard J Gorman
Department of Earth Science and Engineering
Imperial College London
g.gorman@imperial.ac.uk

MS27

Threads: the Wrong Abstraction and the Wrong

Semantic

MPI+OpenMP is frequently proposed as the right evolutionary programming model for exascale. Unfortunately, the evolutionary introduction of OpenMP into existing MPI-only codes is fraught with difficulty. We will describe "The Right Way" to do MPI+OpenMP and ultimately conclude that MPI+MPI is a more effective alternative for legacy codes.

Jeff R. Hammond
Intel Corporation
Parallel Computing Lab
jeff_hammond@acm.org

Timothy Mattson
Intel Corporation
timothy.g.mattson@intel.com

MS28

Recent Directions in Extreme-Scale Solvers

Extreme-scale supercomputers have made extremely large scientific simulations possible, but extreme-scale linear solvers is often a bottleneck. New ideas and algorithms are desperately needed to move forward. In this talk we review several topics of current interest that will be covered in this MS, such as, hierarchical low-rank solvers, fine-grained iterative incomplete factorizations (preconditioners), asynchronous iterations and communication-avoiding (or reducing) iterative solvers.

Erik G. Boman
Center for Computing Research
Sandia National Labs
egboman@sandia.gov

MS28

A New Parallel Hierarchical Preconditioner for Sparse Matrices

We will present a new parallel preconditioner for sparse matrices based on hierarchical matrices. The algorithm shares some elements with the incomplete LU factorization, but uses low-rank approximations to remove fill-ins. The algorithm also uses multiple levels, in a way similar to multigrid, to achieve fast global convergence. Moreover, local communication is needed only for eliminating boundary nodes on every process. We will demonstrate performance and scalability results using various benchmarks.

Chao Chen
Stanford University
cchen10@stanford.edu

Eric F. Darve
Stanford University
Mechanical Engineering Department
darve@stanford.edu

Siva Rajamanickam
Sandia National Laboratories
srajama@sandia.gov

Erik G. Boman
Center for Computing Research
Sandia National Labs
egboman@sandia.gov

MS28

Preconditioning Using Hierarchically Semi-Separable Matrices and Randomized Sampling

We present a software package, called STRUMPACK (Structured Matrices PACKage), to deal with dense or sparse rank structured matrices. Dense rank structured matrices occur for instance in the BEM, quantumchemistry, covariance and machine-learning problems. Our fully algebraic sparse solver or preconditioner uses rank structured kernels on small dense submatrices. The computational complexity and actual runtimes are much lower than for traditional direct methods. We show results for several large problems on modern HPC architectures.

Pieter Ghysels
Lawrence Berkeley National Laboratory
Computational Research Division
pghysels@lbl.gov

Francois-Henry Rouet
Lawrence Berkeley National Laboratory
fhrouet@lbl.gov

Xiaoye Sherry Li
Computational Research Division
Lawrence Berkeley National Laboratory
xsli@lbl.gov

MS28

Preconditioning Communication-Avoiding Krylov Methods

We have previously presented a domain decomposition framework to precondition communication-avoiding (CA) Krylov methods. This framework allowed us to develop a set of preconditioners that do not require any additional communication beyond what the CA Krylov methods have already needed while improving their convergence rates. In this talk, we present our recent extensions to the framework that can improve the performance of the preconditioners.

Ichitaro Yamazaki
University of Tennessee, Knoxville
iyamazak@icl.utk.edu

Siva Rajamanickam, Andrey Prokopenko
Sandia National Laboratories
srajama@sandia.gov, aprokop@sandia.gov

Erik G. Boman
Center for Computing Research
Sandia National Labs
egboman@sandia.gov

Mark Hoemmen, Michael Heroux
Sandia National Laboratories
mhoemme@sandia.gov, maherou@sandia.gov

Stan Tomov
Computer Science Department
University of Tennessee
tomov@eecs.utk.edu

Jack J. Dongarra
Department of Computer Science

The University of Tennessee
dongarra@icl.utk.edu

MS29

Deductive Verification of BSP Algorithms with Subgroups

We present how to perform deductive verification of BSP algorithms with subgroup synchronisation. From BSP programs, we generate sequential codes for the condition generator WHY. By enabling subgroups, the user can prove the correctness of programs that run on hierarchical machines (e.g. clusters of multi-cores). We present how to generate proof obligations of MPI programs that only use collective operations. Our case studies are distributed state-space construction algorithms, the basis of model-checking.

Frédéric Gava
Laboratoire d'Algorithmique, Complexité et Logique
Université de Paris Est
frederic.gava@univ-paris-est.fr

MS29

A BSP Scalability Analysis of 3D LU Factorization with Row Partial Pivoting

Three dimensional matrix multiplication is intuitively no more complex in terms of communication volume than the LU factorization without pivoting. It is known that it is in fact the same. In this presentation, we show that this communication volume lower bound can also be achieved for the LU factorization with "usual" row partial pivoting. The bulk synchronous parallel (BSP) model is used to analyze the scalability of this algorithm and to discuss its tradeoffs.

Antoine Petitet
Huawei Technologies France
antoine.petitet@huawei.com

MS29

Computing Almost Asynchronously: What Problems Can We Solve in $O(logP)$ Supersteps?

The elegance and utility of BSP stem from treating a parallel computation as a sequential one, consisting of supersteps. In many non-trivial computations this number can be independent of the input, depending only on the number of processors p. We give an overview of problems that admit work-optimal BSP solutions using only $O(\log p)$ supersteps: FFT, sorting, suffix array, selection and Longest Common Subsequence, and present general techniques for the design of such almost-asynchronous algorithms.

Alexander Tiskin
Computer Science
University of Warwick
a.tiskin@warwick.ac.uk

MS29

Towards Next-Generation High-Performance BSP

Bulk Synchronous Parallel allows structured programming, ideally making algorithms run predictably, portably, and performantly. This talk focuses on the performance aspect of BSP, in particular its interoperability with ease-of-programming and resiliance, both aspects of BSP which have proven extremely successful in the past decade through MapReduce, Pregel, etc. Can we construct a BSP platform that behaves predictably, is fully portable and resilient, remains easy to program in, while not sacrificing performance?

Albert-Jan N. Yzelman
KU Leuven
Intel ExaScience Lab (Intel Labs Europe)
albert-jan.yzelman@cs.kuleuven.be

MS30

High Performance Computing in Micromechanics

Micromechanics can be characterized as a computation of the global (macrolevel) response of a heterogeneous material possessing a complicated microstructure. It can be implemented via finite element method exploiting information on the microstructure provided by CT. The solution of the finite element systems then requires high performance parallel solution methods. We shall discuss adaptation of an algebraic two level solver exploiting additive Schwarz type approximate inverse for the solution of extreme large finite element systems.

Radim Blaheta
Institute of Geonics, Academy of Sciences of the Czech Republic, Ostrava, Czech Republic
blaheta@ugn.cas.cz

O. Jakl
Institute of Geonics, Academy of Sciences of the Czech Repub
jakl@ugn.cas.cz

J. Stary
Institute of Geonics, Academy of Sciences of the Czech Republic
stary@ugn.cas.cz

I. Georgiev
Institute of Information and Communication Technologies, BAS
Sofia, Bulgaria
georgiev@parallel.bas.bg

MS30

Balanced Dense-Sparse Parallel Solvers for Nonlinear Vibration Analysis of Elastic Structures

The dynamical analysis of elastic structures requires computation of periodic solutions with reliable accuracy, naturally leading to nonlinear models and large-scale systems. A parallel algorithm for implementation of the shooting method is proposed and studied. It involves simultaneously basic matrix operations and solvers for sparse and dense matrices. The computational complexity and parallel efficiency are estimated. Numerical experiments on hybrid HPC systems with accelerators are presented and the achieved scalability is analyzed.

Svetozar Margenov
Institute of Information and Communication Technologies Bulgarian Academy of Sciences, Sofia, Bulgari
margenov@parallel.bas.bg

MS30

Parallel Block Preconditioning Framework and Ef-

ficient Iterative Solution of Discrete Multi-Physics Problems

In this presentation we discuss software engineering issues and parallel performance of the block preconditioning framework - a software middleware designed for simple design and implementation of block preconditioners that arise in common discretisations of multi-physics problems.

Milan D. Mihajlovic
Centre for Novel Computing, Department of Computer Science
The University of Manchester
milan@cs.man.ac.uk

MS30

Simulations for Optimizing Lightweight Structures on Heterogeneous Compute Resources

For lightweight structures in mechanical manufacturing, several parameters have to be determined, which is done in an optimization before production. The optimization is a complex simulation code comprising several CFD and FEM simulation tasks. The challenge is to execute this complex multi-modular code on parallel or distributed hardware environments. Besides others, this leads to a specific scheduling problem of the FEM tasks, which is the main topic of this talk.

Robert Dietze
TU Chemnitz, Germany
dirob@cs.tu-chemnitz.de

Michael Hofmann
Department of Computer Science
Chemnitz University of Technology, Germany
mhofma@cs.tu-chemnitz.de

Gudula Rünger
TU Chemnitz, Germany
ruenger@informatik.tu-chemnitz.de

MS31

Exploiting Synergies with Non-University Partners in CSE Education

The chances and obstacles to involving non-university partners, in our case the federal research center in Jlich, in graduate CSE education will be discussed. Although significant flexibility may be required by both the university and non-university partners in order to chomp with a joint working educational program, the benefits include a sustainable communication platforms, strengthening of the research collaborations, and more opportunities for the graduates. One challenge is dealing with diverse backgrounds of incoming students. Depending on the desired graduate profile, a significant amount of curriculum flexibility may be necessary in order to harmonize the various skill sets. In addition, integrative instruction elements that combine various components of CSE are called for, including custom-designed laboratory modules and projects.

Marek Behr
RWTH Aachen University
Chair for Computational Analysis of Technical Systems
behr@cats.rwth-aachen.de

MS31

Designing a New Masters Degree Program in Computational and Data Science

The Friedrich Schiller University Jena, founded in 1558, is a classical German university with around 20,000 students distributed over ten different departments. In winter 2014/2015, the university started a new master's program in *Computational and Data Science* that is designed to exploit the synergies between computational science for large-scale problems and data-driven approaches to analyze massive data sets. The fundamental methodological basis of this interdisciplinary program consists of parallel processing and scalable algorithms.

Martin Buecker
Mathematics and Computer Science Department
Friedrich Schiller University Jena
martin.buecker@uni-jena.de

MS31

Supporting Data-Driven Businesses Through Talented Students in KnowledgeEngineering@Work

To go beyond the traditional learning experience, KnowledgeEngineering@Work offers its young and talented students in Knowledge Engineering the opportunity to face real and challenging problems on the interface of applied Computer Science and Mathematics, while doing their bachelor. This allows them to become professionals that are tailor made for the fast growing, highly competitive and demanding market in which knowledge of how to handle large data sets is more and more important.

Carlo Galuzzi
Dept. of Knowledge Engineering
Maastricht University
c.galuzzi@maastrichtuniversity.nl

MS31

Blending the Vocational Training and University Education Concepts in the Mathematical Technical Software Developer (matse) Program

The challenging problems in natural and engineering sciences require a close cooperation between mathematicians, computer scientists and end-users. This talk presents the blend of a vocational training as Mathematical Technical Software Developer (MATSE) and the training-accompanying bachelor program Scientific Programming. The graduates acquire the skills to handle mathematical models as well as to find and implement solutions of simulation, optimization and visualization problems.

Benno Willemsen
IT Center
RWTH Aachen University
willemsen@itc.rwth-aachen.de

MS32

Parallel Decomposition Methods for Structured MIP

Large-scale optimization with integer variables is a challenging problem with many scientific and engineering applications, because of the complexity, nonlinearity and size

of the problems. However, certain classes of the problems (e.g., stochastic and/or network optimization) have structures that allow decomposition of the problems. We present decomposition techniques and implementations capable of running on high performance computing systems. Numerical examples are provided using large-scale stochastic unit commitment problems.

Kibaek Kim
Department of Industrial Engineering and Management Sciences
Northwestern University
kimk@anl.gov

MS32

Parallel Optimization of Complex Energy Systems on High-Performance Computing Platforms

We present scalable algorithms and PIPS suite of parallel implementations for the stochastic economic optimization of complex energy systems and parameter estimation for power grid systems under dynamic transients due to unforeseen contingencies. We will also discuss HPC simulations related to the stochastic optimization of electricity dispatch in the State of Illinois' grid under large penetration of wind power.

Cosmin G. Petra
Argonne National Laboratory
Mathematics and Computer Science Division
petra@mcs.anl.gov

MS32

PIPS-SBB: A Parallel Distributed Memory Solver for Stochastic Mixed-Integer Optimization

Stochastic Mixed-Integer Programming (MIP) formulations of optimization problems from real-world applications such as unit commitment can exceed available memory on a single workstation. To overcome this limitation, we have developed PIPS-SBB, a parallel distributed memory Branch-and-Bound based stochastic MIP open-source solver that can leverage HPC architectures. In this talk, we discuss ongoing work on PIPS-SBB, and illustrate its performance on large unit commitment problem instances that cannot be solved to guaranteed optimality using serial implementations.

Deepak Rajan
Lawrence Livermore National Laboratory
rajan3@llnl.gov

Lluis Munguia
Georgia Institute of Technology
lluis.munguia@gatech.edu

Geoffrey M. Oxberry
Lawrence Livermore National Laboratory
oxberry1@llnl.gov

Cosmin G. Petra
Argonne National Laboratory
Mathematics and Computer Science Division
petra@mcs.anl.gov

Pedro Sotorrio, Thomas Edmunds
Lawrence Livermore National Laboratory
sotorrio1@llnl.gov, edmunds2@llnl.gov

MS32

How to Solve Open MIP Instances by Using Supercomputers with Over 80,000 Cores

In this talk, we introduce the Ubiquity Generator (UG) Framework, which is a software framework to parallelize a branch-and-bound based solver on a variety of parallel computing environments. UG gives us a systematic way to develop a parallel solver that can run on large-scale distributed memory computing environments. A success story will be presented where we solve 14 open MIP instances from MIPLIB2003 and MIPLIB2010 using ParaSCIP, a distributed-memory instantiation of UG using SCIP as MIP solver, on up to 80,000 cores of supercomputers.

Yuji Shinano
Zuse Institute Berlin
Takustr. 7 14195 Berlin
shinano@zib.de

Tobias Achterberg
Gurobi GmbH, Germany
achterberg@gurobi.com

Timo Berthold, Stefan Heinz
Fair Issac Europe Ltd, Germany
timoberthold@fico.com, stefanheinz@fico.com

Thorsten Koch
Zuse Institute Berlin
koch@zib.de

Stefan Vigerske
GAMS Software GmbH, Germany
vigerske@gams.com

Michael Winkler
Gurobi GmbH, Germany
winkler@gurobi.com

MS33

Improving Resiliency Through Dynamic Resource Management

Application recovery techniques for exascale systems target quick recovery from node failures using spare nodes. Static resource management assigns a fixed number of spare nodes to an individual job, leading to poor resource utilization and possibly also to insufficiency. Our dynamic resource management scheme allows the execution of non-rigid jobs and can perform on-the-fly replacement of a failed node. Results show better job turnaround times under high failure rates when compared to static resource management.

Suraj Prabhakaran, Felix Wolf
Technische Universität Darmstadt
prabhakaran@cs.tu-darmstadt.de, wolf@cs.tu-darmstadt.de

MS33

Operation Experiences of the K Computer

The K computer is one of the largest supercomputer systems in the world, and it has been operated for more than

three years. Since its release, we have solved a significant number of operation issues such as fixing and improving the system software components, customizing the operation rules according to users needs, and so on. In my talk, I shall talk what we have experienced through the operation of the K computer.

Atsuya Uno
RIKEN Advanced Institute of Computing Science
uno@riken.jp

MS33

The Influence of Data Center Power and Energy Constraints on Future HPC Scheduling Systems

Reducing the power and energy consumption of HPC data centers is one of the essential areas of exascale research. Currently, PC data centers are built and provisioned for expected peak power. But the standard workload power consumption can be as low as 50% of peak power. Therefore, future HPC data centers would then over-provision or right size their data center; meaning that power and cooling is only sufficient for normal operation. Additional power restrictions might come from the power providers (maximum power change in a specific time frame). Other factors could include, for example, peak power vs. energy consumption over the year, real time power consumption adjustments based on renewable energy generation, and moving from CPU hour based accounting to energy consumption. Since scheduling systems are currently the main way to allocate resources for applications, it seems clear that power and energy might need to become a scheduling resource/parameter. This talk will discuss possible data center policies concerning power and energy and their impact on scheduling decisions.

Torsten Wilde
LRZ: Leibniz Supercomputing Centre
Torsten.Wilde@lrz.de

MS34

Resiliency in AMR Using Rollback and Redundancy

Many resiliency strategies introduce redundancy in the applications. Some others use roll-back and recovery. Structured AMR has built-in redundancy and hierarchical structure, both of which can be exploited to build a cost-effective resiliency strategy for AMR. We present an example of a strategy that provides coverage for a variety of faults and the errors they manifest in the context of patch-based AMR for solving hyperbolic PDE's.

Anshu Dubey
Lawrence Berkeley National Laboratory
adubey@anl.gov

MS34

Adaptive Mesh Refinement in Radiation Problems

Radiative heat transfer is an important mechanism in a class of challenging engineering and research problems. A direct all-to-all treatment of these problems is prohibitively expensive on large core counts. Here, we show that radiation calculations can be made to scale within the Uintah framework through a novel combination of reverse Monte Carlo ray tracing techniques combined with adaptive mesh refinement. Strong scaling results from DOE Titan are shown to 256K cores and 16K GPUs.

Alan Humphrey
Scientific Computing and Imaging Institute
University of Utah
ahumphrey@sci.utah.edu

Todd Harman
Department of Mechanical Engineering
University of Utah
t.harman@utah.edu

Daniel Sunderland
Sandia National Laboratories
dsunder@sandia.gov

Martin Berzins
Scientific Computing and Imaging Institute
University of Utah
mb@sci.utah.edu

MS34

Resolving a Vast Range of Spatial and Dynamical Scales in Cosmological Simulations

Solving the equations of ideal hydrodynamics accurately and efficiently is a key feature of state-of-the-art cosmological codes. The AREPO code achieves a high accuracy by utilizing a second-order finite volume method on a moving mesh. Recently, we have developed a second operating mode by adopting a discontinuous Galerkin (DG) scheme on a Cartesian grid, which can be adaptively refined. DG is a higher-order method with a small stencil and thus very suitable for parallel systems.

Kevin Schaal, Volker Springel
Heidelberg Institute for Theoretical Studies
kevin.schaal@h-its.org, Volker.Springel@h-its.org

MS34

A Simple Robust and Accurate a Posteriori Subcell Finite Volume Limiter for the Discontinuous Galerkin Method on Space-Time Adaptive Meshes

We present a class of high order numerical schemes, which include finite volume and discontinuous Galerkin (DG) finite element methods for the solution of hyperbolic conservation laws in multiple space dimensions. For the DG scheme, we employ an ADER-WENO a-posteriori sub-cell finite volume limiter, which is activated only in those regions where strong oscillations are produced. The resulting schemes have been implemented within space-time adaptive Cartesian meshes and MPI parallelization. Several relevant numerical tests are shown.

Olindo Zanotti
Universita degli Studi di Trento
olindo.zanotti@unitn.it

Michael Dumbser
University of Trento
Italy
michael.dumbser@ing.unitn.it

MS35

MPI+MPI: Using MPI-3 Shared Memory As a

MultiCore Programming System

MPI is widely used for parallel programming and enables programmers to achieve good performance. MPI also has a one-sided or put/get programming model. In MPI-3, this one-sided model was greatly expanded, including new atomic read-modify-write operations. In addition, MPI-3 adds direct access to memory shared between processes. This talk discusses some of the opportunities and challenges in making use of this feature to program multicore processors.

William D. Gropp
University of Illinois at Urbana-Champaign
Dept of Computer Science
wgropp@illinois.edu

MS35

How I Learned to Stop Worrying and Love MPI: Prospering with an MPI-Only Model in PFLOTRAN + PETSc in the Manycore Era

The recent rise of "manycore" processors has led to speculation that developers must move from MPI-only models to hybrid programming with MPI+threads to maintain scalability on leadership-class machines. However, careful consideration of communication patterns and judicious use of MPI-3 features presents a viable alternative that retains a simpler programming model. We explore this alternative in the open-source, massively parallel environmental flow reactive transport code PFLOTRAN and the underlying PETSc numerical framework it is built upon.

Richard T. Mills
Oak Ridge National Laboratory
and The University of Tennessee, Knoxville
richardtmills@gmail.com

MS35

MPI+X: Opportunities and Limitations for Heterogeneous Systems

We investigate the use of node-level parallelization approaches such as CUDA, OpenCL, and OpenMP coupled with MPI for data communication across nodes and present comparisons with a conventional flat MPI approach. Our results demonstrate what can be considered an abstraction penalty of MPI+X over flat MPI in the strong scaling limit.

Karl Rupp, Josef Weinbub, Florian Rudolf, Andreas Morhammer, Tibor Grasser
Institute for Microelectronics
Vienna University of Technology
rupp@iue.tuwien.ac.at, weinbub@iue.tuwien.ac.at, rudolf@iue.tuwien.ac.at, morhammer@iue.tuwien.ac.at, grasser@iue.tuwien.ac.at

Ansgar Juengel
Institute for Analysis and Scientific Computing
Vienna University of Technology
juengel@asc.tuwien.ac.at

MS35

The Occa Abstract Threading Model: Implementation and Performance for High-Order Finite-Element Computations

An abstract threading model that simplifies the expression of loop level parallelism will be discussed. This abstraction has been implemented in the OCCA library, including front-end support for C, Fortran, C++, Julia and backend support for OpenMP, CUDA, and OpenCL. Algorithms expressed in the OCCA Kernel Languages are portable across the multiple threading backends. Support is provided for automatic data movement between host and compute devices. Performance of OCCA based finite-element solvers will be presented.

Tim Warburton
Virginia Tech
tcew@vt.edu

David Medina
Rice University
dmed256@gmail.com

MS36

Revisiting Asynchronous Linear Solvers: Provable Convergence Rate Through Randomization

Asynchronous methods for solving systems of linear equations have been researched since Chazan and Miranker's pioneering 1969 paper on chaotic relaxation. The underlying idea of asynchronous methods is to avoid processor idle time by allowing the processors to continue to make progress even if not all progress made by other processors has been communicated to them. Historically, the applicability of asynchronous methods for solving linear equations was limited to certain restricted classes of matrices, such as diagonally dominant matrices. Furthermore, analysis of these methods focused on proving convergence in the limit. Comparison of the asynchronous convergence rate with its synchronous counterpart and its scaling with the number of processors were seldom studied, and are still not well understood. In this talk, we discuss a randomized shared-memory asynchronous method for general symmetric positive definite matrices. We will present a rigorous analysis of the convergence rate that proves that it is linear, and is close to that of the method's synchronous counterpart if the processor count is not excessive relative to the size and sparsity of the matrix. Our work suggests randomization as a key paradigm to serve as a foundation for asynchronous methods.

Haim Avron
Business Analytics & Mathematical Sciences
IBM T.J. Watson Research Center
haimav@us.ibm.com

Alex Druinsky
Computational Research Division
Lawrence Berkeley National Laboratory
adruinsky@lbl.gov

Anshul Gupta
IBM T J Watson Research Center
anshul@us.ibm.com

MS36

Task and Data Parallelism Based Direct Solvers and Preconditioners in Manycore Architectures

Sparse direct or incomplete factorizations are important building blocks in scalable solvers, for example, domain decomposition and multigrid. Due to irregular data access these are difficult to parallelize, especially on manycore architectures. We give an overview of our ongoing parallel

solver projects that exploit either data or task parallelism, or both.

Siva Rajamanickam
Sandia National Laboratories
srajama@sandia.gov

Erik G. Boman, Joshua D. Booth
Center for Computing Research
Sandia National Labs
egboman@sandia.gov, jdbooth@sandia.gov

Kyungjoo Kim
Sandia National Laboratories
kyukim@sandia.gov

MS36

Parallel Ilu and Iterative Sparse Triangular Solves

We consider hybrid preconditioners that are highly parallel and efficient. Therefore, in a first step, we consider stationary iterative solvers for triangular linear equations based on sparse approximate inverse preconditioners. These iterative methods are used for solving triangular systems resulting from parallel ILU preconditioners. For solving general symmetric or nonsymmetric systems we combine sparse approximative inverses with Gauss-Seidel-like splitting methods and the above triangular solvers. This leads to efficient and parallel preconditioners.

Thomas K. Huckle
Institut fuer Informatik
Technische Universitaet Muenchen
huckle@in.tum.de

Juergen Braeckle
TU Muenchen
braeckle@in.tum.de

MS36

Aggregation-Based Algebraic Multigrid on Massively Parallel Computers

To solve large linear systems, it is nowadays standard to use algebraic multigrid methods, although their parallel implementations often suffer on large scale machines. We consider an aggregation-based method whose associated software code (AGMG) has several hundreds of users around the world. To improve parallel performance, some critical components are redesigned, in a relatively simple yet not straightforward way. Excellent weak scalability results are reported on three petascale machines among the most powerful today available.

Yvan Notay
Universite Libre de Bruxelles
ynotay@ulb.ac.be

MS37

Scalable Parallel Graph Matching for Big Data

The BSP style of programming is popular in graph computations, especially in the Big Data field, where routinely graphs of billions of vertices are processed. A fundamental problem here is to match vertices in a large edge-weighted graph. We present a bulk-synchronous parallel approximation algorithm based on local edge dominance, with partial sorting of vertex neighbours and exploitation of data proximity through a suitable vertex partitioning. The effects of partitioning are investigated experimentally.

Rob H. Bisseling
Utrecht University
Mathematical Institute
r.h.bisseling@uu.nl

MS37

From Cilk to MPI to BSP: The Value of Structured Parallel Programming

Parallel programming languages have still not reached the same maturity, clarity and stability as sequential languages. This situation has a strong negative effect on the industry of parallel software. We explain three important styles of parallel programming: Cilk, MPI, and BSP. We explain how their basic features balance expressive power with simplicity, scalability and predictable performance. We also investigate complex language constructs such as types, modules, objects, and patterns.

Gaetan Hains
Huawei Technologies France
gaetan.hains@huawei.com

MS37

Systematic Development of Functional Bulk Synchronous Parallel Programs

With the current generalisation of parallel architectures and increasing requirement of parallel computation arises the concern of applying formal methods, which allow specifications of parallel and distributed programs to be precisely stated and the conformance of an implementation to be verified using mathematical techniques. However, the complexity of parallel programs, compared to sequential ones, makes them more error-prone and difficult to verify. This calls for a strongly structured form of parallelism, which should not only be equipped with an abstraction or model that conceals much of the complexity of parallel computation, but also provides a systematic way of developing such parallelism from specifications for practically nontrivial examples. Program calculation is a kind of program transformation based on the theory of constructive algorithms. An efficient program is derived step-by-step through a sequence of transformations that preserve the meaning and hence the correctness. With suitable datastructures, program calculation can be used for writing parallel programs. This talk presents the SyDPaCC system, a set of libraries for the proof assistant Coq. SyDPaCC allows to write naive (i.e. inefficient) functional programs then to transform them into efficient versions that could be automatically parallelised within the framework before being extracted from Coq as code in the functional language OCaml plus calls to the parallel functional programming library Bulk Synchronous Parallel ML.

Frédéric Loulergue
Laboratoire d'Informatique Fondamentale d'Orléans
Université d'Orléans
frederic.loulergue@univ-orleans.fr

MS37

Parallel Data Distribution for a Sparse Eigenvalue Problem Based on the BSP Model

We discuss data distribution based on the BSP model to load balance a generalised eigenvalue computation arising

from machine learning. Using CG to solve a linear system with a matrix that is the product of three sparse matrices, the main operations are sparse matrix-vector products. The matrices are partitioned based on hypergraph partitioning with net amalgamation and the vectors are partitioned using an (extended) local-bound method. Various variants of these partitioning methods are analysed and compared.

Dirk Roose
KU Leuven
Dept. of Computer Science
Dirk.Roose@cs.kuleuven.be

Karl Meerbergen
K. U. Leuven
karl.meerbergen@cs.kuleuven.be

Joris Tavernier
KU Leuven
Dept. of Computer Science
joris.tavernier@cs.kuleuven.be

MS38

Efficient Parallel Algorithms for New Unidirectional Models of Nonlinear Optics

We are solving the unidirectional Maxwell equation simulating very fast movement of solitons in laser physics

$$\partial_z E + DE + \partial_\tau E^3 = 0,$$

here E is the field component, D describes the linear dispersion operator. The problem is solved by applying the splitting method, the dispersion step is approximated by using the spectral method, the nonlinear term is approximated by using finite difference method. Parallelization of the numerical algorithm is done by applying the domain decomposition method, the algorithm is implemented by using MPI.

Raimondas Ciegas
Department of Mathematical Modelling
University Vilnius, Lithuani
raimondas.ciegis@vgtu.lt

MS38

Parallelization of Numerical Methods for Time Evolution Processes, Described by Integro-Differential Equations

We deal with the iterative solution of a large scale viscoelastic problem arising in geophysics when studying Earth dynamics under load. In addition to the standard equilibrium equation, the problem poses the challenge to handle a constitutive relation between stress and strains that is of integral form. An iterative solution method is applied during each time step. The focus is on efficient and parallelizable preconditioning, related also to the methods used to advance in time.

Ali Dorostkar
Uppsala University
Department of Information Technology
ali.dorostkar@it.uu.se

Maya Neytcheva
Department of Information Technology
Uppsala University
maya.neytcheva@it.uu.se

MS38

Online Auto-Tuning for Parallel Solution Methods for ODEs

In this talk, we consider auto-tuning of parallel solution methods for initial value problems of systems of ordinary differential equations (ODEs). The performance of ODE solvers is sensitive to the characteristics of the input data. Therefore, the best implementation variant cannot be determined at compile or installation time, and offline auto-tuning is not sufficient. Taking this into account, we developed an online auto-tuning approach, which exploits the time-stepping nature of ODE solvers.

Thomas Rauber
Universität Bayreuth, Germany
Rauber@uni-bayreuth.de

Natalia Kalinnik, Matthias Korch
University Bayreuth
Germany
natalia.kalinnik@uni-bayreuth.de, korch@uni-bayreuth.de

MS38

Optimizing the Use of GPU and Intel Xeon Phi Clusters to Accelerate Stencil Computations

MPDATA is a stencil-based kernel of EULAG geophysical model for simulating thermo-fluid flows across a wide range of scales/physical scenarios. In this work, we investigate methods of optimal adaptation of MPDATA to modern hetereogeneous clusters with GPU and Intel Xeon Phi accelerators. The achieved performance results are presented and discussed. In particular, the sustained performance of 138 Gflop/s is achieved for a single NVIDIA K20x GPU, which scales up to 16.1 Tflop/s for 128 GPUs.

Roman Wyrzykowski
Czestochowa University of Technology,
Poland
roman@icis.pcz.pl

Krzysztof Rojek, Lukasz Szustak
Czestochowa University of Technology
Poland
krojek@icis.pcz.pl, lszustak@icis.pcz.pl

MS39

Parallel Implementations of Ensemble Kalman Filters for Huge Geophysical Models

The Data Assimilation Research Testbed (DART) is a community software facility for ensemble Kalman filter data assimilation. The largest DART applications employ ensembles of size O(100) with models with up to one billion variables. DART combines ensemble information from structured model grids with observations from unstructured networks. A scalable implementation of DART must provide parallel algorithms that load balance computation while efficiently moving massive amounts of data. Newly developed algorithms will be presented along with scaling results.

Jeffrey Anderson
National Center for Atmospheric Research
Institute for Math Applied to Geophysics

jla@ucar.edu

Helen Kershaw
NCAR
hkershaw@ucar.edu

Jonathan Hendricks
National Center for Atmospheric Research
hendric@ucar.edu

Nancy Collins
National Center for Atmospheric Research
Boulder, CO
nancy@ucar.edu

MS39

Development of 4D Ensemble Variational Assimilation at Mto-France

The 4D Ensemble Variational assimilation (4DEnVar) has received a considerable attention at Numerical Weather Prediction centres in the very recent years. Indeed, this formalism allows, in such very large size estimation problems related to Ensemble Kalman Filter, consistent error covariances, while being potentially highly parallel on new computer architectures. Such a system is being developed at Mto-France. We will present the different parallelization opportunities offered by this new system and its overall performance.

Gerald Desroziers
Meteo France
gerald.desroziers@meteo.fr

MS39

Parallel 4D-Var and Particle Smoother Data Assimilation for Atmospheric Chemistry and Meteorology

Parallel implemetation strategies building on two different assimilation methods are considered: Firstly, the four dimensional variational technique applying parallel preconditioning of the square root of the forecast error covariance matrix by the diffusion paradigm. This will be presented in the context of an extended advection-diffusion-reaction equation for atmospheric chemistry modelling with aerosols. Secondly, a particle filter/smoother by an ultra-large (up to O(1000) member) ensemble of a mesoscale meteorological model will be given, with a two-level approach on its parallel implementation for short range prediction, where the role of particle weighting algorithms proofs critical.

Hendrik Elbern
Rhenish Institute for Environmental Research
at the University of Cologne, Cologne, Germany
he@eurad.Uni-Koeln.de

Jonas Berndt, Philipp Franke
Institute for Energy and Climate Research
FZ Juelich
jb@eurad.uni-koeln.de, pf@eurad.uni-koeln.de

MS39

Parallel in Time Strong-Constraint 4D-Var

This talk discusses a scalable parallel-in-time algorithm based on augmented Lagrangian approach to perform four dimensional variational data assimilation. The assimilation window is divided into multiple subintervals to facilitate parallel cost function and gradient computations. The scalability of the method is tested by performing data assimilation with increasing problem sizes (i.e. the assimilation window). The data assimilation is performed for Lorenz-96 and the shallow water model. We also explore the feasibility of using a hybrid-method (a combination of serial and parallel 4D-Vars) to perform the variational data assimilation.

Adrian Sandu
Virginia Polytechnic Institute
and State University
sandu@cs.vt.edu

MS40

Parallel Processing for the Square Kilometer Array Telescope

Abstract not available

David De Roure
University of Oxford
david.deroure@oerc.ox.ac.uk

MS40

High Performance Computing for Astronomical Instrumentation on Large Telescopes

Exploiting the full capabilities of the next generation of extremely large ground based optical telescopes will rely on the technology of adaptive optics. This technique requires multiple wave-front sensing cameras and deformable mirrors running at frame rates of 1 KHz. Meeting these requirements is a significant challenge and requires high performance computational facilities capable of operating with very low latency. The proposed solution requires massively parallel systems based on emerging multi-core technology.

Nigel Dipper
Centre for Advanced Instrumentation
Department of Physics, Durham University
n.a.dipper@durham.ac.uk

MS40

Adaptive Optics Simulation for the European Extremely Large Telescope on Multicore Architectures with Multiple Gpus

We present a high performance comprehensive implementation of a multi-object adaptive optics (MOAO) simulation on multicore architectures with hardware accelerators in the context of computational astronomy. This implementation will be used as an operational testbed for simulating the design of new instruments for the European Extremely Large Telescope project (E-ELT), one of Europe's highest priorities in ground-based astronomy. The simulation corresponds to a multi-step multi-stage procedure, which is fed, near real-time, by system and turbulence data coming from the telescope environment. Using modern multicore architectures associated with the enormous computing power of GPUs, the resulting data-driven compute-intensive simulation of the entire MOAO application, composed of the tomographic reconstructor and the observing sequence, is capable of coping with the aforementioned real-time challenge and stands as a reference implementation for the computational astronomy commu-

nity.

Hatem Ltaief
King Abdullah University of Science & Technology (KAUST)
hatem.Ltaief@kaust.edu.sa

Damien Gratadour
Observatoire de Paris
damien.gratadour@obspm.fr

Ali Charara
KAUST
ali.charara@kaust.edu.sa

Eric Gendron
Observatoire de Paris
eric.gendron@obspm.fr

David E. Keyes
KAUST
david.keyes@kaust.edu.sa

MS40

The SKA (Square Kilometre Array) and Its Computational Challenges

The SKA currently being designed will enable radio astronomers to study the Universe at levels of detail and sensitivity well beyond current instruments. The SKA will be truly a software driven instrument, and its computational demands will be unprecedented and will pose challenges that will need to be tackled and overcome. After introducing briefly the aims and overall design of the SKA, this talk will show how computational power will be needed in several areas such as data transport, storage and analysis, and numerical computation in order to achieve the SKA scientific aims: it would indeed be stretching the boundaries of computational capabilities.

Stefano Salvini
OeRC, the University of Oxford
stef.salvini@oerc.ox.ac.uk

MS41

The Missing High-Performance Computing Fault Model

The path to exascale computing poses several research challenges. Resilience is one of the most important challenges. This talk will present recent work in developing the missing high-performance computing (HPC) fault model. This effort identifies, categorizes and models the fault, error and failure properties of today's HPC systems. It develops a fault taxonomy, catalog and models that capture the observed and inferred conditions in current systems and extrapolates this knowledge to exascale HPC systems.

Christian Engelmann
Oak Ridge National Laboratory
engelmannc@ornl.gov

MS41

Resilience Research is Essential but Failure is Unlikely

Once again the high performance computing community is faced with stark predictions of unreliable, failing and unusable leadership computing systems. Previous predictions were false. We argue that unacceptable failure rates will not emerge soon, but will remain at acceptable levels until the computing community is prepared to handle failures via algorithms and applications. In this presentation we provide evidence and arguments for our hypothesis and discuss how resilience research is still essential and timely.

Michael Heroux
Sandia National Laboratories
maherou@sandia.gov

MS41

Memory Errors in Modern Systems

Several recent publications have shown that hardware faults in the memory subsystem are commonplace. These faults are predicted to become more frequent in future systems that contain orders of magnitude more DRAM and SRAM than found in current memory subsystems. These memory subsystems will need to provide resilience techniques to tolerate these faults when deployed in high-performance computing systems and data centers containing tens of thousands of nodes. In this talk, I will focus on learnings about hardware reliability gathered from systems in the field, and use our data to project to whether hardware reliability will become a larger problem at exascale.

Steven Raasch, Vilas Sridharan
AMD Inc.
steven.raasch@amd.com, vilas.sridharan@amd.com

MS41

Characterizing Faults on Production Systems

Understanding and characterizing faults on production systems is key to generate realistic fault models in order to drive flexible fault management across hardware and software components and to optimize the cost-benefit tradeoffs among performance, resilience, and power consumption. In this talk, we will discuss recent efforts to extract fault data from several production systems at LLNL and to store them in a central data repository. In addition to easy access to fault data, this will also enable correlations with other data (incl. power and performance data) to gain a complete picture and to deduce root causes for observed faults.

Martin Schulz
Lawrence Livermore National Laboratory
schulzm@llnl.gov

MS42

From Titan to Summit: Hybrid Multicore Computing at ORNL

The Titan supercomputer at Oak Ridge National Laboratory was one of the first generally available, leadership computing systems to deploy a hybrid CPU-GPU architecture. In this talk I will describe the Titan system, talk about the lessons we have learned in more than three years of successful operation, and how we are applying those lessons in our planning for the Summit system that will be available to our users in 2018.

Arthur S. Bland
Oak Ridge National Laboratory

blandas@ornl.gov

MS42

Hazelhen Leading HPC Technology and its Impact on Science in Germany and Europe

The new HLRS flagship system Hazelhen (a Cray XC40) is Europes fastest system in the HPCG benchmarking list. Designed to provide maximum sustainable performance to the users of HLRS ignoring peak performance numbers is has even climbed to number 8 in the TOP500 list and is one of the fastest systems in the world used both by research and industry. This talk will present insight into the strategy that was at the heart of the HLRS decision. It will further show how Hazelhen can and has already impacted research and industry by providing high sustained performance for real world applications.

Thomas Bönisch
Univerity of Stuttgart
The High Performance Computing Center Stuttgart
tba

Michael Resch
HLRS, University of Stuttgart, Germany
resch@hlrs.de

MS42

Lessons Learned from Development and Operation of the K Computer

The K computer is one of the most powerful supercomputers in the world. The design concept of the K computer includes not only huge computing power but also usability, reliability, availability and serviceability. In this talk, we introduce some functions of the K computer to realize the concept in hardware and software aspects and discuss about lessons learned from the development and operation experiences of the K computer.

Fumiyoshi Shoji
RIKEN
shoji@riken.jp

MS42

Sequoia to Sierra: The LLNL Strategy

Lawrence Livermore National Laboratory (LLNL) has a long history of leadership in large-scale computing. Our current platform, Sequoia, is a 96 rack BlueGene/Q system that is currently number three on the Top 500 list. Our next platform, Sierra, will be a heterogeneous system delivered by a partnership between IBM, NVIDIA and Mellanox. While the platforms are diverse, they represent a carefully considered strategy. In this talk, we will compare and contrast these platforms and the applications that run on them, and provide a glimpse into LLNL's strategy beyond Sierra.

Bronis R. de Supinski
Livermore Computing
Lawrence Livermore National Laboratory
bronis@llnl.gov

MS43

Partitioning and Task Placement with Zoltan2

Zoltan2 is a library for partitioning, task placement, data ordering, and graph coloring. Part of the Trilinos framework, it is our research and development testbed for geometric and graph-based algorithms on emerging high-performance computing platforms. In this talk, we describe Zoltan2 and highlight its partitioning and task placement algorithms. We demonstrate its benefits in the HOMME atmospheric climate model and in particle-based methods.

Mehmet Deveci
The Ohio State University
mndevec@sandia.gov

Karen D. Devine
Sandia National Laboratories
kddevin@sandia.gov

Erik G. Boman
Center for Computing Research
Sandia National Labs
egboman@sandia.gov

Vitus Leung, Siva Rajamanickam, Mark A. Taylor
Sandia National Laboratories
vjleung@sandia.gov, srajama@sandia.gov, mataylo@sandia.gov

MS43

Lessons Learned in the Development of FASTMath Numerical Software for Current and Next Generation HPC

The FASTMath team has experimented with different programming models, communication reducing strategies, and adaptive load balancing schemes to handle both inter-node and intra-node concurrency on deep NUMA architectures. I will give an overview of our experience to date and thoughts on the additional developments needed to achieve performance portability across a range of architectures and software libraries.

Lori A. Diachin
Lawrence Livermore National Laboratory
diachin2@llnl.gov

MS43

Scalable Nonlinear Solvers and Preconditioners in Climate Applications

High resolution, long-term climate simulations are essential to understanding regional climate variation within global change on the decade scale. Implicit time stepping methods provide a means to increase efficiency by taking time steps commensurate with the physical processes of interest. We present results for a fully implicit time stepping method utilizing FASTMath Trilinos solvers and an approximate block factorization preconditioner in the hydrostatic spectral element dynamical core of the Community Atmosphere Model (CAM-SE).

David J. Gardner
Lawrence Livermore National Laboratory
gardner48@llnl.gov

Richard Archibald
Computational Mathematics Group
Oak Ridge National Labratory
archibaldrk@ornl.gov

Katherine J. Evans
Oak Ridge National Laboratory
evanskj@ornl.gov

Paul Ullrich
University of California, Davis
paullrich@ucdavis.edu

Carol S. Woodward
Lawrence Livermore Nat'l Lab
woodward6@llnl.gov

Patrick H. Worley
Oak Ridge National Laboratory
worleyph@ornl.gov

MS43

Parallel Unstructured Mesh Generation/Adaptation Tools Used in Automated Simulation Workflows

This presentation will overview recent advances in the development of a set of unstructured mesh tools that support all steps in the simulation workflow to be executed in parallel on the same massively parallel computer system. These tools to be discussed address the parallel mesh generation/adaptation, parallel mesh structures and dynamic load balancing capabilities needed. The in-memory coupling of these tools with multiple unstructured mesh analysis code will be overviewed.

Mark S. Shephard
Rensselaer Polytechnic Institute
Scientific Computation Research Center
shephard@rpi.edu

Cameron Smith
Scientific Computation Research Center
Rensselaer Polytechnic Institute
smithc11@rpi.edu

Dan A. Ibanez
Rensselaer Polytechnic Institute
SCOREC
ibaned@rpi.edu

Brian Granzow, Onkar Sahni
Rensselaer Polytechnic Institute
granzb@rpi.edu, sahni@rpi.edu

Kenneth Jansen
University of Colorado at Boulder
kenneth.jansen@colorado.edu

MS44

Coupled Matrix and Tensor Factorizations: Models, Algorithms, and Computational Issues

In the era of data science, it no longer suffices to analyze a single data set if the goal is to unravel the dynamics of a complex system. Therefore, data fusion, i.e., extracting knowledge through the fusion of complementary data sets, is a topic of interest in many fields. We formulate data fusion as a coupled matrix and tensor factorization problem, and discuss extensions of the model as well as various algorithmic approaches for fitting such data fusion models.

Evrim Acar
University of Copenhagen
evrim@life.ku.dk

MS44

Parallel Tensor Compression for Large-Scale Scientific Data

Today's computational simulations often generate more data than can be stored or analyzed. For multi-dimensional data, we can use tensor approximations like the Tucker decomposition to expose low-rank structure and compress the data. This talk presents a distributed-memory parallel algorithm for computing the Tucker decomposition for general dense tensors. We show performance results and demonstrate large compression ratios (with small error) for data generated by applications of combustion science.

Woody N. Austin
University of Texas - Austin
austinwn@cs.utexas.edu

Grey Ballard, Tamara G. Kolda
Sandia National Laboratories
gmballa@sandia.gov, tgkolda@sandia.gov

MS44

Parallelizing and Scaling Tensor Computations

Tensor decomposition is a powerful data analysis technique for extracting information from large-scale multidimensional data and is widely used in diverse applications. Scaling tensor decomposition algorithms on large computing systems is key for realizing the applicability of these powerful techniques on large datasets. We present our tool - ENSIGN Tensor Toolbox (ETTB) - that has optimized tensor decompositions through the application of novel techniques for improving the concurrency, data locality, load balance, and asynchrony in the execution of tensor computations that are critical for scaling them on large systems.

Muthu M. Baskaran, Benoit Meister, Richard Lethin
Reservoir Labs
baskaran@reservoir.com, meister@reservoir.com, lethin@reservoir.com

MS44

Parallel Algorithms for Tensor Completion and Recommender Systems

Low-rank tensor completion is concerned with filling in missing entries in multi-dimensional data. It has proven its versatility in numerous applications, including context-aware recommender systems and multivariate function learning. To handle large-scale datasets and applications that feature high dimensions, the development of distributed algorithms is central. In this talk, we will discuss novel, highly scalable algorithms based on a combination of the canonical polyadic (CP) tensor format with block coordinate descent methods. Although similar algorithms have been proposed for the matrix case, the case of higher dimensions gives rise to a number of new challenges and requires a different paradigm for data distribution. The convergence of our algorithms is analyzed and numerical experiments illustrate their performance on distributed-

memory architectures for tensors from a range of different applications.

Lars Karlsson
Umeå University, Dept. of Computing Science
larsk@cs.umu.se

Daniel Kressner
EPFL Lausanne
Mathicse
daniel.kressner@epfl.ch

Andre Uschmajew
University of Bonn
uschmajew.ins.uni-bonn.de

MS45

Multi-ML: Functional Programming of Multi-BSP Algorithms

The Multi-BSP model is an extension of the well known BSP bridging model. It brings a tree based view of nested components to represent hierarchical architectures. In the past, we designed BSML for functional programming BSP algorithms. We now propose the Multi-ML language as an extension of BSML for programming Multi-BSP algorithms.

Victor Allombert, Julien Tesson
Laboratoire d'Algorithmique, Complexité et Logique
Université Paris-est Créteil
victor.allombert@lacl.fr, julien.tesson@lacl.fr

Frédéric Gava
Laboratoire d'Algorithmique, Complexité et Logique
Université de Paris Est
frederic.gava@univ-paris-est.fr

MS45

New Directions in BSP Computing

BSP now powers parallel computing, big data analytics, graph computing and machine learning at many leading-edge companies. It is used at web companies for their analysis of big data from massive graphs and networks with billions of nodes and trillions of edges. BSP is also at the heart of new projects such as Petuum, for machine learning with billions of parameters. At Huawei Paris Research Center we are developing new BSP algorithms and software platforms for Telecomms and IT.

Bill McColl
Huawei Technologies France
bill.mccoll@huawei.com

MS45

Parallelization Strategies for Distributed Machine Learning

A key issue of both theoretical and engineering interest in effectively building distributed systems for large-scale machine learning (ML) programs underlying a wide spectrum of applications, is the bridging models which specify how parallel workers should be coordinated. I discuss the unique challenges and opportunities in designing these models for ML, and present a stale synchronous parallel model that can speed up parallel ML orders of magnitudes over conventional BSP model and enjoy provable guarantees.

Eric Xing
School of Computer Science
CMU
epxing@cs.cmu.edu

MS45

Panel: The Future of BSP Computing

Bulk Synchronous Parallel has had a large influence on parallel computing, both in classical high-performance computing and, more recently, in Big Data computing. Challenges to BSP are increasing latency costs, non-uniform memory hierarchies and network interconnects, the tension between structured/convenient programming versus high performance, and the cost of fault-tolerance and tail-tolerance. Can BSP be relaxed to unify opposing points of view, or will the future hold tens or even hundreds of domain-specific BSP-style frameworks?

Albert-Jan N. Yzelman
KU Leuven
Intel ExaScience Lab (Intel Labs Europe)
albert-jan.yzelman@cs.kuleuven.be

MS46

Linear and Non-Linear Domain Decomposition Methods for Coupled Problems

In this talk, we consider fast and parallel solution methods for linear and non-linear systems of equations, which arise from the discretization of coupled PDEs. We discuss multigrid and domain decomposition based methods for biomechanical applications, including the coupled electro-mechanical simulation of the human heart, as well as for fluid-structure interaction. We focus in particular on the volume coupling of different PDEs, for which we propose an approach based on variational transfer between non-matching meshes.

Rolf Krause
Institute of Computational Science
University of Lugano
rolf.krause@usi.ch

Erich Foster
Institute of Computational Science
USI Lugano
erich.foster@usi.ch

Marco Favino
Institute of Computational Science
University of Lugano
marco.favino@usi.ch

Patrick Zulian
Institute of Computational Science
USI Lugano
patrick.zulian@usi.ch

MS46

Improving Strong Scalability of Linear Solver for Reservoir Simulation

Reservoir simulation plays an important role in defining and optimizing the production strategies for oil and gas

reservoirs. As mesh sizes, complexity and nonlinearity of the models grow highly parallel and scalable linear solvers are of paramount importance for realistic design and optimization workflows. The industry standard Constrained Pressure Residual (CPR) preconditioner consists of various components that are not well tailored for strong scalability. In this talk we will discuss the fundamental underlying problems and review several solution methods on how to overcome this performance barrier.

Tom B. Jonsthovel
Schlumberger
tjonsthovel@slb.com

MS46

Solving Large Systems for the Elastodynamics Equation to Perform Acoustic Imaging

The field of numerical modeling of sound propagation in the ocean, non destructive testing of materials, and of seismic waves in geological media following earthquakes has been the subject of increasing interest of the underwater acoustics and geophysics communities for many years. Computing an accurate solution, taking into account a complex environment, is still the subject of active research focusing on simulations in the time domain for forward and also adjoint/inverse tomography problems in three dimensions (3D). Nowadays, with the stunning increase in computational power, the full waveform numerical simulation of complex wave propagation problems in the time domain is becoming possible. The purpose of the talk is to present some illustrative examples that emphasize the high potential of a high-order finite-element method for forward and adjoint/inverse problems in geophysics as well as in underwater acoustic and non destructive testing applications. We will thus show examples of numerical modeling of seismic wave propagation at very high resolution in large geological models, and summarize recent results for inverse problems in this area. The spectral-element method, because of its ability to handle coupled fluid-solid regions, is also a natural candidate for performing wave propagation simulations in underwater acoustics in the time domain. Moreover, it is also known for being accurate to model surface and interface waves. In this presentation, we will show configurations in which such types of waves, called Stoneley-Scholte waves, can be generated.

Dimitri Komatitsch, Paul Cristini
Laboratory of Mechanics and Acoustics
CNRS, Marseille, France
komatitsch@lma.cnrs-mrs.fr, cristini@lma.cnrs-mrs.fr

MS46

PCBDDC: Robust BDDC Preconditioners in PETSc

A novel class of preconditioners based on Balancing Domain Decomposition by Constraints (BDDC) methods is available in the PETSc library. Such methods provide sophisticated and almost black-box Domain Decomposition preconditioners of the non-overlapping type which are able to deal with elliptic PDEs with coefficients' heterogeneities and discretized by means of non-standard finite element discretizations. The talk introduces the BDDC algorithm and its current implementation in PETSc. Large scale numerical results for different type of PDEs and finite elements discretizations prove the effectiveness of the method.

Stefano Zampini
KAUST
stefano.zampini@kaust.edu.sa

MS47

Title Not Available

Abstract not available

Mike Fisher
ECMWF
mike.fisher@ecmwf.int

MS47

From Data to Prediction with Quantified Uncertainties: Scalable Parallel Algorithms and Applications to the Flow of Ice Sheets

Given a model containing uncertain parameters, (noisy) observational data, and a prediction quantity of interest, we seek to (1) infer the model parameters from the data via the model, (2) quantify the uncertainty in the inferred parameters, and (3) propagate the resulting uncertain parameters through the model to issue predictions with quantified uncertainties. We present scalable parallel algorithms for this data-to-prediction problem, in the context of modeling the flow of the Antarctic ice sheet. We show that the work, measured in number of forward (and adjoint) model solves, is independent of the state dimension, parameter dimension, and data dimension.

Omar Ghattas
The University of Texas at Austin
omar@ices.utexas.edu

Tobin Isaac
The University of Chicago
tisaac@ices.utexas.edu

Noemi Petra
University of California, Merced
npetra@ucmerced.edu

Georg Stadler
Courant Institute for Mathematical Sciences
New York University
stadler@cims.nyu.edu

MS47

Using Adjoint Based Numerical Error Estimates and Optimized Data Flows to Construct Surrogate Models

Surrogate based uncertainty quantification(UQ) using a suitably chosen ensemble of simulations is the method for choice for most large scale simulations with UQ. We have developed novel methods and workflows to inform the construction of these surrogates using adjoint based methods for error estimation in a scalable and parallel way while managing the large data flows.

Abani Patra
SUNY Buffalo
abani@eng.buffalo.edu

MS47

Parallel Aspects of Diffusion-Based Correlation

Operators in Ensemble-Variational Data Assimilation

Ensemble-variational data assimilation requires multiple evaluations of large correlation matrix-vector products. This presentation describes methods to define such products using implicit diffusion-based operators that can be parallelized in the pseudo-time dimension. A solution method based on the Chebyshev iteration is proposed: it does not require global communications and can be employed to ensure the correlation matrix preserves its SPD attribute for an arbitrary convergence threshold. Examples from a global ocean assimilation system are given.

Anthony Weaver, Jean Tshimanga
CERFACS
anthony.weaver@cerfacs.fr, jean.tshimanga@cerfacs.fr

MS48

Parallel Pic Simulation Based on Fast Electromagnetic Field Solver

In this talk, a parallel symplectic PIC program for solving non-relativistic Vlasov-Maxwell system on unstructured grids for realistic applications will be introduced. The implicit scheme is used in our code. We solve Maxwell equations on tetrahedral mesh by finite element method. An efficient and scalable parallel algebraic solver is used for solving the discrete system. Numerical results will be given to show that our method and parallel program is robust and scalable.

Tao Cui
LSEC, the Institute of Computational Mathematics
Chinese Academy of Sciences
tcui@lsec.cc.ac.cn

MS48

Parallel Application of Monte Carlo Particle Simulation Code JMCT in Nuclear Reator Analysis

Massive data for full-core Monte Carlo simulation leads to memory overload for a single core processor. Tree-based domain decomposition algorithm for combinatorial geometry is developed on JCOGIN. An efficient, robust asynchronous transport algorithm for domain decomposition Monte Carlo is introduced. Combination of domain decomposition and domain replication is designed and provided to improve parallel efficiency. Results of two full-core models are presented using the Monte Carlo particle transport code JMCT, which is developed on JCOGIN.

Gang Li
Institute of Applied Physics and Computational Mathematics,
Beijing, 100094, P.R. China
li_gang@iapcm.ac.cn

MS48

Applications of Smoothed Particle Hydrodynamics (SPH) for Real Engineering Problems: Particle Parallelisation Requirements and a Vision for the Future

Recent and likely future HPC developments of SPH are making massively parallel computations viable; consequently there is a real possibility of achieving a step change in the level of complexity and realism in the SPH applications with huge impact on industry and society. A review of the latest SPH developments and its applications will be presented. The challenges of the SPH real engineering simulations and a vision for the future are discussed.

Benedict Rogers
School of Mechanical, Aerospace and Civil Engineering
University of Manchester
benedict.rogers@manchester.ac.uk

MS48

An Hybrid OpenMP-MPI Parallelization for Variable Resolution Adaptive SPH with Fully Object-Oriented SPHysics

Starting from restructuring the original open-source SPHysics, a new SPH programming framework has been developed employing the full functionality of object-oriented Fortran (OOF) to create a new code that is both powerful, highly adaptable to the computing architectures of emerging technology. The novelties of this approach include: (i) Highly modular, program interface (API), (ii) A novel way to generate the neighbour list for particle interactions is proposed, (iii) Hybrid OpenMP-Message Passing Interface (MPI) parallelization.

Ranato Vacondio
University of Parma
rvacondio@gmail.com

MS49

Fault-Tolerant Multigrid Algorithms Based on Error Confinement and Full Approximation

We introduce a novel fault-tolerance scheme to detect and repair soft faults in multigrid solvers, which essentially combines checksums for linear algebra with multigrid-specific algorithm-based fault-tolerance. To improve efficiency we split the algorithm into two parts. For the smoothing stage, we exploit properties of the full approximation scheme to check and repair the results. Our resulting method significally increases the fault-tolerance of the multigrid algorithm and has only a small impact in the non-faulty case.

Mirco Altenbernd
Universität Stuttgart
mirco.altenbernd@mathematik.uni-stuttgart.de

MS49

Soft Error Detection for Finite Difference Codes on Regular Grids

We investigate a strategy to detect soft errors in iterative or time stepping methods where the iterates live on regular grids. The principle idea is to complement iterates by redundant information that allows to determine if memory copies or messages have arrived correctly at the destination. We test our approach by means of the 2D heat equation. Errors are injected asynchronously into the memory or caches for testing.

Peter Arbenz
ETH Zurich
Computer Science Department
arbenz@inf.ethz.ch

Manuel Kohler
ETH Zurich Computer Science Department

manukohl@student.ethz.ch

MS49

Resilience for Multigrid at the Extreme Scale

We consider scalable geometric multigrid schemes to obtain fast and fault-robust solvers. The lost unknowns in the faulty region are re-computed with local multigrid cycles. Especially effective are strategies in which the recovery in the faulty region is executed in parallel with global iterations and when the local recovery is additionally accelerated. This results in an asynchronous multigrid iteration that can fully compensate faults. Excellent parallel performance is demonstrated for systems up to a trillion unknowns.

Ulrich J. Ruede
University of Erlangen-Nuremberg
Department of Computer Science (Simulation)
ulrich.ruede@fau.de

Markus Huber
TU Muenchen
huber@ma.tum.de

Bjoern Gmeiner
University Erlangen-Nuremberg
bjoern.gmeiner@fau.de

Barbara Wohlmuth
Technical University of Munich
barbara.wohlmuth@ma.tum.de

MS49

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

We present a new approach to solve 2D elliptic PDEs that incorporates resiliency to soft and hard faults. Resiliency to hard faults is obtained at the implementation level by leveraging a fault-tolerant MPI implementation (ULFM-MPI) within a task-based server-client model. Soft faults are handled at the algorithm level by recasting the PDE as a sampling problem followed by a resilient regression-based manipulation of the data to obtain a new solution.

Francesco Rizzi, Karla Morris
Sandia National Laboratories
fnrizzi@sandia.gov, knmorri@sandia.gov

Paul Mycek
Duke University
paul.mycek@duke.edu

Olivier P. Le Matre
LIMSI-CNRS
olm@limsi.fr

Andres Contreras
Duke University
andres.contreras@duke.edu

Khachik Sargsyan, Cosmin Safta
Sandia National Laboratories
ksargsy@sandia.gov, csafta@sandia.gov

Omar M. Knio
Duke University
omar.knio@duke.edu

Bert J. Debusschere
Energy Transportation Center
Sandia National Laboratories, Livermore CA
bjdebus@sandia.gov

MS50

Progress in Performance Portability of the XGC Plasma Microturbulence Codes: Heterogeneous and Manycore Node Architectures

The XGC plasma microturbulence particle-in-cell code models transport in the edge region of a tokamak fusion reactor. XGC scales to over 18,000 nodes/GPUs on Cray XK7 Titan and over 64000 cores on Blue Gene/Q Mira. We present our progress and lessons learned in achieving performance portability to heterogeneous GPU and manycore Xeon Phi architectures by using compiler directives, such as nested OpenMP parallelization for manycores, OpenMP SIMD vectorization, and OpenACC for porting to Nvidia GPUs.

Eduardo F. D'Azevedo
Oak Ridge National Laboratory
Computer Science and Mathematics Division
dazevedoef@ornl.gov

Mark Adams
Lawrence Berkeley Laboratory
mfadams@lbl.gov

Choong-Seock Chang, Stephane Ethier, Robert Hager, Seung-Hoe Ku, Jianying Lang
Princeton Plasma Physics Laboratory
cschang@pppl.gov, ethier@pppl.gov, rhager@pppl.gov, sku@pppl.gov, jlang@pppl.gov

Mark Shepard
Scientific Computing Research Center
Rensselaer Polytechnic Institute
shephard@scorec.rpi.edu

Patrick H. Worley
Oak Ridge National Laboratory
worleyph@ornl.gov

Eisung Yoon
Rensselaer Polytechnic Institute
yoone@rpi.edu

MS50

Research Applications and Analysis of a Fully Eddying Virtual Ocean.

We will describe the mathematical and computational techniques used to realize a global kilometer scale ocean model configuration that includes representation of sea-ice and tidal excitation, and spans scales from planetary gyres to internal tides. Computationally we exploit high-degrees of parallelism in both numerical evaluation and in recording model state to persistent disk storage. Using this capability we create a first of its kind multi-petabyte trajectory archive; sampled at unprecedented spatial and temporal frequency.

Chris Hill

Earth, Atmospheric and Physical Sciences
M.I.T.
cnh@mit.edu

MS50

Challenges in Adjoint-Based Aerodynamic Design for Unsteady Flows

An overview of the use of adjoint-based methods for a broad range of large-scale aerodynamic and multidisciplinary design applications is presented. The general motivation for the approach is outlined, along with a summary of the implementation and experiences gained over two decades of concerted research effort. Specific issues associated with applying these methods to unsteady simulations are highlighted, including next-generation computing challenges currently faced by the community.

Eric Nielsen
NASA Langley
eric.j.nielsen@nasa.gov

MS50

Extreme-Scale Genome Assembly

The rate and cost of genome sequencing continues to improve at an exponential rate, making their analysis a significant computational challenge. In this talk we present HipMer, a high performance de novo assembler, which is a critical first step in many genome analyses. Recent work has shown that efficient distributed memory algorithms using the UPC programming model allows unprecedented scaling to thousands of cores, reducing the human genome assembly time to only a few minutes.

Leonid Oliker, Aydin Buluc, Rob Egan
Lawrence Berkeley National Laboratory
loliker@lbl.gov, abuluc@lbl.gov, rsegan@lbl.gov

Evangelos Georganas
EECS
UC Berkeley
egeor@eecs.berkeley.edu

Steven Hofmeyr
Lawrence Berkeley National Laboratory
shofmeyr@lbl.gov

Daniel Rokhsar, Katherine Yelick
Lawrence Berkeley National Laboratory
University of California at Berkeley
dsrokhsar@gmail.com, yelick@eecs.berkeley.edu

MS51

Unstructured Mesh Adaptation for MPI and Accelerators

Conformal simplex meshes can be adapted to improve discretization error and enable Lagrangian simulations of deforming geometry problems. We present new algorithms that perform mesh modifications in batches, allowing OpenMP and CUDA parallelism inside supercomputer nodes and making good use of MPI 3.0 neighborhood collectives between supercomputer nodes. The mesh is stored in a simple array-based structure, allowing low-overhead coupling with physics codes inside device memory.

Dan A. Ibanez
Rensselaer Polytechnic Institute
SCOREC
ibaned@rpi.edu

MS51

Scaling Unstructured Mesh Computations

An unstructured grid, stabilized finite element flow solver (PHASTA) has proven to be a very effective vehicle for development, testing, and scaling demonstrations for a wide variety of parallel software components. These include unstructured grid meshing, adaptivity, load balancing, linear equation solution, and *in situ* data analytics. In this talk, we discuss how each of these components are integrated to keep unstructured grid simulations on a productive path to exascale partial differential equation simulation.

Kenneth Jansen
University of Colorado at Boulder
kenneth.jansen@colorado.edu

Jed Brown
Mathematics and Computer Science Division
Argonne National Laboratory and CU Boulder
jed@jedbrown.org

Michel Rasquin, Benjamin Matthews
University of Colorado Boulder
michel.rasquin@colorado.edu,
benjamin.a.matthews@colorado.edu

Cameron Smith
Scientific Computation Research Center
Rensselaer Polytechnic Institute
smithc11@rpi.edu

Dan A. Ibanez
Rensselaer Polytechnic Institute
SCOREC
ibaned@rpi.edu

Mark S. Shephard
Rensselaer Polytechnic Institute
Scientific Computation Research Center
shephard@rpi.edu

MS51

Comparing Global Link Arrangements for Dragonfly Networks

High-performance computing systems are shifting away from traditional interconnect topologies to exploit new technologies and to reduce interconnect power consumption. The Dragonfly topology is one promising candidate for new systems, with several variations already in production. It is hierarchical, with local links forming groups and global links joining the groups. At each level, the interconnect is a clique. This presentation shows that the intergroup links can be made in meaningfully different ways.

Vitus Leung
Sandia National Laboratories
vjleung@sandia.gov

David Bunde
Knox College

dbunde@knox.edu

MS51

An Extension of Hypres Structured and Semi-Structured Matrix Classes

The hypre software library provides parallel high performance preconditioners and solvers. One of its attractive features is the provision of a structured and a semi-structured interface. We describe a new structured-grid matrix class that supports rectangular matrices and constant coefficients and a semi-structured-grid matrix class that builds on the new class. We anticipate that an efficient implementation of these new classes will lead to better performance on current and future architectures than hypres current matrix classes.

Ulrike M. Yang
Lawrence Livermore National Laboratory
umyang@llnl.gov

Robert D. Falgout
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
rfalgout@llnl.gov

MS52

High Performance Parallel Sparse Tensor Decompositions

We investigate an efficient parallelization of the iterative algorithms for computing CANDECOMP/PARAFAC (CP) and Tucker sparse tensor decompositions on shared and distributed memory systems. The key operations of each iteration of these algorithms are matricized tensor times Khatri-Rao product (MTTKRP) and n-mode tensor-matrix product. These operations amount to element-wise vector multiplication or vector outer product, followed by a reduction, depending on the sparsity of the tensor. We first investigate a fine and a coarse-grain task definition for these operations, and propose methods to achieve load balance and reduce the communication requirements for their parallel execution on distributed memory systems. Then we investigate efficient parallelization issues on shared memory systems. Based on these investigations, we implement the alternating least squares method for both decompositions in a software library. The talk will contain results for both decomposition schemes on large computing systems with real-world sparse tensors. We have a full-featured implementation for the CP-decomposition scheme and development for the Tucker decomposition in progress.

Oguz Kaya
Inria and ENS Lyon
France
oguz.kaya@ens-lyon.fr

Bora Ucar
LIP, ENS-Lyon, France.
Bora.Ucar@ens-lyon.fr

MS52

An Input-Adaptive and In-Place Approach to Dense Tensor-Times-Matrix Multiply

This talk describes a novel framework, called InTensLi ("intensely"), for producing fast single-node implementations of dense tensor-times-matrix multiply (TTM) of arbitrary dimension. Whereas conventional implementations of TTM rely on explicitly converting the input tensor operand into a matrixin order to be able to use any available and fast general matrix-matrix multiply (GEMM) implementation our frameworks strategy is to carry out the TTM in-place, avoiding this copy. As the resulting implementations expose tuning parameters, this paper also describes a heuristic empirical model for selecting an optimal configuration based on the TTMs inputs. When compared to widely used single-node TTM implementations that are available in the Tensor Toolbox and Cyclops Tensor Framework (CTF), InTensLi's in-place and input-adaptive TTM implementations achieve 4 and 13 speedups, showing GEMM-like performance on a variety of input sizes.

Jiajia Li
College of Computing
Georgia Institute of Technology
jiajiali@gatech.edu

Casey Battaglino, Ioakeim Perros, Jimeng Sun, Richard Vuduc
Georgia Institute of Technology
cbattaglino3@gatech.edu, perros@gatech.edu, jimeng.sun@gmail.com, richie@cc.gatech. edu

MS52

Efficient Factorization with Compressed Sparse Tensors

Computing the Canonical Polyadic Decomposition is bottlenecked by multiplying a sparse tensor by several dense matrices (MTTKRP). MTTKRP algorithms either use a compressed tensor specific to each dimension of the data or a single uncompressed tensor at the cost of additional operations. We introduce the compressed sparse fiber data structure and several parallel MTTKRP algorithms that need only a single compressed tensor. We reduce memory by 58% while achieving 81% of the parallel performance.

Shaden Smith
Dept. Computer Science
Univ. Minnesota
shaden@cs.umn.edu

George Karypis
University of Minnesota / AHPCRC
karypis@cs.umn.edu

MS52

Low Rank Bilinear Algorithms for Symmetric Tensor Contractions

We demonstrate algebraic reorganizations that reduce by two the number of scalar multiplications (bilinear algorithm rank) needed for matrix-vector multiplication with a symmetric matrix and the rank-two symmetric update. We then introduce symmetry-preserving algorithms, which lower the number of scalar multiplications needed for most types of symmetric tensor contractions. By being nested with matrix multiplication, these algorithms reduce the total number of operations needed for squaring matrices and contracting partially-symmetric tensors.

Edgar Solomonik
ETH Zurich

solomonik@inf.ethz.ch

MS53

BigDFT: Flexible DFT Approach to Large Systems Using Adaptive and Localized Basis Functions

Density functional theory is an ideal method for the study of a large range of systems due to its balance between accuracy and computational efficiency, however the applicability of standard implementations is limited to systems of around 1000 atoms or smaller due to their cubic scaling with respect to the number of atoms. In order to overcome these limitations, a number of methods have been developed which use localized orbitals to take advantage of the nearsightedness principle and reformulate the problem in a manner which can be made to scale linearly with the number of atoms, paving the way for calculations on systems of 10,000 atoms or more. We have recently implemented such a method in the BigDFT code (http://bigdft.org), which uses an underlying wavelet basis set that is ideal for linear-scaling calculations due to it exhibiting both orthogonality and compact support. This use of wavelets also has the distinct advantage that calculations can be performed in a choice of boundary conditions - either free, wire, surface or periodic. This distinguishes it from other linear-scaling codes, which are generally restricted to periodic boundary conditions, and thus allows calculations on charged systems using open boundary conditions without the need for adding a compensating background charge. We will first outline the method and implementation within BigDFT, after which we will present results showing both the accuracy and the flexibility of the approach, including a demonstration of the improved scaling compared to standard BigDFT calculations.

Thierry Deutsch
Institute of Nanosciences and Cryogeny
CEA
thierry.deutsch@cea.fr

MS53

Design and Implementation of the PEXSI Solver Interface in SIESTA

We have implemented in the SIESTA code an interface to a new electronic structure solver based on the Pole Expansion and Selected Inversion (PEXSI) technique, whose complexity scales at most quadratically with size and can be applied with good accuracy to all kinds of systems, including metals. We describe the design principles behind the interface and the heuristics employed to integrate PEXSI efficiently in the SIESTA workflow to enable large-scale simulations.

Alberto Garcia
ICMAB
albertog@icmab.es

Lin Lin
University of California, Berkeley
Lawrence Berkeley National Laboratory
linlin@math.berkeley.edu

Georg Huhs
Barcelona Supercomputing Center
georg.huhs@bsc.es

Chao Yang
Lawrence Berkeley National Lab
cyang@lbl.gov

MS53

Wavelets as a Basis Set for Electronic Structure Calculations: The BigDFT Project

Dauchechies wavelets are the only smooth orthogonal basis set with compact support. These properties make them ideally suited as a basis set for electronic structure calculations, combining advantages from plane wave and Gaussian basis sets. I will explain the basic principles of wavelet based multi-resolution analysis and the wavelet based algorithms used in our BigDFT electronic structure code both for the solution the Schroedinger's and Poisson's equation. In addition I will present methods implemented in BigDFT to explore the potential energy surface

Stefan Goedecker
Univ Basel, Switzerland
Stefan.Goedecker@unibas.ch

MS53

DGDFT: A Massively Parallel Electronic Structure Calculation Tool

We describe a massively parallel implementation of the recently developed discontinuous Galerkin density functional theory (DGDFT) method, for efficient large-scale Kohn-Sham DFT based electronic structure calculations. The DGDFT method uses adaptive local basis (ALB) functions generated on-the-fly during the self-consistent field (SCF) iteration to represent the solution to the Kohn-Sham equations. DGDFT can achieve 80% parallel efficiency on 128,000 high performance computing cores when it is used to study the electronic structure of 2D phosphorene systems with 3,500-14,000 atoms. This high parallel efficiency results from a two-level parallelization scheme that we will describe in detail.

Chao Yang
Lawrence Berkeley National Lab
cyang@lbl.gov

MS54

Preconditioning of High Order Edge Elements Type Discretizations for the Time-Harmonic Maxwells Equations

Abstract not available

Victorita Dolean
University of Strathclyde
dolean@unice.fr

MS54

Block Iterative Methods and Recycling for Improved Scalability of Linear Solvers.

On the one hand, block iterative methods may be useful when solving systems with multiple right-hand sides, for example when dealing with time-harmonic Maxwells equations. They indeed offer higher arithmetic intensity, and typically decrease the number of iterations of Krylov solvers. On the other hand, recycling also provides a way to decrease the time to solution of successive linear solves, when all right-hand sides are not available at the same time. I will present some results using both approaches, as

well as their implementation inside the open-source framework HPDDM (https://github.com/hpddm/hpddm).

Pierre Jolivet
CNRS
pierre.jolivet@enseeiht.fr

Pierre-Henri Tournier
University Paris VI
tournier@ljll.math.upmc.fr

MS54

Reliable Domain Decomposition Methods with Multiple Search

Domain Decomposition methods are a family of solvers tailored to very large linear systems that require parallel computers. They proceed by splitting the computational domain into subdomains and then approximating the inverse of the original problem with local inverses coming from the subdomains. I will present some classical domain decomposition methods and show that for realistic simulations (with heterogeneous materials for instance) convergence usually becomes very slow. Then I will explain how this can be fixed by injecting more information into the solver. In particular I will show how using multiple search directions within the conjugate gradient algorithm makes the algorithm more reliable. Efficiency is also taken into account since our solvers are adaptive.

Nicole Spillane
Centre de Mathématiques Appliquées
Ecole Polytechnique
nicole.spillane@cmap.polytechnique.fr

Christophe Bovet
LMT, ENS Cachan, CNRS, Université Paris-Saclay,
94235 Cachan
christophe.bovet@lmt.ens-cachan.fr

Pierre Gosselet
LMT, ENS Cachan
gosselet@lmt.ens-cachan.fr

Daniel J. Rixen
TU Munich
rixen@tum.de

François-Xavier Roux
ONERA
Université Pierre et Marie Curie
roux@ann.jussieu.fr

MS54

Parallel Preconditioners for Problems Arising from Whole-Microwave System Modelling for Brain Imaging

Brain imaging techniques using microwave tomography involve the solution of an inverse problem. Thus, solving the electromagnetic forward problem efficiently in the inversion loop is mandatory. In this work, we use parallel preconditioners based on overlapping Schwarz domain decomposition methods to design fast and robust solvers for time-harmonic Maxwell´ s equations discretized with a Nédélec edge finite element method. We present numerical results for a whole-system modeling of a microwave measurement prototype for brain imaging.

Pierre-Henri Tournier
University Paris VI
tournier@ljll.math.upmc.fr

Pierre Jolivet
CNRS
pierre.jolivet@enseeiht.fr

Frédéric Nataf
CNRS and UPMC
nataf@ann.jussieu.fr

MS55

Scaling Up Autotuning

Autotuning has been successful at the node level. As we move towards larger machines, a challenge will be to scale up autotuning systems to handle full applications at the scale of tens to hundreds of thousands of nodes. In this talk, I will outline some of the challenges of running at scale and describe how we have been working to enhance the Active Harmony system to handle these challenges.

Jeffrey Hollingsworth
University of Maryland
hollings@cs.umd.edu

MS55

Auto-Tuning of Hierarchical Computations with ppOpen-AT

We are now developing ppOpen-AT, which is a directive-base Auto-tuning (AT) language to specify fundamental AT functions, i.e., varying values of parameters, loop transformations, and code selection. Considering with expected hardware of Post Moores era, we focus on optimization for computations with deep hierarchy of 3D memory stack. ppOpen-AT provides code selection to optimize code with respect to layers of the memory. Performance evaluation of AT with a code of FDM will be shown by utilizing the Xeon Phi.

Takahiro Katagiri
The University of Tokyo
katagiri@kata-lab.itc.u-tokyo.ac.jp

Masaharu Matsumoto
Information Technology Center
The University of Tokyo
matsumoto@cc.u-tokyo.ac.jp

Satoshi Ohshima
The University of Tokyo
Information Technology Center
ohshima@cc.u-tokyo.ac.jp

MS55

Auto-Tuning for the SVD of Bidiagonals

It is well known that the SVD of a matrix A, $A = USV^T$, can be obtained from the eigenpairs of the matrix $C_V = A^TA$ or $C_U = AA^T$. Alternatively, the SVD can be obtained from the eigenpairs of the augmented matrix $C_{UV} = [0\ A; A^T\ 0]$. This formulation is particularly attractive if A is a bidiagonal matrix and if only a partial SVD (a subset of singular values and vectors) is desired.

This presentation focuses on C_{UV} in the context of bidiagonal matrices, and discusses how tridiagonal eigensolvers can be applied and tuned in that context. The presentation will also discuss opportunities for parallelism, accuracy and implementation issues.

Osni A. Marques
Lawrence Berkeley National Laboratory
Berkeley, CA
oamarques@lbl.gov

MS55

User-Defined Code Transformation for High Performance Portability

Our research project is developing a code transformation framework, Xevolver, to allow users to define their own code transformation rules and also to customize typical rules for individual applications and systems. Use of user-defined code transformations is effective to achieve a high performance while avoiding specialization of an HPC application code for a particular computing system. As a result, without overcomplicating the code, we can achieve high performance portability across diverse HPC systems.

Hiroyuki Takizawa, Shoichi Hirasawa
Tohoku University/JST CREST
takizawa@cc.tohoku.ac.jp, hirasawa@sc.isc.tohoku.ac.jp

Hiroaki Kobayashi
Tohoku University, Japan
koba@isc.tohoku.ac.jp

MS56

Parallelization of Smoothed Particle Hydrodynamics for Engineering Applications with the Heterogeneous Architectures

Smoothed Particle Hydrodynamics (SPH) method is a powerful technique used to simulate complex free-surface flows. However, one of the main drawbacks of the method is its high computational cost and the large number of particles needed in 3D simulations. High Performance Computing (HPC) becomes essential to accelerate SPH codes and perform large simulations. In this study, parallelization using MPI and GPUs (Graphics Processing Units) is applied for executions on heterogeneous clusters of CPUs and GPUs.

Jose M. Domínguez Alonso
EPHYSLAB Environmental Physics Laboratory
Universidade de Vigo
jmdalonso@gmail.com

MS56

H5hut: a High Performance Hdf5 Toolkit for Particle Simulations

Particle-based simulations running on large high-performance computing systems can generate an enormous amount of data. Post-processing, effectively managing the data on disk and interfacing it with analysis and visualization tools can be challenging, especially for domain scientists without expertise in high-performance I/O and data management. We present the H5hut library, an implementation of several data models for particle-based simulations that encapsulates the complexity of HDF5 and is simple to use, yet does not compromise performance

Achim Gsell
PSI
achim.gsell@psi.ch

MS56

Towards a Highly Scalable Particles Method Based Toolkit: Optimization for Real Applications

A large number of industrial physical problems can be simulated using particle-based methods. In this talk, we will present a toolkit for solving a complex, highly nonlinear and distorted flow using Incompressible Smoothed Particle Hydrodynamics in very large scale. This toolkit is implemented cache friendly and specially designed to preserve the data locality. With unstructured communication mechanism in parallel, this toolkit can also be used with the other type of particle methods, such as PIC.

Xiaohu Guo
Scientific Computing Department
Science and Technology Facilities Council
xiaohu.guo@stfc.ac.uk

MS56

Weakly Compressible Smooth Particle Hydrodynamics Performance on the Intel Xeon Phi

DualSPHysics is a Smoothed Particle Hydrodynamics code to study free-surface flow phenomena where Eulerian methods can be difficult to apply, such as waves or impact of dam-breaks on off-shore structures.The IPCC at the Hartree Centre has been porting and optimising DualSPHysics for native execution on the Xeon Phi processor. Threading, vectorization and arithmetic efficiency was increased resulting in speed ups for both the Xeon CPUs and the Xeon Phi.

Sergi-Enric Siso
The Hartree Centre
Science and Technology Facilities Council
sergi.siso@stfc.ac.uk

MS57

Resilience for Asynchronous Runtime Systems and Algorithms

Parallel asynchronous algorithms have shown to be easy to couple with different kinds of low-overhead forward recovery techniques based on tasks. To manage the asynchrony of the parallel workload, a runtime system is typically used. Since its activity may also be sensitive to errors, some runtime specific resilience techniques are required under certain circumstances. This talk will provide a vulnerability comparison between runtime and algorithmic data structures and also suggest low overhead resilience techniques for them.

Marc Casas
Barcelona Supercomputing Center
marc.casas@bsc.es

MS57

Experiments with FSEFI: The Fine-Grained Soft Error Fault Injector

In this talk we will discuss FSEFI, the Fine-Grained Soft

Error Fault Injector and experiments we have conducted with FSEFI to evaluate resiliency, fault-tolerance, and vulnerability of applications and application kernels. FSEFI is a processor emulator-based fault injector where full applications, operating system kernels, drivers, and middleware run. Users can create experiments and conduct trials of hundreds of thousands of fault injections, collect the results, and build a vulnerability/resiliency profile to gain more information about an application. FSEFI has been in construction since 2011 and has users at two government agencies, several national laboratories, and a handful of universities. It is expected to be open sourced by the Department of Energy in the near future.

Nathan A. DeBardeleben
Los Alamos National Lab
ndebard@lanl.gov

Qiang Guan
Los Alamos National Laboratory
qguan@lanl.gov

MS57

What Error to Expect When You Are Expecting a Bit Flip.

We present a statistical analysis of the expected absolute and relative error given a bit flip in an IEEE-754 scalar. We model the expected value agnostic of the standard used, and discuss the implications of using different standards, e.g., double vs quad or single vs double when a bit flip is expected.

James Elliott
North Carolina State University
jjellio3@ncsu.edu

Mark Hoemmen
Sandia National Laboratories
mhoemme@sandia.gov

Frank Mueller
North Carolina State University
mueller@cs.ncsu.edu

MS57

Silent Data Corruption and Error Propagation in Sparse Matrix Computations

In this talk we highlight silent data corruption in several sparse matrix computations. Error rates are expected to impact next generation machines and silent data corruption (SDC) remains a concern as simulations grown size and complexity. Here we investigate the propagation of SDC in sparse matrix operations such as AMG and other sparse solvers. We detail the injection library used and show how errors propagate through different operations, suggesting different algorithm-based mitigation strategies to help reduce the impact of SDC.

Luke Olson, Jon Calhoun
Department of Computer Science
University of Illinois at Urbana-Champaign
lukeo@illinois.edu, jccalho2@illinois.edu

Marc Snir
University of Illinois at Urbana-Champaign
snir@illinois.edu

MS58

Silicon Photonics for Extreme Scale Computing

Silicon Photonics Integrated Circuits will soon enable the direct integration of optical connectivity onto the NICs. Photonic IO through high data rates and wavelength multiplexing is expected to support Tb/s bandwidth per link thereby substantially increasing the available throughput to the node. In the last 5 years, bandwidth per node in dominant HPC systems has only doubled while node compute power increased by a factor of 5 and global computer power by a factor of 13. By suddenly allowing drastically wider bandwidths, and by reducing the IO area and power dissipation, Integrated Photonics is poised to revolutionize inter-node I/O and trigger in depth changes in HPC architectures. In this talk, we show how the main characteristics of these upcoming photonic links can be estimated, how these performance estimates would influence the overall shape of an Exascale-class interconnect, and finally, how these predictions may influence global architectures.

Keren Bergman, Sebastien Rumley
Columbia University
bergman@ee.columbia.edu, sr3061@columbia.edu

MS58

Quantum Computing: Opportunities and Challenges

The core of many real-world problems consists of extremely difficult combinatorial optimization challenges that cannot be solved accurately by traditional supercomputers in a realistic time frame. A key reason is that classical computing systems are not designed to handle the exponential growth in computational complexity as the number of variables increases. In contrast, quantum computers use a physical evolution process that harnesses quantum mechanical effects. In this talk, we will present NASAs experiences with a 1097-qubit D-Wave 2X system that utilizes a process known as quantum annealing to solve this class of hard optimization problems.

Rupak Biswas
NASA Ames Research Center
Rupak.Biswas@nasa.gov

MS58

Task-Based Execution Models and Co-Design

Future HPC machines carry requirements that force changes to hardware and software models. The tension between these models mandates a co-design approach for juggling compromises. In order to effectively address variability, dynamic behavior, power, and resiliency, task-based execution models are a compromise candidate that neither architecture nor software groups particularly desire – but lack viable alternatives. This talk will showcase some tension points, how task models address them, and provide results based on OCR.

Romain Cledat
Intel Corporation
romain.e.cledat@intel.com

Joshua Fryman
Intel

joshua.b.fryman@intel.com

MS58

Emerging Hardware Technologies for Exascale Systems

To achieve exascale computing, fundamental hardware architectures must change. While many details of the exascale architectures are undefined it is clear that this machine will be a significant departure from current systems. Developers will face increasing parallelism, decreasing single threaded performance as well as significant changes to the memory hierarchy. This talk will explore the continuum of emerging hardware technologies, evaluate their viability and briefly discuss the impact on the application programmer.

David Donofrio
Lawrence Berkeley National Laboratory
ddonofrio@lbl.gov

MS59

Additive and Multiplicative Nonlinear Schwarz

Multiplicative Schwarz preconditioned inexact Newton (MSPIN) is naturally based on partitioning degrees of freedom in a nonlinear PDE system by field rather than by subdomain, where a modest factor of concurrency is sacrificed for robustness. ASPIN (2002), by subdomain, is natural for high concurrency and reduction of global synchronization. We augment ASPINs convergence theory for the multiplicative case and show that MSPIN can be significantly more robust than a global Newton method and than ASPIN.

David E. Keyes
KAUST
david.keyes@kaust.edu.sa

Lulu Liu
King Abdullah University of Science and Technology
lulu.liu@kaust.edu.sa

MS59

A Framework for Nonlinear Feti-Dp and Bddc Domain Decomposition Methods

Solving nonlinear problems, e.g., in material science, requires fast and highly scalable parallel solvers. FETI-DP and BDDC domain decomposition methods meet these requirements for linearized implicit finite element problems. Recently, nonlinear versions have been proposed where the nonlinear problem is decomposed before linearization. This can be viewed as a strategy to further localize computational work and to extend parallel scalability towards extreme-scale supercomputers. A framework for different nonlinear FETI-DP and BDDC algorithms will be presented.

Axel Klawonn
Universität zu Köln
axel.klawonn@uni-koeln.de

Martin Lanser
Mathematical Institute
Universit
mlanser@math.uni-koeln.de

Martin Lanser
Mathematical Institute
Universität zu Köln, Germany
mlanser@math.uni-koeln.de

Oliver Rheinbach
Technische Universität Bergakademie Freiberg
oliver.rheinbach@math.tu-freiberg.de

MS59

Convergence Estimates for Composed Iterations

I will define nonlinear preconditioning, a composition technique used to accelerate an iterative solution process, in analogy with the linear case, and demonstrate its effectiveness for representative nonlinear PDE. I will outline the beginning of a convergence theory for these composed iterations.

Matthew G. Knepley
Rice University
knepley@gmail.com *Preferred

MS59

Scaling Nonlinear FETI-DP Domain Decomposition Methods to 786 452 Cores

Nonlinear FETI-DP (Finite Element Tearing and Interconnecting) domain decomposition methods are parallel solution methods for nonlinear implicit problems discretized by finite elements. A recent nonlinear FETI-DP method is combined with an approach that allows an inexact solution of the FETI-DP coarse problem by a parallel Algebraic Multigrid Method. Weak scalability for two TOP 10 supercomputers, i.e., JUQUEEN at FZ Jülich and the complete Mira at Argonne National Laboratory (786,432 cores) are shown.

Axel Klawonn
Universität zu Köln
axel.klawonn@uni-koeln.de

Martin Lanser
Mathematical Institute
Universit
mlanser@math.uni-koeln.de

Martin Lanser
Mathematical Institute
Universität zu Köln, Germany
mlanser@math.uni-koeln.de

Oliver Rheinbach
Technische Universität Bergakademie Freiberg
oliver.rheinbach@math.tu-freiberg.de

MS60

NUMA-Aware Hessenberg Reduction in the Context of the QR Algorithm

Hessenberg reduction appears in the aggressive early deflation bottleneck in parallel multishift QR algorithms. The reduction is memory-bound and performance relies on efficient use of caches and memory buses. We applied the parallel cache assignment technique recently proposed by Castaldo and Whaley and obtained a cache-efficient and NUMA-aware multithreaded algorithm. Numerous tunable parameters were identified and their impact on per-

formance evaluated. The goal is to incorporate runtime auto-tuning capabilities.

Mahmoud Eljammaly
Umeå University
Dept. of Computing Science
mjammaly@cs.umu.se

Lars Karlsson
Umeå University, Dept. of Computing Science
larsk@cs.umu.se

MS60

Low Rank Approximation of Sparse Matrices for Lu Factorization Based on Row and Column Permutations

In this talk we present an algorithm for computing a low rank approximation of a sparse matrix based on a truncated LU factorization with column and row permutations. We present various approaches for determining the column and row permutations that show a trade-off between speed versus deterministic/probabilistic accuracy. We show that if the permutations are chosen by using tournament pivoting based on QR factorization, then the obtained truncated LU factorization with column/row tournament pivoting, LU_CRTP, satisfies bounds on the singular values which have similarities with the ones obtained by a communication avoiding rank revealing QR factorization. We also compare the computational complexity of our algorithm with randomized algorithms and show that for sparse matrices and certain accuracy, our approach is faster.

Laura Grigori
INRIA
France
Laura.Grigori@inria.fr

Sebastien Cayrols
INRIA
sebastien.cayrols@inria.fr

James W. Demmel
University of California
Division of Computer Science
demmel@berkeley.edu

MS60

The Challenge of Strong Scaling for Direct Methods

We present an analysis of the problems of getting near full performance for sparse direct methods, particularly for symmetric systems, on complex modern hardware. Our discussion is informed by experience on both modern CPU and GPU hardware. We focus on getting high performance on small problems that are still important in many communities. The algorithms and techniques described can be used as building blocks to address the solution of larger systems on HPC resources.

Jonathan Hogg
STFC Rutherford Appleton Lab, UK
jonathan.hogg@stfc.ac.uk

MS60

Assessing Recent Sparse Matrix Storage Formats for Parallel Sparse Matrix-Vector Multiplication

The performance of parallel sparse matrix-vector multiplication (SpMV for short, calculating $y \leftarrow Ax$, where A is a sparse matrix, x and y are dense vectors) highly depends on the storage format of the input matrix. Recently, many formats have been proposed for many-core processors such as GPUs and Xeon Phi. This talk will give a comprehensive comparison for SpMV performance and format conversion cost (from the CSR to a new format) of recently developed formats including HYB (SC '09), Cocktail (ICS '12), ESB (ICS '13), BCCOO (PPoPP '14), BRC (ICS '14), ACSR (SC '14), CSR-Adaptive (SC '14, HiPC '15) and CSR5 (ICS '15).

Weifeng Liu
Niels Bohr Institute
University of Copenhagen
weifeng.liu@nbi.ku.dk

MS61

New Building Blocks for Scalable Electronic Structure Theory

We highlight two central building blocks for scalable electronic structure theory on current high-performance computers:(i) New methodological approaches beyond semilocal density-functional theory based on a localized resolution-of-indentity approach for the two-electron Coulomb operator; (ii) a new, general open-source library platform "ELSI" that will provide a seamless interface to address or circumvent the Kohn-Sham eigenvalue problem. Current work focuses on the massively parallel eigenvalue solver ELPA, the orbital minimization method (OMM), and the PEXSI library.

Volker Blum
Duke University
Department of Mechanical Engineering and Materials Science
volker.blum@duke.edu

MS61

MOLGW: A Parallel Many-Body Perturbation Theory Code for Finite Systems

The many-body perturbation theory was mostly developed for condensed matter. The successes of the *GW* approximation to the electron self-energy for the band structure and of the Bethe-Salpeter equation for the absorption spectrum are uncountable for solids. However, very little is known about the performance of many-body perturbation theory for atoms and molecules. The code MOLGW addresses the efficient and accurate calculation within this framework using modern parallelization, based on MPI and SCALAPACK.

Fabien Bruneval
CEA, DEN, SRMP
fabien.bruneval@cea.fr

Samia M. Hamed, Tonatiuh Rangel-Gordillo, Jeffrey B. Neaton
UC Berkeley (USA)
shamed@lbl.gov, trangel@lbl.gov, jbneaton@lbl.gov

MS61

GPU-Accelerated Simulation of Nano-Devices

from First-Principles

This talk will introduce a newly developed algorithm called SplitSolve to take care of the Schrödinger equation with open boundary conditions, as encountered in quantum transport calculations. The strengths of SplitSolve are the significant reduction of the computational time it offers and its ability to fully leverage hybrid nodes with CPUs and GPUs. It thus allows for the transport simulation of nanostructures composed of thousands of atoms at the *ab-initio* (from first-principles) level.

Mathieu Luisier, Mauro Calderara, Sascha Brueck
Integrated Systems Laboratory
ETH Zurich
mluisier@iis.ee.ethz.ch, mcalderara@iis.ee.ethz.ch, bruecks@iis.ee.ethz.ch

Hossein Bani-Hashemian, Joost VandeVondele
Nanoscale Simulations
ETH Zurich
hossein.banihashemian@mat.ethz.ch, joost.vandevondele@mat.ethz.ch

MS61

High-Performance Real-Time TDDFT Excited-States Calculations

From atoms and molecules to nanostructures, we discuss how the FEAST solver has considerably broaden the perspectives for enabling reliable and high-performance large-scale first-principle DFT and real-time TDDFT calculations. Our TDDFT simulations can operate a broad range of electronic spectroscopy: UV-ViS, X-ray (using our all-electron method), and near IR (using our TDDFT-Ehrenfest approach). Various numerical results will be presented and discussed.

Eric Polizzi
University of Massachusetts, Amherst, USA
polizzi@ecs.umass.edu

MS62

Composition and Compartmentalisation As Enabling Features for Data-Centric, Extreme Scale Applications: An MPI-X Approach

Next generation extreme scale programming model is often envisioned as MPI+X, where the second layer X is one of OpenMP, OpenACC, etc. Conversely, we advocate a MPI-X approach, where MPI is only used as communication library and applications are composed assembling in streaming fashion data-parallel kernels described as plain C++11 classes. They are designed to run on compositional accelerators, which can be implemented as cluster of devices, e.g. multi-cores, GPUs, FPGAs or a heterogeneous combination of them.

Marco Aldinucci, Maurizio Drocco, Claudia Misale
University of Torino
Italy
aldinuc@di.unito.it, drocco@di.unito.it, misale@di.unito.it

Massimo Torquati
University of Pisa
Italy
torquati@di.unipi.it

MS62

Recent Advances in the Task-Based StarPU Runtime Ecosystem

The task paradigm has allowed the StarPU dynamic runtime system to achieve scalable performance for various applications over heterogeneous systems and clusters thereof. We will here discuss the latest advances which allow us to proceed a step further: controlling the memory footprint (notably of applications with compressed data), employing local disks, achieving perfectly reproducible experiments through simulation, and showing how far we are from perfect execution.

Samuel Thibault
University of Bordeaux 1, Inria
samuel.thibault@labri.fr

Marc Sergent
INRIA, France
marc.sergent@inria.fr

Suraj Kumar
INRIA
France
suraj.kumar@inria.fr

Luka Stanisic
INRIA Grenoble Rhône-Alpes
luka.stanisic@inria.fr

MS62

Task Dataflow Scheduling and the Quest for Extremely Fine-Grain Parallelism

Task dataflow programming models provide a useful programming abstraction but the underpinning dynamic schedulers carry a burden that imposes limits on strong scaling. This work proposes methods to assess the inherent overhead of schedulers and applies them to assess the burden of industry-standard schedulers. We demonstrate that the scheduler burden explains the performance scalability of a set of memory-intensive and compute-intensive codes. Based on this we develop a novel scheduler that demonstrates increased scalability.

Hans Vandierendonck, Mahwish Arif
Queen's University Belfast
Electronics, Electrical Engineering, and Computer Science
h.vandierendonck@qub.ac.uk, m.arif@qub.ac.uk

MS62

A Single Interface for Multiple Task-Based Parallel Programing Paradigms/Schemes

Task-based parallel programming schemes have proven advantageous regarding ease of programming, efficiency and performance. By employing a hierarchical definition of tasks and data, locality and bandwidth utilization can be maximized for a specific underlying hardware. However, explicitly encoding each level of parallelism becomes cumbersome. We present a unified programming interface that allows an algorithm written in terms of non-hierarchical tasks to be scheduled hierarchically across different shared and distributed memory environments by a task-based run-

time.

Afshin Zafari
Uppsala University
Dept. of Information Technology
afshin.zafari@it.uu.se

Elisabeth Larsson
Uppsala University, Sweden
Department of Information Technology
Elisabeth.Larsson@it.uu.se

MS63

Hardware/Algorithm Co-Optimization and Co-Tuning in the Post Moore Era

In the Post Moore Era performance and power efficiency is obtained from better resource usage, since no free speed up due to technology scaling is helping along any more. In this talk we will present how hardware specialization and design space exploration in conjunction with software synthesis and autotuning provides a technological basis for continued performance and power gains in the Post Moore Era. Users write their programs against well-established math library interfaces like the BLAS and FFTW interface, and behind the scenes code synthesis and hardware/algorithm co-design leverages dark silicon, 3D stacking, and emerging memory technologies to translate substantial functional blocks into highly optimized hardware constructs that can be software configured. This is work in progress; however, we have demonstrated the power of this approach with space time adaptive processing, an FFT and linear algebra heavy kernel.

Franz Franchetti
Department of Electrical and Computer Engineering
Carnegie Mellon University
franzf@ece.cmu.edu

MS63

On Guided Installation of Basic Linear Algebra Routines in Nodes with Manycore Components

Most computational systems are composed of multicore nodes together with several GPUs or Xeon Phi, which might be of different types. This work analyses the adaptation of auto-tuning techniques used for hybrid CPU+GPU systems to these more complex configurations. The study is conducted on basic linear algebra kernels. Satisfactory results are obtained, with experimental execution times close to the lowest experimentally obtained.

Javier Cuenca
Departamento de Ingeniera y Tecnologa de Computadores
University of Murcia
jcuenca@um.es

Luis-Pedro García
Servicio de Apoyo a la Investigación Cient{ı}fica
Technical University of Cartagena
luis.garcia@sait.upct.es

Domingo Giménez
Departamento de Informática y Sistemas
University of Murcia
domingo@um.es

MS63

Auto-Tuning for Eigenvalue Solver on the Post Moore's Era

Dense Real-Symmetric eigenvalue problem consists of three main procedures. Each of three procedures has a unique aspect from the viewpoint of the insensitivity of memory access or floating point operations. In the Post-Moores era, the present tuning strategy, which we reduce data traffic cost by CA (Communication Avoidance), SR (Synchronous Reduction), and AS (Asynchronous algorithm), may change back into or mix up with the era of vector computer. Automatic-tuning helps to find the sub-optimal combination of these strategies.

Toshiyuki Imamura
RIKEN Advance Institute for Computational Science
imamura.toshiyuki@riken.jp

MS63

Automatic Tuning for Parallel FFTs on Intel Xeon Phi Clusters

In this talk, we propose an implementation of a parallel fast Fourier transform (FFT) with automatic performance tuning on Intel Xeon Phi clusters. Parallel FFTs on Intel Xeon Phi clusters require intensive all-to-all communication. An automatic tuning facility for selecting the optimal parameters of overlapping between computation and communication, the all-to-all communication, the radices, and the block size is implemented. Performance results of FFTs on an Intel Xeon Phi cluster are reported.

Daisuke Takahashi
Faculty of Engineering, Information and Systems
University of Tsukuba
daisuke@cs.tsukuba.ac.jp

MS64

Thermal Comfort Evaluation for Indoor Airflows Using a Parallel CFD Approach

The thermal comfort prediction of occupants in rooms and vehicles is of great interest during planning and operation phases. With the evermore increasing computational power available, a great variety of problems can be solved which were unthinkable until recently. This talk will present and discuss coupling procedures and strategies to integrate human thermoregulation based on Fialas model into an in-house CFD code specifically designed with parallel systems in mind to investigate indoor airflows.

Jérôme Frisch
Institute of Energy Efficiency and Sustainable Building
E3D
RWTH Aachen University
frisch@e3d.rwth-aachen.de

Ralf-Peter Mundani
TUM, Faculty of Civil, Geo and Environmental Engineering
Chair for Computation in Engineering

mundani@tum.de

MS64

Fast Contact Detection for Granulates

We study the problem of contact detection between large numbers of triangles resulting from meshed objects. This problem arises in FSI and DEM, e.g., and classic algorithms from computational geometry come along with many case distinctions. We propose an iterative, weak solver that vectorises but may fall back to a classic formulation if the Newton steps do not converge quickly. Based upon a tree domain decomposition, it furthermore is parallelised via MPI.

Tomasz Koziara
Durham University
tomasz.koziara@durham.ac.uk

Konstantinos Krestenitis
Durham University
Durham, UK
konstantinos.krestenitis@durham.ac.uk

Jon Trevelyan, Tobias Weinzierl
Durham University
jon.trevelyan@durham.ac.uk, tobias.weinzierl@durham.ac.uk

MS64

LB Particle Coupling on Dynamically-Adaptive Cartesian Grids

Dynamically-adaptive, tree-structured grids are a well-known way of solving partial differential equations with less computational effort than regular Cartesian grids. We focus on a minimally-invasive approach using p4est, a well-known framework for such grids, which allows reusing huge parts of the existing code. We integrated p4est as a new mesh in ESPResSo, a large MD-simulation package. In our contribution, we present implementational details and performance data for various background flow scenarios.

Michael Lahnert
University of Stuttgart
michael.lahnert@ipvs.uni-stuttgart.de

Miriam Mehl
Universität Stuttgart
miriam.mehl@ipvs.uni-stuttgart.de

MS64

High-Performance Computing in Free Surface Fluid Flow Simulations

Nowadays, more and more questions originate from the so-called emerging sciences such as medicine, sociology, biology, virology, chemistry, climate or geo-sciences. Here, we present an advanced simulation and (visual) exploration concept for various environmental engineering applications such as porous media flows and high water simulations. Based on complex 3D building information and environmental models, we can perform multi-resolution simulations on massive parallel systems combined with an interactive exploration of the computed data during runtime.

Nevena Perovic
Chair for computation in engineering
TU München
nevena.perovic@tum.de

Jérôme Frisch
Institute of Energy Efficiency and Sustainable Building E3D
RWTH Aachen University
frisch@e3d.rwth-aachen.de

Ralf-Peter Mundani
TUM, Faculty of Civil, Geo and Environmental Engineering
Chair for Computation in Engineering
mundani@tum.de

Ernst Rank
Computation in Engineering
Technische Universität München
ernst.rank@tum.de

MS65

A Flexible Area-Time Model for Analyzing and Managing Application Resilience

Traditional error recovery systems such as checkpoint-restart fix key aspects of the error model, detection timing, and recovery approach. As a result, these methods have limited ability to adapt to different error types, statistics, or algorithm behavior. We have designed a new model, area-time, a general, flexible framework for reasoning about errors, detection, and response. We apply this model to several specific application examples to demonstrate its utility.

Andrew A. Chien
The University of Chicago
Argonne National Laboratory
achien@cs.uchicago.edu

Aiman Fang
Department of Computer Science
University of Chicago
aimanf@cs.uchicago.edu

MS65

Local Recovery at Extreme Scales with FenixLR: Implementation and Evaluation

This talk explores how local recovery can be used for certain classes of applications to reduce overheads due to resilience. Specifically, it is centered on the development of programming support and scalable runtime mechanisms to enable online and transparent local recovery under the FenixLR framework. We integrate these mechanisms with the S3D combustion simulation, and experimentally demonstrate (using Titan) the ability to tolerate high failure rates (every 5 seconds) at scales up to 262144 cores.

Marc Gamell
NSF Cloud and Autonomic Computing Center, Rutgers Discovery
Informatics Institute, Rutgers University
mgamell@cac.rutgers.edu

Keita Teranishi, Michael Heroux
Sandia National Laboratories
knteran@sandia.gov, maherou@sandia.gov

Manish Parashar
Department of Computer Science
Rutgers University
parashar@rutgers.edu

MS65
Spare Node Substitution of Failed Node

Spare node substitution for a failed node is not a new idea, but its behavior is not well-studied so far. Spare node substitution may preserve the communication pattern and the number of nodes of a user program. Thus, it can drastically reduce the effort of fault handling at the user program. However, actual communication topology at physical network cannot be the same and this can result in the severe degradation of communication performance. In this talk, recent research result focusing on this topic will be presented.

Atsushi Hori, Kazumi Yoshinaga, Yutaka Ishikawa
RIKEN Advanced Institute for Computational Science
ahori@riken.jp, kazumi.yoshinaga@riken.jp, yutaka.ishikawa@riken.jp

MS65
Will Burst Buffers Save Checkpoint/Restart?

Exascale systems are projected to fail frequently and parallel file system performance is not expected to match the increase in computational speed. Burst buffers can potentially solve this problem by smoothing out bursty I/O requests, and reducing load on parallel file systems. In this talk, we will explore the tradeoffs of using burst buffers as part of storage systems, and how fault tolerance libraries can exploit burst buffers to provide increased performance to user applications.

Kathryn Mohror
Lawrence Livermore National Laboratory
mohror1@llnl.gov

MS66
In Situ Processing Overview and Relevance to the Hpc Community

Abstract not available

E. Wes Bethel
Lawrence Berkeley National Laboratory
ewbethel@lbl.gov

MS66
Overview of Contemporary In Situ Infrastructure Tools and Architectures

Abstract not available

Julien Jomier
Kitware
julien.jomier@kitware.com

Patrick O'Leary
Kitware, Inc., USA
patrick.oleary@kitware.com

MS66
In Situ Algorithms and Their Application to Science Challenges

Abstract not available

Venkatram Vishwanath
Argonne Leadership Computing Facility, Argonne National Lab.
venkatv@mcs.anl.gov

MS66
Complex In Situ and In Transit Workflows

Abstract not available

Matthew Wolf
Georgia Institute of Technology
mwolf@cc.gatech.edu

MS67
Addressing the Future Needs of Uintah on Post-Petascale Systems

The need to solve ever-more complex science and engineering problems on computer architectures that are rapidly evolving poses considerable challenges for modern simulation frameworks and packages. While the decomposition of software into a programming model that generates a task-graph for execution by a runtime system make portability and performance possible, it also imposes considerable demands on that runtime system. The challenge of meeting those demands for near-term and longer term systems in the context of the Uintah framework is considered in key areas such as support for data structures, tasks on heterogeneous architectures, performance portability, power management and designing for resilience by the use of replication based upon adaptive mesh refinement processes is addressed.

Martin Berzins
Scientific Computing and Imaging Institute
University of Utah
mb@sci.utah.edu

MS67
Progress on High Performance Programming Frameworks for Numerical Simulations

This talk reports the progress of high performance numerical simulation frameworks in China. These frameworks are designed to separate parallel algorithm design from parallel programming, which eases challenging application development even with the rapid growth of supercomputers. The domain-specific software architectures and components are introduced on top of general parallel programming systems, namely MPI and OpenMP. A variety of multi-petaflops applications are developed on top of these frameworks, illustrating the success of our approaches.

Zeyao Mo
Laboratory of Computational Physics, IAPCM
P.O. Box 8009, Beijing 100088, P.R. China
zeyao_mo@iapcm.ac.cn

MS67
Can We Really Taskify the World? Challenges for Task-Based Runtime Systems Designers

To Fully tap into the potential of heterogeneous manycore

machines, the use of runtime systems capable of dynamically scheduling tasks over the pool of underlying computing resources has become increasingly popular. Such runtime systems expect applications to generate a graph of tasks of sufficient "width" so as to keep every processing unit busy. However, not every application can exhibit enough task-based parallelism to occupy the tremendous number of processing units of upcoming supercomputers. Nor can they generate tasks of appropriate granularity for CPUs and accelerators. In this talk, I will briefly present the main features provided by the StarPU task-based runtime system, and I will go through the major challenges that runtime systems' designers have to face in the upcoming years.

Raymond Namyst
University of Bordeaux 1, Inria
raymond.namyst@labri.fr

MS67

Exascale Software

Exascale software requries innovation on many levels, such as scalability to beyond a million cores, energy awareness, and algorithm based resilience. Future supercomputers will only deliver their full potential when applications, mathematical models, and numerical algorithms are adapted to ever more complex architectures in a co-design process that is based on the systematic engineering for performance in all stages of the the computational science pipeline.

Hans-Joachim Bungartz
Technische Universität München, Department of Informatics
Chair of Scientific Computing in Computer Science
bungartz@in.tum.de

Ulrich J. Ruede
University of Erlangen-Nuremberg
Department of Computer Science (Simulation)
ulrich.ruede@fau.de

Philipp Neumann
Technische Universität München, Germany
Scientific Computing in Computer Science
neumanph@in.tum.de

MS68

The GraphBLAS Effort: Kernels, API, and Parallel Implementations

I will first give an overview of the GraphBLAS effort for standardizing graph algorithmic building blocks in the language of linear algebra, including its scope, the set of kernels involved and the proposed API. I will then report on recent results in developing high-performance implementations of some of the most important GraphBLAS kernels such as sparse matrix-matrix multiplication and sparse matrix-vector multiplication in clusters of multicore processors.

Aydin Buluc
Lawrence Berkeley National Laboratory
abuluc@lbl.gov

MS68

Parallel Tools to Reconstruct and Analyze Large Climate Networks

Over the last decade, the techniques of complex network analysis have found application in climate research. Here the observation locations serve as nodes and edges (links) are based on statistical measures of similarity, e.g. a correlation coefficient, between pairwise time series of climate variables at these different locations. Many studies in the field focused on investigating correlation patterns in the atmospheric surface temperature and teleconnections as well as the synchronization between climate variability patterns. In most studies above so called interaction networks were used. Up to now, the data sets (both observational and model based) used to reconstruct and analyze climate networks have been relatively small due to computational limitations. As a consequence, coarse-resolution observational and model data have been used with a focus only on large-scale properties of the climate system. This system is, however, known for its multi-scale interactions and hence one would like to explore the interaction of processes over the different scales. Data are available through high-resolution ocean/atmosphere/climate model simulations but they lead to networks with at least $10\tilde{5}$ nodes and hence imposing a compelling need for new algorithms and parallel computational tools that can enable high-performance graph reconstruction and analysis. In this talk we introduce the application of complex graph tools in climate research as well as our complete toolbox ParGraph designed for parallel computing platforms, which is capable of the preprocessing of large number of climate time series and the calculation of pairwise statistical measures, leading to the construction of massive node climate networks. In addition, ParGraph is provided with a set of high-performance network analyzing algorithms designed for symmetric multiprocessing machines (SMPs).

Hisham Ihshaish
UWE Bristol
Faculty of Environment & Technology
hisham.ihshaish@uwe.ac.uk

Alexis Tantet
IMAU
University of Utrecht
a.j.j.tantet@uu.nl

Johan Dijkzeul
VORtech BV
Delft, The Netherlands
johan.dijkzeul@vortech.nl

Henk A Dijkstra
Department of Physics and Astronomy, Utrecht University
The Netherlands
h.a.dijkstra@uu.nl

MS68

Generating and Traversing Large Graphs in External-Memory.

Large graphs arise naturally in many real world applications. The actual performance of simple RAM model algorithms for traversing these graphs (stored in external memory) deviates significantly from their linear or near-linear predicted performance because of the large number of I/Os they incur. In order to alleviate the I/O bottleneck, many external memory graph traversal algorithms have been designed with provable worst-case guarantees. In the talk I

highlight some techniques used in the design and engineering of such algorithms and survey the state-of-the-art in I/O-efficient graph traversal algorithms. I will also report on recent work concerning the generation of massive scale free networks under resource constraints.

Ulrich Meyer
Goethe University Frankfurt am Main
umeyer@cs.uni-frankfurt.de

MS68

(Generating And) Analyzing Large Networks with NetworKit

Complex networks are equally attractive and challenging targets for data mining. Novel algorithmic solutions, including parallelization, are required to handle data sets containing billions of connections. NetworKit is a growing open-source software package for the analysis of such large complex networks on shared-memory parallel computers. It is designed to deliver the best of two worlds: the performance of C++ algorithms and data structures (with parallelism by OpenMP) and a convenient Python frontend for interactive workflows. The current feature set includes not only various fundamental analytics kernels and a collection of basic graph generators. Several novel algorithms for scalable network analysis have been realized with NetworKit as well recently. Of those, we will present network profiling, fast algorithms for computing centrality values (importance of nodes/edges) in different scenarios, and the first generator of hyperbolic random graphs that takes subquadratic time.

Henning Meyerhenke
Karlsruhe Institute of Technology (KIT)
Institute of Theoretical Informatics
henning.meyerhenke@kit.edu

Elisabetta Bergamini
Institute of Theoretical Informatics
Karlsruhe Institute of Technology (KIT)
elisabetta.bergamini@kit.edu

Moritz von Looz, Roman Prutkin, Christian L. Staudt
Karlsruhe Institute of Technology (KIT)
Institute of Theoretical Informatics
moritz.looz-corswarem@kit.edu, roman.prutkin@kit.edu, christian.staudt@kit.edu

MS69

Advancing Algorithms to Increase Performance of Electronic Structure Simulations on Many-Core Architectures

Effective utilization of many-core architectures for large-scale electronic structure calculations requires fast, computationally scalable and efficient computational tools utilizing novel algorithmic approaches on the latest hardware technologies. In this presentation we will discuss results, lessons learned as well as pitfalls resulting from our algorithmic efforts and code optimizations to transform the NWChem computational chemistry software to utilize systems with Intel based multicore architectures connected by fast communication networks.

Wibe A. De Jong
Lawrence Berkeley National Laboratory
wadejong@lbl.gov

MS69

Real-Valued Technique for Solving Linear Systems Arising in Contour Integral Eigensolver

Contour integral eigenvalue solvers such as the Sakurai-Sugiura method and the FEAST algorithm have been developed in the last decade. The main advantage of the contour integral solvers is their inherent parallelism. However they require linear system solving with complex matrices even if the matrices of the problem are real. In this study, we propose a real-valued technique for solving linear systems arising in the contour integral methods, which only involves factorizations of real matrices.

Yasunori Futamura
University of Tsukuba
futamura@cs.tsukuba.ac.jp

Tetsuya Sakurai
Department of Computer Science
University of Tsukuba
sakurai@cs.tsukuba.ac.jp

MS69

Parallel Eigensolvers for Density Functional Theory

Density functional theory (DFT) aims to solve the Schrödinger equation by modelling electronic correlation as a function of density. Its relatively modest $O(N^3)$ scaling makes it the standard method in electronic structure computation for condensed phases containing up to thousands of atoms. Computationally, its bottleneck is the partial diagonalisation of a Hamiltonian operator, which is usually not formed explicitely. Using the example of the Abinit code, I will discuss the challenges involved in scaling plane-wave DFT computations to petascale supercomputers, and show how the implementation of a new method based on Chebyshev filtering results in good parallel behaviour up to tens of thousands of processors. I will also discuss some open problems in the numerical analysis of eigensolvers and extrapolation methods used to accelerate the convergence of fixed point iterations.

Antoine Levitt
Ecole des Ponts ParisTech
antoine.levitt@inria.fr

MS69

A Parallel Algorithm for Approximating the Single Particle Density Matrix for O(N) Electronics Structure Calculations

We present two strategies for approximating the entries of the single particle density matrix that are needed for $O(N)$ solution of the DFT equations. These strategies are based on the use of diagonalization or a Chebyshev polynomial approximation to compute the matrix exponential function using a principal submatrix of the global Hamiltonian matrix. We discuss pros and cons of each strategy within the context of developing a parallel and scalable algorithm for electronics structure calculations.

Daniel Osei-Kuffuor
Lawrence Livermore National Laboratory
oseikuffuor1@llnl.gov

Jean-Luc Fattebert
Lawrence Livermore National Lab.
fattebert1@llnl.gov

MS70

Resilience in Distributed Task-Based Runtimes (2)

Abstract not available

George Bosilca
University of Tennessee - Knoxville
bosilca@icl.utk.edu

MS70

Developing Task-Based Applications with Pycompss

COMP Superscalar (COMPSs) is a task-based programming model which aims to ease the development of parallel applications for distributed infrastructures, such as Clusters, Grids and Clouds, by exploiting the inherent parallelism of applications at execution time. In particular, it is able to handle Python applications through its Python binding (PyCOMPSs), which provides an API and decorators that can be used in order to adapt and manage sequential python applications, hiding the complexities of parallelization.

Javier Conejero
Barcelona Supercomputing Centre
Spain
francisco.conejero@bsc.es

MS70

DuctTeip: A Framework for Distributed Hierarchical Task-Based Parallel Computing

We present a dependency-aware task parallel programming framework for hybrid distributed-shared memory parallelization. Both data and tasks in the framework are hierarchically organized in order to exploit the underlying heterogeneous hardware architecture. A data centric view based on a versioning system is employed. By executing an application code in simulation mode, the choice of data distribution and blocking scheme can be tuned a priori. The performance of the framework is evaluated against other similar efforts.

Elisabeth Larsson
Uppsala University, Sweden
Department of Information Technology
Elisabeth.Larsson@it.uu.se

Afshin Zafari
Uppsala University
Dept. of Information Technology
afshin.zafari@it.uu.se

MS70

Heterogeneous Computing with GPUs and FPGAs Through the OpenMP Task-Parallel Model

The task-based programming model has become one the prominent models for parallelizing software on homogeneous shared-memory architectures. However, because Dennard's scaling no longer holds, we have to leave homogeneity behind and move towards a more heterogeneous (possibly distributed) solution. In this talk, we will show how to extend the OpenMP task-model to support distributed computing through GPUs and how the model further can be used to drive hardware generation for FPGAs.

Artur Podobas
KTH Royal Institute of Technology
podobas@kth.se

MS71

PFLOTRAN: Practical Application of High Performance Computing to Subsurface Simulation

As subsurface simulation becomes more complex with the addition of physical and chemical processes, careful software engineering must be employed to preserve parallel performance. This presentation will discuss a recent refactor of PFLOTRANs high-level process model coupling that has facilitated the integration of new processes (e.g. geophysics, spent nuclear fuel degradation/dissolution) while maintaining scalability. Examples of PFLOTRAN and high performance computing being applied to simulate practical, real-world subsurface problem scenarios will be included.

Glenn Hammond
Sandia National Laboratories
gehammo@sandia.gov

MS71

Parallel Performance of a Reservoir Simulator in the Context of Unstable Waterflooding

We consider the parallel performance of a reservoir simulator in the context of water injection in viscous oils. Such problems usually require very fine grids and high-order schemes in order to capture the complex fingering patterns induced by viscous instabilities. In this context parallel performance is very important and the scalability of the linear solver becomes critical. We compare different strategies for the Algebraic Multigrid method among which a new hybrid aggregation/point-wise coarsening strategy.

Romain De Loubens
Total
romain.de-loubens@total.com

Pascal Henon
TOTAL E&P
pascal.henon@total.com

Pavel Jiranek
CDSP
pavel.jiranek@gmail.com

MS71

Ongoing Development and HPSC Applications with the ParFlow Hydrologic Model

ParFlow is a three-dimensional variably saturated groundwater-surface water flow code that is currently applied at space and time scales up to continents (Europe, USA) and decades. This is made possible by the efficient parallelism and, e.g. an advanced octree space-partitioning algorithm implemented in ParFlow. Ongoing advancements that will be illustrated with subsurface applications include software infrastructure for coupled energy transport; integration of the p4est meshing library; and integration of ParFlow in TerrSysMP for terrestrial systems

modeling.

Stefan J. Kollet
Meteorological Institute
University of Bonn
s.kollet@fz-juelich.de

Klaus Goergen
Jülich Supercomputing Centre
Research Centre Jülich
k.goergen@fz-juelich.de

Carsten Burstedde
Universität Bonn
burstedde@ins.uni-bonn.de

Jose Fonseca
Institute for Numerical Simulation
Bonn University
fonseca.ins.uni-bonn.de

Fabian Gasper
Agrosphere Research Center Jülich
f.gasper@fz-juelich.de

Reed M. Maxwell
Department of Geology and Geological Engineering and IGWMC
Colorado School of Mines
rmaxwell@mines.edu

Prabhakar Shresta, Mauro Sulis
Meterological Institute
Bonn University
pshrestha@uni-bonn.de, msulis@uni-bonn.de

MS71

Hybrid Dimensional Multiphase Compositional Darcy Flows in Fractured Porous Media and Parallel Implementation in Code ComPASS

We consider the Hybrid dimensional multiphase compositional Darcy flows in fractured porous media. The model accounts for the coupling of the mass balance of each component with the pore volume conservation and the thermodynamical equilibrium and dynamically manages phase appearance and disappearance. We implement this model in the framework of code ComPASS (Computing Parallel Architecture to Speed up Simulation), which focuses on parallel high performance simulation (MPI).

Feng Xing
University of Nice
Inria
feng.xing@unice.fr

Roland Masson
University of Nice Sophia Antipolis
roland.masson@unice.fr

MS72

Title Not Available

Abstract not available

Donna Calhoun
Boise State University
calhoun@math.boisestate.edu

MS72

Title Not Available

Abstract not available

Thierry Coupez
MINES ParisTech
thierry.coupez@ec-nantes.fr

MS72

Bb-Amr and Numerical Entropy Production

Abstract not available

Frédéric Golay
Université de Toulon, France
frederic.golay@univ-tln.fr

MS72

Title Not Available

Abstract not available

Pierre Kestener
maisson de la simulation
pierre.kestener@cea.fr

MS73

How Will Application Developers Navigate the Cost/Peformance Trade-Offs of New Multi-Level Memory Architectures?

Data movement and memory bandwidth are becoming the limiting performance factors in many applications. Hybrid memory cubes (HMC), high-bandwidth memory (HBM), and cheap capacity NVARM memories are now becoming available. Performance must balance cost since high bandwidth memory will be far too expensive at high capacity and NVRAM will not provide needed bandwidth, requiring both technologies to be leveraged. Managing complex memory hierarchies will prove a serious programming burden for domain science. Here we present a simulation framework and profiling tools for analyzing memory access profiles to enable runtime tools and algorithm refactoring to fully exploit new memory architectures.

Arun Rodrigues
Sandia National Laboratories
afrodri@sandia.gov

MS73

Performance-Energy Trade-Offs in Reconfigurable Optical Data-Center Networks Using Silicon Photonics

Optical circuits can provide large bandwidth improvements relative to electrical-switched networks in HPC systems. Optical circuit setup, however, can have high latencies. Application-guided, latency-avoiding circuit management techniques for dynamic reconfiguration have shown potential for better achieving the full throughput of photonics-based networks. We extend these ideas to dynamically reconfiguring topologies like dragonfly to match application communication patterns. We further explore where

runtime analysis tools can automatically manage circuits without changing application code.

Sebastien Rumley, Keren Bergman
Columbia University
sr3061@columbia.edu, bergman@ee.columbia.edu

MS73

Helping Applications Reason About Next-Generation Architectures Through Abstract Machine Models

To achieve performance under increasing energy constraints, next-generation architectures will undergo fundamental changes to include more heterogeneous processors, network-on chip topologies, or even multi-level memories. These architectures will inevitably change how programs are written. Detailed specifications do not enable community engagement or design space exploration with application developers. Here we present abstract machine models (AMMs) and corresponding proxy architectures that provide useful starting points for reasoning about applications on extreme-scale hardware. Where appropriate, architecture details are either ignored, abstracted, or fully exposed in the AMMs. We survey the programming implications of different AMMs and illusrate their usefulness in co-designing applications.

John Shalf, David Donofrio
Lawrence Berkeley National Laboratory
jshalf@lbl.gov, ddonofrio@lbl.gov

MS73

When Is Energy Not Equal to Time? Understanding Application Energy Scaling with EAudit

Under future power caps, improving energy efficiency will become a critical component of sustaining improvements in time to solution. However, due to fundamental differences between how energy and time scale with parallelism, time is not a good proxy for the distribution of energy in an application. This talk covers energy analogs to Amdahls law to understand weak/strong scaling of application energy and its use in eAudit a tool for energy debugging of applications.

Sudhakar Yalamanchili
The School of Electrical and Computer Engineering
Georgia Institute of Technology
sudha@ece.gatech.edu

Eric Anger
Georgia Institute of Technology
eanger@gatech.edu

MS74

A Distributed-Memory Advection-Diffusion Solver with Dynamic Adaptive Trees

We present a numerical method to solve the Advection-Diffusion problem. In our approach, we solve the advection equation by using the semi-Lagrangian method. The solution of the diffusion equation is computed by applying an integral transform: a convolution with the fundamental solution of the modified Laplace PDE. For time-marching, the Stiffly-Stable method is used. These methods are implemented on top of a dynamic spatially-adaptive distributed-memory parallelized Chebyshev octree as our data structure. To address the load-imbalance problem for distributed-memory Lagrangian schemes, we introduce a novel and robust partitioning scheme.

Arash Bakhtiari
Technische Universitaet Muenchen
bakhtiar@in.tum.de

Dhairya Malhotra
Institute of Computational Engineering and Sciences
The University of Texas at Austin
dmalhotra@ices.utexas.edu

George Biros
University of Texas at Austin
biros@ices.utexas.edu

MS74

A Volume Integral Equation Solver for Elliptic PDEs in Complex Geometries

We present a new framework for solving volume integral equation (VIE) formulations of Poisson, Stokes and Helmholtz equations in complex geometries. Our method is based on a high-order discretization of the data on adaptive octrees and uses a kernel independent fast multipole method along with precomputed quadratures for evaluating volume integrals efficiently. We will present new performance optimizations and demonstrate scalability of our method on distributed memory platforms with thousands of compute nodes.

Dhairya Malhotra
Institute of Computational Engineering and Sciences
The University of Texas at Austin
dmalhotra@ices.utexas.edu

George Biros
University of Texas at Austin
biros@ices.utexas.edu

MS74

Efficient Parallel Simulation of Flows, Particles, and Structures

The efficient simulation of flow scenarios involving the transport of particles and the interaction with other physical phenomena such as electrostatics or elastic structures usually requires high grid resolutions due to accuracy requirements and different spatial/temporal scales. We present several application cases for the efficient use of octree-like adaptive computational grids in large existing simulation environments. Such grids are an ideal tool in this context due to their minimal memory requirements and flexible adaptivity.

Miriam Mehl
Universität Stuttgart
miriam.mehl@ipvs.uni-stuttgart.de

Michael Lahnert
University of Stuttgart
michael.lahnert@ipvs.uni-stuttgart.de

Benjamin Uekermann
TU Munich

uekerman@in.tum.de

MS74

Communication-Reducing and Communication-Avoiding Techniques for the Particle-in-Dual-Tree Storage Scheme

We study particle-in-cell with suprathermal particles driven by a PDE. The particles do not interact but some jump over cells. It is a HPC worst-case for its low arithmetic intensity and re-sorting which makes it latency- and bandwidth-sensitive. Three ideas help: A tree grammar predicts movements and allows us to skip communication. Run-length encoding for integer data along domain boundaries reduces the memory footprint. A hierarchical floating point representation reduces the byte count further.

Tobias Weinzierl
Durham University
tobias.weinzierl@durham.ac.uk

MS75

Dune - A Parallel Pde Framework

The Distributed and Unified Numerics Environment (DUNE, www.dune-project.de) is a software framework for the numerical solution of partial differential equations with grid-based methods started by several groups in Germany in 2002. Its main goals are combining (1) flexibility on the one hand with (2) efficiency on the other hand while allowing also the (3) reuse of existing finite element software. A central aspect of DUNE is its abstract mesh interface allowing the formulation of generic algorithms operating on a large variety of different types of meshes (any-d, simplicial, cuboid, hierarchical, conforming and non-conforming, different refinement rules, ...). Within the framework of the german priority programme 1648 'Software for exascale computing' DUNE is currently prepared for future exa-scale hardware. The talk will highlight the progress in exploiting parallelism on all different levels (MPI, threading, SIMD) and present results on high-order discontinuous Galerkin methods.

Christian Engwer
INAM, University of Münster
christian.engwer@uni-muenster.de

Steffen Müthing
Heidelberg University
steffen.muething@iwr.uni-heidelberg.de

Peter Bastian
Interdisciplinary Center for Scientific Computing
University of Heidelberg
peter.bastian@iwr.uni-heidelberg.de

MS75

PHG: A Framework for Parallel Adaptive Finite Element Method

PHG (Parallel Hierarchical Grid) is a general framework for developing parallel adaptive finite element method applications. The key feature of PHG includes: bisection based conforming adaptive tetrahedral meshes, various finite element bases support and hp adaptivity, finite element code automatic generation, etc. We also introduce a few applications based on PHG, including earthquake simulation, calculation of electronic structure of materials and simulation of fluid-structure interaction.

Deng Lin
Laboratory of Scientific and Engineering Computing,
Chinese Academy of Sciences, Beijing, China
lind@lsec.cc.ac.cn

Wei Leng
Laboratory of Scientific and Engineering Computing
Chinese Academy of Sciences, Beijing, China
lengweee@gmail.com

Tao Cui
LSEC, the Institute of Computational Mathematics
Chinese Academy of Sciences
tcui@lsec.cc.ac.cn

MS75

Cello: An Extreme Scale Amr Framework, and Enzo-P, An Application for Astrophysics and Cosmology Built on Cello

We present Cello, an extremely scalable AMR software infrastructure for post-petascale architectures, and Enzo-P, a multiphysics application for astrophysics and cosmology simulations built on Cello. Cello implements a forest-of-octrees spatial mesh on top of the CHARM++ parallel objects framework developed at UIUC. Leaf nodes of the octrees are blocks of fixed size, and are represented as chares in a hierarchical chare array. Cello supports particle and field classes and operators for the construction of the explicit finite-difference/volume methods, N-body methods, and sparse linear system solvers called from Enzo-P. We present initial performance and scaling results on a large cosmological test problem.

Michael L. Norman
University of California at San Diego
mlnorman@ucsd.edu

MS75

Peta-scale Simulations with the HPC Software Framework Walberla

We present the HPC software framework waLBerla that has demonstrated excellent parallel performance on different peta-scale supercomputers for simulations based on the lattice Boltzmann as well as the phase-field method. The presentation focuses on the distributed core data structures of the framework that enable block adaptive mesh refinement. We will discuss the performance-driven development process and the continuous integration pipeline that guarantees compatibility with a large number of different compilers and systems through fully automatic testing.

Florian Schornbaum
University of Erlangen-Nuremberg
florian.schornbaum@informatik.uni-erlangen.de

Martin Bauer
University of Erlangen-Nuremberg
Department of Computer Science
martin.bauer@fau.de

Christian Godenschwager, Ulrich Rüde
University of Erlangen-Nuremberg

christian.godenschwager@fau.de, ulrich.ruede@fau.de

MS76

Communication Avoiding Ilu Preconditioner

In this talk, we present a new parallel preconditioner called CA-ILU. It performs ILU(k) factorization in parallel and is applied as preconditioner without communication. Block Jacobi and Restrictive Additive Schwartz preconditioners are compared to this new one. Numerical results show that CA-ILU(0) outperforms BJacobi and is close to RAS in term of iterations whereas CA-ILU, with complete LU on each block, does for both of them but uses lot of memory. Thus CA-ILU(k) takes advantages of both in term of memory consumption and iterations.

Sebastien Cayrols
INRIA
sebastien.cayrols@inria.fr

Laura Grigori
INRIA
France
Laura.Grigori@inria.fr

MS76

Parallel Approximate Graph Matching

We study approximate matching heuristics for undirected graphs by extending our recent work which focused on bipartite graphs. The approach uses a judiciously sampled subgraph of the input graph and search for a maximum matching in it. The approach is highly parallelizable and we are currently working on the theoretical approximation guarantee. It can be used in practice to avoid worst cases while constructing an initial matching. Our experiments demonstrate parallel speed-ups and show that the approximation guarantees do not deteriorate with the increased degree of parallelization.

Fanny Dufossé
Inria Lille, Nord Europe, 59650,
Villeneuve dAscq, France
fanny.dufosse@gmail.com

Kamer Kaya
Sabanci University,
Faculty of Engineering and Natural Sciences
kaya@sabanciuniv.edu

Bora Ucar
LIP, ENS-Lyon, France.
Bora.Ucar@ens-lyon.fr

MS76

Diamond Sampling for Approximate Maximum All-pairs Dot-product (MAD) Search

Given two sets of vectors A and B, our problem is to find the top-t dot products among all possible pairs. This is a fundamental mathematical problem that appears in numerous data applications involving similarity search, link prediction, and collaborative filtering. We propose a sampling-based approach that avoids direct computation of all dot products. We select diamonds (i.e., four-cycles) from the weighted tripartite representation of A and B. The probability of selecting a diamond corresponding to pair (i,j) is proportional to the square of the product of the corresponding vectors, amplifying the focus on the largest-magnitude entries. Experimental results indicate that diamond sampling is orders of magnitude faster than direct computation and requires far fewer samples than any competing approach. We also apply diamond sampling to the special case of maximum inner product search, and get significantly better results than the state-of-the-art hashing methods.

Grey Ballard
Sandia National Laboratories
gmballa@sandia.gov

C. Seshadhri
Sandia National Labs
scomand@sandia.gov

Tamara G. Kolda
Sandia National Laboratories
tgkolda@sandia.gov

Ali Pinar
Sandia National Labs
apinar@sandia.gov

MS76

Scalable and Efficient Algorithms for Analysis of Massive, Streaming Graphs

Graph-structured data in network security, social networks, finance, and other applications not only are massive but also under continual evolution. The changes often are scattered across the graph, permitting novel parallel and incremental analysis algorithms. We discuss analysis algorithms for streaming graph data to maintain both local and global metrics with low latency and high efficiency.

Jason Riedy
Georgia Institute of Technology
School of Computational Science and Engineering
jason.riedy@cc.gatech.edu

David A. Bader
Georgia Institute of Technology
bader@cc.gatech.edu

MS77

A Block Low-Rank Multithreaded Factorization for Dense BEM Operators.

We consider a Block Low-Rank multithreaded LDU factorization for BEM electromagetics applications. In this context, both the computation and accumulation of updates appear to be critical to ensure good performance. In both kernels of the factorization, each matrix may be dense or low rank. First is a matrix-matrix multiply, where outer-orthogonality of the representation is key. Second is a matrix-matrix merge, where there are many strategies in a multithreaded environment.

Julie Anton
LSTC
julie@lstc.com

Cleve Ashcraft
Livermore Software Technology
Corporation
cleve@lstc.com

Clément Weisbecker
LSTC
clement@lstc.com

MS77

Complexity and Performance of Algorithmic Variants of Block Low-Rank Multifrontal Solvers

Matrices coming from elliptic Partial Differential Equations (PDEs) have been shown to have a low-rank property: well defined off-diagonal blocks of their Schur complements can be approximated by low-rank products and this property can be efficiently exploited in multifrontal solvers to provide a substantial reduction of their complexity. Among the possible low-rank formats, the Block Low-Rank format (BLR) is easy to use in a general purpose multifrontal solver and has been shown to provide significant gains compared to full-rank on practical applications. However, unlike hierarchical formats, such as $\mathcal{H}$ and HSS, its theoretical complexity was unknown. In this talk, we present several variants of the BLR multifrontal factorization, depending on the strategies used to perform the updates in the frontal matrices and on the approaches to handle numerical pivoting. We extend the theoretical work done on hierarchical matrices in order to compute the theoretical complexity of each BLR variant. We then provide an experimental study with numerical results to support our complexity bounds. Finally, we compare and analyze the performance of each variant in a parallel setting.

Patrick Amestoy
INP(ENSEEIHT)-IRIT
amestoy@enseeiht.fr

Alfredo Buttari
CNRS-IRIT-Université de Toulouse, France
alfredo.buttari@enseeiht.fr

Jean-Yves L'Excellent
INRIA-LIP-ENS Lyon
jean-yves.l.excellent@ens-lyon.fr

Theo Mary
Universite de Toulouse, INPT ENSEEIHT
theo.mary@enseeiht.fr

MS77

A Comparison of Parallel Rank-Structured Solvers

We present a direct comparison of different rank-structured solvers: STRUMPACK (that implements HSS representations, by the authors), MUMPS (Block Low-Rank representations, Amestoy et al.), and HLIBPro (H-arithmetic, Kriemann et al.). We report on results for large dense and sparse problems from academic and industrial applications, in a distributed-memory context.

Francois-Henry Rouet
Lawrence Berkeley National Laboratory
fhrouet@lbl.gov

Patrick Amestoy
ENINPT-IRIT, Université of Toulouse, France
amestoy@n7.fr or patrick.amestoy@n7.fr

Cleve Ashcraft
Livermore Software Technology Corporation
cleve@lstc.com

Alfredo Buttari
CNRS-IRIT-Université de Toulouse, France
alfredo.buttari@enseeiht.fr

Pieter Ghysels
Lawrence Berkeley National Laboratory
Computational Research Division
pghysels@lbl.gov

Jean-Yves L'Excellent
INRIA-LIP-ENS Lyon
jean-yves.l.excellent@ens-lyon.fr

Xiaoye Sherry Li
Computational Research Division
Lawrence Berkeley National Laboratory
xsli@lbl.gov

Theo Mary
Universite de Toulouse, INPT ENSEEIHT
theo.mary@enseeiht.fr

Clément Weisbecker, Clément Weisbecker
LSTC
clement@lstc.com, clement@lstc.com

MS77

Hierarchical Matrix Arithmetics on Gpus

GPU implementations of hierarchical matrix operations are challenging as they require irregular data structures and multilevel access patterns. Careful design, data layout, and tuning are needed to allow the effective mapping of the rank-structured representations and associated operations to manycore architectures. In this talk, we describe the design of GPU-hosted, tree-based, representations and algorithms that can compress dense matrices into a hierarchical format and perform linear algebra operations directly on the compressed representations in a highly efficient manner.

Wajih Halim Boukaram
KAUST
wajihhalim.boukaram@kaust.edu.sa

Hatem Ltaief
King Abdullah University of Science & Technology (KAUST)
hatem.Ltaief@kaust.edu.sa

George M Turkiyyah
American University of Beirut
george.turkiyyah@kaust.edu.sa

David E. Keyes
KAUST
david.keyes@kaust.edu.sa

MS78

Traversing a BLR Factorization Task Dag Based on a Fan-All Wraparound Map

We consider the factorization of a BLR matrix partitioned by block rows and columns. The atomic factor, solve and MMM tasks merge with submatrices to create a task DAG,

which we embed in a 3D Cartesian mesh. Then the atomic tasks are scheduled to processors block wise using a 3D wraparound map. Using p processors, both message counts and volume grow as $p^{1/3}$.

Julie Anton
LSTC
julie@lstc.com

Cleve Ashcraft
Livermore Software Technology Corporation
cleve@lstc.com

Clément Weisbecker
LSTC
clement@lstc.com

MS78

Scheduling Sparse Symmetric Fan-Both Cholesky Factorization

In this work, we study various scheduling approaches for Sparse Cholesky factorization. We review different algorithms based on a distributed memory implementation of the Fan-Both algorithm. We also study different communication mechanisms and dynamic scheduling techniques.

Mathias Jacquelin
Lawrence Berkeley National Lab
mjacquelin@lbl.gov

Esmond G. Ng, Yili Zheng
Lawrence Berkeley National Laboratory
egng@lbl.gov, yzheng@lbl.gov

Katherine Yelick
Lawrence Berkeley National Laboratory
University of California at Berkeley
yelick@eecs.berkeley.edu

MS78

Optimizing Memory Usage When Scheduling Tree-Shaped Task Graphs of Sparse Multifrontal Solvers

We study the problem of scheduling task trees with large data, which arise in the multifrontal method of sparse matrix factorization. Under limited memory, changing the traversal of the tree impacts the needed amount of memory. Optimal algorithms exist in the sequential case. On parallel systems, the objective is to minimize the processing time with bounded memory, which is much more complex. Heuristics from the literature may be improved to better use the available memory.

Guillaume Aupy
ENS Lyon
guillaume.aupy@ens-lyon.fr

Lionel Eyraud-Dubois
INRIA
University of Bordeaux
lionel.eyraud-dubois@labri.fr

Loris Marchal
CNRS
France
loris.marchal@ens-lyon.fr

Frédéric Vivien
LIP, ENS Lyon, France
frederic.vivien@ens-lyon.fr

Oliver Sinnen
University of Auckland
o.sinnen@auckland.ac.nz

MS78

Impact of Blocking Strategies for Sparse Direct Solvers on Top of Generic Runtimes

Among the preprocessing steps of a sparse direct solver, reordering and block symbolic factorization are two major steps to reach a suitable granularity for BLAS kernels efficiency and runtime management. In this talk, we present a reordering strategy to increase off-diagonal block sizes. It enhances BLAS kernels and allows to handle larger tasks, reducing runtime overhead. Finally, we will comment the resulting gain in the PASTIX solver implemented over STARPU and PARSEC.

Gregoire Pichon
INRIA
gregoire.pichon@inria.fr

Mathieu Faverge
IPB Bordeaux - Labri - Inria
mathieu.faverge@inria.fr

Pierre Ramet
Bordeaux University - INRIA - LaBRI
Bordeaux University
701868

Jean Roman
INRIA
Jean.Roman@inria.fr

MS79

Silent Data Corruption in BLAS Kernels

We summarize different types of fault tolerant parallel algorithms for selected BLAS kernels and compare their behavior under the influence of silent data corruption (bit-flips). In some scenarios, randomized gossip-based algorithms tend to exhibit a much higher degree of resilience than deterministic methods. Nevertheless, existing gossip-based algorithms cannot handle all bit-flips in their data structures, either. We present novel algorithms, both deterministic as well as randomized, with improved resilience against silent data corruption.

Wilfried N. Gansterer
Department of Computer Science
University of Vienna
wilfried.gansterer@univie.ac.at

MS79

Fault Tolerant Sparse Iterative Solvers

Silent data corruption is expected to become an increasing problem as total data set and memory sizes grow in future supercomputers. Application-based fault tolerance (ABFT) techniques are one promising approach to mitigating this problem. In this presentation we will show some promising early results from our work to design a

fault-tolerant sparse iterative solver based on the Conjugate Gradient (CG) method.

Simon McIntosh-Smith
University of Bristol
cssnmis@bristol.ac.uk

MS79

Algorithm-Based Fault Tolerance of the Sakurai-Sugiura Method

The Sakurai-Sugiura (SS) method is a contour integral eigensolver, which have been developed and carefully analyzed in the last decade. The SS method computes eigensubspace by solutions of independent linear systems and this characteristic leads coarse-grained parallel computation. In this presentation we show the algorithm-based fault tolerance of the method and verify the error resilience by several numerical experiments.

Tetsuya Sakurai, Akira Imakura
Department of Computer Science
University of Tsukuba
sakurai@cs.tsukuba.ac.jp, imakura@cs.tsukuba.ac.jp

Yasunori Futamura
University of Tsukuba
futamura@cs.tsukuba.ac.jp

MS79

Soft Errors in PCG: Detection and Correction

Soft errors that are not detected by hardware mechanisms may affect convergence of iterative methods. In this presentation, we discuss the possibility of a self-correction technique for classical Preconditioned Conjugate Gradient (PCG) algorithm in case of soft error in the SpMV calculation. To be able to correct the convergence behavior, the soft errors should be detected immediately. Checksum based detection methods are good candidates for it. But their defect, arising from round-off effects, decreases their reliability. Hence we combine this technique with the well-known results on deviation between the true and computed residuals in finite precision calculation. The combination of these two approaches gives us an opportunity to detect the soft error before the algorithm get stuck in an unrecoverable state. After detection, the convergence behavior can be corrected by applying residual replacement approach. The discussion is illustrated with several fault injection scenarios on different matrices.

Emmanuel Agullo
INRIA
emmanuel.agullo@inria.fr

Siegfried Cools
Antwerp University
wim.vanroose@uantwerpen.be

Luc Giraud
Inria
Joint Inria-CERFACS lab on HPC
luc.giraud@inria.fr

Wim Vanroose
Antwerp University
wim.vanroose@uantwerpen.be

Fatih Yetkin
Inria
emrullah-fatih.yetkin@inria.fr

MS80

Code Verification in Cfd

Abstract not available

Daniel Chauveheid
CEA DAM DIF
Bruyères-le-Chtel , 91297 Arpajon Cedex, France
daniel.chauveheid@cea.fr

MS80

Performance Study of Industrial Hydrocodes: Modeling, Analysis and Aftermaths on Innovative Computational Methods

This talk is the high-performance numerics companion talk given by Mathieu Peybernes in this mini-symposium. From conclusions on the performance of a legacy staggered Lagrange+remap hydrocode, we set a list of requirements for a new numerical solver. It has been possible to derive a second-order Eulerian multimaterial finite volume scheme with collocated variables and geometrical elements defined only from the reference (Eulerian) grid. The resulting scheme shows both SIMD property and strong scalability.

Florian De Vuyst
Ecole Normale Supérieure Cachan
devuyst@cmla.ens-cachan.fr

MS80

Checking the Floating Point Accuracy of Scientific Code Through the Monte Carlo Arithmetic

Exascale supercomputers will provide the computational power required to perform very large scale and high resolution digital simulations. One of the main exascale challenges is to control their numerical uncertainties. Indeed, the floating-point arithmetic is only an approximation of exact arithmetic The aim of this talk is to present a new tool called which can be used for the automatic assessment of the numerical accuracy of large scale digital simulations by using the Monte-Carlo Arithmetic.

Christophe Denis
CMLA - ENS Cachan - University Paris Saclay
UMR 8536
christophe.denis@cmla.ens-cachan.fr

MS80

Code Modernization by Mean of Proto-Applications

With the shift in HPC system paradigm, applications have to be re-designed to efficiently address large numbers of manycores. It requires deep code modification that cannot be practicaly applied at full-application scale. We propose proto-applications as proxy for code-modernization. Our first example features unstructured mesh computation using task-based parallization. The second demonstrates load-balancing for combustion simulation. Both proto-applications are open-source and can serve to the

developement of genuine HPC application.

Eric Petit
PRiSM - UVSQ
45 avenue des Etats-Unis, 78035 Versailles cedex
eric.petit@uvsq.fr

MS81

On the Performance and Productivity of Abstract C++ Programming Models?

Porting production-scale scientific applications to emerging next-generation computing platforms is a significant challenge. Application developers must not only ensure strong performance, they must also ensure code is portable across the wide diversity of architectures available. In order to tackle the significant scale of these codes the community must also take into account developer productivity. In this talk we present results from a recent DOE study which evaluated directive solutions against our Kokkos C++ programming model.

Simon D. Hammond
Scalable Computer Architectures
Sandia National Laboratories
sdhammo@sandia.gov

Christian Trott, H. Carter Edwards, Jeanine Cook, Dennis Dinge
Sandia National Laboratories
crtrott@sandia.gov, hcedwar@sandia.gov, jeacook@sandia.gov, dcdinge@sandia.gov

Mike Glass
Sandia National Labs
mwglass@sandia.gov

Robert J. Hoekstra, Paul Lin, Mahesh Rajan, Courtenay T. Vaughan
Sandia National Laboratories
rjhoeks@sandia.gov, ptlin@sandia.gov, mrajan@sandia.gov, ctvaugh@sandia.gov

MS81

Demonstrating Advances in Proxy Applications Through Performance Gains and/or Performance Portable Abstractions: CoMD and Kripke with RAJA

RAJA is a programming model and C++ implementation thereof which exposes fine-grained parallelism in both loop and task-based fashions. In this talk, we present successful results from a recent study wherein we applied RAJA to enhance the performance of two codes outside of its original design space.

Richard Hornung, Olga Pearce
Lawrence Livermore National Lab
hornung@llnl.gov, pearce8@llnl.gov

Adam Kunen, Jeff Keasler, Holger Jones
Lawrence Livermore National Laboratory
kunen1@llnl.gov, keasler1@llnl.gov, jones19@llnl.gov

Rob Neely
Lawrence Livermore National Lab
neely4@llnl.gov

Todd Gamblin
Lawrence Livermore National Laboratory
tgamblin@llnl.gov

MS81

Future Architecture Influenced Algorithm Changes and Portability Evaluation Using BLAST

BLAST is an arbitrary high-order finite element code for shock hydrodynamics, originally developed for CPUs using MPI. In this talk, we describe the algorithmic and structural changes we made to BLAST to efficiently support multiple emerging architectures, in particular GPU and Xeon Phi. We analyze the key differences between our CPU, GPU, and Xeon Phi versions and our efforts to combine them into a single performance portable implementation to simplify future maintenance.

Christopher Earl
Lawrence Livermore National Laboratory
earl2@llnl.gov

Barna Bihari
Lawrence Livermore National Lab
bihari1@llnl.gov

Veselin Dobrev, Ian Karlin
Lawrence Livermore National Laboratory
dobrev1@llnl.gov, karlin1@llnl.gov

Tzanio V. Kolev
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
tzanio@llnl.gov

Robert Rieben, Vladimir Tomov
Lawrence Livermore National Laboratory
rieben1@llnl.gov, tomov2@llnl.gov

MS81

Kokkos and Legion Implementations of the SNAP Proxy Application

SNAP is a proxy application which simulates the computational motion of a neutral particle transport code. In this work, we have adapted parts of SNAP separately; we have re-implemented the iterative shell of SNAP in the task-model runtime Legion, showing an improvement to the original schedule, and we have created multiple Kokkos implementations of the computational kernel of SNAP, displaying similar performance to the native Fortran.

Geoff Womeldorff, Ben Bergen, Joshua Payne
Los Alamos National Laboratory
womeld@lanl.gov, bergen@lanl.gov, payne@lanl.gov

MS82

Recent Developments on Forests of Linear Trees

Forest-of-octrees AMR is an approach that offers both geometric flexibility and parallel scalability and is being used in various finite element codes and libraries. The core functionality is defined by refinement, size balance, the assembly of ghost layers, and a parallel topology iteration. More general applications, such as semi-Lagrangian and patch-based methods, require additional AMR functionalities. Recently, a tentative tetrahedral variant has emerged.

This talk will summarize the latest theory and applications.

Carsten Burstedde, Johannes Holke
Universität Bonn
burstedde@ins.uni-bonn.de, holke@ins.uni-bonn.de

MS82

Parallel Level-Set Algorithms on Tree Grids

We present results from parallelizing the semi-Lagrangian and Reinitialization algorithms for the level-set method on adaptive Quadtree (2D) and Octree (3D) grids. These algorithms are directly implemented on top of the p4est library which implements high quality parallel refinement/coarsening and partitioning operations. Initial scaling results indicate good scalability up to 4096 cores, our current allocation limit on the Stampede supercomputer. Some applications of our algorithms to modeling the solidification phenomenon and crystal growth are discussed.

Mohammad Mirzadeh
Dept. Mechanical Engineering
UCSB
m.mirzadeh@engineering.ucsb.edu

Arthur Guittet
UCSB
arthur.guittet@gmail.com

Carsten Burstedde
Universität Bonn
burstedde@ins.uni-bonn.de

Frederic G. Gibou
UC Santa Barbara
fgibou@engineering.ucsb.edu

MS82

Exploiting Octree Meshes with a High Order Discontinuous Galerkin Method

High-order discontinuous Galerkin methods are especially attractive for parallel computation due to their high locality. However, the sheer size of modern computing systems requires also the scalable treatment of mesh information. This need drives the design of our Octree mesh implementation TreElM. We will present the concept of Ateles, our modal discontinuous Galerkin solver based on TreElM. In this context, the focus of our presentation will be the high-order representation of geometries.

Sabine P. Roller, Harald Klimach, Jens Zudrop
Universitaet Siegen
sabine.roller@uni-siegen.de, harald.klimach@uni-siegen.de, jens.zudrop@uni-siegen.de

MS82

A Sharp Numerical Method for the Simulation of Diphasic Flows

We present a numerical method for modeling incompressible non-miscible fluids, in which both the interface and the interfacial continuity equations are treated in a sharp manner. Our approach is implemented on adaptive Octree/Quadtree grids, which are highly effective in capturing different length scales. By using a modified pressure correction projection method, we were able to alleviate the standard time step restriction incurred by capillary forces. The solver is validated numerically in two and three spatial dimensions and challenging numerical examples are presented to illustrate its capabilities.

Maxime Theillard
Mechanical and Aerospace Engineering
UCSD
mtheillard@eng.ucsd.edu

Frederic Gibou
UCSB
fgibou@engr.ucsb.edu

David Saintillan
UCSD
dsaintillan@ucsd.edu

Arthur Guittet
UCSB
arthur.guittet@gmail.com

MS83

Optimal Domain Decomposition Strategies for Parallel Geometric Multigrid

A high level cache-directed analysis is used to derive a flexible cache-oblivious heuristic that prunes the process topology search space to obtain families of high-performing domain decompositions for structured 3-D domains. This concept is used to minimize cache-misses in smoothing phases and is combined with process-to-hardware mapping for a subset of active processes at the coarsest correction phase. Our conclusion discusses the impact on domain decomposition strategies resulting from current, and future, multi-core architectures.

Gaurav Saxena
University of Leeds
scgs@leeds.ac.uk

Peter K. Jimack, Mark A. Walkley
School of Computing
University of Leeds
p.k.jimack@leeds.ac.uk, m.a.walkley@leeds.ac.uk

MS83

Design and Implementation of Adaptive Spmv Library for Multicore and Manycore Architecture

Sparse matrix vector multiplication (SpMV) is an important kernel in both traditional high performance computing and emerging data-intensive applications. Previous SpMV libraries are optimized by either application-specific or architecture-specific approach, and they're complicated to be used in real applications. In this work we develop an auto-tuning system (SMATER) to bridge the gap between specific optimizations and general-purpose use. SMATER provides programmers a unified interface based on the compressed sparse row (CSR) format by implicitly choosing the best format and the fastest implementation for any input sparse matrix in the runtime. SMATER leverages a machine learning model and a retargetable backend library to quickly predict the best combination. Performance parameters are extracted from 2386 matrices in UF sparse matrix collection. The experiments show that SMATER achieves the impressive performance while being portable on the state-of-the-art x86 multi-core processors, NVIDIA

GPU and Intel Xeon Phi accelerators. Compared with Intel MKL library, SMATER runs faster by more than 3 times. We demonstrate its adaptivity in an algebraic multigrid solver from hypre library and report above 20% performance improvement.

Guangming Tan
Institute of Computing Technology
Chinese Academy of Sciences
tgm@ict.ac.cn

MS83

Towards Efficient Data Communications for Frameworks on Post-petascale Numa-Multicore Systems

Performance of halo-exchange communications is essential to the scalability of parallel programming frameworks. In this talk, we present our effort on optimizing halo-exchange communications for the JASMIN framework on post-peta scale NUMA multicore systems. We propose a parallelism-exposing communication abstraction, design a scalable scheduling algorithm, and utilize widely available hardware features. Our effort improves typical halo-exchange communication performance by up to 78%. This in turn improves the overall performance of JEMS-FDTD application by 48%.

Zhang Yang
Institute of Applied Physics and Computational Mathematics
China Academy of Engineering Physics
yang_zhang@iapcm.ac.cn

Aiqing Zhang
Institute of Applied Physics and Computational Mathematics
zhang_aiqing@iapcm.ac.cn

MS84

Novo-G#: Strong Scaling Using Fpga-Centric Clusters with Direct and Programmable Interconnects

FPGA-centric clusters with direct and programmable interconnects address fundamental limits on HPC performance by maximizing computational density, minimizing power density, removing the bottleneck between processing and communication, and facilitating intelligent and application-aware processing in the network itself. For example, knowing the routing pattern a priori often enables congestion-free communication. We describe our 100-node publically available cluster, Novo-G#, its software and IP infrastructure, and strong scaling of communication-bound applications such as long time-scale Molecular Dynamics.

Martin Herbordt
Boston University
herbordt@bu.edu

MS84

Benchmarking FPGAs with Opencl

We evaluate the performance of OpenCL benchmarks, including the Rodinia benchmark suite, on FPGAs using Altera's OpenCL SDK. We first evaluate the performance of the original, thread-parallelized OpenCL kernels targeted for GPUs on FPGAs. We study how those kernels can be significantly optimized by exploiting FPGA-specific architectural features such as pipeline parallelism and sliding windows. Our results using an Altera Stratix V FPGA demonstrates significantly higher power efficiency in comparison to GPUs.

Naoya Maruyama
RIKEN Advanced Institute for Computational Science
nmaruyama@riken.jp

Hamid Zohouri
Tokyo Institute of Technology
zohouri.h.aa@m.titech.ac.jp

Aaron Smith
Microsoft Research
aaron.smith@microsoft.com

Motohiko Matsuda
RIKEN Advanced Institute for Computational Science
m-matsuda@riken.jp

Satoshi Matsuoka
Tokyo Insitute of Technology
matsu@is.titech.ac.jp

MS84

Architectural Possibilities of Fpga-Based High-Performance Custom Computing

Recent advancement of FPGA technology provides various possibilities to high-performance custom computing architectures. As a solution to higher scalability and efficiency than the present massively-parallel HPC systems, we focus on data-centric computing architectures which is different from conventional control-centric ones. This talk introduces their microarchitecture to system level approaches for latency-tolerant and decentralized computing, which include bandwidth compression techniques and an overlay architecture for system-wide data-flow computing.

Kentaro Sano
Tohoku University
kentah@caero.mech.tohoku.ac.jp

MS84

Are Datacenters Ready for Reconfigurable Logic?

At Microsoft we have designed and deployed a composable, reconfigurable fabric to accelerate portions of large-scale software services beyond what commodity server designs provide. This talk will describe a prototype deployment of this fabric on a bed of 1,632 servers, and discuss its efficacy in accelerating the Bing web search engine.

Aaron Smith
Microsoft Research
aaron.smith@microsoft.com

MS85

Parallel Performance of the Inverse FMM and Application to Mesh Deformation for Fluid-Structure Problems

In this talk, we present the inverse fast multipole method as an approximate fast direct solver for dense linear systems that performs extremely well as a preconditioner in an iterative solver. The computational cost scales linearly with the problem size. The method uses low-rank approx-

imations to represent well-separated interactions; this is done in a multi-level fashion. Its parallel performance is assessed, and applications related to mesh deformation using radial basis function interpolation are discussed.

Pieter Coulier
Stanford
pcoulier@stanford.edu

Chao Chen
Stanford University
cchen10@stanford.edu

Eric F. Darve
Stanford University
Mechanical Engineering Department
darve@stanford.edu

MS85

Fast Updating of Skeletonization-Based Hierarchical Factorizations

We present a method for updating certain hierarchical factorizations for solving linear systems arising from the discretization of physical problems (e.g., integral equations or PDEs) in 2D. Given a factorization, we can locally perturb the geometry or coefficients and update the factors to reflect this change with complexity poly-logarithmic in the total number of unknowns. We apply our method to the recursive skeletonization factorization and hierarchical interpolative factorization and demonstrate results on some examples.

Victor Minden
Stanford
vminden@stanford.edu

Anil Damle, Ken Ho, Lexing Ying
Stanford University
damle@stanford.edu, klho@stanford.edu, lexing@stanford.edu

MS85

Exploiting H-Matrices in Sparse Direct Solvers

In this talk, we describe a preliminary fast direct solver using HODLR library to compress large blocks appearing in the symbolic structure of the PaStiX sparse direct solver. We will present our general strategy before analyzing the practical gains in terms of memory and floating point operations with respect to a theoretical study of the problem. Finally, we will discuss ways to enhance the overall performance of the solver.

Gregoire Pichon
INRIA
gregoire.pichon@inria.fr

Eric F. Darve
Stanford University
Mechanical Engineering Department
darve@stanford.edu

Mathieu Faverge
IPB Bordeaux - Labri - Inria
mathieu.faverge@inria.fr

Pierre Ramet
Bordeaux University - INRIA - LaBRI
Bordeaux University
701868

Jean Roman
INRIA
Jean.Roman@inria.fr

MS85

Algebraic and Geometric Low-Rank Approximations

We perform a direct comparison between a geometric (FMM) and algebraic (HSS) matrix-vector multiplication of compressed matrices for Laplace, Helmholtz, and Stokes kernels. Both methods are implemented with distributed memory, thread, and SIMD parallelism, and the runs are performed on large Cray systems such as Edison (XC30, 133,824 cores) at NERSC and Shaheen2 (XC40, 197,568 cores) at KAUST. We compare the absolute execution time of both codes, along with their parallel scalability and memory consumption.

Rio Yokota
Tokyo Institute of Technology
rioyokota@gsic.titech.ac.jp

François-Henry Rouet, Xiaoye S. Li
Lawrence Berkeley National Laboratory
fhrouet@lbl.gov, xsli@lbl.gov

David E. Keyes
KAUST
david.keyes@kaust.edu.sa

MS86

Architecture-Agnostic Numerics in GridTools

The rapid evolution of the HPC hardrware has produced major changes in programming techniques and algorithms employed to solve several classical problems, such as the compressible Euler equations. Overcoming the lag between numerical schemes and their efficient implementations on specific hardware requires cross-disciplinary competencies and is often challenging. We propose a set of libraries called GridTools which will provide domain scientists with techniques to specify, discretize and solve (explicitly) PDE problems taking full advantage of modern hardware architectures. We present GridTools and several prototype PDE solvers.

Paolo Crosetto
Ecole Politechnique Federale de Lausanne
paolo.crosetto@epfl.ch

MS86

Generating High Performance Finite Element Kernels Using Optimality Criteria.

We tackle the problem of automatically generating optimal finite element integration routines given a high level specification of arbitrary multilinear forms. Optimality is defined in terms of floating point operations given a memory bound. We show an approach to explore the transformation space of loops and expressions typical of integration routines and discuss the conditions for optimality. Extensive experimentation validates the approach, showing systematic performance improvements over a number

of alternative code generation systems.

Fabio Luporini
Department of Computing
Imperial College London
f.luporini12@imperial.ac.uk

David Ham
Department of Mathematics and Department of Computing
Imperial College London
david.ham@imperial.ac.uk

Paul Kelly
Imperial College London
p.kelly@imperial.ac.uk

MS86

Source-to-Source Translation for Code Validation and Acceleration of the Icon Atmospheric General Circulation Model

The ICON atmospheric model port to accelerators with OpenACC directives illustrates that these are easily inserted and can yield performance similar to CUDA programming, but introduce subtelties which are difficult to optimize and debug. Thus we propose a source-to-source translator to generate accelerator directives based on an abstract description of the accelerator data and kernels, as well as validation code for intermediate results. In this presentation we report our findings and outline possible extensions.

William Sawyer
Swiss Centre of Scientific Computing (CSCS)
Manno, Switzerland
wsawyer@cscs.ch

MS86

Kernel Optimization for High-Order Dynamic Rupture Simulation - Flexibility Vs. Performance

SeisSol simulates seismic wave propagation and dynamic rupture. Its innermost kernels consist of small matrix multiplications (SMMs) where all sparsity patterns are known a-priori. The SMMs are generated in an offline process and make heavy use of vectorization in order to maximize performance. We present our recent efforts to further optimize and automate the code generator. Particularly, we determine the optimal matrix multiplication order, automatically exploit zero blocks, and enable block sparse SMMs.

Carsten Uphoff
TU Munich
uphoff@in.tum.de

Michael Bader
Technische Universität München
bader@in.tum.de

MS87

Bosilca ¡TBA¿

Abstract not available

George Bosilca
University of Tennessee - Knoxville
bosilca@icl.utk.edu

MS87

From System-Level Checkpointing to User-Level Failure Mitigation in MPI-3

In this presentation, I will talk about transparent, system-level fault tolerance and present some (old) results that show the overhead of such mechanisms. Then I will talk about how the user can handle fault tolerance directly within the parallel application, using the (old, deprecated) interface FT-MPI and using a current effort made by a group of the MPI forum in the MPI-3 standard. Finally, I will present how the middleware can handle failures and resist beyond them.

Camille Coti
Université Paris 13
camille.coti@lipn.univ-paris13.fr

MS87

Reinit: A Simple and Scalable Fault-Tolerance Model for MPI Applications

In this talk, we present a global-exception programing model to tolerate failures in MPI applications. The model allows a clean way to reinitialize MPI state after failures and a simple mechanism to clean up application and libraries state via a stack of error handlers. We present an implementation of the model in MVAPICH (a widely used MPI library) and Slurm (an open-source workload manager), and show an initial performance evaluation in HPC benchmarks.

Ignacio Laguna
Lawrence Livermore National Laboratory
ilaguna@llnl.gov

MS87

FA-MPI: Fault-Aware and QoS-Oriented Scalable Message Passing

A Fault-Aware Approach for non-blocking MPI-3.x APIs is presented (FA-MPI), enabled by weak collective transactions This novel model localizes faults, provides tunable fault-free overhead, allows for multiple kinds of faults, enables hierarchical recovery, and is data-parallel relevant. Fault modeling of underlying networks is mentioned as are proposed additions to MPI-4 fully to support resilient parallel programs with basic QoS. Performance and scalability results of prototype FA-MPI are given, keyed to compact applications (e.g., MiniFE, LULESH).

Anthony Skjellum
Auburn University
skjellum@auburn.edu

MS88

Topology-Aware Performance Modeling of Parallel SpMVs

Parallel SpMVs are a key component of many linear solvers. Increasingly large HPC systems yield costly communication; specifically large messages sent long distances. A performance model can associate an approximate cost with the various messages communicated during each SpMV. The common alpha-beta model can be improved through the use of topology-aware methods, allowing for a distance

parameter to be added to the cost of each message, as well as a rough prediction of network contention.

Amanda Bienz
University of Illinois at Urbana-Champaign
bienz2@illinois.edu

Luke Olson
Department of Computer Science
University of Illinois at Urbana-Champaign
lukeo@illinois.edu

William D. Gropp
University of Illinois at Urbana-Champaign
wgropp@illinois.edu

MS88

Cere: Llvm Based Codelet Extractor and Replayer for Piecewise Benchmarking and Optimization

Codelet Extractor and REplayer (CERE) is an LLVM-based framework that finds and extracts hotspots from an application as isolated fragments of code. Codelets can be modified, compiled, run, and measured independently from the original application. Through performance signature clustering, CERE extracts a minimal but representative codelet set from applications, which can significantly reduce the cost of benchmarking and iterative optimization. Codelets have proved successful in autotuning target architecture, compiler optimization or amount of parallelism.

Pablo De Oliveira Castro
PRiSM & ECR - UVSQ
45 avenue des Etats-Unis, 78035 Versailles cedex
pablo.oliveira@uvsq.fr

MS88

Performance Modeling of a Compressible Hydrodynamics Solver on Multicore Cpus

Performance study of a reference Lagrange-Remap algorithm for compressible gas dynamics obtained using the Roofline model and the Execution Cache Memory (ECM) model has been performed. As a result, we are able to predict the whole performance on modern processors with a mean error of 6.5%, which is actually very accurate in this context. The results of these analyses also gave us highlights for designing a new full hydrodynamics solver (Lagrange+flux) with expected better performance.

Mathieu Peybernes
CEA Saclay, 91191 Gif-sur-Yvette cedex, France
mathieu.peybernes@cea.fr

MS88

Using Uncertainties in High Level Performance Prediction

For rapid evaluation of future hardware, we are using high level methodology based on a first approximation of the application behavior, and on several hardware specifications (frequency, efficiency, bandwidth, nb of cores, etc..) that are usually known only a few months before the official launch. The introduction of uncertainties on these parameters is then allowing us to get some tolerance on the prediction results.

Philippe Thierry
2 rue de paris, 92190 Meudon, France
philippe.thierry@intel.com

PP1

Improving Scalability of 2D Sparse Matrix Partitioning Models Via Reducing Latency Cost

Two-dimensional (2D) sparse matrix partitioning models offer better opportunities in reducing multiple communication cost metrics compared to one-dimensional partitioning models. This work makes use of a two-phase partitioning methodology to address three important communication cost metrics: (i) total message volume, (ii) total message count, and (iii) maximum message volume. The first phase only considers the message volume metric, while the second phase considers the other two metrics. We show that the scalability of the existing 2D models can significantly be improved with our methodology.

Cevdet Aykanat, Oguz Selvitopi
Bilkent University
aykanat@cs.bilkent.edu.tr, reha@cs.bilkent.edu.tr

PP1

Hybrid Parallel Simulation of Helicopter Rotor Dynamics

We present a software for the hybrid parallel calculation of rotor forces and control angles using a coupled simulation of rotor dynamics and rotor wake. A big part of the computation time is required to simulate the airflow in the wake using a vortex method. For node-level parallelization, we use OpenACC on GPUs and OpenMP on CPUs, and MPI for parallelization between compute nodes. We focus on the coupling scheme and discuss the performance obtained.

Melven Roehrig-Zoellner, Achim Basermann
German Aerospace Center (DLR)
Simulation and Software Technology
Melven.Roehrig-Zoellner@DLR.de,
Achim.Basermann@dlr.de

Johannes Hofmann
German Aerospace Center (DLR)
Institute of Flight Systems
johannes.hofmann@dlr.de

PP1

Poisson Solver for large-scale Simulations with one Fourier Diagonalizable Direction

We present a Poisson solver for problems with one Fourier diagonalizable direction, such as airfoils. The diagonalization decomposes the original 3D-system into a set of uncoupled 2D-subsystems. Our method focuses on optimizing memory allocations and transactions by considering redundancies on such subsystems. It also takes advantage of the uniformity through the periodic direction for its vectorization, and automatically manages the workload distribution and the preconditioner choice. Altogether constitutes a highly efficient and scalable Poisson solver.

Ricard Borrell
Universitat Politécnica de Catalunya (Barcelona Tech)
ricardb@cttc.upc.edu

Guillermo Oyarzun, Jorje Chiva, Oriol Lehmkuhl, Ivette

Rodríguez, Assensi Oliva
Universitat Politècnica de Catalunya (BarcelonaTech)
guillermo@cttc.upc.edu, jordic@cttc.upc.edu,
oriol@cttc.upc.edu, ivette@cttc.upc.edu,
oliva@cttc.upc.edu

PP1

High Performance Algorithms and Software for Seismic Simulations

The understanding of earthquake dynamics is greatly supported by highly resolved coupled simulations of the rupture process and seismic wave propagation. This grand challenge of seismic modeling requires an immense amount of supercomputing resources, thus optimal utilization by software is imperative. Driven by recent hardware developments the increasing demand for parallelism and data locality often requires replacing major software parts to bring efficient numerics and machine utilization closely together. In this poster we present a new computational core for the seismic simulation package SeisSol. For minimal time-to-solution the new core is designed to maximize value and throughput of the floating point operations performed in the underlying ADER discontinuous Galerkin discretization method. Included are auto-tuned sparse and dense matrix kernels, hybrid parallelization from many-core nodes up to machine-size and a novel high performance clustered local time stepping scheme. The presented computational core reduces time-to-solution of SeisSol by several factors and scales beyond 1 million cores. At machine-size the new core enabled a landmark simulation of the 1992 Landers earthquake. For the first time this simulation allowed the analysis of the complex rupture behavior resulting from the non-linear interaction of frictional sliding and seismic wave propagation at high geometric complexity.

Alexander Breuer
University of California, San Diego
anbreuer@ucsd.edu

Alexander Heinecke
Parallel Computing Laboratory
Intel Corporation, Santa Clara, CA, USA
alexander.heinecke@intel.com

Michael Bader
Technische Universität München
bader@in.tum.de

PP1

Large Scale Parallel Computations in R Through Elemental

R-Elemental is a package for the statistical software R that enables distributed computations through the parallel dense linear algebra library Elemental. R-Elemental offers to the users a straightforward transition from single node to compute clusters, preserving native R syntax and methods; at the same time, it delivers Elementals full performance without any overhead. We discuss the interface design and present several examples common in statistical analysis.

Rodrigo Canales
RWTH Aachen University
rodrigo.canales@rwth-aachen.de

Elmar Peise, Paolo Bientinesi
AICES, RWTH Aachen
peise@aices.rwth-aachen.de, pauldj@aices.rwth-aachen.de

PP1

Parallel Programming Using Hpcmatlab

HPCmatlab is a new, flexible and scalable parallel programming framework for Matlab/Octave applications to scale on distributed memory cluster architectures using the Message Passing Interface (MPI). This includes point-to-point, collective, one sided communication and parallel I/O. It supports shared memory programming within a node using Pthreads. We present scientific application performance results showing the ease of programming for fast prototyping or for large scale numerical simulations.

Mukul Dave
ASU Research Computing
Arizona State University
mukul92@gmail.com

Xinchen Guo
Purdue University
xinchen.guo@asu.edu

Mohamed Sayeed
ASU Research Computing
Arizona State University
msayeed@asu.edu

PP1

Using the Verificarlo Tool to Assess the Floating Point Accuracy of Large Scale Digital Simulations

Exascale supercomputers will provide the computational power required to perform very large scale and high resolution digital simulations. One of the main exascale challenges is to control their numerical uncertainties. Indeed, the floating-point arithmetic is only an approximation of exact arithmetic. This poster will present a new tool called verificarlo which can be used for the automatic assessment of the numerical accuracy of large scale digital simulations by using the Monte-Carlo Arithmetic.

Christophe Denis
CMLA - ENS Cachan - University Paris Saclay
UMR 8536
christophe.denis@cmla.ens-cachan.fr

Pablo De Oliveira Castro
PRiSM & ECR - UVSQ
45 avenue des Etats-Unis, 78035 Versailles cedex
pablo.oliveira@uvsq.fr

Eric Petit
PRISM - UVSQ
eric.petit@prism.uvsq.fr

PP1

A Parallel Fast Sweeping Method on Nonuniform Grids

We present a novel parallel fast sweeping method on nonuniform grids. By constructing a graph of the discretization, we are able to determine the causally independent updates. Each of these independent steps is then computed in parallel. We demonstrate the method on tree based grids and triangulated surfaces. This talk will briefly

describe the algorithm, show an example application, and present parallel scaling results.

Miles L. Detrixhe
University of California Santa Barbara
mdetrixhe@engineering.ucsb.edu

PP1

Chase: A Chebyshev Accellerated Subspace Iteration Eigensolver

ChASE is an iterative eigensolver designed with a rather simple algorithmic structure which rendered easy to parallelize its kernels over virtually all existing range of multi- and many-cores. Such a feature makes ChASE virtually one of the most parallel efficient algorithm for sequences of correlated eigenvalue problems. In particular, ChASE scales almost optimally over a range of processors commensurate to the problem complexity, and it is portable and adaptable to many flavors of parallel architectures.

Edoardo A. Di Napoli
Juelich Supercomputing Centre
e.di.napoli@fz-juelich.de

Mario Berljafa
School of Mathematics
The University of Manchester
m.berljafa@maths.man.ac.uk

Paul Springer
AICES
RWTH Aachen University
paul.springer@rwth-aachen.de

Jan Winkelmann
AICES
RWTH Aachen
winkelmann@aices.rwth-aachen.de

PP1

Evaluation of the Intel Xeon Phi and Nvidia K80 As Accelerators for Two-Dimensional Panel Codes

To predict the properties of fluid flow over a solid geometry is an important engineering problem. In many applications so-called panel methods (or boundary element methods) have become the standard approach to solve the corresponding partial differential equation. Since panel methods in two dimensions are computationally cheap, they are well suited as the inner solver in an optimization algorithm. In this paper we evaluate the performance of the Intel Xeon Phi 7120 and the NVIDIA K80 to accelerate such an optimization algorithm. For that purpose, we have implemented an optimized version of the algorithm on the CPU and Xeon Phi (based on OpenMP, vectorization, and the Intel MKL library) and on the GPU (based on CUDA and the MAGMA library). We present timing results for all codes and discuss the similarities and differences between the three implementations. Overall, we observe a speedup of approximately 2.5 for adding a Intel Xeon Phi 7120 to a dual socket workstation and a speedup between 3.4 and 3.8 for adding a NVIDIA K80 to a dual socket workstation.

Lukas Einkemmer
University of Innsbruck
lukas.einkemmer@uibk.ac.at

PP1

An Impact of Tuning the Kernel of the Structured QR Factorization in the TSQR

In the TSQR algorithm, the so-called structured QR factorization is required for reducing two upper triangular matrix into one and repeated depending on the number of processes in parallel computing. Tuning this part is thus important but not trivial because the input matrix is usually small and has the upper triangular structure. In this poster, we compare several implementations and investigate their impact on the overall performance of the TSQR algorithm.

Takeshi Fukaya
Hokkaido University, RIKEN AICS
fukaya@iic.hokudai.ac.jp

Toshiyuki Imamura
RIKEN Advance Institute for Computational Science
imamura.toshiyuki@riken.jp

PP1

Filters of the Improvement of Multiscale Data from Atomistic Simulations

In multiscale models combining continuum approximations with first-principles based atomistic representations the dominant computational cost is typically the atomistic component of the model. We demonstrate the effectiveness and scalability of a spectral filter for improving the accuracy of noisy data obtained from atomistic simulations. Filtering enables running less expensive atomistic simulations to reach the same overall level of accuracy thus lowering the primary cost of the multiscale model and leading to faster simulations.

David J. Gardner
Lawrence Livermore National Laboratory
gardner48@llnl.gov

Daniel R. Reynolds
Southern Methodist University
Mathematics
reynolds@smu.edu

PP1

Parallel Multigrid and FFT-Type Algorithms for High Order Compact Discretization of the Helmholtz Equation.

In this paper, we analyze parallel implementation of multigrid and FFT-type methods for the solution of a subsurface scattering problem. The 3D Helmholtz equation is discretized by a sixth order compact scheme. The resulting systems of finite-difference equations are solved by different preconditioned Krylov subspace-based methods. The lower order multigrid and FFT-based preconditioners are considered for the efficient implementation of the developed iterative approach. The complexity and scalability of two parallel preconditioning strategies are analyzed.

Yury A. Gryazin
Idaho State University

gryazin@isu.edu

PP1

Computational Modeling and Aerodynamic Analysis of Simplified Automobile Body Using Open Source Cfd (OpenFoam)

Open source softwares have emerged as a cross-disciplinary paradigm for a large spectrum of computing applications including computer aided engineering (CAE). Open Field Operation and Manipulation (OpenFOAM) is an open source computational software capable of solving variety of complex fluid flows. Present research work is focused on investigating the ability of OpenFOAM to simulate external flow around a simplified automobile body. The results are compared with experimental values of commercial CFD packages. These results allow us to gauge the robustness, and adaptability of the OpenFOAM utility.

Javeria Hashmi
National University of Science and technology (NUST)
National University of Sciences and Technology Islamabad
hashmi001@ymail.com

PP1

Performance Evaluation of a Numerical Quadrature Based Eigensolver Using a Real Pole Filter

Numerical quadrature based parallel sparse eigensolvers such as the Sakurai-Sugiura method have been developed in the last decade. Recently, an approach using a numerical quadrature which forms a real pole filter was proposed. This approach potentially reduces the computational and memory cost for solving real symmetric problems. In this presentation, we evaluate the performance of the real pole filter method by comparing it with the often-used complex pole filter method using practical large-scale problems.

Yuto Inoue
Department of Computer Science
University of Tsukuba
yuto@mma.cs.tsukuba.ac.jp

Yasunori Futamura, Tetsuya Sakurai
University of Tsukuba
futamura@cs.tsukuba.ac.jp, sakurai@cs.tskuba.ac.jp

PP1

Synthesis and Verification of BSP Programs

Bulk synchronous parallel (BSP) is an established model for parallel programming which simplifies program conception by avoiding deadlocks and data-races. Nonetheless, there are still a large class of programming errors which plague BSP programs. Our research interest is twofold: correctness by construction where BSP code is automatically generated with a new BSP-automata theory, and automatic verification of existing C programs written with BSPlib using formal methods such as abstract interpretation.

Arvid Jakobsson, Thibaut Tachon
Huawei Technologies France
arvid.jakobsson@huawei.com,
thibaut.tachon@huawei.com

PP1

A Reconstruction Algorithm for Dual Energy Computed Tomography and Its Parallelization

This paper provides a mathematical algorithm to reconstruct the Compton scatter and photo-electronic coefficients using the dual-energy CT system. The proposed imaging method is based on the mean value theorem to handle the non-linear integration coming from the polychromatic energy based CT scan system. We show a numerical simulation result for the validation of the proposed algorithm. In order to reduce the computational time, we use MPI4Py to adopt the parallel processing. We obtain 34 time faster result with 64 cpu cores system.

Kiwan Jeon
National Institute for Mathematical Sciences
Jeonkiwan@nims.re.kr

Sungwhan Kim
Hanbat National University
sungwhan@hanbat.ac.kr

Chi Young Ahn, Taeyoung Ha
National Institute for Mathematical Sciences
chiyoung@nims.re.kr, tha@nims.re.kr

PP1

An Interior-Point Stochastic Approximation Method on Massively Parallel Architectures

The stochastic approximation method is behind the solution of many actively-studied problems in PDE-constrained optimization. Despite its far-reaching applications, there is almost no work on applying stochastic approximation and interior-point optimization, although IPM are particularly efficient in large-scale nonlinear optimization due to their attractive worst-case complexity. We present a massively parallel stochastic IPM method and apply it to stochastic PDE problems such as boundary control and optimization of complex electric power grid systems under uncertainty.

Juraj Kardo
Università della Svizzera italiana
Institute of Computational Science
juraj.kardos@usi.ch

Drosos Kourounis
USI - Universita della Svizzera italiana
Institute of Computational Science
drosos.kourounis@usi.ch

Olaf Schenk
USI - Università della Svizzera italiana
Institute of Computational Science
olaf.schenk@usi.ch

PP1

Nlafet - Parallel Numerical Linear Algebra for Future Extreme-Scale Systems

NLAFET, funded by EU Horizon 2020, is a direct response to the demands for new mathematical and algorithmic approaches for applications on extreme scale systems as identified in the H2020-FETHPC work programme.

The aim is to enable a radical improvement in the performance and scalability of a wide range of real-world applications relying on linear algebra software, by developing novel architecture-aware algorithms and software libraries, and the supporting runtime capabilities to achieve scalable performance and resilience on heterogeneous architectures. The focus is on a critical set of fundamental linear algebra operations including direct and iterative solvers for dense and sparse linear systems of equations and eigenvalue problems.

Bo T. Kågström
Umeå University
Computing Science and HPC2N
bokg@cs.umu.se

Lennart Edblom
Umeå University
Computing Science
edblom@cs.umu.se

Lars Karlsson
Umeå University, Dept. of Computing Science
larsk@cs.umu.se

Dongarra Jack
University of Manchester
School of Mathematics
dongarra@manchester.ac.uk

Iain Duff
Rutherford Appleton Laboratory, Oxfordshire, UK OX11 0QX
and CERFACS, Toulouse, France
iain.duff@stfc.ac.uk

Jonathan Hogg
STFC Rutherford Appleton Lab, UK
jonathan.hogg@stfc.ac.uk

Laura Grigori
INRIA Paris - Rocquencourt
Alpines Group
laura@grigori.inria.fr

PP1

Seigen: Seismic Modelling Through Code Generation

In this poster we present initial performance results for Seigen, a new seismological modelling framework that utilises symbolic computation to model elastic wave propagation using the discontinuous Galerkin finite element method (DG-FEM). Our study is conducted in the context of Firedrake, an automated code generation framework that drastically simplifies in-depth experimentation with multiple problem discretisations, solution methods and low-level performance optimisations.

Michael Lange, Christian Jacobs
Department of Earth Science and Engineering
Imperial College London
michael.lange@imperial.ac.uk, c.t.jacobs@soton.ac.uk

Fabio Luporini
Department of Computing
Imperial College London
f.luporini12@imperial.ac.uk

Gerard J. Gorman
Imperial College London
g.gorman@imperial.ac.uk

PP1

High-Performance Tensor Algebra Package for GPU

Many applications, e.g., high-order FEM simulations, can be expressed through tensors. Contractions by the first index can often be represented using a tensor index reordering followed by a matrix-matrix product, which is a key factor to achieve performance. We present ongoing work of a high-performance package in the MAGMA library to organize tensor contractions, data storage, and batched execution of large number of small tensor contractions. We use automatic code generation techniques to provide an architecture-aware, user-friendly interface.

Ian Masliah
University Paris-Sud
ian.masliah@lri.fr

Ahmad Abdelfattah
University of Tennessee, Knoxville
ahmad@icl.utk.edu

Marc Baboulin
University of Paris-Sud
marc.baboulin@lri.fr

Jack J. Dongarra
Department of Computer Science
The University of Tennessee
dongarra@icl.utk.edu

Joel Falcou
University of Paris-Sud
joel.falcou@lri.fr

Azzam Haidar
Department of Electrical Engineering and Computer Science
University of Tennessee, Knoxville
haidar@icl.utk.edu

Stanimire Tomov
University of Tennessee, Knoxville
tomov@icl.utk.edu

PP1

Multilevel Approaches for FSAI Preconditioning in the Iterative Solution of Linear Systems of Equations

The solution of a sparse linear system of equations with a large SPD matrix is a common and computationally intensive task when considering numerical simulation techniques. This calculation is often done by iterative methods whose performance is largely influenced by the selected preconditioner. In this work, two new preconditioners built upon the dynamic FSAI algorithm in combination to two graph reordering techniques are presented, namely BTF-SAI and DDFSAI.

Victor A. Paludetto Magri, Andrea Franceschini, Carlo Janna, Massimiliano Ferronato
University of Padova

victor.magri@dicea.unipd.it, franc90@dmsa.unipd.it, janna@dmsa.unipd.it, ferronat@dmsa.unipd.it

PP1

Large Scale Computation of Long Range Interactions in Molecular Dynamics Simulations

We consider a physical system of polyelectrolytes in ionic liquids, and perform a molecular dynamics simulation to study the dynamics of electrophoresis. A computational effort of including long range interactions among the particles is made by using a large scale computer cluster. We use an efficient method of distributing computational loads among the nodes, and obtain a new result that is in agreement with experimental findings.

Pyeong Jun Park
School of Liberal Arts and Sciences
Korea National University of Transportation
pjpark@ut.ac.kr

PP1

Accelerated Subdomain Smoothing for Geophysical Stokes Flow

Scalable approaches to solving linear systems arising from geophysical Stokes flow typically involve the use of multigrid methods. Most of the computational effort is thus expended applying smoothing operators, and as such one might naturally ask if this process can be accelerated with the use of increasingly-available per-node coprocessors on modern large-scale computational clusters. Here we discuss several alternate approaches to performing subdomain smoothing on a hybrid CPU-GPU supercomputer, in an attempt to asses the performance gains possible from the use of this new technology. Chebyshev-Jacobi and modern incomplete factorization approaches are considered.

Patrick Sanan
USI - Università della Svizzera italiana
ICS
patrick.sanan@gmail.com

PP1

ExaHyPE - Towards an Exascale Hyperbolic PDE Engine

ExaHyPEs mission is to develop a novel simulation engine for hyperbolic PDE problems at exascale. With our approach - high-order discontinuous Galerkin discretization on dynamically adaptive Cartesian grids - we tackle grand-challenge simulations in Seismology and Astrophysics. We present a first petascale software release bringing together some of ExaHyPE's key ingredients: Grid generation based upon spacetrees in combination with space-filling curves, a novel ADER DG WENO scheme and compute kernels tailored to particular hardware generations.

Angelika Schwarz
Technische Universität München
angelika.schwarz@in.tum.de

Vasco Varduhn
University of Minnesota
varduhn@tum.de

Michael Bader
Technische Universität München
bader@in.tum.de

Michael Dumbser
University of Trento
michael.dumbser@unitn.it

Dominic E. Charrier, Tobias Weinzierl
Durham University
dominic.e.charrier@durham.ac.uk,
tobias.weinzierl@durham.ac.uk

Luciano Rezzolla
Frankfurt Institute of Advanced Studies
rezzolla@itp.uni-frankfurt.de

Alice Gabriel, Heiner Igel
Ludwig-Maximilians-Universität München
alice-agnes.gabriel@geophysik.uni-muenchen.de,
heiner.igel@geophysik.uni-muenchen.de

PP1

Malleable Task-Graph Scheduling with a Practical Speed-Up Model

We focus on the problem of minimizing the makespan of malleable task graphs arising from multifrontal factorizations of sparse matrices. The speedup of each task is assumed to be perfect but bounded. We propose a new heuristic which we prove to be a 2-approximation and which performs better than previous algorithms on simulations. The dataset consists of both synthetic and realistic graphs.

Bertrand Simon
ENS Lyon
bertrand.simon@ens-lyon.fr

Loris Marchal
CNRS
France
loris.marchal@ens-lyon.fr

Oliver Sinnen
University of Auckland
o.sinnen@auckland.ac.nz

Frédéric Vivien
INRIA and ENS Lyon
frederic.vivien@inria.fr

PP1

Overview of Intel Math Kernel Library (Intel MKL) Solutions for Eigenvalue Problems

We outline both existing and prospective Intel Math Kernel Library (Intel MKL) solutions for standard and generalized Hermitian eigenvalue problems. Extended Eigensolver functionality, based on the accelerated subspace iteration FEAST algorithm, is a high-performance solution for obtaining all the eigenvalues and optionally all the associated eigenvectors, within a user-specified interval. To extend the available functionality, a new approach for finding the k largest or smallest eigenvalues is proposed. We envision automatically producing a search interval by a preprocessing step based on classical and recent methods that estimate eigenvalue counts. Computational results are presented.

Irina Sokolova

Novosibirsk State University
Pirogova street 2, Novosibirsk, Russia, 630090
sokolova.mmfd@gmail.com

Peter Tang
Intel Corporation
2200 Mission College Blvd, Santa Clara, CA
peter.tang@intel.com

Alexander Kalinkin
Intel Corporation
alexander.a.kalinkin@intel.com

PP1

Achieving Robustness in Domain Decomposition Methods: Adpative Spectral Coarse Spaces

Domain decomposition methods are a popular way to solve large linear systems. For problems arising from practical applications it is likely that the equations will have highly heterogeneous coefficients. For example a tire is made both of rubber and steel, which are two materials with very different elastic behaviour laws. Many domain decomposition methods do not perform well in this case, specially if the decomposition into subdomains does not accommodate the coefficient variations. For three popular domain decomposition methods (Additive Schwarz, BDD and FETI) we propose a remedy to this problem based on local spectral decompositions. Numerical investigations for the linear elasticity equations will confirm robustness with respect to heterogeneous coefficients, automatic (non regular) partitions into subdomains and nearly incompressible behaviour.

Nicole Spillane
Centre de Mathématiques Appliquées
Ecole Polytechnique
nicole.spillane@cmap.polytechnique.fr

Victorita Dolean
University of Strathclyde
dolean@unice.fr

Patrice Hauret
Michelin
hauret@cmap.polytechnique.fr

Frédéric Nataf
CNRS and UPMC
nataf@ann.jussieu.fr

Clemens Pechstein
University of Linz
clemens.pechstein@oeaw.ac.at

Daniel J. Rixen
TU Munich
rixen@tum.de

Robert Scheichl
University of Bath
r.scheichl@maths.bath.ac.uk

PP1

A Proposal for a Batched Set of Blas.

We propose a Batched set of Basic Linear Algebra Subprograms (Batched BLAS). We focus on multiple independent BLAS operations on small matrices that are grouped together as a single routine. We aim to provide a more efficient and portable library for multi/manycore HPC systems. We achieve 2x speedups and 3x better energy efficiency compared to vendor implementations. We also demonstrate the need for Batched BLAS and its potential impact in multiple application domains.

Pedro Valero-Lara
University of Manchester
pedro.valero-lara@manchester.ac.uk

Jack J. Dongarra
Department of Computer Science
The University of Tennessee
dongarra@icl.utk.edu

Azzam Haidar
Department of Electrical Engineering and Computer Science
University of Tennessee, Knoxville
haidar@icl.utk.edu

Samuel Relton
University of Manchester, UK
samuel.relton@manchester.ac.uk

Stanimire Tomov
University of Tennessee, Knoxville
tomov@icl.utk.edu

Mawussi Zounon
Univeristy of Manchester
mawussi.zounon@manchester.ac.uk

PP1

Spectrum Slicing Solvers for Extreme-Scale Eigenvalue Computations

We present progress on a recently developed new parallel eigensolver package named EVSL (Eigen-Value Slicing Library) for extracting a large number of eigenvalues and the associated eigenvectors of a symmetric matrix. EVSL is based on a divide-and-conquer approach known as spectrum slicing. It combines a thick-restart version of the Lanczos algorithm or a subspace iteration with deflation, and uses a new type of polynomial filters constructed from a least squares approximation of the Dirac delta function to obtain interior eigenvalues. Several computational results are presented emphasizing parallel performance for matrices arising from the density functional theory based electronic structure calculations.

Eugene Vecharynski
Lawrence Berkeley National Laboratory
evecharynski@lbl.gov

Ruipeng Li
Lawrence Livermore National Laboratory
liruipengblue@gmail.com

Yuanzhe Xi
University of Minnesota
yxi@cs.umn.edu

Chao Yang
Lawrence Berkeley National Lab

cyang@lbl.gov

Yousef Saad
Department of Computer Science
University of Minnesota
saad@cs.umn.edu

PP1

Compressed Hierarchical Schur Linear System Solver for 3D Photonic Device Analysis with Gpu Accelerations

Finite-difference frequency-domain analysis of wave optics and photonics requires linear system solver for Helmholtz equation. The linear system can be ill-conditioned when computation domain is large or perfectly-matched layers are used. We propose a direct method named Compressed Hierarchical Schur for time and memory savings, with the GPU acceleration and hardware tuning. These new techniques can efficiently utilize modern high-performance environment and greatly accelerate development of future development of photonic devices and circuits.

Cheng-Han Du
Department of Mathematics, National Taiwan University
b92006@csie.ntu.edu.tw

Weichung Wang
National Taiwan University
Institute of Applied Mathematical Sciences
wwang@ntu.edu.tw

PP1

Parallel Adaptive and Robust Multigrid

Numerical simulation has become one of the major topics in Computational Science. To promote modelling and simulation of complex problems new strategies are needed allowing for the solution of large, complex model systems. After discussing the needs of large-scale simulation we point out basic simulation strategies such as adaptivity, parallelism and multigrid solvers. These strategies are combined in the simulation system UG4 (Unstructured Grids) being presented in the following. We show the performance and efficiency of this strategy in various applications.

Gabriel Wittum
G-CSC
University of Frankfurt, Germany
wittum@techsim.org

PP1

High performance eigenvalue solver for Hubbard model on CPU-GPU hybrid platform

When we solve the smallest eigenvalue and the corresponding eigenvector of the Hamiltonian derived from the Hubbard model, we can understand properties of some interesting phenomenon in the quantum physics. Since the Hamiltonian is a large sparse matrix, we usually utilize an iteration method. In this presentation, we propose tuning strategies for the LOBPCG method for the Hamiltonian on a CPU-GPU hybrid platform in consideration of the property of the model and show its performance.

Susumu Yamada
Japan Atomic Energy Agency
yamada.susumu@jaea.go.jp

Toshiyuki Imamura
RIKEN Advance Institute for Computational Science
imamura.toshiyuki@riken.jp

Masahiko Machida
Japan Atomic Energy Agency
machida.masahiko@jaea.go.jp