

IP1**The Fusion of Supercomputing with Large Scale Data Analytics**

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

William Blake
Cray Inc., USA
bblake@cray.com

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

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

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

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

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

Richard Vuduc
Georgia Institute of Technology
richie@cc.gatech.edu

Jee Whan Choi, Kenneth Czechowski, Marat Dukhan
GA Tech
tba, tba, tba

Aparna Chandramowlishwaran
mit

tba

IP4**Particles, Grids and HPC for Simulations in Fluid Dynamics**

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

Petros Koumoutsakos
Professor of Computational Science
Universitatstrasse 6, ETH Zurich, Switzerland
petros@inf.ethz.ch

IP5**Large-scale GPU Applications for Scientific Computing**

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

Takayuki Aoki
Tokyo Institute of Technology
taoki@gsic.titech.ac.jp

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

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

Jackie Chen
Sandia National Laboratories, USA

jhchen@sandia.gov

IP7

Large-Scale Visual Data Analysis

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

Christopher Johnson
 University of Utah
 Department of Computer Science
 crj@sci.utah.edu

CP1

Parallel Rotor Wake Simulation on Multicore Architectures with Gpgpu Accelerators Using Openacc

The Freewake code from the DLR Institute of Flight Systems simulates the 3d flow around a rotor of a helicopter under active control. It is based on a vortex approach with models for special features of the rotor flow. The method is very well suited for massive parallelism, because the velocity needs to be accumulated independently for each node of the mesh. We will present porting of the existing MPI-parallel implementation to GPGPUs using OpenACC and analyze the performance of hybridly parallelized calculations on multicore CPUs with GPGPU accelerators.

Achim Basermann, Melven Zoellner
 German Aerospace Center (DLR)
 Simulation and Software Technology
 Achim.Basermann@dlr.de, melven.zoellner@dlr.de

Johannes Hofmann
 German Aerospace Center (DLR)
 Institute of Flight Systems
 johannes.hofmann@dlr.de

CP1

Selecting Multiple Order Statistics with a Graphics Processing Unit

We present an algorithm for simultaneously selecting multiple order statistics with a GPU. The algorithm, bucketMultiSelect, significantly reduces computation time compared with sorting the vector on the GPU. For large vectors, bucketMultiSelect returns thousands of order statistics in less time than sorting the vector. For vectors with 2^{28} uniformly distributed doubles, bucketMultiSelect selects the 100 percentile order statistics in approximately 197ms, more than $10\times$ faster than sorting with the GPU optimized radix sort, thrust::sort.

Jeffrey D. Blanchard

Department of Mathematics and Statistics
 Grinnell College
 jeff@math.grinnell.edu

Erik Opavsky, Emircan Uysaler
 Grinnell College
 opavskye@grinnell.edu, uysalere@grinnell.edu

CP1

Cyclic Reduction Type Poisson and Helmholtz Solvers on a GPU

The topics of the presentation are the block cyclic reduction type linear system solvers and how they can be applied to the Poisson and Helmholtz equations on a GPU. The presented implementations are based on a method called radix-q PSCR. A total of three implementations are presented: a simplified radix-2 method, a simplified radix-4 method and a generalized radix-4 method. The presentation will focus on the GPU implementations and the obtained numerical results.

Mirko E. Myllykoski, Tuomo Rossi, Jari Toivanen
 Department of Mathematical Information Technology
 University of Jyväskylä
 mirko.myllykoski@jyu.fi, tuomo.j.rossi@jyu.fi,
 jari.a.toivanen@jyu.fi

CP1

Accelerating Earthquake Simulations on General-Purpose Graphics Processors

This paper describes acceleration of a topologically-realistic numerical simulation of earthquakes, by parallel implementation on Graphics Processing Units (GPUs). The computationally intensive modules include generation of the stress influence matrix (Greens functions) from fault element data, and calculation of stress from the strain vector using matrix vector multiply, during the rupture propagation phase. GPU implementation of these functions results in 45x speedup over the baseline multi-core CPU implementation, for a 30,000-year earthquake simulation.

Monish D. Tandale, Jason Kwan
 Optimal Synthesis Inc.
 monish@optisyn.com, jason@optisyn.com

CP2

Parallel Implementation of An Aggregation-Based Algebraic Multigrid Method

A parallel implementation of an aggregation-based algebraic multigrid method (AGMG) is discussed. The issue involved choosing the optimal strategy to solve the coarsest grid problem is considered. Both conjugate gradient iteration and parallel multifrontal massive direct solver (MUMPS) are applied on coarsest grid solving and the influences on the scalability of AGMG method are compared for several linear systems arising from discretization of 2D or 3D PDEs. For 3D problems, different strategies of domain decomposition are considered and the performance and gained speedup are compared.

Meng-Huo Chen
 Department of Applied Mathematics, University of Washington

mchen01@uw.edu

CP2

Efficient Parallel Adaptive Multi-Grid-Like Solver for CFD Applications

Nowadays, computational fluid dynamics simulations require massive computational resources. In order to efficiently compute physical accurate results an adaptive geometry and/or fluid driven approach is highly beneficial. Unfortunately, a lot of effects influence the accuracy as well as the efficiency. The authors will present an approach for a multi-grid-like concept integrated in a highly efficient data structure in order to boost the parallel efficiency as well as the computational accuracy.

Jérôme Frisch

Chair for Computation in Engineering
Technische Universität München
frisch@tum.de

Ralf-Peter Mundani

TUM, Faculty of Civil Engineering and Geodesie
Chair for Computation in Engineering
mundani@tum.de

CP2

Parallel Algebraic Multigrid and Linear Solvers Suitable for Implicit Particle in Cell Simulations

Unlike explicit particle in cell methods, an implicit particle-in-cell (iPIC3D) method allows time steps at magnetohydrodynamics time scales. However, the implicit formulation requires the solution of two linear systems: a Poisson system related to divergence cleaning, and a system related to a second order formulation of Maxwell equation. Solving these linear systems is the dominant part of the total simulation time. In this talk, we study the scalability and robustness of algebraic multigrid and other linear solvers for both the Poisson and Maxwell solvers for various partitioning and reordering of the domains. We suggest proper choice of partitioning strategies (geometric, or graph partitioning), and parameters related to the linear solvers (coarsening strategies, number of smoothing steps) that helps to reduce the total simulation time.

Pawan Kumar

Katholieke Universiteit Leuven
kumar.work@rediffmail.com

CP2

Performance of Multigrid Smoothers for Large-Scale Finite Element Simulations

Efficient solution of linear systems is critical to achieve parallel scalability for finite element simulations. Multi-core architectures offer unprecedented performance yet require exceedingly high degrees of parallelism exacerbate the situation. Multigrid preconditioners have been a scalable solution approach. However, the performance of domain decomposition smoothers often used can degrade at large scales. We explore the effects of domain decomposition smoothers and the use of a hybrid parallel technique to improve the scalability of multigrid preconditioners.

Paul Lin

Sandia National Laboratories
ptlin@sandia.gov

Eric C. Cyr

Scalable Algorithms Department
Sandia National Laboratories
eccyr@sandia.gov

Sivasankaran Rajamanickam

Sandia National Laboratories
srajama@sandia.gov

CP3

A Symmetry-Based Decomposition Approach to Eigenvalue Problems with Applications to Electronic Structure Calculations

We propose a decomposition approach to differential eigenvalue problems with Abelian or non-Abelian symmetries. The original problem is divided into eigenvalue subproblems which require less eigenpairs and can be solved independently. Such a decomposition made the two-level parallel algorithm and has been applied to electronic structure calculations with several thousands of atoms. Both performance tests and analysis show that our approach is attractive for large-scale quantum eigenvalue problems.

Xingyu Gao

Beijing Institute of Applied Physics and Computational
Mathe
gao_xingyu@iapcm.ac.cn

Jun Fang

Institute of Applied Physics and Computational
Mathematics
fang_jun@iapcm.ac.cn

Aihui Zhou

Chinese Academy of Sciences
azhou@lsec.cc.ac.cn

CP3

CoarrayARPACK: Parallel ARPACK in Fortran 2008

The purpose of our presentation is to introduce a new parallel version of ARPACK implemented within the Fortran 2008 language. This represents a significant departure from PARPACK (Maschoff and Sorensen 1996). Our new implementation also provides numerous software upgrades to ARPACK; in particular, the new Fortran 2008 parallel language construct *coarrays* is exploited - thus obviating the need for MPI.

Richard J. Hanson

High Performance Computing and Analysis
Principal Consultant
richard.koolhans@gmail.com

Richard B. Lehoucq

Sandia National Laboratories
rblehou@sandia.gov

CP3

Parallel Strategy for Computing Eigenvalues of Non-Hermitian Matrices

We demonstrate the efficiency of the upcoming FEAST solver v3.0 for addressing the non-Hermitian eigenvalue problem. The algorithm calculates the subset of eigenpairs that exist within a given contour located in the complex

plane. All main properties of the FEAST algorithm for Hermitian problems are retained. In addition, the highly parallel scheme benefits from our new custom contour utility routine that allows users to partition the complex domain according to their specific problem.

James Kestyn
ECE Department
University of Massachusetts, Amherst
jkestyn@student.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

CP3

Effective Parallel Computation of Eigenpairs to Detect Anomalies in Very Large Graphs

The computational driver for an important class of graph analysis algorithms is the computation of leading eigenvectors of matrix representations of the graph. In this presentation, we discuss the challenges of calculating eigenvectors of modularity matrices derived from very large graphs (upwards of a billion vertices) and demonstrate the scaling properties of parallel eigensolvers when applied to these matrices. Numerical results are presented, along with system recommendations for performing eigenanalysis of very large graphs.

Michael M. Wolf, Benjamin Miller
MIT Lincoln Laboratory
michael.wolf@ll.mit.edu, bamiller@ll.mit.edu

CP4

∇ -Nabla: A Numerical-Analysis Specific Language for Exascale Scientific Applications

This paper tackles the challenging research and development problem to provide portable scientific applications for exascale architectures. It introduces the numerical-analysis specific ∇ language which generates optimized code from a single source for a variety of existing middlewares or hardwares. ∇ 's bottom-up toolset facilitates debugging and performance tuning by a detailed combination of co-designed techniques. The modularity, performance and scalability are demonstrated through several exascale proxy-applications on Nehalem, Fermi and Xeon-Phi architectures.

Jean-Sylvain Camier
CEA, DAM, DIF
Jean-Sylvain.Camier@cea.fr

CP4

Simplifying Heterogeneous Multicore Programming Using Industry Standards

Heterogeneous multicore processors are everywhere in today's well connected world. They employ a wide range of OSes across multiple cores. Even challenging is an accelerator, not running in any form of OS, but needs to interact with the processor. How to program these systems once

and reuse often? Software is an expensive investment. In our study, we design and deploy industry standards creating a shift from proprietary solutions and efficiently harnessing the processor capability.

Sunita Chandrasekaran, Barbara Chapman
University of Houston
sunita@cs.uh.edu, chapman@cs.uh.edu

CP4

Refactorization of a Hybrid, Highly Efficient Parallel PDE Solver

We discuss a refactorization and extension of an existing hybrid parallel code base for solving high-dimensional PDE problems. The code is written for good performance on large-scale clusters. We present measurements of performance and complexity of the two versions of the code. We also compare the amount of work required to add new features.

Malin Källén
Department of Information Technology
Uppsala University
malin.kallen@it.uu.se

CP4

PGAS Programming Model for Manycore Computers

Computers equipped with manycore processors, e.g., GPUs and Intel MICs, have become increasingly important to computational applications. Partitioned Global Address Space (PGAS) programming languages (e.g., UPC, Coarray Fortran, Chapel and X10) have emerged as an attractive alternative and complement to the traditional message-passing programming model (e.g., MPI). In this talk we present our work on designing and implementing a PGAS programming system for manycore processor architectures and share our experience on performance evaluation and optimization.

Yili Zheng
Lawrence Berkeley National Laboratory
yzheng@lbl.gov

CP5

Partitioning and Parallel Computation of Electricity Production Cost Models

Production cost models (PCMs) simulate electric power system operation at hourly (or higher) resolution. While computation times often extend into multiple days, the sequential nature of PCM's makes parallelism difficult. We exploit the persistence of generation unit commitment decisions to select partition boundaries for simulation horizon decomposition and parallel computation. Partitioning strategies consider solution integrity and computational load balancing as factors in determining partition boundaries. Partitioned simulations are benchmarked against sequential solutions for optimality and computation time.

Clayton Barrows
Pennsylvania State University
clayton.barrows@nrel.gov

Marissa Hummon, Wesley Jones
National Renewable Energy Laboratory

marissa.hummon@nrel.gov, wesley.jones@nrel.gov

CP5

Parallel Adaptive Cartesian Level-Set Methods

In this talk we will present parallel algorithms for level-set methods on octree (in 3D) and quadtree (in 2D) grids. The computational domain is first described by a macro-mesh and inside each macro-cell a tree is initialized. The collection of these trees are then managed and distributed in parallel using p4est library. Here, we propose scalable algorithms for the advection and re-initialization problems in the level-set method framework on such distributed adaptive grids.

Mohammad Mirzadeh

Dept. Mechanical Engineering

UCSB

m.mirzadeh@engineering.ucsb.edu

CP5

Polytope: A New Parallel Framework for Computing Voronoi Meshes on Complex Boundaries

Polytope is open-source software designed to generate unstructured mesh data based on the Voronoi diagram for the purposes of simulation and visualization. Polytope wraps existing serial algorithms for generating Voronoi and Delaunay graphs and parallelizes them to handle distributed input data. A clipping algorithm is used to generate boundary-conformal meshes, capable of handling non-convex region containing holes. This talk will describe the methods implemented in Polytope, focusing on the novel parallel algorithm for distributed mesh generation.

David Starinshak, John Owen

Lawrence Livermore National Laboratory

starinshak1@llnl.gov, owen@llnl.gov

Jeffrey N. Johnson

Applied Numerical Algorithms Group

Lawrence Berkeley National Laboratory

jnjohnson@lbl.gov

CP5

Randomized Heuristics for the Bipartite Matching Problem on Shared Memory Parallel Computers

We propose two efficient heuristics for the maximum cardinality bipartite matching problem. One of the heuristics is very simple for parallelization. Apart from a relatively straightforward initialization step, it has virtually no overhead. Furthermore, it has an approximation ratio of around 0.63. The second heuristic is designed to obtain much better results than the first one, with some overhead in a possible parallel implementation. Experiments with the parallel implementations of the proposed heuristics are presented.

Fanny Dufosse

Polytech' Nice - Sophia Antipolis, MODALIS

dufosse@i3s.unice.fr

Kamer Kaya

The Ohio State University

Department of Biomedical Informatics

kamer@bmi.osu.edu

Bora Ucar

LIP, ENS-Lyon, France.

Bora.Ucar@ens-lyon.fr

CP6

Multi Component Polymer Flooding Two Dimensional Oil Reservoir Simulation

We propose a high resolution finite volume scheme for a $(m+1)(m+1)$ system of nonstrictly hyperbolic conservation laws which models multicomponent polymer flooding in enhanced oil-recovery process in two dimensions. In the presence of gravity the flux functions need not be monotone and hence the exact Riemann problem is complicated and computationally expensive. To overcome this difficulty, we use the idea of discontinuous flux to reduce the coupled system into uncoupled system of scalar conservation laws with discontinuous coefficients. High order accurate scheme is constructed by introducing slope limiter in space variable and a strong stability preserving Runge-Kutta scheme in the time variable. The performance of the numerical scheme is presented in various situations by choosing a heavily heterogeneous hard rock type medium. Also the significance of dissolving multiple polymers in aqueous phase is presented

Sudarshan Kumar Kenettinkara

TIFR Centre For Applicable Mathematics, India

sudarshan@math.tifrbng.res.in

CP6

Recent Advances in PHAML

PHAML (Parallel Hierarchical Adaptive MultiLevel) is a parallel finite element code for the solution of 2D and 3D elliptic partial differential equations using h-, p-, and hp-adaptive refinement and multigrid solution techniques. Parallelism can be done through MPI, OpenMP, or hybrid MPI/OpenMP. In this talk we will present some of the recent advances in the development and use of PHAML.

William F. Mitchell

NIST, Gaithersburg, MD

william.mitchell@nist.gov

CP6

Interactive Multi-Scale High-Performance Computing

Due to ever increasing advances in hardware, today's high-performance computing (HPC) systems allow for an interactive treatment of even complex problems stemming from domains such as engineering, medicine, or geosciences. Such an approach not only bridges the gap between HPC and real-time user interaction, it also paves the way for new types of applications in order to obtain insight – not numbers! We will present our approach for interactive computing along with some sample application scenarios.

Ralf-Peter Mundani

TUM, Faculty of Civil Engineering and Geodesie

Chair for Computation in Engineering

mundani@tum.de

Jérôme Frisch

Chair for Computation in Engineering

Technische Universität München

frisch@tum.de

CP6

A Parallel Scattered Node Finite Difference Scheme for the Shallow Water Equations on a Sphere

A parallel solver for the shallow water equations on a sphere with applications in global weather and climate simulations is presented. We use a task-based programming model provided by the SuperGlue library for fine-grained synchronization on shared-memory nodes, and MPI to communicate between nodes. The SuperGlue run-time system provides task dependency management, schedules the tasks at run-time, and uses task stealing for load balancing. Absolute performance and scalability measures are analyzed and presented.

Martin Tillenius

Dept. of Information Technology
Uppsala University
martin.tillenius@it.uu.se

Elisabeth Larsson

Uppsala University, Sweden
Department of Information Technology
Elisabeth.Larsson@it.uu.se

Erik Lehto

National Center for Atmospheric Research
lehto@ucar.edu

Natasha Flyer

National Center for Atmospheric Research
Institute for Mathematics Applied to Geosciences
flyer@ucar.edu

CP7

Parallelization Challenges for Ensemble Data Assimilation

Ensemble Data Assimilation (DA) software combines observations of systems like Earth's atmosphere with an ensemble of forecasts to produce improved initial conditions for subsequent forecasts. The quasi-random locations of observations require irregular communication patterns and add significant complexity to large, parallel models that are already challenging to run on today's high performance machines. We discuss the implementation of the Data Assimilation Research Testbed, a community DA software facility, on emerging parallel architectures.

Helen Kershaw

NCAR
hkershaw@ucar.edu

Nancy Collins

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

Jeffrey Anderson

National Center for Atmospheric Research
Institute for Math Applied to Geophysics

jala@ucar.edu

CP7

Massively Parallel CPU and GPU Implementation of Elliptic Solvers in Geophysical Modelling

Many problems in geophysical modelling require the fast solution of elliptic PDEs in 'flat' domains. To meet operational requirements at ever increasing resolution, solvers must be algorithmically optimal and run efficiently on parallel architectures. We describe the implementation and optimisation of bespoke iterative solvers for typical PDEs encountered in numerical weather- and climate-prediction and demonstrate their scalability for the solution of equations with more than ten billion variables on massively parallel CPU- and GPU- clusters.

Eike H. Mueller, Robert Scheichl

University of Bath
e.mueller@bath.ac.uk, r.scheichl@bath.ac.uk

CP7

Geophysical EULAG Model with Three-Directional MPI Parallelization

Until recently, atmospheric models have been parallelized in two horizontal dimensions. With the advent of shared-memory multicore supercomputer nodes, hybrid MPI/OpenMP parallelization becomes a dominant approach. However, efficient hybrid parallelization of a complete model is usually a tedious task, and does not necessarily resolve intranodal performance bottlenecks like memory bandwidth saturation. Here we report on the performance of a recently accomplished three-dimensional MPI parallelization of the atmospheric model EULAG in geophysical applications, e.g. regional NWP experiment with a set of physical parametrizations on $O(10^6)$ number of cores. EULAG is now being integrated within COSMO weather prediction framework, aiming at Alpine weather prediction at 1 km and better resolution.

Zbigniew P. Piotrowski

National Center for Atmospheric Research
zbigniew.piotrowski@imgw.pl

Piotr Smolarkiewicz

2European Centre for Medium-Range Weather Forecasts
Reading, United Kingdom
piotr.smolarkiewicz@ecmwf.int

Andrzej Wyszogrodzki

NCAR
MMM
andii@ucar.edu

CP7

Large-Scale Parallel Simulation of Urban Flooding Scenarios

While multi-scale modelling and simulation has emerged as new research area having significant impact on many CSE-related disciplines, it is unlike an even bigger challenge to leverage multi-scale approaches for high resolution GIS and BIM (Building Information Modelling) data ranging from kilometre-scales to centimetre-scales, serving as input to a flow solver for multi-scale urban flooding scenarios. Therefore, we will present a framework for large-scale parallel

simulations subject to high accuracy flooding applications.

Vasco Varduhn

Technische Universität München
varduhn@tum.de

Ralf-Peter Mundani

TUM, Faculty of Civil Engineering and Geodesie
Chair for Computation in Engineering
mundani@tum.de

CP8

Modeling Stencil Code Optimizations

Performance models for stencil computation are available, but they miss predict performance after code optimizations are applied. We tackle this problem by extending our model, which now covers all major stencil optimizations techniques such as: temporal and spatial blocking, semi-stencil, loop unrolling. Our extended model accuracy remains high with only platform dependent variations. We truly believe that this tool is a practical and inexpensive complement to auto-tuning, therefore speeding-up the search for optimal implementations.

Mauricio Araya-Polo

Shell
mauricio.araya@shell.com

Raul de la Cruz

Barcelona Supercomputing Center (BSC)
raul.delacruz@bsc.es

CP8

Unstructured Mesh Physics Performance on Current and Future Architectures

In performance studies on new computer architectures, unstructured mesh methods are often bypassed in favor of “easier” structured mesh and particle-based methods. We present performance results from the unstructured mesh mini-app PENNANT, adapted from the LANL production code FLAG, running on current and new architectures. These results provide rough estimates of how various optimizations can improve performance on future architectures, and how performance may suffer if existing codes are run on these architectures without change.

Charles R. Ferenbaugh

Los Alamos National Laboratory
cferenba@lanl.gov

CP8

A Communication Algorithm for the Patch-Based Multiblock Structured Mesh Applications

Multiblock structured mesh allows to handle complex configurations which are widely existed in computational physics applications. A Patch-based data structure is always used in applications with multiblock structured mesh to get satisfying parallel performance. However, such Patch-based data structure seriously challenges the block to block data communications. This talk presents an algorithm for such communication and introduces its integration to JASMIN infrastructure to support the peta-scale simulations while tens of thousands of processors are used.

Performance results show its robustness.

Hong Guo

Institute of Applied Physics and Computational Mathematics
guo_hong@iapcm.ac.cn

Zeyao Mo

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

Aiqing Zhang

Institute of Applied Physics and Computational Mathematics
zhang_aiqing@iapcm.ac.cn

CP8

Numa-Aware Runtime Optimization of Structured Grid Numerical Simulations

Structured grids are widely used in important applications such as complex electromagnetic environment and inertial confinement fusion. However, structured grid simulations, built upon SMP abstractions, fail to catch up in the exascale era where NUMA architecture is widely adopted. To systematically address this issue, we propose a numa-aware programming model and a numa-aware heap manager. Data layout and communications are then optimized upon this model. These approaches improve significantly performance of several realworld applications.

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

Zeyao Mo

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

CP9

A Mixed Precision Eigenvalue Solver on GPU

Jacobi-Davidson type methods (JD) can solve large-scale eigenvalue problems efficiently. Their efficiencies rely on the solutions of the inner-loop linear systems. These systems can be solved in lower accuracy without downgrading the final accuracy of the target eigenvalues. By taking advantage of this algorithmic feature and the computational power of GPU, we develop a mixed precision eigensolver. We demonstrate the efficiency of the solver by solving various eigenvalue problems, including a three-dimensional photonic crystal simulation.

Jihming Huang

Department of Mathematics
National Taiwan University
r01221022@ntu.edu.tw

Tsung-Ming Huang

Department of Mathematics, National Taiwan Normal University
min@ntnu.edu.tw

Pochuan Wang
Institute of Applied Mathematical Sciences, National Taiwan University
r02246011@ntu.edu.tw

Weichung Wang
Institute of Applied Mathematical Sciences, National Taiwan University
wwang@ntu.edu.tw

CP9

X-Ray Laser Imaging of Biomolecules Using Multiple Gpus

Extremely bright X-ray lasers are becoming a promising tool for 3D imaging of biomolecules. By hitting a beam of streaming particles with a very short burst of a high energy X-ray and collecting the resulting scattering pattern, the 3D structure of the particles can be deduced. The computational complexity associated with transforming the data thus collected into a 3D intensity map is very high and calls for efficient data-parallel implementations. We present ongoing work in accelerating this application using multiple GPU nodes. In particular, we look at the scaling properties of the application and give predictions as to the computational viability of this imaging technique.

Jing Liu
Uppsala University, Department of Information Technology,
Division of scientific computing
jing.liu@it.uu.se

Stefan Engblom
Division of Scientific Computing
Uppsala University
stefan.engblom@it.uu.se

CP9

A Scalable and Fast Multi-Gpu Eigensolver for 3D Photonic Crystal Band Structure

Numerical simulation of three-dimensional photonic crystals band structure plays an important role in its physics and applications, such as optimal structure design. However, the computation is very expensive, because it requires solutions of many large-scale generalized eigenvalue problems. Based on newly proposed FFT-based preconditioning and null-space free techniques, we show how these computations can be accelerated via multiple GPUs. Numerical experiments show that the GPU-based eigensolver achieves significant speedup with almost linear scalability.

Weichung Wang
Institute of Applied Mathematical Sciences, National Taiwan University
wwang@ntu.edu.tw

Jiahong Chen
Department of Mathematics, National Taiwan University
r00221034@ntu.edu.tw

Jihming Huang
Department of Mathematics
National Taiwan University
r01221022@ntu.edu.tw

Tsung-Ming Huang
Department of Mathematics, National Taiwan Normal University
min@ntnu.edu.tw

CP10

Compiler Based Generation and Autotuning of Communication-Avoiding Operators for Geometric Multigrid

We describe a compiler approach to introducing communication-avoiding optimizations in geometric multigrid (GMG), one of the most popular methods for solving partial differential equations. Communication-avoiding optimizations reduce vertical communication through the memory hierarchy and horizontal communication across processes or threads, usually at the expense of introducing redundant computation. We focus on applying these optimizations to the smooth operator, which successively reduces the error and accounts for the largest fraction of the GMG execution time. Our compiler technology applies a set of novel and known transformations to derive an implementation comparable to hand-written optimizations. An underlying autotuning system explores the tradeoff between reduced communication and increased computation, as well as trade offs in threading schemes, to automatically identify the best implementation for a particular architecture.

Protonu Basu
University of Utah
School of Computing
protonu@cs.utah.edu

CP10

Reducing Coarse Grids Contention in a Parallel Algebraic Multigrid

Algebraic multigrid solvers are attractive in many scientific applications. However, increasing amounts of communication and data movement on coarser levels present significant challenges. In this talk, we consider a few algorithms improving the scalability. We discuss reducing the number of coarse levels through aggressive coarsening, improving locality of communicating processors through repartitioning, reducing data movement through gathering and storing data redundantly. The results demonstrate substantial speedups on a variety of test problems.

Andrey Prokopenko
Sandia National Laboratories
aprokop@sandia.gov

CP10

Parallel AMG Solver Based on Adaptive Setup Strategy with Applications in Large-Scale Radiation Hydrodynamics Simulations

Coarse-level visitation is the main reason that causes loss of scalability for AMG solver on massive parallel computer. In our presented adaptive setup strategy, coarsening is performed based on the smoothing behavior on each level, instead of constructing via an independent setup phase in

traditional procedure. As a result, doing relaxations on finer-levels as much as possible, while the required coarse-levels as less as possible. Realistic simulations on $O(10^4)$ cores show the improved scalability.

Xiaowen Xu

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

Zeyao Mo

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

CP11

Data Structures and Algorithms for High-Dimensional Structured Adaptive Mesh Refinement (SAMR)

Accurate solution of time-dependent, high-dimensional PDEs requires massive-scale parallel computing and efficient numerical techniques. Spatial decomposition is particularly challenging in higher dimensions, since memory requirements quickly grow prohibitively large and out of reach even for massively parallel computers. We present a framework for parallel SAMR, tailored for spatial decomposition of localized solutions in high dimensional domains. The key data structures and parallelized algorithms for mesh organization and load balancing are demonstrated, along with numerical results.

Magnus Grandin

Uppsala University, Sweden
magnus.grandin@it.uu.se

CP11

Deflation Based Domain Decomposition Preconditioners

Domain decomposition methods are widely used in applied mathematics and regarded as highly scalable algorithms, alongside multigrid methods. Making those methods scalable to thousands of processors is however not a straightforward task. Projection operators are one of the essential tools for achieving scalability: they are used for building deflation preconditioners. We will present a C++ framework accompanied by theoretical results to show how effective it can be to solve problems with billions of unknowns.

Pierre Jolivet

Laboratoire Jacques-Louis Lions
jolivet@ann.jussieu.fr

Frederic Nataf

Laboratoire J.L. Lions
nataf@ann.jussieu.fr

Christophe Prud'homme

Institut de Recherche Mathematique Avancee
prudhomme@unistra.fr

CP11

A Domain Decomposition Method for Unsteady Flow Routing in Complex River Systems

We present an optimization-based domain decomposition for a nonlinear unsteady flow routing problem in complex

river systems. We couple the domain decomposition technique with the performance graphs approach utilizing pre-computed solutions along reaches on a river. While efficient, these performance graphs require extensive memory allocation. Domain decomposition reduces the memory constraint by only requiring those solutions relevant to a specific reach. The performance graphs approach efficiency, in turn, allows an optimization-based scheme to be competitive.

Veronika S. Vasylykivska

Department of Mathematics
Oregon State University
vasylkiv@math.oregonstate.edu

Nathan L. Gibson

Oregon State University
gibsonn@math.oregonstate.edu

Arturo Leon, Luis Gomez

School of Civil and Construction Engineering
Oregon State University
arturo.leon@oregonstate.edu, gomezcul@onid.orst.edu

CP12

Scalability and Performance Analysis for Replica-Exchange Wang-Landau Sampling

Replica-Exchange Wang-Landau sampling (REWL) [Phys. Rev. Lett. 110, 210603 (2013)] is a novel, massively parallel realization for Wang-Landau sampling, a robust Monte Carlo simulation method which demonstrates a broad range of applications in physics, mathematics and statistics. We study the principles for the speed-up and scaling behavior of REWL on high performance machines using different parameter settings. With the parallelization achieved on the CPU level, REWL shows excellent strong and weak scalings.

Ying Wai Li

Oak Ridge National Laboratory
Oak Ridge, TN 37831, USA
ywli@hal.physast.uga.edu

Thomas Vogel

Los Alamos National Laboratory
Los Alamos, NM 87545, USA
thomasvogel@physast.uga.edu

Thomas Wuest

Swiss Federal Research Institute WSL
thomas.wuest@wsl.ch

Markus Eisenbach

Oak Ridge National Laboratory
Oak Ridge, TN 37831, USA
eisenbachm@ornl.gov

David Landau

The University of Georgia
Athens, GA 30602, USA
dlandau@physast.uga.edu

CP12

A Multi-Physics Approach for Time-Parallel Plasma Physics Applications

In the field of plasma physics, vastly diverging time scales

of rapidly moving quasi-free electrons and heavy slow ions render fully resolved molecular dynamics simulations extremely difficult. In this talk, we report on the use of time-parallel methods with our massively space-parallel Barnes-Hut tree code PEPC to overcome the problem's intrinsic strong scaling limit. Following a multi-physics approach, we exploit the system's time scales to resolve fast dynamics while maintaining long physical simulation times.

Mathias Winkel

Juelich Supercomputing Centre
Forschungszentrum Juelich GmbH
m.winkel@fz-juelich.de

Paul Gibbon

Forschungszentrum Juelich GmbH
Juelich Supercomputing Centre
p.gibbon@fz-juelich.de

Michael Minion

ICME Stanford
mlminion@gmail.com

Matthew Emmett

Lawrence Berkeley National Laboratory
Center for Computational Sciences and Engineering
mwemmett@lbl.gov

CP12

Replica-Exchange Wang-Landau Sampling - a Highly Scalable Framework for Stochastic Optimization

Replica-Exchange Wang-Landau sampling (REWL) is a new, massively parallel, Monte Carlo scheme which has proven very promising for the exploration of new frontiers in statistical physics [T. Vogel *et al.*, PRL 110, 210603 (2013)]. The procedure provides a generic framework for simulations of complex systems in various scientific domains. Moreover, REWL features powerful scalability and fault-tolerance properties making it suitable to run on Petascale supercomputers. In addition to benchmark results from statistical physics, we demonstrate the potential of REWL as a novel means of stochastic optimization exemplified on the traveling salesman problem.

Thomas Wuest

Swiss Federal Research Institute WSL
thomas.wuest@wsl.ch

Ying Wai Li

Oak Ridge National Laboratory
Oak Ridge, TN 37831, USA
ywli@hal.physast.uga.edu

Thomas Vogel

Los Alamos National Laboratory
Los Alamos, NM 87545, USA
thomasvogel@physast.uga.edu

David Landau

The University of Georgia
Athens, GA 30602, USA
dlandau@physast.uga.edu

CP13

Parallel Locking Sweeping for Static Hamilton-

Jacobi Problems of Many Dimensions

A large number of minimum time optimal control and path planning problems can be described by a static Hamilton-Jacobi equation of the form $H(\nabla u(\mathbf{x}), \mathbf{x}) = 1$. Often the dimension of the state variable is large, rendering serial methods unsuitable. In this talk, we outline a fast parallel algorithm for such problems. We discuss the parallel Fast Sweeping algorithm and compare to other popular methods. We provide parallel performance results and discuss GPU acceleration.

Miles L. Detrixhe

University of California Santa Barbara
mdetrixhe@engineering.ucsb.edu

CP13

Hashing in the Scientific World

We employ hashing methods and the discrete properties of computational meshes to optimize spatial operations including remap, neighbor-finding, sorting and table lookup in scientific computing applications. Hashing is an intrinsically parallel, non-comparison based search algorithm which requires no inter-thread communication, allowing performance portability across CPU and GPU architectures. We demonstrate consistent speed-up and increased performance across every device tested and explore the ubiquitous application of spatial hashing in scientific computing.

Rebecka Tumblin

Los Alamos National Lab
University of Oregon
rebeckatumblin@gmail.com

Robert Robey

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

Peter Ahrens

UC Berkeley
ptrahrens@gmail.com

Sara Hartse

Brown University
sara.hartse@gmail.com

CP13

Co-Design of Extremely Scalable Algorithms/Architecture for 3-Dimensional Linear Transforms

Based on the proposed 3D decomposition of an $N \times N \times N$ data into $P \times P \times P$ blocks, where $P = N/b$ and $b \in [1, 2, \dots, N]$ is the blocking factor, we systematically design extremely scalable algorithms with a circular data reuse for highly-parallel implementation of any forward/inverse 3D separable transform on the algorithm/architecture-related $P \times P \times P$ torus network of computer nodes. All algorithms require $3P$ 'compute-and-roll' time-steps, where each step is equal to the time of execution in each node b^4 multiply-add operations and concurrent movement of $O(b^3)$ data between nearest-neighbour nodes.

Stanislav Sedukhin

University of Aizu, Japan

sedukhin@u-aizu.ac.jp

blad@nvidia.com

CP14**A Hybrid MPI/openmp 3D FFT Implementation for Parallel Plane Wave First-Principles Materials Science Codes**

First principles electronic structure calculations based on a plane wave (Fourier) expansion of the wavefunctions are the most commonly used approach for electronic structure calculations in materials and nanoscience. The performance of these application codes depends critically on having a 3D FFT that scales efficiently to large processor counts. A specialized hybrid MPI/OpenMP implementation on the Cray XE6 significantly outperforms library routines or a pure MPI version, particularly on large processor counts.

Andrew M. CanningLawrence Berkeley National Laboratory
acanning@lbl.gov**CP14****Exploiting Data Reuse for Reduction of Communication Volume in Parallelization of Multi-Dimensional FFTs**

We present a decomposition method for parallelization of multi-dimensional FFTs with adaptive decomposition and transpose order awareness. Based on a row-wise decomposition that translates the multi-dimensional data into one-dimensional data for allocating to the processes, our method can adaptively decompose the data in the lowest possible dimensions. Also, by analyzing all possible cases, we find out the best transpose orders with minimal communication volumes for 3-D, 4-D, and 5-D FFTs. Numerical results demonstrate good performance of our method.

Truong Vinh Truong DuyResearch Center for Simulation Science
Japan Advanced Institute of Science and Technology
duyvt@jaist.ac.jp

Taisuke Ozaki

Japan Advanced Institute of Science & Technology
t-ozaki@jaist.ac.jp**CP14****Multiple GPU FFT Algorithms**

Cards containing multiple Graphical Processing Units (GPUs) accelerate a range of scientific applications, but few take advantage of more than one GPU per CPU. High performance Fast Fourier Transform (FFT) algorithms are particularly challenging on these distributed memory systems due to all-to-all communication requirements. 1-d, 2-d and 3-d FFT algorithms have been developed for multiple GPUs and performance will be shown for NVIDIA multi-GPU systems.

Kevin Wadleigh, Alex Fit-Florea, Slawomir Kierat,
Lukasz Ligowski, Paul SidenbladNVIDIA
kwadleigh@nvidia.com, afitflorea@nvidia.com,
skierat@nvidia.com, lligowski@nvidia.com, psiden-**CP15****Resource-Aware Scheduling in Task Parallel Frameworks**

A task-based parallel framework needs to be dependency-aware so that tasks are scheduled respecting a certain execution order. Load-balancing is achieved by distribution of tasks over the available compute units combined with task stealing. However, even if the scheduling is successful with respect to load balance, the scaling may be sub-optimal due to resource contention. We present a modified scheduling strategy with resource constraints. Performance gains for both model examples and real applications are shown.

Elisabeth LarssonUppsala University, Sweden
Department of Information Technology
Elisabeth.Larsson@it.uu.se

Martin Tillenius

Dept. of Information Technology
Uppsala University
martin.tillenius@it.uu.se

Rosa M. Badia

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

Xavier Martorell

Barcelona Supercomputing Center, Spain
Computer Sciences
xavier.martorell@bsc.es**CP15****A Comparison of Checkpoint Techniques for HPC Applications**

This talk presents a comparison of current checkpoint protocols for high-performance computing applications: coordinated vs. hierarchical, stable-storage vs. in-memory, etc. We provide a detailed description of each variant, together with analytical models that allow to assess their pros and cons for current and future large-scale platforms.

Yves RobertENS LYON, France
& Institut Universitaire de France
yves.robert@ens-lyon.fr**CP15****Checkpointing Schemes for Adjoint-Based Optimization of Jet Engine Noise**

Jet engine noise is a serious problem for the crew of aircraft carriers. Computationally expensive 3-D simulations of the exhaust are required for accurate measures of noise reduction, thus adjoint-based optimization is an ideal method. JENRE, a fully scalable parallel computational fluid dynamics code, is used. Checkpointing methods have been implemented and optimized to maximize memory usage. This allows for low memory GPU implementation and minimizes overhead.

Zachary B. SpearsMiddle Tennessee State University
zbs2b@mtmail.mtsu.edu

Andrew Corrigan
Naval Research Laboratory Washington DC
acorriga@lcp.nrl.navy.mil

CP16

Timing Performance Surrogates in Auto-Tuning for Qualitative and Quantitative Factors

Auto-tuning problems involving both qualitative and quantitative (Q&Q) factors are common in scientific computing software. But little model-based tuning methods handle such Q&Q factors. To minimize the total runtime, we propose several Kriging-based surrogate methods that manage the Q&Q factors separately or jointly. The proposed approaches are applied to a parallel algebraic multigrid linear system solver simulating bubbles in liquid. Numerical results identify the advantages of these surrogate schemes and show their efficiency.

Jiahong Chen
Department of Mathematics, National Taiwan University
r00221034@ntu.edu.tw

Ray-Bing Chen
Department of Statistics,
National Cheng Kung University
raybingc@gmail.com

Akihiro Fujii
Kogakuin University
fujii@cc.kogakuin.ac.jp

Reiji Suda
Department of Computer Science, The University of
Tokyo
reiji@is.s.u-tokyo.ac.jp

Weichung Wang
Institute of Applied Mathematical Sciences,
National Taiwan University
wwang@ntu.edu.tw

CP16

Performance Analysis of the Pagosa Application

Performance analysis tests were run using Pagosa, a computational fluid dynamics code, on three LANL HPC systems. Tests included small to large size problems and single processor to 32768 processor runs. Pagosa exhibited excellent scaling characteristics and continued to make use of increased number of processors even while approaching very small work loads per processor. This ability to divide a problem into smaller and smaller work loads per processor has promising implications for exascale computing.

Robert L. Kelsey
Los Alamos National Laboratory
rob@lanl.gov

CP16

Towards Parallel Library Generation for Polynomial Arithmetic

In the contributed lecture, we will present our recent work towards parallel library generation for polynomial arithmetic, with a focus on the implementation and performance of the underlying FFT algorithms. We show how to represent the transforms and algorithms in a high-level domain-

specific language. Then we systematically apply symbolic and algebraic optimizations to generate a library with recursive function closure optimized for parallel computing platforms. The resulting scalar and vector codes provide comparable speedup as the fixed size code presented in previous work, which is an order of magnitude faster over the hand-tuned modpn library used by Maple. Additional speedup has been achieved by exploiting multi-core architecture automatically. More transforms and algorithms, such as the Truncated Fourier Transform and Convolution are being experimented in the same fashion, which will be combined to generate a general-size parallel library for polynomial arithmetic.

Lingchuan Meng
Drexel University
Computer Science Department
lingchuanmeng@gmail.com

MS1

Reliability and Resiliency Trends: An Update

Trends in the reliability of leadership-class computing systems continue to cause concerns about their ability to provide a stable platform for applications that run for any significant length of time. While some progress has been made in software solutions for resiliency, more work is still needed in many areas. This talk updates the current status of the system reliability trends, and the survey of the various software solutions for resiliency, originally presented at PP10.

Larry Kaplan
Cray
lkaplan@cray.com

MS1

Fault Tolerance at Exascale: Is Checkpointing Dead?

Exascale systems are projected to fail frequently and parallel file system performance is not expected to match the increase in computational speed. Thus, checkpointing will become less practical in the future. We will present model predictions of the performance of multilevel checkpointing on future systems and research directions that aim to reduce the overhead of checkpointing to acceptable levels. These include checkpoint compression, asynchronous checkpoint movement, and new data movement strategies.

Kathryn Mohror
Lawrence Livermore National Laboratory
mohror1@llnl.gov

MS1

Models for Fault-Tolerance at Very Large Scale

This talk introduces models for fault-tolerance techniques at very large scale. It will cover coordinated and hierarchical checkpoint protocols, fault prediction, and replication. It will also deal with silent error detection through verification mechanisms. The models will be instantiated with realistic scenarios for HPC applications on current petascale and future exascale platforms.

Yves Robert
ENS LYon, France
& Institut Universitaire de France

yves.robert@ens-lyon.fr

MS1

Toward a Local Failure-Local Recovery Resiliency Model

On future extreme-scale HPC systems, Checkpoint-Restart would be a suboptimal resilience solution due to its unproportional response to single process failures. We propose an alternative model called Local Failure Local Recovery (LFLR) that provides application developers with the ability to recover locally and continue execution when a process is lost. We will discuss what features are required from the hardware and software, and what approaches application developers might use in the design of application codes.

Keita Teranishi, Michael A. Heroux
Sandia National Laboratories
knteran@sandia.gov, maherou@sandia.gov

MS2

Performance Modeling for Tiling with PIN

Data movement is costly in both energy and performance, so it is important to model the cache hierarchy memory traffic under different data-structure layouts. In this talk, we will present a tool that with only a minimal set of annotations to the application source, can log the memory accesses of a non-tiled application as if its arrays were actually tiled. Motivation, design, and results obtained will be presented.

John Bachan, Didem Unat
Lawrence Berkeley National Laboratory
jdbachan@lbl.gov, dunat@lbl.gov

Cy Chan
Electrical Engineering and Computer Science Department
Massachusetts Institute of Technology
cychan@lbl.gov

John Shalf
Lawrence Berkeley National Laboratory
jshalf@lbl.gov

MS2

Tiling Dense and Sparse Computations for Parallelism and the Memory Hierarchy of Future Architectures

Tiling is used by programmers and compilers for two purposes: (1) for parallelism, to divide computation evenly among different processors; and, (2) for locality, to reduce the data footprint of a subcomputation to fit within limited-capacity storage such as a register file, cache or software-controlled storage. On GPUs, hierarchical tiling is needed for both purposes, given the parallelism hierarchy across blocks and threads and the complex memory hierarchy. Most compiler-based approaches are restricted to the affine domain, where loop bounds and array subscript expressions are linear functions of loop indices. For computations that employ indirect accesses to arrays through an index arrays, such as for example, sparse matrix-vector multiplication, run-time information is needed to perform aggressive optimization. In this talk, we describe CUDA-CHiLL, which automatically generates hierarchically tiled parallel code for GPUs, for both parallelism and locality. We target both affine and non-affine computations, and

show performance comparable to manually-tuned GPU libraries CUBLAS and CUSP.

Anand Venkat, Manu Shantharam
University of Utah
anandv@cs.utah.edu, manushan@cs.utah.edu

Mary Hall
School of Computing
University of Utah
mhall@cs.utah.edu

Michelle Strout
Colorado State University
michmstrout@cs.colostate.edu

MS2

Two Approaches for Scheduling and Tiling to Very Deep Hierarchies, Implemented and Available in R-Stream

We present our developments in hierarchical iteration space tiling within the R-Stream polyhedral parallelizing compiler. R-Stream's tiling component includes a mixed model-based and empirical search which reduces the tile size search space dramatically. High-level parameters of the search can modify the search space by combining its objectives and constraints. We developed two expressions of hierarchical tiling: the first one expresses the program as a memory-oblivious recursive decomposition, effectively implementing arbitrary-depth tiling with adaptive leaf tile size and adaptive parallelism. The second expression is a memory-conscious approach in which each level is statically tailored to a level of the target machine's memory hierarchy. After presenting these two approaches, we discuss their advantages and current limitations and present our current findings and experimental results.

Richard Lethin, Benoit Meister, Nicolas Vasilache, Muthu Baskaran, David Wohlford
Reservoir Labs
lethin@reservoir.com, meister@reservoir.com, vasilache@reservoir.com, baskaran@reservoir.com, wohlford@reservoir.com

MS2

Tiling for Parallel Execution of Stencil Computations

Stencil computations arise in many computational domains. Achieving high performance requires effective exploitation of data reuse at the different levels of the memory hierarchy, as well as utilization of multiple levels of parallelism in modern architectures, including SIMD parallelism at the lower level and MIMD parallelism at the higher levels. This talk will present recent research on compiler techniques for tiled code generation for parallel execution of stencil computations on multicore CPU and GPU architectures.

P. Sadayappan
Dept. of Computer and Information Science
Ohio State University
saday@cis.ohio-state.edu

MS3

Parallel Discrete Event Simulation of Infectious

Disease Spread in Animal Populations

In this work we study spatio-temporal models of bacterial infection spread among animal populations where epidemiological processes are formulated as continuous-time Markov-chains. To accelerate computations in our discrete-event simulator URDME we divide work among cores of shared memory multiprocessors by using the models inherent parallelism in between transport events gathered from actual data. We will present the epidemiological model and study the overall parallel efficiency of the simulation.

Pavol Bauer
Uppsala University
pavol.bauer@it.uu.se

Stefan Engblom
Division of Scientific Computing
Uppsala University
stefan.engblom@it.uu.se

MS3

Performance of Time Warp on 1,966,080 Cores

Time Warp is an optimistic synchronization protocol for parallel discrete event simulation that coordinates the available parallelism through its rollback and antimessage mechanisms. In this talk we present the results of a strong scaling study of the ROSS massively parallel discrete-event simulator running Time Warp with reverse computation and executing the well-known PHOLD benchmark on Lawrence Livermore National Laboratory's "Sequoia" Blue Gene/Q supercomputer. The benchmark has 251 million PHOLD logical processes and was executed in several configurations up to a peak of 7.86 million MPI tasks running on 1,966,080 cores. At the largest scale it processed 33 trillion events in 65 seconds, yielding a sustained speed of 504 billion events/second using 120 racks of "Sequoia". This is by far the highest event rate reported by any parallel discrete event simulation to date, whether running PHOLD or any other benchmark. ROSS exhibited a super-linear speedup throughout the strong scaling study, with more than a 97x speed improvement from scaling the number of cores by only 60x (from 32,768 to 1,966,080). From these results, we will discuss ROSS' implementation on the Blue Gene/Q system and the implications and potential opportunities for using ROSS to realize kinetic Monte Carlo methods.

Christopher Carothers
Computer Science
Rensselaer Polytechnic Institute
chris@cs.rpi.edu

MS3

Parallelization, Error Analysis and Partial Asynchrony in Lattice Kinetic Monte Carlo

In this talk we address mathematical, numerical and algorithmic issues arising in the parallelization of spatially distributed Kinetic Monte Carlo simulations, by developing a new hierarchical operator splitting as means of decomposing efficiently and systematically the computational load and communication between multiple processors. The spatial decomposition of the Markov operator of the KMC algorithm into a hierarchy of operators, corresponds to processors in the parallel architecture. Based on this operator decomposition, we formulate Fractional Step Approxima-

tion schemes by employing the Trotter Theorem and its random variants; these schemes, (a) are run independently on each processor through a serial KMC simulation on each fractional step time-window, giving rise to partially asynchronous algorithms, and (b) are characterized by the communication schedule between processors. Furthermore, the proposed mathematical framework allows us to rigorously justify the numerical and statistical consistency of the proposed algorithms by showing goal-oriented error estimates for quantities of interest. The approach also provides a systematic evaluation of the balance between error and processor communication in the parallel algorithms.

Markos A. Katsoulakis
University of Massachusetts, Amherst
Dept of Mathematics and Statistics
markos@math.umass.edu

Georgios Arampatzis
University of Crete, Greece
garab@math.uoc.gr

Petr Plechac
University of Delaware
Department of Mathematical Sciences
plechac@math.udel.edu

MS3

Parallelization of Kinetic Monte Carlo Using Synchronous Algorithms: Applications and Differences in Continuum and Discrete Systems

Kinetic Monte Carlo (kMC) is one of the most widely used simulation techniques in a number of scientific disciplines. While other similarly important techniques have kept pace with and taken advantage of advances in computational parallelism, kMC has been slow in adapting to massively parallel computations. In this talk we analyze the causes of this by describing a synchronous parallel algorithm for kMC simulations and analyzing its performance and potential capabilities for computational physics simulations.

Jaime Marian
LLNL
marian1@llnl.gov

MS4

On Predicting Performance on Different Task Mappings Using Supervised Learning

We present the use of supervised learning algorithms (e.g. randomized decision trees) to correlate parallel application performance with communication data, such as the communication graph and network hardware counters. We propose new hybrid metrics that provide high correlation with application performance, and may be useful for accurate performance prediction. For different communication kernels and a production application, we demonstrate a strong correlation between the proposed metrics and the execution time of different task mappings. LLNL-ABS-647552

Abhinav Bhatele
Lawrence Livermore National Laboratory
bhatele@llnl.gov

Nikhil Jain
University of Illinois
nikhil@illinois.edu

Todd Gamblin
Lawrence Livermore National Laboratory
gamblin2@llnl.gov

Laxmikant Kale
University of Illinois at Urbana-Champaign
kale@illinois.edu

MS4

Topology-Aware Task Mapping Using Geometric Partitioning

We present a new method for mapping applications' MPI tasks to cores of a parallel computer such that communication and execution time are reduced. We consider the case of sparse node allocation within a parallel machine. The goal is to assign tasks to cores so that interdependent tasks are performed by "nearby" cores. Our new method applies a geometric partitioning algorithm to both the tasks and the processors, and assigns task parts to processor parts. On a Cray XE6, our mapping method reduced execution time 31% on average on 64K cores for a finite-difference mini-app, and reduced communication time by 26% on average on 6K cores for a molecular dynamics mini-app.

Mehmet Deveci
The Ohio State University
mdeveci@bmi.osu.edu

Siva Rajamanickam, Karen D. Devine, Vitus Leung
Sandia National Laboratories
srajama@sandia.gov, kddevin@sandia.gov, vjleung@sandia.gov

MS4

Task Mapping Stencil Computations for Non-Contiguous Allocations

We examine task mapping algorithms for noncontiguously allocated parallel jobs, such as those on Cray systems. We apply novel and adapted algorithms to this setting and evaluate them using experiments and simulations. Our focus is on jobs with a stencil communication pattern. We evaluate them with a miniApp whose communication behavior mimics CTH, a shock physics application with this pattern. Our strategies improve its running time by as much as 35% over a baseline strategy.

Vitus Leung
Sandia National Laboratories
vjleung@sandia.gov

David Bunde
Knox College
dbunde@knox.edu

MS4

A Comparison of Task Mapping Strategies on Two Generations of Cray Systems

We present an architectural overview of two generations of Cray systems, the XE6 and XC30, with a focus on opportunities for task mapping. The XE6 uses a 3-D torus network topology with heterogeneous link speeds while the XC30 uses a low-diameter Dragonfly topology, which in theory should be less sensitive to task mapping. We describe modifications made to the libtopomap library to support task

mapping on these architectures and present initial empirical results.

Kevin Pedretti
Sandia National Laboratories
ktpedre@sandia.gov

Torsten Hoefler
University of Illinois
htor@illinois.edu

MS5

Recovery-Restart Techniques for Resilient Krylov Methods

The advent of extreme scale machines will require the use of parallel resources at an unprecedented scale, possibly leading to a high rate of hardware faults. In this presentation, we investigate recovery followed by restarting strategies for the resilience of Krylov subspace linear solvers as well as of eigensolvers.

Emmanuel Agullo
INRIA
emmanuel.agullo@inria.fr

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

Salas Pablo
CERFACS
salas@cerfacs.fr

Zounon Mawussi
INRIA
mawussi.zounon@inria.fr

MS5

Increasing the Arithmetic Intensity of Multigrid with Stencil Compilers

Abstract not available at time of publication.

Pieter Ghysels
Lawrence Berkeley National Laboratory
pghysels@lbl.gov

MS5

Acceleration of Domain Decomposition Based Algorithms by Communication Avoiding and Hiding Krylov Method

Abstract not available at time of publication.

Tomas Kozubek
VSB Technical University of Ostrava
Czech Republic
tomas.kozubek@vsb.cz

MS5

Latency Hiding of Global Reductions in Pipelined

Krylov Methods

Abstract not available at time of publication.

Wim I. Vanroose
Universiteit Antwerpen
Belgium
wim.vanroose@ua.ac.be

MS6**Application of Algebraic Multigrid (petsc) for Adaptive Mesh Refinement Applications (Chombo)**

We report on progress in using algebraic multigrid (AMG) methods in the numerical library PETSc for challenging problems in structured grid adaptive mesh refinement (AMR) applications that use the Chombo library. Chombo's built-in geometric multigrid (GMG) solvers are fast for simple operators but for problems with complex geometry GMG is not effective and AMG is an effective solution. We discuss new capabilities in Chombo for constructing matrices, required for AMG, from AMR problems.

Mark Adams
Columbia University
mfadams@lbl.gov

MS6**Scalable Preconditioners for Atmospheric Climate Simulation**

We have developed a preconditioner to improve the efficiency of linear system solves within an implicit formulation of the shallow water equations in the Community Atmospheric Model (CAM-SE). This preconditioner is based on an approximate block factorization and leverages TRILINOS solvers. We will discuss the effects of the preconditioner on scalability of implicit simulations within CAM-SE, and examine the use of CESM-style timers to diagnose and mediate performance bottlenecks within the solver framework.

Aaron Lott
Lawrence Livermore Nat'l Lab
lott3@llnl.gov

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

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

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

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

MS6**Rapid Development of a New Ice Sheet Application****Code Using Albany and Trilinos**

A new Ice Sheet model has been rapidly developed in the Albany code, using dozens of math libraries from Trilinos. The code was born with distributed memory parallelism, unstructured-grid finite element discretizations, scalable linear algebra, robust nonlinear solves, automatic differentiation for analytic Jacobians and sensitivities, load balancing, uncertainty quantification, and more. This high degree of leverage is enabled by object-oriented software and modern software engineering tools and processes.

Andrew Salinger
Applied Computational Methods Dept, Sandia National Labs
Sandia National Labs
agsalin@sandia.gov

Irina Kalashnikova
Sandia National Laboratories
ikalash@sandia.gov

Mauro Perego
Florida State University
mperego@fsu.edu

Stephen Price
Los Alamos National Laboratory
sprice@lanl.gov

Ray S. Tuminaro
Sandia National Laboratories
Computational Mathematics and Algorithms
rstumin@sandia.gov

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

MS6**Nonlinear Solvers and Time Integrators for Dislocation Dynamics**

Strain hardening simulations within the Parallel Dislocation Simulator (ParaDiS) require integrating stiff systems of ordinary differential equations in time with expensive force calculations, discontinuous topological events and rapidly changing problem size. To reduce simulation run times we are incorporating new nonlinear solvers and higher order implicit integrators from the FASTMath Suite of Nonlinear and Differential / Algebraic Equation Solvers (SUNDIALS) and evaluating scalability of the resulting code on the LLNL Sequoia system.

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

Athanasio Arsenlis, Sylvie Aubry
Lawrence Livermore National Laboratory
arsenlis@llnl.gov, sylvie.aubry@llnl.gov

David Gardner
Southern Methodist University
dgardner@mail.smu.edu

Gregg Hommes, Kathryn Mohror
Lawrence Livermore National Laboratory
hombres1@llnl.gov, mohror1@llnl.gov

Daniel R. Reynolds
Southern Methodist University
Mathematics Department
reynolds@smu.edu

MS7

High Performance Solution of Dense Linear Systems with Application to Large 3D Electromagnetics Problems on a Petascale Computer and on a Cluster of Gpus

The numerical treatment of high frequency electromagnetic scattering is very computationally intensive. For scattering, the electromagnetic field must be computed around and inside 3D bodies. Because of this, accurate numerical methods are used to solve Maxwell's equations in the frequency domain, and it leads to solve large dense linear systems. So we have developed on our Petascale supercomputer a hybrid CPU/GPU solver for systems with millions of complex unknowns and thousands of right hand sides.

David Goudin, Cedric Augonnet, Agnes Pujols, Muriel Sesques
CEA/CESTA
david.goudin@gmail.com, cedric.augonnet@cea.fr,
agnes.pujols@cea.fr, muriel.sesques@cea.fr

MS7

On Parallelization of MKL LAPACK SVD

The talk is devoted to modifications of a two-step algorithm for reducing a matrix to bidiagonal form. At the first stage the matrix is reduced to band matrix. A successive band reduction is then used to reduce a band matrix to bidiagonal form. Both stages apply two-sided orthogonal transformations to the matrix based on Householder reflectors. The underlying idea for our modifications is to use speculative computations of components of the Householder transformations. Performance comparisons between the Intel® Math Kernel Library (Intel® MKL), PLASMA bidiagonal reduction and our implementation are provided.

Sergey V Kuznetsov
Intel Corporation
sergey.v.kuznetsov@intel.com

MS7

A Parallel Approach to the Solution of Pde Through Componentwise Approximation of Matrix Functions

Krylov subspace spectral (KSS) methods solve time-dependent PDE by approximating each component of the solution with respect to some orthonormal basis using techniques from "matrices, moments and quadrature" for evaluating entries of matrix functions. This componentwise approach allows Krylov subspace dimension to be independent of the spatial mesh, thus significantly improving scalability. The focus of this talk is the adaptation of these methods to parallel architectures.

James V. Lambers
University of Southern Mississippi
Department of Mathematics

James.Lambers@usm.edu

MS7

Exploring Emerging Manycore Architectures for Uncertainty Quantification Through Embedded Stochastic Galerkin Methods

We explore approaches for improving the performance of embedded stochastic Galerkin uncertainty quantification methods on emerging computational architectures. Our work is motivated by the trend of increasing disparity between floating-point throughput and memory access speed. We describe several new stochastic Galerkin matrix-vector product algorithms and measure their performance on contemporary manycore architectures. We demonstrate these algorithms lead to improved memory access patterns and ultimately greater performance within the context of iterative linear system solvers.

Eric Phipps
Sandia National Laboratories
Optimization and Uncertainty Quantification Department
etphipp@sandia.gov

H. Carter Edwards
Sandia National Laboratories
hcedwar@sandia.gov

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

Jakob Ostien
Sandia National Laboratories
jostie@sandia.gov

MS8

Analyzing Shakespeare's Dramas Using Networks

We discuss how Shakespeares plays can be modeled as networks. The interaction within the plays are more precise therefore the associated social networks differs in a qualitative sense from those seen in social media such as Facebook and Twitter. We demonstrate that by analyzing different networks of the plays, (such as centrality values and driver nodes) we can identify the role of the character, the social context and sometimes the genre of the play.

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

MS8

Using Networks to Model Student Conceptual Understanding of Physics

Cognitive ability as represented in writing may link students language with their ability to do physics. Connections made in writing by a student between physics concepts may signify ways that the student understand physics. Based on student writings on how to solve conceptual problems, this talk outlines a method for investigating these writings using network science and text mining. The end goal is to create networks that capture central parts of

student understanding of physics.

Jesper Bruun
Department of Science Education
University of Copenhagen
jesper.bruun@gmail.com

MS8

On Creating Networks from MRI Data

There are a series of steps taken in turning MRI data into a network that can be analyzed for insight into structural and functional properties of the brain. We explore the impact of some of the choices made in this process on the resultant network model.

Tzu-Yi Chen
Pomona College
tzuyi@cs.pomona.edu

MS8

Stoichiometry of Steroidogenesis: Towards Understanding Optimal Design and Function

We present structural and optimal flux properties of a reconstructed piscine steroidogenic network. Analysis showed network robustness was susceptible to deletions of structurally relevant reactions relative to high flux reactions. Structurally relevant reactions exhibited a power law distribution (degree exponent = 2.3), whereas optimal fluxes exhibited near-random (homogenous) flux distributions (degree exponent = 2.7). This observation is consistent with the demonstrated vulnerability of metrics exhibiting power law distributions to targeted attack (or reaction deletions).

David Hala
Institute of Applied Sciences
University of North Texas
david.hala@unt.edu

Duane Huggett
University of North Texas
dbhuggett@unt.edu

MS9

Fault-tolerant Iterative Linear Solvers with Multi-level Fault Detection

Computer hardware trends may expose incorrect computation or storage. We present a fault model for floating-point numbers. Given this, almost all faults in a GMRES iteration can either be detected cheaply, or cause bounded error. Equilibration sharpens detection. We then apply "detection or bounded error" to GMRES as the inner solver in a fault-tolerant inner-outer iteration. Detection inside inner solves reduces wasted work, and detection after each inner solve with reliable outer computations ensures convergence.

Mark Hoemmen
Sandia National Laboratories
mhoemme@sandia.gov

James Elliott
North Carolina State University

jjellio3@ncsu.edu

MS9

Towards an Unified ABFT Approach for Resilient Dense Linear Algebra

With the number of supercomputers' components on a sharp increase over recent years, it is not an unusual statistic being reported that during its 537 days of operation, an average of 2.33 failures per day occurred at a large super-computer installation at a government lab. Also, less than 10 continuous hours of operation was fault-free. This increases the need for fault-tolerant libraries and this talk will cover the ways of bringing such libraries closer to reality.

Piotr Luszczek
Department of Electrical Engineering and Computer Science
University of Tennessee, Knoxville
luszczek@eecs.utk.edu

MS9

A Fault Tolerant Implementation of Multi-Level Monte Carlo Methods

The theory behind fault tolerant multi-level Monte Carlo (FT-MLMC) methods was recently developed and tested. These tests were made without a real fault tolerant implementation. We implemented a MPI-parallelized fault tolerant MLMC version of an existing parallel MLMC code (ALSVID-UQ). It is based on the User Level Failure Mitigation, a fault tolerant extension of MPI. We confirm our FT-MLMC theory by means of simulations of the two-dimensional stochastic Euler equations of gas dynamics.

Stefan Pauli, Manuel Kohler
ETH Zurich
unknown, manuel.kohler@bsse.ethz.ch

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

MS9

Self-Stabilizing Iterative Solvers

We show how to use the idea of self-stabilization to make fault-tolerant iterative solvers. A self-stabilizing system is one that, starting from an arbitrary state (valid or invalid), reaches a valid state within a finite number of steps. We give two proofs-of-concept of self-stabilizing iterative linear solvers: one for steepest descent (SD) and one for conjugate gradients (CG). We present experiments and analysis to show scalability of this approach for high fault rates.

Piyush Sao, Richard Vuduc
Georgia Institute of Technology
piyush3@gatech.edu, richie@cc.gatech.edu

MS10

Chapel Language Features for Hierarchical Tiling and Exascale Architectures

Chapel is an emerging parallel programming language whose design and development are being led by Cray Inc. Chapel's design avoids locking key performance-oriented

decisions into the definition of the language and compiler in order to maximize its portability across architectures and algorithms. In this talk, I will highlight key features of Chapel that are designed to support hierarchically tiled data structures, tunable tiled iterators, and exascale-era compute nodes.

Bradford L. Chamberlain

Cray Inc.
bradc@cray.com

MS10

Hardware Support for Collective Data Transfers

The scaling of parallel architectures and the resulting power wall calls for innovations through hardware and software co-design, to increase performance without sacrificing power. In this talk, I will present a hardware support mechanism for collective data transfers to maximize main memory and data movement efficiency, as well as the language constructs to make this capability easily available to the programmer. I will demonstrate this concept in distributed array computations.

Georgios Micheliogiannakis

Lawrence Berkeley National Laboratory
mihelog@lbl.gov

MS10

Programming with Tiles

Tiling is an effective mechanism to develop high performance implementations of scientific applications. Tiling can be expressed in either at the loop nest in the code or within the data type. We take the approach to support tiling directly at the type system to make manipulation of tiles easy and centralized. This advanced data type, namely Hierarchically Tiled Arrays (HTAs), can be used to express hierarchical partitioning of data, its layout, and its distribution in parallel computers. The associated communication operations are simply array operations within single threaded programs, facilitating code development. HTA comes with support for dynamic partitioning and overlapped tiling for stencil computations.

David Padua

UIUC
padua@illinois.edu

MS10

Programming Model Support for Tiling

Tiling is an important optimization to introduce parallelism and provide data locality in an application. Tiling can be expressed in various forms in a program. Here, we present two forms, which are implemented in Mint and Tida. Mint is an annotation based model where the programmer can specify tile size and thread mapping by annotating the loops in the program. Mint partitions the iteration space based on this information and maps threads to the iteration space. Its compiler generates CUDA code and optimizes the data movement using software-managed memory and register file to reduce memory traffic. The second approach, Tida, both centralizes and parameterizes the tiling information at the data structure. The programmer expresses the hierarchy and topology of the tiles at the creation of data structures and the loops iterate over the tiles. The tiling information is used by the compiler and runtime to introduce parallelism, enable data locality, and

task scheduling. In this talk, we discuss the advantages and disadvantages of two approaches and present the interface of both Mint and Tida along with performance results.

Didem Unat

Lawrence Berkeley National Laboratory
dunat@lbl.gov

MS11

Advanced Coupling Explorations for Parallel Coupled Neutronic and Thermal-Hydraulic Simulation

This presentation will discuss recent explorations for coupling nuclear reactor neutronics and conjugate heat transfer in parallel, through comparisons of Picard iteration, Anderson acceleration, and a Jacobian-Free Newton-Krylov approach. The problem studied here consists of the k-eigenvalue form of the neutron transport equation governing the distribution of neutrons with the reactor coupled with a conjugate heat transfer model describing the heat removal from fuel pins into neighboring coolant channels.

Kevin Clarno, Mark Berrill

Oak Ridge National Laboratory
clarnokt@ornl.gov, berrillma@ornl.gov

Steven Hamilton

ORNL
hamiltonsp@ornl.gov

Roger Pawlowski

Multiphysics Simulation Technologies Dept.
Sandia National Laboratories
rppawlo@sandia.gov

John Turner

Oak Ridge National Laboratory
Computational Engineering and Energy Sciences
turnerja@ornl.gov

MS11

Physics Based Coupling for Multiscale Full Core Nuclear Reactor Simulation

Simulation of nuclear reactors is a key technology for advances in efficiency, safety and reliability of existing and future reactor designs. Historically, reactor core simulation has been accomplished by linking existing codes. Recent advancements in the Idaho National Laboratory MOOSE framework have enabled a new approach: multiple domain-specific applications, built utilizing the same software framework, are combined to create a cohesive application. A flexible coupling capability allows many different types of data exchanges to occur simultaneously on high performance parallel computers, enabling multiscale, multiphysics simulations of nuclear reactor cores. Examples using the KAIST-3A benchmark core as well as the Westinghouse AP-1000 are demonstrated.

Derek R. Gaston

Idaho National Laboratory
derek.gaston@inl.gov

Cody Permann

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

John Peterson, Andrew Slaughter, David Andrs, Yaqi Wang
Idaho National Laboratory
jw.peterson@inl.gov, andrew.slaughter@inl.gov,
david.andrs@inl.gov, yaqi.wang@inl.gov

Michael Short
Massachusetts Institute of Technology
hereiam@mit.edu

Michael Tonks, Richard Martineau
Idaho National Laboratory
michael.tonks@inl.gov, richard.martineau@inl.gov

MS11

Code Integration Strategies for Large-Scale Reactor Simulation

The Consortium for Advanced Simulation of Light Water Reactors is a U. S. Department of Energy Innovation Hub charged with developing a virtual reactor toolkit that will incorporate science-based models, state-of-the-art numerical methods, modern computational science and engineering practices, and uncertainty quantification (UQ) and validation against operating pressurized water reactors. It will couple state-of-the-art fuel performance, neutronics, thermal-hydraulics (T-H), and structural models with existing tools for systems and safety analysis and will be designed for implementation on both today's leadership-class computers and next-generation advanced architecture platforms. This work will describe activities in developing scalable coupling tools for leadership class machines and their application to a in coupling three stand alone codes for a core simulator capability. The physics discretizations, coupled solution algorithms and convergence requirements will be shown. We will discuss issues on code integration, solution techniques and conservation. Initial scalability results of the individual codes, the coupled codes and the data transfer tools on leadership class platforms will be assessed.

Roger Pawlowski
Multiphysics Simulation Technologies Dept.
Sandia National Laboratories
rppawlo@sandia.gov

Roscoe Bartlett
Oak Ridge National Laboratory
bartlettra@ornl.gov

Kenneth Belcourt
Sandia National Laboratories
kbelco@sandia.gov

Kevin Clarno
Oak Ridge National Laboratory
clarnokt@ornl.gov

Derek R. Gaston
Idaho National Laboratory
derek.gaston@inl.gov

Russell W. Hooper
Sandia National Laboratories
rhooper@sandia.gov

Robert Salko
Oak Ridge National Laboratory

salkork@ornl.gov

MS11

Tools Supporting the Assembly of Multiphysics Simulation Codes from Standalone Parallel Physics Codes

Many parallel codes have been developed for simulating single-physics in areas relevant to nuclear reactor analysis, including thermal-fluid transport, neutron transport, and structural mechanics. Performing multi-physics analysis with these codes requires support both at the top level, for driving the overall simulation, as well as the bottom level, for and transferring solutions between the various physics, often on disparate meshes. We describe tools being developed to support this type of reactor analysis, and results of applying these tools to several reactor types.

Timothy J. Tautges
Argonne National Labory
tautges@mcs.anl.gov

Vijay Mahadevan, Rajeev Jain
Argonne National Laboratory
mahadevan@anl.gov, jain@mcs.anl.gov

MS12

Algorithmic Rethinking and Code Reengineering for Truly Massively Parallel Ab Initio Molecular Dynamics Simulations

We present a new massively parallel implementation of exact exchange in the context of DFT calculations. Novel parallelization schemes allow for millions of threads to be utilized achieving near perfect strong scaling, dramatically extending the practicality of exact exchange calculations. We illustrate the advantages of our approach on the study of Lithium-air batteries.

Costas Bekas
IBM Research - Zurich
BEK@zurich.ibm.com

MS12

Decay Properties of Density Matrices: Rigorous Results

We apply general results from approximation theory and matrix analysis to the study of the decay properties of spectral projectors associated with large and sparse Hermitian matrices. Our theory leads to a rigorous proof of the exponential off-diagonal decay ("nearsightedness") for the density matrix of gapped systems at zero electronic temperature in both orthogonal and non-orthogonal representations, thus providing a firm theoretical basis for the possibility of $O(N)$ methods. We further discuss the case of density matrices for metallic systems at positive electronic temperature.

Michele Benzi
Department of Mathematics and Computer Science
Emory University
benzi@mathcs.emory.edu

MS12

Scalable Density Functional Theory on Blue

Gene/Q

As modern supercomputers continue to grow in size and complexity, it becomes increasingly challenging to scale communication-intensive methods like Density Functional Theory. This talk will highlight recent efforts to optimize the Qbox code on the LLNL Sequoia Blue Gene/Q machine. The impact of threading, tuned math kernels, and communication-aware parallel linear algebra libraries will be discussed, along with the prospects of further scalability on the path to exascale.

Erik W. Draeger
Lawrence Livermore Nat. lab.
draeger1@llnl.gov

MS12**A Scalable and Accurate $O(N)$ Parallel Algorithm for Large-Scale First-Principles Molecular Dynamics Simulations**

Traditional algorithms for First-Principles molecular dynamics simulations only gain a modest capability increase from exascale computers due to their $O(N^3)$ complexity. We present a new algorithm with strict $O(N)$ complexity that exploits sparsity and nearest neighbor communication for large-scale molecular dynamics simulations. We demonstrate excellent weak scaling for up to $O(10^5)$ atoms with $O(1)$ minute per molecular dynamics step, while preserving accuracy with respect to traditional $O(N^3)$ algorithms.

Daniel Osei-Kuffuor
Lawrence Livermore National Laboratory
oseikuffuor1@llnl.gov

Jean-Luc Fattebert
Lawrence Livermore National Lab.
fattebert1@llnl.gov

MS13**Reconstructing Householder Vectors from TSQR**

The Tall-Skinny QR (TSQR) algorithm is much more communication efficient than the standard Householder algorithm for QR decomposition of matrices with many more rows than columns. However, TSQR produces a different representation of the orthogonal factor and therefore requires more software development to add functionality with the new representation. We discuss how to perform TSQR and then reconstruct the original representation (a set of Householder vectors) with the same communication efficiency and some extra computational cost.

Grey Ballard, James Demmel
UC Berkeley
gmballa@sandia.gov, demmel@eecs.berkeley.edu

Laura Grigori
INRIA
France
Laura.Grigori@inria.fr

Mathias Jacquelin
Lawrence Berkeley National Lab
mjacquelin@lbl.gov

Hong Diep Nguyen
University of California-Berkeley
hdnguyen@eecs.berkeley.edu

Edgar Solomonik
University of California at Berkeley
solomon@eecs.berkeley.edu

MS13**Avoiding Synchronization in Geometric Multigrid**

Krylov subspace methods are commonly used as coarse grid solvers in geometric multigrid codes. At scale, however, global communication in each iteration of the Krylov method can dominate overall multigrid solve time. In this talk, we present the communication-avoiding biconjugate gradient stabilized method (CA-BICGSTAB) and discuss its integration into two high-performance geometric multigrid codes. Numerical experiments demonstrate improvements in parallel scalability and performance due to communication-avoiding techniques for problems from a variety of scientific domains.

Erin C. Carson, Nicholas Knight
UC Berkeley
ecc2z@eecs.berkeley.edu, knight@eecs.berkeley.edu

Samuel Williams, Michael Lijewski, Ann S. Almgren
Lawrence Berkeley National Laboratory
swwilliams@lbl.gov, mjlijewski@lbl.gov,
ASAAlmgren@lbl.gov

James Demmel
UC Berkeley
demmel@eecs.berkeley.edu

MS13**High-Bandwidth Communication Avoidance: Oxy-moron Or Recipe?**

Similar parallel programming models at very large and small scales are commonly assumed. Coupling PRAM algorithm for Burrows-Wheeler compression with buildable high on-chip bandwidth hardware to compress/decompress messages sent by MPI nodes traverses that: support large systems MPI by PRAM-like programming of small systems. Google's Snappy compromises compression ratio for speed for MapReduce; however, we demonstrate feasibility of high compression ratio, while beating Snappy on speed using true parallelism (compress/decompress in parallel the entire input).

James A. Edwards
The University of Maryland
jedward5@umd.edu

Uzi Vishkin
University of Maryland
vishkin@umiacs.umd.edu

MS13**Recent Progress in Nested and Communication Avoiding Incomplete Factorization Preconditioners**

In this talk we present several recent advances in preconditioners for solving linear systems that are suitable for matrices arising from the discretization of a system of PDEs on unstructured grids. We focus on communication avoiding incomplete LU factorization and on nested direction preserving preconditioner. We present a set of numerical experiments on matrices with anisotropies and jumps in

the coefficients that show the efficiency and the parallel performance of these preconditioners.

Laura Grigori
INRIA
France
Laura.Grigori@inria.fr

Frederic Nataf
Laboratoire J.L. Lions
nataf@ann.jussieu.fr

Sophie Moufawad
LRI - INRIA Saclay, France
sophie.moufawad@lri.fr

Long Qu
INRIA Saclay / LRI, France
qu@lri.fr

MS14

Performance Optimization of Block Eigensolvers for Nuclear Structure Computations

Existing implementations of SpMM and SpMM_T significantly underperform expectations. We present and analyze optimized implementations of SpMM and SpMM_T. We base our implementation on the compressed sparse blocks (CSB) matrix format and target the Cray XC30 machine at NERSC. We develop a performance model that allows us to understand and estimate the performance characteristics of our SpMM kernel implementations, and demonstrate the efficiency of our implementation on a series of real-world matrices from nuclear structure computations.

H. Metin Aktulga, Sam Williams
Lawrence Berkeley National Lab
hmaktulga@lbl.gov, swwilliams@lbl.gov

Aydin Buluc
Lawrence Berkeley National Laboratory
abuluc@lbl.gov

Chao Yang
Lawrence Berkeley National Lab
CYang@lbl.gov

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

MS14

Automated Massively Parallel Simulations Starting with Circuit Design Data

Large-scale simulation workflows should operate directly from the base problem information that is often not the form needed by the physics simulation package. Tools for massively parallel 3-D integrated circuit simulations, starting from 2-D layout information, will be presented. The layout information must be converted into non-manifold solid models as needed by the meshing procedures. The simulation workflow uses FASTMath and Simmetrix automated, adaptive unstructured mesh technologies integrated with the Albany finite element analysis procedures.

Max Bloomfield
RPI

bloomm2@rpi.edu

Brian Granzow
Rensselaer Polytechnic Institute
granzb@rpi.edu

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

MS14

Progress in Adaptive Computational Mechanics Applications Using the Albany Agile Components Framework

To address the large deformation of ductile materials in fracture and failure regimes, such as laser welds comprised of stainless steel 304L, adaptive meshing and state variable mapping strategies are being developed in Albany. Albany is an analysis framework that leverages Trilinos and RPI SCOREC tools to provide advanced solution, preconditioning, mesh adaptation and load balancing capabilities. Recent results illustrate the parallel considerations required to resolve the dominant mechanisms while satisfying equilibrium and consistent physical state during the adaptive step.

Jakob T. Ostien, Glen Hansen, James W. Foulk III,
Alejandro Mota, Mike Veilleux
Sandia National Laboratories
jtostie@sandia.gov, gahanse@sandia.gov, jw-
foulk@sandia.gov, amota@sandia.gov, mgveill@sandia.gov

MS14

Parallel Infrastructure for Multiscale Simulations

The Automated Multimodel Simulation Infrastructure (AMSI) will be overviewed. The focus of AMSI is to provide developers an effective component-based environment to integrate multiple analysis procedures that operate at different scales with scale linking procedures to create new multiscale simulations for problems of interest. Specific consideration is given to supporting multiple options for the parallel execution of the multiscale simulation steps including support for dynamic load balancing required when models and discretizations are adapted.

William R. Tobin, Daniel Fovargue
Rensselaer Polytechnic Institute
tobinw2@rpi.edu, dfovargue@gmail.com

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

MS15

Hybrid Octree/Quadtree AMR for Anisotropic Domains

Many problems in geosciences, such as ice sheets, oceans, and the atmosphere, involve thin, anisotropic domains. Researchers in these areas are using 3D models where once they used 2D approximations. The p4est library for parallel adaptive mesh refinement (AMR), which uses a forest-of-octrees approach to AMR, only natively supports isotropic mesh refinement, which is often insufficient for

such problems. We present an extension to this library with two refinement modes for these anisotropic problems.

Tobin Isaac

University of Texas at Austin
tisaac@ices.utexas.edu

Carsten Burstedde
Universität Bonn
burstedde@ins.uni-bonn.de

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

MS15

MPI/OpenMP Parallelization of Sam(oa)2 Using Sierpinski Curves

We present sam(oa)², a package for parallel adaptive refinement of triangular grids generated by newest-vertex bisection; target applications include porous media flow and tsunami simulation. sam(oa)² uses a stack-and-stream approach and an element order defined by the Sierpinski space-filling curve to store and process the grid and simulation data. The locality properties induced by the Sierpinski curve are retained even throughout adaptive refinement and coarsening and are exploited for efficient MPI/OpenMP parallelisation. Triangle strips defined by the Sierpinski order are split into grid sections that act as tasks for dynamic load balancing.

Oliver Meister, Michael Bader
Technische Universität München
meistero@in.tum.de, bader@in.tum.de

Kaveh Rahnama
University of Stuttgart
Institute of Parallel and Distributed Systems
rahnama@in.tum.de

MS15

A Nested Partitioning Scheme for Adaptive Meshes on Parallel Heterogeneous Clusters

Modern supercomputers are increasingly requiring the presence of accelerators and co-processors. However, it has not been easy to achieve good performance on such heterogeneous clusters. The key challenge has been to ensure good load balance and that neither the CPU nor the accelerator is left idle. Traditional approaches have offloaded entire computations to the accelerator, resulting in an idle CPU, or have opted for task-level parallelism requiring large data transfers between the CPU and the accelerator. True work-parallelism has been hard as the Accelerators cannot directly communicate with other CPUs (besides the host) and Accelerators. In this work, we present a new nested partition scheme to overcome this problem. By partitioning the work assignment on a given node asymmetrically into boundary and interior work, and assigning the interior to the accelerator, we are able to achieve excellent efficiency while ensure proper utilization of both the CPU and Accelerator resources.

Hari Sundar

Institute for Computational and Engineering Sciences
University of Texas at Austin
hari@ices.utexas.edu

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

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

MS15

Parallel Computing Research Topics in the Context of the Chombo AMR Code

Abstract not available at time of publication.

Brian Van Straalen

Lawrence Berkeley National Laboratory
Computational Research Division
bvstraalen@lbl.gov

MS16

Matrix Factorizations in MapReduce with Applications to Model Reduction

Matrix computations are the heart of many algorithms, but current tools for matrix computations in MapReduce are not efficient for data sets that arise scientific data analysis. In particular, matrices with many more rows than columns, so-called tall-and-skinny matrices, have a structure that allows for efficient MapReduce implementations. In this talk, we will provide an overview of QR and SVD for tall-and-skinny matrices in MapReduce and discuss their application to reduced order modeling.

Paul Constantine

Stanford University
pconstan@mines.edu

Austin Benson
UC Berkeley
arbenson@stanford.edu

MS16

A First and Second Introduction to MapReduce

Gleich and De Sterck will split the time in this first talk to present two introductions to the MapReduce computational model, the Hadoop implementation of MapReduce, and the potential of using Hadoop/MapReduce for scientific computing. These two talks are designed for those who want to learn more about MapReduce and Hadoop and how to design algorithms for them. We'll see a few basic algorithms and tasks in this talk before hearing about more advanced applications in the remainder of the first mini-symposium session.

David F. Gleich

Purdue University
dgleich@purdue.edu

Hans De Sterck
University of Waterloo
Applied Mathematics
hdsterck@uwaterloo.ca

MS16

Scaling Up Tensor Decompositions with MapRe-

duce

Tensor decompositions are increasingly gaining popularity in data science applications. Albeit extremely powerful tools, scalability to truly large datasets for such decomposition algorithms is still a challenging problem. In this talk, we provide an overview of recent algorithmic developments towards the direction of scaling tensor decompositions to big data. In particular, we present ways of leveraging the Map/Reduce framework in order to scale up tensor decompositions in an efficient manner. We showcase the effectiveness of our methods, by providing a variety of real world applications - whose volume previously rendered their analysis very hard, if not impossible- where our algorithms were able to discover interesting patterns and anomalies.

Evangelos Papalexakis
CMU
epapalex@cs.cmu.edu

MS16**Generating Large Graphs with Desired Community Structure**

The analysis of community structure is important for large scale graphs derived from cyber security data, internet-scale social networks, and other applications. Significant community structure can be inferred from the degree distribution and cluster coefficient distribution by degree. We have developed a MapReduce algorithm for estimating such distributions by sampling, with high accuracy and quantifiable error bounds. We use this information in another MapReduce algorithm to generate similar synthetic graphs with billions of edges.

Todd Plantenga
Sandia National Laboratories
tplante@sandia.gov

MS17**Managing Resilience in Exascale Computing**

Future exascale systems will be constructed from devices that will be less reliable than those used today, and faults will become the norm, not the exception. This will pose significant problems for system designers and programmers. In this work, we present an approach based on a range of programming model extensions which is complemented by runtime introspection. Such an approach leverages programmer knowledge and can significantly improve the probability that applications run to successful conclusion.

Saurabh Hukerikar, Pedro Diniz
Information Sciences Institute
University of Southern California
saurabh@isi.edu, pedro@isi.edu

Bob Lucas
ISI
Univ. of Southern California
rflucas@isi.edu

MS17**Tolerable Fault Tolerance**

Fault tolerance is possible with current programming styles, and without additional lines of code that obscure the meaning of algorithms. Furthermore, we should only pay a

recovery cost commensurate with the severity of a fault (i.e. recovering from transient cache line corruption should take milliseconds of time, not a full checkpoint restart). This talk will explain how we can achieve tolerable fault tolerance through a weak form of transactions, in combination with idempotent operations.

Jeff Keasler
Lawrence Livermore National Laboratory
keasler1@llnl.gov

MS17**Evaluating the Impact of Faults and Recovery Mechanisms in Exascale Applications**

As we approach exascale, we expect machines to become less reliable. We must understand the impact of this trend on applications. To what degree do faults manifest themselves in applications and what is the performance hit caused by recovery mechanisms? We are developing an emulation infrastructure, the GREMLINs, to drive a wide range of experiments, incl. recovery blocks, impact of network errors, and recovery from a failed node by reconstructing its state from neighbors.

Martin Schulz, Ignacio Laguna
Lawrence Livermore National Laboratory
schulzm@llnl.gov, lagunaperalt1@llnl.gov

MS17**GVR-Enabled Trilinos: An Outside-In Approach for Resilient Computing**

Resilience has become a major challenge for scientific computing frameworks like Trilinos. The Global View Resilience (GVR) system employs a novel, outside-in approach, which utilizes a global-view data model, a multi-version mechanism and a unified error signalling/handling interface, in order to provide flexible, portable and efficient fault management in large-scale systems. We demonstrate the effectiveness of GVR-enabled Trilinos by adding fault-tolerance functionality to several Trilinos solvers and applications.

Ziming Zheng, Zachary Rubenstein
Department of Computer Science, The University of Chicago
zimingzheng@uchicago.edu, zar1@uchicago.edu

Andrew Chien
The University of Chicago
Argonne National Laboratory
achien@cs.uchicago.edu

MS18**A Methodology for Characterizing the Opportunity and Feasibility of Reconfigurable Memory Hierarchies for Improved Energy Efficiency**

Abstract not available at time of publication.

Laura Carrington
San Diego Supercomputing Center
lcarrington@sdsu.edu

MS18**Holistic Performance Measurement and Analysis**

for High End Applications

Abstract not available at time of publication.

Karen Karavanic
Portland State University
karavan@cs.pdx.edu

MS18**Application of the Pi Theorem from Dimensional Analysis to Computer Performance Modeling**

Abstract not available at time of publication.

Bob Numrich
City University of New York
numri001@umn.edu

MS18**Pattern-Driven Node-Level Performance Engineering**

Abstract not available at time of publication.

Jan Treibig
Erlangen Regional Computing Center
Erlangen, Germany
jan.treibig@rrze.uni-erlangen.de

MS19**The Parallel Full Approximation Scheme in Space and Time (PFASST) on Extreme Scales**

The PFASST algorithm is a method for parallelizing ODEs and PDEs in time. PFASST decomposes the time domain across multiple processors and operates on each time-slice concurrently. LIBPFASST is a Fortran implementation of PFASST that uses either MPI or pthreads for parallelization, and can be used in conjunction with spatial parallelization techniques. Various applications of PFASST, from $O(10)$ cores to $O(100k)$ cores, will be presented and performance aspects of LIBPFASST will be discussed.

Matthew Emmett
Lawrence Berkeley National Laboratory
Center for Computational Sciences and Engineering
mwemmett@lbl.gov

MS19**Hybrid Mpi-OpenMP Implementation of Waveform Relaxation**

In this talk, we review waveform relaxation (a parallel-in-time and parallel-in-space) approach to solving general classes of PDEs. We focus on our hybrid MPI-OpenMP implementation, which allows for high levels of concurrency, providing a path towards exa-scale mathematics.

Benjamin Ong, Scott High
Department of Mathematics
Michigan State University
bwo@math.msu.edu, highscot@msu.edu

Felix Kwok
University of Geneva

Felix.Kwok@math.unige.ch

MS19**Implementation Strategies for Parallel-in-Time Methods**

There are a number of different mathematical approaches available to introduce concurrency in the temporal direction for the numerical solution of initial value problems, examples being parabolic multi-grid, waveform relaxation, Parareal or PFASST. However, there is still little experience how to optimally implement such methods for use in large- and extreme-scale parallel simulations. The talk will give a brief overview of different approaches and discuss previous and new results on using time parallelization in HPC.

Daniel Ruprecht
Institute of Computational Science
University of Lugano
daniel.ruprecht@usi.ch

Robert Speck
Jülich Supercomputing Centre
Forschungszentrum Jülich
r.speck@fz-juelich.de

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

MS19**The Parareal Algorithm - Applications to Fusion Plasma Physics**

Simulating the turbulent behaviour of magnetically confined, fusion plasma is arguably one of the 21st centuries HPC Grand Challenges. Simultaneous ion + electron-scale gyro-kinetic simulations are currently beyond the reach of existing super-computers for example. Such simulations will be key to rapidly achieving thermonuclear breakeven on machines like ITER and the first DEMO fusion reactors. This talk describes the application of the parareal algorithm to simulations of fusion plasma to reduce the wallclock time.

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

MS20**Multiresolution DFT and $O(N)$ Methods**

We describe the use of multiresolution analysis and nonlinear approximation methods to solve the equations of density functional theory in 3D. In particular, the use of a multiwavelet basis produces a hierarchy of localized equations separated by scale. Fast applications of Green's functions for the Poisson equation and the Lippmann-Schwinger's equations will be described. Scaling results of applications to computational chemistry and nuclear structures beyond 100K cores will be presented.

George Fann
Oak Ridge National Laboratory

fanngi@ornl.gov

MS20

Recent Progress on the Pole Expansion and Selected Inversion Method for Solving Kohn-Sham Density Functional Theory

The standard diagonalization method for solving the Kohn-Sham density functional theory (KS-DFT) scales cubically with respect to system size. In the recently developed pole expansion plus selected inversion method (PEXSI), KS-DFT is solved by evaluating the selected elements of the inverse of a series of sparse symmetric matrices. We show that the PEXSI method scales at most quadratically with respect to system size for all materials, and recent progress of the massively parallel PEXSI software which can scale to tens of thousands of processors for solving large scale quantum systems.

Lin Lin

Lawrence Berkeley National Laboratory
linlin@lbl.gov

MS20

Divide-Conquer-Recombine Algorithms for Metascapable Quantum Molecular Dynamics Simulations

We developed a divide-conquer-recombine algorithmic framework for large spatiotemporal-scale quantum molecular dynamics (QMD) simulations. The scheme has achieved parallel efficiency over 0.9 on 786432 IBM BlueGene/Q processors for 1.9 trillion electronic degrees-of-freedom QMD in the framework of density functional theory. We will discuss several applications including: (1) 16616-atom QMD simulation of rapid hydrogen production from water using metallic alloy nanoparticles; and (2) 6400-atom nonadiabatic QMD simulation of singlet fission of excitons for efficient solar cells.

Aiichiro Nakano

University of Southern California
anakano@usc.edu

MS20

High-Order and Enriched Finite Element Methods for Electronic Structure

Over the past few decades, the planewave (PW) pseudopotential method has established itself as the method of choice for large, accurate, density-functional calculations in condensed matter. However, due to its global Fourier basis, the PW method suffers from substantial inefficiencies in parallelization and applications involving highly localized states. Here, we discuss recent high-order, enriched, and discontinuous finite element (FE) based methods for the solution of the Kohn-Sham equations which have made possible order-of-magnitude reductions in basis size relative to PW and calculations of over 4000 atoms.

John Pask

Lawrence Livermore National Lab
pask1@llnl.gov

MS21

A Lower Bound Technique for Communication on BSP with Application to the FFT

Communication complexity is defined, within the *Bulk Sym-*

chronous Parallel (BSP) model of computation, as the sum of the degrees of all the supersteps. A lower bound to the communication complexity is derived for a given class of DAG computations in terms of the *switching potential* of a DAG, that is, the number of permutations that the DAG can realize when viewed as a switching network. The proposed technique yields a novel and tight lower bound for the FFT graph.

Gianfranco Bilardi, Michele Scquizzato, Francesco Silvestri

University of Padova
bilardi@dei.unipd.it, scquizza@dei.unipd.it, silvestri1@dei.unipd.it

MS21

Communication Optimal Loop Nests

We derive lower bounds on data movement for a special class of computations that can be modeled as nested-loop programs with affine array references. Our approach reveals a general strategy for tiling iteration spaces to minimize data movement. In many cases, our lower bounds are tight, and this strategy produces optimal tilings.

Michael Christ

UC-Berkeley
christ@math.berkeley.edu

James W. Demmel

University of California
Division of Computer Science
demmel@cs.berkeley.edu

Nicholas Knight

UC Berkeley
knight@cs.berkeley.edu

Thomas Scanlon

University of California
Department of Mathematics
scanlon@math.berkeley.edu

Katherine Yelick

University of California, Berkeley
Computer Science Division
yelick@cs.berkeley.edu

MS21

Tradeoffs between Synchronization, Communication, and Work in Parallel Linear Algebra Computations

The execution time of any parallel algorithm can be inferred by the longest sequence of computations F , data transfers W , and synchronizations S in the communication schedule. We use graph expansion analysis to show that for certain computations with the dependency structure of a d -dimensional symmetric mesh with n^d vertices, $F \cdot S^{d-1} = \Omega(n^d)$ and $W \cdot S^{d-2} = \Omega(n^{d-1})$. These graph-theoretic results are applied to obtain communication lower bounds for dense matrix factorizations and Krylov subspace methods, which are matched tightly by certain existing algorithms.

Edgar Solomonik

University of California at Berkeley
solomon@eecs.berkeley.edu

Erin C. Carson, Nicholas Knight, James Demmel
 UC Berkeley
 ecc2z@eecs.berkeley.edu, knight@eecs.berkeley.edu, demmel@eecs.berkeley.edu

Skovhede, Brian Vinter
 Niels Bohr Institute, University of Copenhagen
 madsbk@nbi.dk, safl@nbi.dk, blum@nbi.dk,
 skovhede@nbi.dk, vinter@nbi.dk

MS21

A New Sparse Inertia-Revealing Factorization

The talk will describe a new sparse inertia-revealing factorization that can be used as part of a bisection eigensolver for symmetric matrices. The factorization is based on a dense algorithm due to Wilkinson and Martin and on George and Heath's sparse QR algorithm. The talk will explain the algorithm and how sparsity is preserved by symmetric double separators and will present computational and numerical results that indicate that the algorithm is both effective in a sparse eigensolver and numerically stable (under some mild constraints).

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

MS22

Distributing NumPy Using Global Arrays

Global Arrays is a software library from Pacific Northwest National Laboratory that provides an efficient, portable, and parallel shared-memory programming interface to manipulate distributed dense arrays. NumPy is the de facto standard for numerical calculation in the Python programming language. Leveraging both Global Arrays and NumPy, we reimplemented NumPy as a distributed drop-in replacement called Global Arrays in NumPy (GAiN). Serial NumPy applications can become parallel, scalable GAiN applications with minor source code changes.

Jeff Daily
 Pacific Northwest National Laboratory
 jeff.daily@pnl.gov

MS22

Bohrium: Unmodified NumPy Code on CPU, GPU, and Cluster

In this talk we introduce Bohrium, a runtime-system for mapping array-operations onto a number of different hardware platforms, from simple multi-core systems to clusters and GPU enabled systems. As a result, the Bohrium runtime system enables NumPy code to utilize CPU, GPU, and Clusters. Bohrium integrates seamlessly into NumPy through implicit data parallelization of array operations, in NumPy called Universal Functions. Bohrium requires no annotations or other code modifications beside adding the import statement: `import bohrium as numpy`. In order to couple NumPy with the execution back-end, Bohrium uses an intermediate vector bytecode that it generates based on the NumPy array operations. The execution back-end is then able to execute the intermediate vector bytecode without any Python/NumPy knowledge, which, in principle, makes Bohrium usable for any programming language. Additionally, the intermediate vector bytecode solves the Python *import problem* where the `import numpy` instruction overwhelms the file-system in a supercomputer since with Bohrium only a single node needs to run the Python interpreter, the remaining nodes simply execute the intermediate vector bytecode.

Mads Kristensen, Simon Lund, Troels Blum, Kenneth

MS22

ODIN: Bringing NumPy's Strengths to Distributed Computing

ODIN (Optimized Distributed NumPy) is a new project for distributed array computing with a NumPy-like interface. Its goals are usability and interoperability via the distributed array protocol to make common data-parallel array-oriented programming models and distributed libraries easier to use. ODIN provides capabilities for finite differencing calculations, finite element computations, and general array computing that is supported by NumPy. Several optimizations are on the project's roadmap, focused on ensuring memory-efficient operations.

Kurt W. Smith
 Enthought, Inc.
 ksmith@enthought.com

MS22

PyTrilinos: Parallel Solvers and Simulation Tools for Python

Trilinos is a suite of C++ scientific simulation software packages, developed primarily at Sandia National Laboratories, currently totaling 54 in number. These packages cover distributed linear algebra, direct and iterative linear solvers, preconditioning, nonlinear, continuation and eigenproblem solvers, discretization methods, dynamic load balancing, optimization, unstructured meshing, parallel I/O, and high-level interfaces. One of the reasons Trilinos contains such a large number of packages is that it includes both "first generation" and "second generation" packages that address the same problem. First generation packages were developed at a time when C++ features such as namespaces and templates were not reliably portable, and so commonly templated data, such as scalar or ordinal types, were restricted to double precision and integer, respectively. Second generation packages, under active development, support multiple scalar and ordinal types via templates and generic programming techniques. PyTrilinos is a fully parallel Trilinos package that provides Python interfaces to selected Trilinos packages, largely restricted to first generation. Current development efforts are focused on expanding PyTrilinos to support second generation packages. PyTrilinos also serves as a testbed of algorithms for demonstrating the utility of distributed arrays being developed by the Odin project — optimized distributed NumPy, where NumPy is the ubiquitous (and serial) Numerical Python package. This talk will cover current capabilities, the status of the development of second generation interfaces, and the relationship between PyTrilinos and Odin.

William F. Spatz
 Sandia National Laboratories
 wfspatz@sandia.gov

MS23

A Volume Integral Equation Solver for Boundary Value Problems with Highly Heterogeneous Coef-

ficients

Abstract not available at time of publication.

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

MS23**A Block-Structured Parallel Adaptive Lattice-Boltzmann Method for Rotating Geometries**

We have developed a finite volume based dynamically adaptive Lattice Boltzmann method for large eddy simulation of weakly compressible flows in moving complex geometry. The scheme is implemented as a patch solver within the massively parallel block-structured AMR framework AMROC. We will report on details of the method and first scalability results on several hundreds cores. Technically relevant three-dimensional simulations of wind turbines with interacting flow fields will demonstrate the benefit of the overall approach.

Ralf Deiterding
Oak Ridge National Laboratory
Computer Science and Mathematics Division
ralf.deiterding@dlr.de

Stephen L. Wood
University of Tennessee Knoxville
Bredesen Center
swood19@utk.edu

MS23**AMR for Fluid-Structure and Radiation Problems on Recent and Novel Architectures**

Abstract not available at time of publication.

Qingyu Meng
SCI Institute
University of Utah
qymeng@cs.utah.edu

MS23**Parallel Lattice Boltzmann Methods with Static Adaptivity**

The Lattice Boltzmann Method is structurally an explicit time stepping scheme that requires only nearest neighbor communication. It thus shows good scalability when used on uniform grids. However, when used with adaptive refinement, the method requires to adapt the time step simultaneously with the spatial resolution. Thus also temporal interpolation at refinement boundaries becomes necessary and a severe load imbalance may arise. We will introduce a block-structured grid refinement approach within the WaLBerla LBM software framework combined with a forest of octrees as basic data structure. A static load balancing strategy will be presented that optimizes the work distribution subject to main memory constraints. The scalability of the approach will be demonstrated on Peta-Scale systems.

Florian Schornbaum
University of Erlangen-Nuremberg
florian.schornbaum@informatik.uni-erlangen.de

Ulrich J. Ruede

University of Erlangen-Nuremberg
Department of Computer Science (Simulation)
ruede@informatik.uni-erlangen.de

Harald Koestler
University of Erlangen-Nuremberg
harald.koestler@informatik.uni-erlangen.de

David Staubach
Univ. Erlangen
david.staubach@fau.de

MS24**Apache Giraph: Large-Scale Graph Processing Infrastructure on Hadoop**

Analyzing large graphs provides valuable insights for social networking and web companies in content ranking and recommendations. In this talk, we describe usability, performance, and scalability improvements we made to Apache Giraph, an open-source graph processing system, in order to deploy it at Facebook on graphs of up to a trillion edges. We also describe our additions to the original Pregel model that enable a broader range of production applications and improve code reuse.

Avery Ching
Facebook
aching@fb.com

MS24**Traditional and Streaming MapReduce via MPI for Graph Analytics**

Graph algorithms operating in parallel on distributed graphs often require modest computation and lots of communication. The MapReduce paradigm can sometimes encapsulate the needed communication, either for large graphs stored on disk (batch MapReduce) or graphs that arrive edge-by-edge (streaming MapReduce). I'll discuss two open-source tools and illustrate graph algorithms they support: MR-MPI (mapreduce.sandia.gov) and PHISH (www.sandia.gov/~sjplimp/phish.html). Because they're written on top of MPI, the tools can be used on a variety of HPC platforms.

Steve Plimpton, Karen D. Devine, Timothy Shead
Sandia National Laboratories
sjplimp@sandia.gov, kddevin@sandia.gov,
tshead@sandia.gov

MS24**REEF - Beyond MapReduce by Re-Layering the Big Data Stack**

I will present the recent trend towards resource managers like Apache YARN as the base layer for big data applications. Resource managers enable many applications like SQL, BI, Machine Learning and graph processing to share common resources. I will introduce REEF, which is the next layer on top of resource managers. REEF provides functionality needed by many applications with a special focus on Machine Learning and supports pipelines of disparate jobs to be efficiently build.

Markus Weimer
Microsoft Research

mweimer@microsoft.com

MS24

Large-Scale Numerical Computation Using a Data Flow Engine

As computer clusters scale up, data flow models such as MapReduce have emerged as a way to run fault-tolerant computations on commodity hardware. Unfortunately, MapReduce is limited in efficiency for many numerical algorithms. We show how a new data flow engine, Spark, enables much faster iterative and matrix computations, while keeping the scalability and fault-tolerance properties of MapReduce. Spark is open source in the Apache Incubator and has a growing user community with 25 companies contributing.

Matei Zaharia

MIT

matei@mit.edu

MS25

Parallel Methods for Bayesian Network Structure Learning

This talk will focus on parallel exact and heuristic algorithms for Bayesian network structure learning. Exact learning is NP-hard, limiting its use to smaller problem sizes. I will first present a work and space optimal parallel algorithm for exact structure learning, and its extension to the case of restricted node in-degree. This will be followed by a parallel heuristic structure learning algorithm that can scale to larger networks, while retaining close to optimal structure learning.

Srinivas Aluru

Iowa State University

aluru@cc.gatech.edu

MS25

Parallel Algorithms for Point-Correlation Functions

Abstract not available at time of publication.

William March

Georgia Institute of Technology

march@ices.utexas.edu

MS25

Parallel Algorithms For Nearest Neighbor Searches

Abstract not available at time of publication.

Bo Xiao

School of Computational Science and Engineering

Georgia Institute of Technology

boxiao@gatech.edu

MS26

Automating Performance Optimization for Compute Nodes of Hpc Systems

Application performance optimization based on runtime measurement and analysis has four phases: measurement, analysis and diagnosis, recommendation of optimizations and implementation of the recommended optimizations.

PerfExpert completely automates a subset of performance optimizations (currently only local to compute nodes) and return the user an optimized version of her/his application code. This lecture will present and illustrate how each PerfExpert currently implements each phase of automating performance optimization and show case studies of its application.

James C. Browne

Dept. of Computer Sciences

University of Texas-Austin

browne@cs.utexas.edu

Leonardo Fialho

Texas Advanced Computer Center, UT-Austin

leonardofialho@gmail.com

MS26

Performance Analysis of Mpi+openmp Programs on Scalable Parallel Systems

Abstract not available at time of publication.

John Mellor-Crummey

Computer Science Dept.

Rice University.

johnmc@rice.edu

MS26

Customizing Libraries with Dsls and Autotuning

Abstract not available at time of publication.

Boyana Norris

Argonne National Laboratory

norris@mcs.anl.gov

MS26

Mummi: A Modeling Infrastructure for Exploring Power and Execution Time Tradeoffs

Abstract not available at time of publication.

Valerie Taylor

Texas A&M University

taylor@cs.tamu.edu

MS27

Development of a Mini-Driver Application to Test Fem Assembly on Modern Architectures

With manycore technology, unprecedented FLOP counts will be available to the HPC community. Much of the manycore research has gone into developing efficient linear solvers using specific programming models like Cuda or OpenCL. However, the process of assembling the linear system is often overlooked. This talk focuses on efficient assembly techniques on manycore architectures using the Kokkos abstraction layer. Results will be given for assembling fully coupled implicit Navier-Stokes systems on modern hardware.

Matthew Bettencourt

Sandia National Laboratories

US Air Force Research Laboratory

mbetten@sandia.gov

Eric C. Cyr

Scalable Algorithms Department
Sandia National Laboratories
eccyr@sandia.gov

MS27**Portable Manycore Sparse Linear System Assembly Algorithms and Performance Tradeoffs**

Two strategies for lock-free assembly of a sparse linear system from a manycore parallel finite element computation are explored on NVIDIA Kepler GPU and Intel Xeon Phi MIC with a portable mini-application. First, concurrent computations perform random-access atomic sums of the linear system coefficients. Second, coefficient contributions are saved in a conflict-free temporary array and then gathered for summation into the sparse linear system by threads which have exclusive access to coefficients.

H. Carter Edwards

Sandia National Laboratories
hcedwar@sandia.gov

MS27**Open-Source, Multi-Physics Finite Element Analysis Using LibMesh and Grins**

This work describes the development of an open-source framework, based on the libMesh finite element library, for analysis of multi-physics problems. We explain the design of the FEMSystem framework in libMesh and its mechanisms to facilitate multi-physics simulation prototyping. We discuss the modularization and extensibility of various simulation components in the GRINS application, illustrating with examples. We also discuss enabling coprocessor computations using OpenCL and C++ templates to obtain higher performance in hybrid computing environments.

Roy Stogner

University of Texas at Austin
roystgnr@ices.utexas.edu

Paul T. Bauman
ICES

The University of Texas at Austin
pbauman@ices.utexas.edu

MS27**Design Paradigms to Accommodate Architectural Uncertainty in Multiphysics Applications**

This talk describes graph-based simulation software that relies on hierarchical descriptions of the problem to expose and exploit parallelism at multiple levels from coarse-grained distributed (MPI) parallel to fine-grained task and data parallelism suitable for multicore and manycore (e.g., GPU) architectures. We also will discuss how this abstraction not only accommodates present and emerging architectures, but also naturally handles the complexity that plagues many other design approaches.

James C. Sutherland

Department of Chemical Engineering
The University of Utah
James.Sutherland@utah.edu

Christopher Earl
University of Utah
cwearl@cs.utah.edu

Tony Saad
The Institute for Clean and Secure Energy
Department of Chemical Engineering, University of Utah
tony.saad@utah.edu

Matthew Might
The University of Utah
might@cs.utah.edu

MS28**Tensor Hypercontraction and Graphical Processing Units for Electronic Structure and Ab Initio Molecular Dynamics**

It has long been suspected that the amount of useful information in the two-electron repulsion integrals was far less than it would appear at first glance. A number of approaches such as the resolution-of-the-identity (also known as density fitting), pseudospectral, and Cholesky factorization methods have been developed as ansatzes to enforce compact and efficient representations of the Coulomb operator. We show that a more flexible formulation exists (tensor hypercontraction), uncovering hidden structure in the fourth-order two-electron repulsion integral tensor. This tensor hypercontraction approach reduces the scaling of many electronic structure methods by one to two powers of the molecular size. We further show that the same ideas can be applied to wavefunction amplitudes, leading to an $O(N^4)$ scaling coupled cluster single and double excitation method. Finally, we show that the tensor hypercontraction scheme is ideally suited to GPUs. This is joint work with Edward Hohenstein, Robert Parrish, Sara Kokkila, Ivan Ufimtsev and Nathan Luehr.

Todd Martinez

Department of Chemistry
Stanford University
toddmartinez@gmail.com

MS28**Qc4**

Abstract not available at time of publication.

J Ramanujam

Dept. of Electrical and Computer Engineering
Louisiana State University
jxr@ece.lsu.edu

MS28**Distributed Contraction of Symmetric Tensors**

Tensor contractions constitute a computationally significant kernel for many high-accuracy methods in quantum chemistry, such as coupled cluster methods. The tensors in these methods typically have significant degrees of symmetry/anti-symmetry among subsets of dimensions and are therefore represented in a compact form that minimizes redundancy in the storage of the tensors. Production parallel quantum chemistry suites such as NWChem implement distributed contraction algorithms for symmetric tensors, where dynamic load balancing is achieved but structure within the contraction algorithm is not exploited for communication optimization. This talk presents a framework for load-balanced communication optimization for distributed contraction of symmetric tensors.

P. Sadayappan

Dept. of Computer and Information Science
Ohio State University
saday@cis.ohio-state.edu

Samyam Rajbhandari, Akshay Nikam, Kevin Stock,
Pai-Wei Lai
Ohio State University
rajbhand@cse.ohio-state.edu, nikam.5@osu.edu,
stockk@cse.ohio-state.edu, laip@cse.ohio-state.edu

Sriram Krishnamoorthy
Pacific Northwest National Laboratory
sriram@pnnl.gov

MS28

Losing the Barriers: Increased Performance and Parallelism through Control and Data Flow

The chemistry community has been very successful in developing algorithms that scale well on terascale computers. However, very few applications are available and robust enough for use on petascale (let alone proposed exascale) computers. There are multiple challenges including increasing the level of parallelism in the applications and the management of program control and data flow. Another challenge for chemistry applications is that synchronization barriers can cause the full algorithm to scale poorly due to load balance issues. By careful expression and analysis of the control and data flow, several of these synchronizations can be decreased or removed. This talk will focus on expressing a higher level of parallelism within multiple chemistry algorithms and managing the complexity.

Theresa Windus
Pacific Northwest National Laboratory
P.O. Box 999, K8-91, Richland, WA 99352 USA
twindus@iastate.edu

MS29

Overview of Distributed Dense Linear Algebra over StarPU Runtime

In this talk, we give an overview of the Matrices Over Runtime Systems at Exascale (MORSE) project. The goal of the project is to design dense and sparse linear algebra methods that achieve the fastest possible time to an accurate solution on large-scale distributed multicore systems with GPU accelerators, using all the processing power that future high end systems can make available. In this context, we will present recent studies over the StarPU runtime for distributed heterogeneous architectures on dense linear algebra.

Mathieu Faverge
U. of Tennessee - Knoxville
mathieu.faverge@inria.fr

Emmanuel Agullo
INRIA
emmanuel.agullo@inria.fr

MS29

Parallel LU Factorizations on Intel MIC in AORSA

We describe the performance of a parallel dense matrix solver in AORSA, the All ORders Spectral Algorithm fusion application for modeling the response of plasma to

radio frequency waves in a tokamak device, which takes advantage of acceleration on Intel Many Integrated Core (MIC) and is compatible with ScaLAPACK LU factorization routines. A left-looking out-of-core algorithm factors matrices that are larger than the available device memory. We report on the challenges on porting this solver from GPU to Intel MIC.

Ed D'Azevedo, Judith Hill
Oak Ridge National Laboratory
dazevedoef@ornl.gov, hilljc@ornl.gov

MS29

A Performance Study of Solving a Large Dense Matrix for Radiation Heat Transfer Using Intel Xeon Phi Coprocessors

This work adapts out-of-core algorithms for parallel dense matrix factorization for solving large problems on the Intel Many Integrated Core Architecture (MIC). Dense matrix computations arises in diverse applications such as in modeling the response and heating of fusion plasma to radio frequency (RF) waves, modeling radiation heat transfer, and large scale linear least squares problems. The out-of-core algorithms have the advantage solving problems larger than amount of device memory on MIC by transferring data residing in the host memory of a compute node. The dense factorization uses a column-panel oriented left-looking algorithm and right-looking algorithm for in-core computations. Similar to compute on GPU, the performance of the developed solver is limited by the cost of communication between the device and host memory. A number of performance enhancing techniques will be presented and discussed. The results obtained on MIC will be examined and compared to that obtained from GPU.

Kwai L. Wong
Joint Institute for Computational Science
University of Tennessee/ORNL
wong@jics.utk.edu

Ed D'Azevedo
Oak Ridge National Laboratory
dazevedoef@ornl.gov

Shiquan Su
University of Tennessee
ssu2@utk.edu

Ki Sing Chan
Chinese University of Hong Kong
ss.sing@hotmail.com

Yue Qian
Chinese Univesity of Hong Kong
chihyukh@gmail.com

MS29

Optimized GPU Kernels for Sparse Factorization

Dense matrix computational kernels are often used as important building blocks in many sparse linear solvers. In this talk, we investigate the potential of developing and optimizing such computational kernels on a GPU and present their effects on the performance of an existing distributed-memory sparse factorization code.

Ichitaro Yamazaki
UTK

ic.yamazaki@gmail.com

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

MS30

Scaling Up Python with mpi4py

We present an update on mpi4py, a set of full-featured Python bindings for MPI. We discuss the new functionality in MPI-3, and how these features can be used in mpi4py on MPI implementations that conform to the new standard. We also discuss performance challenges in loading and running Pythonic codes on large-scale supercomputers, and discuss both static and dynamic approaches to mitigating them.

Aron Ahmadi

US Army Engineer Research and Development Center
aron@ahmadi.net

Lisandro Dalcin
Centro Int. de Métodos Computacionales en Ingeniería
dalcinl@gmail.com

MS30

Using the IPython Notebook for Reproducible Parallel Computing

The IPython Notebook is an open-source, web-based interactive computing environment that allows users to author documents that combine live code, text, equations, figures and videos. These documents can be version controlled, converted to static HTML/PDF and shared with others. These features open the door for making scientific computing reproducible. By leveraging IPython's builtin parallel computing support, parallel computations can be reproducibly debugged, monitored, run and documented on everything from multicore laptops to supercomputers.

Brian E. Granger

Cal Poly, San Luis Obispo
bgranger@calpoly.edu

Matthias Bussonnier
Physical Chemistry
Institut Curie, UMR168, Paris, France
bussonniermatthias@gmail.com

Jonathan Frederic
Physics
Cal Poly State University, San Luis Obispo, CA
jon.freder@gmail.com

Bradley Froehle
Mathematics
UC Berkeley, Berkeley, CA
bfroehle@math.berkeley.edu

Paul Ivanov, Thomas Kluyver
Brain Imaging Center
UC Berkeley, Berkeley, CA
pivanov314@gmail.com, takowl@gmail.com

Evan Patterson
Enthought, Inc.

Austin, TX
epatters@enthought.com

Fernando Perez
Helen Wills Neuroscience Institute
University of California, Berkeley
Fernando.Perez@berkeley.edu

Min Ragan-Kelley
UC Berkeley AS&T
MinRK@Berkeley.edu

Zachary Sailer
Cal Poly State University
San Luis Obispo, CA
zachsailer@gmail.com

MS30

IPYTHON.PARALLEL, Tools for Interactive Parallel Computing

IPython provides some tools for simple interactive parallel computing. By abstracting the REPL model of execution over ZeroMQ sockets, multiple interactive sessions can be coordinated easily in a multiplexing or load-balanced manner. Engines can be local on a multicore machine, in a cluster via MPI or PBS, distributed across the internet, or even a combination thereof. The architecture and interface will be presented, along with some performance benchmarks.

Min Ragan-Kelley
UC Berkeley AS&T
MinRK@Berkeley.edu

Brian E. Granger
Cal Poly, San Luis Obispo
bgranger@calpoly.edu

Fernando Perez
Helen Wills Neuroscience Institute
University of California, Berkeley
Fernando.Perez@berkeley.edu

MS30

Going from a Python Embedded DSL to a Massively Parallel Heterogeneous AMR CFD Code

Parallel scientific simulations manage an ever growing number of complex software systems and architectures. In most software disciplines one can abstract away the complexities and provide generic systems; unfortunately scientific computing tends to require substantial input throughout making such abstraction difficult. We present a layered approach using Python to create a DSL via the Ignition project. Ignition generates codes driving Riemann solvers in a highly scalable heterogeneous AMR simulation utilizing ForestClaw, p4est, and ManyClaw.

Andy R. Terrel
Institute for Computational Engineering and Sciences
University of Texas at Austin
aterrel@ices.utexas.edu

Kyle T. Mandli
University of Texas at Austin
ICES
kyle@ices.utexas.edu

Donna Calhoun
Boise State University
donnacalhoun@boisestate.edu

Carsten Burstedde
Universität Bonn
burstedde@ins.uni-bonn.de

MS31

Parallel, Adaptive Finite Volume Method for Solving Conservation Laws on Mapped, Multiblock Domains

We present our current efforts to develop ForestClaw, a scalable, adaptive algorithm for solving hyperbolic conservation laws on mapped, multiblock domains. Each block is a tree of fixed size, non-overlapping grids, chosen to adaptively refine the solution features in that block. For our grid management system, we use p4est (C. Burstedde, Univ. of Bonn), a highly scalable tree-code designed for AMR on multiblock domains. Results verifying the code and benchmark tests used from atmospheric sciences will be shown.

Donna Calhoun
Boise State University
donnacalhoun@boisestate.edu

Carsten Burstedde
Universität Bonn
burstedde@ins.uni-bonn.de

MS31

Parallel, Adaptive, Multilevel Solution of Nonlinear Systems Arising in Phase Field Problems

We describe the parallel implementation of an adaptive multigrid solver for nonlinear algebraic systems arising from the discretization of phase-field models for non-isothermal alloy solidification in three space dimensions. Even with mesh adaptivity the number of degrees of freedom requires the use of thousands of cores in order to be able to solve problems to sufficient accuracy. We demonstrate both the accuracy and the efficiency of the proposed approach.

Peter Bollada, Peter K. Jimack
School of Computing
University of Leeds
p.c.bollada@leeds.ac.uk, p.k.jimack@leeds.ac.uk

Andrew Mullis
University of Leeds
a.m.mullis@leeds.ac.uk

MS31

Parallel Strategies for Modeling Storm Surge With Adaptive Mesh Refinement

Coastal hazards related to strong storms, in particular storm surge, are one of the most frequently recurring and wide spread hazards to coastal communities. Efforts to apply adaptive mesh refinement techniques to storm surge events has seen a resultant drastic reduction in computational cost but this drastic may still not be sufficient. We will discuss application of new parallel strategies and new many-core architectures to storm surge numerical models,

in this case GeoClaw.

Kyle T. Mandli
University of Texas at Austin
ICES
kyle@ices.utexas.edu

Marsha Berger
Courant Institute of Mathematical Sciences
New York University
berger@cims.nyu.edu

Clint Dawson
Institute for Computational Engineering and Sciences
University of Texas at Austin
clint@ices.utexas.edu

MS32

Low-Communication Multigrid, with Applications to Time-Dependent Adjoint, in-Situ Visualization, and Resilience

A radical formulation of FAS multigrid can remove "horizontal" communication, leading to a less synchronous algorithm that is attractive for hierarchical architectures and allows compression of a fine-grid solution from data only at the coarse grid. This compressed representation can be recovered up to discretization error using only local processing, enabling local reanalysis for visualization, lower-memory checkpointing for adjoints of nonlinear problems, and local recovery using only checkpointed coarse state.

Jed Brown
Mathematics and Computer Science Division
Argonne National Laboratory
jedbrown@mcs.anl.gov

Mark Adams
Columbia University
mfadams@lbl.gov

MS32

Finite Difference Stencils Robust to Silent Data Corruption

Platform design issues regarding power, heat, and error-correction costs may lead to silent data corruption (SDC) in extreme-scale computing. PDE solvers can be rendered SDC-resilient if the discretization is self-filtering—a criterion linking digital computation with PDE physics. We present an SDC-resilient second-order finite-difference stencil for hyperbolic problems with discontinuities. We show how the stencil filters out isolated data-corruptions while still recognizing and capturing shocks. The degradation of accuracy with data-corruption rate increases is explored.

Jaideep Ray
Sandia National Laboratories, Livermore, CA
jairay@sandia.gov

Jackson Mayo, Rob Armstrong
Sandia National Laboratories
jmayo@sandia.gov, rob@sandia.gov

MS32

Experimentally Tuned Algorithm-Based Fault Tolerance Techniques for Matrix Multiplication and

FFT on GPUs

Radiation experiments performed at ISIS, UK and at LAN-SCE, Los Alamos, NM, USA demonstrate that GPUs are very prone to be corrupted by neutrons. Moreover, both Matrix Multiplication and FFT experience, in the majority of the cases, multiple errors in the output. Experimental results and algorithm analysis can be fruitfully employed to design optimised Algorithm-Based Fault Tolerance strategies and provide pragmatic programming guidelines to increase the parallel code reliability with low computational overhead.

Paolo Rech, Luigi Carro
Federal University of Rio Grande do Sul
prech@inf.ufrgs.br, carro@inf.ufrgs.br

MS32**An Algorithmic Approach to Silent Error Resilience**

We propose a paradigm for detecting errors at the application level: compare, frequently, the primary scheme against a cheaper, checking scheme. We provide, as examples, checking schemes for time-dependent problems. We give examples, for ODEs (Runge-Kutta and Adams methods) and for PDEs. In tests on the heat equation and Navier-Stokes equations, the method effectively reveals errors that would otherwise have been silent, without significant slowdown.

Rob Schreiber
HP
rob.schreiber@hp.com

Austin Benson
UC Berkeley
arbenson@stanford.edu

MS32**Reexamining Algorithm-Based Fault Tolerance for Exascale Architectures**

Most Algorithm-Based Fault Tolerance (ABFT) designs have not considered diverse hardware resilience mechanisms that may be available in future Exascale architectures. As a result, some data structures are over-protected by both ABFT and hardware, which leads to unnecessary costs in terms of performance and energy. We reexamine ABFT with an integrated view that includes mechanisms in both software and hardware with the goal of improving performance and energy efficiency of ABFT-enabled applications.

Jeff Vetter, Dong Li
Oak Ridge National Laboratory
vetter@ornl.gov, lid1@ornl.gov

MS33**Programming Paradigms for Emerging Architectures Applied to Asynchronous Krylov Eigensolvers**

Programming present and future supercomputers require using efficiently every levels of parallelism: from fine grain multi threading, at the many-core level, to the coarse grain task management. We need multi-dimensional programming paradigms in order to incorporate the whole spectrum of parallelism. In this talk we address this challenge for a

Krylov eigensolver ERAM/MERAM using new programming models; integrating YML, a graph of task oriented programming framework, and the PGAS-like language XMP. We describe the optimisation of the fine grain parallelisation of an Arnoldi orthogonalisation process (both multithreading and SIMD) and, then, of the medium parallelisation for the associated Krylov eigensolver. Finally, we use a coarse grain task parallelism with asynchronous communications for the asynchronous co-method oriented MERAM. We present results using standard approaches (MPI and OpenMP/TBB/Cilk+). Then, we analyze the difficulties and limits of existing programming paradigms. We highlight the needs of new programming paradigms that facilitate the expression and management for each parallelism level.

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

France Boillod-Cerneux
CNRS/LIFL
France
boillod.france@gmail.com

Fan Ye
CEA/DEN/DANS/DM2S
Maison de la Simulation
fan.ye@cea.fr

Hervé Galicher
CEA/DEN/DANS/DM2S - CEA Saclay
hervegalicher@yahoo.fr

Serge Petiton
CNRS/Lille
serge.petiton@lifl

MS33**Toward a Portable, Resilient Application Design for Scalable Manycore Computing Systems**

We are now several years into the exploration of manycore and accelerator processors, the scalable systems composed of the processors and the software that must run on these systems. In this talk we present some guidelines and principles that can be helpful when designing next-generation algorithms, applications and libraries for these systems. Topics will include strategies for designing future-scalable algorithms, issues of memory and computation organization and models for resilience.

Michael A. Heroux
Sandia National Laboratories
maherou@sandia.gov

MS33**Preconditioned Iterative Solvers on Manycore Architectures**

GPU and MIC(Xeon Phi) are widely used as manycore processors. We are interested in the optimization of preconditioned iterative solvers on these hardware. In this talk, we will show the current status of our implementation and performance.

Satoshi Ohshima
The University of Tokyo

Information Technology Center
ohshima@cc.u-tokyo.ac.jp

Takahiro Katagiri
The University of Tokyo
katagiri@kata-lab.itc.u-tokyo.ac.jp

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

MS33

Achieving Many-core Performance Portability with Kokkos

To achieve optimal memory access patterns on different many-core architectures one must manage both parallel iterations and data access. Kokkos portably implements such an integrated programming model through C++ template meta programming. We present Kokkos' fundamental abstractions and performance results for sparse linear algebra computations on Intel Xeon CPUs, Intel Xeon Phi, and NVIDIA Kepler GPUs. These computations demonstrate equivalent or better performance than native implementations and vendor's current libraries.

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

MS34

A Hierarchical Parallel Software Package of a Complex Moment Based Eigensolver

We have developed a software package which can solve the standard and generalized eigenvalue problems. The algorithm implemented in our package is based on the contour integral and the complex moment and can solve both symmetric and unsymmetric problems. Our package is implemented with Fortran 90 and MPI and allows us to utilize the hierarchical parallel structure of the algorithm. Some numerical examples will be shown.

Yasunori Futamura, Tetsuya Sakurai
Department of Computer Science
University of Tsukuba
futamura@mma.cs.tsukuba.ac.jp,
rai@cs.tsukuba.ac.jp

saku-

MS34

A Parallel Two-grid Polynomial Jacobi-Davidson Algorithm for Large Sparse PDE Eigenvalue Problems

The polynomial Jacobi-Davidson (PJD) algorithm is a subspace method, which extracts the candidate eigenpair from a search space and the space is updated by embedding the solution of the correction equation. We propose the two-grid PJD algorithm for the dissipative acoustic cubic eigenvalue problem. The coarse grid information is used for constructing good initial basis for the search space and a low-cost effective preconditioner for the correction equation. Some numerical examples are given to demonstrate PJDs robustness and scalability.

Feng-Nan Hwang
National Central University

hwangf@math.ncu.edu.tw

Tsung-Ming Huang
National Taiwan Normal University
Department of Mathematics
min@math.ntnu.edu.tw

Weichung Wang
National Taiwan University
Department of Mathematics
wwang@ntu.edu.tw

Yu-Fen Cheng
National Center for High-performance Computing
yfcheng@nchc.narl.org.tw

MS34

Extreme-Scale Parallel Symmetric Eigensolver for Very Small-Size Matrices Using A Communication-Avoiding for Pivot Vectors

We have developed a parallel eigensolver for very small-size matrices. Unlike conventional solvers, our design policy focusses on nature of non-blocking computations and reduced communications. A communication-avoiding approach for Householder pivot vectors is used to implement part of Householder inverse transformation. In addition to that, we implement some techniques for reducing communications by using non-blocking communications in tridiagonalization part. Performance of the solver with full nodes in the Fujitsu FX10 (76,800 cores) is also presented.

Takahiro Katagiri
The University of Tokyo
katagiri@kata-lab.itc.u-tokyo.ac.jp

Junichi Iwata, Kazuyuki Uchida
Graduate School of Engineering
The University of Tokyo
iwata@ap.t.u-tokyo.ac.jp, kuchida@ap.t.u-tokyo.ac.jp

MS34

Xabclib: Parallel Iterative Linear Solver with Run-Time Auto-Tuning

Many matrix solvers have many parameters as inputs by the user. They include parameters that are difficult to set values, therefore approaches of automatically setting them is needed. In this presentation, we will present a Run-time Auto-tuning method for selecting parameters of parallel iterative linear solver. In addition, we will show results of performance evaluation with FX-10 supercomputer.

Takao Sakurai
Central Research Laboratory
Hitachi Ltd.
takao.sakurai.ju@hitachi.com

Takahiro Katagiri
The University of Tokyo
katagiri@kata-lab.itc.u-tokyo.ac.jp

Mitsuyoshi Igai
Hitachi ULSI Systems Corporation
mitsuyoshi.igai.bf@hitachi.com

Satoshi Ohshima
The University of Tokyo

Information Technology Center
ohshima@cc.u-tokyo.ac.jp

Hisayasu Kuroda
Univ. of Tokyo / Ehime U.
kuroda@cs.ehime-u.ac.jp

Ken Naono
R&D department
Hitachi Asia Malaysia, Malaysia
knaono@has.hitachi.com.my

MS35

Modeling the Performance Repercussions of Hpc Applications on Bandwidth Limited Exascale Systems

Projected designs for Exascale compute nodes call for massive numbers of (simple) cores per processor. An undesirable outcome of such designs is the reduced per core memory bandwidth and its obvious negative performance repercussions on memory intensive HPC workloads. System designers and programmers could benefit immensely if there was a systematic framework that allowed them to identify the types of computations that are sensitive to reduced per core memory bandwidth. Towards that end, we present a framework that combines application and target platform characterizations to understand and predict (via predictive machine learning models) the performance sensitivity of different applications to decreased memory bandwidth and this sensitivity at multiple scales (e.g. application to basic-block level). Understanding this performance impact will allow co-design stakeholders to, for example, re-engineer algorithm implementations to reduce the sensitivity. Our framework fully automates the process of generating the characterization profiles for HPC workloads using state-of-the-art binary analysis and instrumentation tools developed at the PMAc Lab. We evaluate our framework on large scale applications and show that our models are highly accurate (absolute mean error < 5%) in predicting the performance sensitivity of HPC applications.

Laura Carrington
San Diego Supercomputing Center
lcarring@sdsc.edu

MS35

Title Not Available at Time of Publication

Abstract not available at time of publication.

Adolfy Hoisie
Los Alamos National Laboratory
hoisie@lanl.gov

MS35

Title Not Available at Time of Publication

Abstract not available at time of publication.

Joseph Kenny
Sandia National Laboratories
jpkenny@sandia.gov

MS35

Constructing High-Level Application Models for

Exascale Co-Design Simulations

Application-architecture co-design has emerged as a vehicle for exploring the design space of Exascale architectures. The use of application skeletons as has emerged as a candidate macro-scale structural representation of applications representing the core communication and computation properties of the application. However, they rely on accurate high-level models of the timing behavior of applications to augment the application skeleton. We present a methodology and tools that leverage statistical techniques to allow semi-automated construction of these models and require minimal user insight and architectural expertise. We present results using several Exascale mini-applications that leverage programming frameworks for CPU and GPU architectures.

Sudhakar Yalamanchili

The School of Electrical and Computer Engineering
Georgia Institute of Technology
sudha@ece.gatech.edu

MS36

Berkeleygw for Excited States Calculations

Abstract not available at time of publication.

Jack Deslippe
National Energy Research Scientific Computing Center
jrdeslippe@lbl.gov

MS36

Electronic Structure Calculation based on Daubechies Wavelets: BigDFT

BigDFT is a free software code implementation for full electronic structure calculation with a formalism based on Daubechies wavelets in hybrid parallel architectures with graphics processing units. Daubechies wavelets have a number of interesting properties for a basis set being used for electronic structure calculations of complex systems. They form a systematic orthogonal and smooth basis, localized both in real and Fourier space and that allows for adaptivity. Since 2007 the BigDFT code uses this basis for Kohn-Sham Density Functional Theory. The performances of this code meet both the requirements of precision and localization found in many applications. This code may treat traditional and complex environment (e.g. charged systems, electric fields, different boundary conditions...) with a systematic treatment and a mathematically clean description. In this presentation we will revisit the present status of the developments and activities around the BigDFT project (see <http://bigdft.org>). In particular, the discussion will be focused on the advantages that the usage of Daubechies wavelets basis set presents in view of the implementation of a flexible DFT approach. We will describe the developments, carried out in the last two years, related to the usage of wavelet properties for treating large and very large systems. Adaptive localized functions expressed in wavelets are used to define a contracted basis set able to describe Kohn-Sham orbitals and related quantities with the same accuracy of a systematic, unbiased approach.

Laura Ratcliff
Institut Nanosciences et Cryogénie, CEA
laura.ratcliff@cea.fr

Stephan Mohr

University of Basel
stephan.mohr@unibas.ch

Paul Boulanger
Institut Néel, CNRS
paul.boulanger@grenoble.cnrs.fr

Luigi Genovese, Damien Caliste
Institut Nanosciences et Cryogénie, CEA
luigi.genovese@cea.fr, damien.caliste@cea.fr

Stefan Goedecker
Univ Basel, Switzerland
Stefan.Goedecker@unibas.ch

Thierry Deutsch
Institut Nanosciences et Cryogénie
thierry.deutsch@cea.fr

MS36

Using Siesta to Solve Large-Scale Electronic Structure Problem

We describe performance improvement achieved in the widely used SIESTA code for Kohn-Sham DFT calculations through the use of the pole expansion and selected inversion (PEXSI) technique. This method has the novel feature of reducing computational complexity without sacrificing generality or accuracy. We show by examples that the use of PEXSI in SIESTA extends the range of physical problems accessible to simulation at nanoscale, and substantially improves the usage of modern high performance computing facilities.

Georg Huh
Barcelona Supercomputing Center
georg.huhs@bsc.es

MS36

New Development in Nwchem

Abstract not available at time of publication.

Bert de Jong
Lawrence Berkeley National Lab
wadejong@lbl.gov

MS37

Scalable Multilevel Stokes Solver for Mantle Convection Problems

The Hierarchical Hybrid Grids framework is designed to close the gap between flexibility of Finite Element's and the performance of geometric multigrid's by using a compromise between structured and unstructured grids. Recently, we designed a Stokes solver specifically for Earth Mantle Convection simulations within this framework. In the presentation, we will address the treatment of variable viscosity using a the hybrid parallel multigrid solver. Further we investigate and analyze scalability experiments on peta-scale clusters.

Björn Gmeiner
Computer Science 10 - System Simulation
Friedrich-Alexander-University Erlangen-Nuremberg,
Germany

bjoern.gmeiner@informatik.uni-erlangen.de

MS37

Scalable Solvers for Multi-Phase Flow: Algebraic Multigrid for Discontinuous Galerkin and Accelerator Integration

We present a scalable AMG solver for instationary multi-phase flow with heterogeneous parameter fields and a Discontinuous Galerkin discretization using the frameworks DUNE and PDELab. We show that our implementation scales to large numbers of processes. Furthermore, we investigate a partial hybridization of the algebraic multigrid solver by combining a classical MPI domain partitioning with fat, multi-threaded CPU nodes, GPGPUs and recent accelerator architectures like Xeon Phi.

Steffen Müthing
Institute of Parallel and Distributed Systems
University of Stuttgart
steffen.muething@ipvs.uni-stuttgart.de

MS37

Improving the Performance of Algebraic Multigrid Using Structured Coarse Grids

We consider structured coarse grids within an AMG cycle. In particular, we employ an AMG energy minimizing framework which allows fairly general sparsity patterns for the grid transfer operators. This freedom is used to guarantee that the resulting Galerkin projection discretizations correspond to structured grids, even if the fine level grid might be unstructured. This leads to less nonzero growth, reduced communication, and ultimately translates into reduced computation times on large scale machines.

Ray S. Tuminaro
Sandia National Laboratories
Computational Mathematics and Algorithms
rstumin@sandia.gov

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

Andrey Prokopenko
Sandia National Laboratories
aprokop@sandia.gov

MS37

Upscaling Multigrid Towards Exascale Computing

Numerical simulation with supercomputers 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. Crucial issues for such strategies are reliability, efficiency, robustness, scalability, usability, and versatility. 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 novel simulation system UG 4 (Unstructured Grids) being presented in the following. In particular, we discuss scalability and show several scaling studies. In the second part of the talk we show the application of these strategies to the simulation of processes from bio-sciences and environmental sciences. In particular we will

show simulation of permeation through human skin, signal processing in neurons as well as computations of density-driven groundwater flow.

Gabriel Wittum

Goethe Center for Scientific Computing
Goethe University Frankfurt
wittum@gcsc.uni-frankfurt.de

MS38

Challenges and Opportunities in Extreme-Scale Application Software Productivity

We will discuss opportunities and challenges of large-scale computational science software development targeting extreme-scale computing platforms. To start the session, an overview of large-scale science application development case studies will be presented. We will also discuss our experiences integrating a complex multi-scale, multi-physics Fortran application with a high-performance C++ framework and libraries. Overall, we call out the need for approaches that enable portable performance as well as support maintainability of large scientific code bases.

Hans Johansen

Lawrence Berkeley National Laboratory
Computational Research Division
hjohansen@lbl.gov

Jeffrey N. Johnson

Applied Numerical Algorithms Group
Lawrence Berkeley National Laboratory
jnjohnson@lbl.gov

Lois C. McInnes

Mathematics and Computer Science Division
Argonne National Laboratory
curfman@mcs.anl.gov

MS38

NSF Activities to Support Software for Next Generation Systems

The National Science Foundation (NSF) recognizes the challenges faced by computational and data-enabled scientists in dealing with next generation computer systems. In response, NSF is promoting and funding software opportunities that include education, training, research, development, and ongoing maintenance and support. This talk will discuss these activities and some representative projects.

Daniel Katz

National Science Foundation
dkatz@nsf.gov

MS38

Runtime Configurability in PETSc

The ability to dynamically configure a deeply nested hierarchy of objects is a key feature in the design of PETSc. This allows the user to assemble an optimal solver tailored to problem characteristics without changing the application code, and also allows easy comparison among competing methods. We will demonstrate the efficacy of this approach for both linear and nonlinear systems. Moreover, we extend this paradigm to encompass residual evaluation

and other key simulation operations.

Matthew G. Knepley

University of Chicago
knepley@ci.uchicago.edu

Barry F. Smith

Argonne National Lab
MCS Division
bsmith@mcs.anl.gov

Lois C. McInnes, Jed Brown, Peter R. Brune

Mathematics and Computer Science Division
Argonne National Laboratory
curfman@mcs.anl.gov, jedbrown@mcs.anl.gov,
brune@mcs.anl.gov

MS38

Software/Science Co-Development: Software Engineering for Large-Scale in Silico Neuroscience Research

Compared to other scientific fields, neuroscience only more recently picked up on the capabilities offered by in silico approaches. While this late start in principle allows learning sustainable practices from other fields, it also means that today software development in neuroscience is mostly done by scientists and not by professional engineers. Large-scale integrative projects such as the Blue Brain and the Human Brain Project can create an organizational and technological environment for sustainable science/software co-development.

Felix Schürmann, Nenad Buncic, Fabien Delalondre, Stefan Eilemann, Jeffrey Muller

Ecole Polytechnique Federale de Lausanne
felix.schuermann@epfl.ch, nenad.buncic@epfl.ch, fabien.delalondre@epfl.ch, stefan.eilemann@epfl.ch, jeffrey.muller@epfl.ch

MS39

Overview of Parallel Mesh Generation and Optimizations Methods

Parallel mesh generation and optimization (PMGO) is a relatively new research area transcending the boundaries of computational geometry and parallel computing. We will review both the theoretical foundation and the practical aspects related to the implementation of PMGO methods on current and emerging architectures. We will organize the PMGO methods in terms of two basic attributes: the sequential techniques used for solving the individual subproblems, and the degree of coupling between the subproblems.

Andrey Chernikov

Old Dominion University
achernik@cs.odu.edu

Suzanne M. Shontz

Mississippi State University
sshontz@math.msstate.edu

Nikos P. Chrisochoides

Computer Science
College of William and Mary

nikos@cs.wm.edu

MS39

Hybrid MPI/openmp Anisotropic Mesh Generation

Abstract not available at time of publication.

Gerard J Gorman

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

MS39

A Parallel Log-Barrier Algorithm for Untangling and Mesh Quality Improvement

In this talk, we discuss our parallel technique for mesh untangling and mesh quality improvement. Our log barrier algorithm untangles meshes and improves the quality of the worst mesh elements in a global manner on distributed machines. We also describe our edge-based coloring algorithm for synchronizing unstructured communication among processes executing the log-barrier mesh optimization algorithm. The algorithm exhibits greater scaling efficiency when compared to an existing parallel mesh quality improvement technique.

Suzanne M. Shontz

Mississippi State University
sshontz@math.msstate.edu

Shankar Prasad Sastry

University of Utah
sastry@sci.utah.edu

MS39

Multicore CPU or GPU Accelerated Geometry Modeling for Proteins

In this talk, we present an efficient computational framework to construct multi-scale models from atomic resolution data in the Protein Data Bank (PDB), which is accelerated by multi-core CPU and programmable Graphics Processing Units (GPU). In addition to density map construction, three modes are also employed and compared during mesh generation and quality improvement to generate high quality tetrahedral meshes: CPU sequential, multi-core CPU parallel and GPU parallel.

Yongjie Zhang

Department of Mechanical Engineering
Carnegie Mellon University
jessicaz@andrew.cmu.edu

MS40

Parallel Algorithms for Prior Functions in Bayesian Inference

Abstract not available at time of publication.

George Biros

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

MS40

Sparse Inverse Covariance Estimation for a Million

Variables

The L1-regularized Gaussian maximum likelihood estimator has been shown to have strong statistical guarantees in recovering a sparse inverse covariance matrix even under high-dimensional settings. However, it requires solving a difficult non-smooth log-determinant program with number of parameters that scale quadratically with the number of Gaussian variables. State-of-the-art methods thus do not scale to problems with more than 20,000 variables. In this talk, I will describe a quadratic approximation method that can solve 1-million dimensional L1-regularized log determinant problems (which would thus have a trillion parameters) on a single computer. In order to do so, we carefully exploit the underlying structure of the problem. Our innovations include (i) a second-order Newton-like method, (ii) division of the variables into free and fixed sets, (iii) a block co-ordinate descent method, and (iv) a memory efficient scheme that approximates key quantities within the algorithm. Even with the latter approximations, the proposed BIGQUIC algorithm can achieve a quadratic convergence rate. Experimental results using synthetic and real application data demonstrate the improvements in performance over other state-of-the-art methods. This is joint work with Cho-Jui Hsieh, Matyas Sustik and Pradeep Ravikumar.

Inderjit S. Dhillon

University of Texas at Austin
inderjit@cs.utexas.edu

MS40

Parallel Algorithms for Sparse Grids

Abstract not available at time of publication.

Christoph Kowitz

TU Munich
kowitz@in.tum.de

MS40

Scalable Algorithms for Non-Negative Matrix Factorization

Abstract not available at time of publication.

Haesun Park

Georgia Institute of Technology
hpark@cc.gatech.edu

MS41

What Krylov Basis Computation for GMRES on Cluster of Accelerators

We analyse some recent experiments on a cluster of accelerators concerning comparison between orthogonal, incompletely orthogonal and non-orthogonal Krylov basis computing, using communication avoiding technics to compute matrix-vector operations. We target different sparse matrices with several compress formats and non-zero element repartitions. We discuss the predictable stability behaviors and the necessary extra computation (such as orthogonalisation or re-orthogonalisation) for each of these methods when used for linear algebra iterative restarted methods such as GMRES.

Langshi Chen

CNRS/Maison de la Simulation
chenlangshi@gmail.com

Serge G. Petiton
CNRS/LIFL and INRIA
serge.petiton@lifl.fr

Leroy A. Drummond
Lawrence Berkeley National Laboratory
ladrummond@lbl.gov

Maxime Hugues
LIFL
maxime.hugues@gmail.com

MS41

Probabilistic Approaches for Fault-Tolerance and Scalability in Extreme-Scale Computing

We present a novel approach for solving PDEs, using a probabilistic representation of uncertainty in the PDE solution due to incomplete convergence and the effect of system faults. Using domain decomposition, the problem is reduced to solving the PDE on subdomains with uncertain boundary conditions. An iterative approach to solve this problem in a resilient and scalable way, using subdomain computations for sampled values of the subdomain boundary conditions, is demonstrated on elliptic systems.

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

Khachik Sargsyan
Sandia National Laboratories
ksargsy@sandia.gov

Francesco Rizzi
Department of Mechanical Engineering and Materials Science
Duke University
francesco.rizzi@duke.edu

Cosmin Safta, Karla Morris
Sandia National Laboratories
csafta@sandia.gov, knmorri@sandia.gov

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

Habib N. Najm
Sandia National Laboratories
Livermore, CA, USA
hnnajm@sandia.gov

MS41

Parallel H-Matrices with Adaptive Cross Approximation for Large-Scale Simulation

We discuss the hierarchical matrices (H-matrices) with adaptive cross approximation (ACA) for large-scale simulation. An improved method of H-matrices with ACA and a set of parallel algorithms applicable to the method are proposed. Applicability of the proposed method is confirmed through numerical experiments on practical simulation, such as electric field analyses and earthquake cycle simulations. Performance of the implementation of the proposed parallel algorithm is evaluated on SMP cluster

systems.

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

MS41

Optimization of Communications/Synchronizations for Preconditioned Iterative Linear Solvers

Overheads by communication and synchronization will be very critical in preconditioned Krylov iterative solvers for sparse matrices on post-peta/exascale systems. In this talk, recent advancement of communication/synchronization reducing algorithms is overviewed, and details of implementations for preconditioned iterative solvers, which are based on Matrix Power Kernel and on multigrid with hierarchical coarse grid aggregation, will be provided. Finally, recent results using more than 65K cores of Fujitsu PRIMEHPC FX10 are also described.

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

MS42

Spiral on ?(K)

We present the automatic synthesis of FFTs for the K computer system with the autotuning system Spiral. In the talk we will overview the Spiral system and its capability to rapidly be retargeted to a large variety of computer systems and architectures. We then discuss the particular challenges of targeting the K computer system. Finally we will present first results and discuss the potential of specialized auto-generated FFT libraries for such large machines.

Franz Frchetti
Department of Electrical and Computer Engineering
Carnegie Mellon University
franzf@ece.cmu.edu

MS42

A Performance Model based Approach to Auto-tuning Tall and Skinny QR Factorizations

Recently, the TSQR algorithm has been proposed for computing the QR factorization of tall and skinny matrices. In this talk, we construct performance models of the TSQR algorithm and other conventional algorithms to automatically find the effectual one for the given problem size and environment. We also consider to automatically tuning the panel size when using the panel blocking. We evaluate the effectivity of our auto-tuning methods on the K computer.

Takeshi Fukaya
Kobe University
Japan
fukaya@people.kobe-u.ac.jp

Yusaku Yamamoto
Kobe University

yamamoto@cs.kobe-u.ac.jp

MS42

Communication Avoiding-hiding and Auto-tuning for Exteme-scale Eigensolver

We have developed a highly scalable parallel eigensolver for large-scale dense symmetric matrices, EigenExa. The solver works excellently on the K computer. However, the communications between the processes inhibit the performance. In this talk, we address the algorithm and the implementation for EigenExa, and especially discuss the communication avoiding methods for the algorithm.

Toshiyuki Imamura

RIKEN Advance Institute for Computational Science
imamura.toshiyuki@riken.jp

Yusuke Hirota

RIKEN AICS
yusuke.hirota@riken.jp

MS42

BCBCG: Iterative Solver with Less Number of Global Communications

On exa-scale parallel computers, global synchronous communication is becoming a bottleneck. To mitigate the performance drop, several communication avoiding (CA) numerical algorithms are considered. In this talk, we present our work on CA-CG method, which has less number of inner products than the conventional CG algorithm (twice per iteration). Our method use Chebyshev polynomial for generating Krylov subspace, and block CG for further enhancements.

Cong Li

the University of Tokyo
licongngu@is.s.u-tokyo.ac.jp

Reiji Suda

Department of Computer Science, The University of Tokyo
reiji@is.s.u-tokyo.ac.jp

Kohei Shimane, Hongzhi Chen

the University of Tokyo
kohei_smn@is.s.u-tokyo.ac.jp,
hongzhi.nkcs@gmail.com

MS43

HPC from Cluster to Accelerator-level

In this contribution we address two extremes of the parallel computing problem landscape. Reverse Time Migration (RTM) is a widely used wave-equation based seismic imaging algorithm that we have implemented on multiple hardware platforms (Intel Xeon, Nvidia GPUs), and under various parallel programming models (MPI, CUDA, OpenCL, OpenMP, OpenACC). Regular memory access and simple operations on very large 3D floating-point arrays dominate the computation, and impressive parallel speedups can be attained. On the other end of the apparent complexity scale, sparse FFT (sFFT) is a novel algorithm that outperforms the FFT for large but sparse data. sFFT parallelization poses a difficult challenge due to its pseudo-random memory access pattern. We report single- and multi-core performance tuning techniques on Intel Xeon

and Nvidia GPUs that are very different from those used for RTM, with moderate success. These examples demonstrate that, depending on the underlying computational problem and hardware, very different performance optimization approaches must be chosen.

Mauricio Araya-Polo

Shell
mauricio.araya@shell.com

Amik St-Cyr

Computation and Modeling,
Shell International E&P, Inc.
amik.st-cyr@shell.com

Ligang Lu

Shell International Exploration and Production
USA
ligang.lu@shell.com

Detlef Hohl

Shell Global Solutions (US)
detlef.hohl@shell.com

MS43

Integrating Power and Energy Models in Optimization Tools

The SCAPE Laboratory at Virginia Tech has been studying the tradeoffs between performance and power for over a decade. We've developed an extensive tool chain for monitoring and managing power and performance in supercomputers. We've also leveraged advanced modeling techniques to optimize performance and power usage in complex high-performance systems. We will discuss the implications of our findings for exascale systems and current research directions with high impact potential.

Kirk Cameron

Department of Computer Science
Virginia Tech
cameron@cs.vt.edu

MS43

Architectural Performance Analysis of Emerging Workloads

Analyzing and predicting the performance of current and emerging workloads is critical to architecting tomorrows future server silicon and systems. In collaboration with teams across Intel and academia, we lead performance characterization on todays systems and analytical modeling of future systems to drive architectural advances that improve the performance of critical server workloads. We will discuss insights from our recent work analyzing several large-scale, I/O-intensive clustered workloads highlighting the potential impact to future server architectures.

Matthew E. Tolentino

Intel
matthew.e.tolentino@intel.com

MS43

Design and Programming of Application-Specific Multi-Core Architectures

In many application domains the amount of data to be processed increases continuously while simultaneously the

processing time has to be reduced. Besides high performance, other constraints, such as low power consumption, a short development time as well as low system costs, play an important role. These contradictory requirements can be solved using heterogeneous multi-core architectures consisting of processors and application-specific accelerators (e.g. FPGAs, GPUs). This talk presents the challenges faced in designing and programming such systems.

Diana Goehringer
Ruhr University Bochum
diana.goehringer@ruhr-uni-bochum.de

MS44

Parallel Pexsi for Electronic Structure Calculations

Abstract not available at time of publication.

Mathias Jacquelin
Lawrence Berkeley National Lab
mjacquelin@lbl.gov

MS44

Density of States and Eigenvalue Counts via Approximation Theory Methods

We describe algorithms for two related and important problems: 1) computing the Density Of States (DOS) of a Hamiltonian; 2) Finding the number of its eigenvalues in an interval. These algorithms are based on approximating various filter functions or distributions by polynomial or rational functions. A common ingredient to all methods is to estimate the trace of a matrix via random sampling. Collaborators: 1) L. Lin, C. Yang, and 2) E. Di Napoli, E. Polizzi.

Yousef Saad
Department of Computer Science
University of Minnesota
saad@cs.umn.edu

MS44

Parallel Electron Transport Calculation Using Green's Function Methodology

We will present numerical methods that allow fast and robust simulations of coupled electro-thermal transport in ultra-scaled nanostructures. For that purpose the simulation capabilities of the existing quantum transport solver OMEN will be improved by adding electron-phonon and phonon-phonon scattering. We will present numerical methods for computing all diagonal elements of both the retarded and lesser Green's function by using a parallel SPIKE factorization method; As part of this research, the selected inversion approach will be extended to complex and general matrices as needed in the NEGF formalism.

Olaf Schenk
USI - Università della Svizzera italiana
olaf.schenk@usi.ch

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

MS44

A Parallel Orbital-Updating Approach for Elec-

tronic Structure Calculations

In this presentation, we will talk about an orbital-based parallelization approach for electronic structure calculations. This approach is based on finite element discretizations and iterative techniques, and permits us to carry out electronic structure calculations in an orbital separation fashion. We will present some basic analysis and several numerical experiments for the approach. This presentation is based on some joint works with Xiaoying Dai, Xingao Gong, and Jinwei Zhu.

Aihui Zhou
Chinese Academy of Sciences
azhou@lsec.cc.ac.cn

MS45

Multigrid for Structured Grids on 100.000s of Cores

While multigrid methods show excellent scalability, the efficiency of multigrid methods can be improved further, especially on modern computer architectures and large parallel machines, like nowadays supercomputers. These improvements include changes in the actual implementation, like machine specific optimizations or hybrid parallelization, as well as changes in the algorithm, like aggressive coarsening and different smoothers. In this talk recent results for structured grid multigrid obtained on large supercomputers will be presented.

Matthias Bolten
University of Wuppertal
Department of Mathematics
bolten@math.uni-wuppertal.de

MS45

Parallel Time Integration with Multigrid

Since clock speeds are no longer increasing, time integration is becoming a sequential bottleneck. In this talk, we present a multigrid method for parallelizing in time that is based on multigrid reduction (MGR) techniques. The advantage of this approach is that it is easily integrated into existing codes, because it simply calls the users existing time-stepping routine. We demonstrate optimality for parabolic equations and illustrate the speedup potential with both performance models and parallel results.

Robert Falgout
Lawrence Livermore National Laboratory
falgout2@llnl.gov

MS45

Inherently Nonlinear Domain Decomposition and Multigrid Methods for Strongly Nonlinear Problems

By decomposing a large-scale nonlinear problem into nonlinear subproblems, inherently nonlinear domain decomposition methods can be created. By design, they allow for a much more subtle adaption of the solution method to the local behavior of the nonlinear problem. However, questions as convergence, convergence control, and choice of the local nonlinear sub-problems arise. In this talk, we discuss the design of such a nonlinear decomposition method, show its global convergence, and present results illustrating its

quantitative behavior.

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

MS45

Extending Strong-Scaling Limits with Parallel Integration in Time

For time-dependent PDEs, concurrency in the temporal direction is a promising way to accelerate existing space-parallel approaches beyond their strong-scaling limits. With the recent development of the parallel full approximation scheme in space and time (PFASST), a spatial and temporal multi-level structure can be exploited in order to increase parallel speedup. In this talk, we will examine various space-time coarsening strategies and present recent results on successful combinations of classical space-parallel solvers with PFASST.

Robert Speck
Jülich Supercomputing Centre
Forschungszentrum Jülich
r.speck@fz-juelich.de

Daniel Ruprecht
Institute of Computational Science
University of Lugano
daniel.ruprecht@usi.ch

MS46

Tools for Change

Scientific applications are in a state of constant change as the science, algorithms, and underlying implementations advance. After an extended period of relative stability in HPC hardware, we are now entering a period of significant flux, diversity, uncertainty, and perhaps divergence. This, in turn, will force additional changes on applications. I will discuss several projects aimed at developing tools and techniques to support a more systematic, programmer-driven approach to transforming and evolving codes to meet the changing needs of the hardware, the algorithms, and the science.

David E. Bernholdt
Oak Ridge National Laboratory
Computer Science and Mathematics Division
bernholdtde@ornl.gov

MS46

What We Have Learned About Using Software Engineering Practices in Computational Science and Engineering

The uniqueness of the computational science and engineering (CSE) domain requires carefully selected and tailored software engineering (SE) solutions that account for the salient characteristics CSE software development and parallel computing environments. This presentation will discuss the findings from a series of case studies of CSE projects, the results of a community-wide survey, and the outcomes of an ongoing workshop series. In addition, the presentation will briefly illustrate two successful interactions between SE and CSE.

Jeffrey C. Carver

Department of Computer Science
University of Alabama
carver@cs.ua.edu

MS46

Software Lifecycle Models in Trilinos

The Trilinos project emerged as a response to needs for rapid algorithms and software development from a loosely-coupled collection of teams and a simultaneous demand for increased software quality requirements. In this presentation we discuss the progression of software lifecycle models and engineering processes we have used, the lessons learned, what is working well now, and our challenges going forward.

Michael A. Heroux
Sandia National Laboratories
maherou@sandia.gov

MS46

Expressive Environments and Code Generation for High Performance Computing

The development and use of a domain-specific language coupled with code generation has proved to be very successful for creating high-level, high-performance finite element solvers. The use of a domain-specific language allows problems to be expressed compactly in near-mathematical notation, and facilitates the preservation of mathematical abstractions. The generation of low-level code from expressive, high-level input can offer performance beyond what one could reasonably achieve using conventional programming techniques. Moreover, development time can be dramatically reduced and everyday HPC made accessible to domain experts. This presentation will summarise some recent developments in automated modelling for HPC, and a range of challenging modelling problems from different fields that use the presented tools will be shown.

Garth Wells
Department of Engineering
University of Cambridge
gnw20@cam.ac.uk

MS47

Scalable Lattice Cleaving

We explore a parallel and scalable reformulation of the guaranteed quality, multi-material, conforming, tetrahedral mesh generation algorithm known as lattice cleaving. The local scope of the algorithm's operations makes the technique highly amenable to parallelization. The main challenge in reformulating rests on decomposing the background mesh domain into suitable subdomains for each algorithmic phase. Such a reformulation of lattice cleaving enables a streaming, out-of-core solution for memory intensive datasets on architectures with limited resources.

Jonathan Bronson
School Of Computing, Scientific Computing and Imaging Inst.
University of Utah
bronson@sci.utah.edu

Jonathon Nooner
Clemson University
jnooner@clemson.edu

Joshua Levine
University of Utah
levinej@clemson.edu

Tom Peterka
Argonne National Laboratory
tom.peterka@gmail.com

Ross Whitaker
School Of Computing, Scientific Computing and Imaging
Inst.
University of Utah
whitaker@cs.utah.edu

MS47**Parallel Algorithms for Overlay Grid Methods**

Overlay grid methods for automatic all-hex mesh generation provide an ideal hands-off and scalable solution for large-scale modeling and simulation. This talk addresses the specific issues encountered and modifications needed for parallel implementation of these algorithms. Geometry construction, smoothing, boundary layer and pillow insertion and serial-parallel consistency are among some of the areas that require special attention for parallel implementation. Recent results on massively parallel cluster machines at Sandia will also be presented.

Steve J. Owen
Sandia National Laboratories
Albuquerque, NM
sjowen@sandia.gov

MS47**Advances in Parallel Unstructured Mesh Adaptation**

This presentation will address two developments of the MeshAdapt procedure that performs generalized parallel mesh modification to satisfy a given anisotropic mesh size field. The first area is the adaption of mixed element boundary layers where different forms of mesh adaptation are used in semi-structured boundary layer and unstructured mesh in the rest of the domain. The second area considers refactoring MeshAdapt to more effectively apply it on new generations of high core parallel computers.

Onkar Sahni
Rensselaer Polytechnic Institute
sahni@rpi.edu

Daniel Ibanez
RPI
dan.a.ibanez@gmail.com

Kedar Chitale
University of Colorado Boulder
kendar.chitale@colorado.edu

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

MS47**Parallel Mesh Generation and Adaptation with****Distributed Geometry**

The ability to execute large-scale parallel simulations requires all components operate effectively in parallel. In the case of unstructured simulations, this requires the mesh generation and adaptation processes operate in parallel on distributed data. A set of tools where meshes can be generated and adapted in parallel on a geometry that can also be distributed will be presented and their use in multiple simulation workflows will be demonstrated.

Mark Beall
Simmetrix, Inc.
mbeall@simmetrix.com

Saurabh Tendulkar
Simmetrix Inc.
saurabh@simmetrix.com

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

MS48**High-performance and High-productivity Semantic Graph Analysis**

Graph theory is used to model large-scale complex systems in various scientific domains. Filtering on attributed semantic graphs enable a broad set of real applications that maintain important auxiliary information through attributes on individual edges and vertices. We achieve high performance in a high-level Python-based graph library (KDT) by utilizing a highly optimized sparse linear algebra based backend (Combinatorial BLAS), and automatically translating analytic queries via selective just-in-time embedded specialization (SEJITS).

Aydin Buluc
Lawrence Berkeley National Laboratory
abuluc@lbl.gov

John R. Gilbert
Dept of Computer Science
University of California, Santa Barbara
gilbert@cs.ucsb.edu

Adam Lugowski
UC Santa Barbara
alugowski@cs.ucsb.edu

Shoab Kamil
Massachusetts Institute of Technology
skamil@cs.berkeley.edu

Leonid Olikier, Samuel Williams
Lawrence Berkeley National Laboratory
loliker@lbl.gov, swwilliams@lbl.gov

MS48**Managing Centrality in Large Scale Graphs**

Closeness and betweenness centrality are global metrics that quantify how important is a given node in the network. When the network dynamics change, the relative importance of the nodes also changes. The existing algorithms are impractical to recompute them from scratch after each

modification. We provide graph theoretical techniques for faster centrality computation, incremental algorithms to update the closeness centrality values upon changes, and a distributed framework which leverages pipelined and replicated NUMA-aware parallelism.

Ahmet Erdem Sariyuce
The Ohio State University
aerdem@bmi.osu.edu

Erik Saule, Kamer Kaya, Umit V. Catalyurek
The Ohio State University
Department of Biomedical Informatics
esaule@bmi.osu.edu, kamer@bmi.osu.edu,
umit@bmi.osu.edu

MS48

A Scalable Querying System for Large-Scale Graph Pattern Enumeration

Graph patterns such as frequent subgraphs, maximal cliques, and dense subgraphs are useful for scientific discoveries in many disciplines including biology, chemistry, and climatology. Traditional methods for finding such patterns in large-scale graphs are often exploratory in nature, leading to repeated queries and redundant computations over the same graph. To improve repeated pattern enumeration queries, we propose a new strategy that lends itself to parallelization and leverages a partially-expanded search space along with knowledge priors.

Steve Harenberg, Rob Seay, Sriram Lakshminarasimhan,
David Boyuka II, Gonzalo bello, Rada Chirkova
North Carolina State University
harenberg@ncsu.edu, rgseay@ncsu.edu,
slakshm2@ncsu.edu, daboyuka@ncsu.edu,
gabellol@ncsu.edu, rychirko@ncsu.edu

Nagiza Samatova
North Carolina State University
Oak Ridge National Laboratory
samatova@csc.ncsu.edu

MS48

Algorithms for Aligning Massive Networks

What is the best way to map the vertices of two graphs to each other so that they overlap in the largest number of edges? Parallel algorithms employing integer programming and matchings provide heuristic solutions to this problem. By using GMT, a framework supporting a global address space over distributed memory machines, we implement these algorithms to align massive graphs (with billions of edges) on a computer with thousands of cores.

Alex Pothén
Purdue University
Department of Computer Science
apothén@purdue.edu

Arif Khan, David F. Gleich
Purdue University
khan58@purdue.edu, dgleich@purdue.edu

Alessandro Morari
Pacific Northwest National Lab
alessandro.morari@pnl.gov

Mahantesh Halappanavar, Daniel Chavarria, Oreste Villa
Pacific Northwest National Laboratory
mahantesh.halappanavar@pnl.gov,
daniel.chavarria@pnl.gov, oreste.villa@pnl.gov

MS49

High-Performance Gpu Kernels for Multifrontal Sparse Factorization

Sparse matrix factorization involves a mix of regular and irregular computation, which is a particular challenge when trying to obtain high-performance on GPUs. We present a sparse multifrontal QR factorization method that meets this challenge, and is up 11x faster than a highly optimized method on multicore CPU. Our method factorizes many frontal matrices in parallel, and keeps all the data transmitted between frontal matrices on the GPU. A novel bucket scheduler algorithm extends the communication-avoiding QR factorization for dense matrices, by exploiting more parallelism and by exploiting the staircase form present in the frontal matrices of a sparse multifrontal method.

Timothy A. Davis
University of Florida
Computer and Information Science and Engineering
DrTimothyAldenDavis@gmail.com

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

Nuri Yeralan
University of Florida
nuriatuf@gmail.com

MS49

Re-Architecting DFT Kernels for Sustainable Performance

Density Functional Theory (DFT) is the most used method to calculate the electronic structure of materials at the atomic level from first principles. However, trends in computer architectures are resulting in low performance and poor scaling of DFT-based codes. Within DFT, plane-wave pseudopotentials prevail; the electronic wave functions are expanded in Fourier components. We will discuss preliminary work towards minimizing communication and memory accesses, and exposing high levels of parallelism in plane-wave DFT codes.

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

Andrew M. Canning, David Prendergast
Lawrence Berkeley National Laboratory
acanning@lbl.gov, dgprendergast@lbl.gov

Jack Deslippe
National Energy Research Scientific Computing Center
jrdeslippe@lbl.gov

MS49

Accelerating CFD Applications Without Accelerator-Specific Programming

Physis is a framework specifically designed for stencil com-

putations with structured grids, which often appears as an important pattern in computational fluid dynamics applications. It employs a small domain-specific language for expressing stencils in a portable way so as to achieve both high productivity and high performance in current and future generations of HPC systems. This talk will present DSL-based optimizations for CFD applications with several real-world case studies.

Naoya Maruyama
RIKEN Advanced Institute for Computational Science
nmaruyama@riken.jp

MS49

Toward Intelligent Krylov-Based Linear Algebra Methods for Future Extreme Programming and Computing

The convergence and global performance of parallel Krylov methods greatly depends on the choice of several parameters, such as the subspace size, which efficiency are difficult to determine in advance. Avoiding communication strategies may increase locality and minimize energy consumption. Auto-tuning strategies are developed to accelerate convergence, minimize storage space, data movements, and energy consumption for Krylov methods. We survey some auto/smart-tuning strategies for correlated criteria as a milestone toward future intelligent Krylov methods for extreme computing. Several experiments on different platforms are proposed.

Serge G. Petiton
CNRS/LIFL and INRIA
serge.petiton@lifl.fr

MS50

Parameter Selection and Prediction to Tune the Performance of Krylov Subspace Methods

Nowadays, many high-end computer simulations rely on the availability of robust numerical algorithms and corresponding scalable software implementations for a variety of computational platforms. Performance scalability and efficiency in today's hardware requires a more careful introspection of the parameters that influence the numerical behavior of an algorithm as well as the programming model and choices made during the parallel implementation of these algorithms. Our research focuses on identifying key performance parameters, at the algorithmic and implementation level, to obtain a better performance characterization of commonly used numerical schemes. In this talk, we focus on the Explicit Restart Arnoldi Method (ERAM) and key parameters that influence its convergence, robustness and scalability. We aim at efficiently studying and managing these parameters to improve the ERAM convergence at run-time. The low overhead of computing and managing these parameters at run-time is also factored in the runtime tuning process. We present some preliminary results from experiments with different restart strategies, subspace sizes on many-core systems.

France Boillod-Cerneux
CNRS/LIFL
France
boillod.france@gmail.com

Leroy A. Drummond
Lawrence Berkeley National Laboratory
ladrummond@lbl.gov

Herve Galicher
CEA/DEN - Maison de la Simulation
hervegalicher@yahoo.fr

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

Serge G. Petiton
CNRS/LIFL and INRIA
serge.petiton@lifl.fr

MS50

Tuning Asynchronous Co-Methods for Large-scale Eigenvalue Calculations

The restarted Krylovs methods are widely used to solve linear algebra problems such as linear system and eigenproblem. Hybrid versions of these methods called multiple restarted Krylovs methods (MRKM) consist in making use of several instances of a RKM (called also co-methods) allowing accelerating their convergence. More specifically, at the end of an iteration, an instance of RKM, exchanges its interesting information with other instances. This process improves the conditions for restarting an iteration, thus speeding convergence at least one of involved RKMs. Here, we will present multiple restarted Krylovs methods and focus on two representatives examples of MRKM, that is, multiple implicitly/explicitly restarted Arnoldi methods (MIRAM/MERAM). We will illustrate the suitability of these methods for extreme scale computing. This is due to their asynchronous communication schemes, their natural load balancing and their multi-level parallelism. We will also present a description of the corresponding parallel algorithms with theoretical performances and their implementations with YML/XMP frameworks.

Nahid Emad
University of Versailles, PRiSM Laboratory
Nahid.Emad@prism.uvsq.fr

Leroy A. Drummond
Lawrence Berkeley National Laboratory
ladrummond@lbl.gov

Miwako Tsuji
University of Versailles
University of Tsukuba
miwako tsuji |tsuji@hpcs.cs.tsukuba.ac.jp|

Makarem Dandouna
University of Versailles
makarem.dandouna@prism.uvsq.fr

MS50

Dynamic Parallel Algebraic Multigrid Coarsening for Strong Scaling

The algebraic multigrid method is a multi-level method that involves calculations with small problems from the original problem, but this often leads to scalability limitations on large parallel machines. However, it is possible to remap the small problems to fewer processes than all running processes after coarser levels are calculated. We propose a process domain coupling method that dynamically decreases the parallelism according to the matrix size

and simultaneously calculates the coarser level matrix.

Akihiro Fujii, Takuya Nomura, Teruo Tanaka
Kogakuin University
fujii@cc.kogakuin.ac.jp, em13014@ns.kogakuin.ac.jp,
teru@cc.kogakuin.ac.jp

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

MS50

The Impact of Workload Strategies on a UPC-based CG

The CG method is central to a number of applications and has recently been proposed as an HPC benchmark that better characterizes computer workloads. Previous works have sought to minimize memory access time in CG using different storage formats for sparse matrix vector multiplications. We examine four different strategies, with different data distributions and independent of storage formats, for a UPC implementation of CG, and how those strategies impact performance and offer autotuning opportunities.

Jorge González-Domínguez
University of A Coruña
na
jgonzalezd@udc.es

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

María Martín
University of A Coruña
na
mariam@udc.es

Juan Touriño
University of A Coruña
na
Spain
juan@udc.es

MS51

Scalability of Sweep Algorithms in Discrete Ordinates Transport

Discrete Ordinates solutions of the linear Boltzmann transport equation are used to model a variety of physical systems. High performance computing is essential for predictive simulations, for which experiments are often unavailable or prohibitively costly. Sweep algorithms are central to standard discrete ordinates transport solution methods. The scalability of these sweep algorithms will determine whether standard transport iterative methods will scale to petascale architectures and beyond. Contrary to the conventional understanding about limited scalability of sweep algorithms, we have developed scheduling rules and parallel performance models which predict that sweep algorithms will scale to petascale and potentially beyond. We compare our parallel performance models to computational results using two distinctly different discrete ordinates transport codes. Using Sequoia, Lawrence Livermore National Laboratory's petascale supercomputer, we achieved excellent

scaling out to 1,572,864 MPI tasks with over 37.5 trillion unknowns and over 71% parallel efficiency, proving that sweep algorithms scale much further than many believed. Our computational results validate our performance models, giving good indication that sweep algorithms will scale well into exascale sized problems.

Teresa S. Bailey
LLNL
bailey42@llnl.gov

MS51

Design of HACC for Extreme-Scale Simulation

The Hardware/Hybrid Accelerated Cosmology Code (HACC), twice an IEEE/ACM Gordon Bell finalist, has been designed from the ground up for scalable performance on many different supercomputer architectures. HACC, which simulates cosmological large-scale structure formation, runs on both hardware-accelerated systems, such as those with GPUs, and non-accelerated systems, featuring state-of-the-art algorithms for efficiently computing particle-particle gravitational interactions. The gravitational force is split between a long-range component, handled using a grid-based algorithm, and short-range component, handled using some combination of direct computation and a hierarchical approximation (tree) algorithm. Combining this with hierarchical time stepping, an efficient concurrency scheme, and a communication-avoiding neighbor-exchange strategy, HACC has run trillion-particle simulations on both the IBM BG/Q and the Cray XK7.

Hal Finkel, Salman Habib
Argonne National Laboratory
hfinkel@anl.gov, habib@anl.gov

Vitali Morozov
ALCF, Argonne National Laboratory
morozov@anl.gov

Adrian Pope, Katrin Heitmann, Kalyan Kumaran, Tom Peterka, Joe Insley
Argonne National Laboratory
pope@anl.gov, heitmann@anl.gov, kumaran@anl.gov,
peterka@anl.gov, insley@anl.gov

David Daniel, Patricia Fasel
Los Alamos National Laboratory
ddd@lanl.gov, pkf@lanl.gov

Nicholas Frontiere
Argonne National Laboratory and University of Chicago
nfrontiere@anl.gov

Zarija Lukic
Lawrence Berkeley National Laboratory
zarija@lbl.gov

MS51

Scalable Adaptive Mesh Refinement

Patch-based dynamic adaptive mesh refinement recurrently operates on global mesh structures, making it challenging to scale. In the past decade, the SAMRAI project has made significant advancements in this area, including binary box tree searches, distributed mesh management, task-parallel clustering and tree partitioning. Other changes used in conjunction with these advancements also

provided critical improvements. I will present these advancements, which culminated in recent runs using 1M cores.

Brian Gunney
Lawrence Livermore National Lab
gunney1@llnl.gov

MS51

Scalability and Performance of a Legacy Coastal Ocean Model

Several algorithmic factors have limited the performance and scalability of the semi-implicit finite-element/finite-volume coastal circulation model SELFE, which is used by oceanographers and engineers for science and management purposes. We report on a refactoring effort involving the use of PETSc linear solvers, new algorithmic choices for time integration, and improved IO, memory access, and vectorization, resulting in a 2-3x improvement on NERSC and XSEDE machines in strong scaling studies and for research and production configurations.

Jesse Lopez
Oregon Health & Science University
lopezj@stccmop.org

Jed Brown
Mathematics and Computer Science Division
Argonne National Laboratory
jedbrown@mcs.anl.gov

MS52

Real-space Electronic Structure on Graphics Processing Units

We discuss the application of graphics processing units (GPUs) to density functional theory (DFT) and time-dependent DFT. Our approach is based on a finite-difference real-space discretization, implemented in the code Octopus. We obtain a sustained performance of up to 90 GFlops for a single GPU, representing a significant speed-up compared to a CPU. Our implementation can outperform a GPU Gaussian basis-set code, showing that the real-space approach is a competitive alternative for DFT on GPUs.

Xavier Andrade
Harvard University
xandradevalencia@fas.harvard.edu

MS52

Improved Iterative Subspace Methods for Large-Scale Response Calculations

Efficient iterative subspace methods for large linear systems and eigenvalue problems are key for large-scale calculations of molecular response and excited state properties. I will present a simple re-formulation of these methods which takes advantage of sparsity in integral-direct matrix-vector operations. In conjunction with linear scaling exchange techniques, three- to fivefold speedups of molecular response calculations are possible.

Filipp Furche, Brandon Krull, Jake Kwon
University of California, Irvine
Department of Chemistry

filipp.furche@uci.edu,
nana0629@gmail.com

bkrull@uci.edu, jh-

MS52

FEAST Applied to DFT and Real-time TDDFT Calculations

FEAST is presented as a black-box parallel eigenvalue solver for solving both the DFT/Kohn-Sham and the real-time TDDFT problems. Recent developments to augment the parallel capabilities of the FEAST algorithm are presented and discussed. As a result, large-scale electronic structure applications can now benefit from a better load balancing along the eigenvalue spectrum, and real-time TDDFT propagations can operate using a more optimal FEAST spectral-decomposition scheme.

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

MS52

Real-Space DFT for Plane-Wave GW/BSE Calculations

Many-body perturbation theory in the GW approximation and Bethe-Salpeter equation uses input from density-functional theory (DFT) calculations to compute excited states of a condensed-matter system. Many parts of a GW/BSE calculation are most efficiently performed in a plane-wave basis. However, real-space implementations of the DFT step are attractive starting points, since they can take advantage of efficient parallelization in domains for large systems, and are well suited physically for finite systems such as molecules or nanostructures. I will discuss the interfacing of a real-space (TD)DFT code (Octopus, www.tddft.org) with a plane-wave GW/BSE code (BerkeleyGW, www.berkeleygw.org), consider performance issues, and present some applications.

David A. Strubbe
Department of Physics, University of California, Berkeley
Materials Sciences Division, Lawrence Berkeley
Laboratory
dstrubbe@mit.edu

MS53

Fine-Grained Parallel Preconditioning

Sparse linear algebra has been a challenge to parallelize efficiently on architectures such as GPUs and Xeon Phi demanding a very high level of concurrency. We introduce new, non-intuitive parallelizations for two important sparse kernels: triangular solves and incomplete factorizations. The algorithms are based on the idea that concurrency can be improved by trading accuracy. This is possible because for preconditioning operations, only approximations are required – the mathematical approximation does not need to be computed exactly. We show that this approach can lead to much improved overall solution times. In contrast to other approaches, this approach also relies on exploiting numerical properties of the problem being solved, not just sparsity structure.

Edmond Chow
School of Computational Science and Engineering
Georgia Institute of Technology

echow@cc.gatech.edu

MS53

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

The optimal matrix storage format for best performance of the sparse matrix-vector multiply (spMVM) operation is highly system-dependent. We show the advantages and shortcomings of different formats on Intel x86, Xeon Phi, and Nvidia GPGPUs, and suggest a common SIMD-friendly sparse matrix format that shows good performance on all platforms and enables efficient spMVM on heterogeneous systems. Appropriate chip-level performance models are used to get insight into limiting factors and the influence of free parameters.

Moritz Kreutzer, Georg Hager
Erlangen Regional Computing Center (RRZE)
moritz.kreutzer@fau.de, georg.hager@fau.de

Gerhard Wellein
Erlangen Regional Computing Center (RRZE), Germany
Gerhard.Wellein@rrze.uni-erlangen.de

MS53

Accelerating Brownian Dynamics Simulations on Intel MIC

Traditional Brownian dynamics (BD) simulation methods with hydrodynamic interactions have high memory requirements due to the explicit construction of a dense mobility matrix. This limits the scalability of these methods on hardware such as GPUs and Intel MIC that have relatively low memory capabilities. In this talk, we present a 'matrix-free' BD method using a particle-mesh Ewald (PME) approach, relying on 3D FFTs, sparse spreading/interpolation, and a Lanczos method for computing Brownian displacements. We discuss the implementation of this method on multicore and manycore architectures, as well as a hybrid implementation that splits the workload between CPUs and Intel Xeon Phi coprocessors.

Xing Liu, Edmond Chow
School of Computational Science and Engineering
Georgia Institute of Technology
xing.liu@gatech.edu, echow@cc.gatech.edu

MS53

Rethinking Multigrid and Sparse Matrix Computations in High Throughput Environments

Algebraic multigrid methods offer a flexible framework for developing a variety of preconditioners. Yet, while these methods rely on sparse matrix computations, the setup and design of a multigrid method does not normally target high throughput computing. In this talk we highlight some central components of an efficient algebraic method, noting several decisions that utilize highly data parallel primitives. In particular, we discuss the sparse matrix-sparse matrix product and its expected efficiencies for various matrices.

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

lukeo@illinois.edu, dalton6@illinois.edu

MS54

An Overview of FASTMath Technology Developments

The FASTMath SciDAC institute is a collaboration among many DOE applied mathematicians focused on the development of algorithms and software for application use on leadership class computing facilities. In this talk, I present an overview of the institute and highlight key advances in the development of high order and adaptive methods, communication reducing algorithms, and software optimized for hybrid parallelism. I also showcase a few examples of the use of these technologies to advance application science in ice sheet modeling, dislocation dynamics, and nuclear physics.

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

MS54

New Developments in Parallel Unstructured Mesh Methods

The Parallel Unstructured Mesh Infrastructure (PUMI) is a distributed mesh data management system capable of handling general non-manifold domains and supporting automated adaptive analysis. Recent PUMI advances include a two-level partition model where message passing is used at the top level and threads at the lower level, and new services to support specific needs of massively parallel meshes of billions of elements such as the construction of a full partition model from minimal mesh data.

Daniel Ibanez
RPI
dan.a.ibanez@gmail.com

Seegyong Seol, Cameron Smith, Ian Dunn
Rensselaer Polytechnic Institute
seols@rpi.edu, smithc11@rpi.edu, dunni@rpi.edu

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

MS54

PHASTA Unstructured Mesh Scalability to 3.1M Processes

Massively parallel computation provides enormous capacity to perform simulations on a time scale that can change the paradigm of how simulations are used by scientists, engineers and other practitioners to address discovery and design. In this context, the strong scalability of a fully implicit unstructured finite element flow solver is presented with an extreme scale application in aeronautics.

Michel Rasquin
Argonne National Lab
mrasquin@alscf.anl.gov

Cameron Smith
Rensselaer Polytechnic Institute
smithc11@rpi.edu

Kedar Chitale
 University of Colorado Boulder
 kendar.chitale@colorado.edu

Seegyong Seol, Benjamin Matthews
 Rensselaer Polytechnic Institute
 seols@rpi.edu, matthb2@scorec.rpi.edu

Jeffrey Martin
 U. Colorado Boulder
 jeffrey.l.martin@colorado.edu

Onkar Sahni, Mark S. Shephard
 Rensselaer Polytechnic Institute
 sahani@rpi.edu, shephard@rpi.edu

Ray Loy
 Argonne Leadership Computing Facility, Argonne
 National Lab.
 rloy@alcf.anl.gov

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

MS54

Chombo-Crunch: High Performance Simulation of Pore Scale Reactive Transport Processes Associated with Carbon Sequestration

We have added a new embedded boundary-algebraic multi-grid formulation to Chombo to simulate pore scale processes in realistic fluid-rock systems. This formulation exploits recent infrastructure development that supports explicit matrix construction and interoperability with PETSc. Here we describe Chombo-Crunch, which combines this technology in Chombo flow and transport solvers with CrunchFlow, a geochemistry module, to model reactive transport in complex geometries. The new capability has been successfully applied to EFRC-NCGC reactive transport experiments for validation studies.

David Trebotich
 Lawrence Berkeley National Laboratory
 treb@lbl.gov

Mark Adams
 Columbia University
 mfadams@lbl.gov

Brian Van Straalen
 Lawrence Berkeley National Laboratory
 Computational Research Division
 bvstraalen@lbl.gov

Anshu Dubey
 University of Chicago
 dubey@flash.uchicago.edu

Terry J. Ligoeki, Dan Graves
 Lawrence Berkeley Laboratory
 TJJligoeki@lbl.gov, dtgraves@lbl.gov

MS55

High Scalability of Lattice Boltzmann Simulations with Turbulence Models Using Heteroge-

neous Clusters Equipped with GPUs

Our main focus is high scalability on heterogeneous target platforms such as clusters equipped with hundreds or thousands of GPUs. In particular, we are developing a code which implements the Lattice-Boltzmann method augmented by turbulence models. To exploit heterogeneous systems in an optimal way, tailored communication schemes between the computation entities are used to hide communication through computation. A flexible software design of our LBM code enables the easy adaptation to different heterogeneous cluster setups.

Christoph Riesinger
 Technical University of Munich
 riesinge@in.tum.de

MS55

The waLBerla Framework: Multi-physics Simulations on Heterogeneous Platforms

To successfully utilize GPU clusters like Tsubame 2.0 for daily business of a large community, usable software frameworks have to be established on these clusters. The development of such software frameworks is only feasible with maintainable software designs that consider performance and portability as a design objective right from the start. We present software concepts for efficient and scalable multi-GPU parallelization approaches and show how they work in our software framework waLBerla for multi-physics simulations.

Ulrich J. Ruede
 University of Erlangen-Nuremberg
 Department of Computer Science (Simulation)
 ruede@informatik.uni-erlangen.de

Harald Koestler
 University Erlangen-Nürnberg
 harald.koestler@fau.de

MS55

Parallel Computing of Finite-Volume Solver for Euler Equation using a Ghost-cell Immersed Boundary Method using Multiple Graphics Processor Units

We propose an explicit cell-centered finite-volume solver for the Euler equation using a ghost-cell immersed boundary method on multiple graphics processor units. The solver was thoroughly validated with previous simulations. The results show that the speedup can exceed 100 on a single GPU (Nvidia Tesla M2070) compared to a thread of Intel Xeon X5472. Finally, more than 12 times of speedup is obtained using 16 GPUs compared to a single GPU for 9 million cells.

Chongsin Gou
 National Chiao Tung University, Taiwan
 chongsin@faculty.nctu.edu.tw

Jong-Shinn Wu
 National Chiao Tung University
 chongsin@faculty.nctu.edu.tw

MS55

Physics-based Seismic Hazard Analysis on Petas-

cale Heterogeneous Supercomputers

A highly scalable and efficient GPU-base finite-difference code AWP for earthquake simulation is implemented with high throughput, memory locality, communication reduction and communication/computation overlap. A staggered scheme on computational stencils is used. Linear scalability is achieved on Cray XK7 Titan at ORNL and NCAs Blue Waters systems. A speedup by a factor of 110 in key calculations critical to probabilistic seismic hazard analysis is observed, enabling a statewide hazard model achievable with existing supercomputers.

Heming Xu
San Diego Supercomputer Center
h1xu@ucsd.edu

MS56

Large-Scale Metagenomic Sequence Clustering Via Maximal Quasi-Clique Enumeration

This talk will focus on parallel exact and heuristic algorithms for Bayesian network structure learning. Exact learning is NP-hard, limiting its use to smaller problem sizes. I will first present a work and space optimal parallel algorithm for exact structure learning, and its extension to the case of restricted node in-degree. This will be followed by a parallel heuristic structure learning algorithm that can scale to larger networks, while retaining close to optimal structure learning.

Srinivas Aluru
Iowa State University
aluru@cc.gatech.edu

MS56

Characterizing Biological Networks Using Subgraph Counting and Enumeration

We will describe the parallel algorithms and implementation of a new software tool called FASCIA for approximately counting and enumerating network motifs. Subgraph counting is used to detect and characterize local structure in several classes of biological networks, including protein interaction networks and metabolic networks. Exhaustive enumeration and exact counting is extremely compute-intensive, with running time growing exponentially with the number of vertices in the template. In FASCIA, we combine parallelism with the color coding technique to determine approximate counts of non-induced occurrences of a subgraph in the original network. The color coding technique gives a fast approximate algorithm for this problem, using a dynamic programming-based counting approach. Our new contributions are a multilevel shared-memory parallelization of the counting scheme and several optimizations to reduce the memory footprint. We show that approximate counts can be obtained for templates with up to 12 vertices, on networks with up to millions of vertices and edges.

George Slota, Kamesh Madduri
Pennsylvania State University
gms5016@psu.edu, madduri@cse.psu.edu

MS56

Connectome Coding Via Latent Position Estimation

We consider a latent position embedding approach to Con-

nectome Coding, in which adjacency spectral embedding is applied to a large brain graph. The vertices represent neurons and the edges represent axon-synapse-dendrite connectivity. A Euclidean representation of the vertices via the eigendecomposition of the adjacency matrix for such a graph has been shown to estimate latent positions, and we consider subsequent inference (clustering or classification of neurons, for example) in this space.

Carey Priebe
Johns Hopkins University
cep@jhu.edu

Joshua Vogelstein
Duke University
jo.vo@duke.edu

Jacob Vogelstein
Johns Hopkins University APL
jacob.vogelstein@jhuapl.edu

MS56

Fast Clustering Methods for Genetic Mapping in Plants

Graph-analytic methods, such as minimum spanning tree, motif finding, and clustering algorithms, are increasingly finding their way into the domain of genomic science, helping geneticists utilize the large scale and fine granularity of genomic sequence data. In this talk, we will highlight our progress in developing fast, scalable clustering algorithms for linkage group identification in large genetic marker datasets, and we present the challenges and new opportunities for parallel clustering algorithms in genetic applications.

Veronika Strnadova
University of California, Santa Barbara
veronika@cs.ucsb.edu

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

Joseph Gonzalez, Stefanie Jegelka
University of California, Berkeley
joseph.e.gonzalez@gmail.com, stefje@eecs.berkeley.edu

Jarrold Chapman
DOE Joint Genome Institute
jarrod@gmail.com

John R. Gilbert
Dept of Computer Science
University of California, Santa Barbara
gilbert@cs.ucsb.edu

MS57

Runtime Requirements for Scalable Semantic Data Analysis

Scientific discovery, economic competitiveness, and national security depend increasingly on the insightful analysis of web-scale data sets. Unfortunately, traditional database platforms based on relational operations are ineffective at discovering complex relationships among heterogeneous data as they do not support subgraph isomorphism and typed path traversal. In response, we are developing

a multithreaded semantic graph engine. In this talk, I will discuss the engines runtime system requirements to run at scale on HPC clusters.

John T. Feo
Pacific Northwest Laboratory
john.feo@pnl.gov

MS57

The Charm++ Applications Experience: Production Use of an Asynchronous, Adaptive Runtime

Charm++ is a well-established parallel programming system based on over-decomposition and message-driven execution. Programs are written in terms of indexed collection of C++ objects communicating via asynchronous method invocations. Its signature strength is an adaptive runtime system that can migrate objects and schedule execution to automate resource management, optimize communication, automate thermal/power optimizations, and tolerate faults. We illustrate the benefits of this approach via examples from production applications and miniApps, including NAMD, EpiSimdemics, AMR, etc.

Laxmikant Kale
University of Illinois at Urbana-Champaign
kale@illinois.edu

MS57

An Open Community Runtime (OCR) for Exascale Systems

Complex platforms needed to run scientific computing, big data analytics and control (cyber physical systems) workloads stretch the capabilities of conventional runtime systems. New dynamic introspective runtimes designed to provide high efficiency across complex heterogeneous computing platforms by exploiting dynamic scheduling, fine grained power management, and communication reduction techniques may offer significant advantage. This talk introduces the Open Community Runtime (OCR), a modular framework that will help enable community-wide runtime innovation.

Wilfred R. Pinfold
Intel
wilfred.pinfold@intel.com

MS57

Habanero Execution Model and the Open Community Runtime Project

The Habanero execution model is designed to address key challenges that will be faced by exascale systems. It is built on a foundation of asynchronous runtime primitives with strong semantic, scalability, and performance guarantees. This presentation will summarize multiple implementations of the Habanero execution model, one of which is built on the new Open Community Runtime (OCR) framework. We will also discuss migration paths from current execution models to the Habanero execution model.

Vivek Sarkar
Rice University
vsarkar@rice.edu

MS58

Compiler-Automated Communication-Avoiding

Optimization of Geometric Multigrid

We describe a compiler approach to introducing communication-avoiding optimizations in geometric multigrid (GMG), one of the most popular methods for solving partial differential equations. Communication-avoiding optimizations reduce vertical communication through the memory hierarchy and horizontal communication across processes or threads, usually at the expense of introducing redundant computation. We focus on applying these optimizations to the smooth operator, which successively reduces the error and accounts for the largest fraction of the GMG execution time. Our compiler technology applies a set of novel and known transformations to derive an implementation comparable to hand-written optimizations. An underlying autotuning system explores the tradeoff between reduced communication and increased computation, as well as trade offs in threading schemes, to automatically identify the best implementation for a particular architecture.

Protonu Basu
University of Utah
School of Computing
protonu@cs.utah.edu

Samuel Williams
Lawrence Berkeley National Laboratory
swwilliams@lbl.gov

Brian Van Straalen
Lawrence Berkeley National Laboratory
Computational Research Division
bvstraalen@lbl.gov

Anand Venkat
University of Utah
anandv@cs.utah.edu

Leonid Oliker
Lawrence Berkeley National Laboratory
loliker@lbl.gov

Mary Hall
School of Computing
University of Utah
mhall@cs.utah.edu

MS58

Relevant Stencil Structures for Modern Numerics

Abstract not available at time of publication.

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

MS58

Automatic Generation of Algorithms and Data Structures for Geometric Multigrid

Multigrid is one of the most efficient methods for numerical solution of PDEs. However, the concrete multigrid algorithm and its implementation highly depends on the underlying problem and hardware. Therefore, many different variants are necessary to cover all relevant cases. We try to generalize the data structures and multigrid components required to solve elliptic PDEs on Hierarchical Hy-

brid Grids (HHG) in order to formulate them in an intuitive domain specific language and automatically generate them.

Harald Koestler, [Sebastian Kuckuk](#)
 Universität Erlangen-Nürnberg
 harald.koestler@informatik.uni-erlangen.de,
 kuckuk@i10.informatik.uni-erlangen.de

MS58

Performance Engineering for Stencil Updates on Modern Processors

We apply the recently introduced ECM (Execution-Cache-Memory) performance model for multicore processors on stencil- and stencil-like algorithms. ECM is an extension of the Roofline Model and allows a prediction of single-core performance and saturation properties of streaming codes on multicore chips. This leads to deeper insight into performance behavior and energy consumption and enables a model-guided performance engineering approach, in which the concept of "optimal performance" is well-defined. Case studies for several short- and long-range stencil codes are presented.

Jan Treibig
 Erlangen Regional Computing Center
 Erlangen, Germany
 jan.treibig@rrze.uni-erlangen.de

Georg Hager
 Erlangen Regional Computing Center (RRZE)
 georg.hager@fau.de

[Gerhard Wellein](#)
 Erlangen Regional Computing Center (RRZE), Germany
 Gerhard.Wellein@rrze.uni-erlangen.de

MS59

Evolution and Revolution in Massively Parallel Quantum Chemistry Codes

We consider the evolution of NWChem for modern supercomputers with significantly node-level parallelism (e.g. Blue Gene/Q and accelerator-based systems) as well the recently developed Cyclops Tensor Framework and MADNESS that apply fundamentally different algorithmic, mathematical and/or software design approaches from the rest of the field. The tradeoffs between redesigning and rewriting vs. refactoring and using offload directives is evaluated in multiple contexts.

[Jeff R. Hammond](#)
 Argonne National Laboratory
 Leadership Computing Facility
 jhammond@alcf.anl.gov

MS59

Parallel Algorithms for Quantum Many-body Methods: Lessons Learned from NWChem

Two approaches are discussed for improving scalability and reducing on-node execution time in terms of 1) load balance and 2) overlapping communication with computation. The performance of the coupled cluster (CC) module in NWChem is instrumented and shown to suffer from processor starvation and prohibitive overhead due to its use of centralized dynamic load balancing techniques. This work

presents an Inspector/Executor load balancing algorithm based on performance modeling of task kernels and static partitioning of associated data. A novel hybrid method, we call dynamic buckets, is described and shown to be a viable solution for accomplishing good load balance with minimal overhead at scale. New optimization methods are also considered for hiding latency associated with the collection of remote data from the global address space in NWChem. This is done by overlapping time spent waiting on the network with computation time in the form of local computation. Overlap is achieved by utilizing a new library for queuing node-local tasks and consuming them as worker processes become available. While this technique is described within the context of NWChem, it has interesting implications for the future design of performance-optimal partitioned global address space applications.

[David Ozog](#)
 University of Oregon
 ozog@cs.uoregon.edu

Jeff R. Hammond
 Argonne National Laboratory
 Leadership Computing Facility
 jhammond@alcf.anl.gov

Pavan Balaji
 Argonne National Laboratory
 balaji@mcs.anl.gov

James Dinan
 Ohio State University
 dinan@cse.ohio-state.edu

Allen Malony
 University of Oregon
 malony@cs.uoregon.edu

Sameer Shende
 Dept. of Computer and Information Science
 University of Oregon
 sameer@cs.uoregon.edu

MS59

Scaling the Lattice Boltzmann Method on Massively Parallel Systems

Accurate and reliable modeling of cardiovascular hemodynamics has the potential to improve understanding of the localization and progression of heart diseases. However, building a detailed, realistic model of human blood flow is a formidable mathematical and computational challenge. Such simulations can provide insight into factors like endothelial shear stress that act as triggers for the complex biomechanical events that can lead to atherosclerotic pathologies. Currently, it is not possible to measure endothelial shear stress in vivo, making these simulations a crucial component to understanding and potentially predicting the progression of cardiovascular disease. In this talk, I will present and examine our approach for enabling an efficient parallel model of the fluid on IBM Blue Gene/Q architecture and discuss the optimizations that extend the region of the circulatory system that can be modeled. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

[Amanda Randles](#)
 Harvard School of Engineering and Applied Sciences

amanda.randles@gmail.com

Erik W. Draeger
Lawrence Livermore Nat. lab.
draeger1@llnl.gov

Franziska Michor
Harvard University
michor@jimmy.harvard.edu

MS59

Multi-Scale Parallelism in Yt: Lessons from a Community Driven Analysis Package

We describe the parallel environment developed and deployed by yt, and open-source python library for data analysis and visualization. In order to tackle increasingly large data sizes in both single snapshot and time-series analysis, we have developed several different parallel approaches. We leverage Python, Numpy, Cython, and mpi4py to develop a communication system that manages communicators capable of distributing work across and within computational nodes. A large design aspect involves the community-driven nature of yt. We will describe the impact this has had on designing parallel analysis, and in particular our approach of exposing both the data and the parallelism to the user at all levels to allow for as complex or simple of an approach as desired. Finally we will showcase several use examples spanning from massive time-series analysis to volume rendering of some of the largest volumetric data available, as well as discuss future improvements to both user interface and parallel capability.

Samuel W. Skillman
KIPAC
Stanford
samskillman@gmail.com

MS60

A Flexible Programming Environment for Distributed Heterogeneous Resources

Novel system designs with steeply escalating resource counts, burgeoning heterogeneity, and increasing memory accesses unpredictability, calls for novel programming paradigms. PaRSEC, our answer to this challenge, is a layered approach that provides a clear separation of concerns between architecture, algorithm, and data distribution by allowing the algorithm to be decoupled from the data distribution and the underlying hardware. Developers can focus solely on the algorithmic level without the constraints involved with programming for hardware trends.

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

MS60

OmpSs for Modern Cluster Architectures

Cluster architectures are ubiquitous nowadays. In addition the usage of accelerators in such system has become a popular solution to increase the peak performance of these systems. This trend has stressed the need for programming models that can ease the development of applications. OmpSs is a programming model that allows the parallelization of applications by annotating their sequential versions.

The talk will review the features of OmpSs for clusters and present some performance results.

Javier Bueno
Barcelona Supercomputing Center
javier.bueno@bsc.es

MS60

Starpu-MPI: Extending Task Graphs from Heterogeneous Platforms to Clusters Thereof

Production HPC clusters more and more integrate multicore CPUs, GPU accelerators, and now even Xeon Phi accelerators. StarPU has been successfully used on single-node systems to abstract away that complex heterogeneity from the programmer through task graphs, automatically optimizing task scheduling and data transfers. StarPU-MPI extends this to clusters by automatizing inter-node communications, allowing to keep the same application source code for both the debuggable sequential version and the scalable distributed heterogeneous version.

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

MS60

DuctTeip: A Task-Based Parallel Programming Framework with Modularity, Scalability and Adaptability Features

The DuctTeip framework is used for writing task parallel programs for hybrid systems of shared and distributed memory architectures. DuctTeip provides easy interfaces for programmers with different levels of expertise. At the highest level, the technical aspects of parallel programming are hidden, while at lower levels, a range of tools for more advanced users is provided. Key features to be discussed are support for distributed task generation and localized dependency tracking through data versioning.

Afshin Zafari
Uppsala University
Dept. of Information Technology
afshin.zafari@it.uu.se

Martin Tillenius
Dept. of Information Technology
Uppsala University
martin.tillenius@it.uu.se

Elisabeth Larsson
Uppsala University, Sweden
Department of Information Technology
Elisabeth.Larsson@it.uu.se

MS61

Bone Structure Analysis with Multiple GPGPUs

Osteoporosis is a disease that affects a growing number of people by increasing the fragility of their bones. To improve the understanding of bone, large scale computer simulations are applied. A fast, scalable and memory efficient solver for such problems is ParOSol. It uses the conjugate gradient algorithm with a multigrid preconditioner. Modifications of ParOSol are presented that profit from the exorbitant compute capabilities of recent GPGPUs. The fastest achieves speedups beyond five while the

code scales to 256 GPGPUs.

Peter Arbenz, Cyril Flaig
ETH Zurich
Computer Science Department
arbenz@inf.ethz.ch, cyril@flaig.ch

Daniel Kellenberger
ETH Zurich, Computer Science Department
danikel@student.ethz.ch

MS61

Deflated Preconditioned Conjugate Gradient for Bubbly Flows: Multi-GPU/CPU Implementations

We present a multi-GPU/multi-CPU implementation of the Deflated Preconditioned Conjugate Gradient (DPCG) Method for solving a linear system arising in the simulation of Bubbly Flows. We discuss 2 data division schemes for load balancing of our parallel implementation of the DPCG method. Our experiments demonstrate up to 5 times speedup for a system of 16 million unknowns across 8 GPUs connected via MPI when compared to similar number of CPU cores.

Rohit Gupta
Technology University of Delft
rohit.gupta@tudelft.nl

Martin B. van Gijzen, Kees Vuik
Delft University of Technology
m.b.vangijzen@tudelft.nl, c.vuik@tudelft.nl

MS61

Factorized Sparse Approximate Inverses on GPUs

GPUs exhibit significantly higher peak performance than conventional CPUs. However, due to their programming model and microarchitecture, only highly parallel algorithms can exploit their full potential. In this context, the FSAI preconditioner may represent – with its inherent parallelism – an optimal candidate for the conjugate gradient-like solution of sparse linear systems. While its application to a vector involves only matrix-vector products, a GPU-based implementation requires a nontrivial recasting of multiple computational steps.

Massimo Bernaschi
Istituto per le Applicazioni del Calcolo
m.bernaschi@iac.cnr.it

Carlo Fantozzi
Dept. Information Engineering - University of Padova
fantozzi@dei.unipd.it

Massimiliano Ferronato, Giuseppe Gambolati
University of Padova
DMMMSA
massimiliano.ferronato@unipd.it, gambo@dmsa.unipd.it

Carlo Janna
Dept. ICEA - University of Padova
carlo.janna@unipd.it

MS61

Preconditioning Techniques for GPU-Accelerated

Environments

This work discusses our preliminary work to develop iterative linear solvers accelerated by GPUs. We illustrate the potential advantages of GPUs for sparse linear algebra computations as well as the difficulties encountered. Techniques for speeding up sparse matrix-by-vector (SpMV) and finding suitable preconditioners are discussed. A GPU SpMV kernel in the jagged diagonal format is proposed for unstructured matrices. Experiments show that it can be up to 8 times faster than on the host CPU. Sparse triangular solves (SpSV), the basic operations required in incomplete LU (ILU) preconditioners, yield very low performance on GPUs. Although for some cases, SpSV's can still outperform their CPU counterparts with level-scheduling, the overall performance of ILU preconditioned solvers can only achieve a speed-up by a small factor. Better preconditioners for GPUs are therefore needed. We explored several parallel preconditioners, with which the overall performance can be further improved.

Ruipeng Li
Department of Computer Science & Engineering
University of Minnesota
rli@cs.umn.edu

Yousef Saad
Department of Computer Science
University of Minnesota
saad@cs.umn.edu

MS62

Integration of Albany and Mesh Adaptation for Parallel Applications

Dynamic mesh adaptation is an enabling technology for modern multiphysics analysis applications; the technology focuses on maintaining a high quality discretization as the length scales of the coupled physics simulation evolve over time and space. To maintain efficiency in parallel applications, it is critical to maintain an effective load balance as the work performed by a given processor changes relative to its neighbors as the mesh and solution complexity evolves. We present an overview of the approach used to integrate the RPI SCOREC PUMI mesh adaptation library into the Albany multiphysics framework, where the Trilinos Zoltan package is used for dynamic load balancing in the integrated application.

Glen Hansen
Sandia National Laboratories
gahanse@sandia.gov

Brian Granzow, E. Seegyoung Seol, Mark S. Shephard
Rensselaer Polytechnic Institute
granzb@rpi.edu, seols@rpi.edu, shephard@rpi.edu

MS62

Strategies for Reducing Setup Costs in Parallel Algebraic Multigrid

Algebraic multigrid (AMG) preconditioners are often employed in large-scale computer simulations to achieve scalability. AMG construction can be costly, sometimes as much as the solve itself. We discuss strategies for reducing expense through reuse of information from prior solves. The information type depends on the method. For smoothed aggregation AMG this includes aggregation data, whereas for energy minimization it includes sparsity patterns and

prior interpolants. We demonstrate the effectiveness of such strategies in parallel applications.

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

Andrey Prokopenko
Sandia National Laboratories
aprokop@sandia.gov

Ray S. Tuminaro
Sandia National Laboratories
Computational Mathematics and Algorithms
rstumin@sandia.gov

Christopher Siefert
Sandia National Laboratories
csiefer@sandia.gov

MS62

ARKode: A Library of High Order Implicit/explicit Methods for Multi-rate Problems

Modern predictive modeling increasingly requires simulations of interacting processes that evolve at differing rates. Historically, such multi-rate simulations have been treated using low accuracy and potentially unstable operator-splitting approaches. We have recently developed a new solver library, ARKode, for adaptive time evolution of multi-rate systems using stable and high-order-accurate integration methods. As a member of the FASTMath SUNDIALS library, ARKode targets large-scale parallel simulations, and is designed to be easily adapted to user-defined data structures. In this talk, we will present the algorithmic infrastructure for the ARKode library, including specific optimizations for handling simulations with significant spatial adaptivity.

Daniel R. Reynolds
Southern Methodist University
Mathematics Department
reynolds@smu.edu

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

David Gardner
Southern Methodist University
dgardner@mail.smu.edu

Alan Hindmarsh
Lawrence Livermore National Laboratory
Center for Applied Scientific Computing
alanh@llnl.gov

MS62

Algorithmic Advances for Algebraic Multigrid with Reduced Communication

Algebraic Multigrid (AMG) solvers are an essential component of many large-scale scientific simulation codes. The increasing communication complexity on coarser grids combined with the effects of increasing numbers of cores can lead to severe performance bottlenecks for AMG on various computer architectures. We will present recent

progress on several efforts to reduce communication in AMG, including the use of non-Galerkin approaches and additive AMG variants.

Ulrike Meier Yang, Robert Falgout, Jacob Schroder,
Panayot Vassilevski
Lawrence Livermore National Laboratory
yang11@llnl.gov, falgout2@llnl.gov, schroder2@llnl.gov,
vassilevski1@llnl.gov

MS63

Exploiting the Power of Heterogeneous Computing for Kinetic Simulations of Plasmas

In the past decade, advances in computing power have been driven by multi-core processors and, more recently, also by heterogeneous architectures like GPUs and Intel MIC. We will present our work on adapting the kinetic Plasma Simulation Code (PSC) to such architectures. The PSC is an electromagnetic particle-in-cell code that was designed to have run-time swappable computational kernels and data structures that can take advantage of specific processing hardware. The code has been run on up to 5120 Nvidia Kepler K20X GPUs and achieves a performance of > 550 million particles / sec / GPU, compared to 120 million particles / sec on a modern conventional 16-core CPU. We will describe PSC's main GPU algorithms and describe our approach to load balancing.

Kai Germaschewski
Space Science Center and Department of Physics
University of New Hampshire
kai.germaschewski@gmail.com

Narges Ahmadi
Space Science Center
University of New Hampshire
narges.ahmadi@unh.edu

Homayoun Karimabadi
University of California-San Diego
Department of Electrical and Computer Engineering,
MC0407
homakar@gmail.com

MS63

Accelerator Beam Dynamics on Multicore and Gpu Systems

Synergia is a parallel Python/C++ modeling package for collective effects in particle accelerators. The parallelism in Synergia was originally pure MPI, but we have recently extended it to also use OpenMP on multicore architectures and CUDA on GPUs. We describe our efforts to find efficient solutions to the computing challenges on these new architectures and show performance results from our first implementations.

Qiming Lu, James Amundson
Fermi National Accelerator Laboratory
qlu@fnal.gov, amundson@fnal.gov

MS63

Migrating the GTC-P Gyrokinetic Particle-in-Cell Code to Multicore and Manycore Systems

The Gyrokinetic Particle-in-cell (PIC) method has been successfully applied in studies of low-frequency microtur-

bulence in magnetic fusion plasmas. While the excellent scaling of PIC codes on modern computing platforms is well established, significant challenges remain in achieving high on-chip concurrency for the new path to exascale system. In this talk, we will present our efforts in migrating the gyrokinetic particle-in-cell code GTC-P to next generation multicore and manycore systems. This includes optimization techniques to enhance the performance of GTC-P code on large-scale machines and optimization strategy for hybrid programming model involving MPI, OpenMP and CUDA. We show the performance studies of GTC-P on BG/Q (Mira), Cray XK7 (Titan) and Fujitsu K machine up to 524k cores.

Bei Wang

Princeton Institute of Computational Science and Engineering
Princeton University
beiwang@princeton.edu

Stephane Ethier

Princeton Plasma Physics Laboratory
ethier@pppl.gov

William Tang

Princeton University
PPPL
tang@pppl.gov

Khaled Z. Ibrahim

Lawrence Berkeley National Laboratory
Berkeley, CA, USA
kzibrahim@lbl.gov

Kamesh Madduri

Pennsylvania State University
madduri@cse.psu.edu

Samuel Williams, Leonid Olikier

Lawrence Berkeley National Laboratory
swilliams@lbl.gov, loliker@lbl.gov

MS63

Hybrid MPI/OpenMP/GPU Parallelization of XGC1 Fusion Simulation Code

XGC1 is a full-f gyrokinetic particle-in-cell code designed to simulate edge plasmas in tokamaks. We describe how a careful design using MPI, OpenMP, and CUDA Fortran enabled XGC1 to make effective use of a Cray XK7 with GPU accelerated compute nodes, achieving a four times speed-up over the original CPU-only version and excellent weak scalability out to at least 18,624 XK7 nodes, thus enabling heretofore infeasible science studies.

Patrick H. Worley, Ed D'Azevedo

Oak Ridge National Laboratory
worleyph@ornl.gov, dazevedoef@ornl.gov

Jianning Lang, Seung-Hoe Ku, Choong-Seock Chang

Princeton Plasma Physics Laboratory
jlang@pppl.gov, sku@pppl.gov, cschang@pppl.gov

MS64

Streaming Graph Analytics for Real-World Problems

Understanding dynamic graphs is essential to relevant anal-

ysis of near real-time systems, such as a case study we will present on tracking population-scale phenomena. Modeling the evolution of the graph explicitly leads to computational challenges. Hence, the innovation in our approach is using graph statistics in lieu of explicit modeling. Our streaming graph analytics use two phases; in the first phase, we track vertex and/or edge statistics using STINGER, a parallel graph data representation, and in the second phase, we generate big data insights.

David A. Bader, James Fairbanks

Georgia Institute of Technology
bader@cc.gatech.edu, james.fairbanks@gatech.edu

MS64

Scalable Graph Methods for Functional Characterization of Environmental Microbial Communities

Clustering formulations are prevalent in a number of biological applications, including that of determining protein-protein interactions and discovering protein families from metagenomics data. Performing these operations at a large-scale, however, still remains technically challenging. In this talk, I will report on the on-going development of new parallel approaches for identifying protein sequence clusters from large-scale metagenomics data. The parallel design and results of two heuristics on large-scale distributed memory machines will be presented.

Ananth Kalyanaraman

Associate Professor
School of EECS, Washington State University
ananth@eecs.wsu.edu

MS64

Tracking Combustion Features Using Topological Methods Involving Large Graphs

Topological methods have recently become known to capture effectively fundamental structural characteristic of scientific data due to the ability of abstracting with simple graph models multiscale trends in scalar and vector fields. In combustion simulations these graphs represent important features such as burning cells or extinction/reignition kernels. We present results from large scale simulations generating massive graphs that provide full representation of the features of interests while reducing the data size by orders of magnitude.

Valerio Pascucci

SCI Institute - University of Utah
pascucci@sci.utah.edu

Peer-Timo Bremer

Lawrence Livermore National Laboratory
ptbremer@gmail.com

Attila Gyulassy

University of Utah
aggyulassy@ucdavis.edu

Aaditya Landge

SCI Institute - University of Utah
aaditya@sci.utah.edu

Janine Bennet

Sandia National Laboratories
jcbenne@sandia.gov

Jackie Chen
Sandia National Laboratories, USA
jhchen@sandia.gov

MS64

Generating Large Graphs for Benchmarking

Network data is ubiquitous and growing, yet we lack realistic generative network models that can be calibrated to match real-world data. The recently proposed Block Two-Level Erdős-Rényi (BTER) model can be tuned to capture two fundamental properties: degree distribution and clustering coefficients. The latter is particularly important for reproducing graphs with community structure, such as social networks. In this talk, we will describe a scalable implementation of BTER that requires only $O()$ storage where is the maximum number of neighbors for a single node. The generator is trivially parallelizable, and we show results for a modeling a real-world web graph with over 4.6 billion edges. We propose that the BTER model can be used as a graph generator for benchmarking purposes and provide idealized degree distributions and clustering coefficient profiles that can be tuned for user specifications.

Ali Pinar, C. Seshadhri
Sandia National Labs
apinar@sandia.gov, scomand@sandia.gov

Tamara G. Kolda, Todd Plantenga
Sandia National Laboratories
tgkolda@sandia.gov, tplante@sandia.gov

MS65

SWARM: Maximizing Compute Efficiency and Scalability Today While Designing for Reliability and Power Efficiency for Tomorrow

SWARM is a parallel compute framework enabling superior performance and productivity on large-scale systems. SWARM virtualizes distributed heterogeneous systems into a single system image through an asynchronous, dynamic, adaptive runtime system. This presentation will describe SWARM and demonstrate its use to produce significant improvements in efficiency, scalability, and heterogeneous resource utilization. Finally, we will discuss how this framework can be extended to include resilience and power efficiency.

Rishi Khan
ETI International
rishikhan@gmail.com

MS65

Distributed Memory Runtimes for Dense Linear Algebra Software

DPLASMA is a dense linear algebra package for distributed heterogeneous systems. It is designed to deliver sustained performance for distributed systems where each node is featuring multiple sockets of multicore processors and GPU accelerators or Xeon Phi co-processors. DPLASMA achieves this objective through a state of the art PaR-SEC runtime. PaRSEC assigns computation threads to the cores, overlaps communications and computations and uses a dynamic, fully-distributed scheduler based on architectural features such as NUMA nodes and algorithmic

features such as data reuse.

Jakub Kurzak
University of Tennessee Knoxville
kurzak@eecs.utk.edu

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

Jack J. Dongarra
Department of Computer Science
The University of Tennessee
dongarra@cs.utk.edu

MS65

Dynamic Runtime Systems for Scalable Graph Processing

To obtain acceptable parallel performance, large-scale graph algorithms require significant asynchrony. The reliance on synchronous operations and heavily optimized collective operations in traditional approaches based on CSP/MPI preclude their effective use for graph algorithms – suggesting that new execution models (and associated runtime systems) are needed. In this talk we illustrate the productivity and performance advantages that can be realized for graph algorithms when using new execution models that inherently support fine-grained asynchronous message-driven operations.

Andrew Lumsdaine
Open Systems Laboratory
Indiana University
lums@cs.indiana.edu

MS65

ParalleX: Defining the Computer of the Future

ParalleX is an experimental execution model to guide the development and govern the interoperability of hardware and software component layers of future extreme-scale systems. It has served to define the HPX asynchronous runtime system software under development for the DOE X-Stack XPRESS project and the DOE PSAAP-II C-SWARM project among others. This presentation will describe ParalleX and HPX and provide results demonstrating significant improvements for efficiency and scaling with respect to conventional practices.

Thomas Sterling
Indiana University
tron@indiana.edu

MS66

Tuning Sparse and Dense Matrix Operators in SeisSol

In this talk we show recent advances in increasing the computational performance of the software package SeisSol, one of the leading codes for the simulation of earthquake scenarios. SeisSol uses the discontinuous Galerkin method combined with flexible unstructured tetrahedral meshes for spatial and Arbitrary high order DERivatives (ADER) for time discretization. We present our node-level optimization strategy based on hardware-aware code generation for element-local, low-rank matrix operations ranging from sparse to near dense on Intel Sandy Bridge and Xeon Phi.

The involved matrix sparsity patterns are known a priori, which allows us to eliminate all indirect matrix accesses and generate AVX vector instructions in the pre-compile phase utilising the hardware at high performance.

Alexander Breuer, Sebastian Rettenberger
Technische Universität München
Department of Informatics, Scientific Computing
breuera@in.tum.de, rettenbs@in.tum.de

Alexander Heinecke
TU München
heinecke@in.tum.de

Christian Pelties
Ludwig-Maximilians-Universitaet Muenchen
University Munich
pelties@geophysik.uni-muenchen.de

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

MS66

Firedrake: a Multilevel Domain Specific Language Approach to Unstructured Mesh Stencil Computations

How do we enable scientists to specify in a simple and mathematically expressive way the simulation they wish to perform, but still make efficient use of the diverse massively parallel hardware which will dominate supercomputing over the coming years? Firedrake is a multilayer abstraction package which generates unstructured mesh numerical partial differential equation solvers. It builds on the success of the Unified Form Language as an expressive symbolic language for the finite element method, and combines it with the runtime targeting of different parallel architectures provided by the PyOP2 system.

Gheorghe-Teodor Bercea
Department of Computing
Imperial College London
gheorghe-teodor.bercea08@imperial.ac.uk

David Ham
Imperial College
david.ham@imperial.ac.uk

Paul Kelly, Nicolas Lorient
Imperial College London
p.kelly@imperial.ac.uk, n.lorient@imperial.ac.uk

Fabio Luporini, Lawrence Mitchell
Department of Computing
Imperial College London
f.luporini12@imperial.ac.uk,
lawrence.mitchell@imperial.ac.uk

Florian Rathgeber
Imperial College London
f.rathgeber10@imperial.ac.uk

MS66

Evaluating Compiler-driven Parallelization of Stencil Micro-applications on a GPU-enabled Cluster

In this talk we will demonstrate how parallelization and

further optimization of stencil codes for GPUs could be automated by compiler toolchains. By example of wave equation stencil, hand-written naive and optimized for locality versions will be compared against compiler-generated parallel code, presenting the roofline performance, efficiency of tiling, JIT-compilation and other properties. The results of benchmarking KernelGen and PPCG auto-parallelizing compilers as well as one commercial OpenACC compiler will be presented on a set of 10 stencil micro-applications. Finally, we will show how automatically parallelized code of a very large wave propagation problem performs on the Piz Daint supercomputer.

Dmitry Mikushin
USI - Università della Svizzera italiana
Institute of Computational Science
dmitry.mikushin@usi.ch

Olaf Schenk
USI - Università della Svizzera italiana
olaf.schenk@usi.ch

MS66

Stencil Computations: from Academia to Industry

Synthetic benchmarks used in academia may expose false impressions of optimization techniques. This is not an exception on stencil computations. When these techniques are applied to industrial codes their impact can be deceiving. Furthermore, some academia optimizations cannot be easily implemented for industrial numerical codes. The industrial complexity can range considerably: from a simple 2nd order spatial stencil to burdensome cases with 8th order spatial stencils, dozens of variables and staggered grids with many parameters.

Raul de la Cruz, Mauricio Hanzich, Jose Maria Cela
Barcelona Supercomputing Center (BSC)
raul.delacruz@bsc.es, mauricio.hanzich@bsc.es,
josem.cela@bsc.es

MS67

Challenges and Opportunities in Incorporating State-of-the Art Parallel Computation Capabilities in Modern End-User Based Dynamic Reservoir Simulation Systems

The differences in software development characteristics of (parallel) scientific computing and end-user based simulation software systems can often lead to very long time intervals between the availability or publication of new results and the actual implementation in the end-user based simulator. The challenges and opportunities in, and the necessity of bringing those two worlds together are discussed. Examples of attempts to achieve this for an end-user based dynamic simulation system currently under development are given.

Sander Belien
Shell Oil Company
sander.belien@shell.com

MS67

Addressing Domain Decomposition, Load Balancing and Coupled Solution in An Unstructured Parallel Reservoir Simulator

The talk introduces the distributed parallel data infras-

structure of a massively parallel reservoir simulator used to model Middle-East giant reservoirs. It is a solution centric system where the data complexity and communication management are hidden from the rest of the simulator. We highlight how the system ties the parallel I/O space, the simulation space, and coupled implicit solve space to form a consistent, scalable, and extendible platform for on-going feature development.

Larry Fung
Saudi Aramco
larry.fung@aramco.com

MS67

A New Linear Solver Algorithm for Solving Difficult Large Scale Thermal Models

A new preconditioner algorithm, called Strong-Weak Incomplete Factorization with Thresholds (SWIFT), was developed to solve difficult thermal problems. It uses the magnitude of the infill terms to modify the factorization sparsity pattern adaptively to improve accuracy. It combines scaling, equilibration, reordering, and threshold-based incomplete LU factorization. The solver algorithm was applied to several large-scale thermal models with up to 10 million cells. For these models it accelerated run time by between 2 and 10 times.

Gary Li
Chevron
gary.li@chevron.com

John Wallis
Wallis Consulting
blitz2@windstream.net

Gareth Shaw
Schlumberger
gshaw@abingdon.oilfield.slb.com

MS67

Strong Scalability of Reservoir Simulation on Massively Parallel Computers: Issues and Results

We investigate parallel performance for black oil reservoir simulation on multiple massively parallel computing architectures. A deliberate strategy of performance-based development of the major types of computations encountered in reservoir simulation is employed. Even though most operations are memory-bandwidth bound, it is possible with careful implementation, to get excellent parallel efficiency to several 1000s of cores. We discuss numerical issues, scalability and parallel efficiency of reservoir simulator on several very large and geologically challenging examples

Serguei Maliassov, Bret Beckner
ExxonMobil Upstream Research Company
serguei.maliassov@exxonmobil.com,
bret.l.beckner@exxonmobil.com

Vadim Dyadechko
ExxonMobil URC
vadim.dyadechko@exxonmobil.com

MS68

Space, Time, and I/O Efficient Polynomial Factori-

sation Via the Funnel Heap

We present an engineering of the Hensel lifting technique for bivariate polynomial factoring that addresses the expression swell arising in polynomial arithmetic. We overlap all polynomial multiplications and additions required in each lifting step using a priority queue, which reduces the space requirements by an order of magnitude. We use the cache-oblivious funnel heap (priority queue), which is able to process equal degree monomials inherently as part of its inner workings, without destroying the structure of the heap. This makes the resulting implementation suitable for instances that reach levels of the memory hierarchy where the cost of data transfer is inhibiting.

Fatima K. Abu Salem, Khalil El-Harake
Computer Science Department
American University of Beirut
fa21@aub.edu.lb, kme07@aub.edu.lb

Karl Gemayel
Computing Laboratory
University of Oxford
karl@gatech.edu

MS68

Toward High Performance Matrix Multiplication for Exact Computation

Matrix multiplication with exact results (i.e. integers, finite fields, polynomials) is important in computer algebra. Since nowadays processors do not provide such structures, careful implementation must be done to get high performance. I will report on some recent developments on matrix multiplication over integers and polynomials over F_p . I will discuss the use of CRT and Kronecker substitution to reduce to multiplication with small entries.

Pascal Giorgi
LIRMM
Univ. Montpellier 2, France
pascal.giorgi@lirmm.fr

MS68

CUMODP: A CUDA Library for Symbolic Computation with Polynomials over Finite Fields

CUDMOP is a CUDA library for exact computations with dense polynomials over finite fields. A variety of operations like multiplication, division, computation of subresultants, multi-point evaluation, interpolation and many others are provided. These routines are primarily designed to offer GPU support to polynomial system solvers and a bivariate system solver is part of the library. Algorithms combine FFT-based and plain arithmetic, while the implementation strategy emphasizes reducing parallelism overheads and optimizing hardware usage.

Marc Moreno Maza
Computer Science Department
University of Western Ontario
moreno@csd.uwo.ca

Wei Pan
Intel Corp.
wei.pan@intel.com

Sardar Anisul Haque, Farnam Mansouri, Ning Xie
Computer Science Department

University of Western Ontario
shaque4@uwo.ca, mansouri.farnam@gmail.com,
nxie6@csd.uwo.ca

MS68**Accelerating Path Tracking for Polynomial Homotopies**

Polynomial homotopies define deformations of polynomial systems. Path tracking methods apply numerical predictor-corrector techniques to compute solutions starting at known solutions in the homotopy. As the number of variables and degrees of the polynomials grow, the standard double precision becomes insufficient for accurate results. Aiming to compensate for the cost of double double and quad double arithmetic we apply massively parallel algorithms to accelerate the tracking of solution paths on the NVIDIA Tesla C2050 and K20C.

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
xyu30@uic.edu

MS69**Fast Direct Solvers Using Sparse Rocket Matrices**

We will be presenting novel algorithms to solve dense linear systems $Ax=b$ in $O(N)$ floating point operations for matrices A of size $N \times N$. Such matrices arise when computing the LU factorization of sparse matrices. The new algorithms have the potential to replace classical iterative solvers. Their computational cost is competitive and scales linearly with the problem size. In addition, being based on direct LU factorization, they do not require preconditioners.

Eric F. Darve

Stanford University
Mechanical Engineering Department
darve@stanford.edu

MS69**Designing Multifrontal Solvers Using Hierarchically Semiseparable Structures**

We consider the use of hierarchically semi-separable (HSS) representations based on randomized sampling into a multifrontal solver. The potential of these techniques has been shown but many issues need to be addressed in order to design a general-purpose solver. We show how graph-partitioning tools can be used in the clustering phase needed to reveal the HSS forms. We address the problem of pivoting, and we discuss parallelization techniques in a distributed-memory environment.

François-Henry Rouet

Lawrence Berkeley National Laboratory
frouet@lbl.gov

Xiaoye S. Li

LBNL
xsli@lbl.gov

Artem Napov
Université Libre de Bruxelles
anapov@ulb.ac.be

MS69**Multilevel Low-Rank Approximation Preconditioners**

A new class of preconditioning methods for solving linear systems of equations are introduced which are based on exploiting low-rank approximations to certain matrices. They handle indefiniteness quite well and are amenable to highly parallel computations. We first describe the methods for Symmetric Positive Definite model problems arising from Finite Difference discretizations of PDEs. We will then show how to extend them to general situations, by exploiting the domain decomposition framework.

Yousef Saad

Department of Computer Science
University of Minnesota
saad@cs.umn.edu

Ruipeng Li

Department of Computer Science & Engineering
University of Minnesota
rli@cs.umn.edu

MS69**Parallelization and Pivoting in a Block-Low Rank Multifrontal Solver**

The multifrontal factorization of a sparse matrix is achieved through partial factorizations of smaller dense matrices. In many applications, well defined subblocks of these matrices can be approximated by low-rank products. This substantially reduces the memory and operation requirements of the whole factorization. The Block Low-Rank (BLR) format is designed to offer the flexibility needed for a distributed-memory multifrontal solver with threshold partial pivoting. We present results with BLR as a direct solver and preconditioner.

Patrick Amestoy
INP(ENSEEIH)-IRIT
amestoy@enseeiht.fr

Cleve Ashcraft
Livermore Software Technology Corporation
cleve@lstc.com

Olivier Boiteau
EDF R&D
olivier.boiteau@edf.fr

Alfredo Buttari
CNRS-IRIT
buttari@enseeiht.fr

Jean-Yves L'Excellent
INRIA-LIP-ENS Lyon
jean-yves.l.excellent@ens-lyon.fr

Clément Weisbecker

INP(ENSEEIH)-IRIT
clement.weisbecker@enseeiht.fr

MS70

Unified Task-Data-Vector Parallelism on Manycore Architectures

Parallel scalability on emerging manycore computing architectures (Terascale workstations, Petascale clusters, and Exascale supercomputers) requires application codes to exploit all opportunities for manycore parallelism. Sandia's Kokkos library currently provides performance portable data and vector parallelism on CPUs and accelerators (NVIDIA Kepler and Intel Xeon Phi). Our current R&D is enhancing Sandia's Qthreads task scheduling library, integrating it into Kokkos, and extending the Kokkos API to provide a performance portable, unified task-data-vector parallel hierarchy for manycore architectures.

H. Carter Edwards, Stephen Olivier
Sandia National Laboratories
hcedwar@sandia.gov, slolivi@sandia.gov

MS70

Lulesh 2.0 Changes and Initial Experiences

Proxy applications are designed to evolve over time to reflect lessons learned in their use and changes in the application they represent. In this talk we describe changes made to LULESH to update it to better reflect our production applications. We also, discuss some of our initial experiences using the application in hardware and software co-design. In particular we describe experiences porting the code to OpenACC and some pitfalls encountered.

Ian Karlin
Lawrence Livermore National Laboratory
karlin1@llnl.gov

MS70

Algorithmic Acceleration of Ocean Models

We examine a nonlinear elimination method for the free-surface ocean equations based on barotropic-baroclinic decomposition. The two dimensional scalar continuity equation is treated implicitly with a preconditioned Jacobian-free Newton-Krylov method (JFNK) and the remaining three dimensional equations are subcycled explicitly within the JFNK residual evaluation. In this approach, the footprint of the underlying Krylov vector is greatly reduced over that required by fully coupled implicit methods. The method is second order accurate and scales algorithmically with timesteps much larger than explicit methods. Moreover, the heterogeneous nature of the algorithm lends itself readily to emerging architectures.

Chris Newman
Los Alamos National Lab
cnewman@lanl.gov

Geoff Womeldorff, Dana Knoll
Los Alamos National Laboratory
womeld@lanl.gov, nol@lanl.gov

MS70

Exploring Workloads of Adaptive Mesh Refine-

ment

Adaptive mesh refinement provides a dynamic means for focusing computation on particular areas of a problem domain. This approach provides challenges to high performance computers both in terms of computation and inter-process communication. As a means of exploring these issues, we've developed miniAMR, a new mini-application in the Mantevo suite. In this talk we describe miniAMR, present some performance results from various computing environments, and compare these results with a shock hydrodynamics code.

Courtenay Vaughan, Richard Barrett
Sandia National Laboratories
temp@sandia.gov, rbarre@sandia.gov

MS71

A Fully Implicit, Moment Accelerated, Electromagnetic Particle-in-Cell Algorithm

Moment-accelerated fully implicit particle-in-cell methods have been recently established as an approach for accurate electrostatic kinetic plasma simulations. In this talk, we extend the approach to nonradioactive electromagnetic (EM) problems in 1D using the Darwin model. We propose a moment accelerator based on the electron continuity and momentum equations. We will demonstrate its performance via numerical experiment on multiscale EM simulations.

Guangye Chen, Luis Chacon
Los Alamos National Laboratory
gchen@lanl.gov, chacon@lanl.gov

MS71

Multi-rate and Multi-level Spectral Deferred Correction Methods: High-order Integrators with Strong Coupling for Models with Disparate Time Scales and/or Multi-physics

Spectral Deferred Correction (SDC) methods are iterative time-stepping schemes for ODEs and PDEs. For systems that can be described by different physical models of varying computational complexity, multi-level SDC (MLSDC) schemes offer a natural way of accelerating the time-to-solution of fine models by leveraging coarse models through FAS (full approximation scheme) corrections within a hierarchy of SDC schemes. The MLSDC scheme will be presented along with variations including multi-rate and time-parallel versions.

Matthew Emmett
Lawrence Berkeley National Laboratory
Center for Computational Sciences and Engineering
mwemmett@lbl.gov

Michael Minion
ICME Stanford
mlminion@gmail.com

Max Duarte
Center for Computational Sciences and Engineering
Lawrence Berkeley National Laboratory
mdgonzalez@lbl.gov

MS71

An Asymptotic Parallel-in-Time Method for

Highly Oscillatory PDEs

We present a new time-stepping algorithm for nonlinear PDEs that exhibit scale separation in time. Our scheme combines asymptotic techniques (which are inexpensive but can have insufficient accuracy) with parallel-in-time methods (which, alone, can be inefficient for equations that exhibit rapid temporal oscillations). In particular, we use an asymptotic numerical method for computing, in serial, a solution with low accuracy, and a more expensive fine solver for iteratively refining the solutions in parallel. We present examples on the rotating shallow water equations that demonstrate that significant parallel speedup and high accuracy are achievable.

Terry Haut
Los Alamos National Laboratory
terryhaut@lanl.gov

MS71

An Algorithmically Accelerated Ocean Model Framework: Methods and Communication Strategies

We study moment based methods for a free-surface z-level ocean model. In this approach, the full three dimensional equations are coupled to a set of two dimensional moment equations, obtained by vertical integration. This formulation allows relaxation of the timestep by isolating the stiff physics to the reduced system. In addition, the heterogeneous nature of the formulation is amenable to modern architectures. We provide numerical examples to support our study and compare to traditional implementations.

Geoff Womeldorff
Los Alamos National Laboratory
womeld@lanl.gov

Chris Newman
Los Alamos National Lab
cnewman@lanl.gov

Dana A. Knoll, Luis Chacon
Los Alamos National Laboratory
nol@lanl.gov, chacon@lanl.gov

MS72

Computing, Exploring, and Tracking Features in Massive Scientific Data

With the continuing increase in available computing resources scientist are able to study physical phenomena with unprecedented detail and complexity. As a result, many advanced analysis techniques require robust, efficient, and scalable techniques to extract features from the emerging Tera- or Peta-byte sized data sets. We will present a general framework to define, extract, and analyze features from scientific data. It consists of two parts: a streaming and/or massively parallel computing pipeline capable of transforming the data into highly compact, abstract representations and an interactive analysis environment that extracts features from this meta-representation on-the-fly while providing a user with the ability to track feature evolution and explore various statistical summaries.

Peer-Timo Bremer
Lawrence Livermore National Laboratory
ptbremer@gmail.com

Valerio Pascucci
SCI Institute - University of Utah
pascucci@sci.utah.edu

Attila Gyulassy
University of Utah
jediati@sci.utah.edu

Janine C. Bennett
Sandia National Laboratories
jcbenne@sandia.gov

MS72

In-Situ Visualization with Explorable Images

This talk introduces the concept of explorable images as an in situ visualization solution. For extreme-scale simulations, explorable images enable scientists to validate as well as explore their simulation output at low cost using a desktop computer or even a mobile device. Explorable images for both scalar and vector field data are presented. Explorable image as a new media will potentially change how visualization of large-scale simulations is done.

Kwan-Liu Ma
Department of Computer Science
University of California at Davis
ma@cs.ucdavis.edu

Robert Miller, Yucong Ye
University of California at Davis
bobmiller@ucdavis.edu, chrisyeshi@gmail.com

MS72

In-Situ Processing with EAVL, the Extreme-Scale Analysis and Visualization Library

As the high performance computing ecosystem evolves, scientific simulations evolve with it, adopting new data models and programming paradigms to utilize new architectures. EAVL, the Extreme-scale Analysis and Visualization Library, targets these changing requirements with a flexible data model and support for data-parallel architectures including multi-core and graphics processors. This presentation will cover the adaptations EAVL makes to accomplish these goals and how EAVL is used for tightly-coupled in situ analysis.

Jeremy Meredith
Oak Ridge National Laboratory
jsmeredith@ornl.gov

MS72

Approaching Production In-Situ Visualization for Extreme Scale Analysis

In-situ visualization is a critical component of petascale simulation and beyond. Many researchers demonstrate the benefit of in-situ analysis in both development and production environments. However, most in-situ visualizations are designed ad hoc; they perform a singular purpose for a specific use. Our research attempts a more holistic approach. We assimilate existing technology where applicable and build generic in-situ tools that can be leveraged in a wide variety of simulations in addition to our own.

Kenneth Moreland, Ron A. Oldfield, Fabian Fabian
Sandia National Laboratories

kmorel@sandia.gov, raoldfi@sandia.gov,
ndfabia@sandia.gov

Andrew Bauer
Kitware Inc.
andy.bauer@kitware.com

Berk Geveci, David Lonie
Kitware
berk.geveci@kitware.com, david.lonie@kitware.com

MS73

Nebo: A DSL for Numerical Solution of PDEs on Current and Emerging Architectures

This talk discusses Nebo, a domain-specific language embedded within C++ for numerically solving PDEs on multiple architectures. Nebo provides declarative syntax that hides architecture-specific details and allows the same code to run on single-core, multi-core, many-core (GPUs), and/or heterogeneous systems. For single-core execution, Nebo code performs on par with hand-written C++, while for many-core (GPU) execution, Nebo code has shown >35x speedup over single-core CPU.

Christopher Earl
University of Utah
cwearl@cs.utah.edu

James C. Sutherland
Department of Chemical Engineering
The University of Utah
James.Sutherland@utah.edu

Matthew Might
The University of Utah
might@cs.utah.edu

MS73

Harlan: High-Level Programming for GPU Systems

Harlan is a high level language for data parallel computing that supports a rich set of features while shielding the programmer from the complexities of the hardware. A key feature is Harlan's region based memory system, which keeps related data located contiguously in memory and allows objects to move seamlessly between the CPU and GPU. Harlan's higher level features include algebraic data types, first class procedures, and recursive functions within kernels.

Eric Holk
Indiana University
eholk@cs.indiana.edu

Andrew Lumsdaine
Open Systems Laboratory
Indiana University
lums@cs.indiana.edu

Matthew Might
The University of Utah

might@cs.utah.edu

MS73

What Should Yafcc Look Like?

Abstract not available at time of publication.

Robert C. Kirby
Baylor University
Robert_Kirby@baylor.edu

MS73

Layered DSLs for Portable Manycore Scalability

One challenge facing a DSL for high performance computing (HPC) is achieving portable, scalable performance across emerging hardware architectures. A key difficulty is conflicting constraints on memory access. We propose using multiple embedded DSL layers. The lower level is responsible for addressing conflicting constraints in the changing hardware landscape, while higher levels focus on representing advances in the domain. We describe a preliminary finite element DSL layer on top of a parallel device DSL.

Daniel Sunderland, Patrick G. Xavier, H. Carter Edwards
Sandia National Laboratories
dsunder@sandia.gov, pgxavie@sandia.gov, hcedwar@sandia.gov

MS74

Visual Performance Analysis for the Exascale Era

Performance analysis of parallel scientific codes is becoming increasingly difficult due to the rapidly growing complexity of applications and existing tools fall short in providing intuitive views to reveal the root causes of performance problems. We have developed a new paradigm of projecting and visualizing performance data obtained from one domain onto other domains using projections and implemented in our BoxFish tool set.

Peer-Timo Bremer
Lawrence Livermore National Laboratory
ptbremer@gmail.com

MS74

Visual Characterization of High-End Computing

Modern supercomputers are complex, hierarchical systems consisting of huge numbers of cores, systems for disk storage, and nodes for I/O forwarding. These numbers continue to grow and the need for tools to understand the behavior of both the applications and system software becomes paramount. This talk presents visual analytics techniques that address the challenge of understanding the behavior of complex software on very large-scale compute platforms, like the current petascale computers, and shows the resulting tools can potentially tune that software to attain the highest possible degrees of efficiency.

Kwan-Liu Ma
Department of Computer Science
University of California at Davis

ma@cs.ucdavis.edu

MS74**Visions of Tau Dancing in Your Head - Ruminations on Performance Visualization**

Modern performance tools can capture arbitrarily complex information about the execution of parallel applications. Making "sense" of the performance data is as much a process of analysis, to extract high value content based on knowledge of what is "meaningful" about performance, as it is of presentation, to convey the relevant features of the results in forms that enable "understanding" and "insight" for users. The talk explores ideas for performance visualization in the TAU Performance System.

Allen Malony
University of Oregon
malony@cs.uoregon.edu

MS74**Effectively Presenting Application Performance Using Simple Techniques**

Call path profiling and tracing have been shown to provide insight into the performance characteristics of complex parallel programs. However, poor presentation of performance data obscures insight. To enable rapid analysis of an execution's performance bottlenecks, we describe a small set of simple presentation techniques. Using these simple techniques in concert has proven to be very effective. These techniques form the basis of the HPCToolkit performance tools.

Nathan Tallent
PNNL
Nathan.Tallent@pnnl.gov

MS75**Towards Parallel Multi-Scale Materials by Design**

The multiscale modeling of materials raises considerable challenges for the use of parallel computers in that it is necessary to combine different algorithms used at different physical scales. One way to achieve this combination is to make use of a task-based approach using directed acyclic graphs. In this case the methods at different levels may be thought of as sub-graphs. This approach is explored in conjunction with the Uintah software and its scalability considered for the combination of the Uintah material point method and a molecular dynamics code called Lucretius.

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

MS75**Title Not Available at Time of Publication**

Abstract not available at time of publication.

Christopher Johnson
University of Utah
Department of Computer Science

crj@sci.utah.edu

MS75**A Scalable, Efficient Scheme for Evaluation of Stencil Computations over Unstructured Meshes**

Stencil computations are a common class of operations that appear in many computational scientific and engineering applications. Stencil computations often benefit from compile-time analysis, exploiting data-locality, and parallelism. Post-processing of discontinuous Galerkin (dG) simulation solutions with B-spline kernels is an example of a numerical method which requires evaluating computationally intensive stencil operations over a mesh. Previous work on stencil computations has focused on structured meshes, while giving little attention to unstructured meshes. Performing stencil operations over an unstructured mesh requires sampling of heterogeneous elements which often leads to inefficient memory access patterns and limits data locality/reuse. In this talk, we present an efficient method for performing stencil computations over unstructured meshes which increases data-locality and cache efficiency, and a scalable approach for stencil tiling and concurrent execution. We provide experimental results in the context of post-processing of dG solutions that demonstrate the effectiveness of our approach.

Mike Kirby, James King
University of Utah
School of Computing
kirby@cs.utah.edu, king@sci.utah.edu

MS75**Stream Architectures for PDEs on Unstructured Grids**

Abstract not available at time of publication.

Ross Whitaker
School Of Computing, Scientific Computing and Imaging Inst.
University of Utah
whitaker@cs.utah.edu

MS76**Parallel Computation of Multivariate Polynomial Greatest Common Divisors**

Suppose we have two polynomials A and B in $n > 1$ variables with integer coefficients. What's the best way to compute their greatest common divisor G on a multicore computer? Polynomial GCD computation is important in practice because it is the bottleneck of many computations with polynomials. In this talk we will present the results of an experiment to parallelize Brown's GCD algorithm. We evaluate to bivariate polynomials and interpolate the remaining variables. So we do many bivariate GCDs in parallel. We will compare our Cilk implementation with the GCD routines in Magma and Maple.

Michael B. Monagan
Simon Fraser University
mmonagan@cecm.sfu.ca

MS76**Parallel Computation of Echelon Forms and Rank**

Profiles

We present recent advances in the parallelization of exact Gaussian elimination used to compute the echelon form or the rank profiles of a matrix over a finite field. We will first explore how far can these computations rely on the existing parallel Gaussian elimination techniques and which additional constraints are imposed on the pivoting strategies and the block decomposition in order to compute the echelon structure. In doing so we will propose slab and tile block iterative algorithms together with block recursive algorithms. One of these specificities is that some dimensions in the block decomposition are not known a priori but are revealed dynamically during the computation, thus leading to blocks of unevenly balanced dimensions. We then focus on implementations for parallel SMP architectures and try to adapt these algorithms with various work-stealing strategies: parallel loops; tasks and tasks with data-flow dependencies proposed by the runtimes of OpenMP and XKaapi. Experiments show a limited overhead compared to numerical parallel Gaussian elimination and demonstrate the advantage of the dataflow task dependencies approach, especially when dealing with blocks of unevenly balanced dimensions.

Clément Pernet

INRIA/LIG-MOAIS
Grenoble Université
clement.pernet@imag.fr

Jean-Guillaume L. Dumas
MNC-IMAG, UJF Grenoble
jean-guillaume.dumas@imag.fr

Ziad Sultan
INRIA/LIG-MOAIS
Grenoble Univ.
ziad.sultan@imag.fr

MS76

Symbolic Homotopies

I will show how symbolic homotopy techniques, possibly combined with algorithms for structured linear systems, can be used in order to solve systems of polynomial equations.

Eric Schost

Computer Science Department
Univ. of Western Ontario
eschost@uwo.ca

MS76

BPAS: A CilkPlus Library for Basic Polynomial Algebra Subroutines

BPAS provides arithmetic operations (multiplication, division, root isolation, etc.) for univariate and multivariate polynomials over prime fields or with integer coefficients. The currently available distribution focuses on dense polynomials and the sparse case is work in progress. The code is mainly written in CilkPlus targeting multicores. A strong emphasis is put on adaptive algorithms as the library aims at supporting wide range of problem sizes.

Yuzhen Xie

High Performance Engineering and Mathematics Software
Maplesoft
yxie@csd.uwo.ca

Changbo Chen
University of Western Ontario
changbo.chen@gmail.com

Farnam Mansouri
Computer Science Department and Applied Mathematics
Dept.
University of Western Ontario
fmansou3@csd.uwo.ca

Marc Moreno Maza, Ning Xie
Computer Science Department
University of Western Ontario
moreno@csd.uwo.ca, nxie6@csd.uwo.ca

MS77

Randomized Asynchronous Iterative Linear Solver for Spd Matrices

Asynchronous methods for solving systems of linear equations have been researched since Chazan and Miranker published their pioneering paper on chaotic relaxation in 1969. The underlying idea of asynchronous methods is to avoid processor idle time by allowing the processors to continue to work and make progress even if not all progress made by other processors has been communicated to them. Historically, work on asynchronous methods for solving linear equations focused on proving convergence in the limit. How the rate of convergence compares to the rate of convergence of the synchronous counterparts, and how it scales when the number of processors increase, was seldom studied and is still not well understood. Furthermore, the applicability of these methods was limited to restricted classes of matrices (e.g., diagonally dominant matrices). We propose a shared-memory asynchronous method for general symmetric positive definite matrices. We rigorously analyze the convergence rate and prove that it is linear and close to that of our method's synchronous counterpart as long as not too many processors are used (relative to the size and sparsity of the matrix). A key component is randomization, which allows the processors to make guaranteed progress without introducing synchronization. Our analysis shows a convergence rate that is linear in the condition number of the matrix, and depends on the number of processors and the degree to which the matrix is sparse.

Haim Avron

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

MS77

Randomized Algorithms for Dense Linear Algebra

We illustrate how dense linear algebra computations can be enhanced using randomized algorithms by performing less communication and/or computation. A first application concerns the solution of general or symmetric indefinite linear systems where randomization prevents the communication overhead due to pivoting. The second application is related to statistical condition estimation for least squares problems or linear systems. For each application, we provide accuracy and performance results on current hybrid multicore-GPU parallel machines.

Marc Baboulin

INRIA/University Paris-Sud

marc.baboulin@inria.fr

MS77

Employing Random Butterfly Transformation in Sparse Direct Solvers

In parallel sparse direct solvers dynamic pivoting can hamper highly scalable implementation. Various relaxed pivoting strategies, such as static pivoting or threshold pivoting, have been used to mitigate the problem. In this work, we investigate a statistical technique to avoid pivoting; this is based on Random Butterfly Transformation (RBT) using recursive butterfly matrices. We will compare the RBT method with the static pivoting approach in various performance metrics, including robustness, sparsity, and runtime.

Xiaoye S. Li
Lawrence Berkeley National Laboratory
Lawrence Berkeley National Laboratory
xqli@lbl.gov

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

Marc Baboulin
INRIA/University Paris-Sud
marc.baboulin@inria.fr

MS77

Implementing Randomized Matrix Algorithms in Parallel and Distributed Environments

Implementations of recently-developed randomized matrix algorithms have already proven competitive with traditional state-of-the-art algorithms for moderately-large problems stored in RAM. We describe the extension of these methods to computing solutions to L1/L2 regression problems in parallel/distributed environments that are more common in very large-scale data analysis applications. These algorithms take advantage of modern computer architectures to achieve improved communication profiles on MapReduce and on clusters such as Amazon EC2 that have high communication costs.

Michael Mahoney
Stanford University
Applied and Computational Mathematics
mmahoney@cs.stanford.edu

MS78

Co-Design Via Proxy Applications: MiniMD in Chapel

Chapel is an emerging parallel programming language whose design and development are being led by Cray Inc. As part of our effort, we have been porting DOE proxy applications to Chapel in order to study the design and performance of the language. In this talk, I'll describe recent work to express MiniMD in Chapel. I'll describe how MiniMD is expressed in Chapel, report on our findings so far, and list next steps.

Bradford L. Chamberlain
Cray Inc.
bradc@cray.com

Ben Harshbarger

University of Washington
harshb@cs.washington.edu

MS78

C++ Compiler Analysis and RAJA Programming Techniques for Performance Portability

Complex computational science applications typically realize a small fraction of theoretical peak system performance. This gap widens with each new hardware generation as performance increasingly relies on fine-grained multithreading, compiler optimizations, and data locality. The RAJA programming model enables portable high performance by decoupling architecture-dependent compiler directives, data layouts and parallel implementation choices from numerical algorithms using standard C++ language features. To help vendors address C++ compiler deficiencies, we have also developed the LCALS compiler analysis suite. In this talk, we describe hurdles to portable performance and demonstrate how RAJA and LCALS help address these issues.

Richard Hornung
Lawrence Livermore National Lab
hornung@llnl.gov

Jeff Keasler
Lawrence Livermore National Laboratory
keasler1@llnl.gov

MS78

Non-Traditional Approaches to Development of Multi-Scale Simulation Codes

We present strategies and technologies for implementing dynamic, multi-scale physical simulation codes. These implementations are based on fundamental capabilities such as scheduling, communication, asynchrony, concurrency, and caching that are typically provided by single-function, open source, software packages. Example implementations using web and cloud technologies, such as in-memory databases, will be presented.

Allen McPherson, Christopher Mitchell
Los Alamos National Laboratory
mcperson@lanl.gov, mitchell@lanl.gov

Kipton Barros
Center for Nonlinear Studies
Los Alamos National Laboratory
kbarros@gmail.com

MS78

Performance and Energy Comparisons between FPGA and GPU Implementations of Data Assembly

The Data Assembly (DA) stage in Finite Element Method (FEM) is usually the performance bottleneck. Accelerating DA with Field-Programmable Gate Array (FPGA) and Graphics Processing Units (GPUs) presents challenges due to FPGA's programming complexity and DA's mixed compute-intensive and memory-intensive workloads. This talk presents a study which employs a representative finite element mini-application, miniFE, to explore DA acceleration on FPGA and GPU by using OpenCL. FPGA and GPU implementations based on different code partitioning schemes are discussed. The resulting DA performance and

energy consumption are compared with CPU only platforms.

Li Tang, Xiaobo Sharon Hu
University of Notre Dame
ltang@nd.edu, shu@nd.edu

Richard Barrett
Sandia National Laboratories
rfbarre@sandia.gov

MS79

Plasma Simulations on the Intel Xeon Phi Coprocessor

In the past decade, advances in computing power have been driven by multi-core processors and, more recently, also by heterogeneous architectures like GPUs and Intel Xeon Phi. We will present our work on adapting the kinetic Plasma Simulation Code (PSC) to the Intel Xeon Phi technology. The PSC is an electromagnetic particle-in-cell code that was designed to have run-time swappable computational kernels and data structures that can take advantage of specific processing hardware. We will focus on two aspects of the our Xeon Phi port: (a) Implementing a hybrid OpenMP-MPI parallelization approach for parallelization within the cores of one Intel MIC chip and across multiple chips, respectively. (b) Adapting the computational kernels that move particles using interpolated fields, and deposit currents back onto the mesh to take advantage of the LRBng SIMD instructions using intrinsics.

Kai Germaschewski
Institute for the Study of Earth, Oceans, and Space
University of New Hampshire
kai.germaschewski@unh.edu

Narges Ahmadi
Space Science Center
University of New Hampshire
narges.ahmadi@unh.edu

Homayoun Karimabadi
University of California-San Diego
Department of Electrical and Computer Engineering,
MC0407
homakar@gmail.com

MS79

The Beacon Project: Application Experiences

This presentation introduces the Beacon Project, an ongoing project supported by NSF and The University of Tennessee that explores the impact of the Intel Xeon Phi coprocessor on computational science and engineering. The architecture of Beacon, the Cray CS300-AC cluster supercomputer that ranks #1 on the November 2012 Green500 list, is described, application experiences are discussed, and selected results are presented.

R. Glenn Brook
National Institute for Computational Sciences
University of Tennessee - Knoxville
glenn-brook@tennessee.edu

MS79

Lattice Quantum Chromodynamics on the Intel

Xeon Phi Coprocessor

We present a quick overview of the state of our lattice QCD implementation on the Xeon Phi. We will outline the approach taken to achieve good memory bandwidth utilization, vectorization, load balancing and message passing between several Xeon Phi devices running in native mode and present some application results.

Balint Joo
Jefferson Laboratory
bjoo@jlab.org

MS79

Experiences Moving Gpu-Accelerated Applications to the Intel Xeon Phi Coprocessor

This talk presents initial results of porting GPU-accelerated applications that are commonly run on leadership-class DOE supercomputers such as Titan to the Intel Xeon Phi Coprocessor.

Bronson Messer
Oak Ridge National Laboratory
& University of Tennessee
bronson@ornl.gov

MS80

Panel Discussion: Parallel Computing for Big Data

A panel, which will include the previous speakers, will discuss issues in parallel computing for "big data" problems. One topic is whether current distributed-memory architectures are suitable for network science and data analysis of highly irregular data, or whether emerging architectures (such as massively multithreaded systems) are required in this area. A related question is whether data-parallel programming models are still useful.

Erik G. Boman
Sandia National Labs, NM
Scalable Algorithms Dept.
egboman@sandia.gov

MS80

Scaling Techniques for Massive Scale-Free Graphs in Distributed (External) Memory

We present techniques to process large scale-free graphs in distributed memory. Our aim is to scale to trillions of edges, and our research is targeted at leadership class supercomputers and clusters with local NVRAM. We apply an edge list partitioning technique, designed to accommodate high-degree vertices that create scaling challenges when processing scale-free graphs. We present a scaling study with three important graph algorithms: Breadth-First Search, K-Core decomposition, and Triangle Counting.

Roger Pearce
Texas A&M and
Lawrence Livermore Lab
rpearce@tamu.edu

MS80

Scalable Matrix Computations on Large Scale-Free

Graphs Using 2D Graph Partitioning

One-dimensional (row-based) parallel distributions provide low load imbalance and communication costs for PDE-based data. For scale-free graphs, however, 1D distributions produce load imbalance, high communication volume, high message counts, and scalability loss at high core counts. Two-dimensional (nonzero-based) block distributions reduce message counts, but do not address load imbalance or communication volume. We present a new distribution that combines 1D graph partitioning (to reduce load imbalance and communication volume) with 2D distribution (to reduce message count). We demonstrate that these distributions enable scalable, fast sparse matrix-vector multiplication and eigensolves up to 16K cores.

Erik G. Boman
Sandia National Labs, NM
Scalable Algorithms Dept.
egboman@sandia.gov

Karen D. Devine, Siva Rajamanickam
Sandia National Laboratories
kddevin@sandia.gov, srajama@sandia.gov

MS80

Multilevel Graph Partitioning: a Combinatorial and QP-based Method

The graph partitioning problem arises in a variety of scientific and engineering domains such as VLSI design, computer networking, parallel computing, and sparse direct methods. In this talk, we discuss the edge separator graph partitioning problem and present an algorithm which combines aspects of multi-level combinatorial methods with a quadratic programming formulation. We discuss a novel strategies for both graph coarsening and refinement with the goal of producing high quality cut sets.

Nuri Yeralan
University of Florida
nuriatuf@gmail.com

Timothy A. Davis
University of Florida
Computer and Information Science and Engineering
DrTimothyAldenDavis@gmail.com

PP1

Adaptive Sampling for Large Scale Networks

Network sampling is an important strategy for reducing the size of massive graphs while retaining their important properties. Here, we introduce an adaptive sampling technique that takes as input the key properties of the graph and the analysis objective, and creates a sampled graph that best fulfills the analysis requirements. We demonstrate the effectiveness of our sampling technique by estimating parameters such as the degree and clustering coefficient distribution on large-scale networks.

Shehla Ahmed
University of Nebraska at Omaha
shehlaahmed@unomaha.edu

PP1

Fast Sparse Direct Solvers for Systems from Elliptic

Partial Differential Equations

We present a reduced-cost direct method for solving in parallel a large sparse symmetric positive definite system arising from the numerical solution of elliptic partial differential equations. This method provides a new formulation that does not require factorizations of sub-matrices associated with the largest separators of matrices ordered using nested-dissection. Our formulation reduces floating-point operations for a $K \times K$ 5-point finite-difference problem by 20%, and reduces execution time by 30% to 62% relative to Cholesky for our test suite.

Joshua D. Booth
The Pennsylvania State University
jdb5132@cse.psu.edu

Padma Raghavan
The Pennsylvania State Univ.
Dept of Computer Science Engr.
raghavan@cse.psu.edu

PP1

Performance Study of Parallel Octree-based Conforming Tetrahedral Mesh Generation

This work measures the performance impact from tetrahedralization of nonconforming meshes generated by the linear octree with 2:1 balancing constraint. Tetrahedralization technique based on templates is implemented. Performance measurements on the parallel generation of conformal meshes with the present method on 4096 cores show that it is possible to generate a mesh for a real life offshore platform with 9 billion of tetrahedra in less than 40s of CPU time.

Jose J. Camata
Federal University of Rio de Janeiro
camata@gmail.com

Igor Ghisi
ESSS
igor@esss.com.br

Alvaro Coutinho
Dept. of Civil Engineering
COPPE/Federal University of Rio de Janeiro
alvaro@nacad.ufrj.br

PP1

What Can the Roofline Model of Energy Tell Us About How to Build the Next Supercomputer?

The overarching goal of this research is to enable algorithm designers to create more energy efficient algorithms by providing the means of analyzing the relationship between time, energy and power on real systems. Firstly, we provide an analytical cost model for energy and power. This model expresses energy and power for algorithms using a small set of simple parameters (e.g., memory bandwidth, number of FLOPs). We validate our model using highly optimized ubenchmarks and a fine-grained power measurement tool on state-of-the-art systems. Secondly, we iteratively refined our model and extend its validation on over a dozen "candidate compute-node building blocks" for future high performance systems, including the latest x86, GPUs, and mobile SoCs. The purpose of this study is to provide a precise analytical characterization of abstract algorithmic regimes where one building block may

be preferable to others. Together, they provide the means to redesign algorithms for energy efficiency.

Jee W. Choi, Richard Vuduc
 Georgia Institute of Technology
 jee@gatech.edu, richie@cc.gatech.edu

PP1

HPCC and Undergraduate Applied Mathematics

Pairing a series of one-year NSF-funded undergraduate internships with a generous private grant, VMI has created a novel undergraduate interdisciplinary program, centered around High-Performance Computing Clusters (HPCC). We herein describe the motivation, strategy, planning, and implementation of this exciting initiative, which has applied mathematics and parallel computation as its focal points.

Randall E. Cone
 Virginia Military Institute
 conere10@vmi.edu

PP1

A Simple, Parallel Scheme for Support Graph Preconditioning of Networks

We consider the solution of linear systems corresponding to the Laplacian matrices of large unstructured networks. A promising approach to solving these problems is to use combinatorial preconditioners, such as support graph preconditioners. These are difficult to parallelize on distributed memory systems. We show that domain decomposition combined with local support graph preconditioning scales well with increasing number of processors. We implement this preconditioner in the Trilinos software library and provide scaling results.

Kevin Deweese
 UC Santa Barbara
 kdeweese@cs.ucsb.edu

Erik G. Boman
 Sandia National Labs, NM
 Scalable Algorithms Dept.
 egboman@sandia.gov

PP1

Evaluating the Impact of Silent Data Corruption in Numerical Methods

We present work towards sparse iterative solvers that assume the underlying machine is unreliable in unpredictable ways. We present the Skeptical Programmer, a programming technique that trusts no values, while exploiting knowledge of the inputs and numerical methods to derive invariants that allow one to assert whether values are theoretically possible, and reject values that are clearly not. We also demonstrate how to exploit the IEEE-754 floating-point specification to minimize silent errors in dot-products.

James Elliott
 North Carolina State University
 jjellio3@ncsu.edu

PP1

Scaling of Entropy Based Moment Closures Vs Pn

Moment Closures

In radiative transport, moment closures (spectral methods) are used to reduce the number of dimensions of the governing kinetic equations. The classic PN moment closure has several defects, one of which is non-physical negative density distributions of particles. New entropy based (EB) moment closures exist to correct some of these defects, but are computationally expensive to solve. However, EB moment closures have the following two nice properties.

- EB moment closures contain the same number of unknowns as the PN moment closure, so communication costs on compute clusters are the same for both methods.
- EB moment closures are computationally expensive locally on each grid cell, fitting the new paradigm of accelerators on compute clusters.

Therefore, in regimes where communication costs are dominant for the PN moment closure, EB moment closures can be used to get better accuracy for similar computational costs.

Charles K. Garrett
 University of Texas - Arlington
 garrettck@ornl.gov

Cory Hauck
 Oak Ridge National Laboratory
 hauckc@ornl.gov

PP1

Parallel Imperfection Study for Accelerating Cavities

Due to fabrication errors and operating mode tuning procedure, the measured physical parameters of a superconducting cavity differ from those of the designed ideal one. To solve this problem, we developed a parallel imperfection procedure: A mesh distortion method has been developed to generate deformed cavities in which the operating mode has to be tuned back to the designed frequency with field flatness across the cavity cells. This is achieved by adjusting the cavity shape using a shape optimization algorithm. The detail of optimization algorithm will be presented.

Lixin Ge, Liling Xiao, Cho-Kuen Ng
 Stanford Linear Accelerator Center
 lge@slac.stanford.edu, liling@slac.stanford.edu,
 cho@slac.stanford.edu

Zenghai Li
 Stanford Linear Accelerator
 lizh@slac.stanford.edu

PP1

Sparsity in Discontinuous Galerkin on Future Architectures

As scientific computing moves to massively parallel and heterogeneous architectures, managing memory operations bears more weight. In fact, optimizing these memory operations can gain precedence over floating-point operations. As these changes take place, there is concern over the potential performance gains and code complexity. The topic of sparsity addresses these concerns with a drastic impact on code performance. Furthermore, although taking advantage of sparse structure adds conceptual complexity, a

decrease in code length is found.

Zechariah J. Jibben

Arizona State University & Los Alamos National
Laboratory
zjibben@asu.edu

PP1

Treatment of a Lane-Emden Type Equation Via Second Derivative Backward Differentiation Formula Using Boundary Value Technique

This poster presents a second order nonlinear ordinary differential equations of Lane-Emden type solved using the boundary value technique. A second derivative backward differentiation formula is derived from a continuous multistep scheme using the multistep collocation technique. The technique transforms the numerical method to a system of non-linear equations, thereby obtaining numerical solutions concurrently on the entire range of integration by computer solvers. A simulation of the White-Dwarf Equation is demonstrated for a parameter c to show degeneration.

Ehigie Julius, Okunuga Solomon

Department of Mathematics, University of Lagos, Nigeria
jehigie@unilag.edu.ng, sokunuga@unilag.edu.ng

PP1

The Parallel Java 2 Library: Multicore and Cluster Parallel Programming in 100% Java

The Parallel Java 2 Library (PJ2) is an API and middleware for parallel programming on multicore parallel computers and cluster parallel computers, including clusters of multicore nodes, in 100% Java. PJ2 provides high-level parallel programming abstractions, such as parallel loops and reduction variables for multicore programming, and interprocess communication via tuple space for cluster programming. PJ2 also includes a lightweight map-reduce framework for big data computing, fully integrated with PJ2's multicore and cluster parallel computing capabilities. Examples of PJ2 programs, including performance measurements, will be presented.

Alan Kaminsky

Department of Computer Science
Rochester Institute of Technology
ark@cs.rit.edu

PP1

Design and Implementation of a Distributed Spatial Data Structure

The confluence of increasing design sizes and shrinking process node dimensions has reinvigorated research into high-performance layout analysis algorithms and implementations which can analyze full chip mask layout under strict turn-around-time (TAT) constraints. Modern VLSI designs routinely comprise of 10^{13} edges, and analysis often requires parallel processing. In this paper we describe a novel architecture of storing polygon geometry for large VLSI designs which is inspired from Partitioned Global Address Space (PGAS) paradigm. The problem domain which we shall highlight is derived from parasitic extraction, which is a subset of VLSI layout analysis problems. For this problem the PGAS architecture can be optimized for many-core as well as many-thread systems simultane-

ously.

Sandeep Koranne

Mentor Graphics Corporation
sandeep_koranne@mentor.com

PP1

GPU Accelerated Shake and Rattle Algorithms for Systems with Holonomic Constraints

Constraint models are often applied to molecular dynamics simulations. We developed a GPU algorithm for systems satisfying holonomic constraints. In our work we implemented an analog of the original SHAKE algorithm and its velocity version RATTLE using NVIDIA CUDA computer language. Several numerical experiments were performed to examine an acceleration and validity of the algorithms, including a simulation of a stochastic dynamics of lambda DNA. In our simulations, performed using NVIDIA Tesla C2075, the algorithm demonstrates significant acceleration (ten and more times for large systems) in comparison to serial CPU code.

Georgiy Krylov, Nursultan Zarlyk

Nazarbayev University
gkrylov@nu.edu.kz, nursultan.zarlyk@nu.edu.kz

PP1

Paralution - a Library for Iterative Sparse Methods on Multi-Core CPU and GPU

PARALUTION is a library which enables you to perform various sparse iterative solvers and preconditioners on multi/many-core CPU and GPU devices. Based on C++, it provides a generic and flexible design which allows seamless integration with other scientific software packages. PARALUTION contains Krylov subspace and Multigrid solvers, Fixed-point iteration schemes, Mixed-precision schemes and fine-grained parallel preconditioners based on splitting, ILU factorization with levels, multi-elimination ILU factorization and approximate inverse. The software is released as OpenSource.

Dimitar Lukarski

Dept. of Information Technology / Div. of Scientific
Computi
Uppsala University, Sweden
dimitar.lukarski@it.uu.se

Elisabeth Larsson

Uppsala University, Sweden
Department of Information Technology
Elisabeth.Larsson@it.uu.se

PP1

A New Scalable Parallel Algorithm for Fock Matrix Construction

Hartree-Fock (HF) calculations are widely used in quantum chemistry and are the starting point for accurate electronic correlation methods. Existing algorithms, however, may fail to scale for large numbers of cores of a distributed machine, particularly in the simulation of moderately-sized molecules. In existing codes, HF calculations are divided into tasks. Fine-grained tasks are better for load balance, but coarse-grained tasks require less communication. In this poster, we present a new parallelization of HF calculations that addresses this trade-off: we use fine-grained

tasks to balance the computation among large numbers of cores, but we also use a scheme to assign tasks to processes to reduce communication. We specifically focus on the distributed construction of the Fock matrix arising in the HF algorithm, and describe the data access patterns in detail. For our test molecules, our implementation shows better scalability than NWChem for constructing the Fock matrix.

Xing Liu, Aftab Patel, Edmond Chow
School of Computational Science and Engineering
Georgia Institute of Technology
xing.liu@gatech.edu, aypatel@gatech.edu,
echow@cc.gatech.edu

PP1

Scalable Parallel Simulated Annealing

We present a highly scalable parallel simulated annealing algorithm. Our algorithm periodically resamples the states in each processor to keep them as close to “thermal” equilibrium as possible while minimizing communication. Preliminary studies show encouraging results up to 192 processors on a Cray XE6 system with a problem having a 13-dimensional search space. Although motivated by a specific application to the reverse engineering of genetic networks, we expect the algorithm to have general applicability.

Zhihao Lou
Department of Computer Science
The University of Chicago
zhlou@uchicago.edu

Rick Stevens
Argonne National Laboratory
stevens@anl.gov

John Reinitz
Statistics Dept, Dept. of Ecology & Evolution
The University of Chicago
reinitz@galton.uchicago.edu

PP1

Parallel Scaling of Camellia for DPG Using BlueGene/Q

The discontinuous Petrov-Galerkin (DPG) finite element methodology first proposed in 2009 by Demkowicz and Gopalakrishnan offers a fundamental framework for developing robust residual-minimizing finite element methods. Camellia is a toolbox for *hp*-adaptive finite elements which we developed to facilitate DPG research and experimentation. In this poster, we examine the parallel performance of Camellia on the BlueGene/Q in the context of both linear (convection-diffusion) and nonlinear (Navier-Stokes) problems.

Nathan Roberts, Ramesh Balakrishnan
Argonne National Laboratory
nvroberts@alcf.anl.gov, bramesh@alcf.anl.gov

Jesse Chan
Rice University
jesse.chan@caam.rice.edu

Truman E. Ellis
ICES
University of Texas at Austin

truman@ices.utexas.edu

PP1

Feature Extraction for Hyperspectral Data Using Massive Parallel Processors

In this poster we introduce a new group of distributed algorithms for linear unmixing of hyperspectral images. The focus of the work is on the parallelization of recently introduced techniques such as Nonnegative Matrix Factorization (NMF). We present a theoretical classification of hyperspectral data processing algorithms based on the impact their parallel counterparts have. We provide an analysis of the newly developed algorithms and show how they are equivalent to their sequential counterparts. Finally we show how the distributed algorithms provide a significant computational speedup.

Stefan Robila
Montclair State University
stefan.robila@montclair.edu

PP1

Applying Software Agents to Resource Management in Tsunami Modeling

Complexity, versatility and high resource consumption of tsunami modeling systems require designing modular software, access to considerable amount of computational resources, and flexible resource management. Using supercomputers, our approach is focused on application of a Grid-like infrastructure. Specifically, we are adapting the agent-semantic infrastructure developed within the scope of the Agents in Grid (AiG) project to the needs of tsunami modeling. The aim of the presentation will be to summarize current state of our work.

Alexander P. Vazhenin, Kensaku Hayashi, Yutaka Watanobe
University of Aizu
Department of Information Systems
vazhenin@u-aizu.ac.jp, m5161111@u-aizu.ac.jp,
yutaka@u-aizu.ac.jp

Michal Drozdowicz, Maria Ganzha, Marcin Paprzycki
Systems Research Institute
Polish Academy of Sciences
michal.drozdowicz@ibspan.waw.pl,
maria.ganzha@ibspan.waw.pl,
marcin.paprzycki@ibspan.waw.pl

PP1

Performance Evaluation of the Most Program Package on Nvidia Kepler GPUs

Results of the parallelization are presented of the MOST (Method Of Splitting Tsunami) modeling package on the NVIDIA Kepler GPU architectures. The detailed analysis includes performance and bandwidth distributions among package components including inputting source data and output intermediate time slices. Experiments were provided on the detailed Tohoku Bathymetry (2413x2405, 200m) using the Great Japanese Earthquake parameters. The code optimization allowed obtaining the peak performance 1175 Gflop/s (Tesla K20c) and 1570 Gflop/s (Geforce GTX Titan.)

Alexander P. Vazhenin

University of Aizu
Department of Information Systems
vazhenin@u-aizu.ac.jp

Kazuya Matsumoto
University of Aizu, Japan
kazuya-m@u-aizu.ac.jp

Naohiro Nakasato
University of Aizu
Department of Computer Engineering
nakasato@u-aizu.ac.jp

Stanislav Sedukhin
University of Aizu, Japan
sedukhin@u-aizu.ac.jp

Kensaku Hayashi
University of Aizu
Department of Information Systems
m5161111@gmail.com

PP1

Large Scale Multi-Physics Simulation Using the Interoperable Executive Library

The complexity of simulating any system-wide processes in biomedical applications or thermal-fluid systems interactions goes beyond the reach of a single computer code. In this poster we present an Interoperable Executive Library (IEL) that has been designed to run, in parallel, a collection of multi-component physics simulations. The IEL is a light-weight integrator responsible for managing the distribution of data and memory, coordinates communication among parallel processes, and directs execution of a set of loosely coupled numerical and physics tasks. The workflow and computation of these simulations are detailed. The results obtained from running on darter (CRAY XC30 at NICS) and Keeneland (HPSSL250G8 and GPU) will be shown.

Kwai L. Wong
Joint Institute for Computational Science
University of Tennessee/ORNL
wong@jics.utk.edu

Andrew Kail
University of Tennessee
akail@utk.edu

David White
Maryville College
white.david@me.com

PP1

Parallel Implementation of a Contour-Integral-Based Eigensolver for Dense Generalized Eigenvalue Problems on Gpu Clusters

We consider a parallel eigensolver for dense generalized eigenvalue problems (GEP). A contour-integral-based method has been proposed for solving GEP and has coarse-grained parallelism which is suitable for many core clusters such as GPU clusters. In this poster, we show an implementation of the method for GPU clusters with MPI and MAGMA. A performance evaluation on HA-PACS, a GPU cluster equipped with 4 NVIDIA M2090s per node,

is shown.

Takahiro Yano, Yasunori Futamura, Tetsuya Sakurai
Department of Computer Science
University of Tsukuba
yano@mma.cs.tsukuba.ac.jp, futamura@mma.cs.tsukuba.ac.jp, sakurai@cs.tsukuba.ac.jp

PP1

A Highly-Scalable Multi-Bsp Fast Fourier Transform

Multi-BSP is an answer to increasingly hierarchical (NUMA) hardware. It models a parallel computer as a tree with internal nodes representing interconnects and leaves representing compute cores. We show an optimal FFT algorithm in flat BSP (where the Multi-BSP tree has depth one) can be applied recursively to obtain an attractive formulation for a generic, highly scalable, and easily implemented Multi-BSP FFT algorithm that is independent of the actual depth of the NUMA hierarchy.

Albert-Jan N. Yzelman

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

PP1

Investigation on Parallel Performance of Fluid Solvers Based on Different Numerical Methods

Fluid dynamics simulations need efficient numerical methods, which depend on many factors, such as expansion bases, mesh, algorithms and parallel models. In this poster, we will investigate different numerical methods for solving Navier-Stokes equation in some fluid simulations. Their performance will be compared in detail, which includes accuracy, speed, and parallel efficiency. Based on these, advantages and disadvantages of using different methods will be studied. Some simulation results will be presented for some applications.

Xiaohe Zhufu
Institute of Software, ISCACS
Chinese Academy of Science, China
xiaohe@iscas.ac.cn

Xiaotong Dong, Weishan Deng
Institute of Software, Chinese Academy of Science
dongxt0815@hotmail.com, 417408020@qq.com

Jin Xu
Institute of Software, ISCACS
Chinese Academy of Science, China
xu_jin@iscas.ac.cn