IP1

Joint JSIAM – SIAM Lecture: Climbing Steep Mountains: The Quest to Achieve Affordable, Scalable, Reliable and Energy Efficient Exascale Computing

We are embarked on a journey from terascale, through petascale, to exascale computing in the next decade. Twenty years after its early beginnings, massively parallel processing (MPP) is sufficiently mature that most problems can be successfully and effectively parallelized and run on an MPP computer with hundreds to thousands of threads. Today, a few applications even run at hundreds of thousands of threads. We are planning the path to exascale computing in an environment in which energy efficiency is putting unprecedented constraints on hardware and software architectural choices. The 1000-fold per decade increase in the speed of the fastest computers that we have witnessed over the past thirty years has been enabled by but is not been entirely due to Moores Law increases in transistor capabilities. Rather, this speed increase has also been accompanied by increases in system size, power dissipation, and cost over time. Fortunately capability increases have greatly outstripped cost, power, and size increases. By 2018, we will be utilizing process technology with features a couple of dozen silicon atoms across! So, transistor density will be much higher. However, transistor efficiency is already beginning to plateau: voltage scaling is becoming ever more difficult; logic reliability is becoming more challenging, existing DRAM designs will not scale appropriately. Due to costs and purported environmental issues, data centers are requiring that we limit exascale computers to 30 MW or less. When we translate this into requirements on circuit technology, we see that we are being driven to create designs that are extremely simple, that have low inherent reliability, that use slower clock rates than we can achieve, and sustain performance through unprecedented parallelism. An exascale computer will have O(1056) processors, each with O(103) cores for a total of 1089 hardware threads. This means that we will need to develop new software stacks, applications tools, and algorithms that are able to profitably employ billion-way parallelism. In addition, without breakthroughs in reliability, the system would be facing constant application level interrupts. This will indeed be a steep climb. In this talk, I will map a route to successful attainment of this summit in computing.

William J. Camp
Intel Corporation
william.j.camp@intel.com

IP2

Joint JSIAM – SIAM Lecture: GPU Acceleration: A Fad or the Yellow Brick Road onto Exascale?

Since the first commodity x86 cluster Wigraf achieving paltry 10s 100s Megaflops in 1994, we have experienced several orders of magnitude boost in performance. However, the first Petaflop was achieved with the LANL RoadRunner, a Cell-based "accelerated" cluster, and in 2010 we may see the first (GP)GPU-based cluster reaching Petaflops. Do such non-x86accelerator merely push the flops superficially, or are they fundamental to scaling? Based on experiences from TSUBAME, the first GPU-accelerated cluster on the Top500, we show that GPUs not only achieve higher performance but also better scaling, and in fact their true nature as multithreaded massively-parallel vector processor would be fundamental for Exascale. Such results are being reflected onto the design of TSUBAME2.0 and its successors.

Satoshi Matsuoka
Tokyo Insitute of Technology
matsu@is.titech.ac.jp

IP3

Killer Micros II: The Software Strikes Back

High Performance Computing has been completely transformed by The Attack of the Killer Micros, a prediction by Eugene Brooks at Supercomputing 89 that vector and other exotic HPC hardware would soon be supplanted by mainstream microprocessors. He was right, of course, but the story isnt over. Now that parallel computing has entered the mainstream, client-intended parallel software will redefine and dominate the HPC space just as client-intended hardware did. This talk will review the sad state of todays HPC software and discuss some of the changes in store.

Burton Smith
Microsoft
burtons@microsoft.com

IP4

HPC Programming Models: Current Practice, Emerging Promise

In this talk, I will give an overview of parallel programming models for high performance computing (HPC). I'll begin with a brief overview of today's dominant models: MPI and OpenMP. Then I'll introduce Partitioned Global Address Space (PGAS) languages which strive to simplify programming while supporting scalability on large-scale machines. I'll describe traditional PGAS languages as well as emerging ones such as Chapel, Cray's new high-productivity language. For each model, I'll point out what I consider to be its strengths and weaknesses.

Brad Chamberlain
Cray Inc.
bradc@cray.com

IP5

Graph Analysis with High-Performance Computing

Traditionally, High-Performance Computing (HPC) advances have been driven by scientific computing applications. These typically involve physics simulation datasets, and the community has had great success in processing them using distributed-memory supercomputers. These successes have been based upon the validity of certain crucial assumptions. However, we have had to revisit these assumptions in the context of a relatively new domain: HPC graph analysis. This type of analysis is exerting pressures on computer architectures, programming models, and algorithm design. We will discuss these pressures and some current work to address them.

Jonathan Berry
Sandia National Laboratories
jberry@sandia.gov

IP6

Photonic On-Chip Networks for Performance-

Energy Optimized Computing

The emergence of multicore architectures and chip multiprocessors (CMPs) signifies a profound transformation in the drive toward improved computing performance. Whereas in the past progress of computing systems was largely dominated by the underlying acceleration in microprocessor performance, systems today and in the future are increasingly bound by their communications infrastructure and the associated power dissipation of high-bandwidth data movement. Recent advances in chip-scale silicon photonic technologies have created the potential for developing optical interconnection networks that can offer highly energy efficient communications and significantly improve computing performance-per-Watt.

Keren Bergman
Columbia University
bergman@ee.columbia.edu

IP7
Computational Finance on GPUs

Computational finance is one of the fastest-growing areas of scientific computing, with more than 10% of the Top 500 supercomputers being used for financial applications. Also, very recently GPUs have emerged as a powerful new technology for cost-effective and energy-efficient HPC. In this talk I will give an overview of the Monte Carlo and finite difference applications used in computational finance, the reasons why I believe that GPUs will have a major impact in scientific computing for the next decade, and discuss the particular suitability of GPUs for financial Monte Carlo simulations.

Michael B. Giles
Mathematical Institute
Oxford University
Mike.Giles@maths.ox.ac.uk

IP8
Energy Efficient Computing - From Bits to Buildings

In a recent survey by IDC, facilities managers named power and cooling by an overwhelming majority to be the most pressing issues of concern to them. A study of Exaflops computing came to the conclusion that by projecting todays technology, an Exaflops computer might require 120 MW of power, if it can be built at all. A different study commissioned by the EPA estimates that power consumption by servers doubled in the period from 2000 to 2005 worldwide, and that the total amount of electricity consumed by servers world wide now costs about $7.2B. This is already today the same order of magnitude as the investment in HPC technology$(9.2B). We have thus reached a critical threshold that should give us cause to consider the question of power consumption as a potentially limiting factor to the future growth in HPC. I will try to address this very question: what are the power limitations of current technology, and how can we change the equation to assure the future rapid growth of HPC performance without contributing even more to carbon emissions and global warming. In particular, I will discuss several research projects that we have started in Berkeley to address the issue of reducing power consumption in HPC, both at the systems and at the building level.

Horst D. Simon
Lawrence Berkeley Laboratory
hdsimon@lbl.gov

CP1
Comparative Studies of Parallel Preconditioner in FreeFem++

FreeFem++ is a free software which can be used to solve partial differential equation of dimension 1,2 or 3 in parallel computer. In FreeFem++, the parallelism is achieved in the linear solvers used to solve systems of linear equations obtained after discretization of PDEs. Moreover, it is also possible to express domain decomposition algorithms. This talk focuses on the usage of FreeFem++ on several 3-dimensional problems.We also present a comparative study between parallel preconditioners interfaced in FreeFem++ such as RAS(Restricted Additive Schwarz), Schur Complement, AMG (Algebraic MultiGrid), $HILU$ (Multistage Incomplete LU) and PAINV (Parallel Approximate Inverse).

Atenekeng Kahou Guy Antoine
INRIA Saclay, 4 Rue Jacques Monod
91893
ateneken@lri.fr

Morice Jacques, Hecht Frederic
Jacques-Louis Loins Lab, University of Paris 6
jacques.morice@gmail.com, hecht@ann.jussieu.fr

Grigori Laura
NRIA Saclay, 4 Rue Jacques Monod 91893
laura.grigori@inria.fr

Nataf Frederic
Jacques-Louis Loins Lab, University of Paris 6
nataf@ann.jussieu.fr

CP1
Performance and Scalability Analysis of HYCOM Ocean Model on PARAM AMD and PARAM Yuva Cluster

In our research work, we have observed the behavior of HYCOM model as we change the implementation from message passing (MPI) to shared memory (OpenMP) as we change the problem size and with large number of cores on multi-core computing architectures. We have also examined the communication characteristics of the model and the bottlenecks within the code. This will help system designers and application developers to better understand the communication workloads in climate modeling applications.

Kunal G. Rao
University of Florida
kunalrao@ufl.edu

CP1
Numerical Simulation on the SiCortex Supercomputer Platform: a Preliminary Evaluation

We investigate the potential of the SiCortex platform for numerical simulation by analyzing the performance of a set of elementary benchmarks and two fluid dynamics applications executed on the SC5832, SC072 systems and on an Intel Xeon-based commodity cluster. The focus of the evaluation is computational performance, but we also consider the energy consumption for both machines.

Björn Rocker, Staffan Ronnas, Vincent Heuveline
Institute for Applied and Numerical Mathematics

Karlsruhe Institute of Technology, Universität Karlsruhe
bjoern.rocker@kit.edu, staffan.ronnas@student.kit.edu, vincent.heuveline@kit.edu

CP1

Analysis of a Hierarchical Parallelization for Multiphysics Simulations

MPhyScas is a framework specially designed for the automatic development of multi-physics simulators. It is now supporting the development of parallel simulators through a deep hierarchical parallelization of its structural architecture. A prototype was built in order to observe and analyze the behavior of different solution algorithms to the same coupled problem when implemented in simulators produced by this prototype. The analysis is made based on performance issues and their relation to the hierarchical parallelization.

Felix G. Santos
Universidade Federal de Pernambuco
Departamento de Engenharia Mecnica
flxcgs@yahoo.com.br

CP1

Scaling Limits and Strategies for the Parallel Barnes-Hut Tree Code Pepc

We analyze the scalability and communication structure of our Barnes-Hut tree code PEPC on an IBM BlueGene/P. Although the code shows reasonable scaling behavior up to 8192 processors the current parallelization strategy for the tree construction and the exchange of multipoles suffers from intrinsic bottlenecks and algorithmic issues, which are analyzed and discussed in this talk. We define a methodology of evaluating the communication efficiency within the code and discuss its impact on the scalability.

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

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

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

Gudula Ruenger
Chemnitz University of Technology, Germany
ruenger@cs.tu-chemnitz.de

CP2

Fast and Scalable Algorithms for Comparing Genetic Sequences

Comparison of genetic sequences(DNA), used for molecular based classification of organisms, is a computation and data intensive process. Recent research has focused on comparisons using non-alignment methods, such as the longest common subsequence(LCS) and Lempel-Ziv(LZ) complexity, which overcome the drawbacks of the more common sequence alignment approach. However, there are very few parallel algorithms for non-alignment techniques. In this talk, we present efficient scalable algorithms for LCS and LZ and evaluate their performance on mitochondrial genome sequences.

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

Dhundy Bastola
Department of Computer Science
University of Nebraska, Omaha
dkbastola@mail.unomaha.edu

Sriram Surapaneni
Department of Computer Science
University of Nebraska, Omaha
ssurapaneni@unomaha.edu

Hesham Ali
Department of Computer Science
University of Nebraska, Omaha
hesham@unomaha.edu

CP2

Exact Method for Resolving the Q3ap Problem on Calculation Grid

Quadratic assignment to three dimensions (Q3AP) is one of the most difficult of problems Combinatorial Optimization, this problem is NP-complete and has several applications in data transmissions. Its exact resolution pass by the enumeration of a very broad search tree that contains a billion summits for instances of size medium. The resolution optimally large instances of the problem require implementing in Work of the complex methods requiring more than power of calculation. Currently, with the growth of calculation grids, many parallel models for exact methods have been proposed. For a good grids environmental exploitation, the initial problem is divided into many units of work. These are then distributed on a thousand of processors on the grid. The objective is to push the more far as possible the exact resolution Q3AP. Large scale tests on the Q3AP are still in perspective.

Sahraoui Sabrina
Laboratory MISC Mentouri University of Constantine
sahraouipgs@yahoo.fr

DR.KHOLLADI Mohamed-Khireddine
Director of laboratory MISC
Mentouri University of Constantine Algeria
kholladi@yahoo.fr

PROF.MELAB Last Name, PROF Talbi El-Ghazali
LIFL-INRIA Lille France
nouredine.melab@lifl.fr, el-ghazali.talbi@lifl.fr

CP2

Parallel Algorithms for Black-Box Optimization

Black-box optimization consists in optimization involving functions without exploitable properties. Such black-box functions can be nonsmooth, noisy, or expensive to evaluate. They are usually governed by a computer simulation. This talk describes the MADS (Mesh Adaptive Direct Search) method for such problems, and the ways to parallelize the algorithm. The NOMAD software will also

be presented.

Sebastien Le Digabel
GERAD
Sebastien.Le.Digabel@gerad.ca

Charles Audet
École Polytechnique de Montréal - GERAD
Charles.Audet@gerad.ca

John E. Dennis, Jr.
Rice University
dennis@caam.rice.edu

CP3

Using Fault Tolerance as a Basis for a Cooperative/preemptive-Multitasking Environment for Multicore Processors

We show how to extend task-based-fault-tolerance into a cooperative/preemptive-multitasking environment for multicore processors. In this system, each task knows how to redistribute its workload with the loss or the addition of some cores. A scheduler can preempt a task on some of its cores by creating faults on those cores. The preempted task cooperates by redistributing its workload among its remaining cores and continues working. The freed cores are assigned to a new task.

James R. Geraci
Square Enix Research and Development Division
geraci@square-enix.com

CP3

An Evaluation of Graph Partitioning and Ordering on a Massively Parallel Computer

We present the results of an experimental study of ParMETIS and PT-SCOTCH parallel graph partitioning and sparse matrix reordering packages on up to 2048 nodes of Blue Gene/P. We study strong and weak scaling of partitioning and ordering times, memory scaling, and the quality of partitions (in terms of edge-cut and load-balance) and ordering (in terms of factorization Flops). Our results show the relative strengths and weaknesses of the two packages at different levels of parallelism. They provide guidance to potential users and point to directions for improvements in the packages.

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

CP3

Second-Order Treatment of the Interface of Domain Decomposition Method

Estimations for the values on the interface are necessary in a domain decomposition method. However the accuracy of the estimations is of the first order for the most unconditionally stable domain decomposition schemes. In this paper, a second order of accuracy for the estimations on the interface is presented. It has been shown that the new method is unconditionally stable and efficient. Moreover, the optimal number of decomposed subdomains has been proposed in this paper. The optimal over-relaxation parameter is also studied when the SOR method is used for the iterations.

Younbae Jun
University of West Alabama
yjun@uwa.edu

Tsun-Zee Mai
University of Alabama
tmai@bama.ua.edu

CP3

Software Architecture of a Parallel Framework for SPH

We present the software architecture of a parallel framework for Smoothed Particle Hydrodynamics (SPH), written in *Python* and *Cython*. It consists of a *Particle Kernel* which provides an API for particle operations and an *SPH framework* which uses the particle kernel. We show how to setup fairly complex simulations with small Python scripts and how to extend it to incorporate different SPH formulations easily. We will release our work as an *open source* package.

Chandrashekhar P. Kaushik
Department of Computer Science and Engr.
Indian Institute Of Technology, Bombay
shekhar.kaushik@iitb.ac.in

Prabhu Ramachandran
Department of Aerospace Engineering
Indian Institute of Technology, Bombay
prabhu@aero.iitb.ac.in

CP3

Towards Standardizing Parallel Visualization of Scientific Data

Visualization of large quantities of scientific data remains a challenge. Parallel rendering can increase the effectiveness of visualization, but a standard for describing the data to be visualized is also needed so that data from different applications can be visualized consistently. We describe a system of metadata, called VizSchema, that self-describes scientific data in terms of how it can be visualized, with applications to parallel visualization tools such as VisIt.

Seth Veitzer, Svetlana Shasharina
Tech-X Corporation
veitzer@txcorp.com, sveta@txcorp.com

John Cary
Tech-X Corporation
University of Colorado
cary@txcorp.com

Marc Durant, Scott Kruger
Tech-X Corporation
mdurant@txcorp.com, kruger@txcorp.com

CP3

Extending Unified Parallel C for GPU Computing

Unified Parallel C (UPC) is a partitioned global address space (PGAS) parallel programming language for large scale distributed-memory computers. High performance GPUs typically have local memory separate from CPU memory so it is a very challenging task to manage memory

allocation and communicate data efficiently across many GPUs and CPUs. New UPC extensions provide a shared memory address space and a set of unified communication primitives to help scientific computing users program GPU clusters productively.

Yili Zheng, Costin Iancu, Paul Hargrove, Seung-Jai Min
Lawrence Berkeley National Laboratory
yili.zheng@gmail.com, cciancu@lbl.gov,
phhargrove@lbl.gov, sjmin@lbl.gov

Katherine Yelick
Lawrence Berkeley National Laboratory
University of California at Berkeley
yelick@eecs.berkeley.edu

CP4

Semi-Stencil Algorithm: Improving Data Locality and Reuse in Stencil Computation

Stencil computation is widely used in PDE+FD solvers, and it consists on accumulating the contribution of the neighbor points along the cartesian axis. The stencil computation performance is hampered mainly by two problems: the memory access pattern and the low data reuse. We introduce the semi-stencil algorithm which tackle both problems. It is implemented in homogeneous and heterogeneous multicores. Results are promising, for instance the algorithm reaches up to 50% of the Cell/B.E peak performance.

Mauricio Araya, Raul De La Cruz
BSC
mauricio.araya@bsc.es, raul.delacruz@bsc.es

CP4

Parallel Implementation of a Mixing Plane Model on Unstructured Meshes

A parallel implementation of a mixing plane model for unstructured arbitrary meshes is presented. The model performs circumferential averages at interfaces between rotors and stators. The mesh partition does not take into account any geometric aspects and therefore each interface can be scattered over several subdomains. Communication and load balance issues are addressed. Application to the flow calculation in a 5-stage high performance axial-flow compressor is presented.

Dulceneia Becker, Joo Barbosa
Gas Turbine Group
Instituto Tecnológico de Aeronáutica, ITA
dulceneia.becker@gmail.com, barbosa@ita.br

CP4

Multicore-Tailored Implementation of s-Step Methods for High-Dimensional PDEs

Accurate solution of time-dependent, high-dimensional PDEs requires massive-scale parallel computing. We describe an implementation framework for multi-block cartesian grids, optimized for clusters where the nodes have one or more multi-core processors. In particular, we focus on algorithms that minimize the impact of global communication by advancing the gridfunction several steps at once, so-called s-step methods. We present simulation results from quantum chemistry, where the time-dependent Schrdinger equation is solved numerically.

Magnus Gustafsson
Uppsala University, Sweden
magnus.gustafsson@it.uu.se

Sverker Holmgren
Uppsala University
Department of Scientific Computing
sverker@it.uu.se

CP4

An Efficient Parallel Implementation of a New High-Order Method for the Numerical Solution of Heterogeneous Models

A new high-order spatial method is proposed for the numerical solution of heterogeneous parabolic and hyperbolic models in three dimensions. This numerical procedure involves a spectral discretization combined with a Fourier continuation method. Besides, the classical Alternating Direction Implicit technique is used for the approximation of the time derivatives. The accuracy of the discrete algorithm and the efficiency of its parallel implementation are illustrated with some numerical simulations for heat transfer and wave propagation problems.

Andres Prieto, Oscar Bruno
California Institute of Technology
maprieto@acm.caltech.edu, bruno@acm.caltech.edu

CP5

Parallel Solver Coupling Multigrid and Direct Methods for Maxwell Equations

We develop a parallel solver for Maxwell equations using Nedelec's first order finite element discretization. Due to the wavelength restriction, simulation of large 3D objects requires a huge amount of computations which are hardly handled by direct methods. On the other hand, iterative solvers tend to converge slowly or are not accurate enough to meet our requirements. To overcome these problems, an hybrid solution is developed : we rely on a Full Multigrid scheme using a parallel direct solver on the coarsest level and a parallel matrix free iterative smoother (Jacobi) on finer meshes. This setup leads to a solver with both direct method accuracy and iterative method memory requirements.

Mathieu Chanaud
INRIA Bordeaux Sud-Ouest
chanaud@labri.fr

David Goudin, Jean-Jacques Pesque
CEA/CESTA
david.goudin@cea.fr, jean.pesque@cea.fr

Jean Roman
INRIA Bordeaux Sud-Ouest and LaBRI,
Universite de Bordeaux, Talence, France
roman@labri.fr

CP5

Exploiting Parallelism in Matrix-Computation Kernels for Symmetric Multiprocessor Systems: Matrix-Multiplication and Matrix-Addition Algorithm Optimizations by Software Pipelining and

Threads Allocation

We present a simple set of high-level algorithmic principles applied to fast matrix multiply (MM) algorithms for multicore systems. We have four contributions: we present a complete performance overview for the state-of-the-art systems; we introduce new implementations for 3M and Winograd's (MM) (e.g., with up to 20% speedup w.r.t. GotoBLAS MM); we apply software pipelining and threads allocation to all fast MM achieving further 5–7% speedup; we present a practical evaluation of the error analysis.

Paolo D'Alberto
Yahoo! Inc.
pdalbert@yahoo-inc.com

Marco Bodrato
University of Rome tor Vergata
bodrato@mail.dm.unipi.it

Alexandru Nicolau
University of California at Irvine
nicolau@ics.uci.edu

CP5

Block Sparse Householder Decomposition

This paper describes Householder reduction of a rectangular sparse matrix to small band upper triangular form. Using block Householder transformations gives good orthogonality, is computationally efficient, and has good potential for parallelization. The algorithm is similar to the standard dense Householder reduction used as part of the usual dense SVD computation. For the sparse algorithm, the original sparse matrix is accessed only for sparse matrix dense matrix (SMDM) multiplications. For a triangular bandwidth of $k + 1$, the dense matrices are the k rows or columns of a block Householder transformation. Using an initial random block Householder transformation allows reliable computation of a collection of largest singular values. Current implementation is for a large memory multicore processor.

Gary Howell
North Carolina State University
gwhowell@unity.ncsu.edu

CP5

A Block Fsai-Ilu Parallel Preconditioner for Spd Linear Systems

A novel preconditioner is presented coupling a Factorized Sparse Approximate Inverse (FSAI) with an ILU factorization. FSAI is modified so that the $N \times N$ preconditioned matrix acquires a block-diagonal structure with the number of blocks $\geq$ the number of processors n_p. Then, each processor applies an ILU factorization to its own blocks. This preconditoner is fully scalable with an intermediate performance between ILU ($n_p = 1$) and FSAI ($n_p = N$), thus proving always superior to FSAI on any parallel computer.

Carlo Janna
DMMMSA - University of Padova
janna@dmsa.unipd.it

Massimiliano Ferronato, Giuseppe Gambolati
University of Padova
DMMMSA
ferronat@dmsa.unipd.it, gambo@dmsa.unipd.it

CP5

Parallel Numerical Solution of the Time-Harmonic Maxwell Equations

We develop a parallel implementation of a scalable numerical solution to the time-harmonic Maxwell equations. The linear systems are derived from finite element discretization of the mixed formulation. Our approach is based on a diagonal preconditioner, an algebraic multigrid method and a new auxiliary space preconditioning technique. We apply our implementation to complicated domains. Numerical experiments demonstrate the scalability of our implementation.

Dan Li
The University of British Columbia
danli@cs.ubc.ca

Chen Greif
Department of Computer Science
The University of British Columbia
greif@cs.ubc.ca

Dominik Schoetzau
Mathematics Department
University of British Columbia
schoetzau@math.ubc.ca

CP6

Development of a Framework for Adaptive Mesh Refinement

We will present an approach of patch-based adaptive mesh refinement (AMR) for two- and three-dimensional simulations on future machines. Our result is equivalent to cell-based AMR, while physics packages are implemented on uniform structured meshes. The refined regions are able to keep original symmetries of physics problems. The features of our AMR framework include any number of levels of refinement, very minimum communications, and efficient bookkeeping that scales up with the number of computer processors.

William W. Dai
Los Alamos National laboratory
dai@lanl.gov

CP6

Large-Scale Dynamically Adaptive Structured Amr

We present scalable structured AMR algorithms using distributed mesh metadata, which we demonstrate on up to 36K processors. Structured AMR composes meshes from structured grids to avoid the cost of maintaining point connectivity. Metadata (brief descriptions of the grids and their processor owners) helps determine how data is transfered between grids. However, maintaining global metadata grows with the number of processors, leading to poor scalability. Distributed metadata scales but requires entirely new algorithms.

Brian Gunney
Lawrence Livermore National Lab
gunney1@llnl.gov

CP6

A Parallel Adaptive Finite Element Method for Large Scale Computations

We present a general framework for adaptive FEM, suitable for modern multicore and massively parallel architectures, based on highly concurrent and fully distributed algorithms, with local mesh refinement by recursive edge bisection based on a posteriori error estimation. To obtain good parallel efficiency, a dynamic load balancer redistributes the data prior to any mesh adaption. We present a performance study, comparing scaling and efficiency running on a BlueGene/L system.

Niclas Jansson
Computational Technology Laboratory, CSC/NA
Royal Institute of Technology
njansson@csc.kth.se

CP6

Front Tracking Method and Its Various Applications

Moving interfaces are involved in many systems in engineer and biology. Front tracking method is one of the most accurate and efficient computational approaches for studying such systems. A main challenge of developing front tracking algorithms is to capture the interface topological changes, especially in three dimensions. In this talk I shall introduce an improved three-dimensional front tracking method, of which the robustness and accuracy was shown with some testing examples. In addition to its applications in the turbulent mixing, several other potential applications of front tracking method in the developmental biology and cell migration will also be discussed.

Xinfeng Liu
University of South Carolina
xfliu77@gmail.com

CP6

Some Computational Results with *hp*-Adaptive Refinement

The *hp*-adaptive version of the finite element method has been receiving increased attention. In this method, local adaption occurs in both the element size, h, and the element polynomial degree, p. A crucial question is how to determine whether refinement of an element should be by h or p. In this talk we will present some computational results using several strategies for *hp*-adaptive refinement in the parallel program PHAML.

William F. Mitchell
NIST, Gaithersburg, MD
william.mitchell@nist.gov

CP6

Interface Tracking on Dynamic Quadrilateral and Rectilinear Grids

Moving adaptive grids for problems of interface capturing is an active field of research these days. A number of attempts were made to use moving grids in problems of dynamic flows where regions of singularities are moving and need to be zoomed in to be resolved with less numerical error. In two phase flow problems unstructured tetrahedral mesh refinement is popular as well as structured adaptive mesh refinement (AMR). We investigate the possibility of using a dynamic adaptive curvilinear quadrilateral grid and a dynamic adaptive rectilinear grid in the problems of interface capturing. Structured adaptive grids have a number of advantages over unstructured grids: fast grid generation, low memory cost, easy parallelization of numerical methods, multigrid acceleration, trivial to scale, and easy to fit into structured AMR framework. We present interface tracking Level Set (LS) and Conservative Level Set (CLS) methods coupled with dynamic quadrilateral and rectilinear grid adaptation. The Conservative Level Set method is implemented in a more efficient way than it was done before. Elliptic grid generation uses a novel monitor function that minimizes the local truncation error which is based on the numerical scheme for advecting an interface. The LS and CLS methods on dynamic quadrilateral and rectilinear grids are implemented and tested on a number of 2-d interface capturing problems and on incompressible two phase flow problems.

Svetlana Simakhina, Mark Sussman
Florida State University
ssimakhi@math.fsu.edu, sussman@math.fsu.edu

CP7

A Parallel Cartesian Treecode Algorithm for Ewald Summations in Molecular Dynamics Simulations

We will present results on ongoing work to develop a parallel treecode algorithm for Ewald summation. We have adapted and developed algorithms based on both replicated data strategy and domain decomposition. We hope to offer comparisons with parallel smooth particle mesh Ewald as implemented in DL_POLY_2 and DL_POLY_3.

Henry A. Boateng
Department of Mathematics
University of Michigan
boateng@umich.edu

Robert Krasny
University of Michigan
Department of Mathematics
krasny@umich.edu

Eitan Geva
University of Michigan
Department of Chemistry
eitan@umich.edu

CP7

The Parallelized Tree-Code and Its Application to Biomolecular Simulation

The Tree-code is a fast algorithm to calculate particle-particle interaction such as Coulombic and screened Coulombic interactions by using hierarchical tree structures. The algorithm uses a far-field Cartesian Taylor expansion and has several appealing features such as low memory requirement, relatively simple implementation, and capability to parallelization. The tree-code algorithm and its parallelization, as involved in the boundary integral formulation, can significantly improve the performance of the Poisson-Boltzmann solver in biomolecular simulation and dynamics.

Weihua Geng
Department of Mathematics
University of Michgan
geng@umich.edu

Robert Krasny
University of Michigan
Department of Mathematics
krasny@umich.edu

CP7

Plasma Simulation Code Optimization on Petascale Systems

The plasma simulation tool VORPAL, developed by Tech-X Corporation and the University of Colorado at Boulder, is used by many DOE researchers on leadership class supercomputers. In this presentation, we will describe optimization efforts that were undertaken to port this code to petascale systems, ranging from messaging analysis and optimization to single node memory tuning. This work is supported by US Department of Energy, Office of Science Phase II SBIR Grant DE-FG02-07ER84731 and Tech-X Corporation.

Peter Messmer, Paul Mullowney, Keegan Amyx, Ben Cowan
Tech-X Corporation
messmer@txcorp.com, paulm@txcorp.com, amyx@txcorp.com, benc@txcorp.com

Boyana Norris
Argonne National Lab
boyana@mcs.anl.gov

CP7

Biological Neuronal Networks on Nvidia Cuda Graphics Processors

Biological neuronal networks with thousands of neurons create systems with hundreds of thousands of equations that need to be updated at every time step. The structure of the system of ordinary differential equations that govern the cell dynamics lends itself to data parallel programming with CUDA. We are able to run neuron systems with thousands of cells in parallel on the GPU at a fraction of the time required for serial simulations on a workstation.

Joseph A. Rhoads
Florida State University
Department of Mathematics
joseph.rhoads@gmail.com

Richard Bertram
Department of Mathematics
Florida State University
bertram@math.fsu.edu

Joel Tabak
Dept of Biological Sciences
Florida State University
joel@neuro.fsu.edu

Michael Mascagni
Computer Science/Computational Science
Florida State University, Tallahassee
mascagni@.fsu.edu

Gordon Erlebacher
Florida State University
Department of Computational Science
gerlebacher@fsu.edu

CP7

Efficiency and Scalability of Large-Scale Power Grid Simulations

Power grid design is an important stage of VLSI device development which assures efficient power delivery to power drains and signal integrity. Accurate simulation of large ($10^7 - 10^9$ nodes) grids is required for modern devices. It leads to parallel solution of ODE $C\frac{dx}{dt} + Gx = F(t)$ [Chen T., Chen C. 2001 in Proc. Design Automation]. Parallel implementation based on PETSc sparse linear solver is proposed, efficiency and scalability study on BlueGene/P platform is presented.

Vasiliy Y. Voronov, Nina Popova
Department of Computational mathematics and cybernetics
Lomonosov Moscow State University
basrav@angel.cmc.msu.ru, popova@cs.msu.ru

CP7

HPC Modeling of Parasite Dynamics: Chagas' Disease

Around 2 billion people are estimated to harbor parasitic worms. Estimates are that 300 million individuals are severely ill with worms: 50% are school-age children. Because of its complex life-cycle, *T. cruzi* provides one of the most complex model systems for investigation. We have developed an extensible, parallel HPC model of *T. cruzi* parasite dynamics in blood and use that model to study strategies for managing the dynamics of *T. cruzi*, the causal agent in Chagas' disease.

Tarynn M. Witten
VCU, Center for the Study of Biological Complexity
tmwitten@vcu.edu

Samuel Sieg
Department of Computer Science
Virginia Commonwealth University
siegsd@vcu.edu

CP8

Distributed Fast Fourier Transform for Severely Energy and Resource Constrained Systems

In this work we develop an FFT distribution mechanism that enables the computation of a large FFT problem on wireless embedded systems with stringent energy, storage and bandwidth limitations. The idea of the proposed FFT distribution is to recursively partition the given problem into smaller problems which can be accommodated by the storage capacity of the platform. Then push ALL the twiddle factor calculation to that level of partitioning. This not only minimizes the communication overhead tremendously since minimal data exchange is needed, but also makes an 'unsolvable' FFT problem, 'solvable' on such systems. Simulations are done in Simulink Matlab and experiments were done on Telosb sensor nodes. Other implementation optimizations were also proposed.

Sherine Abdelhak, Jared Tessier, Soumik Ghosh, Magdy Bayoumi
University of Louisiana at Lafayette
spa9242@gmail.com, jaredtessier@gmail.com, soumikg1@gmail.com, mab@cacs.louisiana.edu

CP8

Distributed Second Generation Wavelet Transform on Resource Constrained Wireless Sensor Nodes

This work has four folds: 1) implementing a distributed second generation Wavelet transform on a network of IMote2 featuring XScale PXA271 processor, and Telosb sensor nodes, 2) optimizing the implementation on individual nodes to be resource and energy efficient exploiting the advantages of the PXA271 architecture, 3) developing a resource and channel aware scheduling and task allocation, 4) developing a complete communication and computation architecture based on intelligent tags in order to maintain a compact distribution with minimum communication overhead.

Sherine Abdelhak
University of Louisiana at Lafayette
spa9242@gmail.com

CP8

High Performance FFT Algorithms for the Convey Coprocessor

Fast Fourier Transforms (FFTs) perform poorly on standard microprocessors once the problem size exceeds the data cache due to high bandwidth requirements, bad memory access patterns, short inner loop lengths, and complex arithmetic. The Convey HC-1 hybrid-core computer combines an Intel64 processor with a coprocessor that supports multiple instruction set architectures (personalities) that can be tailored to accelerate individual algorithms and applications. This allows the co-development of FFT algorithms and instruction set architectures designed to execute them efficiently. Our solution overcomes the FFT deficiencies mentioned above by a unique memory design and microarchitecture, novel instructions for complex arithmetic, and a large vector register file to support high radix kernels. Complete algorithmic and architecture details as well as the performance of FFTs are presented.

Kevin Wadleigh
Convey Computer Corporation
kwadleigh@conveycomputer.com

CP9

High Performance Computing to Model High Frequency Modes in Proper Orthogonal Decomposition (POD)-Galerkin Model Reduction

We present the application of parallel computing in modeling high frequency modes in POD-Galerkin reduced-order models. We simulate and record three-dimensional turbulent flow field for the flow past a circular cylinder at Re=1000 using an MPI-based parallel CFD code on a 64-processor platform. We compute the POD modes from the 3D snapshot data and develop a reduced-order model. We suggest an LES-type approach and add an eddy viscosity term in the reduced-order model for closure.

Imran Akhtar
Virginia Tech
akhtar@vt.edu

Jeff Borggaard
Virginia Tech
Department of Mathematics
jborggaard@vt.edu

Traian Iliescu
Interdisciplinary Center for Applied Mathematics
iliescu

Zhu Wang
Virginia Polytechnic Institute and State University
wangzhu@vt.edu

CP9

A Parallel, High-Order Numerical Solver for the Navier-Stokes Equations in Two and Three Dimensions

We present a numerical method for solving systems of PDEs, with applications to the direct numerical simulation of the Navier-Stokes equations in two and three dimensions. We cover the computational domain by overlapping curvilinear grids and decompose these grids into sub-blocks for a parallel algorithm. In each sub-block we achieve high-order accuracy using a fast, one-dimensional, periodic continuation method. This allows the use of FFT algorithms to accurately approximate derivatives while avoiding the Gibbs phenomenon.

Nathan Albin, Oscar Bruno
California Institute of Technology
nathan.albin@caltech.edu, bruno@acm.caltech.edu

CP9

A GPU Accelerated Block PCG Pressure Projection Solver on Dynamic Adaptive Grids

In this work we apply a GPU "accelerator" strategy to a preconditioned conjugate gradient (PCG) solver used in the pressure projection step of a Coupled Level Set and Volume of Fluid (CLSVOF) legacy code. The code has been used for numerous incompressible multiphase flow applications which utilize dynamic adaptive grids in the numerical solution. The accelerator strategy consists of mapping data heavy operations within the PCG solver to the GPU for computation, leading to modest performance gains. This minimally intrusive strategy can be seen as an intermediate step to the planned complete replacement of the CPU PCG solver with a GPU PCG solver which will lead to more substantial computational savings.

Austen C. Duffy
Department of Mathematics
Florida State University
aduffy@math.fsu.edu

Mark Sussman
Florida State University
sussman@math.fsu.edu

CP9

An Adaptive Domain Decomposition and Load-Balancing Algorithm for Parallel SPH

We present a domain decomposition and load-balancing algorithm for parallel Smoothed Particle Hydrodynamics (SPH). The algorithm decomposes the simulation domain into fixed size cells only where particles are present. Particles are exchanged among processors and new cells are created as the simulation proceeds. For load balancing, we employ a novel algorithm motivated by the *diffusion equation* to interchange cells between processors. Global

communication is minimized, to ensure good scale up.

Prabhu Ramachandran
Department of Aerospace Engineering
Indian Institute of Technology, Bombay
prabhu@aero.iitb.ac.in

Chandrashekhar P. Kaushik
Department of Computer Science and Engr.
Indian Institute Of Technology, Bombay
shekhar.kaushik@iitb.ac.in

CP9

High-Performance Compressible Turbulence Simulations on the GPU

Modern high-order Godunov schemes for shock capturing require the approximate solution of the Riemann problem for compressible flow many times during the course of the simulation. This "Riemann solve" step is very computationally intensive and reducing its execution time provides an opportunity for significant speedup. We describe a computational framework for compressible turbulence based on the parallel implementation of a Riemann solver on many-core GPU architectures.

George Stantchev, Anil Deane
University of Maryland
gogo@math.umd.edu, deane@umd.edu

CP9

A Structure Preserving Hybrid Direct/iterative Solver for Fluid Flow Problems

We present a novel hybrid direct/iterative solver for fluid flow problems. The algorithm is based on nested dissection combined with an iterative solver for the Schur-complement on the interfaces. A symmetry-preserving incomplete factorization is used as preconditioner for the Schur-complement. The method can be applied recursively and exposes parallelism on each level while preserving important properties of the original problem such as structure, symmetry and definiteness. Nearly grid-independent convergence rates are demonstrated.

Jonas Thies
Dept. of Math and Computer Science
University of Groningen
j.thies@rug.nl

Fred Wubs
University of Groningen
f.w.wubs@rug.nl

CP10

New Results in Numerically Solving Large-Scale Algebraic Sylvester Equations

Two approaches are investigated for constructing numerical solutions of large-scale Sylvester equations. The first is to transform the Sylvester equation into an invariant subspace problem, which can be solved efficiently in real arithmetic. Optimal shift selection is presented to enhance convergence of iterative eigensolvers. This optimal shift selection analysis is then generalized for the computation of multiple optimal real shifts for the ADI approach, resulting in a more compact search region for the optimal shifts.

Hung Nong
Sandia National Laboratories
ryannong@gmail.com

Danny C. Sorensen
Rice University
sorensen@rice.edu

CP10

Scalable solvers for beam dynamic simulations

In this paper, we will report our progress on developing scalable solvers for beam dynamic simulations. Benchmark results will be shown. Different numerical techniques have been adopted for solving Poisson, Vlasov and Maxwell equations. Their performances have been compared. Domain decompositions in high dimension have been used for parallelization. Interesting results will be presented and challenges will be discussed.

Jin Xu
Argonne National Lab.
jin_xu@anl.gov

Peter Ostroumov
Physics Division
Argonne National Lab.
ostroumov@phy.anl.gov

Brahim Mustapha
Argonne National Lab.
mustapha@phy.anl.gov

MS1

Cost Estimation Algorithms for Dynamic Load Balancing of Adaptive Mesh Refined Simulations

Uintah is a parallel computational framework for simulating fluid-structure interactions with fully automated parallelism, check pointing and restarting, and adaptive mesh refinement. Uintah uses a novel asynchronous task-based approach to achieve scalable parallel performance. Achieving scalability with Uintah is complicated by the combination of AMR CFD solvers for the fluid and MPM particle methods for the solid. A new load balancing algorithm using statistical forecasting methods provides accurate estimates of varying workloads and results in both strong and weak scaling results up to 98K cores on ithe NSF NICS Kraken computer.

Justin Luitjens
SCI Institute and School of Computing
University of Utah
luitjens@sci.utah.edu

Martin Berzins
University of Utah
mb@cs.utah.edu

MS1

Parallel Image Registration Algorithms

Abstract unavailable at time of publication.

George Biros
Georgia Institute of Technology

biros@gatech.edu

MS1

Improving the Parallel Scaling of the Block-structured Mesh Adaptation Framework AMROC

AMROC provides patch-based dynamic mesh adaptation to Cartesian finite volume and difference schemes in the Virtual Test Facility software (http://www.cacr.caltech.edu/asc). Current applications include large-scale parallel adaptive simulations of shock-induced combustion, turbulent mixing, and shock-driven fluid-structure interaction. The presentation focuses on assessing weak and strong scalability of the MPI-based implementation on up to several thousands cores IBM BG/P and in resolving uncovered bottlenecks, especially in the rigorous domain decomposition algorithm based on a generalized space-filling curve.

Ralf Deiterding
Oak Ridge National Laboratory
deiterdingr@ornl.gov

MS1

Scalability Challenges for Massively Parallel AMR Application

We describe optimizations to the Chombo AMR framework that enable it to scale efficiently to thousands of processors on the Cray XT4. The optimization process also uncovered OS-related performance variations that were not explained by conventional OS interference benchmarks. Ultimately the variability was traced back to complex interactions between the application, system software, and the memory hierarchy. Once identified, software modifications to control the variability improved performance by 20% and decreased the variation in computation time across processors by a factor of 3. These newly identified sources of variation will impact many applications and suggest new benchmarks for OS-services be developed.

Brian Van Straalen
Lawrence Berkeley National Laboratory
Department N E R S C
bvstraalen@lbl.gov

John Shalf
Lawrence Berkeley National Laboratory
jshalf@lbl.gov

Terry J. Ligocki
Lawrence Berkeley Laboratory
TJLigocki@lbl.gov

Noel Keen
LBNL
ndkeen@lbl.gov

MS2

Dense Linear Algebra on GPU Clusters: One-sided Programming with Global Arrays and CUDA

The Global Arrays toolkit (GA) provides a powerful API for implementing complex numerical algorithms, such as those found in quantum chemical simulation. We report on the extension of the one-sided communication protocols within GA to GPUs using CUDA. Performance of matrix multiplication, the most important linear algebra kernel for quantum chemistry, will be evaluated on GPU clusters with a variety of interconnects. Implementation challenges and numerical precision issues will be discussed.

Jeff R. Hammond
Argonne National Laboratory
Leadership Computing Facility
jhammond@mcs.anl.gov

Manojkumar Krishnan, Oreste Villa
Pacific Northwest National Laboratory
manoj@pnl.gov, oreste.villa@pnl.gov

MS2

Rapid Generalized Molecular Dynamics Simulations on GPUs

Molecular dynamics algorithms are used to investigate phenomena in a wide range of systems, from atomistic models of proteins to coarse-grained models of polymers. Several well known codes have been developed to parallelize MD codes across large numbers of processors. In contrast, the GPU accelerated HOOMD-Blue, Highly Optimized Object-oriented Many-particle Dynamics – Blue Edition, performs general purpose particle dynamics simulations on a single GPU-enabled workstation, but achieves the performance of sixty processor cores. An open-source code, HOOMD-Blue is being expanded to include the array of pair interactions, integrators, and tools requested by MD researchers, such as Brownian Dynamics, and angle, dihedral, and improper forces. In my research in soft matter self-assembly, for example, I have used GPU-accelerated molecular dynamics to rapidly generate and explore phase diagrams of tethered nanoparticles.

Carolyn L. Phillips, Joshua Anderson, Trung Nguyen
University of Michigan
phillicl@umich.edu, joaander@umich.edu,
ndtrung@umich.edu

Sharon C. Glotzer
Chemical Engineering
University of Michigan
sglotzer@umich.edu

MS2

Graph Theoretic Methods for the Orbital Ordering Problem in DMRG

The Density Matrix Renormalization Group (DMRG) method has shown to be a powerful method as applied to quantum chemistry, but still many questions remain unanswered when one considers applications to molecules. DMRG was developed to treat one dimensional systems with primarily local interactions, however, a problem arises of how to represent and map a higher dimensional system to the 1D lattice structure inherent to the current algorithm. We study the use of atomic like molecular orbitals and graph theoretic techniques in order to define this mapping. We compare our ordering of orbitals to those previously proposed, such as band width minimization of the one-electron integral matrix.

Claire Ralph, Garnet Chan
Cornell University
ccr53@cornell.edu, gkc1000@gmail.com

MS2

Exact Ground State Properties of Nearly Infinite Strongly Correlated Electron Systems via the Monte Carlo Power Method

The exact ground state properties of large strongly correlated electron systems have long remained elusive. Deterministically solving for such properties requires diagonalizing matrices too large to store in memory. Stochastically sampling the fermion ground state wave function succeeds only within a limited region of parameter space. In this work, we demonstrate how the novel Monte Carlo Power Method, a stochastic implementation of the deterministic power method, can obtain the ground state and first several excited state energies of systems previously considered too massive to be treated exactly. Our approach can readily compete with current Krylov methods, is trivially parallelizable, and unlike DMRG, can be applied to systems of any dimensionality.

Brenda Rubenstein
Columbia University
rubenstein.brenda@gmail.com

MS3

Combining PLAPACK and Elemental Cyclic Distributions

Between twelve an fifteen years have passed since the analysis of elementally-wrapped 2d matrix distributions and the creation of ScaLAPACK and PLAPACK. In this talk, we review the strengths and weaknesses of the three approaches and propose what we believe to be a simpler, more flexible, and higher performing solution. The key is to combine the formalism of elemental distributions, a careful use of collective communication, and a FLAME-like API. The reason for revisiting the domain of distributed memory libraries is not only to provide a better solution for clusters, but also to prepare for the possibility of distributed memory computing on a single chip.

Jack Poulson
Institute for Computational Engineering and Sciences
The University of Texas at Austin
jack.poulson@gmail.com

MS3

Unleashing the Power of MultiGPU Accelerators with FLAME

In this talk we will provide strong evidence that FLAME solves the programmability problem for developing dense linear algebra operations complex architectures, composed of a multicore processor and multiple hardware accelerators (GPUs, Cell B.E., etc.), each with its own local memory. In particular, we will show that the FLAME programming all complexity is hidden inside the SuperMatrix runtime scheduling mechanism, which incorporates software implementations of standard cache/memory coherence techniques in computer architecture to improve the performance. Our experimental evaluation on a Intel Xeon 8-core host linked to an NVIDIA Tesla S1070 platform with four GPUs delivers peak performances over 1 (single-precision) TFLOPS for the matrix-matrix product.

Enrique Quintana
Universidad Jaume I
Spain
quintana@inf.uji.es

MS3

libflame: What the User Should Know

While it can be argued that the FLAME project has made fundamental contributions to the understanding of dense linear algebra libraries, the impact of the research on scientific computing will likely be measured by the number of users of the libraries that are being produced. In this talk we give an overview of the libflame library. We present some of its features: the application interface, portability to Linux and Windows platforms, functionality, performance, and extensibility. We also discuss an infrastructure that allows us maintain a single representation in code that can target a broad range of architectures, including various options for exploiting parallelism.

Field Van Zee
Department of Computer Sciences
The University of Texas at Austin
field@cs.utexas.edu

MS3

An Overview of the FLAME Project

In this talk, we give an overview of the FLAME project, including its roots, its current status, and the vision for future research. The roots of the project are in a new notation for presenting algorithms and how this facilitates APIs that ease development and maintenance. It will be discussed how this seamlessly supports more complex data storage with better locality, such as hierarchical storage by blocks. This in turn supports algorithms that operate on blocks, algorithms-by-blocks, which enable thread-level parallelism for SMPs and multicore architectures. We discuss how the abstractions will in the future facilitate viewing matrices as a knowledge base of algorithms and expertise about algorithms and architectures rather than instantiations in code. This enables the possibility of mechanically applying optimizing transformations to the algorithms at a much higher level of abstraction, much like an expert would given unlimited time.

Robert A. van de Geijn
The University of Texas at Austin
Department of Computer Science
rvdg@cs.utexas.edu

MS4

Towards an Efficient Implementation of AMG Preconditioners on Massively Parallel Multicore Machines

Classical algebraic multigrid (AMG) has proven to be extremely efficient on distributed-memory architectures. However, the new multi-core architectures require increased small-grain parallelism, and scalability must be improved to take advantage of millions of cores. We will discuss our efforts to improve the design, implementation, and performance of classical AMG on multi-core architectures.

Allison H. Baker
Center For Applied Scientific Computing
Lawrence Livermore National Laboratory
abaker@llnl.gov

Robert Falgout, Ulrike Meier Yang
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory

rfalgout@llnl.gov, umyang@llnl.gov

MS4

A Compressible Navier-Stokes Solver for Heterogeneous Computing Environments

In this talk we cover the extension of a compressible gas dynamics code to handle diffusion transport using basic iterative techniques. The framework for this solver is written using an exciting new standard for accelerated and multicore architectures called OpenCL. This standard presents new challenges for efficiently implementing stencil-based algorithms that will be discussed.

Ben Bergen
Los Alamos National Laboratory
bergen@lanl.gov

MS4

Parallel Smoothed Aggregation Algebraic Multigrid for Multicore Architectures

In this talk we present smoothed aggregation (SA) algebraic multigrid methods for multicore architectures. In particular, we consider how one might address the increasing disparity between on-node CPU performance versus intranode latency. One such approach is a domain-decomposition smoother with multiple communication layers, which allows greater flexibility in the choice of local solves. We show results within an SA framework on a variety of architectures.

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

Chris Siefert
Sandia National Laboratories
csiefer@sandia.gov

MS4

Multigrid for Multicore

Multigrid is a highly parallel algorithm, but its efficient implementation on current multi-core CPUs, such as the IBM Cell or GPU-architectures is not straightforward. We will report on our experience in using these architectures for multigrid algorithms with applications in molecular dynamics and image processing.

Harald Koestler, Daniel Ritter
University of Erlangen-Nuremberg
harald.koestler@informatik.uni-erlangen.de,
daniel.ritter@informatik.uni-erlangen.de

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

Markus Stuermer
University of Erlangen-Nuremberg
markus.stuermer@informatik.uni-erlangen.de

MS5

Software Life-cycle and Integration Issues for CS&E R&D Software and Experiences from Trilinos

This talk will provide an overview of software life-cycle and integration issues for the CS&E R&D community. Important issues will be previewed such as software life-cycle, regulated backward compatibility, regressions in capability, basic software quality, continued maintenance and support, and tighter development and integration models for more aggressive R&D. In addition, past experience and future plans for the Trilinos project that relate to these issues will be discussed.

Roscoe A. Bartlett
Sandia National Laboratories
rabartl@sandia.gov

MS5

Applying Software Engineering Principles to the Development of Scientific and Engineering Software: Lessons Learned from a Series of Case Studies and Workshops

Along with colleagues, I have conducted case studies of CSE software development projects to identify unique characteristics that will help software engineering and CSE researchers create effective CSE software development approaches. I will discuss the case study process along with important characteristics of CSE software development identified. Finally, I will discuss outcomes of a series of workshops designed to facilitate communication between members of the software engineering community and members of the CSE community.

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

MS5

Cross Platform Development and Continuous Integration Testing with CMake, CTest, CPack, and CDash

CMake has been in development since 1999, and has been used on several large open source projects such as ITK, VTK, ParaView, VXL, Trilinos and CMake itself. Further, KDE, one of the largest OSS projects has adopted CMake, demonstrating that CMake is capable of successfully supporting complex and large software systems. Hence CMake usage is growing rapidly with thousands of daily downloads and inclusion in several Linux distributions. Unlike many build systems, CMake is designed to be used in conjunction with native build tools enabling developers to use makefiles, Kdevelop projects, Xcode projects, and even MS Visual Studio projects. A simple input language (included in a CMakeLists.txt file) is used to specify which files to build and what types of system introspection tests need to be performed to build the software. A persistent cache file is used to store the system information and avoid the need for user-defined environment variables. In addition to building software, CMake provides a testing client (CTest) that integrates with the web-based CDash testing server. This server creates dashboards that build a snapshot of the software at a given time. This is critical to cross-platform development since often a change on one platform fails to compile on another one. The testing system provides for nightly builds which use a copy of the software at a specific time each night, experimental tests that can be used

to share build results with other developers before committing source code, and continuous build results that test the build each time files are committed to the source control system. Once the software is built and tested, the CPack tool can be used to package the software. CPack works similar to CMake in that it generates package information for native packaging tools. NSIS, RPM, OSX packages self extracting tar.gz, tar.gz, tar.zip can all be created. CPack information is included as a simple extension to the CMake build files. This session will cover the history and basic usage of CMake, CTest, CDash and CPack in the context of a integrated development environment. Users will learn how to create CMake files to build projects, set up a CDash testing server, execute and submit test results with CTest, and create cross-platform distributions with CPack.

Bill Hoffman
Kitware
bill.hoffman@kitware.com

MS5
PETSc at 15

There is a tension between supporting a long term open source software library with a large user base and simultaneously innovating both implemented algorithms and user interfaces. How can one avoid alienating users with constantly changing calling sequences that require modifications to users' code? With PETSc we made the concious decision never to freeze calling sequences and class hierarchies, but rather to adjust them over time based on our better understanding of what works and what can be made more flexible, general or extensible. We will discuss what we have learned during the 15 years of developing PETSc, how our design has evolved, and how this transition to extended capabilities would simply have been impossible with frozen interfaces.

Barry F. Smith
Argonne National Lab
MCS Division
bsmith@mcs.anl.gov

Lois Curfman McInnes, Mattew Kneply
Argonne National Laboratory
Mathematics and Computer Science Division
mcinnes@mcs.anl.gov, knepley@gmail.com

MS6
An Approach to Automatic Tuning for the Parallel Householder Qr Decomposition

We consider parallel computing of the Householder QR decomposition on SMP machines. In this computation, we partition the target matrix in both column- and row-wise. To achieve efficient computation, we have to partition it adequately depending on the target machines and the size of the matrix. In this talk, we propose an approach to automatic tuning based on the hierarchical structure of the algorithm. We implement our approach and evaluate it through numerical experiments.

Takeshi Fukaya
Department of Computational Science and Engineering
Nagoya University
t-fukaya@na.cse.nagoya-u.ac.jp

Yusaku Yamamoto, Shao-Liang Zhang
Nagoya University, Japan
yamamoto@na.cse.nagoya-u.ac.jp, zhang@na.cse.nagoya-u.ac.jp

MS6
Quality Control Approach for Systematic Performance Evaluations of Numerical Solving Process of Linear Equations

In this talk, we explain the quality control aspect that arises from the solving process of the linear equations. Especially, we introduce our methodology to analyze and evaluate on numerical algorithms systematically by using information visualization. We have carried out the comparing the efficiency between solvers and preconditionings, and we discuss preconditioning directions and their related matters of the Krylov subspace methods.

Shoji Itoh
Advanced Center for Computing and Communication
The Institute of Physical and Chemical Research(RIKEN)
itosho@riken.jp

Masaaki Sugihara
Graduate School of Information Science and Technology, the University of Tokyo
m_sugihara@mist.i.u-tokyo.ac.jp

MS6
High Performance and Low Power GPGPU Computing with Automatic Tuning

We discuss high performance and low power processing using GPU. GPU provides drastically high performance and better power-performance ratio compared with conventional CPU when the computational throughput of GPU is fully exploited. But it is not always straightforward, that is, it requires high MIMD-SIMD parallelism in algorithms, careful design of data structures and memory accesses, and avoidance of excessive branching. In this talk we discuss high performance GPU implementations of discrete structures and algorithms that are big challenge in GPU computing. We also mention our efforts of power optimization: we seek best power-performance ratio through accurate modeling of power dissipation and computing performance on CPU-GPU heterogeneous platforms.

Kamil Rocki
Department of Computer Science
the University of Tokyo
kamil.rocki@is.s.u-tokyo.ac.jp

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

MS6
OpenATLib: A General Auto-tuning Interface for Numerical Solvers

Reusability for implementation of Automatic Tuning facility (RIAT) is needed to establish low-cost construction of auto-tuning software. In this presentation, we propose Auto-tuning interface named OpenATLib to realize RIAT for numerical libraries. To evaluate the effectiveness of the OpenATLib, we have developed sparse iterative solvers named Xabclib_LANCZOS and Xabclib_GMRES with the help of OpenATLib. Performance evaluation of

OpenATLib with T2K Open Supercomputer (U. Tokyo) indicated that the maximum speedup established 22.4x (Xabclib_LANCZOS) and 3.5x (Xabclib_GMRES).

Takahiro Katagiri
Information Technology Center, The University of Tokyo
katagiri@cc.u-tokyo.ac.jp

Ken Naono
Central Research Laboratory, Hitachi, Ltd.
naono@crl.hitachi.co.jp

Hisayasu Kuroda
Information Technology Center, The University of Tokyo
kuroda@cc.u-tokyo.ac.jp

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

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

MS7

From Holistic to Oblivious Fault Tolerance at the Exascale

Today, cutting edge fault tolerance research involves increasing the awareness of faults across all the levels of the system from the hardware all the way up to the application. But when faults become so frequent that they can be considered continuous, awareness and recovery will have to be replaced with algorithms that are oblivious to faults yet still produce the correct result. This talk will discuss the shift from holistic fault tolerance to fault oblivious algorithms and give several examples of such algorithms that we have developed.

Al Geist
Oak Ridge National Laboratory
gst@ornl.gov

MS7

Fault Tolerance and the MPI Standard - Implications to the Standard

As the scale of systems increases, the ability of computer programs to tolerate failures is recognized to be of crucial importance to many users of distributed memory, parallel processing computers. The current version of the MPI standard does not define the response to failure, but work is being done to define such standard support. This talk will describe the proposed MPI standard changes and their prototype implementation within the context of the Open MPI implementation.

Richard Graham
Oak Ridge National Laboratory
rlgraham@ornl.gov

MS7

Moving Towards Coordinated Fault Tolerance in High-end Computing Systems

The need for efficient fault-tolerance continues to increase as emerging high performance systems approach petascale performance. While most systems provide some degree of fault tolerance, fault information is seldom shared across programs or the system. I will present work by the Coordinated Infrastructure for Fault Tolerant Systems initiative that enables holistic fault management by coordinating fault information across the system. As a part of this work, I will describe the Fault Tolerance Backplane, a unified framework for sharing of fault information, and FTB-enabled software.

Rinku Gupta
Argonne National Laboratory
rgupta@mcs.anl.gov

MS7

A Brief Introduction to BLCR (Berkeley Lab Checkpoint/Restart)

Abstract unavailable at time of publication.

Paul Hargrove
Lawrence Berkeley National Laboratory
phhargrove@lbl.gov

MS8

Exa-Scale Volunteer Computing

Peta-scale computing was first reached not by supercomputers, clusters, grids, or clouds, but by volunteers: first the 40,000 Sony Playstation 3 game consoles running Folding@home, and more recently the BOINC network of 600,000 PCs. This achievement is just a waypoint: volunteer computing offers a feasible near-term path to the next three orders of magnitude: Exa-scale. This will involve GPUs and multi-core CPUs. Additionally, volunteer computing engages the public in the scientific process, and changes the mechanisms by which computational resources are divided among scientists.

David Anderson
Space Sciences Lab
U.C. Berkeley
davea@ssl.berkeley.edu

MS8

Desktop Supercomputing

In recent years, the dominant position of traditional CPUs for supercomputing has been challenged by alternative architectures, such as the GPU. For certain types of algorithms, a single GPU can deliver the same performance as dozens of CPUs at a fraction of the cost and power consumption. In this talk I will present our work on combining the power of multiple GPUs, forming a desktop supercomputer that can match the performance of large CPU clusters.

K. Joost Batenburg
University of Antwerp, Belgium
joost.batenburg@ua.ac.be

MS8

Green Supercomputing Comes of Age

In 2002, Green Destiny debuted amidst much ridicule for having squeezed a 240-node cluster into 5 square feet and a mere 3.2 kilowatts. Since then, power and cooling have be-

come first-class design constraints. This talk presents the evolution of green supercomputing from architecturally-based to software-based approaches, and then offers a prognostication of what is yet to come, including future directions for The Green500 List.

Wuchun Feng
Department of CS and ECE
Virginia Tech
feng@cs.vt.edu

MS8

Folding@home: Multi-petaflop Performance from a Distributed Network of Hundreds of Thousands of CPUs, GPUs, and Playstation 3's

Many problems in molecular simulation are fundamentally limited by computational power. I will talk about our efforts to solve this key impasse, highlighting on our novel algorithms for scaling to hundreds of thousands of loosely coupled processors, use of GPUs and Playstation 3's to get significant (up to 1000x) speed increases per node, as well as implications for green computing. Finally, I will discuss how these technologies can be applied on a smaller scale.

Vijay Pande
Stanford University
Pande Lab
pande@stanford.edu

MS9

Parallelisation of Recursively Structured Adaptive Grids Using Space-Filling Curves

We present an approach for numerical simulations on dynamically adaptive grids that is based on recursively structured grids. The grid refinement may follow an octree-type procedure or repeated bisection of grid cells. The respective refinement trees are sequentialised in memory using Peano curves for quadrilateral and Sierpinski curves for triangular grids. A strictly element-oriented processing together with a stack-oriented data access leads to highly memory-efficient algorithms. Parallelisation exploits the space-filling curve as well as the hierarchical refinement tree for efficient partitioning and load balancing.

Michael Bader
University of Stuttgart
Institute of Parallel and Distributed Systems
bader@in.tum.de

Miriam Mehl, Tobias Weinzierl
Technische Universität München
mehl@in.tum.de, weinzier@in.tum.de

MS9

Parallel Adaptive PDE Solvers with Generic Finite Element Libraries

Parallel algorithms to solve partial differential equations are particularly complex if the mesh changes dynamically, for example as part of adaptive mesh refinement. Consequently, few codes accomplish this feat, and most are one-off codes purpose-built for a particular application. Within the deal.II project, we attempt to extend an existing finite element library to support massively parallel computations. We will review the difficulties a generic finite element implementation faces, and what features deal.II currently supports.

Wolfgang Bangerth
Texas A&M University
bangerth@math.tamu.edu

MS9

The ALPS Framework for Parallel, High-order, Multi-octree AMR

We present the ALPS (Adaptive Large-scale Parallel Simulations) framework for parallel adaptive solution of PDEs. ALPS includes the octor and p4est libraries for parallel dynamic mesh adaptivity on single-octree-based and forest-of-octree-based geometries, respectively, and the mangll library for arbitrary-order hexahedral continuous and discontinuous finite/spectral element discretizations on general multi-octree geometries. ALPS has been shown to scale well weakly and strongly to over 130,000 processor cores. We present the ALPS framework and describe applications to several solid earth geophysics problems: global mantle convection with nonlinear rheology, full Stokes models of ice sheet dynamics, and global seismic wave propagation.

Carsten Burstedde, Omar Ghattas, Georg Stadler,
Tiankai Tu, Lucas Wilcox
University of Texas at Austin
carsten@ices.utexas.edu, omar@ices.utexas.edu,
georgst@ices.utexas.edu, tiankai.tu@gmail.com,
lucasw@ices.utexas.edu

MS9

Using Adaptivity and Parallelism to Manage Uncertainty Quantification based Hazard Analysis

In this talk we will present recent work on using adaptivity in model input, parameter and discretization spaces, simultaneously to study the hazard arising from multiple geophysical mass flows. Uncertain input data includes uncertainty about terrain, flow constitutive modeling parameters and flow volume. "Adaptive meshing" of the parameter space allows construction of input-output maps that enable construction of hazard maps. Parallelism in each simulation and in the overall ensemble computations allows us to reasonably account for the uncertainty in a short time frame necessary to support crisis management.

Abani K. Patra
SUNY at Buffalo
Dept of Mechanical Engineering
abani@eng.buffalo.edu

Keith Dalbey
Sandia National Lab,
Albuquerque, New Mexico
abijah2004@gmail.com

Bruce Pitman
SUNY at Buffalo
Dept of Mathematics
pitman@buffalo.edu

Matthew Jones
Center for Comp. Research,
Univ. at Buffalo, Buffalo, NY 14260
jonesm@ccr.buffalo.edu

Eliza Calder

University of Buffalo
ecalder@buffalo.edu

Marcus Bursik
Dept. of Geology
SUNY at Buffalo
mib@buffalo.edu

MS10

Parallel Methods for Food Webs: Bayesian Network Analysis and Stochastic Discrete-Event Simulation

Food webs, networks of feeding relationships in ecosystems, may include hundreds of species and thousands of links in well-sampled data sets, presenting significant computational challenges both for statistical analysis of network structure and for simulation of system dynamics. Past approaches to analyzing food web network structure have focused on comparing summary statistics against null models. Recent work has introduced a likelihood framework for fitting structural models and comparing them via model-selection criteria, and has demonstrated that groups of species with similar feeding relationships are important features of network structure. I will describe an extension of this approach to a Bayesian context using Markov-chain Monte Carlo methods, parallelized at the level of individual likelihood calculations and multiple "heated" chains. Spatial models of food web dynamics also benefit from parallelization. Most past studies of food web dynamics have been based on ODEs, but, due to the spatial distribution and discrete nature of organisms, spatially explicit, discrete-event simulations are often more appropriate. I will introduce a framework for running such simulations in parallel by combining the Gillespie algorithm with an optimistic discrete-event simulator based on Time Warp.

Edward Baskerville
University of Michigan
ebaskerv@umich.edu

Stefano Allesina
University of Chicago
sallesina@uchicago.edu

MS10

Evolving Classical Fields and Metric Perturbations in an Expanding Universe Using a Variational Integrator

Covariant evolution equations, including those for metric perturbations, have complex conversation equations, and in the case of metric perturbations, gauge constraints, which are difficult to satisfy using common numerical evolution schemes. Variational integration, a technique in which evolution is generated by extremizing the variation of a discrete Lagrangian, is particularly well-suited for these kinds of difficult covariantly-conservative systems. The method is multisymplectic-structure preserving, conserving all currents generated by Lagrangian symmetries, and easily handles constraints via the introduction of Lagrange multipliers. The scalability and convergence of a parallel, PETSc-based implementation will be discussed.

Hal Finkel
Yale University
hal.finkel@yale.edu

MS10

Parallelization by Stimulus Leads to Rapid Yield of Sufficient Spikes for Receptive Field Mapping

A primary goal of sensory neuroscience is to be able to predict neural response to any arbitrary stimulus. In vision research a commonly used tool for this goal is the receptive field map, which is considered to be a spatiotemporal filter of a visual scene. The most widely used method for receptive field map construction is reverse correlation where by a neurons spiking behavior is related to stimuli at a certain physiological time lag. Basis set stimuli or white noise are commonly used to span the input space, and on the order of hundreds to thousands of spikes need to be collected for reliable receptive field mapping. For anatomically constrained and large-scale computational models of the primary visual cortex the collection of this magnitude of spikes for each neuron is non-trivial computationally. The corresponding experimental time is on the order of tens of minutes, and most current large-scale models of cortex seek to simulate only seconds of experimental time. This computational demand for constructing receptive field maps in turn necessitates an efficient parallel approach. The first line strategy we take is parallelization by stimulus, which in turn leads to highly efficient and loosely coupled simulations. The receptive field maps are then constructed when the simulation threads converge. The significance of this parallel efficiency is that it allows a battery of different stimuli with different features to probe an anatomically plausible and highly recurrent architecture of cortex.

Armen Kherlopian
Weill Medical College
Cornell University
ark2010@med.cornell.edu

MS10

A New Parallel Dark Matter Halo Finder for Cosmological Simulation

Statistical analysis of cosmological simulations relies upon the ability to identify halos, which are clumps of dark matter that host galaxies, groups, and clusters. We present a new halo finding algorithm built into the particle-mesh cosmology code FLASH. This algorithm scales well and only moderately increases the runtime of a simulation. We compare this algorithm to current halo finders and demonstrate its abilities by using it for an in-code analysis of the growth of supermassive black holes and as a basis for feedback-based subgrid models.

Paul Matthew Sutter
University of Illinois at Urbana-Champaign
psutter2@illinois.edu

Paul M. Ricker
University of Illinois, Urbana-Champaign
Dept. of Astronomy
pmricker@illinois.edu

MS11

Applied Computer Science - An Indispensible Part of Simulation Science Graduate Programs

Educational Programs traditionally tend to emphasize the mathematical and numerical aspects of simulation science. However, in the end, it is software artifacts that brings simulationbehr to life. As a result, concepts from applied computer science must be part of simulation science pro-

grams, as is the case, for example, with the graduate and postgraduate programs offered through the German School for Simulation Sciences (GRS) and the Aachen Institute for Advanced Study in Computational Engineering Science (AICES).

Christian H. Bischof
RWTH Aachen University
Inst of Scientific Computing
bischof@rz.rwth-aachen.de

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

MS11

Teaching Parallelism in an Interdisciplinary Scientific Computing Programme

The advent of multicore processors has brought parallelism to every desktop, and thus parallel programming and parallel algorithm design must be part of modern university curricula in various disciplines. Here, one can benefit from past experience with parallel computing in high-performance oriented programmes, such as the Utrecht MSc programme in scientific computing that has run since 1993. In this talk, we will discuss our experiences with the current programme and our outlook for the future.

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

MS11

3 Big Events at UC Berkeley: the Multicore Revolution, Cloud Computing, and a Graduate Program in Computational Science and Engineering

We describe 3 synergistic events that have occurred at UC Berkeley in the last few years: the establishment of the ParLab, a large research lab with 15 faculty and over 50 graduate students to work on the consequences of the multicore revolution, the establishment of the RadLab, another large research lab focused on cloud computing, and a new Designated Emphasis (graduate minor) in Computational Science and Engineering, with over 110 faculty from 22 departments participating.

Jim Demmel
Division of Computer Science
University of California, Berkeley
demmel@cs.berkeley.edu

MS11

Parallelism Goes Mainstream: What About CS Curricula?

I will outline the motivation for building a new course on HPC at the University of Basel. The University of Basel has been founded in 1460, it is Switzerland's oldest university, but the Department of Computer Science has only been established in 2003. The emphasis of the CS curriculum is on computational science. The high-performance computing course is taught within the master program since 2006. I will talk about how we face the challenge of graduate education for parallel computing.

Olaf Schenk
Department of Computer Science, University of Basel
Switzerland
olaf.schenk@unibas.ch

Helmar Burkhart
University of Basel, Switzerland
helmar.burkhart@unibas.ch

MS12

The Microsoft Parallel Computing Platform

During the next few years we will see an increase of parallel computing resources for Windows. In this session we will give you insights to our vision on how to efficiently develop robust, scalable and responsive applications for manycore, heterogeneous and distributed systems.

Bill Colburn
Microsoft
bill.colburn@microsoft.com

MS12

Managed and Native Code Parallel Programming

Visual Studio and the Microsoft .NET Framework can help you write better performing and more scalable applications. We take a tour of new .NET APIs, including the Task Parallel Library (TPL)and Parallel LINQ (PLINQ)for managed code developers. We also look at how to build more responsive C++ programs that take full advantage of manycore hardware. We demonstrate how the new Parallel Pattern Library (PPL) enables you to express parallelism in your code and how the asynchronous messaging APIs can be used to separate shared state and increase your application's resilience and robustness. Finally, we take a look at some of the new capabilities of C++Ox.

Josh Phillips, Rick Molloy
Microsoft
josph@microsoft.com, rickmol@microsoft.com

MS12

Debugging Enhancements for Parallel Debugging in Visual Studio 2010

This session will cover parallel debugging enhancements in Visual Studio 2010 including the new Parallel Tasks and Parallel Stacks window that helps developers quickly identify concurrency issues and visualize the internal state of their application. We will also discuss the new F5 experience for cluster debugging.

Daniel Moth
Microsoft
danimoth@microsoft.com

MS12

Performance Optimizations for Parallel Code

Learn how to prepare for the new challenges in developing and tuning concurrent applications. Hear about important steps in the creation of or conversion to parallel applications with demonstrations of the parallel performance analysis and optimization tools in the next release of Microsoft Visual Studio. See how to identify opportunities

for parallelism and how to exploit those opportunities by choosing applicable coding patterns using existing or future programming models. Finally, watch a demonstration that shows how to optimize parallel code by focusing on common sources of inefficiency such as I/O and synchronization.

Hazim Shafi
Microsoft Corporation
hshafi@microsoft.com

MS13

Large-Scale Multilingual Document Clustering

Recent years have seen a dramatic increase not just in the volume of text available online, but also the diversity of languages represented therein. This has led to increasing interest in techniques for managing and making sense of multilingual content. We describe a scalable technique for fully-unsupervised multilingual document clustering (grouping documents in different languages according to topic) which both attains precision of around 90% and allows analysis of millions of documents at interactive speed.

Peter A. Chew
Moss Adams, LLC
Peter.Chew@mossadams.com

MS13

ParaText: Scalable Text Analysis and Visualization

ParaText is a text analysis engine for processing and searching large collections of documents. The retrieval method employed is latent semantic analysis using a matrix SVD. In this talk, we demonstrate the fully scalable ParaText analysis pipeline, which utilizes the Trilinos distributed memory linear algebra framework, plus VTK components originally developed for scientific visualization.

Daniel M. Dunlavy
Sandia National Laboratories
Computer Science and Informatics
dmdunla@sandia.gov

Timothy Shead
Sandia National Laboratories
tshead@sandia.gov

MS13

Supporting Interactive Multilingual Document Clustering on the Cray XT3

This talk describes a parallel multilingual document-clustering application designed for the Cray XT3 (RedStorm) system at Sandia National Laboratories. Our application is unique among HPC applications because it provides interactive visualization and analysis capability by spanning three different architectures: the Cray XT3, a visualization cluster, and Netezza Data Warehouse appliance. We will discuss design, scalability challenges, and results for several large multilingual data sets, including the Bible, the Quran, and proceedings of the European Parliament.

Ron A. Oldfield
Sandia National Laboratories
raoldfi@sandia.gov

Brett W. Bader
Sandia National Laboratories
Computer Science and Informatics Department
bwbader@sandia.gov

Peter A. Chew
Moss Adams, LLC
Peter.Chew@mossadams.com

MS13

Parallelization of Multilinear Algebra Applications

A variety of useful tensor or multilinear algebra data models have been proposed over the years with a large number of applications in chemistry, psychometrics, text analysis, cybersecurity, etc. For very large and sparse data sets the computation of these models can be difficult. Here we present results on parallel implementation, scaling, and performance of alternating least-squares algorithms for computation of four popular data models: PARAFAC, Tucker, DEDICOM, and PARAFAC2. We show that these algorithms can be implemented in an efficient and scalable way using C++ and MPI.

Mark P. Sears
Sandia National Laboratories
mpsears@sandia.gov

Brett W. Bader
Sandia National Laboratories
Computer Science and Informatics Department
bwbader@sandia.gov

Tamara G. Kolda
Sandia National Laboratories
tgkolda@sandia.gov

MS14

Scalable Information Fusion for Fault Tolerance in Large-Scale HPC

Fault tolerance enabled by predictive capabilities and the invocation of time-critical response relies on the fusion of large amounts of data in disparate forms from system components, applications, and resource managers. This talk presents strategies and implementations for highly scalable data collection, fusion, exploration, analysis, visualization, and triggering of response addressing actual scenarios in HPC systems.

Jim Brandt, Frank Chen, Vince De Sapio, Ann Gentile, Jackson Mayo, Philippe Pebay
Sandia National Laboratories
brandt@sandia.gov, fxchen@sandia.gov, vdesap@sandia.gov, gentile@sandia.gov, jmayo@sandia.gov, pppebay@sandia.gov

Diana Roe
Sandia National Labs
Livermore CA
dcroe@sandia.gov

David Thompson
Sandia National Labs
dcthomp@sandia.gov

Matthew Wong
Sandia National Laboratories

mhwong@sandia.gov

MS14

Trends in Hardware Failures and Application Resilience Strategies

Trends in scaling up system sizes and chip transistor counts along with scaling down individual component feature sizes and voltages is predicted to reduce overall MTTI. Both detected uncorrectable errors (DUE) and silent data corruption (SDC) are expected to increase. Current strategies for application survival are somewhat limited. This talk discusses the failure rate trends and modes, and surveys some of the application level solutions being used today and pursued for tomorrow.

Larry Kaplan
Cray
lkaplan@cray.com

MS14

Title Not Available at Time of Publication

This talk will be about integrated system resiliency or a system simulator for evaluating system resiliency.

Helia Naeimi
Intel
helia.naeimi@intel.com

MS14

Designing Fault Resilient and Fault Tolerant Systems with InfiniBand

InfiniBand architecture is emerging as a modern interconnect for designing large-scale HEC systems. This architecture provides many novel mechanisms for performance, network fault-tolerance and Quality of Service (QoS). Examples of some of these mechanisms include: Remote DMA (RDMA), Automatic Path Migration (APM), Service Levels (SL) and Virtual Lanes (VL). We will present novel schemes using these mechanisms to provide various levels of fault resilience and fault tolerance at the MPI layer on InfiniBand clusters.

Dhabaleswar K. Panda
The Ohio State University
panda@cse.ohio-state.edu

MS15

Collaborative Autotuning of Scientific Applications

In this talk, we will describe collaborative autotuning tools that work with the application programmer to automate their performance tuning tasks and permit them to express their algorithms in architecture-independent code. We focus on compiler-based autotuning tools and an interface that provides application developers the opportunity to concisely express a range of implementations, from which the compiler technology can automatically generate highly optimized code. We present performance improvements for two case study large-scale scientific applications.

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

MS15

Annotation-Based Empirical Performance Tuning of Numerical Kernels

In many scientific applications, developers spend significant time tuning codes for a particular high-performance architecture. Tuning approaches range from the relatively nonintrusive (e.g., by using compiler options) to extensive code modifications that attempt to exploit specific architecture features. Intrusive techniques often result in code changes that are not easily reversible, which can negatively impact readability, maintainability, and performance on different architectures. We describe an extensible annotation-based empirical tuning system called Orio, which is aimed at improving both performance and productivity by enabling software developers to insert annotations in the form of structured comments that trigger a number of low-level performance optimizations on a specified code fragment. Given the annotated code as input, Orio generates many tuned versions of the same operation and empirically evaluates the versions to select the best performing one for production use. We have also enabled the use of the PLuTo automatic parallelization tool in conjunction with Orio to generate efficient OpenMP-based parallel code. We describe our experimental results involving a number of computational kernels, including dense array and sparse matrix operations.

Boyana Norris
Argonne National Laboratory
norris.mcs.anl.gov

Albert Hartono
Department of Computer Science and Engineering
Department of Computer Science and Engineering
hartonoa@cse.ohio-state.edu

MS15

Tiling and Parallelism in the Build to Order BLAS

The Build to Order BLAS (BTO) is a domain-specific compiler that translates high-level specifications, written in MATLAB, into high-performance C++. The compiler employs auto-tuning to decide which optimizations are profitable. In this talk we describe ongoing work to create a uniform framework for loop tiling and parallelism. The framework is based on a notion of hierarchically partitioned matrices and vectors. The existing BTO generic algorithms automatically adapt to the introduced partitions.

Jeremy Siek
Department of Electrical and Computer Engineering
University of Colorado at Boulder
jeremy.siek@colorado.edu

Geoffrey Belter
Dept. of Electrical, Computer, and Energy Engineering
University of Colorado at Boulder
geoffrey.belter@colorado.edu

Ian Karlin
Department of Computer Science
University of Colorado at Boulder
ian.karlin@colorado.edu

Elizabeth Jessup
University of Colorado at Boulder
Department of Computer Science
jessup@cs.colorado.edu

MS15
Cray Automatic Tuning Framework

Cray's Automatic-Tuning Framework CrayATF has been used to release real vendor-tuned scientific libraries to the Cray community. While much academic research has influenced the design of CrayATF, there are also new concentrations that are specific to the autotuning of libraries in the industrial sector. This talk expounds some of those differences and challenges, and describes what challenges remain unsolved, describing the key development methodologies in our approach to acheive tuned sparse matrix kernels and HPL.

Adrian Tate, Keita Teranishi
Cray Inc.
adrian@cray.com, keita@cray.com

MS16
Dynamically Load-balanced Anisotropic Mesh Adaptivity

A novel approach to parallel anisotropic mesh adaptivity. Rather than modifying the serial adaptive algorithm, the method focuses on dynamic load-balancing in response to the local refinement and coarsening of the mesh. In essence, the domain partition mesh partition is perturbed away from mesh regions of high relative error, while also balancing the computational load across processes. The new method is applied to a transient computational fluid dynamics problem. The parallel scaling properties of the method are discussed. It is shown that the new method is dominated by the cost of the serial adaptive mesh procedure and that the parallel overhead, interprocess data migration, represents only a small fraction of the overall cost.

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

Christopher Pain
Imperial College
c.pain@imperial.ac.uk

Matthew D Piggott
Department of Earth Science and Engineering
Imperial College London
m.d.piggott@imperial.ac.uk

Adrian Umpleby
Department of Earth Science and Engineering
Imperial College London
adrian@imperial.ac.uk

James Maddison
Atmospheric, Oceanic and Planetary Physics
Oxford University
maddison@atm.ox.ac.uk

Patrick Farrell
Department of Earth Science and Engineering
Imperial College London
patrick.farrell@imperial.ac.uk

MS16
Adaptive Hierarchical Grids with a Trillion Tetrahedra

The Hierarchical Hybrid Grid data structure has been used to solve finite element systems on tetrahedral meshes on up to 10 000 processor cores. Recently the method has been augmented to support patch-wise adaptive mesh refinement and to use multigrid with tau-extrapolation, a technique that can raise the approximation order at minimal extra cost.

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

Tobias Gradl
Universität Erlangen
tobias.gradl@informatik.uni-erlangen.de

Björn Gmeiner
Universität Erlangen
Chair for Simulation
bjgmeiner@web.de

MS16
Unstructured Mesh Adaptation on Massively Parallel Computers

The adaptive simulation of unstructured meshes on massively parallel computers requires the adaptation of the meshes in parallel using either all, or a large fraction, of the computing cores. This presentation will discuss improvements and extensions being made to our parallel mesh adaptation procedure for large core counts including improvements to the parallel mesh data structures, more scalable communication tools and incremental dynamic load balancing procedures to provide a better partition for scalable simulations.

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

Aleksandr Ovcharenko, TIng Xie
Rensselaer Polytechnic Institute
shurik@scorec.rpi.edu, txie@scorec.rpi.edu

Seegyoung Seol
Rensselaer Polytechnic Institute
Scientific Computation Research Center
seol@scorec.rpi.edu

Min Zhou
Rensselaer Polytechnic Institute
zhoum@scorec.rpi.edu

Onkar Sahni
Rensselaer Polytechnic Institute
Scientific Computation Research Center
osahni@scorec.rpi.edu

Kenneth Jansen
Rensselaer Polytechnic Institute
jansen@rpi.edu

MS16

An Unstructured H-P Mesoscale Flow Solver

This work concerns the efficient implementation of a h-p discontinuous Galerkin method for the solution of atmospheric flows. The solver is capable of explicit or implicit time-integration and, it relies on a fast mesh database technology developed for the Earth System Modeling Framework (ESMF). Various strategies for estimating the error are presented and tested on standard test cases.

Amik St-Cyr
National Center for Atmospheric Research
Institute for Mathematics Applied to the Geosciences
amik@ucar.edu

MS17

A Fast Parallel Poisson Solver on Irregular Domains Applied to Beam Dynamic Simulations

We discuss the scalable parallel solution of the Poisson equation within a Particle-In-Cell (PIC) code for the simulation of electron beams in particle accelerators of irregular shape. The problem is discretized by Finite Differences. Depending on the treatment of the Dirichlet boundary the resulting system of equations is symmetric or 'mildly' nonsymmetric positive definite. In all cases, the system is solved by the preconditioned conjugate gradient algorithm with smoothed aggregation (SA) based algebraic multigrid (AMG) preconditioning. We investigate variants of the implementation of SA-AMG that lead to considerable improvements in the execution times. We demonstrate good scalability of the solver on distributed memory parallel processor with up to 2048 processors. We also compare our SAAMG-PCG solver with an FFT-based solver that is more commonly used for applications in beam dynamics.

Andreas Adelmann
Paul Scherrer Institut
andreas.adelmann@psi.ch

Peter Arbenz
Swiss Federal Institute of Technology (ETH)
Institute of Computational Science
arbenz@inf.ethz.ch

Yves Ineichen
ETHZ
PSI
yves.ineichen@psi.ch

MS17

High-performance Adaptive Methods for Elliptic PDE

We will discuss the issues that arise in developing high-performance implementations of elliptic solvers on block-structured adaptive meshes. These include methodological issues, such as defining suitable metrics for scalability; algorithmic issues, such as locality-preserving load balancing; and implementation issues, such as minimizing overheads due to irregular computation at refinement boundaries. We will present scaling results for Poisson's equation to more than 10K processors, corresponding to more than 20B unknowns; and discuss the new issues that are arising in scaling to 100K processors and beyond.

Phillip Colella
Lawrence Berkeley National Laboratory
PColella@lbl.gov

Brian Van Straalen
Lawrence Berkeley National Laboratory
Department N E R S C
bvstraalen@lbl.gov

MS17

Poisson Solver on GPUs

We have recently ported an electrostatic Particle-in-Cell code to the NVIDIA GPUs. The field solver in this code is spectral, making using the FFT provided by NVIDIA. Boundary conditions other than periodic can be handled by extending the FFT to perform Fast Sine/Cosine Transforms. Performance results and lessons learned using this new architecture will be discussed.

Viktor K. Decyk, Tajendra Singh
UCLA
decyk@physics.ucla.edu, tvsingh@ucla.edu

MS17

A Linear Scaling and Massively Parallel Solver for the Electrostatic Problem in Quantum Chemical Calculations

The brute force calculation of the electrostatic interaction energy of manyelectron systems is $\mathcal{O}(\mathcal{N}^{\triangle})$. The Gaussian and plane waves (GPW) algorithm, which essentially transforms a density in a Gaussian basis to a density on a regular grid, reduces this scaling to $\mathcal{O}(\mathcal{N}\ln\mathcal{N})$ using Fast Fourier Transforms. Here we discuss the mixed MPI/OMP implementation of the GPW algorithm.

Joost VandeVondele
University of Zurich
Physical Chemistry Institute
vondele@pci.uzh.ch

Juerg Hutter
University of Zurich
Physical Chemistry Institute
hutter@pci.uzh.ch

Iain Bethune
The University of Edinburgh
ibethune@epcc.ed.ac.uk

Matthew Watkins, Ben Slater
University College London
Department of Chemistry
matthew.watkins@ucl.ac.uk, b.slater@ucl.ac.uk

MS18

Runtime Data Flow Scheduling of Matrix Computations

We investigate the scheduling of matrix computations expressed as directed acyclic graphs for shared-memory parallelism. Because of the data granularity in this problem domain, even slight variations in load balance or data locality can greatly affect performance. We provide a flexible framework for scheduling matrix computations, which we use to empirically quantify different scheduling algorithms. We have developed a scheduling algorithm that addresses both load balance and data locality simultaneously and

show its performance benefits.

Robert A. van de Geijn
The University of Texas at Austin
Department of Computer Science
rvdg@cs.utexas.edu

Ernie Chan
The University of Texas at Austin
Department of Computer Sciences
echan@cs.utexas.edu

MS18

Programming for Locality and Parallelism with Hierarchically Tiled Arrays

Programming for Locality and Parallelism with Hierarchically Tiled Arrays In this talk I will present Hierarchically Tiled Arrays (HTAs), a data type with features for the management of locality and data parallelism. We have tested the expressiveness of HTAs for sequential, distributed and shared memory computations with excellent results in terms of readability and performance. The main reasons for the effectiveness of HTAs are their operators and their ability to directly represent tiles which are pervasive in high-performance computing.

Maria Garzaran
University of Illinois at Urbana-Champaign
garzaran@uiuc.edu

MS18

Matrix Abstractions and Storage Schemes in the PLASMA Library

Abstract unavailable at time of publication.

Jakub Kurzak, Jack Dongarra
University of Tennessee Knoxville
kurzak@eecs.utk.edu, dongarra@eecs.utk.edu

MS18

Microsoft Numerical Development Platform

In this talk we will discuss some challenges to enable numeric applications on current and future Windows platforms for many-core architectures. We will present new matrix abstractions that increase programmer productivity and how to use those abstractions for the implementation of scalable parallel matrix computations. Several examples demonstrating our matrix framework will be presented.

Mary Beth Hribar, Laurent Visconti
Microsoft
marybeth@microsoft.com, laurv@microsoft.com

MS19

Efficient FPGA Implementations of High-dimensional Cube Testers on the Stream Cipher Grain-128

Cube testers are highly parametrizable and parallelizable techniques for attacking ciphers. Grain is a stream cipher recently selected by the ECRYPT Network of Excellence as a promising cipher for hardware architectures, which can be efficiently implemented by using a x32 parallel datapath. Grain is thus an attractive target for a parallel implementation of cube testers. In this talk, I will present such an implementation on a Xilinx Virtex-5 FPGA, for the version of Grain with 128-bit key. This implementation computes high-complexity cube testers by running 256 instances of Grain simultaneously, and allowed us to run experiments involving 2^{54} clockings of the Grain-128 mechanism. Our results suggest that Grain-128 does not offer 128-bit security as expected, but rather (conjecturally) 83-bit security.

Jean-Philippe Aumasson
School of Computer and Communication Sciences
Ecole Polytechnique Federale de Lausanne
jeanphilippe.aumasson@gmail.com

Luca Henzen
Dept. of Information Technology and Electrical Engineering
Eidgenössische Technische Hochschule Zürich
henzen@iis.ee.ethz.ch

MS19

Parallel Cube Tester Analysis of the CubeHash One-way Hash Function

A cube tester is an algorithm that reveals information about the polynomial structure of a cryptographic primitive, needing only black-box access to the primitive. Although they require extensive computation, cube testers can be implemented in a massively parallel fashion. Using the Parallel Java Library, a cube tester for the CubeHash SHA-3 candidate one-way hash function was implemented to attempt to distinguish CubeHash from a random polynomial. The cube tester Java program was run on a 40-processor hybrid SMP cluster parallel computer. This paper reports the performance of the program and the results of the distinguishing attack on CubeHash.

Alan Kaminsky
Department of Computer Science
Rochester Institute of Technology
ark@cs.rit.edu

MS19

Parallel Performance of Some SHA-3 Second Round Candidate One-way Hash Functions

The National Institute of Standards and Technology is currently holding the second round of a competition to select SHA-3, the next federal hashing standard. Fourteen candidate algorithms remain in the contest, and in this presentation we consider the parallel performance of several of the algorithms. To create a fair comparison of the algorithms we place all of their compression functions inside of an identical tree-based structure and investigate their performance on a large multi-processor system when hashing extremely large messages.

Jason Martin
Department of Mathematics and Statistics
James Madison University
martinjw@jmu.edu

MS19

Accelerating Cryptographic Applications and Attacks with Multi-core Game Processors

In this work we evaluate the use of the Cell broadband engine and Graphics Processing Unit (GPU) as crypto-

logic accelerators. These processors are widely available in low-cost devices such as Sony's PlayStation 3 (PS3) video game consoles and NVIDIA graphics cards; the multi-core Cell (6 synergistic processing elements on the PS3) and many-core GPU (480 scalar processors on the GTX 295) can process many streams simultaneously, using single instruction, multiple data and single instruction, multiple threads techniques, respectively. We evaluate the performance of the AES block cipher, MD5 hash function, and Blake SHA-3 candidate on the PS3 and the GTX 295. Our results outperform the fastest FPGA implementations and are comparable to ASIC designs. Moreover, we discuss the inherent parallelizable nature of cryptanalytic attacks which allows for the use of a cluster of PS3s and graphics cards to launch full-scale practical attacks—e.g., creating rogue X.509 certificates by generating MD5 collisions.

Joppe Bos, Arjen Lenstra
School of Computer and Communication Sciences
Ecole Polytechnique Federale de Lausanne
joppe.bos@epfl.ch, arjen.lenstra@epfl.ch

Deian Stefan
Department of Electrical Engineering
The Cooper Union
stefan@cooper.edu

MS20

Experiences Developing and Using the Common Component Architecture

The Common Component Architecture Forum is a grassroots effort to develop tools and methodologies (inspired by the commercial component software approach) intended to improved software engineering of scientific applications in HPC environments. Ten years on, we have a variety of experiences in development and use of component software tools to report.

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

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

MS20

Software Life-cycle and Integration Issues Panel Discussion

This is a panel discussion to wrap up the two sessions of the mini-symposium. The panel members will include the various speakers and as a group with the audience will discuss issues such as software life-cycle, regulated backward compatibility, regressions in capability, basic software quality, continued maintenance and support, and tighter development and integration models for more aggressive R&D. Other issues raised during the sessions will be discussed as well.

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

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

Lois Curfman McInnes
Argonne National Laboratory
curfman@mcs.anl.gov

Bill Hoffman
Kitware
bill.hoffman@kitware.com

Mike Glass
Sandia National Labs
mwglass@sandia.gov

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

Roscoe A. Bartlett
Sandia National Laboratories
rabartl@sandia.gov

Barry F. Smith
Argonne National Lab
MCS Division
bsmith@mcs.anl.gov

Tony A. Drummond
Lawrence Berkeley National Laboratory
LADrummond@lbl.gov

Mattew Kneply
Argonne National Laboratory
Mathematics and Computer Science Division
knepley@gmail.com

MS20

Software Sustainability in The DOE ACTS Collection

The DOE Advanced Computational Software (ACTS) Collection Project has pursued and implemented mechanisms to bring a set of advanced and portable software tools that can be used in the development of high performance simulation codes to a broad community of computational scientists. ACTS efforts are geared at leveraging application development using ACTS tools and long-term support for the functionality available in the collection. Here we discuss our software sustainability infrastructure.

Leroy A. Drummond
Computationa Research Division
Lawrence Berkeley National Laboratory
LADrummond@lbl.gov

MS20

Agile Software Development for the SIERRA Project

The Sierra modeling and simulation software project has successfully adopted an agile software development lifecycle. Here I will describe our adoption process and the advantages of the agile development life-cycle for our environment — especially in the area of collaboration, soft-

ware quality, testing and the integration with other software projects.

Mike Glass
Sandia National Labs
mwglass@sandia.gov

MS21
Enabling Resilient Peta-Scale Systems

As large-scale systems grow to millions of components, they become increasingly unreliable. Indeed, systems such as the BlueGene/L and ASCI Q fail multiple times per day. A component failure may make the system unusable, cause performance degradation or even introduce subtle errors into the application's output. To enable productive use of large-scale systems into the peta- and exa-scale era we must develop techniques to enable systems and applications to run with unreliable components. To this end we need to

- Help users develop resilient applications by identifying for them the critical application and system components most in need of hardening and the most significant failure modes, and
- Given the manifestation of a fault as an application error, identify the system component that is the root cause of this error.

I will talk about our ongoing research on tools that enable the above capabilities and how they can be applied to support the continued high productivity of large-scale computers.

Greg Bronevetsky, Bronis R. de Supinski
Lawrence Livermore National Laboratory
greg@bronevetsky.com, bronis@llnl.gov

MS21
Effects of Soft Errors on Scientific Application Outcome: Will They Result in Bad Science?

In scientific computation humans are typically in the loop and thus offer some inherent fault tolerance. Scientists who rely on computation already observe occasional non-fatal computational glitches, possibly from soft errors, and simply throw out results that "look bad.' With this human-in-the-loop model, the soft-error question becomes "What sort of computational error will escape human notice and result in erroneous scientific conclusions?' We present initial experimental results that attempt to make progress on that question.

Jonathan Cook
New Mexico State University
joncook@nmsu.edu

MS21
Some Experiments with Fault Tolerent Linear Algebra Algorithms

Abstract unavailable at time of publication.

Jack Dongarra
University of Tennessee Knoxville
dongarra@eecs.utk.edu

MS21
Soft Errors in Linear Solvers as Integrated Components of a Simulation

In this work, we investigate the sensitivity of linear solvers to soft errors. In many simulations, linear algebra accounts for more than 80% of the computational time, making it a critical component in over-all fault tolerance. However, the linear solver does not function in isolation and must be considered as an integrated component of the algorithmic stack in a simulation. We investigate the effects of soft errors in the linear solver on the overall simulation.

Victoria E. Howle
Texas Tech University
victoria.howle@ttu.edu

MS22
Autotuning Multigrid with PetaBricks

Algorithmic choice is essential in any problem domain to realizing optimal computational performance. We present a programming language and autotuning system that address issues of algorithmic choice, accuracy requirements, and portability for multigrid applications in an efficient manner. We search the space of algorithmic choices and cycle shapes efficiently by utilizing a novel dynamic programming method to build tuned algorithms from the bottom up. Our implementation uses PetaBricks, an implicitly parallel programming language and autotuning system where algorithmic choices are exposed in the language.

Cy Chan
Electrical Engineering and Computer Science Department
Massachusetts Institute of Technology
cychan@mit.edu

MS22
Tuning Dense Linear Algebra Libraries on Multicore and GPU Architectures

Tuning numerical libraries has become more difficult over time, as modern systems get more sophisticated. We present the current research on tuning the PLASMA and MAGMA dense linear algebra libraries. These libraries are designed to run efficiently on correspondingly homogeneous multicore and hybrid multicore+GPU systems. The approaches use combination of empirical and model driven optimization techniques. This work can be a basis for autotuning these libraries and enabling easy performance portability across hardware systems.

Rajib Nath
Department of Electrical Engineering and Computer Science
University of Tennessee, Knoxville
rnath1@eecs.utk.edu

Jack Dongarra
University of Tennessee Knoxville
dongarra@eecs.utk.edu

Stanimire Tomov
Innovative Computing Laboratory
University of Tennessee; Knoxville
tomov@eecs.utk.edu

Emmanuel Agullo
Department of Electrical Engineering and Computer Science
University of Tennessee, Knoxville

eagullo@eecs.utk.edu

MS22

Model-driven Autotuning of Sparse Matrix-vector Multiply for GPUs

We present a performance model-driven framework for automated performance tuning (autotuning) of sparse matrix-vector multiply (SpMV) on systems accelerated by graphics processing units (GPU). Our best implementations achieves up to 29.0 Gflop/s in single-precision on the NVIDIA Tesla T10 multiprocessor on standard sparse matrix benchmark suites, or up to 1.8x faster than prior art. Our model-driven approach can select an implementation with performance within 15% of the best found by exhaustive search.

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

Jee Whan Choi
School of Electrical and Computer Engineering
Geogia Institute of Technology
jee@ece.gatech.edu

Amik Singh
Indian Institute of Technology, Roorkey
amik.iitr@gmail.com

MS22

Memory-Efficient Optimization of Gyrokinetic Particle-to-Grid Interpolation for Multicore Processors

In this talk, we compare a series of multicore parallelization strategies for the particle-to-grid interpolation step in the Gyrokinetic Toroidal Code (GTC). This step involves particles scattering their charge onto a 3D toroidal mesh with the possibility of collisions. As such, efficient multicore synchronization becomes essential. We implement thirteen different variants for this kernel and identify the best-performing ones given typical PIC parameters. We find that our best strategies can be 2x faster than the reference optimized MPI implementation while consuming as little as 1% of the DRAM capacity.

Samuel Williams
Lawrence Berkeley National Laboratory
swwilliams@lbl.gov

Kamesh Madduri
Lawrence Berkeley National Lab
kmadduri@lbl.gov

Stephane Ethier
Princeton Plasma Physics Laboratory
ethier@pppl.gov

Leonid Oliker, John Shalf, Eric Strohmaier, Katherine Yelick
Lawrence Berkeley National Laboratory
loliker@lbl.gov, jshalf@lbl.gov, estrohmaier@lbl.gov, kayelick@lbl.gov

MS23

Communication-Avoiding Linear Algebra

We give an overview of recent developments in Communication-Avoiding Linear Algebra, whose goal is to develop algorithms that reduce, and ideally provably minimize, the movement of data (between levels of a memory hierarchy, or between processors on a network). We discuss lower bounds on data movement for dense and sparse linear algebra that conventional algorithms do not attain, and outline new algorithms that do.

James Demmel
UC Berkeley, USA
demmel@eecs.berkeley.edu

MS23

Communication Avoiding LU and QR Factorizations on Multicore Architectures

We discuss communication-avoiding algorithms based on CAQR and CALU (recently introduced for distributed memory architectures) that are adapted to multicore architectures. For matrices that are tall and skinny, that is, with many more rows than columns, multithreaded CALU and multithreaded CAQR outperform the corresponding routines dgetrf and dgeqrf from Intel MKL library by up to a factor of 2 and 7 respetively on a two-socket, quad-core machine based on Intel Xeon EMT64 processor.

Simplice Donfack
INRIA, France
sidonfack@gmail.com

Laura Grigori
INRIA
France
Laura.Grigori@inria.fr

MS23

Implementing Communication-Optimal Parallel and Sequential Qr Factorizations

We present parallel and sequential dense QR factorization algorithms for tall and skinny matrices and general rectangular matrices that both minimize communication, and are as stable as Householder QR. The sequential and parallel algorithms for tall and skinny matrices lead to significant speedups in practice over some of the existing algorithms, including LAPACK and ScaLAPACK, for example up to 6.7x over ScaLAPACK. The parallel algorithm for general rectangu- lar matrices is estimated to show significant speedups over ScaLAPACK, up to 22x over ScaLAPACK.

Julien Langou
Department of Mathematical & Statistical Sciences
University of Colorado Denver
julien.langou@ucdenver.edu

MS23

Lower Bounds on Communication

A lower bound on communication (words of data moved) needed to perform dense matrix multiplication is known to be $\Omega(f/\sqrt{M})$ where f is the number of multiplications $O(n^3)$ and $M$is the fast memory size. We extend this result to all direct methods of linear algebra (BLAS, LU, QR, eig, svd, etc), to sequential or parallel algorithms, and to

dense or sparse matrices. We show large speedups over algorithms in LAPACK and ScaLAPACK.

Grey Ballard
UC Berkeley
ballard@cs.berkeley.edu

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

Olga Holtz
UC Berkeley and
TU Berlin
holtz@math.berkeley.edu

Oded Schwartz
TU Berlin
oded.schwartz@gmail.com

MS24

Optimizing and Tuning the Fast Multipole Method for Multicore and Accelerator Systems

We discuss single-node architectural and algorithmic tuning of the parallel kernel-independent fast multipole method [Ying, Biros, Zorin, and Harper (2003)]. We consider several modern multicore platforms, including those based on quad-core Intel Nehalem, quad-core AMD Barcelona, multithreaded eight-core Ultra Sparc T2+, and NVIDIA T10 GPU processors. We present various optimization and scalable parallelization strategies that achieve high performance without loss of accuracy for both single and double precision implementations.

Aparna Chandramowlishwaran, IIya Lashuk, Aashay Shringarpure
Georgia Institute of Technology
aparna@cc.gatech.edu, ilashuk@cc.gatech.edu, aashay.shringarpure@gmail.com

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

George Biros, Richard Vuduc
Georgia Institute of Technology
gbiros@acm.org, richie@cc.gatech.edu

MS24

High Performance Computing Using MPI and OpenMP in Multi-core and Many-core Environments

The rapidly increasing number of cores in modern processors are pushing the current high performance computing (HPC) systems into petascale and exascale era. The hybrid nature of these systems distributed memory across nodes and NUMA shared memory within each node poses a great challenge to application developers. In this talk, we describe our experiences with the two traditional programming models, MPI and OpenMP and combinations thereof, for large-scale computations on these systems. We will also present new data locality extensions to OpenMP to better match the hierarchical structure of multi-core architectures. We will discuss our runtime library implementation of these extensions and their evaluation using benchmarks.

Haoqiang Jin
NASA
haoqiang.jin@nasa.gov

Rupak Biswas, Piyush Mehrotra
NASA Ames Research Center
Rupak.Biswas@nasa.gov, piyush.mehrotra@nasa.gov

Lei Huang, Barbara Chapman
University of Houston
leihuang@cs.uh.edu, chapman@cs.uh.edu

MS24

Experiences from the Roadrunner Petascale Hybrid System

The combination of flexible microprocessors (AMD Opterons) with high-performing accelerators (IBM PowerXCell 8i) resulted in the extremely powerful Roadrunner system. Many challenges in both hardware and software were overcome to achieve its goals. In this talk we detail some of the experiences in achieving performance on the Roadrunner system. In particular we examine several implementations of the kernel application, Sweep3D, using a work-queue approach, a more portable Thread-building-blocks approach, and an MPI on the accelerator approach.

Darren Kerbyson, Scott Pakin, Michael Lang, Jose Sancho, Kei Davis, Kevin Barker, Josh Peraza
Los Alamos National Laboratory
djk@lanl.gov, pakin@lanl.gov, mlang@lanl.gov, jcsancho@lanl.gov, kei@lanl.gov, kjbarker@lanl.gov, jperaza@lanl.gov

MS24

Hierarchical Auto-Tuning of a Hybrid Lattice Boltzmann Computation

This work quantifies the application-level performance benefits of auto-tuning local computation, domain decomposition, and division of computational resources among processes via a hybrid MPI/pthreads implementation. Our study examines a Lattice Boltzmann Magneto-hydrodynamics computation and shows that on the massively-parallel Cray XT4 — although the bulk of performance gains come from single-thread optimization — an additional 17% boost is achieved through tuning the domain decomposition and balance between threads per process and processes per node.

Leonid Oliker, Samuel Williams, John Shalf, Jonathan Carter
Lawrence Berkeley National Laboratory
loliker@lbl.gov, swwilliams@lbl.gov, jshalf@lbl.gov, jtcarter@lbl.gov

MS25

Parallel Adaptive Multiscale Simulation, Error Control and Uncertainty Quantification via Libmesh

We consider parallel adaptive simulation of multiscale problems, parameter sensitivity and uncertainty quantification. Recent extensions of the Libmesh software framework to incorporate adjoint approaches for error control,

parameter sensitivity and uncertainty analysis will be summarized and results of multiscale adaptive simulations presented. The re-entry space vehicle simulation being undertaken by our PECOS Center is utilized as an illustrative large-scale application.

Graham F. Carey
Institute for Computational Engineering and Sciences (ICES)
University of Texas at Austin
carey@cfdlab.ae.utexas.edu

Roy Stogner
ICES, UT Austin
roy@stogners.org

Vikram Garg
CFDLab, U Texas, Austin
vikramvgarg@mail.utexas.edu

Varis Carey
Colorado State University
carey@math.colostate.edu

Ben Kirk
NASA-JSC
benjamin.kirk-1@nasa.gov

Paul Baumann
CES,UT
pbauman@ices.utexas.edu

MS25

Reordering Effects on Block-Ilu Preconditioners for Parallel Amr/c

The effects of reordering the unknowns on the convergence of ILU preconditioned Krylov subspace methods are investigated. Of particular interest is the resulting preconditioned iterative solver behavior when adaptive mesh refinement and coarsening (AMR/C) are utilized for parallel simulations. In this case, reordering is applied to local subdomains for block ILU preconditioning and subdomains are repartitioned dynamically as mesh adaptation proceeds. As representative schemes, we consider the familiar Reverse Cuthill-McKee (RCM) and Quotient Minimum Degree (QMD) algorithms applied with ILU preconditioners Krylov solvers. Numerical studies for representative applications are conducted using the object oriented AMR/C software system libMesh linked to the PETSc solver Library. Experimental results indicate the combined effects of local reordering for subdomain block ILU preconditioning and of dynamic repartitioning due to AMR can improve overall computational effciency.

Alvaro Coutinho
Dept. of Civil Engineering
COPPE/Federal University of Rio de Janeiro
alvaro@nacad.ufrj.br

Jose Camata
Center for Parallel Computing
COPPE/Federal University of Rio de Janeiro, Brazil
camata@gmail.com

Andre Rossa
Dept of Cvil Engineering
COPPE/Federal University of Rio de Janeiro, Brazil
andre_rossa@hotmail.com

Andrea Valli, Lucia Catabriga
Computer Science
Federal University of Espirito Santo, Brazil
avalli@inf.ufes.br, luciac@inf.ufes.br

Graham F. Carey
Institute for Computational Engineering and Sciences (ICES)
University of Texas at Austin
carey@cfdlab.ae.utexas.edu

MS25

Hybrid Parallelism for Preconditioned Jacobian-Free Newton-Krylov Multiphysics Applications

Increasingly, processor manufacturers are turning to multicore packages as a way of increasing compute power. This trend has lead to hybrid cluster architectures where each traditional distributed compute node now contains as many processors as small scale clusters from just a few years ago. For parallel algorithm developers this presents new challenges in efficiently utilizing these resources. This study presents a technique for exploiting the mathematical structure of Jacobian Free Newton-Krylov to leverage existing solver frameworks while taking advantage of hybrid architectures. This is accomplished through the simultaneous use of both MPI and threading. The efficacy of this approach is explored through the use of Idaho National Laboratorys MOOSE (Multiphysics Object Oriented Simulation Environment) framework [D. Gaston et al., MOOSE: A parallel computational framework for coupled systems of nonlinear equations, Nucl. Engrg. Design 239 (2009) 1768] to study several multiphysics applications in different parallel environments (i.e. shared memory, distributed memory and hybrid).

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

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

MS25

Leveraging the libMesh Library for SUPG Finite Element Simulations of Hypersonic Flows in Thermochemical Nonequilibrium

We present recent results for predictive simulations of hypersonic flows in thermochemical nonequilibrium. Specific implementation details of the SUPG algorithm that enhance its stability for this problem class will be discussed. The parallel adaptive solution algorithm will be presented, and in particular we will overview how the libMesh framework is used as an enabling technology for the scheme. Results will be presented for a number of hypersonic aerothermodynamics applications, including h-adaptive simulations of dissociating, high-enthalpy flows.

Ben Kirk
University of Texas at Austin
benjamin.kirk-1@nasa.gov

Graham F. Carey
Institute for Computational Engineering and Sciences (ICES)
University of Texas at Austin

carey@cfdlab.ae.utexas.edu

Roy Stogner
ICES, UT Austin
roy@stogners.org

Vikram Garg
CFDLab, U Texas, Austin
vikramvgarg@mail.utexas.edu

Varis Carey
University of Texas at Austin
not submitted

MS26

Parallel Graph Matching and Coloring for Petascale Scientific Computing

We present new massively parallel algorithms for two prototypical graph problems: a parallel half-approximation algorithm for edge- weighted matching, and a framework for parallel distance-1 vertex coloring. We use approximation, partitioning, randomization, speculation, and optimized communication to achieve scalability. We describe implementations on massively parallel processors and provide experimental results on the IBM BlueGene/P. We demonstrate scalable results for up to 16,384 cores for model problems and problems from circuit simulations.

Mahantesh Halappanavar
Old Dominion University
MHalappa@odu.edu

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

Florin Dobrian
Columbia University
dobrian@cs.odu.edu

Assefaw H. Gebremedhin
Purdue University
agebreme@purdue.edu

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

MS26

Amorphous Data Parallelism in Graph Algorithms

Over the past thirty years, most research in the programming languages and compilers community has focused on regular algorithms, which manipulate dense arrays and matrices; the parallelism in these algorithms is well understood. However, irregular algorithms, which manipulate graphs, are less understood; in such algorithms, it is much harder to identify and exploit parallelism. This talk presents a new type of parallelism, *amorphous data parallelism* that is prevalent in irregular algorithms and can be readily exploited.

Milind Kulkarni
School of Electrical and Computer Engineering
Purdue University
milind@purdue.edu

MS26

Active Messages for Parallel Graph Computations

Graph algorithms are difficult to parallelize because they typically involve irregular, non-local, structures with fine-grained dependencies. Logically, graph computations can be represented as messages that are passed between structural elements of the graph. In actual implementations based on message-passing, individual messages tend to be small and frequent relative to computation, leading to poor performance. In this talk we discuss the application of the active-message paradigm to graph computations. This approach assumes that numerous small messages will be sent, will arrive asynchronously, and will each perform a simple operation at the receiver. These characteristics match the needs of parallel graph computations and promise better parallel efficiency because of lower synchronization and greater concurrency. Unfortunately, current implementations of the active-message paradigm are too limited to fully capture the fine-grained dependency structure of graph algorithms. We describe a generalization of the active message paradigm that provides high performance on fine-grained algorithms for memory-latency bound problems (graph algorithms in particular). We also demonstrate that our generalization provides high performance even on current networking hardware.

Nicholas Edmonds, Jeremiah Willcock,
Andrew Lumsdaine
Open Systems Laboratory
Indiana University
ngedmond@osl.iu.edu, jewillo@osl.iu.edu,
lums@cs.indiana.edu

MS26

Simulating Diffusion Processes on Very Large Complex Networks

The spread of transmitted phenomena referred to as contagion (e.g. opinions, attitudes, beliefs, worms) in social and communication networks any actual or virtual contagion transmitted in biological, social, organizational, computing and communication networks is a well-known complex problem. The recent fear of H1N1 (swine-flu) epidemics and the cascading failures in national and global financial sectors serve as excellent examples. The rapidly-growing social and communication networks such as Twitter and Facebook are emerging classes of networks over which diverse contagion spreads. The talk focuses on developing high performance computing oriented highly scalable simulations of contagion like diffusion processes on very large complex networks. Developing highly scalable simulations of contagions on large complex networks is complicated and scientifically challenging due to their size, their co-evolutionary nature and the need for simulating multiple contagion processes simultaneously. In this talk, we will describe a collection of methods for simulating contagions that are specifically geared towards overcoming the above challenges. Empirical results show that the methods scale quite easily to 10-100 million nodes and 500-5 billion edge networks. Practical usefulness of these methods will be described via well chosen case studies.

Madhav Marathe, Keith Bisset, Jiangzhou Chen, Xizhou Feng
Virginia Bioinformatics Institute, VA Tech
mmarathe@vbi.vt.edu, kbisset@vbi.vt.edu,

chenj@vbi.vt.edu, fengx@vbi.vt.edu

Anil Vullikanti
Dept. of Computer Science, and Virginia Bioinformatics Inst.
Virginia Tech
akumar@vbi.vt.edu

MS27

Parallel Adaptive Simulation of Low Mach Number Flows

Numerical simulation of reacting low Mach number flows can be computationally demanding, often requiring high-fidelity modeling of the fluid mechanics, complex reaction networks, and high spatial and temporal resolution. In this talk we discuss how we combine a low Mach number formulation with a local adaptive mesh refinement algorithm to model flows such as the convective period preceding ignition of a Type Ia supernova.

Ann S. Almgren
Lawrence Berkeley National Laboratory
ASAlmgren@lbl.gov

John B. Bell
CCSE
Lawrence Berkeley Laboratory
jbbell@lbl.gov

Mike Lijewski, Andy Nonaka
Lawrence Berkeley National Laboratory
mjlijewski@lbl.gov, ajnonaka@lbl.gov

Mike Zingale
Stony Brook University
mzingale@mail.astro.sunysb.edu

MS27

Modeling of ICF Experiments using a New Multi-physics Code, ALE-AMR

Advances in high performance computing are allowing the simulation of new phenomena such as late-time aspects of ICF experiments to control debris/shrapnel impacts. We describe a new 3D multimaterial, multi-physics code, ALE-AMR, which models the needed wide range of spatial and temporal scales. We discuss numerical issues associated with treating hot radiating plasmas and cold fragmenting solids in the same simulation. While developed for ICF applications, the code is appropriate to a range of applications.

David C. Eder, Aaron Fisher, Nathan D. Masters
Lawrence Livermore National Laboratory
eder1@llnl.gov, fisher47@llnl.gov, masters6@llnl.gov

Alice Koniges
Lawrence Berekely Laboratory
aekoniges@lbl.gov

MS27

Multimaterial Remapping in ALE-AMR

The combination of Adaptive Mesh Refinement and Arbitrary Lagrangian-Eulerian hydrodynamic simulation in the ALE-AMR hydrocode provides computational efficiency while maintaining accuracy, but also poses a number of challenges. We will present our recent work on interface reconstruction, remapping, and refinement in ALE-AMR.

Nathan D. Masters, Robert W. Anderson, David C. Eder, Aaron C. Fisher
Lawrence Livermore National Laboratory
masters6@llnl.gov, anderson110@llnl.gov, eder1@llnl.gov, fisher47@llnl.gov

Alice E. Koniges
National Energy Research Scientific Computing Center
Lawrence Berkeley National Laboratory
aekoniges@lbl.gov

MS27

Asymptotics and Computation of Microphase Separation of Diblock Copolymers

Diblock copolymer melts, dubbed "designer materials', have the remarkable ability for self-assembly into various ordered structures. These structures are key to the many properties that make diblock copolymers of great technological interest. The density functional theory of Ohta and Kawasaki leads to a nonlocal variational problem, and presents an excellent setting for the analysis of microphases. In this talk, I will discuss some asymptotic and computational techniques and results on the phase diagram of these materials in this setting.

J.F. Williams
Simon Fraser University
Dept. of Mathematics
jfw@math.sfu.ca

Rustum Choksi
Department of Mathematics
Simon Fraser University
choksi@cs.sfu.ca

MS28

Building a Large Parallel Quantum Computer

Quantum computers made up of a few quantum bits (qubits) have recently been built and tested. The field of quantum computation is now gearing up for the challenge of scaling these technologies to a quantum computer consisting of many hundreds to thousands of qubits. Here I will discuss one such proposal, MUSIQC, which allows for a high degree of parallel operation: something which is necessary for building a fault-tolerant quantum computer.

David Bacon
Computer Science and Engineering
University of Washington
dabacon@cs.washington.edu

MS28

Threads Should Not Play Dice: Determinism for Multithreaded Programs

Current multicore systems are nondeterministic. This frustrates debugging and limits the ability to properly test multithreaded code. In this talk I will present fully deterministic shared memory multiprocessing (DMP). The behavior of an arbitrary multithreaded program on a DMP system is only a function of its inputs. Previous approaches focused

on replay, useful only for debugging. In contrast, while DMP systems are directly useful for debugging, we argue that parallel programs should execute deterministically in the field as well.

Luis Ceze
Computer Science and Engineering
luisceze@cs.washington.edu

MS28

Automatic Generation of Code-Centric Graphs for Understanding Shared-Memory Communication

With multicore architectures leading to more and more complex pieces of software using explicit threads, tools to understand the structure of such software are increasingly important. This talk presents a novel tool that automatically builds a graph that describes how parallel shared-memory programs communicate. Compared to prior work, the communication graphs are code-centric: nodes represent pieces of code (e.g., functions) and edges represent inter-thread shared-memory communication via the pieces of code.

Daniel Grossman
Computer Science and Engineering
University of Washington
djg@cs.washington.edu

MS28

New Language Constructs for Parallel Abstractions

Traditional methods for writing parallel programs such as message passing, threading, etc. rely heavily on a procedural interface. As a result the conceptual payload of each operation a programmer writes is small. Procedure calls just list arguments; so programming is difficult. Control structures like for each for element-wise data parallelism help, but abstractions like divide and conquer, and parallel prefix need new language constructs. The talk presents new constructs for the parallel prefix abstraction.

Larry Snyder
Dept of Computer Science and Engineering
snyder@cs.washington.edu

MS29

Dynamic Community Structure in Time-varying Biological Networks

As organisms grow, biological pathways turn on and off. We describe new methods for analyzing dynamic biological networks through time-evolution of a hierarchical model [Park Moore Bader 2009 PLoS ONE, in press]. Analysis of experimental data for the growing plant root reveals hierarchical levels of organization for protein complexes and pathways, which change as the network ages in the maturing root. These methods are generally applicable to networks that evolve over space and and time.

Joel S. Bader, Yongjin Park
Johns Hopkins University
Department of Biomedical Engineering
joel.bader@jhu.edu, ypark28@jhu.edu

MS29

Degeneracies, Hierarchies, Resolution Limits and the Community Detection Problem

Identifying modular structure in complex networks is a fundamental task for understanding the function, dynamics, robustness and evolution of complex biological, technological and social systems. However, the accuracy of the most popular technique, which is based on optimizing the quantity called modularity, remains poorly characterized in practical situations. In this talk, I'll show that the modularity function Q exhibits extreme degeneracies: that is, an exponential number of high-modularity but structurally dissimilar solutions. This fact implies several cautionary points: (i) the more modular a network, the more rugged the modularity function and the less distinct the function's peak, (ii) high-modularity solutions can disagree strongly on the composition of even the largest modules, and (iii) the maximum modularity score depends strongly on the size of the modular network. Thus, the results of maximizing modularity should be interpreted conservatively. I'll conclude with a brief discussion alternative avenues for accurately identifying modular structure, e.g., by combining information from many degenerate solutions or using generative models, and the relevance of degeneracies to their use in scientific applications.

Benjamin Good
Swarthmore College
conkerll@gmail.com

Yves-Alexandre de Montjoye
Universite Catholique de Louvain
yvesalexandre.de.montjoye@gmail.com

Aaron Clauset
Santa Fe Institute
aaronc@santafe.edu

MS29

Parallelizing Community Detection Algorithms For Large-Scale Networks

In this talk, we will speculate on parallel algorithms that might considerably increase the size of networks for which we can detect community structure. We will consider both (approximately) computing sophisticated currently-serial statistically-rigorous methods and (approximately) computing simpler heuristics. Our goal is to apply rigorous methods to moderate-sized graphs and to heuristically detect communities in huge graphs.

Jonathan Berry, Cynthia Phillips
Sandia National Laboratories
jberry@sandia.gov, caphill@sandia.gov

Ali Pinar
Sandia National Labs
apinar@sandia.gov

MS29

Networks in Modeling and Simulation

As networks are becoming the standard tool to model interactions between entities of complex, distributed interconnected systems, simulations with networks draw increasing interest. Simulations on a network, simulations of a network, as well as simulations to infer a network all have important applications. A common challenge for all these applications is characterization of network structure. This

talk will summarize our efforts in this area.

Ali Pinar
Sandia National Labs
apinar@sandia.gov

MS30

Ultra High Resolution and High Performance Simulation of Compressible Atmospheric Flows over Urban Area

An efficient fully compressible scheme is applied to simulation of ultra high-resolution atmospheric boundary layer over urban area resolving buildings. Computational acceleration rate compared to horizontal explicit vertical implicit, HEVI method is estimated, and comparison with horizontal implicit vertical implicit, HIVI method is discussed. Several validation tests are also performed. The results indicate the compressible approach is proper under the condition coupled to mesoscale dynamics where buildings wall is highly heated.

Yuya Baba
JAMSTEC
babay@jamstec.go.jp

Keiko Takahashi
The Earth Simulator Center, JAMSTEC
takahasi@jamstec.go.jp

MS30

Parallel Regridding Schemes in the Distributed Coupling Toolkit

The Distributed Coupling Toolkit (DCT) is a coupling library that supports the data management and exchange of model data between parallel computational models. The DCT has a simple and flexible interface and its functionality is implemented in a way that minimizes synchronization points and avoids the needs for a central coupler entity. Here we discuss recent functionality that has been added and compared to other centralized approaches.

Leroy A. Drummond
Computationa Research Division
Lawrence Berkeley National Laboratory
LADrummond@lbl.gov

Dany DeCechis
Computational Science Research Center
San Diego State University
dececchi@sciences.sdsu.edu

MS30

Framework for Development of Parallel Codes in Integrated Predictive Simulation System for Earthquake and Tsunami Disaster

In this talk, framework for development of parallel codes in the project ofIntegrated Predictive Simulation System for Earthquake and Tsunami Disaster will be described. This is a 5-year project from FY.2005, supported by Japanese Government, and will be the first integrate simulation system for prediction of earthquake and tsunami disasters, which covers entire multi-scale processes such as plate deformation, dynamic fault rupture, seismic wave/tsunami propagation, and oscillation of buildings.

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

Tsuyoshi Ichimura
Earthquake Research Institute
The University of Tokyo
ichimura@eri.u-tokyo.ac.jp

Toshio Nagashima
Sophia University
nagashim@me.sophia.ac.jp

MS30

Constrained Interpolation for Remapping in Esmf

We present a correction algorithm for conserving integral quantities when a field is interpolated between grids. Many grid types are supported, including 2D grids composed of triangles and quadrilaterals or 3D grids composed of tetrahedrons and hexahedrons. The interpolation matrix and conservation correction terms are computed in parallel, either offline or during runtime if using the Earth System Modeling Framework (ESMF), and the correction reflects the locality and order of the interpolation method.

Ryan O'Kuinghttons
ESMF-NCAR
rokuingh@ucar.edu

Amik St-Cyr
National Center for Atmospheric Research
Institute for Mathematics Applied to the Geosciences
amik@ucar.edu

Robert Oehmke
ESMF-NCAR
oehmke@ucar.edu

MS31

Chapel: the Cascade High Productivity Language

Chapel is a new programming language being developed by Cray Inc. as part of the DARPA-led High Productivity Computing Systems program (HPCS). Chapel strives to increase productivity for parallel users by greatly improving programmability as compared to today's languages/libraries while matching or surpassing their performance and portability. In this talk, I will provide an overview of the Chapel language including motivating philosophies and recent results. I'll also mention several opportunities for collaboration and future work.

Brad Chamberlain
Cray Inc.
bradc@cray.com

MS31

UPC at Scale

The shared memory programming model found in UPC and other PGAS programming models provide a simple way to build up complicated distributed data structures, thereby aiding programming productivity. The focus of this talk will be techniques in UPC to get applications

to scale out to massive levels of parallelism. We first show how the one-sided communication model found in UPC allows applications to scale better than their two sided counterparts. We then focus on applications that exploit nonblocking communication to maximize communication/computation overlap. We finally discuss the collective communication found in UPC and demonstrate how a combination of automatic tuning to tune these operations and transforming applications to take advantage of these operations can lead to significant performance improvements.

Rajesh Nishtala
UC Berkeley
rajeshn@cs.berkeley.edu

MS31

Combining Object-Oriented Techniques with Co-arrays in Fortran 2008

Fortran 2003 is an object-oriented language. Fortran 2008 incorporates the co-array parallel programming model. Declaring objects as co-array objects gives the programmer a powerful new tool for writing modern, efficient parallel applications entirely within the Fortran language. Examples related to weather and climate models illustrate how to combine these two new features of the language.

Robert Numrich
University of Minnesota
rwn@msi.umn.edu

MS31

Determinate Parallel Programming

We show that the X10 finish/async/at concurrency and distribution framework is particularly well-suited to develop safe parallel programs. A program is safely parallel if its semantics is identical to its serial elision (obtained by replacing finish S by S, and async S by S). Safety can be established for many HPC programs through a compositional, symblic effect analysis which establishes the commutativity of two steps which may happen in parallel. The analysis framework is the basis for a concurrency refactoring engine we are building.

Vijay Saraswat, Robert Fuhrer, Olivier Tardieu, Emmanuel Geay
IBM Research
vsaraswa@us.ibm.com, rfuhrer@us.ibm.com, tardieu@us.ibm.com, egeay@us.ibm.com

MS32

Massively Parallel Low Cost Uncertainty Quantification

Abstract unavailable at time of publication.

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

MS32

Parallel Computing for Accelerator Cavity Modeling

Abstract unavailable at time of publication.

Lie-Quan (Rich) Lee
Stanford Linear Accelerator Center
liequan@slac.stanford.edu

MS32

A Parallel Algorithm and Implementation of Selected Inversion of a 2D Kohn-Sham Hamiltonian

We present an efficient parallel algorithm and implementation for computing the diagonal of the inverse of a 2D Kohn-Sham Hamiltonian. Such an algorithm can be used to speed up electronic structure calculations for quantum mechanical systems. We show how an elimination tree is used to organize the parallel computation and how synchronization overhead is minimized. We will present the performance characteristics of our implementation and show that it is much faster than the existing approaches.

Lin Lin
Princeton University
linlin@math.princeton.edu

MS32

Large-scale Parallel Flow Simulations for Nuclear Fuel Performance

Design and testing of nuclear fuel for reactors is currently an iterative experimental process with limited computational support and a cycle time of 15-20 years. Through US DOE NEAMS (Nuclear Energy Advanced Modeling and Simulation) program, a set of tools are being developed to build predictive models targeted to simulate nuclear fuel and reduce the design cycle time. Initial results from the flow simulations along with the parallel scalability will be presented.

Sreekanth Pannala
Computer Science and Mathematics Division
Oak Ridge National Laboratory
pannalas@ornl.gov

Srikanth Allu, Emilian Popov
Oak Ridge National Laboratory
allus@ornl.gov, popove@ornl.gov

MS33

Fully Implicit Domain Decomposition Methods for a Global Shallow Water Model on the Cubed Sphere

In this talk, we discuss the parallel performance of a PETSc-based implementation of an overlapping domain decomposition method for solving the global shallow water equations on the cubed-sphere. Most existing techniques for climate modeling are explicit and semi-implicit, which often have some restrictions on the time step size. We present a fully implicit approach with a well-balanced high order finite volume discretization, and report the parallel scalability of a Newton-Krylov-Schwarz method on machines with thousands of processors.

Chao Yang
Software Institute, Chinese Academy of Sciences
P. R. China
yangchao.china@gmail.com

Xiao-Chuan Cai
University of Colorado, Boulder
cai@cs.colorado.edu

MS33

A PETSc-based Jacobi-Davidson Approach for Large Sparse Polynomial Eigenvalue Problems with Application in Quantum Dot Simulation

Abstract unavailable at time of publication.

Feng-Nan Hwang
National Central University
hwangf@math.ncu.edu.tw

MS33

UNIC - Massively Parallel Reactor Core Analysis Code

UNIC code is being developed at Argonne National Laboratory to carry out detailed, high-resolution simulations of neutron transport in fast reactor cores. When the seven-dimensional neutron transport equation is discretized, several large-scale linear systems need to be solved. In this presentation, we will describe our experience of using the PETSc library to do some scalable simulations of neutron transport on the leadership computing platforms (such as Blue Gene/P at Argonne and XT5 at ORNL).

Dinesh K. Kaushik
Argonne National Laboratory
2151, Bldg. 240, MCS Div
kaushik@mcs.anl.gov

Micheal Smith, Allan Wollaber
Argonne National Laboratory
masmith@mcs.anl.gov, awollaber@anl.gov

Barry F. Smith
Argonne National Lab
MCS Division
bsmith@mcs.anl.gov

Won Sik Yang
Argonne National Laboratory
wyang@anl.gov

MS33

High Fidelity Simulations of Groundwater Flow and Reactive Transport with Pflotran + Petsc on Ultrascale Computers

To provide true predictive utility, subsurface simulations often must accurately resolve complicated, multi-phase flow fields in highly heterogenous geology with numerous chemical species and complex chemistry, incorporating processes operating over a very wide range of spatial and temporal scales. We describe some experiences with PFLOTRAN, a research code built on top of the PETSc framework that is being developed to solve such problems using the largest-scale supercomputers in the world.

Richard Mills
Oak Ridge National Laboratory
rmills@climate.ornl.gov

Glenn Hammond
Pacific Northwest National Laboratory
glenn.hammond@pnl.gov

Peter Lichtner
Division of Earth and Environmental Sciences
Los Alamos National Laboratory
lichtner@lanl.gov

Chuan Lu
University of Utah
clu@egi.utah.edu

Barry F. Smith
Argonne National Lab
MCS Division
bsmith@mcs.anl.gov

Bobby Philip
Oak Ridge National Laboratory
philipb@ornl.gov

MS34

Modeling and Simulation of Particle Laden Flow

We consider the problem of particle laden flow on an incline. Experimental results show that there is a transition from flows where the particles settle to the substrate, vs flows where the particles settle to the contact line, with markedly different behaviors in each regime. Moreover, the settling to the contact line appears to suppress the well-known fingering instability of the interface. I will present current state of art in modeling and simulation of this complex process, including comparison with recent physical experiments.

Andrea L. Bertozzi
UCLA Department of Mathematics
bertozzi@math.ucla.edu

MS34

Diffusion Equation for Heat Conduction and Radiation Transport in ALE-AMR

We discuss a diffusion equation solver that can operate in the ALE frame, handle AMR, and work with multiple materials in a single cell. This solver uses FEM with transition elements to handle the interface between coarse and fine elements. The degrees of freedom in the FEM solution are nodal and require a mapping from the cell-averaged quantities in ALE-AMR. Heat conduction and radiation transport capabilities are built on this finite element diffusion solution.

Aaron Fisher, Nathan D. Masters, David Bailey
Lawrence Livermore National Laboratory
fisher47@llnl.gov, masters6@llnl.gov, dsb@llnl.gov

Brian Gunney
Lawrence Livermore National Lab
gunney1@llnl.gov

Alice Koniges
Lawrence Berekely Laboratory
aekoniges@lbl.gov

David C. Eder, Robert Anderson, Velimir Mlaker
Lawrence Livermore National Laboratory
eder1@llnl.gov, rwa@llnl.gov, mlaker1@llnl.gov

MS34

Simple Method for Computing the Motion of Interfaces on Surfaces

Abstract unavailable at time of publication.

Steve Ruuth
Simon Fraser University
Department of Mathematics
sruuth@sfu.ca

MS34

Remapping - Conservative Interpolation

We present review of remapping methods for Arbitrary Lagrangian-Eulerian (ALE) methods. We first consider different approaches for remapping of cell-centered quantities. It includes methods based on exact intersection (overlays), swept region integration, flux-corrected remapping, and error compensation. Then we describe remapping of nodal quantities like velocity. Finally, we present remapping for multimaterial case.

Michail J. Shashkov
Los Alamos National Laboratory
shashkov@lanl.gov

MS35

Fast Calculation of Potential: The FMM on GPUs

Abstract unavailable at time of publication.

Ramani Duraiswami
University of Maryland
ramani@umiacs.umd.edu

MS35

Fast MR Image Formation with NuFFT on Multicore Processors

The low arithmetic complexity of nonuniform Fast Fourier Transform (NUFFT) algorithms and high latency in practical use of NUFFTs with large data sets, –especially in multi-dimensional domains–, pose a large problem in sequential and parallel implementations. With this work we introduce an effective approach to reducing the uneven latency distribution in memory references across different levels in the memory hierarchy by exploiting the geometric features of the samples. The restructured NUFFT algorithms demonstrate substantial performance gains in sequential as well as in parallel processing of radially encoded magnetic resonance image data.

Nikos Pitsianis
Department of Electrical & Computer Engineering
Duke University
nikos.p.pitsianis@duke.edu

Xiaobai Sun
Department of Computer Science
Duke University
xiaobai@cs.duke.edu

MS35

Parallelization of the Adaptive Fast Multipole Method on Multicore Architectures

We present a new parallelization scheme of the new version of Fast Multipole Method (FMM) based on the spatio-temporal attributes of the FMM interaction graph. The scheme was implemented on multicore architecture using POSIX threads to calculate the long range screened Coulomb interactions of N particles and compared with the sequential solver at http://www.fastmultipole.org/. The preliminary results suggested the scheme was efficient and made large scale simulation possible on workstations.

Bo Zhang
University of North Carolina at Chapel Hill
zhangb@email.unc.edu

MS36

Performance Evaluation of Software Framework for Memory Fault Tolerance in GPU Accelerators

Our software-only framework achieves DRAM reliability in commodity off-the-shelf GPUs by combining data coding and checkpointing. This presentation discusses performance aspects of the software approach by using several widely-used GPU kernels. We show that error checking can be effectively interleaved with original computation, resulting in negligible to moderate performance overheads.

Naoya Maruyama
Tokyo Institute of Technology
naoya@matsulab.is.titech.ac.jp

Akira Nukada, Satoshi Matsuoka
Tokyo Insitute of Technology
nukada@matsulab.is.titech.ac.jp, matsu@is.titech.ac.jp

MS36

Heterogeneous Many-core Computing Trends: Past, Present and Future

Heterogeneous computing is a recent trend that will have significant impact on the scientific software community. In this paper we explore the reasons behind this trend and analyse the effect this will have on scientific software development in the short, medium and long term. We also review the likely technology trends taking us towards Exascale systems in the 2020 timeframe with the goal of enabling long-term planning for scientific software design and development.

Simon N. Mcintosh-Smith
Department of Computer Science
University of Bristol
simonm@cs.bris.ac.uk

MS36

The MAGMA Matrix Algebra Library for Heterogeneous Multi- and Many-core Architectures

We present work on a new generation of linear algebra libraries (MAGMA) that achieve the fastest possible time to an accurate solution on heterogeneous systems, using all available processing power. The main focus is dense linear algebra for multicore+multi-GPU systems. MAGMA is designed to be similar to LAPACK in functionality, data storage, and interface, in order to allow scientists to effortlessly port any of their LAPACK-relying software components to take advantage of the new architectures.

Stanimire Tomov
Innovative Computing Laboratory, Computer Science Dept

University of Tennessee, Knoxville
tomov@cs.utk.edu

Jack J. Dongarra
Department of Computer Science
The University of Tennessee
dongarra@cs.utk.edu

MS36

The Open Computing Language (OpenCL) Standard for Programming Heterogeneous Many-core Systems

In this talk we shall provide a technical overview of OpenCL (Open Computing Language), the first open, royalty-free standard for general-purpose parallel programming of heterogeneous systems. OpenCL provides a uniform programming environment for software developers to write efficient, portable code for high-performance compute servers, desktop computer systems and handheld devices using a diverse mix of multi-core CPUs, GPUs, Cell-type architectures and other parallel processors such as DSPs.

Neil Trevett
OpenCL consortium
ntrevett@nvidia.com

MS37

Communication Bounds for Sequential and Parallel Eigenvalue Problems

We propose a set of divide-and-conquer algorithms for eigenvalues and eigenvectors that minimize communication (between levels of memory on a sequential computer and between processors on a parallel computer) in an asymptotic sense. The algorithms use as building blocks matrix multiplication and QR decomposition, each of which can be implemented by communication-optimal algorithms. The algorithms depend on a randomized rank-revealing QR decomposition to minimize the communication complexity while keeping the algorithms stable.

Grey Ballard, Jim Demmel
UC Berkeley
ballard@cs.berkeley.edu, demmel@eecs.berkeley.edu

Ioana Dumitriu
University of Washington, Seattle
dumitriu@math.washington.edu

MS37

Communication Avoiding Algorithms in Plasma and Magma

We will discuss some of the communication avoiding algorithms in Plasma and Magma.

Jack Dongarra
University of Tennessee Knoxville
dongarra@eecs.utk.edu

MS37

Communication-Avoiding Iterative Methods

Many iterative methods for solving sparse linear systems and eigenvalue problems have communication-bound performance. This makes them run much slower than machine peak on current computers, and architectural trends suggest this will only get worse over time. We present new iterative methods, equivalent to existing ones in exact arithmetic, that require asymptotically fewer messages in parallel and move asymptotically less data between levels of the memory hierarchy. The algorithms invoke new computational kernels, which replace the communication-bound dense and sparse kernels of current iterative methods. Even though the algorithms and kernels may perform redundant floating-point operations, performance models and benchmarks show significant speedups over existing algorithms. We also discuss numerical stability and preconditioning.

James Demmel
UC Berkeley, USA
demmel@eecs.berkeley.edu

Mark Hoemmen, Marghoob Mohiyuddin
UC Berkeley
mhoemmen@cs.berkeley.edu, marghoob@eecs.berkeley.edu

MS37

Partitioned Triangular Tridiagonalization of Symmetric Matrices

We present a partitioned algorithm for reducing a symmetric matrix to a tridiagonal form, with partial pivoting. That is, the algorithm computes a factorization $PAP^T = LTL^T$ where P is a permutation matrix, L is lower triangular with a unit diagonal and entries magnitudes bounded by 1, and T is symmetric and tridiagonal. The algorithm is based on the basic (non partitioned) methods of Parlett and Reid and of Aasen. We show that our factorization algorithm is component-wise backward stable (provided that the growth factor is not too large), with a similar behavior to that of Aasens basic algorithm. Our implementation also computes the QR factorization of T and solves linear systems of equations using the computed factorization. The partitioning allows our algorithm to exploit modern computer architectures (in particular, cache memories and high-performance BLAS libraries). Experimental results demonstrate that our algorithms achieve approximately the same level of performance as the partitioned Bunch-Kaufman factor and solve routines in LAPACK.

Gil Shklarski
Microsoft
gil.shklarski@microsoft.com

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

Miroslav Rozloznik
Czech Repub Academy of Science
Institute of Computer Science
miro@cs.cas.cz

MS38

FACETS Support for Parallel Coupled Core-Edge Fusion Simulations

Full-device modeling of large test fusion reactors such as ITER requires the self-consistent coupling of codes simulating different physical processes. The multi-institutional FACETS project is developing a multiphysics, parallel application that couples core and edge plasma simulations,

along with transport and wall interactions. We will discuss the design of FACETS software, including issues involved in developing concurrent component parallelism, and we will present preliminary results of coupled core-edge simulations.

John Cary
Center for Integrated Plasma Studies
University of Colorado at Boulder
cary@colorado.edu

Lois Curfman McInnes
Argonne National Laboratory
curfman@mcs.anl.gov

the FACETS team
Tech X Corporation
Argonne National Laboratory
cary@txcorp.com

MS38

MOOSE: A Parallel Solution Framework for Multiphysics Applications

MOOSE (Multiphysics Object Oriented Simulation Environment) is a software framework targeted at the development of implicit, tightly coupled multiphysics applications. It is designed to support the solution of finite element problems and provides an element library, input and output capabilities, mesh adaptation, and a set of parallel nonlinear solution methods. MOOSE is based on a physics-based preconditioned Jacobian-free Newton Krylov (JFNK) approach to support rapid application development and provide the robustness needed for engineering analysis and design. The JFNK abstraction results in a clean and extensible architecture for addressing a variety of multiphysics and multiscale problems.

Glen A. Hansen, Derek Gaston
Idaho National Laboratory
Glen.Hansen@inl.gov, derek.gaston@inl.gov

Chris Newman
Idaho National Laboratory
christopher.newman@inl.gov

Dana A. Knoll
Idaho National Laboratory
Dana.Knoll@inl.gov

MS38

Parallel Finite Element Analysis for Assembly Structures under Hierarchical Gridding

A finite element structural analysis software called FrontISTR is being developed. For realizing analyses of very large scale models with up to 10 billion elements, multilevel data structure is implemented and hierarchical mesh refinement is automatically done there. Mutigrid type convergence acceleration for the Krylov solvers is considered as well. To treat assembled structures, a new algorithm imposing multipoint constraint conditions without the penalization in the Krylov solvers has been developed.

Hiroshi Okuda
Research into Artifacts, Center for Engineering (RACE)
The University of Tokyo
okuda@race.u-tokyo.ac.jp

Takeshi Takeda
Institute of Industrial Science
The University of Tokyo
ktakeda@iis.u-tokyo.ac.jp

Kazuya Goto
dvanceSoft Corporation
goto@advancesoft.jp

MS38

Parallel CFD Applications and Its Performance on a Middleware

An Object-Oriented framework 'SPHERE' was designed to enhance the software development, to manage various physics simulators and to offer better usability for the applications in engineering and scientifice fields. This framework also delivers high-level conceptual parallel programming environment based on various parallelism of domain decomposition including AMR on octree, Building-Cube method, heierarchical nesting grid, and also a unique multibox division. In a presentation, the concept of SPHERE and the performance results will be demonstreated.

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

Satoshi Ito
RIKEN
sito@riken.jp

MS39

Hadoop Architecture and API and its Usage at Facebook

We talk about a new breed of applications that generate a lot of data. We describe how the Hadoop map-reduce framework and Hive can be used to process these large datasets in near-real-time fashion. Then we describe the type of applications that are using Hadoop in Facebook,the configuration of hardware and software of the Facebook Hadoop clusters, size and volume of datasets, characteristics of jobs and the processes built on top of Hadoop to keep the data pipeline alive and active.

Dhruba Borthakur
Facebook
dhruba@gmail.com

MS39

GPU Computing: Architecture and Programming Models for Massively Parallel Commodity Hardware

Modern GPUs easily qualify as the world's most successful parallel architecture. In this talk I will briefly motivate GPU computing and explore the transition it represents in massively parallel computing: from the domain of supercomputers to that of commodity "manycore" hardware available to all. I will discuss the goals, implications, and key abstractions of the CUDA architecture and programming model, and survey languages that run on CUDA such as C, Fortran, and OpenCL. Finally I will close with a discussion of future workloads in high-performance computing, games, and consumer applications, and their implica-

tions for future GPU architectures.

David Luebke, Jonathan Cohen
NVIDIA
dluebke@nvidia.com, jocohen@nvidia.com

MS39

Yada: A Deterministic Parallel Programming Language

To take advantage of new multicore chips, programmers need a much easier way of exploiting parallel processing than offered by the current dominant parallel processing paradigms: threads with shared memory or processes communicating by message-passing. To address this problem, we are designing Yada, a new parallel language featuring deterministic execution. Deterministic programs, i.e. programs whose behavior is identical from run-to-run, are much easier to develop, test and debug than the usual nondeterministic parallel programming paradigms.

David Gay
Intel Research
david.e.gay@intel.com

MS39

DryadLINQ: General-purpose Distributed Computing using a High-level Language

LINQ is a set of extensions to .NET languages that allows developers to manipulate and transform collections of data items. When actions on sets of items are explicitly described at a high level, the system can automatically exploit parallelism inherent in many algorithms. DryadLINQ automatically distributes a LINQ program over a cluster of computers, and I will describe our experiences implementing a variety of data-mining and machine learning algorithms using the system.

Michael Isard, Yuan Yu
Microsoft Research
misard@microsoft.com, yuanbyu@microsoft.com

MS40

Parallel Algorithms for Social Network Analysis

The explosion of real-world graph data poses a substantial challenge: How can we analyze constantly changing graphs with billions of vertices? Our approach leverages the Cray XMT's fine-grained parallelism and flat memory model to scale to massive graphs. On the Cray XMT, our static graph characterization package GraphCT summarizes such massive graphs, and our ongoing STINGER streaming work updates clustering coefficients on massive graphs at a rate of tens of thousands updates per second. This work is supported in part by Pacific Northwest National Laboratorys Center for Adaptive Supercomputing Software for Multithreaded Architectures.

David A. Bader
Georgia Institute of Technology
bader@cc.gatech.edu

MS40

Building Blocks for Scalable Graph and Data Mining Software

Software development for large scale graph and data mining applications is a formidable task that requires an enormous amount of human expertise. In contrast to numerical computing, a scalable software stack that eases the application programmers job does not exist for computations on graphs. We will describe the Parallel Combinatorial BLAS, which consist of a small but powerful set of linear algebra primitives specifically targeting graph and data mining applications. Our focus will be on the algorithmic techniques that make the primitives scalable, and software engineering techniques that make the combinatorial BLAS generic and extendible. Finally, we will talk about the performance and productivity of applications that are implemented using the combinatorial BLAS.

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

MS40

Comparing Programming Paradigms for Graph Algorithms

We describe the implementation of common graph algorithms such as Pagerank, Connected Components, and Single-Source Shortest Paths on distributed-memory architectures. Our implementations use the MR-MPI MapReduce library (our implementation of MapReduce using MPI), the Hadoop MapReduce framework, and the Trilinos Matrix-Vector library. We compare performance of the implementations for large power-law graphs on distributed memory clusters.

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

Steve Plimpton, Greg Bayer, Jonathan Berry
Sandia National Laboratories
sjplimp@sandia.gov, gbayer@sandia.gov, jberry@sandia.gov

MS40

A Model and a System for Large-Scale Graph Processing

We present a computational model and a system for processing large graphs. Programs are expressed as a sequence of iterations. In each, any vertex can receive messages sent to it in the previous iteration, send messages to other vertices, modify its own and its outgoing edges' states, and mutate the graph topology. The model is realized as a fault-tolerant and scalable system on clusters of commodity computers. The result is an expressive and easy to program framework for efficient processing of graphs.

Grzegorz Malewicz
Google
malewicz@google.com

MS41

Many Core Architectural Challenges for Parallel Computing of Fluids

Changes in computer architecture and languages are providing new challenges for traditional high-performance computing. In this talk we describe some of these changes and show how they currently and in the future will affect

many of the applications described in this series.

Alice Koniges
Lawrence Berekely Laboratory
aekoniges@lbl.gov

MS41

A Second Order Virtual Node Method for Poisson Interface Problems

I will present a second order accurate, symmetric positive definite, geometrically flexible and easy to implement method for solving the variable coefficient Poisson equation with interfacial discontinuities. We discretize the equations using an embedded approach employing virtual nodes at interfaces and boundaries. A variational method is used to define numerical stencils near these special virtual nodes and a Lagrange multiplier approach is used to enforce jump conditions and Dirichlet boundary conditions. Numerical experiments indicate second order accuracy in L-infinty.

Joseph Teran
University of California, Los Angles
joseph.teran@gmail.com

Jacob Bedrossian, James von Brecht, Siwei Zhu, Eftychios Sifakis
University of California, Los Angeles
jacob.bedrossian@math.ucla.edu, jub@math.ucla.edu, siwei@math.ucla.edu, sifakis@math.ucla.edu

MS41

Computation of Time-periodic Solutions of Interface Problems in Fluid Mechanics

I will describe a spectrally accurate numerical method for finding non-trivial time-periodic solutions of nonlinear PDE. We minimize a functional (of the initial condition and the period) that is positive unless the solution is periodic, in which case it is zero. We solve an adjoint PDE to compute the gradient of this functional with respect to the initial condition. We apply our method to the Benjamin-Ono equation and the vortex sheet with surface tension.

Jon Wilkening
UC Berkeley Mathematics
wilken@math.berkeley.edu

MS42

High Performance and High Scalable Eigenvalue Solver on a Peta-scale Computing Environment

Toward a peta-scale supercomputer system, we are developing a high performance and high scalable eigenvalue solver. The problem of numerical linear algebra on a multi-core processor is low memory bandwidth for each core. Moreover, an enough parallelism cannot be secured only by handing a single problem. We solve the first problem by introducing the block version of the Householder tridiagonalization and the divide-and-conquer method for band matrices. In addition, high parallelism is enabled by progressive implementation that we can use the solver as an asynchronous collective operation with multiple process groups. Preliminary results on a T2K open supercomputer system and a Nehalem core cluster will be presented.

Toshiyuki Imamura, Huu Phuong Pham
The University of Electro-Communications
imamura@im.uec.ac.jp, phuong@i.im.uec.ac.jp

Susumu Yamada, Masahiko Machida
Japan Atomic Energy Agency
yamada.susumu@jaea.go.jp, machida.masahiko@jaea.go.jp

MS42

Parallel Incomplete Factorization Preconditioning based on Algebraic Block Multi-Color Ordering Method

We introduce a technique for parallel processing of IC preconditioning for a linear system of equations having a random sparse coefficient matrix. We propose algebraic block multi-color ordering, which is an enhanced version of block multi-color ordering for a structured finite difference analysis. Blocking and coloring procedures for improving a cache hit ratio and convergence are presented. Three numerical tests are conducted for examining the proposed method, and it is shown that the computational time is successfully reduced to half of that of the conventional multi-color ordering method.

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

Yasuhito Takahashi
Graduate School of Informatics
Kyoto University
ytakahashi@i.kyoto-u.ac.jp

MS42

A High-level Programming Framework for Efficient Hybrid-architecture Computing

This work explores runtime auto-tuning methodologies enabling a high-level GPGPU programming environment which does not need to consider underlying hardware. To this end, there are two important things: generating an efficient code from a high-level code, and selecting a processor for each task. The optimal selection depends on runtime factors. Despite the importance of runtime processor selection, it is not considered even in state-of-the-art frameworks. Therefore this paper proposes a framework with history-based processor selection.

Kazuhiko Komatsu, Kentaro Koyama, Katsuto Sato, Hiroyuki Takizawa, Hiroaki Kobayashi
Tohoku University
komatsu@sc.isc.tohoku.ac.jp, kentaro@sc.isc.tohoku.ac.jp, katuto@sc.isc.tohoku.ac.jp, tacky@isc.tohoku.ac.jp, koba@isc.tohoku.ac.jp

MS42

Automatic Tuning for Parallel 3-D FFT with 2-D Decomposition

In this talk, an automatic performance tuning method for parallel 3-D fast Fourier transform (FFT) with 2-D decomposition is presented. An blocking algorithm for parallel FFTs utilizes cache memory effectively. Since the optimal block size may depend on the problem size, we propose a method to determine the optimal block size that minimizes the number of cache misses. Performance results of parallel

3-D FFTs on clusters of multi-core processors are reported.

Daisuke Takahashi
Graduate School of Systems and Information Engineering
University of Tsukuba
daisuke@cs.tsukuba.ac.jp

MS43

Extreme Scale Implicit PDE Simulations with Multigrid: Parallelization, Applications and Performance

Multigrid methods are introduced with discussion of their optimal complexity. Performance techniques for modern computers that optimized data locality are discussed in general and two examples of extreme scale applications are presented. Newton Krylov, algebraic multigrid solvers for unstructured nonlinear elasticity are presented with optimizations for extremes scale architectures, along with geometric nonlinear multigrid for problems in magnetohydrodynamics, are presented. We present performance results extreme scale architectures, include 32K cores of a Cray XT-5.

Mark Adams
Columbia University
mark.adams@columbia.edu

Ravi Samtaney
Princeton Plasma Physics Laboratory
samtaney@pppl.gov

Achi Brandt
The Weizmann Institute of Science
achi.brandt@weizmann.ac.il

MS43

Computational Challenges in Fast Reactor Core Simulations

The seven-dimensional neutron transport equation leads to extremely memory intensive and computationally intensive problems when applied to complex nuclear reactor geometries with large number of material compositions. We present a discrete-ordinate method for fast reactor core simulations that has scaled to large processor counts. The primary objective of these high-fidelity simulations is to reduce the uncertainty in important reactor design and operational parameters, facilitating the development of nuclear reactors that are safe and environment friendly.

Dinesh K. Kaushik
Argonne National Laboratory
2151, Bldg. 240, MCS Div
kaushik@mcs.anl.gov

Micheal Smith, Allan Wollaber
Argonne National Laboratory
masmith@ra.anl.gov, wollaber@mcs.anl.gov

Barry F. Smith
Argonne National Lab
MCS Division
bsmith@mcs.anl.gov

Won Sik Yang
Argonne National Laboratory
wyang@anl.gov

MS43

Scalable Massively Parallel Implicit Simulations of Fluid Flows to over 250,000 Processor-cores

Implicit methods for PDEs using unstructured meshes allow for an efficient solution strategy for many real-world problems (e.g., simulation-based virtual surgical planning). Scalable solvers employing these methods not only enable solution of extremely-large problems but also lead to dramatic compression in time-to-solution. We present a parallelization paradigm and associated procedures that enable our implicit, unstructured flow-solver to achieve strong scalability. We consider fluid-flow examples to demonstrate the effectiveness of our procedures that yield near-linear strong-scaling on various near-petascale systems (including over 250,000 processor-cores).

Onkar Sahni, Kenneth Jansen
Rensselaer Polytechnic Institute
Scientific Computation Research Center
osahni@scorec.rpi.edu, kjansen@scorec.rpi.edu

Min Zhou
Rensselaer Polytechnic Institute
zhoum@scorec.rpi.edu

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

MS43

Experiences in Scaling Multiphysics FLASH Applications on Largest Available Supercomputers

The FLASH code is a multiphysics AMR software with a wide user base. Applications using the FLASH code have been run on some of the largest available supercomputing platforms including UP and BG/L at LLNL and Intrepid at ANL. Here we present some of the challenges we have faced in running our simulations for supernova research, at scale, on some of these platforms.

Klaus Weide, Chris Daley
University of Chicago
klaus@flash.uchicago.edu, tba@chicago.edu

Dean Townsley
University of Arizona
tba@arizona.edu

Carlo Graziani, Anshu Dubey, Don Lamb
University of Chicago
tba@chicago.edu, tba@chicago.edu, tba@chicago.edu

MS43

Improved Unstructured Mesh Partitions for Parallel Simulations at Extreme Scale

Parallel simulation requires the mesh distributed with equal work load and minimum inter-part communications. Graph/hyper-graph based methods meet these goals. However, the global partition can fail on really large core counts, which is resolved by combination of global and local procedures. The usage of one type of mesh entities as graph nodes in graph/hyper-graph partition make the balance of other mesh entities not optimal. Two algorithms

are developed to provide improved partitions.

Min Zhou
Rensselaer Polytechnic Institute
zhoum@scorec.rpi.edu

Onkar Sahni
Rensselaer Polytechnic Institute
Scientific Computation Research Center
osahni@scorec.rpi.edu

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

Mark S. Shephard, Kenneth Jansen
Rensselaer Polytechnic Institute
Scientific Computation Research Center
shephard@scorec.rpi.edu, kjansen@scorec.rpi.edu

MS44

Epidemics on social contact networks: impact of structural properties and designing effective interventions

The structure of the underlying network is known to have a significant impact on the spread of epidemics. This has motivated interventions based on the network structure, and in this talk, we study the effectiveness of such strategies. We introduce a new measure called the vulnerability. In addition to being a fundamental quantity related to disease propagation, we find it leads to more effective interventions than those based on static graph measures.

Madhav Marathe
Virginia Bioinformatics Institute, VA Tech
mmarathe@vbi.vt.edu

Anil Vullikanti
Dept. of Computer Science, and Virginia Bioinformatics Inst.
Virginia Tech
akumar@vbi.vt.edu

Jiangzhuo Chen
Virginia Tech
Virginia Bioinformatics Institute
chenj@vbi.vt.edu

MS44

Networks, Communities and Kronecker Products

One of the principal challenges in analysis of large complex network is to build models and understanding of the structure of such networks. I will present our work on clustering and community structure in large networks, where clusters are thought of as sets of nodes that are better connected internally than to the rest of the network. We find that large networks have very different clustering structure from well studied small social networks and graphs that are well-embeddable in a low-dimensional structure.

Jure Leskovec
Stanford University
jure@cs.stanford.edu

MS44

High-Performance Simulations for Capturing Feedback and Fidelity in Complex Networked Systems

In a variety of complex networked systems, simulation is a powerful method to capture critical feedback effects among inter-dependent processes. Network-based phenomena in areas such as cyberinfrastructure, transportation, epidemiology, and social networks, all offer important analysis problems that need such feedback effects to be accurately captured. However, accurate modeling of feedback effects requires increased levels of model fidelity. Moreover, such high-fidelity, feedback-heavy models are especially characterized by very high computational needs. In this backdrop, the need for high-fidelity simulations is illustrated, with examples of how they are driving new high-performance computing-based solutions in the aforementioned areas. Our parallel computing approaches are described in the context of very large-scale, high-fidelity simulations in regional-scale transportation network simulations, nation-scale epidemiological simulations, and Internet simulations with detailed models of millions of nodes.

Kalyan Perumalla
Oak Ridge National Laboratory
perumallaks@ornl.gov

MS44

Dynamical Systems Models of Gene Regulatory Networks

Regulatory networks of biomolecular interactions in cells govern virtually all cellular behaviors and functions. Models of these dynamical systems can be used to understand biological function and enable rational control strategies for therapeutic intervention in disease. Emergent systems-level behavior can be studied to gain insights into cellular decision making, robustness, and ability to coordinate complex behavior. Information theoretic approaches combined with elements of dynamical systems theory, such as phase transitions and structure dynamics relationships, are promising frameworks for studying these fundamental principles of living systems.

Ilya Shmulevich
University of Texas M.D. Anderson Cancer Center
is@ieee.org

MS45

Analysis of Hybrid Applications on Modern Architectures

Commodity clusters as well as custom-built supercomputers are today exclusively of hierarchical type, meaning that more than one parallel programming model may be appropriate to use, at the same time, within an application. Starting from a taxonomy of options for MPI+OpenMP hybrid programming, which is still the most attractive option for the application programmer, we will analyze the performance of hybrid codes on modern parallel systems. Benchmarks and application programs from fluid dynamics and condensed matter physics will exemplify which programming models are appropriate for practical use. We conclude with an outlook to advanced hybrid architectures beyond "commodity' multi-core/multi-socket clustering.

Gerhard Wellein, Georg Hager
Erlangen Regional Computing Center (RRZE), Germany

Gerhard.Wellein@rrze.uni-erlangen.de,
georg.hager@rrze.uni-erlangen.de

MS45

Experiences with Hybrid MPI/OpenMPI Parallelization

Today most large scale parallel systems are clusters of multi-core nodes where distributed as well as shared memory programming can be applied. We have parallelized a 3D flow solver using MPI for coarse grain parallelism based on a 1D domain decomposition. Within each sub-domain we exploit additional fine grain parallelism on the undistributed dimensions by employing OpenMP. We describe how this approach can increase performance and scalability of the application.

Gabriele Jost
Texas Advanced Computing Center
The University of Texas at Austin
gjost@tacc.utexas.edu

Bob Robins
Northwest Research Associates, Inc
bob@nwra.com

MS45

A Closer Look at the Parallel Pattern Library, Asynchronous Agents Library and Concurrency Runtime in Visual Studio 2010

The aim of this session is to provide deeper insight into the design and motivation for Microsofts Concurrency Runtime and Parallel Pattern Library for building both compute intensive programs and more general purpose applications, where appropriate similarities and differences with existing programming models like OpenMP will be explained.

Rick Molloy
Microsoft
rick.molloy@microsoft.com

MS45

OpenMP: Beyond Specification 3.0

OpenMP has a long history of supporting loop level, fork-join style parallelism and coarse grain task parallelism through its sections construct. The OpenMP 3.0 specification added support for fine grained task parallelism. However, many challenges remain and the OpenMP Language Committee is actively developing the next set of extensions in 3.1 and 4.0. This talk will detail these extensions, including user defined reductions, further tasking refinements, a detailed error model and support for accelerators.

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

MS46

Parallel Dense Polynomial Arithmetic on Multi-Cores

We aim at multicore-enabling FFT-based dense polynomial arithmetic over finite fields. We show that balanced input data can maximize parallel speedup and minimize cache complexity for bivariate multiplication. We present effective techniques to reduce multivariate (and univariate) multiplication to balanced bivariate multiplication. This parallel multiplication also provides an efficient kernel for fast division and normal form computation so that composition of multi-level of parallelism succeeds. Our implementation in Cilk++ demonstrates good speedup on multicores.

Marc Moreno Maza
Computer Science Department
University of Western Ontario
moreno@csd.uwo.ca

Yuzhen Xie
Massachusetts Institute of Technology
yxie@csail.mit.edu

MS46

Fast Algorithms for Real Solving Polynomial Systems of Inequalities/inequations

Let $(f_1, \ldots, f_s) \subset \mathbf{Q}[\mathbf{X_1}, \ldots, \mathbf{X_n}]$ and $D \geq \deg(f_i)$. We consider the problem of deciding the emptiness of the semi-algebraic set defined by $f_i \, \sigma_i \, 0 \, , 1 \leq i \leq s$ with $\sigma_i \in \{>, <, \neq\}$. We present an algorithm solving this problem. We first show that its complexity is in the state of the art (i.e. $s^{n+1} D^{O(n)}$). We find that its implementation can tackle important and challenging problems which are out of reach of standard algorithms/implementations, even if they are in the same complexity class.

Mohab Safey El Din
UPMC, Univ Paris 6
Mohab.Safey@lip6.fr

MS46

Using Plain, Middle and Short Products to Speed-up Newton Iteration

We discuss various forms of Newton iteration, for computing power series solutions of differential or polynomial equations. We show how to apply fast multiplication techniques such as short product or middle product in a systematic manner, generalizing some previous approaches known for e.g. power series inverse or square root.

Eric Schost
Computer Science Department
Univ. of Western Ontario
eschost@uwo.ca

Ling Ding
The University of Western Ontario
lding6@gmail.com

MS46

Polynomial Homotopies on Multicore Workstations

Homotopy continuation methods to solve polynomial systems scale very well on parallel machines. In this talk we examine its parallel implementation on multiprocessor multicore workstations, using threads. Preliminary timings indicate good speedups for basic pleasingly parallel path tracking jobs. The use of multitasking will lead to more efficient parallel implementations in a multi-tiered approach. Developing parallel code for complicated homotopies on multicore workstation is more convenient with

threads than with MPI.

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

MS47

Performance Evaluation of Preconditioned Krylov Subspace Methods on GPU

Many preconditioning algorithms for Krylov subspace methods are developed and well tested on parallel computing environment like PC clusters, but performance of those methods on GPU differs from those on clusters. By using C for CUDA environment, we have implemented and evaluated various kinds of preconditioning methods for Krylov subspace solvers on GPU, including SOR variants, Multigrid, Line-by-Line and incomplete Cholesky decomposition.

Hidetoshi Ando
University of Yamanashi
ando@yamanashi.ac.jp

MS47

Challenges and Progress in Effective Use of Scalable Multicore and Manycore Systems

Multicore and manycore nodes have been available for some time now and our understanding of how to best design and implement algorithms for these nodes has grown. In this presentation we present some of the lessons we have learned and describe the current challenges we face.

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

MS47

A Study on OpenMP/MPI Hybrid Programming Models for Density-functional Calculations

Real space density function theory (RSDFT) is used to investigate the electronic structure of many-body systems. In order to find a better approach to exploit the hybrid of OpenMP and MPI on large-scale multicore clusters, the hybrid-parallelism of the RSDFT is evaluated and explored. We show the characteristics of the OpenMP/MPI hybrid programming models in RSDFT using the PAPI performance counter interface.

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

Mitsuhisa Sato
Center for Computational Sciences
University of Tsukuba, Japan
msato@cs.tsukuba.ac.jp

MS47

Multicore and Manycore Performance Studies using a Finite Element Miniapplications

Miniapplications–small, portable performance proxies for large-scale applications–provide an effective way of rapidly co-designing software and hardware. MiniFE is finite element miniapplication with features that support the use of mixed MPI, threads and CUDA. MiniFE also provides several different ways of supporting sparse collective communications and allows tunable work and load imbalance for testing the performance of scalable networks. In this presentation we show results using MiniFE on a variety of system platforms and draw conclusions from these results regarding the best approaches to effectively use multicore and manycore systems.

Alan B. Williams
Sandia National Laboratory
Distributed Systems Research Department
william@sandia.gov

MS48

Analyzing and Exploring Massive Scientific Data Using Topology

Abstract unavailable at time of publication.

Timo Bremer
LLNL
ptbremer@llnl.gov

MS48

The Challenges Ahead for Visualizing and Analyzing Massive Data Sets

Abstract unavailable at time of publication.

Hank Childs
Lawrence Berkeley National Laboratory
hchilds@lbl.gov

MS48

Moving Analysis to the Data: Scalable Visualization Using Simulation Resources

Increasingly, the cost of data movement dominates analysis and simulation of large-scale scientific data. This is especially true of accessing storage. The pressure on I/O can be alleviated by bringing more numerical and visual analysis closer to the data and executing these operations directly on supercomputing resources. We will examine using these machines for parallel visualization, in postprocessing and in situ, and scaling parallel analysis algorithms with simulations.

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

MS48

Scaling Up Algorithms: Contouring, Volume Rendering, and Streamlines

Abstract unavailable at time of publication.

David Pugmire
Oak Ridge National Laboratory
pugmire@ornl.gov

MS49

Development of a Parallel Energy Discretization

for Eigenvalue Neutron Transport

We will detail our effort on extending an existing parallel transport solver called Denovo, to develop a first-of-a-kind, mathematically consistent, two-level approach to the multiscale challenge. Denovo is a significant advance over current technology, because it allows fully consistent multi-step approaches to high-fidelity nuclear reactor simulations that cannot be performed with current technology. New parallel algorithms for the energy dimension will also be presented.

Kevin Clarno, Tom Evans
Oak Ridge National Laboratory
clarnokt@ornl.gov, evanstm@ornl.gov

MS49

Parallel Algorithms for Particle Transport with the Method of Characteristics

We discuss theoretical and numerical results from our recent work in long-characteristic methods for particle transport problems. We present methods with very small spatial and temporal discretization errors for problem types ranging from streaming to diffusive. We discuss new massively parallel algorithms that are enabled by long-characteristic methods and offer analyses showing that these algorithms will outperform existing parallel algorithms, especially for large processor counts.

W Hawkins, T Pandya, M Adams
Texas A&M UNiversity
dhawkins@tamu.edu, tarapandya@gmail.com, mladams@tamu.edu

MS49

Algorithmic Scalability for the Boltzmann Transport Equation

This presentation will describe a new multigrid method applied to the most common Sn discretizations (Petrov-Galerkin, diamond-differenced, corner-balanced, and discontinuous Galerkin) of the mono-energetic Boltzmann transport equation in the optically thick and thin regimes, and with strong anisotropic scattering. Unlike methods that use scalar DSA diffusion preconditioners for the source iteration, this multigrid method is applied directly to an integral equation for the scalar flux.

Barry Lee
Pacific Northwest National Lab
barry.lee@pnl.gov

MS49

A Scalable Solver for the Even-Parity Form of Neutron Transport Equation

As part of UNIC code at Argonne, a scalable solution methodology for the discrete Ordinates, even-parity form of the neutron transport equation is being developed. For high-fidelity descriptions of complex reactor geometries (respecting spatial heterogeneities and large number of energy groups), the memory requirements are huge. We will present a novel approach (using p-multigrid) to contain the memory requirement in the context of Zero Power Reactor (ZPR) Experiment 6/6a simulations on Blue Gene/P and XT5.

Allan Wollaber
Argonne National Laboratory
awollaber@anl.gov

Dinesh K. Kaushik
Argonne National Laboratory
2151, Bldg. 240, MCS Div
kaushik@mcs.anl.gov

Micheal Smith
Argonne National Laboratory
masmith@anl.gov

Barry F. Smith
Argonne National Lab
MCS Division
bsmith@mcs.anl.gov

Won Sik Yang
Argonne National Laboratory
wyang@anl.gov

MS50

Addressing Software Complexity in a Multiphysics Parallel Application: Coupled Core-Edge-Wall Fusion Simulations

Coupled core-edge-wall fusion simulations must be run in a tightly coupled fashion because of the need for in-memory coupling for speed and to accommodate the complex communication patterns. Larger time steps require implicit and versatile couplings, e.g., by fields at different points and/or fluxes. Finally, one must accommodate the potential absence of particular components and multiple implementations of certain functionalities. This talk discusses the software patterns in the FACETS project that meet these different requirements.

John Cary
Tech-X Corporation
University of Colorado
cary@txcorp.com

Lois Curfman McInnes
Argonne National Laboratory
Mathematics and Computer Science Division
mcinnes@mcs.anl.gov

Tom Epperly
Lawrence Livermore National Lab
epperly2@llnl.gov

Ammar Hakim, Scott Kruger, Mahmood Miah, Alex Pletzer, Sveta Shasharina
Tech-X Corporation
ammar@txcorp.com, kruger@txcorp.com, mmiah@txcorp.com, pletzer@txcorp.com, sveta@txcorp.com

MS50

Automating the Self-Assembly of Multiphysics Simulators from Components

There are now many high-quality component libraries which provide building blocks for simulations, however, integrating and coordinating these components to form a complex multi-physics simulator is still usually done by human programmers. We show how the mathematical information encoded in a symbolic problem specification may

be used to automate the self-assembly of a simulator from a suite of components.

Kevin Long
Texas Tech University
kevin.long@ttu.edu

MS50

A Multiphysics Assembly Engine for Advanced Analysis Techniques

Advanced solution and analysis techniques such as Newtons method, optimization, bifurcation analysis and uncertainty quantification are typically invasive and require auxiliary information such as sensitivities. This creates a barrier for users to add new physics without learning and implementing the additional analysis algorithm requirements. This talk will discuss the design concepts and software that decouples the equation description from the solution/analysis algorithms through the use of embedded technology. We will show results from large-scale applications including fluid flow, magnetohydrodynamics, and semiconductor device modeling. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Roger Pawlowski, Patrick Notz
Sandia National Labs
rppawlo@sandia.gov, pknotz@sandia.gov

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

MS50

Automated Algorithm Generation and Thread Paralellism in Complex Scientific Computing Applications

Multiphysics simulation software is plagued by complexity stemming from nonlinear coupling. In addition to nonlinear coupling, such software often must support many models, each of which may require a different set of transport equations, constitutive models, and equations of state. Strong coupling, together with a multiplicity of models, leads to complex algorithms and rigid software. The rigid, complex software is due to design that focuses on algorithms. In this talk, we will discuss an alternative programming paradigm where programmers focus not on the algorithm, but on the data. Mathematical expressions are reflected in software in a way that directly exposes data dependencies, and graph theory is employed to automatically generate an algorithm. This allows programmers to avoid the complexity of ordering operations, and allows problems with very complex dependencies to become entirely tractable, and removes virtually all logic from the algorithm itself. Changes are highly localized, allowing model developers to implement code without requiring a detailed understanding of any algorithms. It provides a natural framework to achieve model adaptivity, where one model suite may be substituted for another dynamically, allowing for addition or removal of transport equations, modifying constitutive models, etc. and allows us to automatically generate a new algorithm. Unlike traditional programming models, this approach naturally handles complexity associated with multiple modeling options (which may imply different nonlinear coupling and may require solution of different sets of PDEs, equations of state, or constitutive relationships). Furthermore, this approach enables efficient algorithmic parallelization via threads. By exposing dependencies in the algorithm explicitly, thread-based parallelism can be implemented through algorithm decomposition, thus providing a basis for exploiting parallelism that is independent from and complimentary to domain decomposition approaches.

James C. Sutherland
University of Utah
James.Sutherland@utah.edu

MS51

Exploiting Geometry in Tree-based Hybrid Preconditioners

Abstract unavailable at time of publication.

Padma Raghavan
The Pennsylvania State Univ.
Dept of Computer Science Engr.
raghavan@cse.psu.edu

MS51

Mapping Parallel Applications on Volatile Resources

We have an iterative application made of parallel tasks. We execute this application on a volatile platform, whose processors obey independent but different UP/DOWN Markov processes (either fail-and-reboot scenarios, or preemption during cycle-stealing episodes). We investigate the complexity of the off-line problem (combinatorial instances, reduction to bi-clique problems), and we report the behavior of various resource selection and migration algorithms through simulations for the on-line setting.

Henri Casanova
University of Hawai'i
henric@hawaii.edu

Fanny Dufossé
LIP, ENS Lyon, France
fanny.dufosse@ens-lyon.fr

Yves Robert
ENS LYon, France
& Institut Universitaire de France
yves.robert@ens-lyon.fr

Frédéric Vivien
LIP, ENS Lyon, France
frederic.vivien@ens-lyon.fr

MS51

The Effect of Overlapping Partitionings for Massively-parallel PDE-constrained Optimization

The availability of large-scale computing platforms comprised of thousands of multicore processors motivates the need for the next generation of highly scalable sparse linear system solvers. These solvers must optimize parallel performance, processor performance, as well as memory requirements, while being robust across broad classes of applications and systems. In this talk, we discuss the effect of overlapping partitionings that can be used to build par-

allel solvers with applications in massively-parallel PDE-constrained optimization.

Olaf Schenk
Universität Basel
olaf.schenk@unibas.ch

MS51
Partitioning Regular Meshes for Minimizing the Total Communication Volume

We investigate one-dimensional partitioning of sparse matrices that arise after the discretization of square domains with the five-point stencil. The objective is to minimize the total communication volume and to have balance among the processors, when the partition is used to parallelize sparse matrix-vector multiplies. The problem is known to be equivalent to the NP-complete hypergraph partitioning problem.We propose a geometry-based partitioning method which obtains consistently better results than a hypergraph partitioning-based method.

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

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

MS52
Towards a High-Level Unifying Programming Model for Heterogeneous Systems

High-end computing systems are recently integrated with heterogeneous multicores and/or programmable hardware accelerators. OpenCL and CUDA are current popular programming paradigms, which, however, are time consuming to program, hard to debug and tune, and high maintenance. As an alternative, we discuss extensions to the popular OpenMP model targeted towards heterogeneous systems. We present language extensions designed to offload tasks and handle communication. We also describe an implementation based on the OpenUH compiler.

Lei Huang, Barbara Chapman
University of Houston
Lei.Huang@mail.uh.edu, bmchapman@earthlink.net

MS52
Experiences with Co-Array Fortran

This talk will explore the advantages of using Co-Arrays when scaling applications to larger and larger processor counts. By using CAF put and gets, the user will have the power to overlap numerous communication calls with each other as well as computation at a level never possible with MPI. This capability will allow applications like adaptive mesh refinement (AMR) with numerous gather/scatters to scale beyond current limits.

John M. Levesque
Cray
levesque@cray.com

MS52
On Hybrid Programming

MPI is a standard way of moving data among address spaces. As the memory available for separate address spaces per core drops on future machines, programmers need a mechanism for expressing parallelism within a single address space that is compatible with MPI. OpenMP is a current approach, but PGAS languages, notably UPC, offer intriguing alternatives. We explore the motivation for the MPI + UPC model and what it might look like in practice.

Ewing (Rusty) Lusk
Argonne National Laboratory
lusk@mcs.anl.gov

William Gropp
University of Illinois
groppp@uiuc.edu

James Dinan
Ohio State University
dinan@cse.ohio-state.edu

Rajeev Thakur, Pavan Balaji
Argonne National Laboratory
thakur@mcs.anl.gov, balaji@mcs.anl.gov

P Saddayapan
Ohio State University
saday@cse.ohio-state.edu

MS52
Getting Multicore Performance with UPC

Unified Parallel C (UPC) enables users to express data locality for parallel programs, which is very important for achieving scalable performance on NUMA multicore architectures. The Berkeley UPC implementation has end-to-end optimizations in compiler, runtime library and operating system support for multicore computers. Programming techniques and software tools for optimizing UPC application performance on multicore systems are described including examples with different tradeoffs in performance and productivity.

Yili Zheng, Filip Blagojevic
Lawrence Berkeley National Laboratory
yili.zheng@gmail.com, fblagojevic@lbl.gov

Dan Bonachea
University of California at Berkeley
bonachea@cs.berkeley.edu

Paul Hargrove, Steven Hofmeyr, Costin Iancu, Seung-Jai Min
Lawrence Berkeley National Laboratory
phhargrove@lbl.gov, shofmeyr@lbl.gov, cciancu@lbl.gov, sjmin@lbl.gov

Rajesh Nishtala
UC Berkeley
rajeshn@cs.berkeley.edu

Katherine Yelick
Lawrence Berkeley National Laboratory
University of California at Berkeley

yelick@eecs.berkeley.edu

MS53

Direct Numerical Simulation of Red Blood Cells in Shear Flow

Simple constitutive flow models inadequately describe the rheology of blood and cannot show the microscale cellular structures. Computational simulation of blood flow at a cellular level is only possible with massively parallel processing systems in conjunction with a highly scalable computational approach such as the hybrid lattice-Boltzmann/finite element (LB/FE) method to be presented. The scalability results presented in this study are based on the *Blue Gene/P* (BGP) Intrepid configuration at Argonne National Lab (ANL).

Cyrus Aidun
G.W. Woodruff School of Mechanical Engineering
Georgia Institute of Technology
cyrus.aidun@me.gatech.edu

Jonathan Clausen, Daniel Reasor
Georgia Inst Technology
clausen@gatech.edu, daniel.reasor@gatech.edu

MS53

Parallel Algorithms for Fluid-Structure Interaction and Applications to Cardiovascular Flows

In fluid-structure interaction for haemodynamics applications, parallelization and scalability are key issues. In this talk we introduce a new class of parallel preconditioners for the FSI problem obtained by exploiting the block-structure of the linear system. We provide a bound in the condition number of the preconditioned system in terms of the preconditioned blocks. The class is then tested on benchmarks and physiological 3D geometries; we report some comparisons and scalability results.

Simone Deparis
Ecole Politechnique Federale de Lausanne
simone.deparis@epfl.ch

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

Paolo Crosetto, Gilles Fourestey
Ecole Politechnique Federale de Lausanne
paolo.crosetto@epfl.ch, gilles.fourestey@epfl.ch

MS53

Medical Image Registration on TeraGrid: New Potentials for Improved Accuracy of Clinical Decisions

Spatial alignment, or registration, of medical images is a crucial step while analyzing normal changes in the anatomy, diagnosing disease and studying progression of the pathology. Existing image registration algorithms do not provide bounds on the registration error. Deviations of the derived transformation from the true solution may affect the analysis accuracy. Current approaches to registration in 3D Slicer will be presented together with the preliminary results of integrating them to leverage the Grid resources.

Andriy Fedorov
Brigham and Women's Hospital
Harvard Medical School
fedorov@bwh.harvard.edu

Nikos P. Chrisochoides
Computer Science
College of William and Mary
nikos@cs.wm.edu

Ron Kikinis
Brigham & Women's Hospital
Department of Radiology
kikinis@bwh.harvard.edu

MS53

A Multilevel Newton-Krylov-Schwarz Method for the Bidomain Model of Electrocardiology

A multilevel Newton-Krylov-Schwarz (NKS) solver is constructed and analyzed for implicit time discretizations of the Bidomain reaction-diffusion system. This model describes the bioelectrical activity of the heart by coupling two degenerate parabolic equations with several ordinary differential equations. The proposed NKS Bidomain solver employs an outer Newton iteration to solve the nonlinear finite element system originating at each time step of the implicit discretization. The Jacobian update during the Newton iteration is solved by a Krylov method employing a multilevel overlapping Schwarz preconditioner. A convergence rate estimate is proved for the resulting preconditioned operator, showing that its condition number is independent of the number of subdomains (scalability) and bounded by the ratio of the subdomains characteristic size and the overlap size. This theoretical result is confirmed by several parallel simulations employing up to more than 2000 processors for scaled and standard speedup numerical tests in three dimensions.

Simone Scacchi
Universita' degli Studi di Milano
simone.scacchi@unimi.it

Marilena Munteanu, Luca F. Pavarino
Dipartimento di Matematica
Universita' degli Studi di Milano
marilena_munteanu@hotmail.com, luca.pavarino@unimi.it

MS54

Iterative Methods for Constrained Optimization Problems in Multibody Dynamics Simulation

We present an iterative method for the resolution of the inequality and complementarity constrained sub problems that appear in time stepping methods for rigid multi-body dynamics with contact and friction. Such methods use a hard constraint approach by means of an iterative projection algorithm. Our method builds on recent results of the authors that show that an accurate solution of the constrained dynamics can be obtained while solving convex optimization sub problems.

Mihai Anitescu
Argonne National Laboratory
Mathematics and Computer Science Division
anitescu@mcs.anl.gov

Alessandro Tasora
Department of Mechanical Engineering
University of Parma, Italy
tasora@ied.unipr.it

Dan Negrut
University of Wisconsin Madison
Department of Mechanical Engineering
negrut@engr.wisc.edu

MS54

A Large Scale Rigid Body Dynamics Algorithm

In this talk we present a novel approach for large scale rigid body dynamics simulations. The presented algorithm enables rigid body simulations of more than one billion interacting rigid bodies. We describe in detail the parallel rigid body algorithm and its necessary extensions for a large scale MPI parallelization and analyze the parallel algorithm by means of a particular simulation scenario.

Klaus Iglberger, Ulrich Rüde
Computer Science 10 - System Simulation
Friedrich-Alexander-University Erlangen-Nuremberg, Germany
klaus.iglberger@informatik.uni-erlangen.de, ulrich.ruede@informatik.uni-erlangen.de

MS54

A Scalable Parallel Method for Large Scale Collision Detection Problems

The presentation outlines a scalable collision detection algorithm on the GPU that demonstrates a fortyfold speedup over state-of-the art sequential implementations when handling multi-million object collision detection tasks. The algorithm can be used to detect collisions between five million objects in less than two seconds and scales to handle problems with more than a billion contacts. The proposed methodology is expected to positively impact a range of granular flow dynamics and smoothed particle hydrodynamics applications.

Dan Negrut, Hammad Mazhar, Toby Heyn
University of Wisconsin Madison
Department of Mechanical Engineering
negrut@engr.wisc.edu, hmazhar@wisc.edu, heyn@wisc.edu

MS54

Real-time Simulation of Complex Mechanical Systems by Means of GPU Parallel Computation

The simulation of large numbers of interacting mechanical parts requires high computational efforts, especially when unilateral contacts and friction come into play. In some cases, such as in driving simulators, videogames, augmented reality or virtual reality, the real-time constraint requires very fast and robust algorithms that, within the few milliseconds of each simulation frame, must perform collision detection, solve the complementarity problem caused by frictional contacts and perform the time integration. When using traditional solution schemes, based on serial algorithms, the upper limit in the number of bodies can be quite limited, hence the real-time simulation of complex scenarios such as tracked vehicles running over debris or rocks is impossible. Thus, by leveraging the new computational paradigm provided by GPU parallel architectures, we developed a simulation library that can perform simulation of mechanical systems with large amounts of bilateral and unilateral constraints. Aiming at real-time performance, our iterative complementarity solver supports premature termination, robust stabilization, matrix-less O(n) data structures and linear-time overhead per iteration. Custom optimizations and improvements are discussed.

Alessandro Tasora
Department of Mechanical Engineering
University of Parma, Italy
tasora@ied.unipr.it

Dan Negrut
University of Wisconsin Madison
Department of Mechanical Engineering
negrut@engr.wisc.edu

MS55

Tile QR Factorization with Parallel Panel Processing for Multicore Architectures

We present a new fully asynchronous method for computing a QR factorization on shared-memory multicore architectures. Our contribution is to adapt an existing algorithm performing a panel factorization in parallel (Communication-Avoiding QR) to the context of tile algorithms using asynchronous computations. An experimental study shows a significant improvement (up to almost 10 times faster on tall and skinny matrices) compared to state-of-the-art approaches. We aim to eventually incorporate this work into the PLASMA library.

Emmanuel Agullo
Department of Electrical Engineering and Computer Science
University of Tennessee, Knoxville
eagullo@eecs.utk.edu

Jack Dongarra
University of Tennessee Knoxville
dongarra@eecs.utk.edu

Bilel Hadri, Hatem Ltaief
Department of Electrical Engineering and Computer Science
University of Tennessee, Knoxville
hadri@eecs.utk.edu, ltaief@eecs.utk.edu

MS55

GPGPU Programming to Solve the Boltzman Neutron Transport Equation

In the frame of the Neutron Physics code development, we study the use of GPGPU accelerators. We take advantage of the higher computing power and memory bandwidth available on GPUs to speed-up computations of a neutron transport equation solver, MINOS. We will discuss the programming models and algorithms to accelerate the MINOS solver using GPGPU and heterogeneous computing using domain decomposition method mixed with GPGPU acceleration on each sub-domain.

Christophe Calvin, Erell Jamelot, Jérôme Dubois
CEA-Saclay/DEN/DANS/DM2S/SERMA/LLPR
christophe.calvin@cea.fr, erell.jamelot@cea.fr, jerome.dubois@cea.fr

Serge Petiton
LIFL / MAP
serge.petiton@lifl.fr

MS55

Parallel Multigrid Solvers using OpenMP/MPI Hybrid Programming Models on Multi-Core/Multi-Socket Clusters

OpenMP/MPI hybrid parallel programming models were implemented to 3D finite-volume based simulation code for groundwater flow problems through heterogeneous porous media using parallel iterative solvers with multigrid preconditioning. Performance and robustness of the developed code has been evaluated on the T2K Open Supercomputer and Cray-XT4 using up to 1,024 cores. Detailed discussion on optimization and performance evaluation of OpenMP/MPI hybrid parallel programming models will be presented.

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

MS55

Toward a Multi-Target Linear Algebra Library for GPUs and Multicore CPUs.

Highly structured sparse matrices arise frequently from numerical discretizations of partial differential equations. Legolas++ is a C++ generic library designed for describing and manipulating such multi-level blocked matrices with the corresponding blocked vectors and algorithms. Legolas++ allows a very detailed description of the linear systems to be solved that can be used to generate efficient parallel implementations. We are working on enlarging the scope of Legolas++ from multi-core target to GPUs and cluster of GPUs.

Laurent Plagne, Wilfried Kirschenmann
EDF
laurent.plagne@edf.fr, wilfried.kirschenmann@edf.fr

Stéphane Vialle
ALGORILLE (INRIA Lorraine - LORIA)
stephane.vialle@supelec.fr

MS56

Scalable Preconditioning for Two-fluid Extended MHD Systems

Recently, a scalable preconditioning approach for 3D MHD has been proposed [L. Chacón, *Phys. Plasmas*, 2008]. The method is based on a conceptual Schur-complement decomposition, which exploits the nature of the hyperbolic couplings in MHD to produce a block diagonally dominant PDE system, amenable to classical multilevel techniques. In this talk, we describe the extension of the method to two-fluid extended MHD systems, which support fast dispersive waves, and thus are significantly stiffer numerically.

Luis Chacon
Oak Ridge National Laboratory
chaconl@ornl.gov

MS56

A Mixed Finite Element Method for Incompressible Magnetohydrodynamics

We present a mixed finite element method for nonlinear incompressible MHD problems. Our method is proved *inf-sup* stable and can correctly capture strong singularities in non-convex domains. We carry out an a-priori error analysis for the proposed mixed method, and develop a block preconditioning approach that takes into consideration the coupling between the magnetic and the fluid flow terms. A series of numerical convergence tests on MHD benchmark problems are presented to highlight the practical performance of our method.

Chen Greif
Department of Computer Science
The University of British Columbia
greif@cs.ubc.ca

Dan Li
The University of British Columbia
danli@cs.ubc.ca

Dominik Schoetzau
Mathematics Department
University of British Columbia
schoetzau@math.ubc.ca

Xiaoxi Wei
The University of British Columbia
weixiaoxi@math.ubc.ca

MS56

Physics-based Preconditioning of Jacobian-Free Newton-Krylov Methods for Problems in Energy and Climate

Many partial differential equation systems modeling problems in energy systems or climate simulation can be categorized as multiphysics systems. Often such systems can also be categorized as stiff wave systems. Stiff wave systems are systems which exhibit a slow dynamical time scale while possessing fast wave phenomena. The physical effects of this fast wave may be important to the system, but resolving the fast time scale may not be required. When simulating such phenomena one would like to use time steps on the order of the dynamical scale for time integration. Historically, Semi-Implicit (SI) methods have been developed to step over the stiff wave time scale in a stable fashion and operator splitting methods have been used to decompose multiphysics aspects of such problems. However, these methods require some linearization and time splitting, and both of these can produce additional time integration errors. In this presentation, the concept of using SI methods and operator splitting methods as preconditioners to Jacobian-Free Newton-Krylov (JFNK) methods is developed. Applications in the areas of nuclear reactor simulation and ocean modeling will be discussed.

Dana A. Knoll
Idaho National Laboratory
Dana.Knoll@inl.gov

MS56

Parallel Block-Oriented Preconditioners for the Solution of the Semiconductor Drift-Diffusion Equa-

tions with Defect Physics

We apply block-oriented preconditioners to the implicit solution of the drift-diffusion equations with defect physics for semiconductor device modeling. These subblocks are solved by algebraic multigrid methods. The equations are discretized by a finite element method to produce the nonlinear coupled system, then solved with a parallel preconditioned Newton-Krylov method. Preliminary results will be presented to demonstrate the performance of these preconditioners relative to preconditioners that handle the algebraic system in a fully coupled manner.

Paul Lin
Sandia National Laboratories
ptlin@sandia.gov

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

MS57
HPC for Primary Breakup of a Diesel Jet

Massive computing is needed to resolve spray and jet breakup. We address a series of problems associated in the application of HPC to the problem. These include a robust description of the tracked interface, elimination of serial program segments to extreme scaling and code revision for thread-safe computing. Scientific conclusions of this study include identification of the mechanism for breakup (Kelvin-Helmholtz surface instability) and prediction of droplet size distribution.

Wurigen Bo, Tulin Kaman, Xingtao Liu
SUNY Stony Brook
bowrg@ams.sunysb.edu, tkaman@ams.sunysb.edu, liuxt12@ams.sunysb.edu

James G. Glimm
SUNY at Stony Brook
Dept of Applied Mathematics
glimm@ams.sunysb.edu

MS57
Simulation of a Reacting Jet in Crossflow on the Cray XT5 at ORNL

Direct numerical simulation of turbulent combustion requires leadership class computational capabilities. The Sandia DNS code S3D solves the governing equations for a reactive flow on a finite-difference grid in parallel using solution algorithms which scale well and production runs with O(30k) MPI-tasks are now routine. With supercomputers now providing over 100k computational cores, although the underlying algorithm continues to scale well, careful attention needs to be applied to the implementation details to ensure good performance.

Ray Grout
Sandia National Laboratories
rwgrout@sandia.gov

A. Gruber
SINTEF Energy Research, Trondheim, Norway
tba

Chun Sang Yoo
Combustion Research Facility
Sandia National Laboratories
csyoo@sandia.gov

Jacqueline Chen
Sandia National Laboratories
jhchen@sandia.gov

MS57
Scalability of Fluid-structure Interaction Solver Uintah using Optimized Task Graphs

The Uintah software uses asynchronous communication and a task-graph-based approach to solve challenging fluid-structure interaction problems using a combination of adaptive mesh refinement and MPM particle methods. Uintah originally ran computational tasks in predefined order; in order to achieve scalability new dynamic scheduler allows better overlapping of communication and computation and dynamic task rescheduling, including out-of-order execution. The effectiveness of the new approach is shown on large-scale fluid-structure examples through an analysis of the performance of the software.

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

Justin Luitjens
SCI Institute and School of Computing
University of Utah
luitjens@sci.utah.edu

Martin Berzins
University of Utah
mb@cs.utah.edu

MS57
High-Order Accurate Solution of Coupled Acoustic-Elastic Wave Propagation Problems on Adapted Meshes using a Discontinuous Galerkin Method on Massively Parallel Computers

Our goal is to develop scalable methods for global full-waveform seismic inversion. The first step we have taken towards this goal is the creation of a high-order accurate discontinuous Galerkin solver for the numerical simulation of wave propagation in media with fluid-solid interfaces. Along with convergence studies, we report strong and weak parallel scaling results for mesh generation and solution of the wave equations on adaptively resolved global Earth models.

Lucas Wilcox, Carsten Burstedde, Georg Stadler, Omar Ghattas
University of Texas at Austin
lucasw@ices.utexas.edu, carsten@ices.utexas.edu, georgst@ices.utexas.edu, omar@ices.utexas.edu

MS58
OpenCurrent: Solving PDEs on Structured Grids with CUDA

This talk will have two parts. In the first part, I will discuss the essential architectural features of a modern GPU and how throughput-oriented GPUs differ from multicore

CPUs. Starting from a straw man single-core CPU design, I will describe the design decisions that lead to increasing levels of on-chip parallelism, ultimately resulting in a many-core GPU. In the second part, I will describe OpenCurrent, an open source library for solving PDEs over structured grids using CUDA. OpenCurrent is designed to take advantage of the throughput-oriented nature of GPUs, and has been benchmarked against a comparable Fortran code running Rayleigh-Benard convection problems under a variety of different regimes. I will conclude with directions for future research, including data parallel numerical methods, and extensions to OpenCurrent for multi-GPU configurations.

Jonathan Cohen
NVIDIA
jocohen@nvidia.com

MS58

A Framework for Parallel Unstructured Grid Applications on GPUs

GPUs offer great increases in computational performance over multicore CPUs, but their programming is more complex and so the challenge is how to provide users with the computational benefits with the least programming effort. In this talk I will discuss progress in developing a parallel framework for unstructured grid calculations on GPUs. This builds on previous work using MPI distributed memory computing, and uses a set-based abstract view of a general class of unstructured grid computations.

Michael B. Giles
Mathematical Institute
Oxford University
Mike.Giles@maths.ox.ac.uk

MS58

Parallel Stochastic Simulation Using Graphics Processing Units for the Systems Biology Toolbox for Matlab

Graphics processing units (GPU) are well suited to decrease the computational intensity of stochastic simulation of chemical reaction systems. We compare Gillespie's Direct Method and Gibson-Bruck's Next Reaction Method on GPUs. The gain of the GPU implementation of these algorithms is approximately 120 times faster than on a CPU. Furthermore our implementation is integrated into the Systems Biology Toolbox for Matlab and acts as a direct replacement of its Matlab based implementation. The software is open source (GPL) and available at http://www.maths.ox.ac.uk/cmb/CUDA.

Guido Klingbeil
Centre for Mathematical Biology
University of Oxford
klingbeil@maths.ox.ac.uk

Radek Erban
University of Oxford
Mathematical Institute
erban@maths.ox.ac.uk

Philip K. Maini
Centre for Mathematical Biology
University of Oxford
maini@maths.ox.ac.uk

MS58

Massively Parallel Population-based Monte Carlo Methods with Many-core Processors

A recent trend in desktop computer architecture is the move from single-core processors to multi-core processors and further to many-core or massively multi-core processors. Therefore, statistical methods that can take advantage of many-core architectures can make the best use of the latest technology. A particularly promising avenue in this regard is the implementation of statistical algorithms for execution on graphics processing units (GPUs) since they are dedicated, low cost, low maintenance, energy-efficient devices that are becoming increasingly easy to program and are representative of a class of many-core architectures that is likely to be prevalent in the future. I present a case study on the suitability of using GPUs for three population-based Monte Carlo algorithms - population-based MCMC, sequential Monte Carlo samplers and the particle filter - with speedups ranging from 35 to 500 fold over conventional single-threaded computation.

Anthony Lee
University of Oxford
Department of Statistics
lee@stats.ox.ac.uk

MS59

Non-intrusive Uncertainty Quantification Using Random Fields in Parallel Finite Element Codes

Implementing uncertainty quantification in parallel finite element codes often requires a non-intrusive approach, such as a standalone, non-intrusive capability for generating random field (RF) realizations based on techniques such as polynomial chaos (PC) or Karhunen-Loeve (KL) expansions. This includes both an interface to characterize the uncertainty of the RF, as well as tools to generate the PC or KL realizations. We present details on large scale implementation and applications from solid mechanics and thermal/fluids codes.

Brian Carnes
Sandia National Laboratories
bcarnes@sandia.gov

John Red-Horse
Validation and Uncertainty Quantification Processes
Sandia National Laboratories
jrredho@sandia.gov

MS59

Parallel MCMC for Bayesian Inference in Inverse Scattering Problems

We present a new Parallel MCMC method for the solution of the Bayesian statistical inverse problem. Local Hessian and gradient information is used to adaptively construct a radial basis function approximation (RBF) of the posterior, which is used as proposal distribution for the Metropolis-Hastings algorithm across several parallel chains. Parallelism allows for rapid convergence of this approximation, and thus minimal sample correlation in the resulting MCMC chains.

Tan Bui-Thanh
The University of Texas at Austin
tanbui@ices.utexas.edu

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

James R. Martin
University of Texas at Austin
Institute for Computational Engineering and Sciences
jmartin@ices.utexas.edu

Lucas Wilcox
University of Texas at Austin
lucasw@ices.utexas.edu

MS59

Uncertainty Quantification in the Simulations of a Hypersonic Vehicle Flight

Flight conditions uncertainties create inherent difficulties in assessing hypersonic propulsion performance. We use numerical simulations to investigate the correlation of wind-tunnel and flight measurements for the HyShot vehicle. Simulations are validated against both reacting and non-reacting conditions in a ground-based facility. Next, we focus on reproducing the flight scenario using Bayesian inversion to infer the trajectory from noisy pressure measurements. The estimated conditions are then used to predict the thermal fields in the combustor.

Alireza Doostan
Stanford University
doostan@stanford.edu

Parviz Moin
Center for Turbulence Research
Stanford University/ NASA Ames
moin@stanford.edu

Gianluca Iaccarino
Stanford University
jops@stanford.edu

Rene Pecnik
Center for Turbulence Research - Stanford University
renep@stanford.edu

Vincent Terrapon
Stanford University
terrapon@stanford.edu

MS59

Parallel Multilevel Sampling Algorithms for Multimodal Distributions

In this talk we present a parallel algorithm for sampling multimodal distributions along with an illustrative example of its application. Sampling multimodal distributions is challenging because it requires correctly capturing the volume proportions among different modes. Parallel algorithms have to consider load balancing as well. The presented algorithm has the feature of sampling a sequence of distributions (levels) that converge to the final distribution of interest. It also carefully communicates sampling information between consecutive levels.

Ernesto E. Prudencio, Sai Hung Cheung
Institute for Computational Engineering and Sciences
prudenci@ices.utexas.edu, saihung@ices.utexas.edu

Karl W. Schulz
Institute for Computational Engineering and Sciences
The University of Texas at Austin
karl@ices.utexas.edu

MS60

A Parallel Algorithm for Computing a "Heavy" Perfect Matching in Bipartite Graphs

The problem of computing a maximum weight perfect matching in a bipartite graph has applications in various areas including the determination of good pivoting strategies in LU factorization. We present a parallel algorithm suitable for distributed memory computers for computing a "heavy" perfect matching in a bipartite graph. The algorithm is based on first using a parallel push-relabel algorithm to obtain a perfect matching. The obtained solution is then improved by searching for weight increasing augmenting cycles.

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

MS60

Recent Advances in Two-dimensional Sparse Matrix Partitioning

Sparse matrix-vector multiplication is a common kernel in many computational science simulations. An important combinatorial problem in parallel computing is distributing the matrix and the vectors among processors to minimize the communication cost. In this talk, we outline improvements made to a new 2D partitioning method, the nested dissection partitioning algorithm. We show empirical results that demonstrate an advantage of 2D partitioning over 1D partitioning for certain types of matrices.

Michael M. Wolf
Sandia National Laboratories
mmwolf@sandia.gov

Erik G. Boman
Sandia National Labs, NM
Scalable Algorithms Dept.
egboman@sandia.gov

Cedric Chevalier
Sandia National Laboratories
Scalable Algorithms Department
ccheval@sandia.gov

MS60

Partitioning, Load Balancing, and Matrix Ordering in a Parallel Hybrid Solver

We have been developing a parallel hybrid solver based on the Schur complement method to solve large highly-indefinite linear systems of equations. This method has the potential of combining the robustness of a direct method with the scalability of a preconditioned iterative method. However, to obtain high-performance on a large number of processor, its implementation needs to be carefully designed. In this talk, we discuss the effects of matrix partitioning and ordering on the load-balance and performance of our hybrid solver.

Ichitaro Yamazaki

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

Sherry Li
Lawrence Berkeley National Laboratory
xsli@lbl.gov

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

MS60

Reordering Sparse Matrices for Cache-Oblivious Computations

The sparsity structure of matrices directly influences cache-behaviour of matrix computation algorithms, e.g., sparse matrix–vector multiplication. Reordering input matrices can improve memory access patterns. We introduce the Separated Block Diagonal matrix structure which enhances cache use while remaining oblivious of any actual cache properties. Reordering a sparse matrix this way is closely related to sparse matrix partitioning methods (e.g., Mondriaan) as used in parallel computation. We also discuss methods for improving the reordering speed.

Albert-Jan Yzelman
Utrecht University
Department of Mathematics
a.n.yzelman@uu.nl

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

MS61

Language Requirements for Large-Scale Shared Memory Programming

Irregular applications can benefit significantly from global shared memory. Low overhead, single word transfers and dynamic thread scheduling independent of locality improve performance. Many memory operations in irregular codes access immutable values and private data. Storing such data in global memory rather than local data stores unnecessarily increases latency and consumes bandwidth. We examine the benefit of including local data stores in global memory machines and describe language extensions to facilitate their use.

John Feo
Pacific Northwest Laboratory
john.feo@pnl.gov

MS61

Parallel Motifs on the Microsoft Platform

The Motifs represent algorithm families that are important to the High Performance Computing community. We have selected a representative algorithm from each of the Motifs and implemented serial, shared memory, and distributed versions for each of the 13 Motifs, using both managed and native environments. We have also created an application to facilitate runs and comparisons between Motif implementations. This work represents a starting point from which Motif instances can be created, compared, and improved.

Robert Palmer, Shahrokh Mortazavi, Matej Ciesko
Microsoft
robert.palmer@microsoft.com, smortaz@microsoft.com, v-macies@microsoft.com

Vladimir Safonov, Nikolay Viskov, Grigory Fedyukovich, Ilya Yuneev
St. Petersburg University
vosafonov@gmail.com, nikolay.viskov@gmail.com, grigory.feduykovich@gmail.com, yuneev.ilya@gmail.com

MS61

The Architectural Requriements driving Language Evolution

The current MPI+Fortran ecosystem has sustained HPC application software development for the past decade, but was architected for coarse-grained concurrency largely dominated by bulk-synchronous algorithms. Future hardware constraints and growth in explicit on-chip parallelism will likely require a mass migration to new algorithms and software architecture that is as broad and disruptive as the migration from vector to parallel computing systems that occurred 15 years go. The challenge is to efficiently express intranode parallelism.

John Shalf
Lawrence Berkeley National Laboratory
jshalf@lbl.gov

Alice Koniges
Lawrence Berekely Laboratory
aekoniges@lbl.gov

MS61

Parallel Motifs as Paradigms for Comparing Language Performance

Parallel kernels or motifs provide a means for exploring new languages and performance issues key to a variety of areas of modern high performance computing. We describe a research project to develop these motifs and study their performance. The list of motifs is based on an idea called the parallel "dwarfs" from a paper by Phil Collela and later expanded by David Patterson.

Eric Strohmaier
Lawrence Berkeley National Laboratory
estrohmaier@lbl.gov

MS62

Multi-GPU Scalabilities for Mesh-based HPC Applications

Wonderful performances have been often achieved when HPC applications are successfully ported on a GPU. To overcome the memory-size limitation and pursuit further acceleration, we study multi-GPU scalabilities on multi-node GPU cluster (TSUBAME) for three mesh-based applications; the tsunami simulation, the Lattice Boltzmann method, the Phase Field model. Because of unbalance between the GPU computational performance and the interconnection speed, a technique overlapping the GPU-to-GPU communication with the computation has to be in-

troduced.

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

MS62

Developing Large-Scale Scientific Software For Generic Multi-core Nodes

Portable scientific applications must support a wide range of architectures, even at the rarefied level of high-performance computing. Modern supercomputers employ an increasing variety of processor types, from relatively mundane multi-core CPUs to hybrid architectures with more exotic attached processors. We describe some approaches taken in the Trilinos project in writing portable scientific libraries capable of supporting diverse platforms, including our custom communication interfaces and retargetable parallel node API.

Christopher G. Baker
Oak Ridge National Laboratory
cgbaker@gmail.com

Mike Heroux
Sandia National Laboratories
maherou@sandia.gov

MS62

(Pragmatically) Programming on the (Portable) Path to Productive Exa-scale

Within the context of some scientific application codes that have grown from giga- to tera- and now to the peta-scale, we will examine the means by which we intend to exploit exa-scale architectures. We are preparing for distinctly different programming and runtime requirements, requiring increased concurrency and granularity, both through a reorganization of our codes and reconfiguring our algorithms

Richard F. Barrett
Oak Ridge National Laboratory
rbarrett@ornl.gov

MS62

Performance Concerns for Coupling Hybrid-Parallel Kernels

Complex algorithms are typically implemented through a sequence of calls to computationally simple kernels to operate on algorithm-specified data. Hybrid-parallelization of these kernels on clusters of node with GPGPUs has demonstrated performance gains for individual kernels. In order for algorithms to realize a similar performance gain the programming model for coupling these kernels must not significantly penalize performance. Performance concerns for coupling hybrid-parallel kernels are analyzed against different programming model strategies.

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

MS63

A Block Preconditioner for the 2D Vector Potential MHD Equations

The Magnetohydrodynamics (MHD) equations model the fluid dynamics of a plasma in the presence of a magnetic field. Because of the inherent multiphysics behavior, MHD poses a formidable challenge to developing effective solvers. In this talk, we propose a Newton-Krylov method using a novel block preconditioner for the incompressible 2D vector potential MHD formulation. The preconditioner is an approximate block factorization of the Jacobian operator that localizes the coupling between the physics into an approximate Schur complement. This decomposition facilitates the use of existing solver technology on the individual blocks while accounting for coupling inherent in the problem. Performance and scaling results of the block preconditioner compared to a fully coupled algebraic multilevel preconditioner will be presented.

Eric Cyr
Scalable Algorithms Department
Sandia National Labs
eccyr@sandia.gov

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

Ray S. Tuminaro
Sandia National Laboratories
rstumin@sandia.gov

MS63

Fast Solvers for Models of Microfluidic Flows

We demonstrate the performance of fast algorithms for modeling the design of a microfluidic mixing device. The device uses an electrokinetic process, *induced charge electroosmosis*, by which flow through the device is driven by a set of polarizable obstacles in it. Its design is realized by manipulating the shape and orientation of the obstacles to maximize fluid mixing within the device. The computation entails solution of a constrained optimization problem in which function evaluations require the numerical solution a potential equation, the incompressible Navier-Stokes equations, and a mass transport equation. Using preconditioners that take advantage of problem structure, together with a derivative-free pattern search algorithm for optimization, we determine optimal configurations of microfluidic devices.

Howard C. Elman
University of Maryland, College Park
elman@cs.umd.edu

Robert Shuttleworth
Exxon
robert.shuttleworth@gmail.com

Kevin Long
Texas Tech University
kevin.long@ttu.edu

Jeremy Templeton
Sandia National Laboratories
jatempl@sandia.gov

MS63

Block Preconditioners and Advanced Software for Multiphysics Problems in Solid Earth Geodynamics

Solid Earth Geodynamics offers a host of multiphysics problems ranging from thermo-chemical convection to reactive fluid flow in deformable porous media ("magma dynamics"). Traditionally, these problems have been solved with various splitting/lagging schemes with little control on the convergence of the full non-linear problem. Here we describe our experience with the magma dynamics problem using these recipes as preconditioners for the inner solve of a Newton-Krylov method. In particular, we describe methods using available software for composition and assembly of exact Jacobians (using FFC/Dolfin from the FEniCS project) and FieldSplit preconditioners from PETSc.

Marc Spiegelman
Columbia University
Lamont-Doherty Earth Obs.
mspieg@ldeo.columbia.edu

MS63

A Solver for Bifurcation Analysis of Ocean-Climate Models

In [A. de Niet et al., A tailored solver for bifurcation analysis of ocean-climate models, J. Comput. Phys. 2007, 227(1),654-679] we developed a block preconditioner for the 3D ocean equations. In the construction of it, we made use of the physics of the problem. In fact there is a part for fluid motion, which is forced by wind and boyancy, and a tracer part for heat and salinity which is forced by differences in temperature and salinety at the sea surface. The former is a kind of incompressible Navier-Stokes equation with the special property that the momentum part is dominated by the Coriolis force. We used a SIMPLE preconditioner for the fluid motion part, where an important adaptation is needed to handle the Coriolis term. The tracer part is of convection-diffusion type, which allows for a standard method. In [J.Thies et al., Bifurcation analysis of 3D ocean flows using a parallel fully implicit ocean model, Ocean Modelling, 2009, DOI information: 10.1016/j.ocemod.2009.07.005] we parallized this code making use of Trilinos. This lead to an almost optimal speedup for up to 20 processors. In the talk we will discuss the method in more detail.

Fred Wubs
Dept. of Math and Computer Science
University of Groningen, Groningen, The Netherlands
wubs@math.rug.nl

Jonas Thies
Dept. of Math and Computer Science
University of Groningen
j.thies@rug.nl

MS64

The Roadrunner Computing Architecture: System Overview and Programming Models

Roadrunner is the first supercomputer to use a hybrid processor architecture, which is based on both AMD Opteron x86_64 processors and IBM PowerXCell 8i processing elements. As such, this architecture presents unique challenges to developers of codes for scientific computing applications. This talk will cover the basic system design, and will discuss some of the approaches and progarmming models used in developing codes for the Roadrunner Open Science Project.

Ben Bergen
Los Alamos National Laboratory
bergen@lanl.gov

MS64

SPaSM: Large-Scale Molecular Dynamics Studies of Material Dynamics on Roadrunner

Large-scale classical molecular dynamics simulations with 10^6 to 10^{12} atoms are providing unprecedented insight into material deformation processes under high strain-rate mechanical loading. We will describe the algorithm redesign motivated by the evolution in computer architectures over the past decade, specifically the heterogeneous LANL Roadrunner platform, the resulting performance, and initial scientific applications to understand the response of copper single- and poly-crystals to shock compression and release.

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

MS64

Direct Numerical Simulations of Compressible Reacting Turbulence with Type Ia Supernovae Microphysics

Abstract unavailable at time of publication.

Daniel Livescu
Los Alamos National Laboratory
livescu@lanl.gov

MS64

Understanding the Onset and Saturation of Laser Backward Stimulated Raman Scattering Through Large-Scale Plasma Kinetic Simulations on a Hybrid Supercomputer

Achieving inertial confinement fusion (ICF) ignition will have far-reaching consequences. However, laser-plasma interaction (LPI) may jeopardize ignition. To assess LPI risk, at-scale 3D VPIC particle-in-cell simulations have been performed on Roadrunner. In this talk, a description of the VPIC algorithm will be given, as well as challenges in its deployment. Discussion will highlight unique science enabled and special difficulties encountered for modeling LPI on memory bandwidth starved platforms.

Lin Yin, Brian Albright, Kevin Bowers, Ben Bergen
Los Alamos National Laboratory
lyin@lanl.gov, balbright@lanl.gov, kevin.j.bowers@ieee.org, bergen@lanl.gov

MS65

Blendenpik: Supercharging LAPACK's Least-squares Solver

Several innovative random-sampling and random-mixing techniques for solving problems in linear algebra have been proposed in the last decade, but they have not yet made a significant impact on numerical linear algebra. We show that by using an high quality implementation

of one of these techniques we obtain a solver that performs extremely well in the traditional yardsticks of numerical linear algebra: it is significantly faster than high-performance implementations of existing state-of-the-art algorithms, and it is numerically backward stable. More specifically, we describe a least-square solver for dense highly overdetermined systems that achieves residuals similar to those of direct QR factorization based solvers (LAPACK), outperforms lapack by large factors, and scales significantly better than any QR-based solver.

Haim Avron
School of Computer Science
Tel-Aviv University
haim.avron@gmail.com

Petar Last Maymounkov
MIT
petarm@gmail.com

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

MS65

Randomized Algorithms in Linear Algebra: From Approximating the Singular Value Decomposition to Solving Regression Problems

The introduction of randomization in the design and analysis of algorithms for matrix computations (such as matrix multiplication, least-squares regression, the Singular Value Decomposition (SVD), etc.) over the last decade provided a new paradigm and a complementary perspective to traditional numerical linear algebra approaches. These novel approaches were motivated by technological developments in many areas of scientific research that permit the automatic generation of large data sets, which are often modeled as matrices. In this talk we will outline how such approaches can be used to approximate problems ranging from matrix multiplication and the Singular Value Decomposition (SVD) of matrices to approximately solving least-squares problems and systems of linear equations.

Petros Drineas
Computer Science Department
Rennselaer Polytechnic Institute
drinep@cs.rpi.edu

MS65

A Randomized Algorithm Minimizing Communication for Parallel and Sequential EIG/SVD

Abstract unavailable at time of publication.

Ioana Dumitriu
University of Washington, Seattle
dumitriu@math.washington.edu

MS65

Randomized Algorithms for Matrices and Large-Scale Data Applications

Abstract unavailable at time of publication.

Michael Mahoney
Stanford University
Applied and Computational Mathematics
mmahoney@cs.stanford.edu

PP1

Automatic Sparse Preconditioners for the Spectral Element Method

The spectral element method (SEM) is a high-order discretization method that joins spectral method accuracy with finite element method (FEM) sparsity. The SEM has been found to model physics that the low-order FEM fails to capture. Fast linear solvers are needed for SEM so simulations can be maximally efficient. We introduce a sparse preconditioner that will be added to modern solver packages so SEM users have an efficient and sparse preconditioner.

Travis M. Austin
Tech-X Corporation
austin@txcorp.com

Thomas Manteuffel
University of Colorado
tmanteuf@colorado.edu

Marian Brezina
U. of Colorado, Boulder
Applied Math Department
marian.brezina@colorado.edu

John Ruge
University of Colorado at Boulder
ruge@colorado.edu

PP1

Real-Time Signal Processing Using Multi-Core Devices

The availability of multi-core processor chips for commodity implementations offers potential acceleration for real-time and embedded signal processing platforms. Key metrics in these contexts include energy consumption, as well as processing time. Motivated by these concerns, we present transverse-vectorized implementations of the FFT on several multi-core architectures, including the IBM Cell and Nvidia Tesla. We additionally report on progress to utilize the low-power many-core architecture of the Coherent Logix HyperX for real-time signal processing.

Travis Humble, Jacob Barhen
Oak Ridge National Laboratory
humblets@ornl.gov, barhenj@ornl.gov

Pramita Mitra
Oak Ridge National Laboratory and University of Notre Dame
mitrap@ornl.gov

PP1

An Adaptive Checkpoint/Restart Library for Large Scale HPC Applications

The effective use of large scale HPC systems requires checkpoint/restart enabled scientific applications to incorporate, often system specific, advanced storage and fault notification techniques. The application level checkpoint/restart library presented encapsulates and adapts to the availability of such techniques. When available, this library transparently provides fault notification, system level check-

point/restart, and scalable storage services. We present our experiences developing this library with the LAMMPS molecular dynamics simulator, CIFTS Fault Tolerance Backplane and Open MPI library.

Joshua Hursey
Open Systems Laboratory
Indiana University
jjhursey@osl.iu.edu

Scott Hampton
Oak Ridge National Laboratory
hamptonss@ornl.gov

Pratul Agarwal
ORNL
agarwalpk@ornl.gov

Andrew Lumsdaine
Indiana University
lums@osl.iu.edu

PP1

Fully Implicit, Jacobian-Free, Newton-Krylov Methods in Production Level MHD Fusion Codes

Recently, a framework for efficiently constructing a fully-centered, fully-implicit method in magnetohydrodynamics (MHD) simulations was introduced by Chacn. This work allows for the accurate computation of nonlinear terms and large time-steps for advancing the MHD system on massively parallel systems. We report on initial progress at implementing a fully-centered, fully-implicit method in NIMROD, a DOE MHD simulation tool, using the PETSc toolkit which enables access to high-performance, scalable solvers provided by the SNES library.

B. Jamroz, T. Austin, S. Kruger
Tech-X Corporation
jamroz@txcorp.com, austin@txcorp.com, kruger@txcorp.com

PP1

SPIKE: A Parallel Scalable Sparse Linear System Solver

Solving large sparse linear systems is the most time consuming operation in the simulation of Micro-Electro-Mechanical-Systems (MEMS). We compare the parallel scalability of our SPIKE algorithm with Trilinos-ML for solving a large sparse system arising in the finite volume discretization of a MEMS device. This experiment effectively compares the parallel scalability of an algebraic multigrid preconditioning scheme (Trilinos-ML) with the same Krylov method preconditioned via a particular central band of the coefficient matrix (SPIKE).

Murat Manguoglu
Purdue University
mmanguog@cs.purdue.edu

Faisal Saied
fsaied@purdue.edu
purdue university

Ahmed Sameh
Purdue University
sameh@cs.purdue.ed

PP1

GPU Acceleration of Parallel Out-Of-Core Dense Linear Solvers

Method of Moments codes used for electromagnetic scattering simulations require the solution of large, dense linear systems. These matrices often cannot fit into main memory and out-of-core solvers are required. GPUs have the potential to reduce the computational cost of such solvers, but the complex memory hierarchy presents new challenges. In this poster we will describe a GPU accelerated, parallel out-of-core dense linear solver and present theoretical scaling and measured performance resutls. This work was supported by DOD/NAVY SBIR Grant N68335-09-C-0247.

Peter Messmer, Matthew Koch, Paul Mullowney
Tech-X Corporation
messmer@txcorp.com, mkoch@txcorp.com, paulm@txcorp.com

PP1

High Performance Computing for Groundwater Reactive Transport Modeling

Understanding and predicting fate and transport of radionuclide in the environment entails groundwater reactive transport modeling. Due to complexity of the subsurface environment, reactive transport modeling at the field scale is computationally demanding. In this study, we proposed a physically-based parallelization mechanism, based on which reaction equations are simulated simultaneously at each node of the computational grid. Based on the HydroGeoChem 5.0 (HGC5) code, we implemented this parallel mechanism using the Message Passing Interface (MPI).

Geoffery L. Miller
Florida State University
Department of Scientific Computing
geoffery.miller@gmail.com

Ming Ye
Department of Scientific computing
Florida State University
mye@fsu.edu

Guoping Tang, Fan Zhang
Environmental Sciences Division
Oak Ridge National Laboratory
tangg@ornl.gov, zhangf@ornl.gov

Jack Parker
Department of Civil and Environmental Engineering
University of Tennessee
jparker@utk.edu

David Watson, Philip Jardine
Environmental Sciences Division
Oak Ridge National Laboratory
watsondb@ornl.gov, jardinepm@ornl.gov

PP1

Parallel High Order Integrators

In this work we discuss a class of parallel high order time integrators, ideally suited for developing methods which can be order adaptive in time. The method is based on Integral Deferred Correction (IDC), which was itself motivated by

Spectral Deferred Correction (SDC) by Dutt, Greengard and Rokhlin (BIT-2000). The method presented here is a revised formulation of explicit IDC, dubbed Revisionist IDC, which can achieve p^{th} order accuracy in wall-clock time comparable to a single forward Euler simulation, on problems where the time to evaluate the right hand side of a system of differential equations is greater than latency costs of inter-processor communication, such as in the case of the -body problem. The key idea is to rewrite the defect correction framework so that, after initial startup costs, each correction loop can be lagged behind the previous correction loop in a manner that facilitates running the predictor and correctors in parallel. Various RIDC schemes are shown to be significantly faster than the popular fourth-order Runge–Kutta method on an example N-body calculation. The ideas behind RIDC extend to implicit and semi-implicit IDC and have high potential in this area.

Benjamin Ong
Department of Mathematics
Michigan State University
bwo@math.msu.edu

Andrew J. Christlieb
Michigan State Univerity
Department of Mathematics
andrewch@math.msu.edu

Colin Macdonald
Oxford University
cbm@math.ucla.edu

PP1

A Simulator for Large-scale Parallel Computer Architectures

Efficient hardware and software design for large-scale parallelism requires detailed understanding of interactions between the application, computer, and network. We developed a macroscale simulator (SST/macro) that permits coarse-grained study of distributed-memory applications. The simulator is driven from either a trace file or a skeleton application. Currently, applications using MPI are simulated; however, the modular simulator architecture allows inclusion of other programming models, as well as novel network models, trace file formats, and detailed processor models.

Ali Pinar
Sandia National Labs
apinar@sandia.gov

Curtis Janssen
Sandia National Laboratories
cljanss@sandia.gov

Helgi Adalsteinsson
Sandia National Laboratories, Livermore CA
hadalst@sandia.gov

Scott Cranford, Joseph Kenny, David Evensky, Jackson Mayo
Sandia National Laboratories
scranfo@sandia.gov, jpkenny@sandia.gov, evensky@sandia.gov, jmayo@sandia.gov

PP1

Hopspack Software Framework for Parallel Derivative-Free Optimization

HOPSPACK (Hybrid Optimization Parallel Search PACKage) solves derivative-free optimization problems using an open source C++ framework. The framework enables parallel operation using MPI and multithreading. Multiple algorithms can be hybridized to run simultaneously, sharing a cache of computed objective and constraint function evaluations that eliminates duplicate work. HOPSPACK comes with a Generating Set Search algorithm, and is easily extended by developers to add new algorithms.

Todd Plantenga, Tamara G. Kolda
Sandia National Laboratories
tplante@sandia.gov, tgkolda@sandia.gov

PP1

Open CL Implementation of the Kullback-Leibler Divergence for Weighted Samples

Graphic cards provide enormous amounts of parallel computing power. However, actually using it is sometimes rather technical. We consider the Kullback-Leibler divergence, a measure for the similarity between two sample sets, that is studied in the context of particle filters. Using the OpenCL standard, we were able to add a GPU accelerated implementation to the open source Bayesian Filtering Library (BFL) in a portable fashion and observed speedups of up to 250 (relative to C++).

Koen Poppe
K.U.Leuven, Belgium
Koen.Poppe@cs.kuleuven.be

Ronald Cools
Dept. of Computer Science
K.U.Leuven
ronald.cools@cs.kuleuven.be

Herman Bruyninckx
K.U.Leuven, Belgium
herman.bruyninckx@mech.kuleuven.be

PP1

Algebraic Distance on (Hyper)graphs with Applications to Combinatorial Scientific Computing

Measuring the connection strength between a pair of vertices in a (hyper)graph is one of the most vital concerns in many graph applications. We present a measure of the connection strength (called the algebraic distance) defined from stationary iterative processes. Its computation is simple, linear, and easily parallelized. We demonstrate rigorous analysis and practical effectiveness of the proposed measure through several optimization problems on (hyper)graphs and multiscale methods.

Ilya Safro
Argonne National Laboratory
safro@mcs.anl.gov

Achi Brandt
UCLA
abrandt@math.ucla.edu

Jie Chen
Department of Computer Science and Engineering

University of Minnesota
jchen@cs.umn.edu

Dorit Ron
The Weizmann Institute of Science
dorit.ron@weizmann.ac.il

PP1

Multigrid Scalability for Least Squares Finite Element Problems

Least-squares finite element methods are an attractive class of methods for the numerical solution of certain div-curl systems. We consider a discretization where we approximate either H(curl) or H(div) spaces conformally and use a discrete approximation of the other. We derive a multigrid method from previous work of Bochev, Hu, Siefert, Tuminaro, Xu and Zhu and show good parallel scalability up to 5000 processors.

Chris Siefert
Sandia National Laboratories
csiefer@sandia.gov

Pavel Bochev
Sandia National Laboratories
Computational Math and Algorithms
pbboche@sandia.gov

Kara Peterson
Sandia National Laboratories
kjpeter@sandia.gov

PP1

Remeshing for Vortex Methods in Parallel Tree Codes

We present a new remeshing approach for parallel vortex tree codes. The necessity of particle resets in Lagrangian methods can now be met using a mesh-free tree algorithm. Our sorting approach avoids the use of an underlying grid and therefore preserves the inherent mesh-free character of vortex particle model and tree code concept. Since sorting is necessary in every tree code and already efficiently implemented in parallel, the remeshing overhead is limited.

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

PP1

A Performance-Based Load Balancing Scheme for Madness

We discuss a novel load-balancing algorithm for MADNESS (Multiresolution Adaptive Numerical Environment for Scientific Simulation). Computations in MADNESS generate octrees with millions of nodes distributed across tens or hundreds of thousands of processors, making load balancing crucial for performance. This scheme monitors the computational load on each tree node, allowing adaptive balancing of the octree based on actual performance data as the simulation progresses. We show the impact of this scheme on performance.

Paul Matthew Sutter
University of Illinois at Urbana-Champaign
psutter2@illinois.edu

Rebecca J. Hartman-Baker
Oak Ridge National Laboratory
hartmanbakrj@ornl.gov

Robert Harrison
University of Tennessee
Oak Ridge National Laboratory
harrisonrj@ornl.gov

PP1

Practical l1 Optimization Using Parallel Computing

There are multiple applications of l1-optimization, a sparsity inducing regularization technique. In our project arising in Global Seismic Tomography, we are faced with the problem of reconstructing sparse solutions in the wavelet domain from very large under-determined linear systems of equations. We use recently developed iterative thresholding techniques which utilize simple linear operations and yield well to parallel computation. We present these issues and the use of parallel packages PETSc and SLEPc in our work.

Sergey Voronin
Princeton University
svoronin@princeton.edu

Ingrid Daubechies
Princeton University
Department of Mathematics
ingrid@math.princeton.edu

Guust Nolet
Geosciences Azur
nolet@princeton.edu

Frederik Simons
Princeton University
fjsimons@princeton.edu

Ignace Loris
Vrije Universiteit Brussel, Belgium
igloris@vub.ac.be

Stephen Judd, Cédric Vonesch, Philip Vetter
Princeton University
judd@math.princeton.edu, cvonesch@princeton.edu,
pvetter@math.princeton.edu

Jean Charlety
Geosciences Azur, France
charlety@geoazur.unice.fr