CP1

Efficient Parallel Simulation for Stochastic Simulation of Biochemical Systems on the Graphics Processing Unit

The small populations of some reactant species in biological systems formed by living cells can result in inherent randomness that cannot be captured by traditional deterministic (ordinary differential equation) simulation. A more accurate simulation can be obtained by using the Stochastic Simulation Algorithm (SSA). Many stochastic realizations are required to obtain accurate probability density functions. This carries a very high computational cost. The current generation of general-purpose graphics processing units (GPU) is well-suited to this task. Computational experiments illustrate the power of this technology for this important and challenging class of problems.

Hong Li
Department of Computer Science
University of California, Santa Barbara
hongli@cs.ucsb.edu

Linda Petzold
UCSB
petzold@cs.ucsb.edu

CP1

From HPC to Grid Computing: CSE Scenarios with GridSFEA

With the simulation framework GridSFEA, we intend to provide the HPC application developer and user with a complementary instrument for working in grid. GridSFEA services, application library, and wrappers simplify tasks such as parameter studies and long running simulations. This is important, since the HPC community lacks well-suited and usable grid tools, and, thus, remains reluctant to exploit computing grids. We demonstrate several successful scenarios such as molecular dynamics and traffic simulations in grid.

Ioan L. Muntean, Hans-Joachim Bungartz, Martin Buchholz, Michael Moltenbrey, Ekaterina Elts, Dirk Pflüger
Technische Universität München, Department of Informatics
Chair of Scientific Computing in Computer Science
muntean@in.tum.de, bungartz@in.tum.de, buchholm@in.tum.de, moltenbr@in.tum.de, elts@in.tum.de, pflueged@in.tum.de

Ralf-Peter Mundani
TUM, Faculty of Civil Engineering and Geodesie
Chair for Computational Civil and Environmental Engineering
mundani@tum.de

CP1

Plasma Turbulence Computation and Visualization on the GPU

We present techniques to accelerate plasma turbulence simulations on modern Graphics Processing Units (GPUs). We contrast and compare the performance of two classes of methods: spectral methods used in MHD models and particle methods used in gyrokinetic models. We demonstrate a prototype of a scalable computational steering framework based on a tight coupling of plasma turbulence simulation and visualization on the GPU.

George Stantchev
University of Maryland
gogo@math.umd.edu

CP1

Parallel Numerical Methods for Solving Nonlinear Evolution Equations that Model Optical Fiber Communication Systems

Nonlinear evolution equations of the Nonlinear Schrdinger types are of tremendous interest in both theory and applications. Various regimes of pulse propagation in optical fibers are modeled by some form of the NLS and CNLS equations. In this talk we introduce parallel algorithms for numerical simulations of these equations. The parallel methods are implemented on the IBM p655 multiprocessor computer. Our numerical experiments have shown that the used methods give accurate results and considerable speedup.

Thiab R. Taha
Professor at the University of Georgia
thiab@cs.uga.edu

CP1

Efficient Parallel Algorithm for 2D2V Vlasov Equation with High Order Spectral Element Method

For decades kinetic space plasma simulation has been dominated by PIC (Particle-In-Cell) codes. Due to its inherent noise, solving the Vlasov equation directly becomes more promising. In this work, we are developing tera-scalable scheme of 2D2V Vlasov solver using high order spectral element method. For this 4 diemnsions problem, efficient parallel algorithm is necessary due to memory and speed requirements. The results have been virified with PIC codes.

Jin Xu
Physics Division
Argonne National Lab.
jin_xu@anl.gov

CP1

Using Multi-Core, Multi-CPU, PC Clusters for Statistical 3-D Virus Reconstructions from Cryo Electron Microscopy Images

Cryo electron microsopy images of viruses provide roughly projection data at unknown projection angles and signal to noise ratios less than 1/2. Statistical approaches to 3-D reconstruction (e.g., Doerschuk and Johnson, IEEE Transactions Information Theory, 2000) require extensive computation. A system to perform such computations on a PC cluster with multi-cpu and multi-core nodes using C, MPI, and OpenMP is described including experiments with different algorithms and different mixtures of parallelism methodologies.

John Johnson
The Scripps Research Institute
Department of Molecular Biology
jackj@scripps.edu

Yili Zheng
Electrical and Computer Engineering
Purdue University
yz247@cornell.edu

Peter C. Doerschuk
Cornell University
Biomedical Engineering and Electrical and Computer Engr.
pd83@cornell.edu

CP2

Electronic Structure Calculations for Large Nanosystems on Parallel Computers

Electronic Structure calculations based on the density functional theory approach have become one of the biggest consumers of cycles on high performance computers around the world. In this talk I will discuss this approach, as used in nanoscience applications on high performance computers as well as new methods that go beyond density functional theory and allow us to simulate much larger systems with first principles accuracy. Performance of new parallel solvers for these methods on high performance parallel computers such as the IBM BGL, Cray XT4, NEC Earth Simulator and PC clusters will also be discussed. I will discuss some applications of these methods to nanosystems and detector materials. Work done in collaboration with L-W Wang, O. Marques, S. Tomov and C. Voemel.

Andrew M. Canning
Lawrence Berkeley National Laboratory
acanning@lbl.gov

CP2

Parallelized Topography Simulation for Electronic Device Manufacturing

Three-dimensional topography simulation for electronic device manufacturing is a demanding task regarding computational resources. We developed a parallelized method using a Monte Carlo algorithm for flux calculation and level-set algorithms for front tracking. The computational complexity scales optimally with surface size. The flux calculation is parallelized using MPI and OpenMP. Communication is minimized by storing the compressed structure geometry on each node, which makes our method applicable for cheap infrastructure consisting of conventional PCs with LAN connection.

Otmar Ertl, Johann Cervenka, Siegfried Selberherr
Vienna University of Technology
Institute for Microelectronics
ertl@iue.tuwien.ac.at, cervenka@iue.tuwien.ac.at, selberherr@iue.tuwien.ac.at

CP2

Parallel Block-Oriented Preconditioners for Fem Modeling of Semiconductor Devices

Various approximate block factorization and physics-based preconditioners are applied to the drift-diffusion equations for modeling semiconductor devices. The resulting scalar subsystems are solved by various iterative methods including AMG-type techniques. We employ a stabilized finite element discretization of the drift-diffusion equations on unstructured meshes. The nonlinear coupled system is solved with a parallel preconditioned Newton-Krylov method. Preliminary results will be presented demonstrating the performance of the block-oriented preconditioners compared with one-level and multilevel preconditioners. This work was partially funded by the DOE NNSA's ASC Program and the DOE Office of Science AMR Program, and was carried out at Sandia National Laboratories operated for the U.S. Department of Energy under contract no.DE-ACO4-94AL85000

Paul Lin
Sandia National Laboratories
ptlin@sandia.gov

Gary Hennigan
Sandia National Labs
glhenni@sandia.gov

Robert J. Hoekstra
Sandia National Laboratories
rjhoeks@sandia.gov

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

CP2

A Massively Parallel Schroedinger Solver for Nano-Electronics

NEMO3D is a scalable simulator for nano-electronic devices such as quantum dots that uses a quantum mechanical description of the device. A key computational challenge is to compute degenerate eigenstates in the interior of the spectrum for a very large Hamiltonian matrix, with up to 10^9 degrees of freedom. We compare several parallel eigensolvers and present detailed performance results that demonstrate the petascale potential of NEMO3D on state of the art parallel platforms.

Maxim Naumov
Purdue University
Department of Computer Sciences
naumov@purdue.edu

Faisal Saied, Hansang Bae, Steve Clark, Ben Haley, Gerhard Klimeck, Sunhee Lee
Purdue University
fsaied@purdue.edu, baeh@purdue.edu, clarks@purdue.edu, bhaley@ecn.purdue.edu, gekco@purdue.edu, lee509@purdue.edu

CP2

Parallel Methods for Electronic Transport Through Nanoscale Devices

Nanoelectronics is a fast developing field. Therefore understanding of electronic transport at the nanoscale is currently of great interest. In this talk, we present new parallel algorithms to calculate the electronic transport through large nanoscale devices consisting of many thousand atoms. Applying the semi-empirical Extended Hückel Theory to model nanowire junctions, we compare the parallel performance of our recently developed direct and iterative approaches, the latter of which uses preconditioned Krylov subspace techniques.

Hans Henrik B Sørensen
Informatics and Mathematical Modelling
Technical University of Denmark
hhs@imm.dtu.dk

Martin van Gijzen

Numerical Analysis Group
Delft University of Technology
m.b.vangijzen@ewi.tudelft.nl

Dan Erik Petersen, Stig Skelboe
Department of Computer Science
University of Copenhagen
danerik@diku.dk, stig@diku.dk

Kurt Stokbro
NBI and Department of Computer Science
University of Copenhagen
kurt.stokbro@gmail.com

Per Christian Hansen
Technical University of Denmark
Informatics and Mathematical Modelling
pch@imm.dtu.dk

CP3

Cortically-Inspired Parallel Processing

Symbolic logic and serial computations have proven inadequate for modeling human cognition and solving hard cognitive problems. On the other extreme, the recent use of a Blue Gene supercomputer for the molecular simulation of a single cortical column, abeit a remarkable leap for molecular neuroscience, is still very far from explaining cognition. I will present an implementation of a massively parallel, highly scalable, heteroassociative network of attractor networks that abstracts the functionality of millions of cortical columns and explains phenomena of visual cognition. This is a paradigm of how brain-like computing might look in the future.

Socrates Dimitriadis
Brown Univeristy
Department of Cognitive and Linguistic Sciences
socrates@brown.edu

CP3

Uc-geowave: A Stereo-distributed-parallel Application for Seismic Modeling in Oil Exploration

This paper describes a stereo–distributed–parallel application for the simulation of two–dimensional acoustic wave propagation on heterogeneous media, called UC–geoWave. This application has three modules: pre–processing, processing and post-processing. Using the pre–processing module the user can create synthetic terrain in 3D and finally view it in stereoscopic way. In the processing module, the application computes a parallel reverse time migration (RTM) of a seismic data set to obtain a depth imaging. This migration technique is based on the solution of the acoustic wave equation in 2D using a finite difference scheme. This phase is executed on a cluster machine using MPI. In the post-processing module, the user can load the resulting data and see the graphics. The application can be accessed over the internet using any commercial browser (internet explorer, netscape, etc.) and run on different operating systems (Windows, MacOS X, Linux and Solaris).

Juan Medina
Universidad de Carabobo
juanruben@gmail.com

German A. Larrazabal
University of Carabobo
Valencia-Venezuela
glarraza@uc.edu.ve

CP3

Ratio-Based Parallel Time Integration (RaPTI) for Satellite Trajectories

Chartier, Philippe (1993), Erhel, Rault (2000) have intoduced Parallel time integration for satellites trajectories. We apply a version of RaPTI algorithm to solve this problem. RaPTI is a predictor-corrector scheme based on automatic generation of time slices (Nassif et al (2005)) at end of which solution values exhibit a ?ratio phenomenon?. Such approach leads to parallel time integration schemes used in previous authors works (2006-2007). The present paper extends RaPI to a J2 perturbed satellite trajectory.

Nabil R. Nassif
Mathematics Department
American University of Beirut
nn12@aub.edu.lb

Jocelyne Erhel, Noha Makhoul-Karam
IRISA,
UNIVERSITE DE RENNES
erhel@irisa.fr, noha.makhoul@irisa.fr

Yeran Soukiassian
Mathematics Department
American University of Beirut
ys07@aub.edu.lb

CP3

A Scalable Parallel Classification Algorithm for Remote Sensing

Previous work on the parallel IGSCR (iterative guided spectral class rejection) classification algorithm for remote sensing resulted in good speedup through 64 processors, however, speedup began deteriorating beyond 16 processors. This work will tackle scalability issues with parallel clustering that will be essential in a scalable distributed memory version of IGSCR. These issues include clustering schemes that are more amenable to a parallel environment and employing load balancing methodologies that will increase overall parallel efficiency.

Layne T. Watson
Virginia Polytechnic Institute and State University
Departments of Computer Science and Mathematics
ltw@cs.vt.edu

Rhonda D. Phillips
Virginia Polytechnic Institute and State University
rdphllps@vt.edu

Randolph H. Wynne
Virginia Polytechnic Institute and State University
Department of Forestry
wynne@vt.edu

CP3

Parallel Implementation of Data Mining Algorithms

This paper discusses the parallel or distributed implementation of key data mining algorithms in the areas of col-

laborative filtering and latent semantic analysis. Both optimization problems are elegantly captured by matrix representations which are usually sparse and with very large dimensions. We discuss the implementation of these algorithms in architectures of different levels of granularity such as dedicated highly coupled parallel processors, loosely coupled parallel processors and a distributed platform.

Yosef G. Tirat-Gefen
Castel Research Inc. and George Mason University
yosefgavriel@computer.org

CP4

Higher-Level Abstractions and Patterns for Designing Data-Parallel Applications

The level of abstraction provided by the Message Passing Interface (MPI) are too low-level and enormous amounts of time and effort is spent in refactoring existing code to include MPI primitives. This presentation described higher-level abstractions and design patterns to encapsulate data distribution, communication, and load balancing. Several applications developed using a single pattern will be presented and their performance comparisons with hand-written versions of the applications will also be provided.

Purushotham Bangalore
Univ. of Alabama at Birmingham
Dept. of Computer and Information Sciences
puri@cis.uab.edu

CP4

Parallel Io and Data Management for Data Structures in Applications

We have developed a middle layer between parallel file systems/MPI-IO and applications, through which applications are able to efficiently use the most efficient part of MPI-IO and parallel file systems for millions of distributed un-aligned small datasets. For a variety of data structures in applications, such as unstructured mesh and variable, the middle layer provides sustainable, interoperable, efficient, scalable, and convenient tools for parallel IO and data management. The IO performance of the middle layer for high-level data structures in applications is almost the same as the performance of MPI-IO for large datasets. The IO performance of either collective or non-collective calls for millions of distributed non-aligned small datasets is comparable to the performance of MPI-IO for large datasets.

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

CP4

Using GPUs From High-level Programming Languages

This talk describes a high-productivity development model for general purpose computing on Graphics Processing Units (GPUs). This is accomplished by exposing the capabilities of NVIDIA's CUDA architecture to high-level programming languages, such as Python, IDL, MATLAB and Java. In this talk we describe an array-based programming interface that hides the details of CUDA and the GPU from the user, allowing them to perform GPU accelerated computations with little effort. The result is a development model that makes the high performance of GPUs accessible to end users and working scientists. This talk will focus on the Python implementation of the interface, but will also briefly cover IDL, MATLAB and Java versions. We acknowledge Nathaniel Sizemore and Dave Wade-Stein for help with the build system for this project.

Dan Karipides, Paul Mullowney, Michael Galloy, Peter Messmer, Brian E. Granger
Tech-X Corporation
karipid@txcorp.com, paulm@txcorp.com,
mgalloy@txcorp.com, messmer@txcorp.com,
ellisonbg@gmail.com

CP4

Grids and Clusters with Multi-Core Nodes: A Genetics Application Perspective

The introduction of multicore processor implies that algorithms which are parallelized at an outer, coarse grain level should possibly be revisited to examine if multithreading should also be used at an inner, fine grain level. In this paper we discuss parallel versions of the tightly coupled global optimization algorithm DIRECT. We examine how both coarse grained and fine grained parallelism can be exploited using a hybrid programming model. We show that excellent performance can be archived when using the hybrid algorithm on loosely-coupled systems like clusters and grids with multicore nodes.

Henrik Löf
Stanford University
Department of Energy Resources Engineering
henrik.loef@stanford.edu

Mahen Jayawardena
Uppsala University
Department of Scientific Computing
mahen@it.uu.se

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

CP4

Group Locality Based Performance Analysis of Triplet Architecture A Static Direct Interconnection Network for Multi-Processor (mp-SoC)

We propose a new criterion in performance evaluation based on the concept of group locality in interconnection networks, the lower layer complete connect i.e., how completely a node in a subset of processing nodes is connected to its neighbors. Triplet Based Architecture, TriBA - a new idea in MP-SoC architectures is compared with three static interconnection networks from three orthogonal entities physical (chip area, dissipation), computational speed (message delay) and cost (chip yield, layout cost.)

Haroon-Ur-Rashi Khan, Shi Feng, Ji Wei Xing
School of Computer Science and Technology
Beijing Institute of Technology
haroon65@gmail.com, shifengyoujian@tom.com,
pass@bit.edu.cn

Kamran Kamran
Departemnt of Electrical Engineering
Univerity of Engg. & Tech., Lahore, Pakistan

kamran.uet@gmail.com

CP4

A Data-Distributed Massively Parallel Design of DIRECT

A data-distributed massively parallel implementation is developed for the optimization algorithm DIRECT, favored for its deterministic nature and global convergence property. Sharing data across multiple machines reduces the local memory burden. Multilevel parallelism boosts the concurrency and mitigates the data dependency, thus improving the load balancing and scalability. Also, user-level checkpointing is integrated as a fault-tolerance feature. On large-scale systems, the design was evaluated using benchmark functions and real-world applications.

Rhonda D. Phillips
Virginia Polytechnic Institute and State University
rdphllps@vt.edu

Layne T. Watson
Virginia Polytechnic Institute and State University
Departments of Computer Science and Mathematics
ltw@cs.vt.edu

Jian He
Department of Computer Science
Virginia Tech
jihe@vt.edu

Masha Sosonkina
Ames Laboratory/DOE
Iowa State University
masha@scl.ameslab.gov

CP5

Load Distribution in Madness

Load balancing is vital to the efficiency of MADNESS (Multiresolution Adaptive Numerical Environment for Scientific Simulation), an environment for prototyping and developing scientific applications being developed to run on leadership computing resources. We propose the melding algorithm to load balance the computational work in MADNESS. In this presentation, we describe the method, discuss its theoretical advantages over alternative load balancing techniques for this problem, and present preliminary results from runs on leadership computing resources.

Rebecca J. Hartman-Baker, George Fann
Oak Ridge National Laboratory
hartmanbakrj@ornl.gov, fanngi@ornl.gov

Robert Harrison
University of Tennessee
Oak Ridge National Laboratory
harrisonrj@ornl.gov

CP5

A Benchmark Study of Compiler Performance for Sparse Kernels on Multicore Processors

Obtaining optimal performance for scientific applications on modern computer architectures continues to be a challenge. This study presents an empirical comparison of the impact of hardware architecture, compilation options, data structure and coding technique on algorithm performance for a small set of representative mathematical kernels including sparse matrix-vector products on a set of multicore processor-based HPC platforms. Numerical results are presented, and implications for the optimization of numerical software codes are considered.

Wayne Joubert
U.S. Army Engineer Research and Development Center (ERDC)
Major Shared Resource Center (MSRC)
Wayne.D.Joubert@erdc.usace.army.mil

CP5

Database Components for Support of Computational Quality of Service for Scientific CCA Applications

While component-based design has proven helpful in managing the complexity of parallel scientific simulations, many challenges remain in selecting and configuring components during runtime to improve performance. This presentation introduces a new aspect of our infrastructure in computational quality of service (CQoS), namely database components that manage historical performance data and metadata. We illustrate their use in selecting appropriate parallel solver components.

Li Li
Argonne National Laboratory
likli@mcs.anl.gov

Boyana Norris
Argonne National Laboratory
Mathematics and Computer Science Division
norris@mcs.anl.gov

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

CP5

Computational Forces in the Linpack Benchmark

The efficiency of parallel algorithms can be explained as a balancing act between computational forces. These forces, also called computational intensities, are determined by the particular algorithm and the particular machine running the algorithm. For a timing formula describing the Linpack benchmark from Greer and Henry, we show that different machines follow different paths along a single efficiency surface.

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

CP5

Performance Comparison Between Square-to-Hemisphere and Cubed Sphere Projections of a Global Shallow-Water Model on a Toroidal Interconnect Architecture

Motivated by limited scalability issues encountered with the cubed sphere projection implemented in global shallow-water models over a toroidal interconnect, we propose a square-to-hemisphere projection. We argue that the square-to-hemisphere projection is superior in optimizing processor communication and decreasing complexity

of computational load balancing. We present a performance comparison for a numerical shallow-water model under both projections using a discrete Galerkin Runge-Kutta (DGRK) method on the IBM BlueGene/L system over 1024 nodes.

Marcus Waldman
Undergraduate, University of Colorado at Boulder
Student Research Assistant, NCAR
marcus.waldman@colorado.edu

Siddhartha Ghosh
CISL/NCAR
sghosh@ucar.edu

CP6

Why Column Pivoting Should Be Used for Performance

This talk shows new research for doing parallel dense linear algebra with implicit column pivoting to improve load balancing. After showing the performance on a cluster of workstations, we discuss heterogeneous clusters where the dynamic load balancing helps the most. We show how the same new idea can be applied to hybrid OpenMP/MPI problems and sparse problems. We also show how this research is being integrated into the latest Intel Math Kernel Librarys cluster products.

Greg Henry
Intel Corporation
greg.henry@intel.com

CP6

Block Householder Reduction of Sparse Matrices to Small Band Upper Triangular Form

Bidiagonalization can be accomplishing by accessing a sparse matrix A only to perform sparse matrix dense vector multiplications Ax and y^tA. Only a moderate number of leading rows and columns are eliminated. The computations Ax and y^tA are predominant, especially when x and y are too large to fit in cache memory. If the reduction is to bandwidth k, the multiplications can instead be AX and Y^TA, A sparse, X, Y dense with k columns. Blocking A gives further speedup. On a cache based architecture, the resulting algorithm is fast and stable. It adapts easily to multi-core architectures.

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

CP6

Divide and Conquer Eigenvalue Solver Parallelization

The Divide and Conquer algorithm is very great to be parallel by idea: division of a big task to smaller ones that can be solved in parallel. But in fact it is not so easy because small solutions should be merged in a big one and in addition they impact each other on solving stage. This work describes problems and their solutions that appeared in eigenvalue solver parallelization.

Alexander V. Kobotov
Intel Corp.; Institute of Computational Mathematics and Mathematical Geophysics SB RAS
alexander.v.kobotov@intel.com

CP6

Weighted Matrix Reordering and Parallel Banded Preconditioners for Non-Symmetric Linear Systems

With the emergence of petascale architectures, the role of preconditioning techniques that can scale well on large number of processors have become crucial. We present a reordering scheme that allows the extraction of a central dominant band that can be used as a preconditioner. Our results demonstrate excellent scalability and robustness for a large class of problems for which other black-box preconditioners, such as ILU and varieties, are poorly scalable.

Murat Manguoglu
Purdue University Department of Computer Science
mmanguog@cs.purdue.edu

Ahmed Sameh
Department of Computer Science
Purdue University
sameh@cs.purdue.edu

Mehmet Koyuturk
Case Western Reserve University
Depertment of Electrical Engineering and Computer Science
koyuturk@eecs.case.edu

Ananth Grama
Purdue University
Department of Computer Science
ayg@cs.purdue.edu

CP6

One World, One Matrix

We propose a new parallel algorithm, called Directed Transmission Method (DTM), to solve the sparse linear system whose coefficient matrix is symmetric-positive-definite (SPD). DTM is a fully scalable, asynchronous, distributed and continuous-time iterative algorithm, which is quite different from the traditional discrete-time iterative algorithms. It is proved to be convergent. DTM is able to be efficiently running on any kind of homogeneous or heterogeneous parallel computers, e.g. multicore and manycore microprocessors, SMP, clusters, supercomputers, grids, clouds and WWW. By means of DTM, we are capable of solving arbitrarily-large sparse SPD linear systems, as long as we have enough processors and memories. Furthermore, we may unite the supercomputers all over the world to solve an unprecedented, extremely large sparse linear system, and the dream of "One World, One Matrix" would come true at that time. Besides, DTM would be a persuasive benchmark to test the performance of the parallel computers, especially the supercomputers and the manycore microprocessors.

Huazhong Yang, Fei Wei
Department of Electronic Engineering
Tsinghua University, Beijing, China
yanghz@tsinghua.edu.cn, weifei00@mails.tsinghua.edu.cn

CP6

New Algorithms for Sparse Matrix Partitioning

We discuss how to partition a sparse matrix to reduce communication in parallel sparse matrix computations. We focus on sparse matrix-vector multiplication, which is an

important kernel in scientific computing. We consider two-dimensional distributions, and present a new algorithm based on vertex separators and nested dissection. Empirical results on real application matrices show our method is better than the traditional 1-d (row) distribution, and competitive with other 2-d distributions.

Erik G. Boman
Sandia National Labs, NM
Scalable Algorithms Dept.
egboman@sandia.gov

Michael Wolf
Univ. of Illinois, Urbana-Champaigne
mmwolf@uiuc.edu

CP7

High Performance Solution of Sparse Linear Systems Using Direct Methods with Application to Electromagnetic Problems

The numerical treatment of high frequency electromagnetic scattering in inhomogeneous media is very computationally intensive. For scattering, the electromagnetic field must be computed around and inside 3D complex bodies. Because of this, accurate numerical methods must be used to solve Maxwell's equations in the frequency domain, and it leads to solve very large linear systems. In order to solve these systems, we have combined on our TERAscale computer modern numerical methods with efficient parallel algorithms.

Katherine Mer-Nkonga, Michel Mandallena, Jean-Jacques Pesque, David Goudin
CEA/CESTA
katherine.nkonga@cea.fr, michel.mandallena@cea.fr, jean.pesque@cea.fr, david.goudin@cea.fr

CP7

Parallel Subspace Newton Methods for Algebraic Systems with Local High Nonlinearities

We present locally refined Newton type methods for large nonlinear systems of algebraic equations, arising from the discretization of nonlinear partial differential equations. We focus on the type of systems that have local high nonlinearities. In other words, the nonlinear system may have many equations, but only a small percentage of them are highly nonlinear compared to the rest of the equations. Global Newton methods may be used to solve the system, but often the computing time is wasted since all equations are treated equally as if they were all highly nonlinear. We introduce subspace Newton methods to remove the local high nonlinearities and therefore improve the efficiency and the effectiveness of the outer global Newton method, which performs well on equations with roughly the same level of nonlinearities. We prove the convergence of this new method under certain assumptions. We also discuss the parallel implementation of the new method using PETSc and provide some numerical results from solving several different nonlinear differential equations.

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

Xuefeng Li
Loyola University New Orleans
Li@Loyno.edu

CP7

Fully Coupled Two-Level Domain Decomposition Algorithms for Inverse Problems

In this talk, we discuss multilevel domain decomposition methods for solving some coupled nonlinear systems of equations obtained from the discretization of inverse problems. We focus on a fully coupled Newton-Krylov algorithm with two-level Schwarz type domain decomposition methods as the preconditioner. We study the parallel performance of the algorithms on supercomputers with hundreds of processors for solving some difficult inverse problems arising from the modeling of ground water flows.

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

Si Liu
Department of Applied Mathematics
University of Colorado, Boulder 80309-0526
sliu@colorado.edu

CP7

A Parallel Multigrid Preconditioner for High-Order and *hp*-Adaptive Finite Elements

The *hp* version of the finite element method is an adaptive finite element approach in which adaptivity occurs in both the size, h, of the elements and in the order, p, of the approximating piecewise polynomials. An optimal order parallel linear system solver is needed to get the best efficiency of these methods. We present a parallel multigrid preconditioner whose rate of convergence is independent of both h and p.

William F. Mitchell
National Institute of Standards and Technology
Mathematical and Computational Sciences Division
william.mitchell@nist.gov

CP7

Impact of Dual-Core Processors on the Performance of Parallel Krylov Subspace Linear Solvers and Preconditioners for Porous Media Flow Applications

Data from finite element modeling of porous media flow were used to solve linear systems of equations using 12 Krylov subspace parallel linear solvers with five preconditioners (60 scenarios) using PETSc to test for efficiency and accuracy of the different options. The Cray XT3 used in this study has been recently upgraded to 4160 dual core nodes. This presentation will highlight the performance of the linear solvers before and after the dual-core processors were installed.

Thomas Oppe
Engineer Research and Development Center
Waterways Experiment Station
thomas.c.oppe@erdc.usace.army.mil

Sharad Gavali
NASA Ames Research Center
gavali@nas.nasa.gov

Fred T. Tracy
Engineer Research and Development Center
Waterways Experiment Statiuon
fred.t.tracy@erdc.usace.army.mil

CP7

Multi-Length Scale Preconditioned Iterative Solver for Parallel Hybrid Quantum Monte Carlo Simulation

The hybrid quantum Monte Carlo (HQMC) method of the Hubbard model is a powerful method used to study the electron interactions that characterize the properties of materials, such as magnetism and superconductivity. The bottleneck of the method is on the repeated solutions of the underlying multi-length-scale linear systems of equations. In this talk, we present a preconditioning technique and its parallelization for solving the linear systems. The preconditioned solver demonstrates the optimal linear scaling complexity of the HQMC method for moderately-correlated materials.

Zaojun Bai
Department of Computer Science
University of California, Davis, USA
bai@cs.ucdavis.edu

Richard Scalettar
Department of Physics,
University of California, Davis, USA
scalettar@physics.ucdavis.edu

Wenbin Chen
School of Mathematical Science,
Fudan University, China
wbchen@fudan.edu.cn

Ichitaro Yamazaki
Department of Computer Science
University of California, Davis
yamazaki@cs.ucdavis.edu

CP8

A Parallel Algorithm for Optimization-Based Smoothing of Unstructured 3-D Meshes

Serial optimization-based smoothing algorithms are computationally expensive. Using Metis (or ParMetis) to partition the mesh, the parallel algorithm moves (or does not move) a processor's internal nodes based on a cost function derived from the Jacobians and condition numbers of surrounding elements. Ghost cells are used to communicate new positions, the lower processor on a boundary uses the new information to move boundary nodes, and the process repeats. The result is a ready-to-use decomposed mesh.

Vincent C. Betro
University of Tennessee at Chattanooga
vincent-betro@utc.edu

CP8

Distributed Transpose for 3D Fft: The Effects of Machine Geometry and Process Mapping on Blue Gene/L

We describe how to extend the scalability 3D-FFT using 2D-decomposition on thousands of BlueGene/L processors. The communication cost of carrying out the data transposes required by the 3D-FFT is very high and dominates the computation cost at the limits of scalability. This motivated us to focus on performance measurements of the distributed transpose alone. We report performance data on two communication protocols, MPI and BG/L-ADE. The proposed approach is effective in improving performance for Particle-Mesh-based N-body simulations.

T.J.C. Ward
IBM Software Group,
Hursley Park, Hursley, UK
tjcw@uk.ibm.com

Philiph Heidelberger
IBM Thomas J. Watson Research Center
Yorktown Heights, NY 10598-0218, USA
philiph@@us.ibm.com

Robert S. Germain, Blake Fitch, Aleksandr Rayshubskiy, Maria Eleftheriou
IBM Thomas J. Watson Research Center
rgermain@us.ibm.com, bgf@us.ibm.com, arayshu@us.ibm.com, mariae@us.ibm.com

CP8

New Parallel Techniques for Bvps in Ords

The main objective of this paper is the devlopment of a new parallel integration algorithms for solving boundary value problems (BVPs) in ordenary deffirential equations (ODEs). the idea of new techniques is combinning the parallel integration processes with parallel interpolation processes suitable for running on MIMD (Multiple instruction streams with multiple data streams) computing systems. The stablity of the developed algorithms are anylsed. We also studied the treatment of stiff BVPs by the devloped techniques.

Bashir M. Khalaf
Professor of Scientific Computing
bmskhalaf@yahoo.co.uk

CP8

Programming with Large Scale Edge-Node Simulator on BlueGene/L: A Case Study of 3D Fft

We designed a network simulator for rapid specification of complex networks, such as those required to model neural tissue. Here we demonstrate a more general use of the network simulator for implementing generic parallel algorithms, with a case study of the 3D-FFT. We demonstrate scaling of the 128x128x128 FFT network to 4,096 BG/L processors, and compare performance against the original algorithm (Eleftheriou et al, 2006). Strategies for automatically mapping network calculations to BG/L are discussed.

Robert S. Germain, Blake Fitch, Maria Eleftheriou
IBM Thomas J. Watson Research Center
rgermain@us.ibm.com, bgf@us.ibm.com,
mariae@us.ibm.com

James Kozloski
IBM TJ Watson Research Center
kozloski@us.ibm.com

Charles Peck
Biometaphorical Computing Research
IBM T.J. Watson Research Center
cpeck@us.ibm.com

CP8

Improving the Scalability of Adaptive Mesh Refinement

In many large scale adaptive simulations scalability is hindered due to costs associated with the changing mesh. Algorithmic improvements to the mesh changing processes have led to a significant reduction in these costs. In addition, the frequency of remeshing can be reduced through the use of dilation. These changes have led to large improvements in overall scalability of the Uintah simulation framework. Results up to 4096 processors will be shown.

Justin P. Luitjens, Tom Henderson
University of Utah
luitjens@cs.utah.edu, tch@cs.utah.edu

Martin Berzins
SCI Institute
University of Utah
mb@sci.utah.edu

CP8

Finite Element Assembly on Arbitrary Meshes

One goal of automating Finite Element Methods (FEM) is to allow arbitrary element types and orders on arbitrary meshes. A challenge to this goal is separating local element definitions from the mesh definition. We show our conceptual paradigm for this separation using the PETSc Sieve library, a library based on representing meshes as Grothendieck topologies, and demonstrate results with a grade-2 fluid application.

Andy R. Terrel
University of Chicago
Department of Computer Science
aterrel@uchicago.edu

Matthew G. Knepley
Argonne National Laboratory
knepley@mcs.anl.gov

MS1

Towards General Auto-tuning Description Language on Advanced Computing Systems

The description of auto-tuning is crucial, but time-consuming work for developing numerical libraries with auto-tuning facility. In this presentation, a description language for auto-tuning, named ABCLibScript, is explained with several examples of numerical computation. Although the target of ABCLibScript was vector supercomputers, but we show the effectiveness on it to software development process on embedded systems. The effect on the advanced computer environment, which is supercomputer with multi-core processor, will be also shown.

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

MS1

Proposal of Run-time Parameter Auto-Tuning Approach for Restarted Lanczos Method

Many input parameters in matrix solvers are difficult to predict the best values before runtime. This paper proposes an automatic tuning approach for the restarted Lanczos method, which explores the best projection matrix size from the history of residual value at runtime. The numerical experiments show the proposed approach is 100 times faster than the original method in the best case. The result implies the runtime automatic tuning is effective for iterative matrix solvers.

Takao Sakurai, Ken Naono, Masashi Egi
Central Research Laboratory
Hitachi Ltd.
takao.sakurai.ju@hitachi.com, ken.naono.aw@hitachi.com, masashi.egi.zj@hitachi.com

Mitsuyoshi Igai, Hiroyuki Kidachi
Hitachi ULSI Systems Corporation
mitsuyoshi.igai.bf@hitachi.com, hiroyuki.kidachi.yd@hitachi.com

MS1

A Bayesian Approach to Automatic Performance Tuning

Code tuning has been done based on models, experiments or their combinations, but the combinations are mostly of heuristics. In this talk it is shown that Bayesian statistics can provide a convenient mathematical framework to combining model and experiments for code tuning. The example problem here is online selection of several unrolled codes for matrix-matrix multiply, and some sequential experimental designs based on a simple performance model are proposed and evaluated.

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

MS1

Automatic Tuning for Parallel FFTs

In this talk, an automatic performance tuning method for parallel fast Fourier transforms (FFTs) is presented. A 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 FFTs on a PC cluster are reported.

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

MS2

Neutral Territory Methods for Efficient Parallelization of Molecular Dynamics Simulations

The majority of the computational workload in molecular dynamics simulations involves interactions between nearby particles. We will describe a class of algorithms for parallelization of range-limited particle interactions, the neutral territory methods, some of which confer significant practical advantages over traditional parallelization algorithms. We will illustrate specific neutral territory methods introduced by other researchers and by ourselves, and we will

discuss the tradeoffs that led us to select different neutral territory methods for different molecular dynamics implementations.

Ron O. Dror
D. E. Shaw Research
dror@deshaw.com

David E. Shaw
D. E. Shaw Research
Columbia University
david@deshaw.com

Kevin J. Bowers
D. E. Shaw Research
bowersk@deshaw.com

MS2

Scaling NAMD to Large Parallel Machines

NAMD's parallel design, circa 1996, has stood the test of time. The basic parallel structure includes (a) decomposition into cells, and force-computation objects for each pair of interacting cells, (b) implementation using message-driven objects in Charm++, and (c) assignment of objects to processors using measurement-based load balancers that also reduce communication. This talk will review recent optimizations to scale NAMD to over 32,000 processors for small and large biomolecular systems.

Laxmikant V. Kale
University of Illinois at Urbana-Champaign
kale@cs.uiuc.edu

James C. Phillips
Beckman Institute, U. Illinois at Urbana-Champaign
jim@ks.uiuc.edu

Chao Mei
University of Illinois at Urbana-Champaign
chaomei2@ks.uiuc.edu

Abhinav Bhatele, Gengbin Zheng, Sameer Kumar
Beckman Institute, U. Illinois at Urbana-Champaign
bhatele2@uiuc.edu, gzheng@ks.uiuc.edu,
sameer@ks.uiuc.edu

Klaus Schulten
University of Illinios at Urbana Champaign
kschulte@ks.uiuc.edu

MS2

Nanoparticle and Colloidal Simulations with Molecular Dynamics

Modeling nanoparticle or colloidal systems in a molecular dynamics (MD) code requires coarse-graining on several levels to achieve meaningful simulation times for study of rheological and other manufacturing properties. These include treating colloids as single particles, moving from explicit to implicit solvent, and capturing hydrodynamic effects. These changes also impact parallel algorithms for tasks such as finding neighbor particles and interprocessor communication. I'll describe enhancements we've made to our MD code LAMMPS to make nanoparticle simulations more efficient, highlighting its flexible design that has enabled the new capabilities.

Steve Plimpton
Sandia National Laboratories
sjplimp@sandia.gov

MS2

A Summary of the Performance and Scaling of AMBER 10 and the Challenges Ahead

This talk will present a summary of the current performance and scaling of the soon to be released version 10 of the AMBER software on a range of NSF and DOE high performance computing systems. In addition it will include an overview of the supported methods and the approaches used to obtain the level of performance seen. Finally some of the challenges that may face the molecular dynamics community in the near future will be discussed.

Ross C. Walker
San Diego Supercomputer Center
rcw@sdsc.edu

Robert E. Duke
NIEHS and UNC-CHapel Hill
rduke@email.unc.edu

David A. Case
The Scripps Research Institute
case@scripps.edu

MS3

Creating Interoperability for Parallel Meshing Tools

Mesh technology, such as mesh generation, database queries, and adaptivity, plays a critical role in scientific simulations. While many frameworks providing mesh technology exist, their incorporation into applications requires significant effort and learning by application developers. Interfaces allowing interoperable use of mesh tools greatly simplify this process while providing a wider range of technology than a single framework. In this talk, we discuss interoperable mesh interfaces and, in particular, their extension to parallel mesh services.

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

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

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

Tim Tautges
Argonne National Laboratory
tautges@mcs.anl.gov

Carl Ollivier-Gooch
University of British Columbia
cfog@mech.ubc.ca

Vitus Leung

Sandia National Laboratories
vjleung@sandia.gov

MS3

Algorithms for Parallel Mesh Smoothing Using Mesquite

We discuss the development of an infrastructure that supports the use of Mesquite mesh quality improvement algorithms in distributed memory applications. We start with the application's decomposition of the mesh data and use an iterative process to select independent sets of vertices to resposition in each pass. We experiment with a mix of local and global techniques from Mesquite and report on the scalability and performance or our methods.

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

Martin Isenburg
Lawrence Livermore National Lab
isenburg1@llnl.gov

MS3

Zoltan Load Balancing Approaches

Dynamic load-balancing is a data-management service that is critical to a wide range of unstructured and/or adaptive parallel applications. The Zoltan Library provides a suite of dynamic load-balancing tools. Access to Zoltan is now available through a common interface that supports interoperability within the ITAPS data model. In this presentation, we give a brief overview of the dynamic load-balancing approaches available through Zoltan's ITAPS interface.

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

MS3

A Partition Model for Massively Parallel Mesh-Based Computations

The Interoperable Technologies for Advanced Petascale Simulations DOE SciDAC center is designing and implementing an interoperable partition model to support parallel mesh-based operations including adaptive computations accounting for the complexities that arise due to the changing computational load and communications of adapted meshes. The presentation will first discuss the overall partition model design. Consideration will then be given to its implemented and relation to adaptive mesh control and Zoltan load balancing procedures.

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

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

Kenneth Jansen, TIng Xie
Rensselaer Polytechnic Institute
jansen@scorec.rpi.edu, txie@scorec.rpi.edu

MS4

Issues in Exploiting the Power of Multiple Methods

In this presentation, we will discuss some of the issues in multimethod implementation. While the focus of using multimethods is mapping a "single" method to a simulation stage, for certain problems,several suitable methods might be combined to produce more effective results. The trade-off related to the frequency of changing methods is another issue, as it might not be practical to switch methods at every opportunity for adaptivity. Yet another challenge is the efficient identification of adaptivity in the simulation.

Sanjukta Bhowmick
Department of Computer Science and Engineering
Pennsylvania State University
bhowmick@cse.psu.edu

MS4

Machine Learning Support for Numerical Decision Making

We present the SALSA (Self-Adapting Large-scale Solver Architecture) software system for intelligent multimethods. The system is based on a modular architecture for composite algorithms (for instance, choice of scaling/preconditioner/iterator in iterative linear system solvers) and uses machine learning techniques for adaptively choosing the component algorithms. We will discuss various learning techniques we have explored, and the high level of accuracy obtained.

Victor Eijkhout
The University of Texas at Austin
Texas Advanced Computing Center
eijkhout@tacc.utexas.edu

MS4

Evaluation of a Meta-partitioner for Simulations Using Block-structured Adaptively Refined Meshes

High parallel efficiency for structured adaptive mesh refinement (SAMR) applications requires repeated data partitioning and distribution. We present a performance evaluation of a framework for adaptive partitioning. Considering computational load, communication volume, synchronization delays, and data movement, the framework selects, configures and invokes the most efficient partitioning algorithm. We show that adaptive partitioning can significantly improve parallel efficiency for SAMR applications.

Henrik Johansson
Department of Information Technology
Uppsala University
henrikj@it.uu.se

MS4

Adaptive Partitioning for Unstructured AMR Applications

Improving performance of large scientific adaptive applications is non-trivial due to their inherent dynamics and wide spectrum of properties. Performance is limited by the partitioner's ability exploit computer resources given the application state. No single partitioning configuration can generally achieve high performance; partitioning must be dynamically adaptive. In this talk, we describe the meta-

partitioner: a framework for selecting and configuring the most suitable partitioner based on run-time state.

Johan Steensland
Sandia National Laboratories
jsteens@sandia.gov

MS5
Parallel Programming in MATLAB: Best Practices

Matlab is one of the most commonly used languages for scientific computing with approximately one million users worldwide. The Lincoln pMatlab library (http://www.ll.mit.edu/pMatlab), The Mathworks DCT, and StarP from ISC have brought parallel computing to the this community using the distributed array programming paradigm. This talk provides an introduction to distributed array programming and will describe the best programming practices for using distributed arrays to produce well performing parallel Matlab programs.

Jeremy Kepner
MIT Lincoln Laboratory
kepner@ll.mit.edu

MS5
Parallel MATLAB in Production Supercomputing with Applications in Signal and Image Processing

Parallel MATLAB enables the large community of MATLAB users to harness the increased computing capacity and memory of distributed memory clusters. At the Ohio Supercomputer Center we provide our users with three varieties of Parallel MATLAB. In this talk, we will describe how we run these Parallel MATLAB environments within a traditional batch oriented queuing system. We will also describe our experiences in developing three signal and image processing applications within this environment.

Ashok Krishnamurthy
Ohio Supercomputing Center
ashok@osc.edu

David Hudak, John Nehrbass, Siddharth Samsi, Vijay Gadepally
Ohio Supercomputer Center
dhudak@osc.edu, nehrbass@osc.edu, samsi@osc.edu, vijayg@osc.edu

MS5
Parallel Computing Toolbox (PCT) and Parallel Programming in MATLAB

Parallel Computing Toolbox addresses computationally and data-intensive problems using MATLAB and Simulink in a multiprocessor computing environment. The toolbox allows both several independent tasks or a single parallel computation by harnessing computing clusters and a variety of batch queuing software implementation. The toolbox provides high-level constructs, such as parallel loops and algorithms, and MPI-based functions. Also, low-level constructs for resource management are included. The Parallel Command Window provides interactive environment for developing parallel applications.

Piotr Luszczek
The MathWorks, Inc.
luszczek@cs.utk.edu

MS5
Interactive Data Exploration with Star-P

High performance applications increasingly combine numerical and combinatorial algorithms. Past research on high performance computation has focused mainly on numerical algorithms, and there is a rich variety of tools for high performance numerical computing. On the other hand, few tools exist for large scale combinatorial computing. We describe our efforts to build a common infrastucture for numerical and combinatorial computing by using parallel sparse matrices to implement parallel graph algorithms.

Viral B. Shah
Interactive Supercomputing
vshah@interactivesupercomputing.com

MS6
Integrated Air/Ocean/Wave Modeling Using ESMF

Development of an integrated air/ocean/wave modeling system is described. The single executable system is built from mature stand-alone models using the Earth System Modeling Framework (ESMF). The framework provides the required functionality for treating each model as a separate component and for the redistribution and remapping of data between them. An exchange grid approach is implemented to simplify the interface between models that use telescoping nests. In addition to describing the implementation details, preliminary results will be presented for two regional test cases.

Sue Chen
Naval Research Laboratory
Monterey, CA
sue.chen@nrlmry.navy.mil

Hao Jin
SAIC, Naval Research Laboratory
Monterey, CA
hao.jin@nrlmry.navy.mil

Rich Hodur
Naval Research Laboratory
Monterey, CA
rich.hodur@nrlmry.navy.mil

Sasa Gabersek
UCAR, Naval Research Laboratory
Monterey, CA
sasa.gabersek@nrlmry.navy.mil

Tim Campbell
Naval Research Laboratory
Stennis Space Center
Tim.Campbell@nrlssc.navy.mil

MS6
Algorithms for a Scalable Earth System Model

Abstract not available at time of publication.

John Drake
Oak Ridge National Laboratory
drakejb@ornl.gov

MS6

A Coupled Watershed-Nearshore Model Using DBuilder

Coupling of independent models involves implementation of synchronization and data-exchanging algorithms. Also, coupling may be along a shared edge of two meshes or an overlapped region between two meshes. The latter can be difficult in terms of spatially mapping nodes/elements between two meshes. DBuilder, a parallel data management toolkit, provides users with APIs such as element searching and data synchronization routines to accomplish these tasks.

Robert M. Hunter
U.S. Army Engineer Research & Development Center
robert.m.hunter@erdc.usace.army.mil

MS6

Parallel Rendezvous Regridding in ESMF

In coupled multiphysics simulations often each physics is modelled by a distinct, specialized code; to combine these codes into a coupled solver, it is necessary to transfer fields from one code to another (often called *regridding*). In the Earth Sciences (and other disciplines) each individual physics code will likely be a massively parallel code, with a unique parallel decomposition of the physical domain. We discuss the ESMF implementation of the Parallel Rendezvous algorithm of Stewart et al, which creates a geometric rendezvous mesh to perform the search and interpolation. We discuss the application and extension of this algorithm to interpolation of high order finite elements with non-nodal interpolation rules (e.g. Hierarchical elements). We also demonstrate a smoothing interpolation method that is based on finite element patch recovery techniques.

David Neckels
National Center for Atmospheric Research
dneckels@ucar.edu

MS7

Solving Rank Deficient Linear-Least Squares Problems Using Sparse QR Factorizations

We address the problem of solving linear least-squares problems $min||Ax - b||$ when A is a sparse m-by-n rank deficient or highly ill-conditioned matrix. Since A is rank-deficient or highly ill-conditioned the factorization $A = QR$ is not useful because the computed R is ill-conditioned. We have developed a new method that uses a regular QR factorization instead of a rank-revealing QR factorization. The goal of this work is to implement and test the algorithm in an high performance QR factorization.

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

Haim Avron
School of Computer Science
Tel-Aviv University
haima@tau.ac.il

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

MS7

Computing the Conditioning of Dense Linear Least Squares with (Sca)LAPACK

We define condition numbers that can assess the accuracy of the components of the least squares solution. We interpret them in terms of statistical quantities. We show that the ratio of the variance of one component of the solution by the variance of the right-hand side is exactly the condition number. We also propose codes based on (Sca)LAPACK for computing the variance-covariance matrix. Finally we present experiments from the space industry with real physical data.

Julien Langou
University of Colorado at Denver and Health Sciences Center
julien.langou@cudenver.edu

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

Marc Baboulin
CERFACS
baboulin@cerfacs.fr

Serge Gratton
CNES
serge.gratton@cnes.fr

MS7

Sparse QR Rank-revealing Factorization

We discuss an algorithm for computing a rank revealing sparse QR factorization. First, a QR factorization with no pivoting is performed, that allows to obtain efficiently a sparse triangular factor R. Second, an incremental condition number estimator (ICE) is used iteratively on R to identify redundant columns. We also introduce a block formulation of ICE algorithm. Numerical tests show that block ICE leads to approximations close to those obtained by successive runs of ICE.

Bernard Philippe
IRISA-INRIA
Rennes France
philippe@irisa.fr

Laura Grigori
INRIA
France
Laura.Grigori@inria.fr

Frederic Guyomarch
IRISA
France
Frederic.Guyomarch@irisa.fr

MS7

Blocked Bidiagonal Reduction of Sparse Matrices Using Givens Rotations

Computing the singular value decomposition involves bidiagonalization of a sparse upper triangular matrix R. Conventional methods do not exploit the sparsity of R. We introduce a method to bidiagonalize R using a sequence of Givens rotations while preserving the "mountainview" pro-

file of R. A dynamic blocking scheme extends the method to two blocked variations which generate no more fill than the unblocked version. We present performance results comparing all the different methods.

Gene H. Golub
Stanford University
Department of Computer Science
golub@sccm.stanford.edu

Timothy A. Davis
University of Florida
Computer and Information Science and Engineering
davis@cise.ufl.edu

Sivasankaran Rajamanickam
Dept. of Computer and Information Science and Engineering
Univ. of Florida, USA
srajaman@cise.ufl.edu

MS8

NVIDIA CUDA Software and GPU Parallel Computing Architecture

In the past, graphics processors were special purpose hardwired application accelerators, suitable only for conventional rasterization-style graphics applications. Modern GPUs are now fully programmable, massively parallel floating point processors. This talk will describe NVIDIAs massively multithreaded computing architecture and CUDA software for GPU computing. The architecture is a scalable, highly parallel architecture that delivers high throughput for data-intensive processing. Although not truly general-purpose processors, GPUs can now be used for a wide variety of compute-intensive applications beyond graphics.

Michael Garland
NVIDIA
mgarland@nvidia.com

MS8

Implementation of the Navier-Stokes Stanford University Solver (NSSUS) on a GPU

Current graphics processing units are capable of over 300 Gflops peak performance and this typically doubles every year. We have ported some of the capabilities of the Navier-Stokes Stanford University Solver (NSSUS), a multi-block structured code with a provably stable and accurate numerical discretization which uses a vertex-based finite-difference method and multigrid for convergence acceleration. Speed-ups of over 40x were demonstrated for simple test geometries and up to 20x for realistic geometries of engineering interest.

Patrick LeGresley
Stanford University, now with NVIDIA
plegresley@nvidia.com

Erich Elsen
Stanford University
eelsen@stanford.edu

Eric F. Darve
Stanford University
Center for Turbulence Research
darve@stanford.edu

MS8

Performance and Productivity of Graphics Processing Units for a Quantum Monte Carlo Application

The increased programmability and performance of Graphics Processing Units (GPUs) can have profound positive impact on developer productivity. In this talk, we discuss the acceleration of a Quantum Monte Carlo application using GPUs. Topics include the impact of GPU features on performance, tradeoffs in using a library approach versus a more informed hand optimized acceleration path, and the feasibility of combining these approaches.

Jeremy Meredith
Oak Ridge National Laboratory
jsmeredith@ornl.gov

MS8

Accelerating Molecular Modeling Applications with Graphics Processors

State-of-the-art graphics processing units (GPUs) can perform over 500 billion arithmetic operations per second, a powerful computational resource that can now be harnessed for use by scientific applications. We present an overview of recent advances in programmable GPUs, with an emphasis on their application to biomolecular modeling applications and the programming techniques required to obtain optimal performance in these cases. Performance and implementation details are presented for several applications. The calculations include runs on multiple GPUs.

John Stone, James C. Phillips, Peter Freddolino,
Leonardo Trabuco
Beckman Institute
Univ of Illinois at Urbana-Champaign
johns@ks.uiuc.edu, jim@ks.uiuc.edu,
petefred@ks.uiuc.edu, ltrabuco@ks.uiuc.edu

Klaus Schulten
University of Illinios at Urbana Champaign
kschulte@ks.uiuc.edu

David J. Hardy
Theoretical Biophysics Group, Beckman Institute
University of Illinois at Urbana-Champaign
dhardy@ks.uiuc.edu

MS9

Communication Avoiding Algorithms for Linear Algebra: Motivation, Approach

We survey results to be presented in this minisymposium on designing numerical algorithms to minimize the largest cost component: communication. This could be bandwidth and latency costs between processors over a network, or between levels of a memory hierarchy; both costs are increasing exponentially compared to floating point. We describe novel algorithms in sparse and dense linear algebra, for both direct methods (like QR and LU) and iterative methods that can minimize communication.

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

MS9
Communication Avoiding Gaussian Elimination

We present CALU, a Communication Avoiding LU factorization algorithm for dense matrices distributed in a two-dimensional (2D) cyclic layout. The new algorithm leads to an important decrease in the number of messages exchanged during the factorization, and thus it overcomes the latency bottleneck of the LU factorization as implemented in ScaLAPACK. We also discuss the stability of the pivoting strategy used in CALU and present performance results on several computational platforms.

Julien Langou
University of Colorado at Denver and Health Sciences Center
julien.langou@cudenver.edu

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

Laura Grigori
INRIA
France
Laura.Grigori@inria.fr

Hua Xiang
INRIA, France
Hua.Xiang@inria.fr

MS9
AllReduce Algorithms: Application to Householder QR Factorization

QR factorizations of tall and skinny matrices with their data partitioned vertically across several processors arise in a wide range of applications. Various methods exist to perform the QR factorization of such matrices: Gram-Schmidt, Householder, or CholeskyQR. In this talk, We present the Allreduce Householder QR factorization. This method is stable and performs, in our experiments, from four to eight times faster than ScaLAPACK routines on tall and skinny matrices. The idea of Allreduce algorithms can be extended to 2D block-cyclic LU or QR factorization.

Julien Langou
University of Colorado at Denver and Health Sciences Center
julien.langou@cudenver.edu

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

Laura Grigori
INRIA
France
Laura.Grigori@inria.fr

Mark Hoemmen
UC Berkeley, USA
mhoemmen@cs.berkeley.edu

MS9
A Low Latency Approach for Parallel Sparse LU Factorization

We present a new scheme for computing the sparse LU factorization. Our goal is to decrease the number of messages, hence decreasing the time spent in communication. The reduction in the number of messages is obtained by using a heuristic pivoting strategy, which is shown by numerical experiments to be stable in practice. The parallel algorithm is based on a hypergraph reordering strategy, and an associated separator tree is used for distributing the data.

Laura Grigori
INRIA
France
Laura.Grigori@inria.fr

Hua Xiang
INRIA, France
Hua.Xiang@inria.fr

MS10
Anton: A Special-purpose Machine for Molecular Dynamics Simulation

Anton is a massively parallel machine which should make practical millisecond-scale classical molecular dynamics (MD) simulations of proteins in explicit solvent. The machine, which is scheduled for completion by the end of 2008, is based on 512 identical MD-specific ASICs that interact in a tightly coupled manner using a specialized high-speed communication network. Anton has been designed to use both novel parallel algorithms and special-purpose logic to dramatically accelerate those calculations that dominate the time required for a typical MD simulation. The remainder of the simulation algorithm is executed by a programmable portion of each chip that achieves a substantial degree of parallelism while preserving the flexibility necessary to accommodate anticipated advances in physical models and simulation methods.

Ron O. Dror
D. E. Shaw Research
dror@deshaw.com

David E. Shaw
D. E. Shaw Research
Columbia University
david@deshaw.com

Martin M. Deneroff, Jeffrey S. Kuskin, Richard H. Larson, John K. Salmon, Cliff Young
D. E. Shaw Research
deneroff@deshaw.com, kuskinj@deshaw.com, larsonr@deshaw.com, salmonj@deshaw.com, youngc@deshaw.com

MS10
Scaling Classical Molecular Dynamics to O(1) Atom per Node

We will describe some of the issues involved in scaling biomolecular simulations onto massively parallel machines as well as some of the science that we have been able to achieve using the Blue Matter molecular simulation application on Blue Gene/L. Our experiences in scaling to order one atom/node on BG/L should provide some insights into the challenges involved in scaling biomolecular simulations

onto larger peta-scale platforms.

Blake G. Fitch, Christopher Ward, Michael C. Pitman, Robert S. Germain
IBM. T.J. Watson Research Center
bgf@us.ibm.com, tjcw@uk.ibm.com, pitman@us.ibm.com, rgermain@us.ibm.com

Aleksandr Rayshubskiy, Maria Eleftheriou
IBM Thomas J. Watson Research Center
arayshu@us.ibm.com, mariae@us.ibm.com

MS10
Accelerating NAMD with Graphics Processors

Commodity graphics processors allow a single workstation to achieve teraflop performance on certain workloads. Working with the NVIDIA CUDA programming system, we have adapted our molecular dynamics code NAMD (www.ks.uiuc.edu/Research/namd/) to offload the most expensive calculations to graphics processors while maintaining its parallel capability (J. Comp. Chem., 28:2618-2640, 2007). This talk will present recent work and parallel performance results for CUDA-accelerated NAMD.

John E. Stone, James C. Phillips
Beckman Institute, U. Illinois at Urbana-Champaign
johns@ks.uiuc.edu, jim@ks.uiuc.edu

Klaus Schulten
University of Illinios at Urbana Champaign
kschulte@ks.uiuc.edu

MS10
Petascale Special-Purpose Computer for Molecular Dynamics Simulations: MDGRAPE-3 and Beyond

We have developed the MDGRAPE-3 system, a petaflops special-purpose computer for molecular dynamics simulations. The MDGRAPE-3 is a PC cluster equipped with accelerators of 4,778 ASICs that calculate nonbonded interactions between atoms. Currently serial Amber-8 and in-house parallel MD software has been ported for the system. We will present the architecture and performance of the system as well as the next-generation project to develop a tile processor with special-purpose engines over TFLOPS performance.

Tetsu Narumi
Genomic Sciences Center, RIKEN
and Keio University
narumi@a7.keio.jp

Duraid Madeina
University of Tokyo
madeina@u-tokyo.ac.jp

Makoto Taiji, Yosuke Ohno
Genomic Sciences Center, RIKEN
taiji@gsc.riken.jp, ohno@gsc.riken.jp

Takashi Ikegami
The Graduate School of Arts and Sciences
University of Tokyo
ikeg@sacral.c.u-tokyo.ac.jp

MS11
Irregular Algorithms on the Cell Broadband Engine

The Sony-Toshiba-IBM Cell Broadband Engine is a heterogeneous multicore architecture that consists of a traditional microprocessor (PPE) with eight SIMD co-processing units (SPEs) integrated on-chip. Noting that while the Cell processor is architected for multimedia applications with regular processing requirements, we are interested in its performance on problems with non-uniform memory access patterns. In this talk, we present a case study of list ranking, a fundamental kernel for graph problems, that illustrates the design and implementation of parallel combinatorial algorithms on Cell. List ranking is a particularly challenging problem to parallelize on current cache-based and distributed memory architectures due to its low computational intensity and irregular memory access patterns. To tolerate memory latency on the Cell processor, we decompose work into several independent tasks and coordinate computation using the novel idea of Software-Managed threads (SM-Threads). We apply this generic SPE work-partitioning technique to efficiently implement list ranking, and demonstrate substantial speedup in comparison to traditional cache-based microprocessors. For instance, on a 3.2 GHz IBM QS20 Cell blade, for a random linked list of 1 million nodes, we achieve an overall speedup of 8.34 over a PPE-only implementation.

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

MS11
Linear Algebra Algorithms on the IBM Cell

This talk describes the design concepts behind implementations of some linear algebra routines targeted for the Cell processor and multicore in general. It describes in detail the implementation of code to solve linear system of equations using Gaussian elimination in single precision with iterative refinement of the solution to the full double precision accuracy. We will also look at the PlayStation 3 for use in scientific computations.

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

MS11
The Implementation of FFTW on Cell

FFTW is a library for computing Fourier transforms of complex, real, and real-symmetric multi-dimensional sequences. In this talk, I describe the port of FFTW to the Cell Broadband Engine, which was completed at the beginning of 2007 by the IBM Austin Research Lab. The bulk of FFTW runs on the Cell PPE, treating the SPE's as accelerators. The SPE's execute a specialized program capable of executing one-dimensional DFT's of two-dimensional vectors. While the capabilities of the SPE program are restricted, FFTW can reduce an arbitrary multi-dimensional DFT to this restricted form, thus taking advantage of the SPE's in most cases.

Matteo Frigo
Cilk Arts
athena@fftw.org

MS11
Dealing with the Memory Bandwidth Bottleneck on the Cell Processor

The computational workload on the Cell processor is handled by co-processors called SPEs. They have small local stores, which makes it necessary to store the application data in main memory. In many applications, especially those requiring O(n) computational effort, the bandwidth to main memory limits application performance. We will discuss the effectiveness of data compression in dealing with this limitation in a few important applications, such as matrix-vector multiplication.

Ashok Srinivasan
Department of Computer Science,
Florida State University
asriniva@cs.fsu.edu

Gunaranjan Gunaranjan, T Nagaraju, Ramprasad Ramprasad, T.V. Sivakumar
Sri Sathya Sai University
India
gunaranjan.84@gmail.com, t.nagaraj@gmail.com, ramprasad85@yahoo.com, tvsivakumar@gmail.com

MS12
Fluid Dynamics Simulations on Massively Parallel Computers

To achieve the goal of reliable flow simulations for realistic problems requires methods that are extensible to levels of parallelism that scale on 100,000s of processors and that can attain petaflop performance. We present a framework to perform massively parallel simulations where work is partitioned into balanced parts with well-controlled communications. We demonstrate scalability on 30,000 processors on an IBM BlueGene/L for the case of blood flow in a patient-specific arterial system.

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

Ken Jansen
Rensellaer Polytechnic Institute
kjansen@scorec.rpi.edu

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

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

MS12
Load Balancing in FronTier/AMR Computation

Computation of fluid physics with dynamically evolving fronts embedded in the PDE solution domain poses a great challenge for load balancing in parallel computing, particularly when adaptive mesh refinement is used. In this presentation, we present our adjusted rectangular domain decomposition algorithm and AMR patch redistribution for computation on a parallel platform with large number of processors. These algorithms enhances the parallel efficiency and scaling substantially. Our implementation conforms with the ITAPS interoperability requirements.

Ryan Kaufman, Brian Fix
SUNY at Stony Brook
rkaufman@ams.sunysb.edu, bfix@ams.sunysb.edu

Xiaolin Li
Department of Applied Math and Stat
SUNY at Stony Brook
linli@ams.sunysb.edu

MS12
Sets and Tags in the ITAPS Data Model

The data model used in the ITAPS iMesh and iGeom interfaces includes sets (an arbirary collection of entities and other sets) and tags (application-defined data assigned to entities, sets, and the interface itself). The combination of sets and tags is a powerful mechanism for embedding data with a variety of sources and semantics in the ITAPS interfaces. However, in practice this same variety makes it difficult to find data through such an abstract interface. This issue is discussed in the context of parallel scientific computing for both geometry and mesh data.

Tim Tautges
Argonne National Laboratory
tautges@mcs.anl.gov

Mark Miller
Lawrence Livermore National Laboratory
miller86@llnl.gov

MS12
Parallel, Scalable Unstructured Mesh Generation and Computation Physics Tools on Petascale Computing Architectures

As a scientific mesh based modeling community we are making steady progress toward petascale computing hardware. The computing hardware that we will be dealing with over the next several years is going to get more complex for mesh based algorithms in terms of multi/many-core processors, hierarchal memories and distributed I/O because the relationships between CPU speed, memory bandwidth, memory latency and communication latency are going to change dramatically. We need to make sure that our modeling and simulation tools keep pace with these hardware developments, such that the algorithms are scalable to 100s of thousands of processors and data is partitioned properly with respect to memory hierarchies. This presentation will describe our efforts to maintain parallel, scalable, efficient software tools and technology for mesh generation and continuum/discrete simulation on advanced computing architectures by making use interoperable tools such as the ITAPS mesh/field interfaces; Zoltan data/graph partitioning tools; and complex mesh generation and solvers, like NWGrid/NWPhys and FronTier on applications such as multiscale subsurface transport in complex geometries.

Yilin Fang, Harold E. Trease
Pacific Northwest National Laboratory
yilin.fang@pnl.gov, het@pnl.gov

Bruce Palmer
Pacific Northwest National Lab

bruce.palmer@pnl.gov

MS13

Challenges and Achievements in Computational Electromagnetics

Under consideration are problems in the vicinity of existing and future high current and high brightness particle accelerators such as high power proton drivers and 4th generation light sources i.e. x-ray free electron lasers. One can distinguish between two broad classes of problems in this field: real or complex eigenvalue problems in the context of cavity designs and relativistic particle tracking in 3D time dependent electromagnetic fields from Maxwells equation. Here the source terms, i.e. the time dependent charge distribution must be explicit modeled with high accuracy. Also of great importance is the efficient handling of datasets in the terra byte region, our HDF5 based Ansatz will be drafted. Our workhorses are a massive parallel particle in cell code and a finite element based eigenmode solver. I will talk about our implementations and show some results. Ongoing projects are time dependent hp-finite element based particle codes; here I will sketch our ideas.

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

MS13

Parallel Particle-In-Cell (PIC) Simulation on Hybrid Meshes

Particle-In-Cell (PIC) codes have become an essential tool for the numerical simulation of many physical phenomena involving charged particles, in particular beam physics, space and laboratory plasmas including fusion plasmas. Genuinely kinetic phenomena can be modeled by the Vlasov-Maxwell equations which are discretized by a PIC method coupled to a Maxwell field solver. Todays and future massively parallel supercomputers allow to envision the simulation of realistic problems involving complex geometries and multiple scales. To achieve this efficiently we propose to couple a Finite Element Maxwell solver with particles on hybrid grids with several homogeneous zones having their own structured or unstructured mesh type and size. This allows in particular fast particle tracking in zones having a structured mesh, but needs a fine analysis of load balancing issues for efficient parallelization. Our latest progress towards this goal will be presented.

Latu Guillaume
Strasbourg University
atu@labri.fr

MS13

Parallel Smoothed Aggregation Multigrid for Large Scale Electromagnetic Simulations

We present a new AMG preconditioner for linear systems arising from edge element discretization of the eddy current equations. The linear system is implicitly transformed into a 2×2 block system whose diagonal blocks are an edge Hodge Laplacian and a nodal scalar Laplacian, respectively. Solving the edge Hodge Laplacian involves matrix-free smoothing and a specialized restriction to a coarse nodal problem. We present three-dimensional computational results on twenty thousand Cray XT3 processors.

Pavel Bochev
Sandia National Laboratories
Computational Math and Algorithms
pbboche@sandia.gov

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

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

Chris Siefert
Sandia National Laboratories
csiefer@sandia.gov

MS13

Parallel Auxiliary Space AMG for Maxwell Problems

In this talk we will discuss the implementation and performance of an *auxiliary space* based algebraic solver for definite Maxwell problems, discretized with edge elements. The algorithm is based on a recent theoretical result by Hiptmair and Xu, and utilizes two internal Algebraic Multigrid (AMG) V-cycles: one for a scalar and one for a vector Poisson-like matrix. The parallel scalability of this approach is directly tied to the AMG performance on Poisson problems.

Tzanio V. Kolev, Panayot S. Vassilevski
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
tzanio@llnl.gov, panayot@llnl.gov

MS14

Scalability Infrastructure for the Lustre File System

This paper describes low-level infrastructure in the Lustre file system that addresses scalability in very large clusters. The features deal with I/O and networking, lock management, recovery and failure, and other scalability-related issues.

Peter Braam
Cluster File Systems
braam@clusterfs.com

MS14

High End Computing File Systems and I/O (HEC FSIO): Coordinating the US Government Research Investments

The High End Computing Interagency Working Group (HEC IWG) is chartered with coordinating US Government investments in Research and Development (R&D) for HEC. The HEC FSIO Technical Advisory Group (TAG) is chartered with providing guidance to the HEC IWG in the area of File Systems and I/O (FSIO). The HEC FSIO research needs and priorities will be discussed. Also, the currently portfolio of 28 research projects will be reviewed. Additionally, the future direction for the HEC FSIO area

will be outlined including programs for taking research outcome into products and discussion of a new round of government sponsored research to continue to feed the R&D pipeline in this area will be outlined.

Gary Grider
Los Alamos National Laboratory
ggrider@lanl.gov

MS14

I/O Architectures for Petascale Computing

Production high-performance storage systems today are typically constructed from enterprise storage hardware, with a parallel file system such as GPFS, Lustre, or PVFS tying this hardware into a coherent whole. As we move into the petascale regime, constructing storage systems in this way is becoming problematic. In this talk we will discuss some of the challenges in storage at petascale, particularly in reliability and performance, and examine hardware and software options that can help us construct effective storage systems at this extreme scale.

Rob Ross
Argonne National Laboratory
rross@mcs.anl.gov

MS14

Structured Streams: Data Services for Petascale Science Environments

The challenge of meeting the I/O needs of petascale applications is exacerbated by an emerging class of data-intensive HPC applications that requires annotation, reorganization, or even conversion of their data. We introduce an end-to-end approach to meeting these requirements. The Structured Streaming Data System (SSDS) enables high-performance data movement or manipulation between the compute and service nodes of the petascale machine and between/on service nodes and ancillary machines. This talk describes the SSDS architecture, motivating its design decisions and intended application uses. Performance claims are supported with experiments benchmarking the underlying software layers of SSDS, as well as application-specific usage scenarios.

Karsten Schwan
College of Computing
Georgia Institute of Technology
schwan@cc.gatech.edu

MS15

An Empirical Investigation of Generating Parallel Quasirandom Sequences by Using Different Scrambling Methods

Quasi-Monte Carlo (QMC) methods are now widely used in scientific computation. The use of randomized QMC methods, where randomness can be brought to bear on quasirandom sequences through scrambling and other related randomization techniques, brings more wide applications for QMC. Scrambling QMC offers a natural way to generate parallel sequences. QMC applications have high degrees of parallelism, can tolerate large latencies, and usually require considerable computational effort, making them extremely well suited for grid computating. Parallel computations using QMC require a source of quasirandom sequences, which are distributed among the individual parallel processes. However, the integration variance can depend strongly on the scrambling methods. Much of the work dealing with scrambling methods has been aimed at ways of linear scrambling methods. In this paper, we take a close look at the quadratic scrambling method for Halton sequences in generating parallel sequences.

Hongmei Chi
Computer Science
Florida A&M University
hchi@cis.famu.edu

MS15

Estimation of Migration Rates and Effective Population Numbers by Using Importance Sampling

After coalescence theory is widely used to explore diversity among populations and species in population genetics (phylogenetics), the computation of likelihood or posterior distribution of the population genetics parameters are common tasks in computational biology. The numerical results of these approaches can be achieved by Monte Carlo simulations. This paper focuses on exploring the use of uniform random sequences, more specifically, completely uniformly distributed sequences to calculate the likelihood with the help of importance sampling. We demonstrate by examples that quasi-Monte Carlo can be a viable alternative to the Monte Carlo methods in population genetics. Analysis of a simple one-population problem in this paper showed that quasi-Monte Carlo methods achieve the same or better parameter estimates as standard Monte Carlo, but have the potential to converge faster and so reduce the computational burden.

Peter Beerli
Florida State University
School of Computational Science
beerli@scs.fsu.edu

Hongmei Chi
Computer Science
Florida A&M University
hchi@cis.famu.edu

MS15

Hybrid Parallel Tempering and Simulated Annealing Method in Rosetta Practice

We applied our recently developed hybrid Parallel Tempering (PT)/ Simulated Annealing (SA) method to the Rosetta program. The hybrid PT/SA method is an effective global optimization algorithm to overcome the slow convergence in low-temperature protein simulation by initiating multiple systems to run at multiple slowly decreasing temperature levels (SA scheme) and randomly switch with neighbor temperature levels (PT scheme). The PT scheme can significantly enhance the relaxation rate in the SA search. With hybrid PT/SAs fast barrier-crossing capability, we expect to achieve resolution improvement compared to the original Rosetta program. Our preliminary results show that the Rosetta fragment assembly implementation using hybrid PT/SA method has a broader exploration of the protein folding scoring function landscape and exhibits a 0.2 4.0A shift toward the native structure in most of the Rosetta benchmark proteins. Our analysis and computational results strongly indicate that the key to further improve ab initio protein structure prediction resolution is to develop a more sensitive scoring function.

Yaohang Li

Computer Science
North Carolina A&T University
yaohang@ncat.edu

MS15

Novel Stochastic Methods in Biochemical Electrostatics

We will present a Monte Carlo method for solving boundary value problems (BVPs) involving the Poisson-Boltzmann equation (PBE). Such BVPs arise in many situations where the calculation of electrostatic properties of solvated large molecules. The PBE is one of the implicit solvent models, and has accurately modeled electrostatics over a wide range of ionic solvent concentrations. With the new method we compare the algorithmic and computational properties of this algorithm to more commonly used, deterministic, techniques, and we present some computational results. This work is part of an ongoing collaboration with several Florida State faculty members and a collaborator at the Russian Academy of Sciences.

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

MS16

A Programming Model for Hybrid Parallelism with Consistent Numerical Results

The anticipated use of manycore processors in next generation HPC systems has re-opened the dialogue on distributed memory versus shared memory versus hybrid parallel programming models. If thread based parallelism becomes necessary then the existing problem of parallel algorithms producing a different numerical result for different parallel decompositions will be exacerbated by thread race conditions. A hybrid parallel programming model is proposed to (1) manage complexity of hybrid parallelism and (2) generate consistent numerical results.

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

MS16

Challenges in Programming Next-generation Parallel Computer Systems

The coming generations of scalable systems will present a new dimension of parallelism: multicore. Although similar in some ways to shared memory multiprocessor (SMP) nodes, chip multiprocessors (CMPs) are different in a number of important ways that will challenge algorithms and application developers to consider new techniques. In this presentation we discuss a brief history of the challenges that are ever-present for shared memory parallel computing, survey the conclusions made a decade or more ago with the first generation of SMP nodes, and highlight how CMPs are different. We then discuss the current challenges that we face in scalable computing when effective use of CMPs is a requirement for scalability.

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

MS16

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

Fortran 2008 is a modern object-oriented programming language that also contains, for the first time, a parallel programming model. Fortran 2003, the current standard, allows programmers to develop applications using object-oriented abstractions without sacrificing performance. Fortran 2008, the next standard, contains the co-array parallel programming model, a simple yet powerful model based on Fortran syntax and semantics. Combining these two new Fortran features into application-specific class libraries provides a new direction for application development.

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

MS16

Multicore Extensions to the Petra Object Model for Parallel Linear Algebra

Petra provides a foundation of algebraic matrix and vector infrastructure in the Trilinos framework. Petra includes data structures, methods for performing operations such as matrix-vector products, updates, and parallel redistributions, etc. Previously, Petra has been targeted at distributed-memory platforms, using MPI. We will discuss efforts to extend Petra to multicore platforms using a threaded programming model. We will survey related efforts in this field, describe the approach weve taken in Petra, and show performance results.

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

MS17

Overview of the Vertical Integration of Trilinos Solver Algorithms in a Production Application Code

In this talk, I will provide an overview of for the solver integration work. I will describe the Thyra interface effort which provides the foundation for interoperability of the numerical algorithms. I will describe the Thyra ModelEvaluator interface that supports a variety of different numerical problem types ranging from square nonlinear equations to transient optimization. As for solved problems, I will describe a steady-state parameter estimation optimization problem using MOOCHO and transient sensitivities using Rythmos. I will also describe our work to develop a new collaboration between application and algorithm developers that is based on constant software integration and nightly testing.

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

MS17

Analysis Tools for Large-scale Simulation with Application to the Stationary Magnetohydrodynamics Equations

In many cases, the ability to perform forward solves of large-scale systems of partial differential equations is not

enough to be a useful as a design and decision support tool. As simulations become more complex, application tools need to be able to traverse larger and larger parameter spaces in an efficient manner. Bifurcation, stability, and optimization tools are meant to identify important design phenomena/criteria without the need to perform tedious and time-consuming forward simulations over the entire parameter space. This talk will focus on the application of bifurcation and linearized stability analysis algorithms to identify regions of instability for the well-known Hydromagnetic Reyleigh-Bernard problem. A new stabilized-finite element formualtion for the single-fluid MHD equations will be presented along with pertinent analysis algorithms. Efficiency and robustness of new solvers/algorithms including block-eigensolvers and pitchfork tracking algorithms will be discussed. Typical pitfalls related to moving algortihms from a "research state" to a "production state" will be addressed.

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

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

Roger Pawlowski
Sandia National Labs
rppawlo@sandia.gov

Heidi K. Thornquist
Sandia National Laboratories
hkthorn@sandia.gov

MS17

Analytic Sensitivities in Large-scale Production Applications via Automatic Differentiation with Sacado

A fundamental task in many nonlinear solver, optimization, and stability analysis algorithms is accurately and efficiently computing derivatives of simulation responses with respect to state variables and parameters. In this talk, a hybrid symbolic/automatic differentiation procedure for computing analytic derivatives in large-scale, element-based application codes will be presented along with its connection to vertical integration of Trilinos solvers in Charon. The automatic differentiation calculations are implemented by a new Trilinos package called Sacado.

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

MS17

Solving Linear Systems with Multiple Right-hand Sides using Belos

The solution of very large and sparse linear systems with multiple right-hand sides is often required in modeling and simulation. While many numerical methods for solving these linear systems have been proposed in the literature, Belos provides an iterative linear solver framework for developing and delivering these methods to application codes. In this talk we will discuss interacting with Belos and the available numerical methods for solving both single and multiple right-hand side linear systems.

Heidi K. Thornquist
Sandia National Laboratories
hkthorn@sandia.gov

MS18

Retro Dataflow: The Accomplishments of the Sisal Language Project

The Sisal Language Project was a joint effort by Lawrence Livermore National Laboratory, Colorado State University, and University of Manchester to develop a high performance functional programming language for conventional parallel computer systems. The main project began in the early '80s and concluded in 1996. The project was successful in that many Sisal programs ran as fast as their Fortran or C equivalent on parallel systems of the time such as the SGI Challenge, multi-processor VAXs, and Cray vector computers. In this talk, we review the language's syntax and semantics, optimizing compiler, and high-performance runtime system. We describe two critical optimizations: memory pre-allocation and update-in-place. We conclude with comparative performance numbers, and give the availability of reports and source code.

Thomas DeBoni
Lawrence Livermore National Laboratory
tdeboni@llnl.gov

John Feo
Microsoft Inc.
jofeo@microsoft.com

Pat Miller
D.E. Shaw Research, LLC
patrick.miller@gmail.com

MS18

Ct: Channeling NeSL and SISAL in C++

Abstract not available at time of publication.

Anwar Ghuloum
Intel Corporation
2200 Mission College Blvd, Santa Clara ,CA 95054
anwar.ghuloum@intel.com

MS18

Multithreading for Linear Algebra in Distributed Memory Environments

We describe our work implementing multithreaded versions of linear algebra kernels in distributed memory environments. We chose to implement these using a data driven approach mainly for its latency tolerance properties. Performance and correct execution were primarily influenced by scheduler decisions. The primary decision of task graph execution order directly impacts memory usage and processor cache behaviour leading to varied performance. We present our techniques for approaching these problems in the context of dense LU and sparse Cholesky decompositions.

Parry Husbands
Lawrence Berkeley National Laboratory
pjrhusbands@lbl.gov

MS18

Implementing Dense Linear Algebra Algorithms on Multi-Core Processors Using Dataflow Execution Model

We identify dynamic execution of parallel tasks as a necessary component for achieving high performance for dense linear algebra algorithms executed on modern multi-core processors. We utilize fine granularity of parallelization and dynamic scheduling of tasks based on data dependencies expressed by a direct acyclic graph of the computation. We present results for LU, Cholesky and QR factorizations on an Intel dual-socket, quad-core system, which are competitive with Math Kernel Library implementations.

Jakub Kurzak
Innovative Computing Laboratory
University of Tennessee
kurzak@cs.utk.edu

John Shalf
Lawrence Berkeley National Laboratory
jshalf@lbl.gov

MS18

Computational Space-times and Domain Flow (TM)

Today's microprocessors take more time to load an operand from memory than to compute a basic instruction. Furthermore, performance predictability of modern microprocessors that use caching structures for data and virtual address translations has become essentially unsolvable. Together, these two problems lead to a very unproductive environment for high-performance software development. We will introduce and use a computational space-time to derive a conceptual programming model that takes communication delays into account, and apply this model on fine-grained parallel algorithms to demonstrate more predictable and efficient execution environments as compared to multi-core or distributed memory machines.

Theodore Omtzigt
Stillwater Supercomputing Inc.
theo@stillwater-sc.com

MS19

Accelerators, Cell Broadband Engine, Graphics Processors, and FPGAs

While we are still witnessing Moore's Law by the steady production of chips that mass billions of transistors, clearly we have reached plateaus on clock frequency, power, and single stream performance. This new era has caused a rethinking of microprocessor design in search of innovations that will allow the continued performance improvement of scientific applications at an exponential rate. One technology that holds promise combines traditional microprocessors with special-purpose, very high performance, low-power chips such as the IBM Cell Broadband Engine, Graphics processors, and FPGAs, to accelerate the performance of computational science and engineering applications. The use of these chip accelerators will likely be a path forward, yet new challenges await such as system-level design, partitioning applications to accelerators, and tools for designing applications.

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

MS19

Looking Beyond Familiar Programming Paradigms to Efficiently Program Chip Multiprocessors

Our attempts to make chip multiprocessors (CMPs) conform to familiar programming paradigms such as message passing or treating them as an SMP-on-a-chip using OpenMP represents a failure of imagination. Neither approach exploits the unique features of the CMP architecture such as 10-100x lower latencies and 10-100x higher bandwidth between cores on the chip. This talk will examine alternatives to conventional programming approaches that are better able to exploit the capabilities of CMPs and their ramifications for programming models and algorithm design.

John Shalf
Lawrence Berkeley National Laboratory
jshalf@lbl.gov

MS19

MegaAlgorithms: Challenges and Approaches to Applications Development

It appears that general-purpose Exascale computer systems may be technologically feasible within the next 15 years. There are however significant challenges to accomplishing an Exascale performance goal, including algorithm scalability, power efficiency, systems reliability, and chip-level packaging. In this talk I will sketch some approaches to attacking the challenges of developing scientific applications for multi-million core systems with a special focus on algorithm development and applications programming models.

Rick Stevens
Argonne National Laboratory
stevens@anl.gov

MS19

Tightly-Coupled Heterogeneous Computing - Deus Ex Machina?

Commodity architectures are rapidly moving to homogenous multicore processors as a strategy to sustain performance improvements while confronting the constraints of power, heat, signaling, and instruction level parallelism. Enter tightly-coupled heterogeneous architectures. Numerous such architectures have recently appeared on the supercomputing scene: FPGAs, Cell, GPUs, etc. Our initial investigations have revealed that heterogeneous architectures can dramatically improve the performance of specific algorithms; however, these improvements must be balanced against programmer productivity and unstable performance reactivity.

Jeffrey S. Vetter
Oak Ridge National Laboratory
Georgia Institute of Technology
vetter@computer.org

MS20

Parallel Issues in Combinatorial Scientific Computing

The design and implementation of efficient parallel al-

gorithms in the area of combinatorial scientific computing requires dealing with issues such as data distribution, workload balancing, and communication reduction, making complicated computations even more challenging. In this introductory talk, we will present several commonly encountered parallel issues, such as choosing between 1D and 2D distributions and determining the granularity of the computation. We illustrate this with examples from parallel partitioning, matching, and sparse LU decomposition.

Rob H. Bisseling
Dept. Mathematics
Utrecht University
Rob.Bisseling@math.uu.nl

MS20

A Scalable Algorithmic Framework for Parallel Graph Coloring on Distributed Memory Computers

As part of the SciDAC-sponsored CSCAPES Institute's effort to develop parallel algorithms and software for combinatorial scientific computing problems, we have recently developed a framework for parallelizing greedy graph coloring algorithms on distributed-memory computers. The framework blends a variety of parallelization techniques, including exploiting features of the initial data distribution, the use of speculative coloring and randomization, and a BSP-style organization of computation and communication. We discuss the framework and present experimental results that demonstrate its scalability.

Erik G. Boman
Sandia National Labs, NM
Scalable Algorithms Dept.
egboman@sandia.gov

Fredrik Manne
Department of Informatics
University of Bergen
fredrikm@ii.uib.no

Doruk Bozdag
Ohio State University
bozdag@ece.odu.edu

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

Assefaw Gebremedhin
Old Dominion University
Compter Science Department
assefaw@cs.odu.edu

MS20

Efficient and Scalable Parallel Graph Partitioning Methods

The realization of efficient parallel graph partitioners requires the parallelization of the multi-level framework which is commonly used in sequential partitioners to improve quality and speed. While parallel matching algorithms are now efficient and un-biased enough to yield coarsened graphs of good quality, the local optimization algorithms used in the refinement step during uncoarsening are still an issue. This talk will address this problem and present scalable parallel diffusive methods which can advantageously replace classical Fiduccia-Mattheyses-like algorithms for this purpose.

Jun-Ho Her, Francois Pellegrini
INRIA Futurs, ScAlApplix project
LaBRI
jhher@labri.fr, pelegrin@labri.fr

MS20

Parallel Partitioning Techniques in Circuit Simulation

The linear systems generated through circuit simulation provide a challenge for matrix ordering and load balancing due to their heterogeneous matrix structure. We will discuss the graph techniques employed by the Xyce circuit simulation code and their effect on direct and iterative linear solvers with a particular emphasis on parallel partitioning techniques. A performance comparison of the hypergraph and graph partitioning algorithms in Zoltan/Issoropia applied to circuit problems in Xyce will be presented.

Robert J. Hoekstra, Heidi K. Thornquist
Sandia National Laboratories
rjhoeks@sandia.gov, hkthorn@sandia.gov

MS21

Enabling Highly-Scalable Ultra-High Resolution Climate Simulations Using the Community Climate System Model

Performing climate simulations offers a number of challenges. These include scaling to a large enough process set to make the calculation tractable, and negating high-frequency deleterious computational modes that limit the time step while not contributing meaningfully to the solution. Extending to ultra-high resolution exacerbates these issues while exposing scaling bottlenecks and computational phenomena that were previously hidden. We present our latest progress using the Community Climate System Model with Lin-Rood dynamical core.

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

Arthur A. Mirin
Lawrence Livermore Nat'l Lab.
mirin@llnl.gov

MS21

Computational Performance of Parallel Air Quality Models

Air quality models are complex and incorporate a wide variety of processes including transport, diffusion, chemistry, particle interaction and thermodynamics, emissions and deposition. We discuss the parallel performance of several data decomposition strategies applied to the STEM air quality model. The results are obtained on Virginia Tech's System X supercomputer.

John Linford
Virginia Polytechnic Institute and State University
jlinford@vt.edu

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

Emil Constantinescu
Virginia Polytechnic Institute and State University
emconsta@vt.edu

MS21
A High-order Adaptive Mesoscale Model

We present a non-conforming adaptive discontinuous Galerkin model intended for meso-scale atmospheric simulations. The goal of the project is to show that high-order methods can effectively resolve complex structures in mesoscale flows. Adaptivity in both h (elemental) and p (polynomial) is used to detect and resolve such structures efficiently. The model employs explicit, and implicit Rosenbrock W-methods. Preconditioning is of utmost importance and the performance of various strategies are presented: Block Jacobi, p-Multigrid.

Cathy Mavriplis
NOAA National Severe Storms Laboratory
University of Oklahoma (CIMMS)
catherine.mavriplis@noaa.gov

Amik St-Cyr
National Center for Atmospheric Research
amik@ucar.edu

Louis Wicker
NOAA, NSSL
louis.wicker@noaa.gov

MS21
Performance and Performance Scalability of the Community Atmosphere Model

The Community Atmosphere Model (CAM) has recently undergone significant development, with changes in and additions to the subgrid parameterizations, increases in resolution of the target computational grids, and new target computing platforms. In this work we describe performance of CAM for a new benchmark problems that more accurately represent current production usage. Minimally, we will examine both performance and scalability on the Cray XT4 with dual-core AMD Opteron processors, focusing on evaluating the parallel performance of the current numerical algorithm and their parallel implementations on this platform. If possible, we will also present results from a quad-core Cray XT4 and a quad-core IBM BG/P, determining the sensitivity of the current algorithms to system performance characteristics. The presentation will finish with a brief discussion of a benchmarking effort that will compare performance of the current CAM dynamical core with that of a number of proposed, potentially more scalable, alternatives.

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

Arthur A. Mirin
Lawrence Livermore Nat'l Lab.
mirin@llnl.gov

MS22
Parallel Mesh Generation for Medical Image Computing

Parallel mesh generation codes based on known sequential mesh generation methods. Given that it takes very long time to develop the software infrastructure for sequential industrial strength mesh generation libraries, it is clear that traditional parallelization approaches deliver technology that is outdated. This problem becomes more serious if one considers that improvements of sequential codes in terms of quality, speed, and functionality are open ended. In this talk we present a COTS (commercial of-the-shelf) based approach to parallel mesh generation for addressing this serious problem. We will present our experience from different parallel meshing methods and in addition we will present their extensions to meet medical image computing requirements like patient/image specific input and conformity of the mesh between different tissues.

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

MS22
Mesh Generation Using the CGAL Library

CGAL, the Computational Geometry Algorithms Library, is an established C++ framework for performing various geometric computations. It provides tools for the generation and processing of meshes of volumes and surfaces, which can be useful for meshing 3D medical images among other things. We will present the existing algorithms of CGAL, mostly based on the Delaunay refinement and filtering paradigms, as well as the ongoing developments, in particular concerning the parallelization of the meshing algorithms.

Sylvain Pion
INRIA
Sylvain.Pion@sophia.inria.fr

Pierre Alliez
INRIA Sophia-Antipolis
pierre.alliez@sophia.inria.fr

Mariette Yvinec
INRIA
mariette.yvinec@sophia.inria.fr

MS22
Mesh Generation and Parallel Mesh Generation for Biomedical Applications

Historically, mesh generation has been a critical bottleneck in efforts to conduct biomedical simulations. In this presentation, we will discuss a variety of tools and methodologies for generating meshes suitable for biomedical simulations consisting of a wide array of geometries and phenomena. We will also discuss current efforts to extend these methods for parallel environments.

Jason Shepherd, Philippe Pebay
Sandia National Laboratories
jfsheph@sandia.gov, ppebay@sandia.gov

Michael Stephenson
Computation Simulation Software

mbsteph@sandia.gov

MS22

Parallel Mesh Quality Improvement Techniques for Blood Clot Entrapment

Meshes are generated from medical images by: assembling the 2D patient scans into 3D, extracting the 3D geometry, and creating an initial volume mesh for the geometry. Meshes obtained via automatic mesh generation of the 3D geometry often lack suitable quality due to low image resolution, difficulty in performing the extraction, or geometric complexity. We employ parallel mesh quality improvement techniques to improve the quality of the initial mesh in the blood clot entrapment problem.

Suzanne M. Shontz, Nicholas Voshell
Department of Computer Science and Engineering
The Pennsylvania State University
shontz@cse.psu.edu, njv116@cse.psu.edu

MS23

Some Issues in Dense Linear Algebra Algorithms for Multicore

The advent of multicore processors will considerably influence future research in High Performance Computing. To take advantage of these architectures in dense linear algebra operations, new algorithms are proposed that use finer granularity and minimize synchronization points. After presenting some of these algorithms, we address the issue of implementing pivoting strategies and investigate randomization techniques to avoid pivoting in some cases. A study on adapting SuperLU for multicore systems is also presented.

Marc Baboulin
CERFACS
baboulin@cerfacs.fr

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

Alfredo Buttari
University of Tennessee
buttari@cs.utk.edu

Riadh Fezzani
ONERA
rfezzani@onera.fr

MS23

Fault-tolerant Algorithms for Dense Linear Algebra

In this talk we propose to use ABFT techniques for having efficient fault tolerance of a large scale applications at the algorithm level. We will only focus on linear algebra algorithms. ABFT techniques enable us to maintain a consistent checksum of the matrices along the algorithms. The overhead is that we need to use $2n$ processors from the n^2 originally available. Some experiments using dense linear algebra kernels are provided and we report performance of more than 1 TeraFLOP/sec on a 500 nodes cluster for a node failure and the overhead of the method goes to zero as the number of processors increases as predicted.

Remi Delmas
France
delmas@cs.utk.edu

Julien Langou
University of Colorado at Denver and Health Sciences Center
julien.langou@cudenver.edu

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

MS23

Precise Solutions for Linear Systems and Least Squares Problems in (Sca)LAPACK

We present the algorithms, error bounds, and numerical results of extra-precise iterative refinement for linear systems and linear least squares problems. The algorithms require only limited use of extra precision and add only $O(n^2)$ work to the $O(n^3)$ cost of LU, or $O(mn)$ work to the $O(mn^2)$ cost of QR. Our algorithms reduce the forward normwise and componentwise errors to $O(\varepsilon)$ unless the system is too ill-conditioned. The extra precision calculation is facilitated by the new extended-precision BLAS standard in a portable way.

Jason Riedy
University of California, Berkeley
Division of Computer Science
ejr@cs.berkeley.edu

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

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

Yozo Hida
Univ. of California
Berkeley
yozo@cs.berkeley.edu

MS23

Special-purpose Hardware and Algorithms for Accelerating Dense Linear Algebra Computations

High performance scientific computing using not only multicores but also special-purpose hardware, like GPUs and FPGAs, is becoming pervasive. We address challenges in programming and developing innovative DLA algorithms for FPGAs and GPUs. In particular, this will include issues related to LU, QR, and Cholesky factorizations; QR vs LU in solving Ax=b; mixed-precision iterative refinement; pivoting (e.g. ways to avoid it). Finally, of interest is to compare parallel algorithm designs for multicores, FPGAs and GPUs.

Greg Peterson
Electrical and Computer Engineering
The University of Tennessee
gdp@utk.edu

Shirley Moore
Department of Computer Science
The University of Tennessee
shirley@cs.utk.edu

Vincent Natoli
Stone Ridge Technology
vnatoli@stoneridgetechnology.com

Dilip Patlolla
Electrical and Computer Engineering
The University of Tennessee
dpatloll@utk.edu

Marc Baboulin
CERFACS
baboulin@cerfacs.fr

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

MS24

Accelerating Combinatorial Scientific Computing with the Cell Broadband Engine Processor

The Sony-Toshiba-IBM Cell Broadband Engine is a heterogeneous multicore architecture that consists of a traditional microprocessor (PPE), with eight SIMD coprocessing units (SPEs) integrated on-chip. Because of the performance capabilities of the Cell BE, it is considered as an application accelerator for next-generation petascale supercomputers. In this talk, we discuss the use of Cell for solving challenging combinatorial scientific computing applications. Our algorithms are tested on a cluster of 14 IBM QS20 Cell Blades with 24 Cell BE processors.

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

MS24

Some Experiment with Linear Algebra Algorithms on the Sony Playstation 3

This talk describes some experiment and limitations of using the Sony PlayStation3 for scientific computations. We will also look at issues in clustering PS3s.

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

MS24

Multi-core Optimization of Sparse Matrix Vector Multiplication

Industry has moved to chip multiprocessor (CMP) system design in order to better manage trade-offs among performance, energy efficiency, and reliability. To fully unleash the potential of these systems, the HPC community must develop multicore specific optimization methodologies for important scientific computations. In this work, we examine multicore-specific optimization of sparse matrix-vector multiply (SpMV) – one of the most heavily used kernels in scientific computing – across a broad spectrum of CMP designs.

Lenny Oliker
Lawrence Berkeley National Laboratory
loliker@lbl.gov

Sam Williams
UC Berkeley
samw@cs.berkeley.edu

MS24

High-Performance String Search in Multi-core Processors

The explosion of digital data promises a profound impact in both the quality and rate of discovery and innovation in science and engineering, as well as in other societal contexts. String search is at the core of the tools we use to curb this explosion, such as search engines, network intrusion detection systems, spam filters, and anti-virus programs. But as communication speed grows, our capability to perform string matching seems to fall behind. Multi-core architectures promise enough computational power to cope with the incoming challenge, but it is still unclear which algorithms and programming models to use to unleash this power. In this talk, we will discuss several algorithms and techniques to perform high-performance searches on large data sets with multi-core processors.

Fabrizio Petrini
IBM TJ Watson
fpetrin@us.ibm.com

MS25

Automatic Tuning for the LOBPCG Eigenvalue Solver

Abstract not available at the time of publication.

Toshiyuki Imamura
The University of Electro-Communications
imamura@im.uec.ac.jp

MS25

Auto-tuning Effect of Iterative Method Library on Windows CCS

Current computer systems become more and more complex. Therefore, numerical libraries which achieve high performance on any platforms are highly recommended. Our libraries select an optimal code automatically and run it in order to reduce computation time as much as possible. For example, an appropriate preconditioning method and communication scheme are automatically determined on each platform. We show the performance of our iterative method library on Windows Computer Cluster Server with MS-MPI.

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

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

MS25

Performance of the Complex Hessenberg QR Algorithm on the CSX600 Processor

The small-bulge multishift QR algorithm for complex nonsymmetric eigenproblems performs most of the computation with the level-3 BLAS and is therefore suited to floating-point accelerators such as the CSX600. However, to exploit the performance of the accelerator, we have to tune several parameters such as the number of shifts and blocking sizes in matrix multiplications. In this talk, we optimize these parameters based on a semi-empirical performance model and show how much speedup can be obtained.

Shao-Liang Zhang
Nagoya University
qzy13306@nifty.com

Takafumi Miyata
Nagoya University, Japan
miyata@na.cse.nagoya-u.ac.jp

Yusaku Yamamoto
Nagoya University
yamamoto@na.cse.nagoya-u.ac.jp

Yoshimasa Nakamura
Graduate School of Informatics, Kyoto University
ynaka@i.kyoto-u.ac.jp

MS25

Automatic Selection of Parameters in Parallel Preconditioners for Ill-conditioned Problems

Selection of optimum preconditioning methods with appropriate parameters for parallel computing is a difficult task, especially for ill-conditioned problems beacause of many parameters. In real applications, convergence of iterative solvers is often affected by local heterogeneity and/or discontinuity of the field. Our strategy utilizes both of global information obtained from derived coefficient matrices and local information, such as information of each mesh in finite-element applications. In this presentation, progress of this work will be discussed.

Kengo Nakajima
The University of Tokyo
Department of Earth & Planetary Science
nakajima@eps.s.u-tokyo.ac.jp

MS26

Charm++ for Cells and Clusters of Cells

Charm++ programming model is a good fit to the Cell BE processor. We have developed an "Offload API", which can be used by Charm++ applications, to offload work to the SPEs. We are extending the Charm++ translator to automatically generate code for offloading safe entry methods to the SPEs, enabling Charm++ applications to be portable between Cell-based and non-Cell-based platforms. We are also studying new load balancing problems faced on multiple-cell clusters.

Laxmikant Kale, David Kunzman
University of Illinois at Urbana-Champaign
kale@uiuc.edu, kunzman2@uiuc.edu

MS26

Unified Scheduling of Polymorphic Parallelism on the Cell Processor

Programming the Cell presents unique challenges, stemming from heterogeneity and the need to orchestrate parallelism across many layers. Though approximate guidelines for harnessing parallelism at any given layer exist, programmers lack models, methods and abstractions to synthesize multiple forms of parallelism, into optimal program decomposition execution plans. This talk presents a unified framework for synthesizing and scheduling polymorphic parallelism on Cell, based on an event-driven execution model and dynamic allocation of processor resources to events.

Dimitris Nikolopoulos
Virginia Tech.
dsn@cs.vt.edu

MS26

Making Cell Do Hard Things: Where is the Cell's Sweet Spot?

Technological obstacles have prevented the microprocessor industry from achieving increased performance through increased chip clock speeds. In a reaction to these restrictions, the industry has chosen the multicore processors path. Multicore processors promise tremendous GFLOPS performance but raise the challenge of how one programs them. In this talk, I will discuss the motivation for multicore, the implications to programmers and how the Cell/B.E. processors design addresses these challenges. As an example, I will review one or two applications that highlight the strengths of Cell.

Michael Perrone
IBM Watson
mpp@us.ibm.com

MS26

Opportunities and Challenges for Cell Broadband Engine in Life Sciences

This talk will focus on the evaluation of the Cell Broadband Engine (CBE) for biology workloads. A significant number of life sciences workloads show both fine and coarse-grained parallelism which can benefit immensely from multiple vectorized cores on the CBE. In this talk, we will detail our experiences with the Cell using a variety of programming platforms. We also discuss the programming and performance issues involved from a programmer perspective.

Vipin Sachdeva
IBM Research, Austin
vsachde@us.ibm.com

MS27

Parallel Algorithms for Small-world Network Analysis and Partitioning

We present SNAP, a parallel library for exploratory anal-

ysis and partitioning of small-world networks. SNAP includes fast, parallel implementations of fundamental graph-theoretic kernels and topological analysis metrics (e.g., breadth-first search, connected components, betweenness centrality), as well as novel community structure detection algorithms. In the talk, we highlight algorithmic details and optimizations to achieve scalable parallel performance on small-world graph instances.

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

MS27

The MultiThreaded Graph Library

The MultiThreaded Graph Library (MTGL) is a set of open-source codes for graph algorithms that target emerging massively multithreaded architectures. This library is based upon a kernel of the Boost Graph Library, though it does not require Boost. MTGL codes run on serial workstations, and can run on the Cray MTA/XMT supercomputers. Much of the detail inherent in the programming model of the latter machines is abstracted away from the application programmer.

Jonathan Berry
Sandia National Laboratories
jberry@sandia.gov

MS27

Tensor Decompositions for Analyzing Multi-Link Graphs

Link analysis of data represented by a graph typically focuses on a single type of connection. However, often we want to analyze data that has multiple linkages between objects. The goal of this talk is to show that tensors and their decompositions provide a set of tools for multi-link analysis. We provide examples of how tensors can be used to understand the structure of document spaces and define similarities based on multiple linkages.

Philip Kegelmeyer, Tammy Kolda
Sandia National Laboratories
wpk@sandia.gov, tgkolda@sandia.gov

Danny Dunlavy
Computer Science and Informatics Department
Sandia National Laboratories
dmdunla@sandia.gov

MS27

Analytic Theory of Power Law Graphs

An analytical theory of power law graphs is presented based on the Kronecker graph generation technique. The analysis uses Kronecker exponentials of complete bipartite graphs to formulate the sub-structure of such graphs. This allows various high level quantities (e.g. degree distribution, betweenness centrality, diameter, eigenvalues, and isoparametric ratio) to be computed directly from the model parameters. The implications of this work on "clustering" and "dendragram" heuristics are also discussed.

Jeremy Kepner
MIT Lincoln Laboratory
kepner@ll.mit.edu

MS28

Parallel Performance Issues for Multiphysics Adaptive Mesh-refinement Algorithms

Abstract not available at the time of publication.

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

MS28

Distributed Coupling Toolkit

Recent developments in computational sciences have enabled scientists to create sophisticated software tools and techniques that have contributed to the development of high accuracy numerical models that are used to study physical phenomena. The next level of sophistication addresses the integration or coupling of one or more computational models to simulate a more complex physical system. These coupled systems, are gen-erally multidisciplinary in nature and are now emerge in a broad spectrum of fields in science and engineering such are Earth Sciences, Fusion Energy, Structural Engineering and Astrophysics.

Vicent Vidal
Polytechnic University of Valencia
vvidal@upv.es

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

MS28

An Integrated Simulation of Seismic Wave and Tsunami Propagation

An integrated simulation code for seismic waves and tsunami is developed for mitigating earthquake disasters expecting for future earthquakes. The ground motions and deformation of seafloor are firstly evaluated by solving equation of motions in 3D, and the Navier-Stokes equation is used for evaluating tsunami. The simulation code is implemented on a cluster of PC and the Earth Simulator for examining strong ground motion and tsunami caused by the Sanriku earthquake in 1896.

Takashi Furumura, Tatsuhiko Saito
Earthquake Research Institute
University of Tokyo
furumura@eri.u-tokyo.ac.jp, saito-ta@eri.u-tokyo.ac.jp

MS28

Large Scale Coupling Simulations through MxN Parallel Data Redistribution

A prototype of framework for large-scale parallel coupling simulations through MxN parallel data redistribution under SPMD environment has been developed for "Integrated Predictive Simulation System for Earthquake and Tsumami Disaster". Coupled finite element simulations for seismic response of multiple oil-storage tanks with fluid-structure interaction are performed by the developed system on PC cluster and the results are demonstrated.

Hiroshi Okuda
RACE, University of Tokyo

okuda@race.u-tokyo.ac.jp

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

Tsuyoshi Ichimura
TIT/JST, Japan
ichimura@cv.titech.ac.jp

Shin'ichi Ezure
MRI Research Associates
ezure@mri-ra.co.jp

Massaki Matsumoto
Mitsubishi Research Institute, Inc.
matsumot@mri.co.jp

Kengo Nakajima
The University of Tokyo
Department of Earth & Planetary Science
nakajima@eps.s.u-tokyo.ac.jp

MS29

Parallel Eigenvalue Reordering in Real Schur Forms

Parallel block algorithms for eigenvalue reordering in the standard and generalized real Schur forms are presented, filling one of the gaps of the ScaLAPACK eigenproblem functionalities. Experimental results on HPC2Ns Linux clusters demonstrate the algorithms superior serial performance and good scalability. The algorithms are applied in parallel computation of stable invariant subspaces of Hamiltonian matrices arising from algebraic Riccati matrix equations. Moreover, generalizations of the reordering algorithms to product eigenvalue problems are outlined.

Daniel Kressner
SAM - Seminar for Applied Math
ETH-Zentrum, Zurich
kressner@math.ethz.ch

Robert Granat
Dept. of Computing Science and HPC2N
Umeå University, Sweden
granat@cs.umu.se

Bo T. Kågström
Umeå University
Computing Science and HPC2N
bokg@cs.umu.se

MS29

Distributed Square Blocked Packed Cholesky Factorization Algorithms with Near-Optimal Scheduling

Three algorithms for the Cholesky factorization (Basic, Static, and Dynamic), all using the Distributed Square Block Packed (DSBP) storage format, are designed and evaluated on a HPC cluster. The packed storage format saves memory and allows for level 3 performance via the BLAS. Two of the variants overlap communication with computation, virtually eliminating the communication overhead. This requires frequent polling of the MPI library as well as the use of both nonblocking sends and receives.

Fred Gustavson
IBM T.J. Watson Research Center,
Yorktown Heights NY, 10598, USA
fg2@us.ibm.com

Lars Karlsson
Umeå University
Computing Science and HPC2N
larsk@cs.umu.se

Bo T. Kågström
Umeå University
Computing Science and HPC2N
bokg@cs.umu.se

MS29

New Parallel Variants of Dense Eigenvalue Solvers

Recent improvements to the QZ algorithm for solving generalized eigenvalue problems are summarized. Among the major modifications are novel multishift QZ iterations based on chasing chains of tiny bulges and an extension of the so called aggressive early deflation strategy aiming at reducing the overall computation time on modern computing systems. A novel parallel implementation of a multishift QZ algorithm is presented. In addition, some inherent bottlenecks and possible remedies and improvements are discussed.

Björn Adlerborn
Umeå University
Computing Science and HPC2N
adler@cs.umu.se

Daniel Kressner
SAM - Seminar for Applied Math
ETH-Zentrum, Zurich
kressner@math.ethz.ch

Bo T. Kågström
Umeå University
Computing Science and HPC2N
bokg@cs.umu.se

MS29

RECSY and SCASY: Recursive Blocked and Parallel Algorithms for Sylvester-type Matrix Equations with Some Applications

RECSY uses recursive blocked algorithms for solving various matrix equations on serial and SMP/multicore nodes. The recursive approach leads to an automatic variable blocking amenable for deep memory hierarchies. SCASY uses explicit matrix blocking and 2D-block-cyclic mapping, wave-front algorithms, and RECSY node solvers to create ScaLAPACK-style algorithms for solving similar equations on distributed memory platforms. Both RECSY and SCASY distinguish between standard and generalized, one-sided and two-sided matrix equations, which are all ubiquitous in control system design.

Robert Granat
Umeå University
Computing Science and HPC2N
ganat@cs.umu.se

Bo T. Kågström
Umeå University
Computing Science and HPC2N
bokg@cs.umu.se

Isak Jonsson
Umeå University
Dept of Computing Science and HPC2N
isak@cs.umu.se

MS30

A Parallel SAMR Framework for Strongly Driven Fluid-structure Interaction Problems

The numerical simulation of transient shock- and detonation-driven fluid-structure interaction problems requires the coupling of time-explicit shock-capturing schemes to Lagrangian solid solvers for large deformations and/or fracture. In order to combine numerical accuracy with data locality, we employ a Cartesian embedded boundary approach supplemented with dynamic block-structured mesh adaptation in the computationally dominating fluid solver. The parallel performance of several applications will be discussed; the framework design of the SAMR software will be sketched briefly.

Ralf Deiterding
Oak Ridge National Laboratory
deiterdingr@ornl.gov

MS30

Parallel Computation of Three-dimensional Flows Using Overlapping Grids with Adaptive Mesh Refinement

Overlapping grids can be used to accurately and efficiently solve partial differential equations in complex geometry. We have developed an AMR toolkit, part of the Overture framework, to support adaptive mesh refinement on overlapping grids. Recently we have developed a parallel version of these AMR capabilities. In this talk we describe the features of this approach. We discuss the validation of the method and present some results from solving the reactive Euler equations.

Donald W. Schwendeman
Rensselaer Polytechnic Institute
Department of Mathematical Sciences
schwed@rpi.edu

William D. Henshaw
CASC
Lawrence Livermore National Lab
henshaw@llnl.gov

MS30

Weak Scaling Performance of Locally-structured Elliptic and Hyperbolic Solvers

In this paper we discuss weak scaling results for the Chombo AMR Helmholtz equation solver and AMR PPM Hyperbolic Solver algorithms. Both are SPMD MPI-based implementations running from 1 to 8192 processors. Good scaling results are obtained after a small set of critical optimizations. A methodology of replication-scaling is presented to avoid the usual problems with scaling studies and adaptivity. We have also attempted to distinguish between processor peak performance and a more realistic attainable performance.

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

Phillip Colella
Lawrence Berkeley National Laboratory
PColella@lbl.gov

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

Noel Keen
LBNL
ndkeen@lbl.gov

MS30

Implementing Block-structured AMR in High Level Programming Languages

AMR applications are challenging to implement due to the irregularity introduced by local mesh refinement. For modern parallel programming languages designed to deliver high productivity, expressing computation at logical level meanwhile achieving good performance is a great challenge. In this talk, we examine the programmability and performance difficulties posed by AMR through a case study. Different programming models will be addressed. Particularly, we will focus on optimizing the exchange of ghost values where communication is expressed at fine-grained level.

Tong Wen
IBM Research
tong.wen@gmail.com

Katherine Yelick
UC Berkeley
yelick@cs.berkeley.edu

Jimmy Su
Dept. of computer science
U. C. Berkeley
jimmysu@eecs.berkeley.edu

Phillip Colella
Lawrence Berkeley National Lab
colella@hpcrd.lbl.gov

MS31

Domain-Specific Languages as a Tool for HPC Application Development: Lessons Learned from the Tensor Contraction Engine

I will examine the experience of the Tensor Contraction Engine (TCE) project in developing a domain-specific language (DSL) targeting both high productivity and high performance for a class of problems in quantum chemistry. With current programming environments and tools, we found the task extremely challenging, but rewarding. Next generation general-purpose programming languages, such as those sponsored by the DARPA High Productivity Computing Systems program, may provide a much stronger

supporting environment for future DSLs.

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

MS31

Automatic Creation of Memory-Efficient Linear Algebra Routines

The performance of many linear algebra routines is limited by data retrieval from memory. One way to lower memory costs is to fuse loops of separate routines that access the same data. The performance improvement of the resulting composed routines is directly proportional to the reduction in memory traffic. However, hand coding composed routines is time consuming. In this presentation, we show how the automatic creation of memory-efficient routines reduces program development and execution time.

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

MS31

Using Annotations to Portably Tune Codes for BG/L and BG/P

This talk presents the design of an annotation system for semi-automated performance tuning of arbitrary computations (in Fortran, C, and C++). Multiple code optimization transformations are applied based on user-provided tuning annotations, and an empirical search for the best performing version is carried out. We demonstrate performance improvement achieved through annotation-guided code generation on the Blue Gene/L and Blue Gene/P.

William Gropp
Department of Computer Science
University of Illinois
wgropp@uiuc.edu

Boyana Norris
Argonne National Laboratory
Mathematics and Computer Science Division
norris@mcs.anl.gov

MS31

Automated Data Structure Transformations for Sparse Matrix Kernels

We present an overview of how sparse matrix data structures are represented and transformed in OSKI, a BLAS-like library of automatically tuned sparse matrix kernels. These data structures are hierarchical, and are designed to reduce memory traffic and to expose parallelism. OSKI also provides an embedded scripting language, OSKI-Lua, for finding out what data structure transformations OSKI has applied to a given matrix, and for enabling users to explicitly request a desired data structure.

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

Katherine Yelick
UC Berkeley
yelick@cs.berkeley.edu

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

MS32

Parallel Higher-Order Finite Element Methods for Accelerator Modeling

SLAC's Advanced Computations Department has developed a suite of parallel higher-order finite element codes for high-fidelity, high-accuracy modeling of large, complex accelerator structures. Incorporating advanced numerical algorithms and computational techniques, these codes aim at solving some of the most challenging accelerator problems, and enable virtual prototyping of accelerator components on supercomputers. Application examples including the design of the LCLS rf gun and wakefield calculations of the ILC cryomodule will be presented.

Arno Candel
Stanford Linear Accelerator Center
candel@slac.stanford.edu

MS32

High Performance 3-D Image Reconstruction for Molecular Structure Determination

The single particle reconstruction process from cryo-electron microscopy (cryo-EM) consists of taking a collection of 2-D projection images from various angular orientations and recovering a 3-D volume representation. Accurate volume reconstruction can provide important information on complex molecular structures and their assemblies. However, the reconstruction process can be computationally challenging for large-volume structures, due to massive data and memory requirements. Parallel computing has been crucial in allowing researchers to develop high performance reconstruction algorithms for achieving large-scale molecular structures.

Julianne Chung
Emory University
jmchung@emory.edu

MS32

Solving 3-D and 6-D Schrodinger's Equations Using Multiresolution Adaptive Pseudo-spectral Methods on Leadership Class Computers

Abstract not available at the time of publication.

George Fann
Oak Ridge National Laboratory

fanngi@ornl.gov

MS32
High Performance Nuclear Structure Computation

The configuration interaction (CI) methodology for determining many-body nucleon wavefunctions leads to an extremely large eigenvalue problem. To solve the problem efficiently on a massively parallel computer, the construction of the Hamiltonian and the matrix vector multiplication required in an iterative eigensolver must be performed in an optimal fashion. The algorithmic and implementation issues involved a nuclear CI calculation will be discussed. The performance of some of these calculations will be reported.

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

Chao Yang
Lawrence Berkeley National Lab
CYang@lbl.gov

Philip Sternberg
Lawrence Berkeley National Laboratory
PSternberg@lbl.gov

James Vary, Pieter Maris
Iowa State University
jvary@iastate.edu, pmaris@iastate.edu

MS33
Some Experiment with Mixed Precision on the Cell Processor

This talk describes the design concepts behind implementations of some linear algebra routines targeted for the Cell processor. It describes in detail the implementation of code to solve dense linear system of equations in single precision with iterative refinement of the solution to the full double precision accuracy.

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

MS33
Multifrontal Computations on Graphics Coprocessors

This talk explores the hypothesis that Graphics Processing Units (GPUs) can be exploited to accelerate the speed of a multifrontal linear solver. We demonstrate that even when only factoring a small number of the largest frontal matrices, GPUs can more than double the throughput of the overall sparse matrix factorization.

Robert Lucas
Information Sciences Institute
University of Southern California
rflucas@isi.edu

MS33
Phase-Aware Hardware Adaptivity for Scientific Applications

Scientific applications consist of long lasting memory and cache bound phases, with varying utilization of hardware functional units. We develop an adaptive phase aware control scheme to improve energy efficiency by selecting hardware features and power control modes to match application characteristics. Our phase aware schemes reduce memory access latencies to improve performance and they reduce power through dynamic voltage and frequency scaling. We will report on observed energy savings and performance improvements for a suite of scientific codes.

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

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

Konrad Malkowski
Penn State University
USA
malkowsk@cse.psu.edu

MS33
System Design Constraints for Power Efficient Scientific Computing

The past few years has seen a sea change in computer architecture that will impact every facet of our society as every electronic device from cell phone to supercomputer will need to confront parallelism of unprecedented scale. Whereas the conventional multicore approach (2, 4, and even 8 cores) adopted by the computing industry will eventually run out of steam, the highest performance per watt and per chip area is achieved using manycore technology (hundreds or even thousands of cores). However, fully unleashing the potential of the manycore approach to ensure future advances in sustained computational performance will require fundamental advances in computer architecture and programming models that are nothing short of reinventing computing. This presentation will examine how controlling power efficiency is the driving motivation for manycore systems, and its ramifications for system design, applications and programming models.

John Shalf
Lawrence Berkeley National Laboratory
jshalf@lbl.gov

MS34
Bandwidth Avoiding Stencil Operations

Stencil-based kernels constitute the core of many important scientific kernels on block-structured grids. Our work targets cache reuse methodologies across single and multiple stencil sweeps, examining cache-aware algorithms as well as cache-oblivious techniques on the Intel Itanium 2, AMD Opteron, and IBM Power5. Additionally, we consider stencil computations on the heterogeneous multi-core design of the Cell processor, a machine with an explicitly-managed memory hierarchy. We show that a cache-aware implementation is significantly faster than a cache-oblivious approach, while the explicitly managed memory on the Cell

enables the highest overall efficiency.

Kaushik Datta
UC Berkeley, USA
kdatta@eecs.berkeley.edu

MS34

Communication Avoiding Krylov Subspace Methods

The exponential growth of communication costs relative to computation on modern computers motivates revisiting a previously dismissed set of algorithms: s-step Krylov subspace methods. One iteration of an s-step method has almost the same communication cost as one iteration of its related standard Krylov method, but accomplishes the same work as s of these iterations. We address concerns which precluded the earlier acceptance of these algorithms: performance, numerical stability, and preconditioning.

Marghoob Mohiyuddin, Mark Hoemmen, James Demmel
UC Berkeley, USA
marghoob@eecs.berkeley.edu,
mhoemmen@cs.berkeley.edu, demmel@eecs.berkeley.edu

MS34

Fast Implementations of Akx Kernel

We describe communication avoiding parallel and sequential algorithms for the Akx kernel. Our performance models predict near linear speedups for high latency parallel machines like the Grid even for moderately large problem sizes. For the sequential case, our algorithms are expected to gain impressive speedups for all problem sizes. Finally, our implementation shows at least 3x improvement over the conventional algorithm on a sequential machine.

Katherine Yelick
UC Berkeley
yelick@cs.berkeley.edu

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

Mark Hoemmen
UC Berkeley
USA
mhoemmen@eecs.berkeley.edu

MS34

Communication Avoidance for Sparse Applications Using Full Sparse Tiling

Sparse computations occur in several important science and engineering domains: molecular dynamics simulations, finite element analysis, etc. Due to irregular memory references such as A[B[i]], run-time reordering transformations (RTRT) are required to improve the data locality and parallelism of such applications. We present an RTRT called full sparse tiling and show how it improves the data locality and enables parallelism in various sparse computations.

Michelle Strout
Computer Science Department,Colorado State University
mstrout@cs.colostate.edu

MS35

Turbulent Convective Clouds Simulated with a Spectral Bin Method in MSSG-A

We have implemented a spectral bin cloud microphysics in the Multi-Scale Simulator for the Geoenvironment, MSSG, which is designed for atmosphere and ocean coupled simulation. The cloud microphysics scheme implemented here is highly optimized for large-scale simulation of turbulent convective clouds since turbulence is coupled to cloud collisions. Numerical simulation results of turbulent clouds have confirmed that turbulence enhances the cloud development and consequently increases rainfall.

Takeshi Sugimura
Earth Simulator Center
sugi@jamstec.go.jp

Satoru Komori
Kyoto University
komori@mech.kyoto-u.ac.jp

Ryo Onishi
Earth Simulator Center
JAMSTEC, JAPAN
onishi.ryo@jamstec.go.jp

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

MS35

Atmospheric Flow Simulation Including the Convective Effect over Complex Terrain

We focus on the convective turbulence effect over complex terrain with ultra high resolution. Wind turbulence influenced by complex terrain is simulated with LES model based on the Smagorinsky-Lilly type in Multi-Scale Simulator for the Geoenvironment (MSSG). Resolved complex terrain affects vertical heat/momentum flux transfer due to resolved heterogeneous turbulence. Physical characteristics of vertical fluxes transfer coupled to complex terrain will be discussed based on simulation results under different resolution conditions.

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

Takeshi Sugimura
Earth Simulator Center
JAMSTEC
sugi@jamstec.go.jp

MS35

Multi-scale Simulation for the Geoenvironment (MSSG) with Dynamic Adaptive Mesh Refinement

Climate/weather system is a tightly coupled system with interactions between large-scale and small-scale phenomena. Furthermore earth system is not only multi-scale systems but also a coupled system among atmosphere, ocean and land. Multi-scale simulation for the Geoenvironment (MSSG) is a coupled atmosphere-ocean-land model designed to resolve multi-scale events with dynamic adaptive mesh refinement. Impacts to scientific issues and its implementation with high performance computing on the Earth

Simulator will be introduced in this presentation.

Hiromitsu Fuchigami
NEC Informatec Systems LTD
fuchi@es.jamstec.go.jp

Koji Goto
NEC Corporation
kgoto@jamstec.go.jp

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

MS35

Large Scale Numerical Simulation of Air/Water Interfacial Flows

We have constructed a fully parallelized numerical model for direct simulations of multi-phase fluid dynamics with large problem size. The code is mainly characterized by the following numerical aspects: a robust fluid solver, VSIAM3; an efficient free surface capturing scheme, THINC; and an algebraic multi grid conjugate gradient (AMGCG) method for solving the pressure Poisson equation. We have conducted large-scale simulations for the wind-driven waves and the turbulent structures across the air/water interfaces by using the Earth simulator.

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

Susumu Yamashita
Department of Energy Sciences
Tokyo Institute of Technology
yamashita.s.aa@m.titech.ac.jp

Feng Xiao
Tokyo Institute of Technology
xiao@es.titech.ac.jp

MS36

On the Solution Phase of a Parallel Sparse Solver in a Limited Memory Environment

We consider the parallel solution of sparse linear systems of equations in a limited memory environment. In this context (parallelism and limited memory environment) the performance of the solution phase is critical for many applications. Experiments based on the parallel sparse multifrontal solver MUMPS (MUltifrontal Massively Parallel Solver) will be used to illustrate our discussion.

Patrick Amestoy
ENSEEIHT, Toulouse, France
amestoy@enseeiht.fr

Abdou Guermouche
LaBRI-INRIA futurs
abdou.guermouche@labri.fr

Tzvetomila Slavova, Iain S. Duff
CERFACS (Toulouse)
mila.slavova@cerfacs.fr, iain.duff@cerfacs.fr

MS36

Toward Jacobian-free Preconditioning

We consider the iterative solution of a large sparse system of linear equations whose coefficient matrix is a Jacobian of some function given in the form of a program. Automatic differentiation is capable of efficiently computing Jacobian-vector and transpose Jacobian-vector products in a matrix-free fashion. For preconditioning, however, there is currently no remedy to avoid the assembly of the Jacobian. We introduce a bipartite graph model for parallel preconditioning partially avoiding the Jacobian assembly.

Andreas Wolf, Martin Buecker
Scientific Computing
RWTH Aachen University
wolf@sc.rwth-aachen.de, buecker@sc.rwth-aachen.de

Michael Luelfesmann
Institute for Scientific Computing
RWTH Aachen University
luelfesmann@sc.rwth-aachen.de

MS36

Parallel Preconditioners Using Non-symmetric Permutations

Nonsymmetric matrix pre-orderings, along with scaling, proved to be an effective tool for solving difficult unstructured linear systems of equations. When using this tool to precondition large problems, less expensive alternatives have been proposed that permute matrix rows *and* columns such that their pre-selected subset forms an upper-left submatrix that is well-conditioned. The process is then repeated in a multilevel setting. For those parallel multilevel preconditioners available in the pARMS package, the non-symmetric permutations are applied on local subdomains and their applicability in the global Schur Complement system is explored. Experiments with some difficult circuit and device simulation matrices will be reported.

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

Masha Sosonkina
Ames Laboratory/DOE
Iowa State University
masha@scl.ameslab.gov

MS36

Computing a Class of Bipartite Matchings in Parallel

We consider matchings yielding maximum diagonal products in sparse matrices. An optimal matching is invariant under suitable scalings. We use scaling algorithms to obtain at least one 1.0 per each row and column. We find a maximum cardinality matching using only the entries of magnitude 1.0. If the matching is perfect, we are done. If not, we select a special entry and rescale the matrix. We repeat this process until a perfect matching is found.

Daniel Ruiz
ENSEEIHT, Toulouse, France
daniel.ruiz@enseeiht.fr

Iain S. Duff

CERFACS, FRANCE and RAL, ENGLAND
duff@cerfacs.fr

Bora Ucar
CERFACS
Toulouse, France
ubora@cerfacs.fr

MS37

Continuing Development of Models Based on Icosahedral Geodesic Grids at CSU

The atmospheric modeling group at Colorado State University (CSU) has adopted the icosahedral geodesic grid with Vorononi cells for its global models. The icosahedral grid eliminates the difficulties present in cartesian grids at the poles due to the convergence of meridians in spherical coordinates. The icosahedral grid provides approximately homogenous resolution and isotropic distributing of grid points over the entire sphere. Our approach to parallelization is based on a decomposition of the spherical domain. These pieces of the global domain are called subdomains. We will discuss our strategy to efficiently distribute subdomains to processes and communication between subdomains in the parallel model. In particular we will discuss our approach to hybrid OpenMP and MPI parallelization. We are preparing a Global Cloud Resolving Model (GCRM) with a resolution of about 2 km. The resulting grid will have about 167 million grid points in the horizontal with roughly 50 levels in the vertical. We will discuss our plans for parallel I/O and archiving of data from this model.

Ross Heikes, Celal S. Konor, David Randall
Colorado State University
ross.heikes@mac.com, konor@csu.edu,
randall@atmos.colostate.edu

MS37

Toward a Multi-Purpose, Unstructured Mesh, Finite Element, Marine Model Slim

The multi-purpose finite element marine model SLIM (Second-generation Louvain-la-Neuve Ice-ocean Model) is currently employed to simulate a large number of two-dimensional and three-dimensional flows on the shelves, basin and global scales on the sphere. SLIM resorts to parallel computing and finite element methods on unstructured meshes. The three-dimensional barotropic version of SLIM has been validated and preliminary baroclinic results have already been obtained.

Eric Deleersnijder, Vincent Legat, Jean-Francois Remacle, Richard Comblen, Olivier Gourgue, Sébastien Blaise, Laurent White
Université catholique de Louvain
Belgium
ericd@uclouvain.be, vincent.legat@uclouvain.be,
jean-francois.remacle@uclouvain.be,
richard.comblen@uclouvain.be,
olivier.gourgue@uclouvain.be,
sebastien.blaise@uclouvain.be,
laurent.white@uclouvain.be

Jonathan Lambrechts
U. Louvain
Belgium
jonathan.lambrechts@uclouvain.be

MS37

GPUs, Cells, and Other New Directions for the Weather Research and Forecast Model

The Weather Research and Forecast (WRF) model is a state-of-the art community atmospheric model serving several thousand users in academia, research, and operational forecasting. Designed specifically for high-end systems, the WRF software has demonstrated portability, high performance, and efficient scaling over a range of platforms and several generations of hardware since it was first released in December 2000. The next generation of systems entering petascale computing will present new challenges: overcoming load imbalance, communication overhead, and memory bandwidth limitations on hundreds of thousands of conventional CPU cores crowded onto chip-multiprocessors (CMPs). WRF will also need to exploit new heterogeneous architectures that feature exotic "non-traditional" coprocessors – graphics processing units (GPUs), the Cell Broadband Engine (Cell), and field programmable gate arrays (FPGA) which means dealing with the myriad arcane programming models associated with these processors. In this presentation we will discuss the current status of WRF, performance of WRF as well as foreseeable pressure points moving to petascale - memory footprint, computational intensity, load imbalance, and I/O.

Manish Vachharajani
University of Colorado at Boulder
manish

Michael O. McCracken
San Diego Supercomputer Center
mike

John Michalakes
NCAR
michalak@ucar.edu

MS37

The Cubed-Sphere Finite-Volume Dynamical Core in the GEOS Global Atmospheric Model

The cubed-sphere implementation of the finite-volume dynamical core (FVCORE) has been included as a component within the NASAGoddard Earth Observing System Model (GEOS). GEOSis a system of models integrated using the Earth System Modeling Framework (ESMF) such that components may be interchanged or augmented under a common modeling framework. The cubed-sphere FVCORE was developed as an efficient extension to the original lat-lon dynamical core to overcome the parallel limitations of the algorithm on traditional grids and allow for scalability at petascale levels through 2-dimensional horizontal domain decomposition. A description of the cubed-sphere FVCORE will be presented, along with a discussion of the use of modeling frameworks for parallel implementation and coupling of the cubed-sphere FVCORE.

William M. Putman
NASA
William.M.Putman@nasa.gov

Shian-Jiann Lin
NOAA-GFDL
shian-jiann.lin@noaa.gov

Max Suarez
NASA-GSFC

suarez@nasa.gov

MS38

Reconfigurable Computing: the Emerging Paradigm for High-Performance Computing

High-performance reconfigurable computing holds tremendous promise in addressing the needs of a broad range of emerging applications. Concomitant with technical advantages in speed, power, size, cooling, versatility, etc. are logistical advantages, leveraging the rich heritage and technologies in field-programmable logic. However, research challenges remain in realizing the potential of this new computing paradigm. This presentation will explore challenges and share results from several inaugural projects in a new national center and consortium in this field.

Alan D. George
University of Florida
NSF Center for High-Perf. Reconfigurable Computing (CHREC)
george@hcs.ufl.edu

MS38

Accelerating Cosmology Applications: from 80 Mflops to 8 Gflops in 4 Steps

In this work we outline the general steps necessary to implement an existing application on a reconfigurable computing platform. As the test case example, we use an algorithm for calculating the two-point angular correlation function implemented on an SRC-6 reconfigurable computer. Our implementation utilizes a microprocessor and two reconfigurable processors to accelerate the computationally intensive kernel of the application. The four major steps involved in the algorithm implementation on the reconfigurable computer include exploring deep and wide parallelism, exploiting fixed-point numerical types, and dynamic job scheduling. The end-to-end algorithm execution speedup achieved by our implementation is over 90x as compared to a single microprocessor.

Robert J. Brunner, Adam D. Myers
Department of Astronomy, University of Illinois
rb@astro.uiuc.edu, admyers@astro.uiuc.edu

Volodymyr V. Kindratenko
National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign
kindr@ncsa.uiuc.edu

MS38

FPGA Programming Made Easy

This talk describes the vision that has gone into the design and implementation of CHiMPS (Compiling High level languages Into Massively Pipelined Systems). CHiMPS is a system for implementing applications onto multiple FPGAs and processors. The overriding directive of this vision is 'ease of use' since the target user (HPC scientist) is more interested in science than hardware (FPGA) implementation details. The ability to easily compile and implement existing or legacy languages and code in an FPGA system with sufficent GOPs per Watt is considered much more important than sheer performance.

Jeff Mason
Xilinx
jeff.mason@xilinx.com

MS38

Accelerating Science Applications up to 100X with FPGAs

The performance of two Cray XD1 FPGA systems was evaluated for human genome (DNA and protein) sequence comparisons. Scalable FPGA speedups of 50X (Virtex-II Pro 50) and 100X (Virtex-4 LX160) over a 2.2 GHz Opteron reduced solution times from 8 hours to 5 minutes with speedups increasing with larger query sizes. Results were also obtained for Molecular Dynamics, climate modeling and sparse matrix equation applications.

Olaf O. Storaasli
Oak Ridge National Laboratory
olaf@ornl.gov

MS39

Parallelization of Adaptive Mesh Refinement Algorithms for Multiphysics Application

Adaptive mesh refinement approaches based on structured, hierarchical grids, originally developed for gas dynamics, are now used in a number of multiphysics applications. The added algorithmic complexity of AMR algorithms combined with the requirements of multiphysics applications makes parallelization of these types of algorithms a challenging topic. Key issues in developing efficient parallel implementations include load balancing, data distribution, non-numerical metadata manipulations and linear solvers. In this talk, we will discuss each of these issues in the context of several applications drawn from combustion, astrophysics and subsurface flow. For each issue we will discuss approaches for addressing each of these issues and the impact on overall scaling behavior.

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

MS39

Discussion: Implementing SAMR Methods on Parallel Manycore Systems

SAMR methods are built around a direct (matrix-free) recursive numerical update on sequences of cache-coherent subgrids that overlay each other hierarchically. While this approach ensures maximal computational efficiency, organization and synchronization of subgrids have turned out to be major scalability obstacles. The discussion will gather the presenters' opinions how the upcoming generation of manycore systems might possibly be employed to overcome such problems, while avoiding a re-implementation of existing SAMR software from scratch.

Ralf Deiterding
Oak Ridge National Laboratory
deiterdingr@ornl.gov

MS39

Parallel Implicit Block-based Adaptive Mesh Refinement Schemes for Multi-scale Physically-complex Flows

A somewhat novel computational framework is described

for the solution of multi-scale physically complex flows using parallel block-based adaptive mesh refinement (AMR) schemes on body-fitted multi-block computational mesh. Key elements of the proposed framework will be discussed and include a hierarchical data structure that permits local anisotropic mesh refinement, efficient and scalable parallel implementation via domain decomposition, a parallel implicit time-marching scheme based on a Newton-Krylov-Schwarz (NKS) approach, and a high-order central essentially non-oscillatory (CENO) spatial discretization procedure.

Scott Northrup, Xinfeng Gao
University of Toronto
Institute for Aerospace Studies
northrup@utias.utoronto.ca, gao@utias.utoronto.ca

Clinton P. Groth
University of Toronto
Canada
groth@utias.utoronto.ca

MS39

Scaling and Convergence in Structured AMR for Astrophysical Radiation Diffusion

We present performance results for parallel solution of linear systems arising from astrophysical radiation diffusion equations on meshes created using block-structured adaptive mesh refinement (AMR). The simplest systems are symmetric positive definite and couple grids on a single level of refinement. More complicated nonsymmetric systems arise locally due to higher-order boundary conditions and interfaces between multiple refinement levels, and globally through inclusion of additional physics such as corrections for material motion.

Michael Singer, Louis Howell
Lawrence Livermore National Laboratory
singer3@llnl.gov, howell4@llnl.gov

MS40

New Developments in a State-wide Undergraduate Grid Computing Course

In 2004, we developed an undergraduate Grid computing course that crosses organizational boundaries using resources at several North Carolina universities. We now describe a new more top-down approach to teaching Grid computing beginning with a production-style Grid portal, registration process, and job submission, then leading into infrastructure details. We incorporated seven assignments, including a "capstone" mini-project. We describe our experiences in this new course structure.

Clayton Ferner
Dept. of Computer Science
University of North Carolina at Wilmington
cferner@uncw.edu

Barry Wilkinson
Department of Computer Science
The University of North Carolina at Charlotte
abw@uncc.edu

MS40

Teaching Parallel Computing in 100% Java with the Parallel Java Library

Since more and more computing students are learning Java as their principal programming language, it is becoming imperative to teach parallel programming in Java. Patterned after the OpenMP and MPI standards, the Parallel Java Library (PJ) is a unified API for writing object oriented, 100% Java parallel programs that run on SMP, cluster, and hybrid SMP cluster parallel computers. In this talk I describe PJ's capabilities and the parallel computing courses I teach using PJ.

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

MS40

Teaching Parallel Programming to Lower Division Undergraduates

We discuss a class in parallel programming that we have been teaching for the last four years. The primary audience is freshmen and sophomores with one or two semesters of experience in conventional, serial programming. Students learn the basics of parallel program development in C with MPI, Pthreads, and OpenMP. We discuss content of the lectures, homework, and programming assignments. We also discuss the basic approach to parallel computing which allows us to introduce it to relatively inexperienced programmers.

Peter Pacheco
Department of Computer Science
University of San Francisco
peter@cs.usfca.edu

MS40

Teaching Parallel Processing to Upper-Level Undergraduates

This talk describes a parallel processing course targeted at upper-level undergraduate computer science students. Through a series of readings and programming assignments, students are challenged to think about parallel approaches to problems, and learn three mainstream tools to implement parallel programs: POSIX threads, OpenMP and MPI. Partitioning and dynamic load balancing algorithms are also studied. The assignments culminate with a parallel heat solver using an adaptive quadtree structure and a final project of the students' choosing.

James D. Teresco
Department of Computer Science
Mount Holyoke College
jteresco@mtholyoke.edu

MS41

Dense Linear Algebra on Multicore Architectures: What Kind of Parallelism?

As multicore processors are becoming commodity hardware, linear algebra libraries need to be tuned not only for asymptoticly large matrices, but across the entire spectrum of problem sizes. In this talk we present two different approaches for generating high-performance dense linear algebra libraries for multicore and shared memory systems. High performance is obtained through blocking and careful scheduling in one case, and via multithreaded BLAS in the

other. We present evidence that both solutions have merit in different scenarios.

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

Paolo Bientinesi
Department of Computer Science
Duke University
pauldj@cs.duke.edu

Enrique S. Quintana-Orti
Depto. de Ingeniera y Ciencia de Computadores
Universidad Jaume I
quintana@icc.uji.es

Gregorio Quintana-Orti
Departamento de Ingeniera y Ciencia de Computadores
Universidad Jaume I
gquintan@icc.uji.es

MS41
Spiral: Tackling Parallelism

Spiral (www.spiral.net) is a program and hardware design generation system for linear transforms such as the discrete Fourier transform, discrete cosine transforms, filters, and others. For a user-selected transform, Spiral autonomously generates different algorithms, represented in a declarative form as mathematical formulas, and their implementations to find the best match to the given target platform. Besides the search, Spiral performs deterministic optimizations on the formula level, effectively restructuring the code in ways unpractical at the code or design level.

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

MS41
Matrix Matrix Multiplication for Multi-core Architecture

Multi-core architectures hold the promise of substantially greater computational power. Unfortunately, the historic discrepancy between the speed of processor and memory continues, and thus the quest for improving the reuse of data and reducing of data movement continues. In this talk I discuss the fundamentals of what kinds of data movement can be overlapped with computation and how to reduce the adverse effects of data movement that cannot be overlapped.

Kazushige Goto
The University of Texas at Austin
kgoto@tacc.utexas.edu.

MS41
Multi-dimensional FFTs on Multicore Architectures

We present an algorithm, experiments and an optimization model for the computation of a sequence of multi-dimensional fast Fourier transforms on multicore architectures. There are two phases in each single FFT. In the coupling phase, pair-wise communications are dovetailed with local arithmetic operations. In the uncoupled phase, partitioned and smaller FFTs local to each processing element are executed, leaving the communication network to transfer data from and to the main memory for streaming FFTs.

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

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

Paolo Bientinesi
Department of Computer Science
Duke University
pauldj@cs.duke.edu

MS42
Parallel Primitives for Computation with Large Graphs

Large combinatorial graphs appear in many applications of high-performance computing, including computational biology, informatics, analytics, web search, dynamical systems, and sparse matrix methods. High-performance combinatorial computing is much less well understood than high-performance numerical computing, where there are standard algorithmic primitives and a deep understanding of effective mappings of problems to computer architectures. Here we describe the usage and implementation of sparse generalized matrix-matrix multiplication as a primitive operation for high-performance computing on large graphs.

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

MS42
Kronecker Graphs

Given a large, real graph, how can we generate a synthetic graph that is similar, i.e., it has similar degree distribution, diameter, spectrum, etc.? I will introduce "Kronecker graphs" model, which naturally generates graphs with above properties, and present a fast and scalable algorithm for fitting Kronecker model to real networks. Experiments on large real networks show that Kronecker mimics well the patterns found in target graphs. Once fitted, the model can be used for anonymization, extrapolations, and graph summarization.

Jure Leskovec
Carnegie Mellon University
jure@cs.cmu.edu

MS42
High Performance Combinatorial Techniques for Processing Dynamic Interaction Networks

Graph abstractions are extensively used to model temporal data streams from socio-economic interactions, the world-wide web, and communication networks. For tractable analysis of massive temporal data sets, we require holistic schemes that couple HPC techniques, dynamic graph algorithms, and social network analysis kernels. In this talk, we present a computational framework for the topological analysis of dynamic networks: we experiment with several graph representations, identify key analysis kernels to be optimized, and discuss parallel algorithms for large-scale graph analysis.

Kamesh Madduri
Georgia Institute of Technology
kamesh@cc.gatech.edu

MS42
Array Based Betweenness Centrality

The betweenness centrality metric measures the importance of a node in a graph by examining the number of shortest paths through it. A natural translation of a sequential BC algorithm using linear algebraic notation is shown. A batching parameter shifts this algorithm between $O(E)$ and $O(V^2)$ space, possibly providing constant speedup for small batch values. When implemented in pMatlab, both sequential and parallel algorithms are shown to have reasonable performance compared to their C counterparts.

Eric Robinson
Northeastern University
tivadar@ccs.neu.edu

MS43
Meros: Specialized Preconditioners for Problems with Coupled Simultaneous Solution Variables

Meros is a segregated preconditioning package that provides scalable block preconditioning for problems that couple simultaneous solution variables. Our initial focus has been on methods for solution of the incompressible Navier-Stokes equations. Meros is available as part of the Trilinos Project (http://trilinos.sandia.gov/). This talk will give an overview of the capabilities currently available in Meros as well as plans for future development.

Robert Shuttleworth
Exxon
robert.shuttleworth@gmail.com

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

Victoria Howle
Texas Tech
victoria.howle@ttu.edu

MS43
Model Coupling Toolkit

Abstract not available at the time of publication.

Robert Jacobs
Argonne National Laboratory
jacob@mcs.anl.gov

MS43
Ultra High-resolution Non-hydrostatic Ocean-ice Coupling Simulation Using MSSG-O

Coupling non-hydrostatic ocean physics to sea ice production plays quite important role to intermediate water formation in the Okhotsk Sea and the North Pacific. Ultra high resolution simulation is required to understand the mechanism of the intermediate water. We conduct ultra high resolution simulation with MSSG-O on the Earth Simulator. Impact of dense water production mechanism coupled by intensive vertical mixing in intermediate water will be provided.

Genki Sagawa
Faculty of Engineering
University of Tokyo
sagawa@fluidlab.naoe.t.u-tokyo.ac.jp

Hajime Yamaguchi
Faculty of engineering, university of Tokyo
yama@fluidlab.naoe.t.u-tokyo.ac.jp

Kouji Goto
NEC
goto@jamstec.go.jp

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

MS43
Multiscale Application in Nuclear Engineering

The purpose of this work is the assessment of the coupled codes to study thermal-hydraulic oscillations in nuclear reactors, for instance Boiling Water Reactors, where a strong non-linear coupling exists between neutronic and thermal-hydraulic processes via the void feedback reactivity. We will demostrate the techniques and codes that have been used to produce these couple simulations.

Vicent Vidal
Polytechnic University of Valencia
vvidal@upv.es

Rafael Miro, Gumersindo Verdu
Polytechnich University of Valencia
Valencia, Spain
vvidal@upv.es, gverdu@iqn.upv.es

MS44
Implementing Dense Linear Algebra Algorithms on the STI Cell Processor

The STI CELL processor offers unique computational capabilities through the use of multiple processing cores in a single chip, powerful SIMD instruction set architecture, scratchpad memories and high performance internal bus.

We exploit the capabilities of the CELL processor for dense matrix factorizations by utilizing fine granularity of parallelization, along with mechanisms like block data layout, algorithms by tiles and pipelining of operations. We report performance in excess of 170 Gflop/s from a single chip.

Jakub Kurzak
Innovative Computing Laboratory
University of Tennessee
kurzak@cs.utk.edu

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

MS44

Blocked and Recursice Algorithms for Triangular Tridiagonalization

We present partitioned (blocked) algorithms for reducing a symmetric matrix to a tridiagonal form, with partial pivoting. That is, the algorithms compute a factorization $PAP^T = LTL^T$ where P is a permutation matrix, L is lower triangular with a unit diagonal, and T is symmetric and tridiagonal. The algorithms are based on the column-by-column methods of Parlett and Reid and of Aasen. Our implementations also compute the QR factorization of T and solve linear systems of equations using the computed factorization. The partitioning allow our algorithms 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 factorization and solve in LAPACK.

Gil Shklarski
School of Computer Science
Tel-Aviv University
shagil@tau.ac.il

Sivan A. Toledo
MIT (on sabbatical from Tel Aviv University)
stoledo@tau.ac.il

MS44

Block Gram–Schmidt Orthogonalization

Abstract not available at the time of publication.

Pete Stewart
Department of Computer Science
University of Maryland
stewart@cs.umd.edu

MS44

Programming Algorithms for Matrices Stored by Blocks Made Easy

There has been considerable interest in storing matrices by blocks, possibly recursively, in an effort to achieve better data locality for dense linear algebra libraries. While success has been achieved for individual operations, this interest has invariably led to frustration due to programmability issues when entire libraries are to be created. And thus little practical progress has been made.

In this talk we provide evidence that by embracing the hierarchical nature of matrices (recursively) stored by blocks, formulating algorithms-by-blocks, and coding at a high level of abstraction, the problem becomes relatively simple. Moreover, this allows leaf blocks to become units of data and operations with blocks units of computation, which in turn simplifies the scheduling of operations on multithreaded architectures. These technique are being incorporated in the FLAME dense linear algebra library in combination with a run-time system (SuperMatrix) for scheduling operations.

This research is in collaboration with Paolo Bientinesi, Ernie Chan, Tze Meng Low, Enrique Quintana-Ortí, Gregorio Quintana-Ortí, and Field Van Zee.

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

MS45

A
Charge-Conserving Alternating-Direction-Implicit Maxwell Solver

Over the last few years Zheng et al. have developed an unconditionally stable solver, based on the Alternating Direction Implicit (ADI) algorithm, for the FDTD Maxwell equations. Recent work at Tech-X Corporation has resulted in a charge-conserving ADI Maxwell solver. To achieve good scalability, we are using Concurrent Divide & Conquer to bypass the inherent sequential bottleneck of ADI. Preliminary results with the new solver used by the VORPAL code will be presented.

John R. Cary
Tech-X Corporation
cary@txcorp.com

David N. Smithe, Johan Carlsson
Tech-X Corporation
smithe@txcorp.com, johan@txcorp.com

MS45

Scalable Multigrid Preconditioning for 3D Extended MHD (XMHD)

XMHD supports fast, dispersive waves ($\omega \sim k^2$), which make explicit approaches intractable. Implicit methods can step over such frequencies, but must invert a large algebraic system of stiff nonlinear equations. For this, we employ Newton-Krylov methods, preconditioned with multigrid techniques for scalability. A crucial element of our strategy is the parabolization of XMHD to render it multigrid-friendly. In this talk, we will describe the approach, and provide results of its performance (serial and parallel). In particular, we will demonstrate optimal algorithmic and parallel scalability up to 4096 processors.

Luis Chacon
Los Alamos National Laboratory
chacon@lanl.gov

MS45

An Approach to Parallel Implicit Solvers in Edge Plasma Transport

This presentation discusses benefits and challenges of par-

allel implicit solution strategies in large-scale magnetic fusion simulations, which include a broad range of time and space scales with strong nonlinearities. Experience with preconditioned Newton-Krylov methods in parallel edge plasma transport, with emphasis on scalable algorithms and implementations, is highlighted. Multi-component edge plasmas play key rolls in establishing conditions for high fusion gain and distributing exhaust power to material surfaces.

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

Hong Zhang
Argonne National Lab
hzhang@mcs.anl.gov

Thomas Rognlien
Lawrence Livermore Nat. Lab
trognlien@llnl.gov

MS45

Operator-Based Preconditioning of Stiff Waves in Implicit MHD

Implicit MHD poses algorithmic and computational challenges due to its wide range of spatio-temporal scales, strong anisotropy and nonlinearity. A significant difficulty in such computations is the presence of stiff hyperbolic effects, that may dominate over diffusive processes even at large scale. In this talk, we provide an attractive preconditioning approach for combating such hyperbolic stiffness and describe its parallel implementation. Demonstrations of its efficacy and scalability are shown on representative MHD test problems.

Carol S. Woodward
Lawrence Livermore National Laboratory
Center for Applied Scientific Computing
cswoodward@llnl.gov

Daniel R. Reynolds
University of California, San Diego
Mathematics Department
drreynolds@ucsd.edu

Ravi Samtaney
Princeton Plasma Physics Laboratory
samtaney@pppl.gov

MS46

Towards Fully Adaptive PDE Solvers on Parallel Computers

Devising parallel algorithms to solve partial differential equations is a complex task. It is possible whenever a mesh can be partitioned statically once at the beginning of a calculation. However, it is more complicated when meshes change dynamically, for example between time steps or refinements. We review the problems that occur in this case as well as strategies to deal with the situation currently implemented in the finite element library deal.II and other software packages.

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

MS46

Terascale and PetaScale Adaptive Mesh Refinement Projections and Challenges

The Utah code UIntah developed as part of the CSAFE DOE-funded project has shown good scalability to thousands of processors. An analysis of this code and its components is used to consider what the limit of scalability is for advanced terascale and potential petaflop architectures. The analysis will be validated as far as possible with computational experiments and measurements.

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

Justin Luitjens
SCI Institute and School of Computing University of Uta
jl@cs.utah.edu

Thomas Henderson
SCI Institute and School of Computing University of Utah
tch@cs.utah.edu

MS46

Adaptivity and Parallelism in Constructing Model Based Hazard Maps for Geophysical Mass Flows

We will describe in this talk our use of adaptivity and necessary technologies for it in enabling a class of computations needed for analysis of hazardous geophysical mass flows. Adaptivity and parallel computing are crucial in enabling the ensemble of large scale runs needed to construct hazard maps. We will use adaptivity firstly to capture the flow outlines correctly, secondly to reduce the computation cost by orders of magnitude and finally to select the sampling of input parameters to establish useful output statistics at minimum cost.

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

Keith Dalbey
Dept of Mechanical Engineering
SUNY at Buffalo
kdalbey@eng.buffalo.edu

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

Dinesh Kumar
Dept of Mechanical Engineering
SUNY at Buffalo, Buffalo, NY 14260
dkumar@eng.buffalo.edu

MS46

Scalable Parallel Octree-Based AMR

Adaptive mesh refinement (AMR) is an essential tool for modeling physical problems that possess wide range of scales in space and exhibit dynamic behavior over time. Building on our prior work in scalable multi-thousand-processor octree mesh generation algorithms, we have developed new algorithms for parallel octree-based AMR with

a focus on solving the problem of mantle convection. In this talk, we present the design, implementation, and evaluation of our new techniques.

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

MS47

Scalable On-Line Performance Measurement and Analysis

In recent years, the size of large-scale, distributed-memory supercomputers has risen exponentially. The largest existing supercomputer, IBM's Blue Gene/L, has 212,992 processors. Detecting and diagnosing scalability problems on these system will require tools that can work efficiently at this scale. Previous work has covered scalable techniques for aggregating performance data to tool front-ends. Here, I describe recent work using parallel wavelet transforms and clustering for timely analysis of performance data in situ.

Todd Gamblin
University of North Carolina at Chapel Hill
tgamblin@cs.unc.edu

MS47

Autotuning of Scientific Applications (Final Title TBD)

Abstract not available at the time of publication.

Mary Hall
ISI/USC
mhall@isi.edu

MS47

Sampling-based Measurement and Analysis of Parallel Program Performance

Abstract not available at the time of publication.

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

MS47

Measuring Multiple Counter Domains with PAPI-C

Abstract not available at the time of publication.

Dan Terpstra
Innovative Computing Laboratory
University of Tennessee, Knoxville
terpstra@eecs.utk.edu

MS48

Data Analysis and Visualization for Emerging Architectures

Many simulation and modeling efforts are continuing to produce data at ever increasing rates that are overwhelming our capabilities to explore, hypothesize, document, and thus fully interpret the underlying details. The computer architectures that support these computations, including the tasks for data analysis and visualization, are undergoing a revolutionary change as manufacturers transition to building chips that use an increasing number of processor cores. In addition, graphics hardware that was once designed entirely for the rendering of polygonal primitives has rapidly evolved into a powerful general-purpose processor. While these trends have the ability to provide new capabilities and increase performance, they will do so in a disruptive manner, potentially placing a significant strain on software development activities. In this talk we present the preliminary results of exploiting these architectures to provide an integrated environment for high-performance data analysis and visualization. This effort provides scientists with a simplified programming model for these processor architectures, enabling them to perform both the numerical and visualization tasks necessary for the study of large and complex data sets.

Patrick McCormick
Los Alamos National Laboratory
pat@lanl.gov

MS48

Visualization on Leadership Class Computing Resources

When software rendering is balanced by high numbers of cores and parallel i/o, supercomputers may be a better choice for large-scale parallel data visualization than graphics clusters. Besides performance advantages over hardware rendering at sufficiently large data sizes, this method is suited for in-situ visualization and massive parallelization of software algorithms such as ray casting. We will discuss parallel polygon raster rendering and distributed volume rendering on the Blue Gene L and P machines.

Rob Ross, Tom Peterka
Argonne National Laboratory
rross@mcs.anl.gov, tom.peterka@gmail.com

MS48

An I/O Conscious Data Management Strategy for Parallel Volume Rendering

To allow efficient visualization of ultra scale data sets, one important step is to develop strategies that can increase I/O efficiency and reduce unnecessary data access. In this talk, we describe a novel volume rendering method that can simultaneously achieve the above two goals by taking advantage of the access pattern of data blocks. We consider both the visibility ordering and distribution of transfer functions to group data blocks that are most likely to be retrieved together for any given view. Features in parallel message passing libraries such as MPI-2 are utilized to further increase the rate of data I/O and communication so that the overall volume rendering speed can be optimized.

Han-Wei Shen, Yuan Hong
The Ohio State University
hwshen@cse.ohio-state.edu, hongy@cse.ohio-state.edu

MS48
Scalable Parallel Vector Field Visualizations

Massively parallel supercomputers enable scientists to simulate complex phenomena in unprecedented detail. When scientists attempt to analyze and understand the data generated by the large-scale simulations, the sheer size of the data is a major challenge. To address this challenge, many advances have been made for large-scale data visualization. However, most of the techniques were developed for the visualization of scalar field data, regardless of the fact that vector fields in the same data sets are equally critical to the understanding of the modeled phenomena. Existing vector field visualization methods are mostly base on particle tracing to portray the structure and direction of a flow vector field. When an appropriate set of seed points are used, we can construct paths and surfaces from the traced particles to effectively characterize the flow field. Visualizing a large time-varying vector field on a parallel computer using particle tracing, however, presents some unique challenges. Even though the tracing of each individual particle is independent of other particles, a particle may drift to anywhere in the spatial domain over time, demanding interprocessor communication. Furthermore, as particles move around, the number of particles each processor must handle varies, leading to uneven workloads. In this talk, I will present the design and performance of a scalable parallel pathline construction method for visualizing large time-varying 3D vector fields.

Chaoli Wang
University of California Davis
chawang@ucdavis.edu

Kwan-Liu Ma
Department of Computer Science
University of California at Davis
ma@cs.ucdavis.edu

MS49
Hybrid vs. Flat MPI ? : Experiences in Preconditioned Iterative Linear Solvers for Unstructured Grids

Appropriate parallel programming model is a very critical issue in large-scale scientific computing in this multicore era. In this presentation, optimization issues of preconditioned Krylov iterative solvers for finite-element applications on SMP/multicore cluster architectures will be discussed. ILU-type preconditioning methods are optimized through multicoloring approach, and examples of SAI (Sparse Approximate Inverse) preconditioning are also shown. Experiences in recent architectures, such as *Earth Simulator*, IBM SP series, Hitachi SR series, and *TSUBAME*, are described.

Kengo Nakajima
The University of Tokyo
Department of Earth & Planetary Science
nakajima@eps.s.u-tokyo.ac.jp

MS49
A Mixed Precision Krylov Subspace Method for QCD Simulation

We consider preconditioned Krylov subspace methods for sparse systems of linear equations from large-scale QCD simulations. Our new scheme combines single and double precision arithmetic in order to reduce computation and memory traffic overheads while maintaining double precision accuracy. Some numerical experiments illustrate the efficiency of our method in QCD simulation on multi-core architectures such as the Cell, Intel and AMD processors.

Tetsuya Sakurai
Department of Computer Science
University of Tsukuba
sakurai@cs.tsukuba.ac.jp

Hiroto Tadano
Graduate School of Infomatics
Kyoto University
tadano@amp.i.kyoto-u.ac.jp

Takanori Kihara
Department of Computer Science
University of Tsukuba
kihara@mma.cs.tsukuba.ac.jp

MS49
Auto-tuning Cray LibSci on Multi-core AMD Opteron Processors.

This talk details the drastic changes and improvements made to the Cray XT-LibSci product, the scientific library for Cray XT series systems, in light of on-chip multiprocessing era. We describe how the inter-play of distributed memory parallelism, thread-level parallelism and SIMD short vectorization leads to a complex library design phase and a more challenging optimization problem for even well-understood algorithms. We describe how the focus on automatically tuned scientific libraries, and particularly the automatic generation of sparse BLAS programs provides a better vehicle for the improvement of sparse iterative solvers, the main emphasis of future LibSci releases.

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

MS49
Sparse Matrix Kernels Tuned for Multicore Platforms

We describe on-going work to implement sparse matrix kernels tuned for multicore platforms within OSKI, an automatically tuned library with sparse BLAS-like functionality. We specifically highlight aspects of OSKI's internal architecture designed to support transformations known to be effective for sparse matrix-vector multiply (Williams, et al., in Supercomputing '07). We compare how explicit tuning for multicore compares to an "off-the-shelf" implementation in which we plug OSKI's tuned serial implementations into the PETSc solver framework.

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

MS50
Panel Discussion and Community Input on Addressing the Computational Challenges

Abstract not available at the time of publication.

Eduardo F. D'Azevedo

Oak Ridge National Laboratory
Mathematical Sciences Section
e6d@ornl.gov

MS50

Parallel Matrix Computational Kernels for Petascale Quantum Mechanical Simulations

Parallelization on the computational kernels of multi-length scale structured matrix is essential to reach petascale quantum simulations for strongly correlated materials of technological importance. General-purpose numerical packages, like ScaLAPACK, do not provide an immediate solution for the underlying matrix computations. In this talk, we will present stabilized matrix computation algorithms and the parallel prefix-like techniques for achieving massive parallelization on the calculation of the matrix kernels in quantum simulations.

CheRung Lee
Department of Computer Science
University of California at Davis
cherung@gmail.com

MS50

Parallelization of a Multi-Scale Materials Modeling Code

The multi-scale many body (MSMB) approach being developed in our SciDAC-2 project circumvents current problems in scaling simulations of complex correlated materials by using different algorithms for modeling interactions at different length scales. We will present the computational challenges in employing Petascale computing for such simulations. In particular we will discuss the hybrid parallelization approach, memory requirements and performance bottlenecks of the quantum Monte Carlo (QMC) and Parquet solvers used in the MSMB simulation.

Karen Tomko
Ohio Supercomputer Center
ktomko@osc.edu

Shuxiang Yang
University of Cincinnati
yangphysics@gmail.com

Eduardo F. D'Azevedo
Oak Ridge National Laboratory
Mathematical Sciences Section
e6d@ornl.gov

Andrew Dittes, Mark Jarrell
University of Cincinnati
adittes@gmail.com, mark.jarrell@uc.edu

Thomas Maier
Oak Ridge National Laboratory
maierta@ornl.gov

Cengiz Sen
University of Cincinnati
cengizsen@gmail.com

MS51

Fully Nonlinearly Implicit Resistive MHD Solvers with Textbook Multigrid Efficiency

Magnetohydrodynamic simulations of tokamak fusion plasmas exhibit a large separation of temporal scales. To overcome the temporal stiffness associated with the fast compressive and Alfven waves in MHD, we consider the development of optimal implicit algorithms. We emphasize achieving "textbook" multigrid efficiency in which the set of nonlinear equations is solved to discretization accuracy with a cost equivalent to a few residual calculations. We present preliminary results for some canonical MHD problems: magnetic reconnection and tilt mode instability.

Mark F. Adams
Columbia University
adams@pppl.gov

Ravi Samtaney
Princeton Plasma Physics Laboratory
samtaney@pppl.gov

Achi Brandt
UCLA
abrandt@math.ucla.edu

MS51

Linear Iterative Solves and the Fusion RF Problem

Radio frequency waves in the ion cyclotron frequency range are used to heat and drive currents in magnetized fusion plasmas, including ITER. Conceptually, the problem is simple: Maxwell's equations with a linear plasma response to the applied electric fields. However, the plasma response is nonlocal, and results in large, dense, ill-conditioned matrices. Effective iterative techniques are elusive. Our experience in this quest will be presented along with a discussion of the issues.

E. F. Jaeger
Oak Ridge Natinal Laboratory
jaegeref@ornl.gov

E. F. D'Azevedo, Lee Berry
Oak Ridge National Laboratory
dazevedoef@ornl.gov, berryla@ornl.gov

MS51

Preconditioning and Scalability of Implicit Extended MHD Plasma Simulation by FETI-DP Domain Substructuring

Recent analytical work has proven parallel scalability for the application of the FETI-DP (Finite Element Tearing and Interconnecting, Dual-Primal) domain substructuring to a class of elliptic PDEs, associated with large, sparse, symmetric-positive-definite (SPD) matrices. This presentation describes the extension of this work to a larger class of extended MHD simulations, using empirical, computational methods rather than analytical proofs to demonstrate scalability, applied to a 2D spectral element code.

Alan H. Glasser
Los Alamos National Laboratory
ahg@lanl.gov

MS52

An Octree Method for Multi-source Inverse Prob-

lems with Application to Maxwell's Equations

In this talk we discuss the solution of inverse problems that arise from Maxwell's equations. For many of these problems thousands of sources are used. To obtain computationally feasible algorithms we combine octree discretization of the forward and inverse problem with parallelization. As a result we are able to solve geophysical problems never solved before.

Eldad Haber
Emory University
Dept of Math and CS
haber@mathcs.emory.edu

MS52

Multigrid PDE Solvers on PetaScale Systems

Abstract not available at the time of publication.

Ulrich Ruede
Computer Science
Ulrich.Ruede@informatik.uni-erlangen.de

MS52

Multigrid on 2:1 Balance-constrained Octrees for Finite Element Calculations with Billions of Unknowns

Octree based meshes are used in solving large scale elliptic or evolution partial differential equations with variable coefficients and complex geometries. They offer both the flexibility of unstructured meshes and the simplicity of regular grid discretizations. We propose new parallel algorithms for the construction and 2:1 balance refinement of large linear octrees and using them for the construction of conforming, second-order accurate, trilinear finite element discretizations of self adjoint operators on distributed memory machines. The cost of applying the Laplacian is comparable to that of applying it using a direct indexing regular grid discretization with the same number of elements. Our algorithm has scaled up to four billion octants on 4096 processors on a Cray XT3 at the Pittsburgh Supercomputing Center. The overall tree construction time is under a minute in contrast to previous implementations that required several minutes; the evaluation of the discretization of a variable-coefficient Laplacian takes only a few seconds.

George Biros
University of Pennsylvania
biros@seas.upenn.edu

Rahul Sampath
Mechanical Engineering
Univeristy of Pennsylvania
rahulss@seas.upenn.edu

MS52

Multilevel Solvers on GPUs

We describe data structures and computational abstractions for defining and operating on locally refined octree meshes on GPUs. This abstraction layer allows algorithms to be expressed at a higher level simplifying the implementations of multilevel solvers and precondioners that use the GPU to accelerate their computations. We describe examples using these abstractions and show that their performance compares favorably to that of low-level handwritten solvers.

George Turkiyyah
University of Washington
georgeturk@gmail.com

MS53

Measuring Productivity in a Graduate Parallel Computing Course

High-performance parallel programming is often considered an arcane skill. In our course "Applied Parallel Computing," graduate students from several science and engineering departments study parallel architectures, algorithms, programming, and performance tuning. Over the past four years, we have collected (with student permission) highly detailed data on the students programming process. This data enables us to replay the entire programmer experience, yielding comparative analyses of programming models in terms of effort and performance.

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

Vikram Aggarwal
University of California, Santa Barbara
vikram@mayin.org

Imran Patel
Dept of Computer Science
University of California, Santa barbara
imran@cs.ucsb.edu

Viral Shah
University of California, Santa Barbara
Interactive Supercomputing
viral@cs.ucsb.edu

MS53

HPC Curriculum Integration: The Pathways to Supercomputing and SC Education Program's Impact on Curriculum

The coalition of the National Computational Science Institute (NCSI), the Supercomputing Conference's (SC) Education Program, the SC Pathways efforts and the Computational Science Education Reference Desk (CSERD), have promoted curriculum development for high performance computing and computational methods in computer science, chemistry, biology, and physics. This talk focuses on identifying gaps in the existing curriculum structure, particularly as they are a part of the materials necessary to teach scientists and students how to leverage the next generation of Petascale computational resources for research and education.

Charlie Peck
Earlham College
charliep@cs.earlham.edu

Dave Joiner
Kean University
Union, NJ
djoiner@kean.edu

Tom Murphy

Contra Costa College
San Pablo, CA
tmurphy@contracosta.edu

Paul Gray
University of Northern Iowa
Computer Science
gray@cs.uni.edu

MS53
Teaching Concurrency

Many parallel algorithms employ restricted forms of concurrency such as barrier synchronization or message-passing on a regular grid. Yet programmers need to understand general problems of process coordination and communication to maximize parallelism, avoid subtle timing problems, and control costs. This talk describes 25 years of teaching concurrency, and the tradeoffs between focusing on implementation details of synchronization primitives for multitasking and on higher-level issues of their use.

David Hemmendinger
Department of Computer Science
Union College
hemmendd@union.edu

MS53
Parallel Scientific Computing: Remembrances and Reflections on Supercomputing

Reflections on a career in research and teaching scientific supercomputing. Remembrance of modifying code and adapting algorithms for use on supercomputers as they evolved from serial to vector to parallel to parallel-vector to clusters. As well as, comments on teaching scientific parallel computing to undergraduate students.

David R. Kincaid
Univ. of Texas at Austin
TICAM
kincaid@cs.utexas.edu

MS54
Petaflop Challenges in Lattice QCD

Lattice QCD is arguably the only feasible approach to evaluate quantum chromodynamics (QCD), the theory which explains the quarks as the constituents of matter. Lattice QCD requires very heavy computer simulations which can take up to several months on current parallel supercomputers. In this talk we will motivate how this enormous demand for computing powers arises and why these computations are suitable for highly parallel computers. We will also address recent progress in modelling and algorithmic development.

Andreas J. Frommer
Bergische Universitaet Wuppertal
Fachbereich Mathematik und Naturwissenschaften
frommer@math.uni-wuppertal.de

MS54
Towards Petascale Adaptive Simulations of Mantle Convection

Mantle convection is the principal control on the thermal and geological evolution of the Earth. It is central to our understanding of the origin and evolution of tectonic deformation, the evolution of the thermal and compositional states of the mantle, and ultimately the evolution of the Earth as a whole. Mantle convection is an important driver for petascale computing, due to the wide range of length and time scales involved. Our goal is to conduct high resolution mantle convection simulations that can resolve thermal boundary layers and faulted plate boundaries, down to 1 km scales. To enable this (local) resolution, we are developing `Rhea`, a new generation mantle convection code incorporating parallel adaptive mesh refinement/coarsening algorithms designed to scale to hundreds of thousands of processors. We discuss parallel performance on *Ranger*, the new 500 Teraflops system at TACC.

Tiankai Tu
DE Shaw Research
tiankai.tu@gmail.com

Georg Stadler, Carsten Burstedde, Omar Ghattas
University of Texas at Austin
georgst@ices.utexas.edu, carsten@ices.utexas.edu,
omar@ices.utexas.edu

Shijie Zhong
University of Colorado, Boulder
shijie.zhong@colorado.edu

Michael Gurnis
Caltech
gurnis@gps.caltech.edu

George Biros
University of Pennsylvania
biros@seas.upenn.edu

Lucas Wilcox
University of Texas at Austin
lucasw@ices.utexas.edu

MS54
High Performance Computing Strategies for Adaptive FEM Simulations

We propose the relocation of grid points preserving the topological structure of the mesh ("grid deformation") as technique for grid adaptation in FEM. Our approach allows for maintaining locally structured grids during the adaptation process in contrast to the widespread elementwise h-adaptivity and requires little additional numerical effort only. We discuss the derivation and realisation of our method, its asymptotic complexity and its accuracy. We show applications to examples in CFD and CSM.

Stefan Turek
Universität Dortmund
stefan.turek@mathematik.uni-dortmund.de

M. Grajewski
Universität Dortmund, Germany

MS55
Supporting Visual Data-mining by Dimension Reduction

To support visualization of high dementional data, we introduce algorithm "HyperMap" for mapping data to 3-dimensional space, and develop a novel parametric visu-

alization. HyperMap holds the linear computation complexity, but releases the restriction that each dimension has only two pivot objects. Having more than 2 pivots to represent one dimension enables less information-loss. Tunning the weights of pivots, the scatter of data can be observed in various viewpoints, the impacts of bad pivot object is reduced, too.

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

Tetsuro Ogi
Academic Computing and Communications Center
University of Tsukuba
tetsu@cc.tsukuba.ac.jp

Hanxiong Chen
S.I.E.,
University of Tsukuba
chx@cc.tsukuba.ac.jp

MS55

InfoSpace Governance Technologies for People Working Together

Information spaces created by cooperative knowledge workers are difficult to manage, because such an information space spreads over distributed and heterogeneous systems and data integrity and circulation is not supported in systematic ways. We approach to the problem by developing technologies to manage metadata on the information spaces to control the status and circulation of information in the spaces. This talk addresses the technological issues and their roles in the visual data-mining project.

Keishi Tajima
Graduate School of Informatics
Kyoto University
tajima@i.kyoto-u.ac.jp

Atsuyuki Morishima
Graduate School of Library, Information and Media Studies,
University of Tsukuba
mori@slis.tsukuba.ac.jp

MS55

VINDAM: Data Visualization and Data Mining in Tele-immersion Environment

In this study, the tele-immersion environment was constructed by connecting several CAVE-like immersive projection displays through the broad-band network. In this environment, remote users can communicate with each other using the video avatar in the shared virtual world. In addition, they can retrieve the necessary data from the database through the network and handle them synchronously. This environment was applied to the collaborative data visualization and the visual data mining in the immersive virtual world.

Takashi Furumura
Earthquake Research Institute
University of Tokyo
furumura@eri.u-tokyo.ac.jp

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

Tetsuro Ogi
Academic Computing and Communications Center
University of Tsukuba
tetsu@cc.tsukuba.ac.jp

Hanxiong Chen
S.I.E.,
University of Tsukuba
chx@cc.tsukuba.ac.jp

Osamu Tatebe
Center for Computational Sciences,
University of Tsukuba
tatebe@cs.tsukuba.ac.jp

Atsuyuki Morishima
Graduate School of Library, Information and Media Studies,
University of Tsukuba
mori@slis.tsukuba.ac.jp

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

Kengo Nakajima
The University of Tokyo
Department of Earth & Planetary Science
nakajima@eps.s.u-tokyo.ac.jp

MS55

Gfarm Grid File System for Visual Distributed Data-Mining

Gfarm Grid file system has been developed for facilitating reliable file sharing and high-performance data computing in a Grid across administrative domains. It is a scalable virtual file system federating local file systems of cluster nodes. This talk describes the design and implementation of a secure, robust, scalable and high-performance Gfarm file system.

Osamu Tatebe
Center for Computational Sciences,
University of Tsukuba
tatebe@cs.tsukuba.ac.jp

PP0

A Parallel Algorithm for Elliptic Eigenvalue Problems on Polygonal Domains Using Spectral Method

Here we show a technique to solve the elliptic eigenvalue problems on polygonal domains using *h-p* spectral element method on parallel computers. We develop a MPI-based parallelization algorithm, to seek a solution which minimizes a weighted squared norm of the residuals in PDEs and a norm of the residuals in the boundary conditions and enforce continuity by adding a term which measures the jump in the function and its derivatives at inter element boundaries to the functional being minimized.

Pravir Dutt
Department of Mathematics,
IIT Kanpur, India

pravir@iitk.ac.in

Lokendra K. Balyan
Indian Institute of Technology Kanpur
balyan@iitk.ac.in

R.K.S. Rathore
IIT Kanpur, India
rksr@iitk.ac.in

PP0

Parallel Processing of Ultrasound Propagation in Structures Using the Finite Integration Technique for Use in Structural Health Monitoring

Structural health monitoring combines sensors with extensive knowledge of the waveguide physics and automated signal processing algorithms in order to accurately characterize damage progression in structures. In this poster we discuss using the finite integration technique to simulate the propagation of ultrasonic waves through plates, pipes and aircraft stiffeners. We split the multidimensional simulation space across parallel processors alleviating previous memory shortage and extended runtimes issues, which allows testing of numerous structural conditions without experimentation.

Kevin Rudd
NRL
kerudd@wm.edu

Jill P. Bingham, Chris Bording, Mark Hinders
The College of William and Mary
jxbing@wm.edu, rcbord@wm.edu, hinders@wm.edu

PP0

Visual Data-mining and Tele-immersive Communication System for Seismic Wave Analysis

In many fields of numerical simulations and data analyses, it is very important how to analyze data visually. As one of application, research for "Analysis of Seismic Wave" is carried out enormously to wide area. To this application, especially, our system can provide "3D visual analysis by data-mining under tele-immersive communication environment" based on individual data management systems and provide "automatic performance tuning" of adjustment parameter in order to reduce noisy data.

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

Takashi Furumura
Earthquake Research Institute
University of Tokyo
furumura@eri.u-tokyo.ac.jp

Kengo Nakajima
The University of Tokyo
Department of Earth & Planetary Science
nakajima@eps.s.u-tokyo.ac.jp

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

Tetsuro Ogi
Academic Computing and Communications Center
University of Tsukuba
tetsu@cc.tsukuba.ac.jp

Hanxiong Chen
S.I.E.,
University of Tsukuba
chx@cc.tsukuba.ac.jp

Osamu Tatebe
Center for Computational Sciences,
University of Tsukuba
tatebe@cs.tsukuba.ac.jp

Atsuyuki Morishima
Graduate School of Library, Information and Media Studies,
University of Tsukuba
mori@slis.tsukuba.ac.jp

PP0

Tools and Techniques for Exploiting the Vector Multimedia Engine on the Power6 Architecture

This presentation, in poster format, will present various tools and techniques now available in the POWER6 and POWERPC 970 architectures to further accelerate HPC (High Performance Computing)application codes. We present examples of using the VMX (Vector Multimedia eXtensions) engine to provide SIMD vectorization capabilities for speeding up floating point, integer, and character operations. This presentation will include, among other features, use of automatic vectorization, profiling and performance results.

John Martine
I.B.M. Corporation
jmartine@us.ibm.com

PP0

Intelligent Domain Partitioner for Parallel Finite-Element Applications

In parallel finite-element applications, strategy of domain decomposition strongly affects the convergence of preconditioned iterative linear solvers. Convergence becomes worse in problems with local heterogeneity and/or discontinuity of the field. In this poster, outline of an intelligent domain partitioner, which can consider heterogeneity of element shapes, sizes and material properties, is described. This type of partitioner is also very important for automatic selection of parameters in parallel preconditioning methods for ill-conditioned problems.

Kengo Nakajima
The University of Tokyo
Department of Earth & Planetary Science
nakajima@eps.s.u-tokyo.ac.jp

PP0

Accurate Computations of Accelerator Cavity Frequencies Using VORPAL

Efficient and accurate computations of frequencies in complex accelerator cavities from data generated by a Finite Difference Time Domain (FDTD) code VORPAL are presented. The results obtained through massively parallel

computations show that the calculation of frequencies to a few parts in 10^5 is achievable. This is more accurate than manufacturing variations, and so inverse methods can now be applied to designs of such systems as the crab cavities for the proposed International Linear Collider.

Chet Nieter
Tech-X Corporation
nieter@txcorp.com

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

Serguei Ovtchinnikov
Tech-X Corporation
serguei@txcorp.com

PP0

Exploring the Potential for Using Thread Level Speculation Hardware for Overlapping Communication and Computation in Scientific MPI Applications

We propose to leverage mechanisms typically used for Thread Level Speculation (TLS) and/or Transactional Memory (TM), to tolerate communication delays, by keeping track of pending receive buffers and allowing the application to execute past blocking *MPI_Recv* calls. We divide the application code into different regions (or epochs), and impose restrictions on when speculation is allowed. We characterize LAMMPS, HYCOM, and POP to determine representative sizes of read/write sets, and potential lengths of speculative instruction regions.

Peter Kogge
Department of Computer Science and Engineering
University of Notre Dame
kogge@cse.nd.edu

Jeffrey Vetter
ORNL
vetter@ornl.gov

Srinivas Sridharan
Department of Computer Science and Engineering
University of Notre Dame
ssridhar@nd.edu

Collin McCurdy
Computer Science and Mathematics Division
Oak Ridge National Laboratory
cmccurdy@ornl.gov

PP0

Parallel Recurrence Calculation on Cell Be

Although Cell BE (Broadband Engine) was originally designed for a processor suitable for streaming applications such as media processing, it also can be applied to other high performance computing applications. We implement a recurrence equation solver on the Cell BE by using the P-scheme algorithm and achieves the speedup of about 4 with 6 SPEs compared with 1 PPE using the normal solver.

Akiyoshi Wakatani
Konan university
wakatani@konan-u.ac.jp