CP1

A Cartesian Structured AMR Framework for Parallel Fluid-Structure Interaction Simulation

Time-accurate fluid-structure interaction simulations of shock waves impinging on deforming solid structures benefit significantly from the application of dynamic mesh adaptation in the fluid. We present an extension of the patch-based parallel fluid solver framework AMROC for this problem class. Special attention is given to a scalable implementation of the hierarchical mesh refinement method, efficient parallel inter-solver communication routines, and the effective transformation of the evolving solid boundary into a Cartesian level set function.

Ralf Deiterding, Sean P. Mauch
Center for Advanced Computing Research
California Institute of Technology
ralf@cacr.caltech.edu, sean@cacr.caltech.edu

CP1

Comparison of Parallel Decomposition Strategies for Turbulent Channel Flow Simulations

As we push the record of Reynolds number higher and higher in turbulent channel flow simulations, finding efficient parallel decomposition strategies becomes more and more important. Most of large computations are now using the message-passing model such as MPI, where the communication performance is a critical factor in the overall code performance. Conventional decomposition strategies that used to be adequate on small- to mid-size clusters can incur a disastrous communication cost at Tera-scale clusters, where typically hundreds or thousands processors are connected through complex network topology. Typical turbulent channel flow simulations have two homogenous directions, where efficient FFTs can be used. However, using domain decomposition across those directions is discouraged to minimize the communication cost. In this study, we will compare several decomposition strategies and provide guidelines to achieve the best overall performance under various conditions.

Junwoo Lim
Pittsburgh Supercomputing Center
Carnegie Mellon University
jlim@psc.edu

CP1

Parallelization of Fourier Spectral Element Code

In this paper, we implemented new parallel models for Fourier Spectral Element code in order to use large number of processors. New models can use more processors than before by doing domain decomposition in two directions, while the old one only in one direction. This overcomes the bottleneck of Fourier planes limit in parallel computing. Detailed implementation and benchmark results will be given. The new models can be run on thousands of processors in practical simulation. Some results of large scale turbulent simulations using new parallel models will be given.

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

CP2

Performance Evaluation of a 3D Electrodynamics Simulation on Massively Parallel Computers

We evaluate the performance for a 3D Maxwell FDTD code on two novel massively parallel computers, the 12,288 processor QCDOC at Brookhaven National Laboratory and a 1024 node BlueGene/L at Argonne National Laboratory and the Alpha cluster at Pittsburgh Supercomputing Center. QCDOC is the more distinctive of these architectures (e.g., higher-dimensional mesh) designed for QCD applications. The FDTD application will be one of the first non-QCD applications evaluated from a scaling perspective on this machine.

Robert E. Bennett
Computation Science Center
Brookhaven National Laboratory
robertb@bnl.gov

Nicholas D'Imperio
Computational Science Center
Brookhaven National Laboratory
dimperio@bnl.gov

CP2

Fast Electrostatic Force Calculation

Fast Multipole Method (FMM) and Particle Mesh Ewald (PME) are well known fast algorithms to evaluate long range electrostatic interactions in Molecular Dynamics. However, those algorithms suffer from poor efficiency on large parallel machines especially for problem size under 100,000 atoms. A variation of FMM based on Plane Wave Expansions is introduced and its parallel efficiency is compared with PME through detail time measurements.

Amirali Kia
PhD Student
Stanford University
akia@stanford.edu

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

Daejoong Kim
PhD Student
Stanford University
daejoongkim@stanford.edu

CP2

Simulation of Laser Propagation in a Plasma with a Frequency Maxwell Equation

The aim is to perform numerical simulations of the propagation of a laser in a plasma. At each time step, one has to solve a variable coefficient Helmholtz equation in a domain consisting in one or two hundreds of millions of cells. The algorithm is based on a blend of a domain decomposition method with a fast variable coefficient Helmholtz solver and a QD diagonalization of an auxiliary matrix. We show some results which are obtained on a parallel architecture at CEA-DAM (France).

Frédéric Nataf
CMAP, Ecole Polytechnique
France

nataf@cmap.polytechnique.fr

CP2

Parallel Performance of An Iterative Solver for Helmholtz Problems

We solve the linear system obtained from the discretization of the time-harmonic wave equation in a heterogeneous medium with a Bi-CGSTAB iterative method, using a complex shifted-Laplace operator as preconditioner. A multigrid method is used to approximately invert the preconditioning operator. We investigate the parallel performance of a matrix-free algorithm for this solver and show results for large scale problems in seismic wave propagation in the Earth in the time and frequency domain.

Kees Vuik
Department of Applied Mathematical Analysis
k.vuik@ewi.tudelft.nl

Christina D. Riyanti
Delft Institute of Applied Mathematics
Delft University of Technology
C.D.Riyanti@ewi.tudelft.nl

Alex Kononov
Computational Physics Group, PCMT
Delft University of Technology
a.v.kononov@tnw.tudelft.nl

Kess Oosterlee
Delft Institute of Applied Mathematics
Delft University of Technology
c.w.oosterlee@math.tudelft.nl

CP3

Parallelisation and Scalability Issues in a Multilevel EHL Solver

The computation of numerical solutions to elastohydrodynamic lubrication problems is only possible on fine meshes by using a combination of multigrid and multilevel techniques. In this paper we show how the parallelisation of both multigrid and multilevel multi-integration for these problems may be accomplished and discuss the scalability of the resulting code. A performance model of the solver is constructed and used to perform an anlysis of the results obtained. Results are shown with good speed-ups and excellent scalability for both shared and distributed memory architectures and in agreement with the model.

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

Christopher E. Goodyer
University of Leeds
ceg@comp.leeds.ac.uk

CP3

Application of Machine Learning to Selecting Solvers for Sparse Linear Systems

We demonstrate how machine learning techniques, like boosting, can be used to select iterative linear solvers that match given application characteristics. We discuss the research issues involved and our current implementation of this automatic selection process. These include,creation, storage, and retrieval of a representative database; mapping the selection problem into a binary classification problem commensurate with machine learning algorithms; and software to bridge between different codes and provide a user-friendly interface.

Sanjukta Bhowmick
Department of Applied Physics and Applied Mathematics
Columbia University
bhowmick@cse.psu.edu

Erika Fuentes
University of Tennessee
efuentes@cs.utk.edu

Victor Eijkhout
University of Texas at Austin
Texas Advanced Computer Center
eijkhout@cs.utk.edu

Yoav Freund
Department of Computer Science and Engineering,
University of California, San Diego
yfreund@ucsd.edu

David E. Keyes
Columbia University
Department of Applied Physics & Applied Mathematics
kd2112@columbia.edu

CP3

Performance Studies of a Parallel Global Search on System X

We evaluate a parallel implementation for a global search algorithm DIRECT on System X, a 2200 processor cluster. The parallel scheme addresses design challenges in memory requirement and data dependency. The performance studies focus on data structure efficiency, program scalability, resource utilization, and load balancing. We analytically explore the strength and weaknesses of DIRECT in this setting, identify several key sources of inefficiency, and experimentally evaluate a number of effective solutions in the parallel implementation.

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

Jian He
Virginia Polytechnic Institute and State University
Department of Computer Science
jihe@vt.edu

Alex Verstak
Department of Computer Science,
Virginia Polytechnic Institute and State University
averstak@vt.edu

Tom Panning
Virginia Polytechnic Institute and State University
tpanning@vt.edu

CP3

Parallelising An Existing Object-Oriented Simula-

tion Framework

COOLFluiD is a state-of-the art object oriented framework, written in C++. It uses advanced software techniques to combine modularity, flexibility and speed, and was recently parallelized for distributed memory machines. COOLFluiD now provides transparent parallelisation for both end-users and code developers, while still allowing compilation without parallelisation. Input and output is fully abstracted and uses parallel IO whenever possible. This paper discusses the methods and algorithms used, and elaborates on our experiences with the framework for solving engineering problems on high performance platforms.

Tiago Quintino
von Karman Institute for Fluid Dynamics
quintino@vki.ac.be

Andrea Lani
Aerospace Depart.
von Karman Institute for Fluid Dynamics
lani@vki.ac.be

Stefan Vandewalle
Scientific Computing Research Group
KULeuven
stefan.vandewalle@cs.kuleuven.be

Stefaan Poedts
Centre for Plasma Astrophysics
KULeuven
stefaan.poedts@wis.kuleuven.be

Dries Kimpe
Scientific Computing Research Group
KULeuven
Dries.Kimpe@cs.kuleuven.be

CP4

Architecutre and Desing of Gridsat: A Distributed Large Scale Satisfiability Solver

GridSAT is a sophisticated and complex grid application which can simultaneously run on a large set of geographically distributed computational resources including several National Science Foundation computing centers and desktop machines and smaller scale, locally managed clusters. GridSAT was able to solve especially hard and previously unsolved problems. The GridSAT solving power is currently accessible to users through a simple portal interface at http://orca.cs.ucsb.edu/sat_portal that automatically employs a variety of nationally distributed resources on behalf of each user.

Wahid Chrabakh
University of California, Santa Barbara
chrabakh@cs.ucsb.edu

CP4

A Generative Programming Approach for Master/Slave Scientific Applications

The purpose of this work is to simplify the development of parallel scientific applications that follow a master/slave pattern and use the MPI library. Generative Programming is used to provide high-level constructs which enable programmers to focus on writing the specific code for the master and slaves while the generated communication details are hidden. We provide test cases to evaluate the programmer effort required, the scalability, and the performance of applications that use our approach.

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

Francisco A. Hernandez, Kevin Reilly
University of Alabama at Birmingham
hernandf@cis.uab.edu, reilly@cis.uab.edu

CP4

Task-Based Software Development for Scientific Computing

We extend the model of task-parallel executions so that the same coordination program can be executed on different heterogeneous platforms. The extended model is particularly suited for large applications consisting of independent modules which can be mapped onto different parts of distributed platforms. We show that a suitable representation of the execution activities is crucial for combining a flexible multi-level specification with a dynamic scheduling that can be adapted to a dynamically changing execution environment.

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

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

CP4

Integrated Performance Monitoring of Highly Parallel HPC Workloads

Integrated Performance Monitoring (IPM) is an infrastructure for profiling HPC workloads. The software is designed to be sufficiently lightweight and easy to use that profiles can be taken in the course of normal production runs. Analysis of application profiles making up a workload yields valuable insight into parallel applications and architectures. IPM is open source, relies on portable software technologies and is scalable to tens of thousands of tasks.

David E. Skinner
NERSC/LBL
dskinner@nersc.gov

CP4

Memory Adaptation of Scientific Applications Via An Application-Level, Remote Memory Library

MMlib is an application-level, runtime library that controls the DRAM allocations of specified objects, enabling scientific applications to execute optimally under variable memory availability. We extend MMlib so that managed objects are not stored on disk, but on several remote servers, and are fetched transparently when needed through MPI. Advances in network speeds, limited local disk on MPPs, and large granularity suggest the benefits and often the need for our approach. To our knowledge, our remote memory

approach is the first entirely at the application-level.

Dimitrios Nikolopoulos
Dept. of Computer Science
College of William and Mary
dsn@cs.wm.edu

Andreas Stathopoulos
College of William & Mary
Department of Computer Science
andreas@cs.wm.edu

Chuan Yue
Department of Computer Science
College of William and Mary
cyue@cs.wm.edu

CP5

Evaluating the Implementation of An Iterative Solver in An Ocean Model

We present a tool aimed at improving the design and implementation of memory efficient sparse linear algebra algorithms through automated memory analysis. Our automated memory analysis uses a language processor to predict the data movement required for an iterative algorithm based upon a Matlab implementation of that algorithm. We demonstrate how automated memory analysis is applied to evaluate implementation choices in the conjugate gradient solver in the Parallel Ocean Program (POP).

John M. Dennis
Scientific Computing Division
National Center for Atmospheric Research
dennis@ucar.edu

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

CP5

Testing Unsaturated Flow Finite Element and Finite Volume Groundwater Programs on High Performance Computers Using Analytical Solutions

Three-dimensional steady-state and transient analytical solutions of Richards equation for unsaturated groundwater flow have recently been derived, and these solutions are being used to test performance and accuracy of a parallel unstructured mesh finite element program and a structured grid finite volume program. Variables in this study include grid/mesh size, time step size, Picard versus Newton iterations, and the number of processors. The analytical solutions and the performance/accuracy results will be presented.

Fred T. Tracy
Engineer Research and Development Center
Waterways Experiment Statiuon
tracyf@wes.army.mil

CP5

Homme: A High-Performance Scalable Atmospheric Modeling Framework

The High Order Method Modeling Environment is a framework that provides the tools necessary to create a high-performance scalable Atmospheric Model. It currently supports spectral element and discontinuous Galerkin schemes. We provide scaling results for a spectral element based 3D moist primitive equations simulation on a 10 km horizontal resoultion on 32K processors and a discontinuous Galerkin based shallow-water equations problem on a 40 km horizontal resolution on 2K processors of a Blue Gene/L system.

Ramachandran Nair
Scientific Computing Division
National Center for Atmospheric Research
rnair@ucar.edu

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

John M. Dennis, Jim Edwards
Scientific Computing Division
National Center for Atmospheric Research
dennis@ucar.edu, jedwards@ucar.edu

Michael Levy
Department of Applied Mathematics
University of Colorado at Boulder
levym@colorado.edu

Richard Loft
Scientific Computing Division
National Center for Atmospheric Research
loft@ucar.edu

Henry Tufo
NCAR
University of Colorado at Boulder
tufo@ucar.edu

Theron Voran
Department of Computer Science
University of Colorado at Boulder
voran@ucar.edu

Steve Thomas
National Center for Atmospheric Research
thomas@ucar.edu

CP6

A Multi-Method Ode Software Component for Parallel Simulation of Diesel Engine Combustion

Modeling of Diesel engines requires the solution of a large number of large systems of ODEs arising from detailed models of combustion. In these models, the concentrations of the chemical species change with very different time scales, introducing high degrees of stiffness in the critical phase where combustion starts. In this work we propose a parallel combustion solver based on a multi-method ODE software component for effective solution of combustion in simulations of Diesel engines.

Pasqua D'Ambra
Institute for High-Performance Computing and Networking
(ICAR-CNR)
pasqua.dambra@na.icar.cnr.it

Paola Belardini, Claudio Bertoli

Engine Institute (IM-CNR)
p.belardini@im.cnr.it, c.bertoli@im.cnr.it

Stefania Corsaro
Dept. of Statistics and Mathematics for Economic Research
University of Naples "Parthenope"
stefania.corsaro@uniparthenope.it

CP6

Multiple Projection Algorithms on the Peer to Peer Lake/yml Environment

The LAKe/YML environnement is constituted by the integration of the LAKe component-oriented numerical library to the peer to peer YML system. YML is a framework allowing workflow graph description adaptable to several middeware. We present some multiple projection methods to compute a few eigen-elements of a very large non-Hermitian matrix on LAKe/YML. The open problems posed by these methods from the numerical and/or environmental viewpoint will be pointed out.

Olivier J. Delannoy
PRISM Laboratory
olivier.delannoy@firstream.net

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

CP6

Parallelization of the Diffusion Equation Method for Global Optimization

We use a discrete version of the Diffusion Equation Method (DEM) for global optimization to solve a problem from magnetotelluric geoprospecting. The main computational effort for this application is in repeatedly evaluating the objective function, which fortunately can be done in a highly parallel manner. We discuss our multilayer parallel implementation of the objective function and discrete DEM, and analyze the performance and scalability of our approach.

Rebecca J. Hartman-Baker
Oak Ridge National Laboratory
hartmanbakrj@ornl.gov

Michael Heath
University of Illinois at Urbana-Champaign
heath@cs.uiuc.edu

CP6

Parallel Fem Modeling of Semiconductor Devices

This talk presents preliminary results for an unstructured mesh stabilized finite element (FE) formulation of the drift diffusion equations to model semiconductor devices. A brief motivation, overview of the formulation and preliminary results are presented. Performance and scaling results for domain decomposition and multilevel solution methods on large-scale simulations will be presented. Unique challenges in modeling this class of problems will be discussed, including the presence of very strong source terms and gradients. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Paul Lin
Sandia National Laboratories
ptlin@sandia.gov

Gary Hennigan, Rob Hoekstra, Joe Castro, Debbie Fixel
Sandia National Labs
glhenni@sandia.gov, rjhoeks@sandia.gov,
jpcastr@sandia.gov, dafixel@sandia.gov

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

Eric Phipps
Sandia National Laboratories
etphipp@sandia.gov

Roger Pawlowski
Sandia National Labs
rppawlo@sandia.gov

CP7

Algorithm-Based Checkpoint-Free Fault Tolerance for Parallel Matrix Computations on Volatile Resources

This talk presents an algorithm-based checkpoint-free fault tolerance approach in which, instead of taking checkpoints periodically, a coded global consistent state of the critical application data is maintained in memory by modifying applications to operate on encoded data. Although the applicability of this approach is not so general as the typical checkpoint/rollback-recovery approach, in parallel linear algebra computations where it usually works, because no periodical checkpoint or rollback-recovery is involved in this approach, partial process failures can often be tolerated with a surprisingly low overhead.

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

Julien Langou
The University of Tennessee
langou@cs.utk.edu

Zizhong Chen
University of Tennessee, Knoxville
zchen@cs.utk.edu

CP7

Scalable Parallel Infrastructure for Unstructured Data

A new interface and implementation for unstructured parallel data is based upon a concise, geometrically-inspired interface. Abstract covering relationships encode decompositions of computational spaces into components, using distributed graph data structures termed *Sieves*. Arbitrary data over Sieves is elegantly represented using a dual fiber-bundle-like construction. Sieve can efficiently support many hierarchical algorithms (multigrid, FMM), independently of dimension, element shape, and finite element. A reference C++ implementation and several applications

are demonstrated.

Matthew G. Knepley, Dmitry Karpeev
Argonne National Laboratory
knepley@mcs.anl.gov, karpeev@mcs.anl.gov

CP7
Directive-Driven Parallelization of 4GL Codes

Fourth Generation Languages, like RSI/IDL or Matlab, are very popular tools for image processing and scientific data analysis. Many algorithms written in these languages could benefit from distributed memory parallel processing. However, writing explicit message passing code in these languages contradicts their ease-of-use and ad-hoc coding philosophy. We here present a tool to simplify message-passing parallelization of IDL scripts using directive-driven pre-processing. Supported by NASA SBIR Grant #NNG05CA85C

Peter Messmer, Seth Veitzer
Tech-X Corporation
messmer@txcorp.com, veitzer@txcorp.com

CP7
Distributed Out-of-Core Parallel Linear Algebra on Grid5000 Heterogeneous Platform

Several experimentations already shown that data optimizations, such as data persistence and data migration anticipations, are crucial to develop efficient parallel large scale linear algebra methods on heterogeneous computing GRID. Nevertheless, we also have to use out-of-core programming on each computing resource to obtain adapted granularities and, thus, interesting efficiencies. We present results obtained on the GRID5000 heterogeneous computing platform for a few linear algebra methods, such as Block Gauss-Jordan and Givens-Bisection ones. As a conclusion, we propose a Disk-to-Disk (D2D) programming paradigm for parallel linear algebra on Heterogeneous GRID.

Serge G. Petiton
INRIA and LIFL, Universite de Lille
serge.petiton@inria.fr

Lamine Aouad
INRIA and LIFL
lamine.aouad@lifl.fr

CP8
Improving the Performance of Scientific Applications Involving a Large Number of Time Steps

Irregularities in scientific applications involving a large number of time steps, and those in the underlying computing systems may evolve unpredictably, requiring runtime selection of load balancing algorithms to improve application performance. Fortunately, the large number of time steps typical in such applications provide ample opportunities for a reinforcement learning (RL) agent to learn how to automatically select an appropriate load balancing algorithm during runtime. This contribution describes a framework integrating RL and dynamic load balancing in scientific applications to improve performance.

Ioana Banicescu
Dept. of Computer Science and Engineering
Mississippi State University
ioana@cse.msstate.edu

Ricolindo Carino
ERC Center for Computational Sciences
Mississippi State University
rlc@erc.msstate.edu

CP8
Preconditioners for Transient Problems on a Space-Time Domain

We present the results of using various preconditioners in a Newton-Krylov algorithm for solving transient problems as boundary value problems (as opposed to initial value problems). In this formulation, the entire solution over the space-time domain is solved simultaneously, with parallelism over both the time and space domains. Preconditioners designed to take advantage of the structure of Jacobian matrix used in solving the resulting system will be compared to general-purpose preconditioners.

Daniel M. Dunlavy
Sandia National Laboratories
Optimization and Uncertainty Estimation
dmdunla@sandia.gov

Andrew Salinger
Computational Sciences Department, Sandia National Labs
agsalin@sandia.gov

CP8
Parallel Finite-Difference Time-Domain Computations Aided by Modal Decomposition

High frequency electromagnetic simulation of electrically-small objects requires computationally-expensive methods such as FDTD. Parallel processing is one means to reduce simulation time. A shortcoming of the parallel approach is the communication overhead that is responsible for reduced speedup efficiency but inevitable at every time step. We approach FDTD from the signal processing point of view and consider concurrency with transient and steady-state excitation scenarios. We rely on modal decomposition and convolution as the techniques to aid in our research. We develop communication-intensive as well as communication-free implementations for each.

Dmitry Gorodetsky, Philip Wilsey
University of Cincinnati
goroded@email.uc.edu, philip.wilsey@ieee.org

CP8
Application of Dimensionality Reduction Techniques to Time-Parallelization

We recently developed a new data-driven parallelization strategy that decomposes the time domain, instead of the conventional space domain. It depends on dimensionality reduction for identifying important modes of system behavior. We will report on the application of using these techniques in important molecular dynamics simulations in materials and biology.

Lei Ji
Department of Computer Science
Florida State University
ji@cs.fsu.edu

Yanan Yu, Ashok Srinivasan

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

CP9

MCell-K: A Highly Scalable Cell Microphysiology Simulator on Blue Gene L

MCell-K is a scalable variant of MCell, a highly successful Monte Carlo simulator of cellular microphysiology, which uses a random walk to model reaction/diffusion systems. We present new scalability results on hundreds of processors of Blue Gene and IBM Power4 systems, demonstrating exceptional scalability on a realistic simulation of neurotransmission. We discuss performance issues, including run time support and algorithmic changes needed to achieve high scalability.

Gregory T. Balls
CSE Department
Univ of California, San Diego
gballs@cs.ucsd.edu

Terrence Sejnowski
Salk Inst. for Biological Studies
CSE Dept., Univ of California, San Diego
terry@salk.edu

Scott B. Baden
Dept of Computer Science and Engineering
University of California, San Diego
baden@cs.ucsd.edu

Thomas M. Bartol
The Salk Institute
bartol@salk.edu

CP9

Runtime System for Asynchronous Cellular Microphysiology Simulation

Simulations of dynamic environments are hard to parallelize efficiently, because the need to preserve causality exacerbates the impact of load imbalance. We discuss Tarragon, a runtime system that offers asynchronous data-driven execution, and supports relaxed synchronization and dynamic load balancing. Tarragon's API hides lower level system dependent details, taking advantage of any architectural features that can support its data driven execution model. We discuss recent results with a cell microphysiology simulator.

Pietro Cicotti
University of California, San Diego
Department. Computer Science and Engineering
pcicotti@cs.ucsd.edu

Scott B. Baden
University of California, San Diego
Dept of Computer Science and Engineering
baden@cs.ucsd.edu

CP9

Large Scale Parallel Implementation of Replica Exchange on Blue Gene/L

Replica exchange is an enhanced sampling algorithm used in thermodynamic simulations of proteins and other systems. We will present an efficient and scalable implementation of replica exchange that is tuned to exploit the physical network topology and hardware features of BlueGene/L. Performance measurements of the Blue Matter Molecular Dynamics application using the replica exchange method on thousands of nodes of Blue Gene/L will be presented.

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

Jed M. Pitera
Science and Technology, IBM Almaden Research Center, San Jos
pitera@us.ibm.com

Aleksandr Rayshubski, Robert S. Germain, Ruhong Zhou
IBM Thomas J. Watson Research Center
arayshu@us.ibm.com, rgermain@us.ibm.com, ruhongz@us.ibm.com

CP9

Strong Scaling of Biomolecular Simulation on Blue Gene/L

To support simulations of proteins and membranes over long time scales, we have developed a molecular dynamics code, Blue Matter, that also serves as a test bed for exploitation of Blue Gene/L hardware features and for explorations of parallel decompositions and algorithms suitable for massively parallel machines. We will describe the parallel decomposition and distributed 3D-FFT that have enabled us to achieve continued speedup through 16,384 Blue Gene/L nodes on a 43,222 atom system.

Blake Fitch, Aleksandr Rayshubskiy, Maria Eleftheriou, Michael Pitman
IBM Thomas J. Watson Research Center
bgf@us.ibm.com, arayshu@us.ibm.com, mariae@us.ibm.com, pitman@us.ibm.com

Chris Ward
IBM Software Group
tjcw@uk.ibm.com

Mark Giampapa, Frank Suits, Robert S. Germain
IBM Thomas J. Watson Research Center
giampapa@us.ibm.com, suits@us.ibm.com, rgermain@us.ibm.com

CP10

High Order Solutions to Self-Adjoint Elliptic Equations by Reduction to Constant Coefficients

We develop a solver for variable coefficient, self-adjoint, elliptic equations. The solution of an auxiliary constant coefficient equation induces a reduction to constant coefficient equations plus small inhomogeneous deviations. A highly accurate, fast, Fourier-spectral algorithm can solves such equations. A small number of steps achieves high accuracy. The procedure becomes more efficient as the size of the domains of solution decreases. A highly parallelizable, hierarchical, decomposition and reconstruction procedure gives an accurate global solution.

Moshe Israeli
Technion Institute of Technology, Israel
Israeli@cs.technion.ac.il

CP10
Parallel Domain Decomposition Methods for Some Stochastic

We present a parallel multilevel domain decomposition preconditioned recycling Krylov subspace method for the numerical solution of some elliptic partial differential equations with stochastic uncertainties in the operator. Karhunen-Loeve expansion and finite element method with double orthogonal polynomial basis are used to transform the stochastic problem into a sequence of deterministic equations. A PETSc based parallel software is developed for a recycling domain decomposition preconditioned iterative method and its parallel scalability will be reported in this talk.

Chao Jin, Congming Li
Dept. of Applied Mathematics
University of Colorado, Boulder
chao.jin@colorado.edu, cli@colorado.edu

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

CP10
New Scalability and Accuracy Results for a Latency-Tolerant Elliptic Solver

We present new scalability and accuracy results for a second-order accurate elliptic free space solver. Employing a method of local corrections, we reduce communication costs by representing far-field effects at a coarser resolution. Improvements in parallel scalability for the global coarse calculation and the introduction of adaptivity allow us to scale efficiently to thousands of processors. Numerical overheads incurred are independent of the number of processors for a wide range of problem sizes.

Peter Mccorquodale
Lawrence Berkeley National Laboratory
PWMcCorquodale@lbl.gov

Scott B. Baden
Dept of Computer Science and Engineering
University of California, San Diego
baden@cs.ucsd.edu

Phillip Colella
Lawrence Berkeley National Laboratory
PColella@lbl.gov

Gregory T. Balls
CSE Department
Univ of California, San Diego
gballs@cs.ucsd.edu

CP10
The Addition of Hp-Adaptivity to a Parallel Adaptive Finite Element Program

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. The objective is to determine a distribution of h and p that minimizes the error using the least amount of work in some measure. In this talk we discuss the addition of hp-adaptivity to PHAML, a parallel adaptive multigrid finite element program.

William F. Mitchell
National Institute of Standards and Technology
mitchell@cam.nist.gov

CP11
Multiresolution Analysis and Implicit Deferred Correction on Parallel Computers

We describe the application of the deferred correction procedure to a multiwavelet based multiresolution analysis framework for solving integro-differential equations for quantum electronics structures and time dependent density functional theory computations. The deferred correction method is implemented in the multiresolution solver package Multiresolution Adaptive Evaluation for Scientific Simulations (MADNESS). We also report on preliminary parallel performance on the Cray-XT3, the Cray XD-1 and the IBM Blue Gene/L machines.

George Fann
Oak Ridge National Laboratory
fanngi@ornl.gov

CP11
Scalability Study of Accelerator Simulations on Large-Scale Clusters

In this paper, we present a detailed scalability analysis of advanced accelerator simulations on various TOP500 systems at NCSA. In particular, we study the performance of Synergia, a state-of-the-art accelerator modeling code funded by the DOE SciDAC program. The goal is to understand the performance and scaling behavior of accelerator simulations on different CPU architectures and message layers and further to investigate the performance bottlenecks associated with the existing accelerator simulations.

Zhiling Lan, Jungmin Lee
Illinois Institute of Technology
lan@iit.edu, leejung6@iit.edu

Panagiotis Spentzouris
Fermi National Accelerator
spentz@fnal.gov

James Amundson
Fermi National Accelerator Laboratory
amundson@fnal.gov

CP11
A Parallel Software Framework for Solving Radiative Heat Transfer Problems by the Photon Monte Carlo Method

Radiative heat transfer is extremely difficult to model because of its highly nonlocal and nonlinear nature. In this talk we present a parallel software framework that implements the photon Monte Carlo method of ray tracing to simulate radiative effects. We demonstrate the scalability of this framework and address load balancing issues inherent with this approach. The framework allows for the incorporation of user-specified radiative properties which enables its use within a variety of combustion applications.

Paul E. Plassmann

Virginia Tech
plassmann@vt.edu

Ivana Veljkovic
Penn State University
Department of Computer Science and Engineering
icotaki@psu.edu

CP12

Decomposing Solution Sets of Polynomial Systems in Parallel

Consider the numerical irreducible decomposition of a positive dimensional solution set of a polynomial system into irreducible factors. With path tracking techniques we can track loops around singularities and may connect points on the same irreducible components. The computation of a linear trace for each factor certifies the decomposition. Using the concepts of monodromy and linear trace, we present a new parallel monodromy breakup algorithm. Our implementation is written in C using MPI and is linked to the library of PHCpack, a software package for polynomial homotopy continuation. It has exhibited good speedups and the fastest decomposition times on several classical benchmarks, especially on solution sets of large degree.

Anton Leykin
University of Illinois at Chicago
leykin@math.uic.edu

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

CP12

Parallel Space-Filling Curve Generation for Dynamic Load Balancing

Adaptive mesh refined codes require a dynamic load balancer. Load balancers based on space-filling curves can generate reasonable partitions but it has not been clear how to generate the space-filling curves in parallel. We have developed a parallel merge sort method to generate these curves that scales well up to thousands of processors. Complexity analysis and experimental results are given.

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

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

CP12

An Improved Bi-Level Algorithm for Load-Balancing Structured Refinement Hierarchies

Simulations on evolving structured hierarchical meshes typically suffer from poor parallel efficiency due to load imbalances. We analyse this problem with a set of test configurations derived from AMROC. In particular, we discuss a singularly converging Richtmyer-Meshkov instability for which the native load balancer partially exhibits imbalances greater than 100 percent. This presentation proposes an improved bi-level partioning algorithm that automatically switches between two approaches and reduces the imbalances to less than 30 percent.

Johan Steensland
Sandia National Laboratories
jsteens@ca.sandia.gov

Ralf Deiterding
Center for Advanced Computing Research
California Institute of Technology
ralf@cacr.caltech.edu

Jaideep Ray
Sandia National Laboratories
jairay@ca.sandia.gov

CP12

A Scalable Approach for Adaptive Embedded Boundary Cartesian Grid Generation

This talk describes an adaptive patch-based AMR scheme for complex geometry embedded boundaries where both the flow solution and the grid generation are done completely in parallel. The grid generation steps are distributed across the processors in a load-balanced way. The work integrates the SAMR package SAMRAI and the Cartesian adaptive grid generator in Cart3D. We demonstrate scaling performance using a model of aerosol dispersion in urban environments. We also discuss methods for determining grid generation costs and how this was used for optimizing load balance.

Andrew M. Wissink
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
awissink@llnl.gov

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

CP13

Distributed Memory Parallel Implementation of the Block Shift and Invert Lanczos Algorithm

The Block Shift and Invert Lanczos Algorithm is the workhorse of commercial finite element analysis packages for performing vibration and buckling analyses. Eigenanalysis problems are reaching 50 million rows in these applications. This paper will provide an overview of what we believe to be the first distributed memory parallel implementation of this algorithm for use in a commercial finite element package LS-DYNA.

Roger Grimes
Livermore Software Technology Corporation
grimes@lstc.com

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

CP13

Block Locally Optimal Preconditioned Eigenvalue

Xolvers

Block Locally Optimal Preconditioned Eigenvalue Xolvers (BLOPEX) is a package, written in C, that at present includes only one eigenxolver, Locally Optimal Block Preconditioned Conjugate Gradient Method (LOBPCG). BLOPEX supports parallel computations through an abstract layer. BLOPEX is incorporated in the HYPRE package from LLNL and is availabe as an external block to the PETSc package from ANL as well as a stand-alone serial library. Main LOBPCG features: a matrix-free iterative method for computing several extreme eigenpairs of symmetric positive generalized eigenproblems; a user-defined preconditioner; robustness with respect to random initial approximations, variable preconditioners, and ill-conditioning of the stiffness matrix; apparently optimal convergence speed. Numerical comparisons suggest that LOBPCG may be a genuine block analog for eigenproblems of the standard preconditioned conjugate gradient method for symmetric linear systems.

Merico Argentati
University of Colorado at Denver
rargenta@math.cudenver.edu

Andrew Knyazev, Ilya Lashuk
CU-Denver
andrew.knyazev@cudenver.edu,
ilashuk@math.cudenver.edu

Evgueni Ovtchinnikov
U. Westminster, London
eovtchin@math.cudenver.edu

CP13
An Efficient Cluster-Based Parallel FFT Algorithm

The paper presents a parallel radix-2 Cooley-Tukey FFT algorithm on a cluster. A new data distribution method is invented and parallelized to distribute data points among processors such that no message passing is needed by the bit-reversal step in which each processor only performs the bit-reversal on its local vector with a fast linear-time recursive bit-reversal algorithm. The experimental result shows that the parallel FFT algorithm achieves significant speedup over the serial Stockham FFT algorithm.

Jingyu Zhang
University of Illinois at Springfield
Department of Computer Science
jzhan2@uis.edu

Michael Stults
Legal Files Software Inc.
Springfield, IL 62704
mikes@legalfiles.com

CP14
A Parallel Priority Queueing System with Balking and Reneging

We pose and analyze a novel model of a priority queueing system with balking and reneging that is designed to minimize the idleness of the system. An algorithmic solution for the general two-station case and closed-form solutions for special cases of this model are given. We find distribution of the queue length; probabilities of the system idleness, lost to the system, and stations being busy; and the first two moments. Numerical examples are also given.

Aliakbar M. Haghighi, Dimitar Mishev
Prairie View A&M University
amhaghighi@pvamu.edu, dimichev@pvamu.edu

CP14
Feedback-Based Guided Self-Scheduling

Parallel loops account for the greatest amount of parallelism in numerical and scientific applications. Scheduling parallel loops, i.e., the way iterations are mapped on to different processors, plays a critical role in the efficient execution of such applications, on multiprocessor systems. In applications where problem dimension (and hence execution time) is dependent on run-time data, loop iterations also tend to be of variable length – this variability affects both sequential and parallel loops and in particular nested loops and it is quite prevalent in sparse matrix solvers. In this work, we propose a feedback-directed approach for dynamic scheduling of irregular parallel loops on heterogeneous multiprocessor systems. Results show that our technique performs (upto) 18% better than the existing dynamic scheduling schemes.

Arun Kejariwal, Alexandru Nicolau
University of California at Irvine
arun_kejariwal@computer.org, nicolau@ics.uci.edu

Constantine Polychronopoulos
University of Illinois at Urbana-Champaign
cdp@csrd.uiuc.edu

CP14
A Data Driven Model for Tolerating Communication Delays

Hiding communication latency is a difficult challenge due to the high software overheads incurred. We discuss Thyme, a data-driven programming model and run time library that manages communication pipelining and scheduling through task graph, actor-like execution. Thyme frees the user from many low-level system dependent details involved in masking communication, easing the reformulation of applications to tolerate delays. We demonstrate results with some practical applications and discuss compiler support using the ROSE source-to-source translator.

Scott B. Baden
University of California, San Diego
Dept of Computer Science and Engineering
baden@cs.ucsd.edu

Jacob Sorensen
University of California, San Diego
Computer Science and Engineering Department
jcsorens@cs.ucsd.edu

MS1
Cache-Oblivious Algorithms

Cache-oblivious algorithms and data structures are platform-independent. They run efficiently on a hierarchical memory, even though they avoid any memory-specific parameterization, such as cache-block sizes or access times. We give an overview of the work on designing cache-oblivious algorithms and data structures. We then give theoretical and experimental results on searching and al-

gorithmic problems. Finally, we outline future directions in the area.

Michael Bender
Stony Brook University
bender@cs.sunysb.edu

MS1
Adaptive Programming with Hierarchical Multiprocessor Tasks

The talk considers the programming with hierarchically structured multiprocessor tasks (M-tasks). An M-task can be mapped onto a group of processors and can be executed concurrently to other independent tasks. Internally, an M-task can consist of a hierarchical composition of smaller M-tasks or can incorporate different kinds of parallelism. Depending on the structure of the application, significant performance improvements can be obtained. As examples we consider specific ODE solvers and hierarchical algorithms for matrix multiplication.

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

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

MS1
Adaptive Algorithms: Theory and Applications

The minisymposium is firstly introduced by a taxonomy of adaptive methods in scientific computing. Then, the talk focuses on adaptive parallel computing on non-uniform and shared resources. Minimizing the completion time depends on two antagonists complexity measures: number of operations and critical time. We present a dynamic cascading technique based on the coupling of two algorithms, sequential and parallel fine-grain. It is illustrated by practical applications in combinatorial optimization and linear algebra.

Van Dat Cung
GILCO Lab., INPG Grenoble, France
cung@gilco.inpg.fr

Guillaume Huard
ID Laboratory (Grenoble, France)
supported by UJF / CNRS / INRIA / INPG
Guillaume.Huard@imag.fr

Jean-Louis Roch
MOAIS Project (CNRS-INRIA)
ID-IMAG Lab., Grenoble, France
Jean-Louis.Roch@imag.fr

Denis Trystram
MOAIS Project
ID-IMAG lab. INPG Grenoble
denis.trystram@imag.fr

Jean-Guillaume Dumas
MNC-IMAG, UJF Grenoble
jean-guillaume.dumas@imag.fr

Thierry Gautier
MOAIS Project, ID-IMAG, France
thierry.gautier@imag.fr

Bruno Raffin
MOAIS Project, ID-IMAG Lab
INRIA Rhone-Alpes, France
bruno.raffin@imag.fr

Christophe Rapine
Lab GILCO, INP Grenoble, France
christophe.rapine@imag.fr

MS1
Hybrids in Exact Linear Algebra

For a remarkable range of types and sizes of integer and rational matrix problems, there is now sophisticated high performance software. Among those problems are linear system solution, nullspace basis, rank, determinant, and characteristic polynomial. To reach high performance, engineered hybrids of multiple algorithms (close-to-the-hardware BLAS, iterative and exact methods) play an increased role. We will discuss hybrids in more detail, illustrating several ways to gain and revealing some dilemmas too.

Dumas Jean-Guillaume
LMC-IMAG Grenoble University
jean-guillaume.dumas@imag.fr

David Saunders
Dept of Computer and Information Sciences
University of Delaware, Newark
saunders@cis.udel.edu

MS2
Identification of Severe Multiple Contingencies in Electric Power Networks

The August 14, 2003 blackout in the Northeast has emphasized on the need to achieve improved operational efficiency and reliability standards for electric power systems. Such attempts are however hindered by the lack of advanced computational tools to better perform security evaluations. Present practices and policies are limited by the inability to evaluate multiple simultaneous outages, and as observed by 2003 Northeast blackout, the impact of triple contingencies can be enormous. This talk will focus on a two-stage screening and analysis process for identifying multiple contingencies that may result in very severe disturbances and blackouts. In a screening stage, an optimization problem is formulated to find the minimum change in the network to move the power flow feasibility boundary to the present operating point, which will cause the system to separate with a user-specified power imbalance. The lines identified by the optimization program are used in a subsequent analysis stage to find combinations that may lead to a blackout. Application of this approach to the IEEE 30-bus system and 179-bus system that yields encouraging results will be presented.

Bernard Lesieutre
LBNL
BCLesieutre@lbl.gov

Vanessa Lopez
Computational Research Division
Lawrence Berkeley National Laboratory

vllopez@lbl.gov

Vaibhav Donde
Environmental Energy Technologies Division
Lawrence Berkeley National Laboratory
vdonde@lbl.gov

Chao Yang
Lawrence Berkeley National Lab
CYang@lbl.gov

Juan C. Meza
Lawrence Berkeley National Laboratory
JCMeza@lbl.gov

Ali Pinar
Lawrence Berkeley Lab
apinar@lbl.gov

MS2

Computational Challenges in the Power Industry

Not available at time of publication.

Lawrence Jones
Transmission Optimization & Partnering (TOP) Program
AREVA T&D
lawrence.jones@areva-td.com

MS2

A Social Intelligence Approach for Identifying the Most Effective Locations for Terrorist Attack in the National Power Network

This presentation focuses on research related to the identification of critical nodes within the national bulk power system. These are nodes that, if disrupted through natural events or terrorist action would cause the most widespread, immediate damage. A number of traditional methods were investigated to identify these nodes: integer-based optimization, genetic algorithms, network theory (small world), polyhedral dynamics, and cultural-based methods. In addition, a number of modeling approaches were investigated including linear (dc) and non-linear (ac) power flow models. The final approach is based on the use of an ac power flow model coupled with an adaptive cultural model and collective intelligence as a means of characterizing the reliability of bulk power networks. The method involves a collective of virtual power engineers acting as a cell of potential terrorists attempting to attack the bulk power system in the most optimal manner. The methodology is technology independent: it can be applied not only for reliability analysis of bulk power systems, but also other energy systems or transportation systems or a combination of these.

David Robinson
Sandia National Laboratories
drobin@sandia.gov

MS2

Using Stochastic Partitioning Algorithms for Contingency Identification in Electric Power Systems

Several failure analysis and mitigation tasks for electric power networks require identification of network partitions with either large or zero generation-load imbalance. One unique feature of these tasks is the need for finding not only the extremal partition (i.e., the one with the largest or smallest absolute imbalance) but also other partitions that are near the extreme. Here, we advance the perspective that stochastic partitioning algorithms - including a variant of Karger's min-cut algorithm and an influence model-based algorithm - are well-suited to search for all partitions with large or small generation-load imbalance.

Sandip Roy
Washington State University
School of Electrical Engineering and Computer Science
sroy@eecs.wsu.edu

MS3

Code Coupling in the Center for Simulation of Waves in Magnetohydrodynamics

CSWIM is a DoE project for integrating different physics codes together for better simulation of magnetically confined plasmas. Physics, mathematics, high performance computing, and software systems must be addressed in a tightly coupled manner. The codes model the same region of space but with different time-scales, discretizations, symmetries, and dimensionality reduction for an inherently 7-dimensional problem. This talk is an overview of CSWIM problems shared by many related fusion energy simulation integration efforts.

Randall B. Bramley
Indiana University
bramley@cs.indiana.edu

MS3

Not Available at Time of Publication

Abstract not available at time of publication.

Phil Colella
Lawrence Berkeley National Laboratory
pcolella@lbl.gov

MS3

An Overview of The Distributed Coupling Toolkit (DCT): Functionality and Applications

The Distributed Coupling Toolkit (DCT) offers a functional interface for flexible and dynamic coupling. Among key features of DCT are its mechanisms for handling coupling in a distributed manner thus exploiting scalability in large parallel systems. In DCT there is no coupler-process(or) or centralized control. Instead coupling is formulated in a more natural way to multiphysics models. Here, we present basic DCT functionality, scientific applications deploying DCT, and comparisons with centralized coupling approaches.

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

MS3

The Model Coupling Toolkit and Parallel Coupling Infrastructure

Parallel multiphysics and multiscale models have inherent mutual parallel data exchanges, leading to a *parallel coupling problem* (PCP). The architectural and data process-

ing parts of the PCP lead to a set of requirements we used to create a performance-portable package–the Model Coupling Toolkit–to ease construction of bespoke solutions to the PCP. We will present a roadmap for development of a more general parallel coupling infrastructure (PCI) that will address more general parallel coupling problems.

Jay W. Larson
Math and Computer Science Division
Argonne National Laboratory
larson@mcs.anl.gov

MS4

A Quantitative Characterization of the Structure of Toolkits Designed Using the Common Component Architecture Paradigm

Component-based software designs, when used for developing codes for scientific simulation, are supposed to enhance productivity by subsetting software complexity into modular components and promoting component resuse. We investigate this supposition quantitatively by devising various metrics and evaluating them based on statistics drawn from two toolkits (for combustion and computational chemistry simulations), both of which have adopted the Common Component Architecture paradigm. We also investigate how such modularization enables novel techniques such as adaptive codes.

Benjamin Allan, Joseph Kenny
Sandia National Laboratories
baallan@ca.sandia.gov, jpkenny@sandia.gov

MS4

The Cactus Framework: Design, Applications and Future Directions

The Cactus Code is an open source, modular framework for collaborative high-performance computing. In addition to including different computational and community toolkits, Cactus provides a range of abstract interfaces to application developers through which a variety of third party libraries and tools, including HDF5, PAPI, SAMRAI, PETSc, and Globus, can be easily leveraged. In this talk, we review the design of Cactus, describe how it is used by different application communities, and discuss ongoing development plans.

Gabrielle Allen
Center for Computation & Technology
Louisiana State University, Baton Rouge, LA 70803
gallen@cct.lsu.edu

MS4

An Introduction to Framework-Based Approaches for Parallel Scientific Computing

Research in parallel scientific computing continues to increase in complexity, as multi-physics, multi-scale, and multi-institutional projects are becoming widespread. Component-based frameworks aim to address many of the software and social challenges that arise in incorporating the combined expertise of diverse groups in such projects. This presentation will introduce the main concepts of framework-based approaches and will highlight the use of parallel numerical components in several scientific applications.

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

MS4

Enabling Cross-Organization Coupling of Climate Models with the Earth System Modeling Framework

Cross-organization coupling of climate models is one efficient way of discovering new science, increasing forecasting skills, and validating models. NASA's Earth-Sun System Technology Office/Computational Technologies Project has funded development of the Earth System Modeling Framework (ESMF) to support cross-organization couplings within a standardized software environment. This talk will discuss the ESMF, which so far has been used to couple major climate models from NASA, NOAA, NSF/NCAR, DOE/LANL, MIT, and UCLA.

Shujia Zhou
Software Integration and Visualization Office
NASA Goddard Space Flight Center
szhou@pop900.gsfc.nasa.gov

MS5

Software-Directed Disk Power Minimization for Scientific Workloads

Energy-consumption has become a very important issue for high-performance machines that target scientific workloads. In particular, disk power consumption of data-intensive scientific applications stands as one of the major problems for application designers and system maintainers. While most of the prior approaches to disk power minimization focused exclusively on hardware based techniques, our work explores the possibility of employing an optimizing compiler to control disk power. In this work, we study both data layout organization and code restructuring techniques for the single processor and multi-processor cases. We also compare the performance of our approach to pure hardware-based schemes.

Mahmut Kandemir
Dept of Computer Science Engr.
The Pennsylvania State Univ.
kandemir@cse.psu.edu

Seung Son, Guilin Chen
Computer Science & Engg.
The Pennsylvania State University
sson@cse.psu.edu, guilchen@cse.psu.edu

MS5

Maximizing Multinode Performance on Massively Parallel Computers

This talk concerns recent progress in exploring the frontiers of supercomputing including architecture, software and tools for massively parallel scientific computing with low cost, space and power consumption. We will discuss how a large number of low power processors, multiple integrated interconnection networks and a novel system software architecture enables applications from a broad variety of scientific disciplines to effectively scale to tens of throusands of processors.

Jose Moreira
IBM Research

jmoreira@us.ibm.com

MS5
High-Throughput Low Power Scientific Computing

In this talk we explore how recent advances in high-throughput architectures that concentrate on maximizing the amount of thread-level parallelism at the expense of instruction-level parallelism can achieve very high-performance on scientific applications. This style of architecture, which was originally targeted at commercial server applications, can achieve high-performance while maximizing the performance/watt on a wide range of scientific applications. We expect these high-throughput architectures to dramatically reduce the power requirements of scientific computing.

Kunle Olukotun
Electrical Engg. and Computer Science
Stanford University
kunle@csl-mail.stanford.edu

James Laudon
Sun Microsystems
ames.laudon@sun.com

MS5
Energy-Aware Architectural Optimizations for Fast Sparse Scientific Computations

We utilize current trends in power-aware high-performance architecture design to explore the impact of memory subsystem optimizations on sparse scientific codes. We propose a variety of features such as memory and cache prefetching and load miss prediction to enable high performance, sparse scientific applications at reduced power/energy levels. We will discuss the use of cycle-accurate simulations with tools such as SimpleScalar and Wattch to characterize performance and power profiles. Our results for an optimized sparse-matrix-vector multiplication kernel and two sparse kernels from the NAS benchmark indicate that in certain cases, our optimizations can reduce application time by 80%, and the energy-time product by 90%, at system power levels less than or equal to those in the original configuration.

Mary Jane Irwin
Compter Science & Engineerin
Pennsylvania State University
mji@cse.psu.edu

Ingyu Lee
Computer Science & Engineering
The Pennsylvania State University
inlee@cse.psu.edu

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

Padma Raghavan
Computer Science & Engineering
The Pennsylvania State University
raghavan@cse.psu.edu

MS6
Performance Evaluation of the Columbia Constellation

Since its debut in spring of 2002, the Earth Simulator (ES) has attracted worldwide attention – occupying the number one spot on the Top500 for a record two and a half years. The ES uses a dramatically different architectural approach than most conventional cache-based supercomputers, utilizing powerful vector processors connected via a fast single-stage crossbar. In this talk, we examine ES performance on several ultra-scale scientific applications and demonstrate the tremendous potential of this architectural paradigm.

Rupak Biswas
NASA Ames Research Center
rbiswas@nas.nasa.gov

MS6
An Evaluation of MareNostrum Performance

This talk presents an evaluation of the MareNostrum PowerPC cluster from the individual components to the final performance delivered for real applications. Beyond just reporting the achieved performance metrics, the talk aims at developing models of the machine and applications. The results of microbenchmarks and parameter fits based on simulations with Dimemas are used to generate the models. Such characterizations prove very useful to properly predict and understand the observed behaviour and extrapolate expectations.

Rosa Badia, Gerard Rodríguez, Jesus Labarta
Barcelona Supercomputing Center
rosab@ac.upc.es, gerard@ac.upc.es, jesus@cepba.upc.edu

MS6
Leading Computational Methods on the Earth Simulator

Since its debut in spring of 2002, the Earth Simulator (ES) has attracted worldwide attention – occupying the number one spot on the Top500 for a record two and a half years. The ES uses a dramatically different architectural approach than most conventional cache-based supercomputers, utilizing powerful vector processors connected via a fast single-stage crossbar. In this talk, we examine ES performance on several ultra-scale scientific applications and demonstrate the tremendous potential of this architectural paradigm.

Leonid Oliker, Jonathan Carter
Lawrence Berkeley National Laboratory
loliker@lbl.gov, jtcarter@lbl.gov

MS6
Scalable Application Performance on BG/L, a Truly Massively Parallel Platform

Unprecedented levels of performance have been achieved on BlueGene/L (BG/L) designed by IBM Research in partnership with the Advanced Simulation and Computing (ASC) Program, part of the U.S. Department of Energy's National Nuclear Security Administration. The order of magnitude improvement in performance was enabled by an order of magnitude increase in processors. This unprecedented scale has presented unique challenges for applications, from

diverse domains such as molecular dynamics and breadth-first search. This talk will present results for several applications, many of which have achieved record levels of performance, as well as the BG/L system architecture.

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

MS7

Nonlinear Adaptive Smoothed Aggregation Multigrid

The increasing demands of large-scale complex solid mechanics simulations are placing greater emphasis on the challenges associated with the efficient solution of the set of nonlinear equations. Here, the solution of large systems of equations with material, geometric and contact nonlinearities in a parallel framework is adressed. Instead of applying widely used Newton- or Newton-Krylov type methods that involve the derivation of a stiffness matrix and a sequence of linear solves, the presented work details the implementation of a variational nonlinear algebraic multigrid algorithm based on the full approximation scheme (FAS) applied to solve this set of nonlinear equations [2]. The multigrid hierachy is constructed using an adaptive smoothed aggregation multigrid approach [1],[5] applied to a fine grid Jacobian that will also be used to construct nonlinear smoothers on all grids. As the variational evaluation of the nonlinear residual function plays a crucial role for the cost of the overall method, reducing the number of residual evaluations by encountering the extra effort of a (linear) adaptive smoothed aggregation procedure can reduce overall solution times significantly. The algorithm is implemented within Sandia National Laboratories' freely available parallel 'Trilinos' linear algebra framework[3] and makes use of its smoothed aggregation multigrid library 'ML' [4] and its nonlinear solver library 'NOX'. The outline of the algorithm and implementation are given together with examples demonstrating the advantages of this new approach. Several variants of the algorithm will be discussed and compared. References [1] Brezina, M. et al., "Adaptive Smoothed Aggregation (αSA) Multigrid", SIAM Review, 47, 317-346, 2005. [2] Gee, M.W., Tuminaro, R.S., "Nonlinear algebraic multigrid for constrained solid mechanics problems using Trilinos", technical report, Sandia National Laboratories. 2005. [3] Heroux, M.A., Willenbring, J.M., "Trilinos Users Guide", Sandia Report SAND2003-2952, 2003. [4] Sala, M., Hu, J.J., Tuminaro, R.S., "ML 4.0 Smoothed Aggregation User's Guide", Sandia Report SAND2004-2195, 2004. [5] Vanek, P., Mandel, J., Brezina, M., "Algebraic Multigrid by Smoothed Aggregation for Second and Fourth Order Elliptic Problems", Computing, v. 56, p. 179-196, 1996.

Ray S. Tuminaro, Michael W. Gee, Kendall H. Pierson
Sandia National Laboratories
rstumin@sandia.gov, mwgee@sandia.gov, khpiers@sandia.gov

MS7

Smoothed Aggregation Based on Energy Optimization and Aggressive Coarsening

Algebraic multigrid (AMG) based on smoothed aggregation (SA) was originally developed around the central idea of minimizing the energy of interpolation operators that also interpolate certain null space vectors exactly [1]. Mandel, Brezina, and Vanek later proposed a more general method that seeks to minimize the energy of the interpolation operators by explicitly solving a constrained optimization problem, subject to a specified nonzero pattern and certain null space interpolation constraints [1]. In this talk we will begin with an overview of the SA approach to AMG. We will then explore extensions of the energy optimization method to SA with aggressive coarsening (i.e., larger aggregates) and explicit detection of anisotropies. We will end with parallel numerical results demonstrating the effectiveness of the AMG preconditioning. [1] P. Vanek, J. Mandel, and M. Brezina, "Algebraic multigrid based on smoothed aggregation for second and fourth order problems." *Computing*, 1996 (56), p. 179-196. [2] J. Mandel, M. Brezina, and P. Vanek, "Energy Optimization of Algebraic Multigrid Bases". *Computing*, 1999 (62), p. 205-228.

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

Alejandro Cantarero
UCLA Mathematics Department
Los Angeles, CA
cantarer@gmail.com

MS7

A Parallel Implementation of a Fast Adaptive Composite Mesh Solver (FAC) in HYPRE

The FAC Solver is an extension of the multigrid method for solving AMR problems. Parallelization of this solver involves two types of communication, between the AMR hierarchy and within one AMR level. Since our FAC algorithm involves algebraically constructing the coarse patch operators, both of these types of communication must be considered in the setup and solve phases. This talk will discuss the parallel implementation of our FAC solver in the Semi-Structured Interface of HYPRE.

Barry Lee
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
lee123@llnl.gov

MS7

A new framework to define smoothed aggregation preconditioners for the parallel solution of non-symmetric linear systems

A new variant of the smoothed aggregation multigrid method suitable for solving non-symmetric linear systems is presented. We extend the classical smoothed aggregation algorithm for symmetric linear systems, which first defines a tentative prolongator operator, applies an operator of type $I - \omega A$ to it to produce the prolongator operator, then takes the transpose of this final operator to construct the restriction operator. We introduce a new framework, that can be used for both symmetric and non-symmetric problems, to compute the damping parameter ω. The framework also allows the construction of Petrov-Galerkin coarse level corrections. Numerical results will be presented, to show the validity of the proposed approach on highly non-symmetric systems, obtained with finite differences and finite elements methods.

Marzio Sala

ETHZ Computational Laboratory
marzio@inf.ethz.ch

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

MS8

Object-Oriented Generic Programming for Abstract Numerical Algorithms via Thyra

Object-oriented generic programming techniques allow software developers to express algorithms using abstract interfaces where details of how work is accomplished are not prescribed *a priori*. Thyra defines the interfaces needed for a large class of abstract numerical algorithms on parallel computers. By using Thyra to access basic linear algebra objects and solvers, users can access broad sets of interchangeable solver components. In this talk we present Thyra and how it can be used for high-performance parallel numerical software.

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

MS8

The Trilinos Package Architecture and Infrastructure

Trilinos is a collection of independent, interoperable and reusable software libraries for constructing and solving systems of linear, nonlinear, eigen, and time-dependent equations that arise in complex, high-fidelity multi-physics applications on serial and parallel computers. In this presentation, we discuss the Trilinos package architecture and infrastructure, showing how Trilinos impacts software development efficiency and quality. Finally, we discuss how package interoperability is accomplished and its impact on solver capabilities.

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

MS8

PyTrilinos: A Parallel Python Interface to Trilinos

PyTrilinos is a collection of Python modules to use the solver capabilities in Trilinos. PyTrilinos is attractive to applications written in Python, and to others who want a rapid prototyping environment. Its performance is remarkably good and it executes in parallel via MPI. In this presentation we give an overview of PyTrilinos and its design and use. We also compare it to using Trilinos via native C++, illustrating ease of programming and performance results.

William F. Spotz
Sandia National Laboratories
wfspotz@sandia.gov

MS8

Anatomy of a Package: Anasazi

Anasazi is an extensible and interoperable framework for large-scale eigenvalue algorithms, designed to provide a generic interface to a collection of algorithms for solving large-scale eigenvalue problems. Anasazi eigensolvers enlist generic programming techniques to give the user the ability to leveraging any existing software investment. This talk will discuss the design of Anasazi, from top to bottom, illustrating its capabilities for interoperability within the Trilinos framework and beyond.

Heidi Thornquist
Sandia National Laboratories
hkthorn@sandia.gov

MS9

A Hybrid Software Framework for Parallel Tsunami Simulations

Efficient tsunami simulations require different mathematical models, numerical algorithms, and local meshes in different regions of an ocean. We therefore present a hybrid framework based on subdomains, where a subdomain solver can vary from a Fortran FDM code working on a uniform local mesh to a C++ FEM code on an unstructured mesh. Serial wave codes are thus incorporated into a parallel hybrid tsunami simulator. Preliminary simulation results of the Indian Ocean Tsunami will also be reported.

Xing Cai
Simula Research Laboratory
1325 Lysaker, Norway
xingca@simula.no

MS9

Component-Based Quantum Chemistry: Examining Coarse and Fine-Grained Scientific Component Applications

Component-based molecular structure optimization, integral evaluation, and numerical Hessian evaluation have been implemented by the quantum chemistry working group associated with the Common Component Architecture. Serving as examples of both coarse and fine-grained component applications, the development of these capabilities will be described in an effort to evaluate the advantages and disadvantages of component-based development in typical scientific applications. Topics of focus will include components and frameworks from a scientific programmer's perspective, development costs, benefits of interoperabilty, and performance.

Joseph Kenny
Sandia National Laboratories
jpkenny@sandia.gov

MS9

Requirements for an End-to-End Solution for the Center for Plasma Edge Simulation Fusion Simulation Project

The Center for Plasma Edge Simulation, CPES, is one of the two funded SciDAC Fusion Simulation Projects. A major emphasis in this project is in coupling components of XGC, a massively parallel Monte Carlo particle-in-cell (PIC) code designed to model plasma in the edge region of tokamaks, to the MHD code M3D. Working with the Scalable Data Management center, we are creating new enhancements to a workflow system (called Kepler) to create a data model for this project that allows for fast NXM data streaming, collaborative data monitoring, and distributed

storage.

Scott Klasky
Oak Ridge National Laboratory
Oak Ridge, TN
klasky@ornl.gov

MS9
The Systems Biology Problem: A Component-based Approach to Computational Sciences

The major challenge for computational science in biology is dealing with disparate forms of data that are increasingly massive in volume and diverse in format. A component architecture approach is increasingly being used to address complex query problems in the domain of biological science by coupling applications and software into powerful new technologies. I will discuss the current advances in applications of computational science for biological challenges and provide a view toward where computing needs to evolve to address future directions.

George Michaels
Pacific Northwest National Laboratory
george.michaels@pnl.gov

MS10
Application Performance on the TeraGrid Systems

Abstract not available at time of publication.

Robert Pennington
National Center for Supercomputing Applications
robp@ncsa.uiuc.edu

MS10
Preliminary Performance Evaluation of the Power5-based Purple System

Abstract not available at time of publication.

Tom Spelce
Lawrence Livermore National Laboratory
spelce1@llnl.gov

MS10
Performance Evaluation of the Cray XT3

Oak Ridge National Laboratory recently received delivery of a 25 TF Cray XT3. The XT3 is Cray's third-generation massively parallel processing system. The system builds on a single processor node-the AMD Opteron-and uses a custom chip-called SeaStar-to provide interprocessor communication. In addition, the system uses a lightweight operating system on the compute nodes. This presentation describes our initial experiences with the system, including micro-benchmark, kernel, and application benchmark results. In particular, we provide performance results for important Department of Energy applications areas including climate modeling and fusion simulation.

Jeffrey S. Vetter, Patrick H. Worley, Mark Fahey
Oak Ridge National Laboratory
vetterjs@ornl.gov, worleyph@ornl.gov, faheymr@ornl.gov

Sadaf Alam, Tom Dunigan, Philip Roth
ORNL
alamsr@ornl.gov, duniganth@ornl.gov, rothpc@ornl.gov

MS10
Performance Characterization of the Cray X1E

The Cray X1E has a number of novel architectural features that affect performance. While utilizing a vector processor, it also has a vector cache and an additional level of fine grain parallelism called streaming. It supports low latency distributed shared memory programming models such as Co-Array Fortran, but has relatively high MPI latency. This talk presents an overview of the architecture and describes performance data that illustrate the promise and peculiarities of the system.

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

MS11
A Framework for Efficient Multigrid on High Performance Architectures

The hierarchical hybrid Grids (HHG) framework attempts to remove limitations on the size of problem that can be solved using a finite element discretization of a partial differential equation (PDE) by using a process of regular refinement, of an unstructured input grid, to generate a nested hierarchy of patch-wise structured grids that is suitable for use with geometric multigrid. The regularity of the resulting grids may be exploited in such a way that it is no longer necessary to explicitly assemble the global discretization matrix. This drastically reduces the amount of memory required for the discretization, thus allowing for a much larger problem to be solved. Here we present a brief description of the HHG framework, detailing the principles that led to solving a finite element system with 1.7x1010 unknowns, with an overall performance of 0.96 TFLOP/s.

Ben Bergen
Los Alamos National Laboratory
bergen@lanl.gov

MS11
Scalable Algebraic Multigrid on 3500 Processors

A parallel algebraic multigrid linear solver method is presented which is scalable to thousands of processors on significant classes of two- and three-dimensional problems. The algorithm is entirely algebraic and does not require prior information on the physical problem. Scalability is achieved through the use of an innovative parallel coarsening technique in addition to aggressive coarsening and multipass interpolation techniques. Details of this algorithm are presented together with numerical results on up to several thousand processors.

Wayne Joubert
Xiylon Software
joubert@pobox.com

MS11
Parallel Algebraic Methods for Discretization and Solution of Finite Element Problems

While classical AMG works well on many finite element matrices, complicated problems such as higher-order elements, elasticity and electromagnetics require additional

information, that often can be extracted only from the local stiffness matrices. We will describe a parallel element-based algebraic multigrid method which uses local trace minimization to construct high-quality vector-preserving interpolation operators. We will also discuss the assembling of the stiffness matrix, the parallel element agglomeration procedure and some numerical results.

Panayot Vassilevski
Larence Livermore National Laboratory
panayot.llnl.gov

Tzanio Kolev
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
kolev1@llnl.gov

MS11

The Performance of Parallel AMG Methods for Systems of PDEs on BlueGene/L

Algebraic multigrid (AMG) is a very efficient, scalable algorithm for solving large linear systems on unstructured grids. When solving linear systems derived from systems of partial differential equations (PDEs) often a different approach is required than for those derived from scalar PDEs. We will give a brief overview of AMG and present several parallel implementations of various approaches for solving systems of PDEs with AMG. Advantages and disadvantages of these approaches are discussed, and numerical results on BlueGene/L are presented.

Ulrike Meier Yang
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
umyang@llnl.gov

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

MS12

On the Use of Inexact Subdomain Solvers for BDDC Algorithms

The BDDC preconditioner is shown to be equivalent to a domain decomposition preconditioner written in terms of solving a partially subassembled finite element model. The connection of another with the FETI–DP algorithm with a lumped preconditioner is also considered. Multigrid methods are used to replace exact solutions in the preconditioners and under certain assumptions, it can be established that the iteration count essentially remains the same as when using exact solvers.

Jing Li
Department of Mathematics
Kent State University
li@math.kent.edu

Olof Widlund
Courant Institute
widlund@cims.nyu.edu

MS12

FETI-DP for Higher Order Methods

Nonoverlapping domain decomposition methods of the FETI-DP and BDDC type have been widely in use. We present sequential and parallel performance results of FETI-DP and BDDC algorithms including an inexact variant of the FETI-DP method for spectral elements. This is joined work with Axel Klawonn, Essen and Luca Pavarino, Milan.

Oliver Rheinbach
University of Duisburg-Essen
oliver.rheinbach@uni-essen.de

MS12

BDDC Algorithms for Flow in Porous Media

BDDC algorithms are developed for the linear systems arising from flow in porous media discretized with mixed and hybrid finite elements. Our methods are proven to be scalable and the condition numbers of the operators with our BDDC preconditioners grow only polylogarithmically with the size of the subdomain problems.

Xuemin Tu
Courant Institute
xuemin@cims.nyu.edu

MS12

An Introduction to FETI-DP and BDDC Algorithms and Some Recent Results

The FETI-DP and BDDC family of algorithms in many ways represent the state of the art of iterative substructuring methods, i.e., the domain decomposition algorithms which work with non-overlapping subdomains. The methods will be introduced and their close connections will be explored; recent theory suggests that the rates of convergence are very similar and that other considerations come into play when choosing to use FETI-DP or BDDC. Several extensions to new families of problems will also be mentioned

Olof Widlund
Courant Institute
widlund@cims.nyu.edu

MS13

Operating System Interference Effects At Extreme Scale

For parallel applications which are highly synchronous, the performance of collective operations such as barriers and reductions are important considerations to scalability. The cumulative nature of operating system interference effects is particularly troublesome for large-scale parallel applications. This presentation describes recent investigations undertaken by Lawrence Livermore National Laboratory (LLNL) and IBM to understand and quantify operating system interference effects. Results on LLNL's 12,000+ processor ASC Purple supercomputer are presented.

Terry Jones
NERSC/LBL
dskinner@nersc.gov

MS13

Handling OS Interference Via Migratable Message-Driven Objects

We examine three sources of interference from our experience with applications: those due to handling of subnormal floating point numbers on many machine/OS combinations, those due to OS daemons, and those due to variable communication latencies and overheads. Such interference can damage performance significantly. We show how Charm++ and Adaptive MPI handle this challenge effectively and automatically in many cases, with its asynchronous character, multiple objects per processor, automatic latency tolerance and dynamic load balancing.

Gengbin Zheng
University of Illinois at Urbana-Champaign
gzheng@uiuc.edu

Issaac Dooley
UIUC
idooley2@cs.uiuc.edu

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

MS13

Understanding the Causes of Performance Variability in HPC Workloads

Variability of parallel application performance has broad implications for how much useful work can be produced by a particular HPC system. This work examines performance variations in parallel applications on time scales of years to microseconds, with a focus on understanding the causes of performance variability in multi-user production environments. Performance variability is often the result of contention and reducing variability nearly always equates with increased performance. An evaluation of actions taken to achieve both of these goals is presented.

Bill Kramer, David E. Skinner
NERSC/LBL
kramer@nersc.gov, dskinner@nersc.gov

MS13

Exploring Performance Sensitivity of Distributed Memory Parallel Programs to System Interference

The degradation of parallel program performance due to operating system interference is recognized as an important, but tunable, factor in choosing and configuring parallel platforms. This talk discusses a methodology and prototype tool for studying the sensitivity of message passing code performance to perturbations due to operating system interference and variability in interconnect latency. This work uses messaging traces from MPI-based programs and microbenchmark derived system parameters to explore the performance characteristics of real codes.

Matthew Sottile
Los Alamos National Laboratory
matt@lanl.gov

MS14

Performance Annotations on the BlueGene/L

This talk introduces the idea of using source code annotations for addressing performance issues by permitting (but not requiring) access by the programmer to low-level domain-specific performance optimizations that are not normally performed by compilers. Performance annotations are implemented through source transformation, without directly extending the programming language thus enabling rapid development of optimized code for domain-specific abstractions. We will demonstrate performance improvement achieved through annotation-guided code generation on the Blue Gene.

William D. Gropp
Argonne National Laboratories
gropp@mcs.anl.gov

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

Dinesh K. Kaushik
Argonne National Laboratory
D-247, Bldg. 221, MCS Div
kaushik@mcs.anl.gov

MS14

Performance and Memory Evaluation using the TAU Performance System

High-End applications and systems are evolving towards more sophisticated modes of operation, higher levels of abstraction, and larger scales of execution. This talk will discuss our current and future work on improving the utility of parallel performance tools and evolving their capabilities to meet high-end application requirements for the IBM BlueGene/L platform using the TAU performance system. Advances in instrumentation techniques, development of performance metrics for measurement of application performance using hardware performance counters (PAPI), measurement of available heap memory and headroom available for a process to grow, and visualization techniques for analysis of performance data from tens of thousands of processors will be presented.

Sameer Shende
Dept. of Computer and Information Science
University of Oregon
sameer@cs.uoregon.edu

MS14

Performance Measurement and Analysis on BlueGene/L Using SvPablo

SvPablo is a graphical environment for source code instrumentation and performance analysis. It generates statistical performance data including software performance data and hardware counter metrics. The GUI provides a convenient means for high-level bottleneck detection, multi-level detailed analysis, and scalability and load balancing studies. We will present SvPablo toolkit and the performance analysis results for a scientific application on BG/L, and also the insights on the performance comparisons of BG/L and other HPC systems.

Ying Zhang

University of North Carolina
yingz@unc.edu

MS15
Dynamic Computations in Large-Scale Graphs

Graph abstractions are used in many computationally challenging science and engineering problems. For instance, the minimum spanning tree (MST) problem finds a spanning tree of a connected graph G with the minimum sum of edge weights. MST is one of the most studied combinatorial problems with practical applications in VLSI layout, wireless communication, and distributed networks, recent problems in biology and medicine such has cancer detection, medical imaging, and proteomics, and national security and bioterrorism such as detecting the spread of toxins through populations in the case of biological/chemical warfare, and is often a key step in other graph problems. Graph abstractions are also used in data mining, determining gene function, clustering in semantic webs, and security applications. For example, studies have shown that certain activities are often suspicious not because of the characteristics of a single actor, but because of the interactions among a group of actors. Interactions are modeled through a graph abstraction where the entities are represented by vertices, and their interactions are the directed edges in the graph.

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

MS15
Dynamic Data Driven Finite Element Modelling of Brain Shape Deformation During Neurosurgery

Key challenges for neurosurgeons during tumor resection are to remove as much tumor tissue as possible, minimize removal of healthy tissue and know when to stop resection. These rise due to the intra-operative brain shape deformation that happens as a result of surgical procedures. We are implementing an intra-operative MRI data driven simulation process that will capture the intra-operative brain shape deformation in real time. This will help to improve intra-operative surgical navigation.

Petr Krysl
Structural Engineering Dept
Univ of California San Diego
pkrysl@ucsd.edu

Kim Baldridge
San Diego Supercomputer Center
Univ of California San Diego
kimb@sdsc.edu

Amitava Majumdar
San Diego Supercomputer Center
majumdar@sdsc.edu

DongJu Choi
San Diego Supercomputer Center
Univ of California San Diego
dchoi@sdsc.edu

Simon Warfield, Neculai Archip
Radiology Dept, Brigham and Women's Hospital
Harvard Medical School
warfield@bwh.harvard.edu, narchip@bwh.harvard.edu

MS15
Tacking Obesity in Children

This talk will focus on work in progress in datamining and parametric modeling on medical data in a high end computing environment to study and understand the current student health screening in children and reporting practice. Plans for developing a tracking system to assimilate the screened data, analyze them and evaluate the efficiency and utility of the system based on feedback and development of guidelines for the future implementation of the system, including potential economies of scale will be discussed.

Weimo Zhu
University of Illinois at Urbana-Champaign
weimozhu@uiuc.edu

Maryann Mason
Northwestern University
mmason@childrensmemorial.org

Radha Nandkumar
National Center for Supercomputing Applications
radha@ncsa.uiuc.edu

MS15
Data-Driven Parallelization in Multi-Scale Applications

Conventional parallelization typically decomposes the state space. It scales well when the computational effort arises from simulating a large state-space, but not when the effort arises from simulating for long times. We show how we can parallelize the time domain instead, using data from related simulations. This enables parallelization to scale up to three orders of magnitude greater numbers of processors than conventional parallelization, on important applications involving multiple time-scales in nanotechnology and computational biology.

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

MS16
A Graph Infrastructure for Multithreaded Architectures

The next generation of supercomputers will include massively multithreaded architectures specifically designed to tolerate latency. This is of direct interest to graph algorithm designers since these algorithms typically have poor memory locality. Sandia is developing a small software infrastructure to support algorithm development on the Cray MTA-2/Eldorado multithreaded machines. We will describe this infrastructure and some algorithmic idioms suggested by multithreaded architectures. We will also present some performance results for a few core graph algorithms

Jonathan Berry
Sandia National Laboratories
jberry@sandia.gov

MS16
The Parallel Boost Graph Library

The Parallel Boost Graph Library is a research platform for the development, evaluation, and dissemination of distributed-memory graph algorithms. The core of the Parallel BGL is a set of concepts that state the abstract requirements on graphs, auxiliary data structures, and the underlying communication layers. By implementing algorithms using only these requirements, any component of the Parallel BGL can be replaced with another that implements the same, allowing independent experimentation and optimization of library components.

Douglas Gregor, Andrew Lumsdaine
Open Systems Laboratory
Indiana University
dgregor@cs.indiana.edu, lums@cs.indiana.edu

MS16
Computing Approximate Matchings in Parallel

Computing a (possibly weighted) matching in a graph is an important task in many scientific applications. Although there exist exact polynomial time algorithms for the matching problem these algorithms are often to time consuming to used in practice. Lately several linear time approximation algorithms with bounds as good as 2/3 within the optimal matching have been proposed. In this talk we present results from parallelizing these algorithms.

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

MS16
Scalable Graph-Theoretical Approaches to Biological Network Analysis

Many biological objects are naturally represented as graphs. Examples include metabolic, signaling and regulatory pathways, protein interaction networks, and chemical compound graphs. The elucidation of genome-scale structure-function relationships between these objects necessitates the development of more efficient and effective methods for the comparison of their underlying graphs. We present scalable graph-theoretical approaches to this problem including graph matching, maximum clique finding and maximal clique enumeration. Performance benchmarks on advanced hardware architectures will be also presented.

Yun Zhang, Michael Langston
Computer Science Department
University of Tennessee, Knoxville
yzhang@cs.utk.edu, langston@cs.utk.edu

Nagiza F. Samatova
Oak Ridge National Laboratory
Computer Science & Mathematics Division
samatovan@ornl.gov

Faisal Abu-Khzam
Division of Computer Science and Mathematics
Lebanese American University, Beirut,Lebanon
abukhzam@cs.utk.edu

Henry Suters
epartment of Mathematics and Computer Science
Carson-Newman College
suters@cs.utk.edu

MS17
Parallel Tet Mesh Generation: An Overview

Parallel mesh generation is a relatively new research area between the boundaries of two scientific computing disciplines: computational geometry and parallel computing. In this talk we present a survey of parallel unstructured mesh generation methods. Parallel mesh generation methods decompose the original mesh generation problem into smaller subproblems which are meshed in parallel. We organize the parallel mesh generation methods in terms of two basic attributes: (1) the sequential technique used for meshing the individual subproblems and (2) the degree of coupling between the subproblems. This survey shows that without compromising in the stability of parallel mesh generation methods it is possible to develop parallel meshing software using off-the-shelf sequential meshing codes. However, more research is required for the efficient use of the state-of-the-art codes which can scale from emerging chip multiprocessors (CMPs) to clusters built from CMPs.

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

MS17
Mesh Generation for Non-Regid Registration

Adaptive mesh generation and refinement impose uncommon requirements on runtime support substrate. Conventional communication libraries like MPI are not well-suited for development of mesh generation codes. In this talk we will present the communication library, Clam, which aims the applications of interest, and describe our experience using this code on CoWs and on TeraGrid.

Andriy Fedorov
Harvard Medical School
fedorov@bwh.harvard.edu

MS17
Challenges in Parallel Meshing of Evolving Geometries

The Center for Simulation of Advanced Rockets (CSAR) conducts research on whole-system simulation of solid propellant rockets on highly-parallel computers. Important subsystems include solid mechanics of the propellant and fluid dynamics of the interior flow and exhaust plume. As the geometry of the solid propellant and interior flow evolves, increasingly aggressive mesh adaptation is required, including mesh smoothing, local mesh repair, and in some cases global remeshing. Mesh smoothing relocates vertices to maintain mesh quality without changing element connectivity. The smoothing procedure must be efficient enough to be invoked at every few time steps by every processor and must respect displacement limits imposed by the ALE (Arbitrary Lagrangian Eulerian) formulation of our physics solvers. Local mesh repair is applied for severe geometric deformation, eliminating poor elements in the problematic area while preserving good elements elsewhere. Parallelization of local repair is challenging because the area in need of repair can be shared by an arbitrary number of processors. Global remeshing of the entire do-

main has the advantage of being independent of deterioration of the old mesh, but has the disadvantage of requiring global solution transfer, which entails an expensive geometric search. For a mesh of moderate size, sequential remeshing may suffice, but for a very large one, parallel mesh generation is necessary. For each of these three levels of mesh improvement, our implementations involve some use of third-party software. We will address software integration issues as well.

Damrong Guoy
University of Illinois at Urbana-Champaign
guoy@uiuc.edu

MS17

Parallel 3D mesh generation Algorithms Applied to Complex Biology Problems Using Global Arrays

Abstract not available at time of publication.

Harold E. Trease
Pacific Northwest National Laboratory
het@pnl.gov

MS18

An Approximate BDDC Preconditioner with Applications

BDDC (Balancing Domain Decomposition by Constraints) preconditioners require direct solutions of two linear systems for each substructure and one linear system for a global coarse problem. The computations and memory needed for these solutions can be prohibitive if any one system is too large. An approach is presented for addressing this issue based on approximating the direct solutions by preconditioners. Example parallel computations demonstrate the usefulness of the approach in terms of reduced memory requirements and decreased computing times.

Clark R. Dohrmann
Sandia National Laboratories
crdohrm@sandia.gov

MS18

A Parallel FETI-DP Mortar Method for Second and Fourth Order Elliptic PDEs

We discuss a parallel implementation and performance results on up to 1000 processors of a FETI-DP method for second order elliptic partial differential equations with highly discontinuous coefficients on subdomains with nonmatching triangulations. We also present preliminary results on the extension of this method for fourth order plate problems with mortar interface conditions for complete and reduced Hsieh-Clough-Tocher finite element discretizations.

Nina N. Dokeva
University of Southern California
dokeva@usc.edu

MS18

Parallel FETI Based Algorithms for Numerical Solution of Variational Inequalities

We shall first briefly review our theoretical results concerning optimal algorithms for the solution of bound and/or equality constrained quadratic programming problems. Then we shall review our results on scalability of our FETI-DP and FETI based algorithms for numerical solution of model coercive and semicoercive variational inequalities, including mortar contact discretization. Finally we shall extend the scalability results to the elasticity including 2D problems with given (Tresca) friction and give numerical results with parallel implementation.

Zdenenk Dostal
Ostrava, The Czech Republic
zdenek.dostal@vsb.cz

D. Horak
Technical University Ostrava, The Czech Republic
zdenek.dostal@vsb.cz

D. Stefanica
Baruch College
not available

MS18

BDDC and FETI-DP for Mortar Finite Element Methods

A BDDC (balancing domain decomposition by constraints) algorithm is developed for elliptic problems with mortar discretizations for geometrically non-conforming partitions in both two and three spatial dimensions. The coarse component of the preconditioner is defined in terms of one mortar constraint for each edge/face which is an intersection of the boundaries of a pair of subdomains. A condition number bound of the form $C \max_i\{(1+\log(H_i/h_i))^3\}$ is established. In geometrically conforming cases, the bound can be improved to $C \max_i\{(1+\log(H_i/h_i))^2\}$. This estimate is also valid in the geometrically non-conforming case under an additional assumption on the ratio of mesh sizes and jumps of the coefficients. In addition, this BDDC algorithm is closely related to the FETI–DP algorithm with the Neumann-Dirichlet preconditioner. It is proved that these BDDC and FETI–DP algorithms share the same spectra except possibly for an eigenvalue equal to 1. This fact has been also established in BDDC algorithms applied to conforming finite element approximation. Numerical results present that the BDDC algorithm has 1 as its minimum eigenvalue and the FETI–DP algorithms has all of its eigenvalues grater than 1.

Hyea Hyun Kim
Courant Institute, NYU
hhk2@cims.nyu.edu

Olof Widlund
Courant Institute
widlund@cims.nyu.edu

MS18

Approximate FETI-DP preconditionng for Nonlinear Solid Mechanics

The FETI algorithms are numerically scalable iterative domain decomposition methods for solving equations arising from the Finite Element discretization of second- or fourth-order elasticity problems. The one level FETI method equipped with the Dirichlet preconditioner was shown to be numerically scalable for second-order elasticity problems while the two level FETI method was designed to be numerically scalable for fourth-order elasticity problems. Most recently, the dual-primal FETI method (FETI-DP) has emerged as an attractive alternative to the first FETI

algorithms. Implicit nonlinear solid mechanics applications have a long history at Sandia. Adagio is the latest implicit solid mechanics code that uses FETI-DP as a preconditioner within the framework of a nonlinear conjugate gradient solver. I will report on serial, parallel, approximate, and adaptive algorithmic extensions to FETI-DP that have enabled a larger class of problems to be solved.

Kendall H. Pierson
Sandia National Laboratories
khpiers@sandia.gov

MS19

Using Parallel Coupling to Evolve Multi-Physics Simulations

There is little disagreement that component architecture is a revolutionary step away from traditional software practices and it is also well-accepted that most software evolves from other software. On the face of it, the transition from disparate to integrated multi-physics component-oriented applications requires a "leap". This presentation will focus on using parallel coupling between erstwhile separate applications to evolve a successively more integrated component-oriented code. Examples from practice will be cited.

Rob Armstrong
Sandia National Laboratories
rob@sandia.gov

MS19

Comparison of Approaches to Using the Earth System Modeling Framework

The ESMF is emerging as an open standard in the climate and weather domain. Its component-based, hierarchical architecture is being adopted by a diverse group of modelers, who have experimented with ways to organize their applications using the framework. In this talk we examine level of componentization, modes of execution and sequencing, and use of framework services and data structures in ESMF-based applications from UCLA, MIT, NASA, GFDL, and elsewhere.

Atanas Trayanov
NASA GMAO
trayanov@janus.gsfc.nasa.gov

C. Roberto Mechoso
UCLA
mechoso@atmos.ucla.edu

Max Suarez, Arlindo da Silva
NASA GMAO
max.j.suarez@nasa.gov, arlindo.dasilva@nasa.gov

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

V. Balaji
NOAA GFDL
v.balaji@noaa.gov

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

Cecelia DeLuca
Computing Science Section
National Center for Atmospheric Research
cdeluca@ucar.edu

MS19

Cpl6: Software for Building Parallel Coupled Climate Models

Cpl6 is the coupling software for the Community Climate System Model, version 3. Cpl6 provides a reusable main program and four spokes to which atmosphere, ocean, sea ice and land models are attached to create a coupled climate model. Coupling is achieved by placing calls to the cpl6 library within the component models. All cpl6 data transfer and interpolation routines are fully parallel. Using cpl6, user's have successfully added new models to the CCSM3 system.

Tony Craig, Brian Kauffman
National Center for Atmospheric Research
tcraig@ucar.edu, kauff@ucar.edu

Jay W. Larson
Math and Computer Science Division
Argonne National Laboratory
larson@mcs.anl.gov

Robert Jacob
Mathematics and Computer Science Division
Argonne National Laboratory
jacob@mcs.anl.gov

MS19

Component-Based Multi-Physics Simulations of Fires and Explosions

We describe the challenges in creating detailed simulations of fires and explosions at the University of Utah. We discuss how we address those challenges using a component-based infrastructure that it designed for integrated multi-physics applications, including highly parallel simulations of complex fluid-structure interaction problems. This approach has been used to simulate the response of energetic devices subject to harsh environments such as hydrocarbon pool fires.

Steven G. Parker
Scientific Computing and Imaging Institute
University of Utah
sparker@cs.utah.edu

MS19

Flexible Control over Inter-Component Data Transfers for Multiphysics Simulations

Allowing loose coupling between the components of complex multiphysics simulations has many advantages, including being able to easily incorporate new application components and flexible specification of how the components are connected to transfer data between them. To facilitate efficient data transfers between application components, the InterComm environment enables specifying when data transfers are desired outside the individual components. The talk focuses on the design and implementa-

tion of efficient methods for deciding at runtime when data transfers should occur.

Alan Sussman
University of Maryland
Department of Computer Science
als@cs.umd.edu

MS20

Whole Genome, Physics-based Sequence Alignment for Pathogen Signature Design

Two computational obstacles currently impede hybridization based pathogen detection assay design: the rapid growth in the number of DNA sequences in GenBank and the reliance on heuristic metrics of sequence similarity poorly suited to DNA hybridization. We have developed a high performance computational pipeline incorporating a temperature-dependent, free energy based measure of sequence similarity and assay specific physical constraints. DNA signatures have been designed for both national security and food safety applications.

Wu Feng, Murray Wolinsky, Jason D. Gans
Los Alamos National Laboratory
feng@lanl.gov, murray@lanl.gov, jgans@lanl.gov

MS20

ScalaBLAST: Scalable High-Performance Sequence Alignment

ScalaBLAST is a high-performance BLAST engine built using the Global Array toolkit. This shared memory interface makes it possible to efficiently operate on large sequence databases on distributed and shared memory architectures. ScalaBLAST is implemented using a combination of query scheduling and database sharing. This hybrid approach has enabled large-scale sequence alignments to be performed at the whole-genome scale, opening the door for informatics-driven science discovery using data-driven techniques such as genome context mining.

Jarek Nieplocha
Pacific Northwest National Laboratory
MS-IN: K1-85, P.O. Box 999, ichland, WA 99352
jarek.nieplocha@pnl.gov

Christopher S. Oehmen
Pacific Northwest National Laboratory
Christopher.Oehmen@pnl.gov

MS20

PARALIGN: Rapid Sequence Similarity Searches Utilizing Several Levels of Parallelism

PARALIGN is a rapid and sensitive tool for sequence similarity search tool for identification of distantly related genes and proteins in sequence databases. Two methods are implemented: accelerated Smith-Waterman and ParAlign. Both methods utilize parallelism at several levels. Single-Instruction Multiple-Data (SIMD) technology is used to greatly speed up sequence comparison at the sub-sequence level. This technology, also known as multimedia technology, is available in most modern processors, but rarely used by other bioinformatics software. At a higher level, the database is split and searched in parallel by multiple threads and multiple nodes communicating using MPI. A public search service is available at www.paralign.org

Torbjørn Rognes
University of Oslo
Center for Molecular Biology and Neuroscience
torbjorn.rognes@medisin.uio.no

Jon Myrseth
Center for Molecular Biology and Neuroscience (CMBN)
Rikshospitalet-Radiumhospitalet and University of Oslo
j.r.myrseth@medisin.uio.no

Sten Morten Andersen
Sencel Bioinformatics AS
PO Box 180 Vinderen, NO-0319 Oslo, Norway
sten.morten.andersen@sencel.com

Jon Lærdahl
Center for molecular biology and neuroscience (CMBN)
Institute of Medical Microbiology, and University of Oslo
j.k.lardahl@medisin.uio.no

Per Eystein Sæbo
Center for Molecular Biology and Neuroscience (CMBN)
Institute of Medical Microbiology, and University of Oslo
p.e.sabo@labmed.uio.no

MS20

Efficient Data Handling in Comparative Genome Analysis Applications

Comparative genome analysis is a routine but computationally demanding task in biology. While parallelization of an analysis engine addresses this computation-intensive problem, data-intensive initial preparation and result merging tasks become a performance bottleneck with the exponential growth in sequence database sizes. We present a set of techniques for efficient and flexible data handling in parallel sequence analysis applications. Applying these techniques to mpiBLAST leads to an order of magnitude overall performance and scalability improvement.

Nagiza F. Samatova
Oak Ridge National Laboratory
Computer Science & Mathematics Division
samatovan@ornl.gov

Heshan Lin, Xiaosong Ma
Department of Computer Science
North Carolina State University
hlin2@unity.ncsu.edu, xma@unity.ncsu.edu

Wu-chun Feng
Advanced Computing Laboratory
Los Alamos National Laboratory
feng@lanl.gov

Al Geist
Oak Ridge National Laboratory
Computer Science and Mathematics Division
hlin2@unity.ncsu.edu

MS21

Combinatorial Preprocessing of Matrices for Direct and Iterative Solution of Large Sparse Systems

We discuss a simple technique that is a very powerful tool

for preprocessing large sparse matrices. This technique obtains a permutation that maximizes the modulus of the product of entries on the diagonal of the permuted matrix. A consequent scaling ensures that the permuted diagonal is all one and all off-diagonal entries have modulus less than or equal to one. Since the MC64 code for doing this was included in HSL in 1999, many people have used the code in a wide range of applications some of which we discuss in this talk.

Iain Duff
Rutherford-Appleton Laboratory and CERFACS
duff@cerfacs.fr

MS21

A Parallel Version of GMRES Preconditioned by the Multiplicative Schwarz Method

Domain decomposition provides a class of divide-and-conquer methods suitable for the solution of linear or nonlinear systems of equations arising from the discretization of partial differential equations. For linear systems, domain decomposition methods can be viewed as preconditioners for Krylov subspace techniques. Traditionally, there are two classes of iterative methods that derive from domain decomposition with overlapping: Additive Schwarz and Multiplicative Schwarz. The former can easily be parallelized but the latter is usually a better preconditioner. We present a parallel version of the combination of the GMRES method and of the Multiplicative Schwarz preconditioner. The algorithm is based on a parallel construction of a non orthogonal basis of the Krylov subspace. Parallelism is obtained by pipelining two recursions, namely the Multiplicative Schwarz recursion and the Krylov basis construction : it occurs through a front-wave of these two recursions. Orthogonalization of the basis is performed by a fully parallel algorithm. One of the main goal of the implementation is to avoid almost all the scalar products of the vectors distributed over the processors in order to forbid global communications and therefore to reach a better scalability. The developed code is based on PETSc and MPI standards. In order to obtain a full chain operating as a black box linear system solver, an automatic partitioner is also under development.

Guy Antoine Atenekeng-Kahou
University of Yaounde I and INRIA/IRISA
guy.atenekeng_kahou@irisa.fr

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

Emmanuel KAMGNIA
University of Yaounde I
erkamgnia@yahoo.fr

Roger Blaise Sidje
University of Queensland
rbs@acmc.uq.edu.au

MS21

Preconditioning with Non-Symmetric Permutations

This talk will discuss a preconditioning method based on combining two-sided permutations with a multilevel approach. The nonsymmetric permutation exploits a greedy strategy to put large entries of the matrix in the diagonal of the upper leading submatrix. The method can be regarded as a complete pivoting version of the incomplete LU factorization. This leads to an effective incomplete factorization preconditioner for general nonsymmetric, irregularly structured, sparse linear systems. The algorithm is implemented in a multilevel fashion and borrows from the Algebraic Recursive Multilevel Solver (ARMS) framework. Numerical experiments will be reported on problems that are known as difficult to solve by iterative methods.

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

MS21

A Parallel Hybrid Banded System Solver: the SPIKE Algorithm

We describe a robust hybrid parallel solver "the SPIKE algorithm" for banded linear systems. Two versions of SPIKE with their built-in-options are provided: (a) Recursive SPIKE for handling non-diagonally dominant systems, and (b) Truncated SPIKE for diagonally dominant systems. These SPIKE schemes can be used either as direct solvers, or as preconditioners for outer iterative schemes. Both versions are faster than the direct solvers currently used in ScaLAPACK on parallel computing platforms. The SPIKE scheme is also quite competitive in terms of achieved accuracy for handling systems that are dense within the band. Further, we provide effective versions of the SPIKE algorithm for handling systems that are sparse within the band.

Eric Polizzi
Department of Computer Science
Purdue University
epolizzi@cs.purdue.edu

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

MS22

Query-Driven Visualization

As datasets have grown larger and more complex, many research projects have focused on techniques for "scalable visualization." While these technologies are more capable of processing larger datasets, they don't necessarily improve scientific insight. Query-driven visualization is a different approach that focuses visualization and analysis processing only on data deemed to be "interesting." Our approach to query-driven visualization leverages state of the art indexing and query technology to extract "interesting" subsets of data that are then used as input to a visualization or analysis processing pipeline. This approach holds the potential to foster greater scientific insight by focusing processing resources only on "interesting data."

Wes Bethel
Lawrence Berkeley National Laboratory
ewbethel@lbl.gov

MS22

Approaches for Large-Scale Data Analysis and Vi-

sualization

Placeholder text: Over the years, numerous approaches have been explored for use in high performance, large scale data analysis and visualization. These include parallel implementations of common visualization algorithms (isosurfacing), parallel rendering algorithms (object order and image order) along with techniques to accelerate the underlying algorithms themselves based upon qualities of the data. This presentation will present a survey of these techniques along with a discussion of "lessons learned."

Charles Hansen
University of Utah
hansen@cs.utah.edu

MS22

Efficient Streaming Multiresolution Data Models for Visualization and Analysis

Placeholder text. Of the many approaches to high performance visualization and analysis, data organization and access strategies offer the ability to quickly interrogate, access, analyze and visualize data on platforms ranging from a lowly laptop to a powerful dedicated parallel visualization system. Our strategy is to exploit the coupling between time-critical algorithms and progressive multi-resolution data-structures to realize an end-to-end optimized flow of data from the original source, such as remote storage or large scientific simulation, to the rendering hardware.

Valerio Pascucci
Lawrence Livermore National Laboratory
pascucci1@llnl.gov

MS22

Data Analysis Challenges for Terascale Supernovae Modeling

Large scale radiation-hydrodynamic simulations of core collapse supernovae produce copious amount of data for several reasons. First, the required spatial and energy resolution means that the number of of variables in the problem is quite large. Secondly, this phenomena involves timescales that require long temporal evolutions which produce a lengthy time series of data. These characteristics provide challenges for parallel I/O, data transfer, data management, post-run analysis, and scientific visualization. In this talk we discuss these challenges in detail and describe some of the solutions that we have employed to address them.

Doug Swesty
SUNY at Stony Brook
Dept. of Physics & Astronomy
dswesty@mail.astro.sunysb.edu

MS23

Parallel Methods for Nano/Materials Science Applications

The nano/materials science community is one of the largest users of computer cycles in computer centers all over the world. They were one of the first communities to parallelize and redesign their methods to run efficiently on parallel computers. In recent years through the funding of computational projects such as the DOE scidac projects there has been renewed effort in developing new parallel algorithms for larger nanoscale applications that scale to larger processor counts. The new parallel algorithms developed for these codes will be discussed as well as new parallel methods for solving the physics equations. Performance of these methods/codes will also be discussed on the leading HPC parallel platforms such as the IBM BG/L, Cray X1 and Earth Simulator.

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

MS23

Large-Scale First-Principles Molecular Dynamics for Nanoscience

We present recent applications of first-principles molecular dynamics (FPMD) to nanoscience. FPMD simulations provide a consistent and simultaneous description of atomic trajectories and electronic structure. Recent advances in the development of large-scale parallel computers have considerably increased the size of systems that can be modeled with FPMD. We present examples of applications of FPMD to nanoscience and discuss the current limits of this approach. We also discuss some implementation challenges on large parallel computers, including the BlueGene/L supercomputer.

Francois Gygi
University of California, Davis
Lawrence Livermore National Laboratory
fgygi@ucdavis.edu

MS23

Direct Parallel Quantum Simulations of the Properties of Magnetic

It is becoming increasingly clear that, with continued theoretical and algorithmic development and more powerful computers, direct first-principles simulation of the properties - magnetic moments, magnetic anisotropy energies, thermal stability, and switching modes and times of magnetic nanostructures can play a key role in development of future magnetic storage, magnet sensor, and permanent magnet devices. Here I will outline the theory, algorithms, and our first principle multiple scattering methods for simulating the magnetic properties of surface, bulk, and isolated magnetic nanostructures on parallel computers. To illustrate what can now be achieved I will show results for some magnetic nanostructures that are of direct experimental interest Fe and Co nanowires at Pt-surface step-edges, as well as Fe and FePt-nanoparticles containing many thousands of atoms. Work supported by the DOE-OS through the Offices of Basic Energy Sciences (BES),Division of Materials Sciences and Engineering. Calculations were performed at the National Energy Research Supercomputer Center, the Pittsburg Supercomputer Center, and the ORNL-Center for Computational Sciences. Oak Ridge National Laboratory is managed by UT-Battelle, LLC under Contract No. DE-AC05-00OR22725. Work performed in collaboration with Yang Wang (PSC), Aurelian Rusanu and J. S. Faulkner (Florida Atlantic University), Markus Eisenbach, Don Nicholson (ORNL), Balazs Ujfalussy, Lazlo Szunyogh , Bence Lazarovits (Budapest), Balazs Gyorffy (Bristol) and Peter Weinberger (Vienna)

Malcolm Stocks
Theory Group, M&C Division
Oak Ridge National Laboratory

stocksgm@ornl.gov

MS23

Accelerating Parallel Electronic Nanostructure Calculations Through Bulk Band Preconditioning

The need for efficient iterative parallel eigensolvers in the field of computational nanotechnology is well recognized. This is because most electronic nanostructure calculations lead to discrete eigenvalue problems of very large size. An obstacle to the efficiency of the iterative solvers is that their convergence depends on the condition number of the system at hand. We have developed a preconditioning technique, called bulk band (BB preconditioner), that alleviates this problem and gives a significant computational speed improvement compared to the best currently used preconditioner. The basic idea behind BB is to use the electronic properties of the bulk materials constituent for the nanostructure in designing a preconditioner. The bulk band states can be computed inexpensively from 2 atom unit cells. Since the nanostructure eigenstates are consistent with these bulk band states, the bulk band can be used to construct a preconditioner which would be more efficient than using the plane-wave basis.

Stanimire Tomov
Innovative Computing Laboratory, Computer Science Dept
University of Tennessee, Knoxville
tomov@cs.utk.edu

MS24

Managing Complexity in Parallel Adaptive Multiphyiscs Applications

Massively parallel processing and adaptive methods are being used by advanced scientific and engineering applications to perform simulations with high fidelity models and multiple coupled physics. Managing complexity is challenging and critical to the success of such parallel adaptive multiphysics applications. The Sierra Framework uses mathematical abstractions from set theory to manage this complexity for a family of such applications. This theory and its application to parallel dynamic distributed mesh data management will be presented.

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

MS24

Techniques for Bandwidth Management in the Loci Framework

The Loci framework employs a logic-relational model for the description continuum mechanics solution algorithms. The framework has been used in the development of a wide range of applications ranging from computational combustion to aircraft icing simulations. In this talk we describe techniques that have been used to increase performance of Loci applications on bandwidth limited architectures. These transformations include segmenting work into cache sized blocks and replicating fast computations to reduce bandwidth limited communications.

Edward Luke
Dept. of Computer Science and Engineering
Mississippi State University
luke@cse.msstate.edu

MS24

Efficient Massively Parallel Adaptive Algorithms for Time-Dependent Transport

This talk presents the implementation a general infrastructure for the simulation of physical systems whose evolutions depend on the transport of subatomic particles coupled with other complex physics. We will present the discrete-ordinates method and its most time consuming step, the ordered traversal of the spatial grid in each direction of particle travel. We have implemented the parallel version of this simulation using STAPL, a parallel version of STL under development in our group.

Lawrence Rauchwerger
Computer Science Department
Texas A&M University
rwerger@cs.tamu.edu

MS24

Effective Indexing of Distributed Multidimensional Scientific Datasets

Applications that query into very large distributed multidimensional datasets are becoming more common. These datasets are becoming truly enormous, and are often stored across multiple sites. Therefore, we have been developing indexing algorithms that work well in a distributed computing environment, employing both hierarchy and replication to provide scalability and maximize performance of both searches and updates. We also describe a completely decentralized scheme that takes advantage of the reliability mechanisms inherent in unstructured peer-to-peer-systems.

Alan Sussman
University of Maryland
Department of Computer Science
als@cs.umd.edu

MS25

parpcx: A Parallel Solver for Linear Programming

Surprisingly, there appears to be no publicly available LP solver suitable for solving large problems in parallel on distributed-memory machines (clusters). We describe the design of parpcx, a parallel LP solver based on PCx, which may fill this gap. parpcx goes beyond the original PCx in that it implements more choices of linear algebra methods, allowing experiments and easy comparisons. We present some preliminary results with LPs from sensor placement. Also, we discuss how hypergraph partitioning may be used in LP decomposition methods.

Ojas Parekh
Emory University
ojas@mathcs.emory.edu

Erik G. Boman
Sandia National Labs, NM
Discrete Math and Algorithms
egboman@sandia.gov

MS25

Direct Parallel Solution of Linear Systems of Dimension 10^9

Linear systems of equations resulting from very large op-

timization problems are usually solved by iterative methods. However this needs careful crafting of an appropriate preconditioner, and might still fail to reach adaequate accuracy. Our object-oriented parallel solver OOPS can use direct factorizations for these systems by exploiting block sparsity. We have recently ported the code to massively parallel architechture and are now able to solve stochastic QP problems exceeding 10^9 variables.

Andreas Grothey
School of Mathematics
University of Edinburgh
A.Grothey@ed.ac.uk

Jacek Gondzio
The University of Edinburgh, United Kingdom
J.Gondzio@ed.ac.uk

MS25

A Convex QP Solver Based on Block-LU Updates

Active-set methods have an advantage over interior methods in permitting warm starts. We describe a convex QP method intended for use within SNOPT (a sparse SQP package for constrained optimization). An initial KKT system is factorized by any available method (such as LUSOL or PARDISO). Active-set changes are implemented by block-LU updates that retain sparsity while leaving the original KKT factors intact.

Michael Saunders
Systems Optimization Laboratory (SOL)
Dept of Management Sci and Eng, Stanford
saunders@stanford.edu

Hanh Huynh
Stanford Univ.
hanh@sccm.stanford.edu

MS25

Combinatorial Approaches to the Solution of Saddle-Point Problems in Large-Scale Parallel Interior-Point Optimization

Interior-point methods are among the most efficient approaches for solving large-scale nonlinear programming problems. At the core of these methods, highly ill-conditioned symmetric saddle-point problems have to be solved. We present combinatorial methods to preprocess these matrices in order to establish more favorable numerical properties for the subsequent solution step. We demonstrate the competitiveness of this approach within IPOPT on a large set of test problems from the CUTE and COPS sets, as well as large optimal control problems involving 2D partial differential equations. The largest nonlinear optimization problem solved has more than 12 million variables and 6 million constraints.

Andreas Wächter
IBM T.J. Watson Research Center
Yorktown Heights, NY
andreasw@watson.ibm.com

Michael Hagemann
Computer Science Departmenrt Universitg of Basel
michael.hagemann@unibas.ch

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

MS26

A Hybrid 2D Method for Sparse Matrix Partitioning

Two-dimensional matrix partitioning has been shown to be effective for many irregular applications, including data mining, circuit simulation and electrophoresis. 2D approaches include recursive bipartitioning in the Mondriaan partitioner, fine-grained partitioning provided by PaToH, and coarse-grained partitioning based on multiconstraint hypergraph partitioning (also in PaToH). We present a hybrid adaptive approach combining the recursive and fine-grained techniques to reduce communication volume, and demonstrate its use in non-PDE applications.

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

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

Tristan van Leeuwen
Utrecht University
tleeuwen@students.math.uu.nl

MS26

Parallel Hypergraph Partitioning for Irregular Problems

Parallel simulations of electrical circuits and inhomogeneous fluids do not perform well with standard partitioning techniques. Similarly, many partitioners do not apply naturally to linear programming and network-based simulations. Partitioners that accommodate non-symmetry, heterogeneity, and high connectivity are needed. Hypergraph models naturally represent such applications and provide decompositions with lower communication volume and higher parallel performance. We present parallel hypergraph partitioning algorithms and demonstrate their performance for circuit and density functional theory applications.

Robert Heaphy
Sandia National Laboratories
rheaphy@sandia.gov

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

Erik G. Boman
Sandia National Labs, NM
Discrete Math and Algorithms
egboman@sandia.gov

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

MS26

Parallel Algorithms for Inhomogeneous Fluids Models using Molecular Density Functional Theories

Molecular DFT problems typically involve tens to one hundred or more unknowns per mesh point, requiring the solution of inherently discrete equations, short-to-medium range integral equations and PDEs. In this presentation, we describe a new family of algorithms that provides orders of magnitude improvement in performance and robustness over previous general-purpose PDE-based approaches. We highlight how these problems are unlike PDEs and how we addressed these differences to achieve robust, scalable solutions.

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

Andrew Salinger
Computational Sciences Department, Sandia National Labs
agsalin@sandia.gov

Laura Frink
Sandia National Labs
ljfrink@sandia.gov

MS26

Scalable Electrical Circuit Simulation: Why PDE Algorithms Don't Fit

Unlike mesh based PDE simulation, electrical circuit simulation involves networks of nonlinear device models with no basis in spatial operators. This means many assumptions made in linear and nonlinear algorithms for solution of PDEs may not be appropriate for this class of problems. Development of the Xyce(TM) simulator has lead to development of new algorithms in these areas which can substantially improve robustness and massively parallel performance for this class of problems.

Michael A. Heroux, David Day
Sandia National Laboratories
maherou@sandia.gov, dmday@sandia.gov

Roger Pawlowski
Sandia National Labs
rppawlo@sandia.gov

Robert J. Hoekstra, Eric Keiter
Sandia National Laboratories
rjhoeks@sandia.gov, erkeite@sandia.gov

MS27

Parallelism and Linear Solver Performance in MHD and PIC Simulations of Burning Plasmas

Abstract not available at time of publication.

Mark F. Adams
Columbia University
adams@pppl.gov

MS27

Exploring Mixed MPI/OpenMP Parallelization and Matrix Reordering in implicit MHD Computations

In many applications, the stiffness of extended-magnetohydrodynamics requires implicit time-stepping when applying numerical computation. However, the range of dynamical scales leads to ill-conditioned matrices that have dense blocks resulting from high-order spatial methods. Combining direct methods for approximate matrices with matrix-free iterative solves has proven successful, but parallel efficiency is needed for different and sometimes conflicting aspects of the algorithm. Efforts combining OpenMP and MPI parallelization and applying sparse-matrix reordering using nested dissection are reported.

Christopher Carey, Carl R. Sovinec
University of Wisconsin-Madison
cscarey@wisc.edu, sovinec@engr.wisc.edu

MS27

Numerical Challenges in Extended Magnetohydrodynamics Modeling

Abstract not available at time of publication.

Jin Chen
Princeton Plasma Physics Laboratory
jchen@pppl.gov

MS27

Performance and Experiences on Porting XGC to the Cray X1E and XT3

Abstract not available at time of publication.

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

MS28

DUNE – The Distributed and Unified Numerics Environment

The main aspect of this talk is the DUNE Parallel Grid Interface. Using C++ techniques DUNE allows to use very different grid implementation under a common interface with a very low overhead. We will give an overview over the techinques used and present numerical simulations done on different grids using the same discretization. The abilities of these grids range from sequential structured grid to parallel adaptive grids with load balance and will include conforming as well as non conforming grids.

Peter Bastion
University of Heidelberg, Germany
Peter.Bastian@iwr.uni-heidelberg.de

Christian Engwe
University of Heidelberg
christi@hal.iwr.uni-heidelberg.de

MS28

Efficient Data Management in Parallel Finite Element Applications through Object-Oriented Design

Often times the bookkeeping in parallel unstructured finite

element applications becomes cumbersome. In order to alleviate this burden from the application developer we have created the Simple Parallel Object-Oriented Computing Environment for the Finite Element Method (SPOOCEFEM). The SPOOCEFEM library routines manage the interface between global and local data, mesh partitioning, and synchronizing data. This frees the developer to concentrate on application development without the bookkeeping headache or parallel performance concerns.

Dale Shires
Computer Scientist, High Performance Computing Division
U. S. Army Research Laboratory
dshires@arl.army.mil

Brian Henz
US Army Research Laboratory
bhenz@arl.army.mil

MS28

Generic Interface for Parallel Mesh Infrastructure

Generic programming enables to develop generalized software components that are easily reusable and adaptable in a wide variety of situations. In the development of a SCOREC parallel mesh infrastructure, the FMDB, the design of generic interface for parallel mesh database including the mesh migration and dynamic mesh load balancing procedure which enables the mesh infrastructure reusable and easily refinable as possible with various needs of parallel mesh-based applications is presented.

E. Seegyoung Seol
SCOREC, RPI
seole@scorec.rpi.edu

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

MS28

Scalable Parallel Unstructured Meshing in ParFUM

Unstructured meshes are used in many computationally intensive engineering applications. The dynamic behavior of such applications makes the development of scalable parallel software challenging. The Charm++ Parallel Framework for Unstructured Meshes (ParFUM) allows applications to use parallel unstructured meshes with minimal parallel computing knowledge, and scale to hundreds of processors. Charm++'s message-driven model enables adaptive communication/computation overlap, and incorporates dynamic load balancing. ParFUM includes advanced meshing capabilities, such as parallel mesh adaptivity and collision detection.

Terry Wilmarth
Computational Science and Engineering
University of Illinois at Urbana-Champaign
wilmarth@uiuc.edu

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

MS29

Asynchronous Parallel Pattern Search in the Context of a Globally Convergent Augmented Lagrangian Method

Kolda, Lewis, and Torczon have proposed an augmented Lagrangian algorithm that uses a linearly-constrained pattern search method to solve the subproblem. This algorithm is globally convergent but does not require derivatives for the objective or the nonlinear constraints. In this talk, we discuss an asynchronous parallel implementation using APPSPACK. Parallelism is important for real-life problems where the function and constraint evaluation are often computationally expensive. We will describe the parallel algorithm and present numerical results.

Josh Griffin
Sandia National Labs
jgriffi@sandia.gov

MS29

An Introduction to Parallel Optimization Methods

Abstract not available at time of publication.

Juan Meza
Lawrence Berkeley National Laboratory
JCMeza@lbl.gov

MS29

Using Task Priorities to Design Dynamic, Distributed Numerical Optimization Algorithms for Non-Dedicated, Heterogeneous Computing Environments

Many algorithms for solving numerical optimization problems proceed by dynamically generating, and modifying, a pool of tasks used to direct the search. Such algorithms are attractive candidates for distributed computing since they generate a large number of computationally intensive tasks. The challenge is that the dynamic nature of the task pool makes it difficult, a priori, to devise either external scheduling schemes or internal load balancing schemes that are effective. We discuss a strategy to determine task placement by combining forecasts about the system with priorities assigned to tasks. We give a simple example of such an algorithm and report some preliminary results.

Virginia J. Torczon
College of William & Mary
Department of Computer Science
va@cs.wm.edu

Anne Shepherd
College of William and Mary
plshep@cs.wm.edu

MS29

Parallel Interior-Point Methods Newline for Large Semi-Definite Programming

Semi-Definite Programming (SDP) is an optimization with a linear objective function over positive semidefinite constraints of symmetric matrices. SDP solvers provide efficient methods to solve Linear Matrix Inequality and to compute the ground-state energy in quantum chemistry. To solve extremely large SDPs, the parallel SDP solver (SDPARA) mainly reduces the computation time of

the Schur complement matrix in primal-dual interior-point methods owing to parallel computation power. We also introduce SDPA online solver which enables users without parallel computation environment to solve SDPs via the Internet.

Masakazu Kojima, Mituhiro Fukuda
Tokyo Institute of Technology
Dept of Math & Comp Sciences
kojima@is.titech.ac.jp, mituhiro@is.titech.ac.jp

Katsuki Fujisawa
Tokyo Denki University
fujisawa@r.dendai.ac.jp

Kazuhide Nakata
Tokyo Institute of Technology
knakata@is.titech.ac.jp

Makoto Yamashita
Kanagawa University
Department of Industrial Management and Science
makoto.yamashita@ie.kanagawa-u.ac.jp

MS30

Parallel Multilevel Preconditioners for Simulations on Adaptively Refined Meshes

Adaptive mesh refinement (AMR) is important for problems where the gradient of the solution varies significantly in space and time. AMR offers accurate solution against modest cost since we use high resolution only where necessary. However, adaptive mesh refinement poses a number of challenges for (parallel) sparse iterative solvers, in particular for the preconditioner. While frequent local changes in the mesh and load balancing make most preconditioners impractical, sparse approximate inverse preconditioners (SPAI) remain efficient. They are easily updated for local changes in the mesh and the cost of multiplying by the preconditioner does not change much. However, SPAI needs a large number of nonzero coefficients per column to address the global coupling of the equations and reduce the number of iterations sufficiently. Unfortunately, this tends to make the preconditioner expensive. Therefore, we are developing multilevel SPAI preconditioners that exploit the AMR structure to provide global coupling at very low cost.

Shun Wang
University of Illinois at Urbana-Champaign
shunwang@uiuc.edu

Eric De Sturler
University of Illinois at Urbana-Champaign
Department of Computer Science
sturler@cs.uiuc.edu

MS30

Parallel Multilevel Iterative Linear Solvers for Heterogeneous Field with Adaptive Mesh Refinement

Multilevel method is a promising scalable method for large-scale scientific simulations. Although Gauss-Seidel method is a typical way for smoothing at each level, some preconditioners such as ILU(0) and SPAI can be applied as smoothers. We study various types of smoothers for our multilevel method applied to ill-conditioned groundwater problems with heterogeneous field properties and adaptive mesh refinement. We show the robustness and scalability of these smoothers, including comparison with the single-level method.

Kengo Nakajima
The University of Tokyo
Department of Earth & Planetary Science
nakajima@tokyo.rist.or.jp

MS30

On Preconditioners for Complex Symmetric Systems of Linear Equations in a Parallel Eigensolver

We consider efficiency and effectiveness of preconditioners for complex symmetric linear equations that arise from sparse eigenvalue problems of large scale molecular orbital computation. We implement an incomplete factorization preconditioner with diagonal shifting only for imaginary part, and a complex version of sparse approximate inverse preconditioner. We report the performance of our methods applied to our master-worker type parallel eigensolver, which runs on multiple PC clusters connected through a hybrid of MPI and GridRPC system.

Hiroto Tadano, Masayuki Okada
Graduate School of Systems and Information Science
University of Tsukuba
tadano@mma.cs.tsukuba.ac.jp,
okada@mma.cs.tsukuba.ac.jp

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

Umpei Nagashima
Research Institute of Computational Science
National Institute of Advanced Industrial Science and Techno
u.nagashima@aist.go.jp

MS30

Parallel Incomplete Cholesky Preconditioner with Selective Sparse Approximate Inversion

We consider drop-threshold incomplete Cholesky (ICT) preconditioners for the Conjugate Gradient method on distributed memory multiprocessors. Our preconditioner combines sparse factorization with selective use of sparse approximate inversion. The latter is used in certain distributed supernodal submatrices to enable latency tolerant preconditioner construction and application. We present results on the performance of our preconditioner on large sparse linear systems from a variety of practical applications.

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

Keita Teranishi
Department of Computer Science and Engineering
The Pennsylvania State University
teranish@cse.psu.edu

MS31

A Model for Dynamic Load Balancing on Hetero-

geneous and Non-dedicated Clusters

As clusters become increasingly popular alternatives to custom-built parallel computers, they expose a growing heterogeneity in processing and communication capabilities. Performing an effective load balancing on such environments can be best achieved when the heterogeneity factor is quantified and appropriately fed into the load balancing routines.

In this work, we discuss an approach based on constructing a tree model that encapsulates the topology and the capabilities of the cluster. Thedifferent components of the execution environment are dynamically monitored and their processing and communication capabilities are collected and then used to assign partition sizes in load balancing.

We used the model, called DRUM, to guide load balancing in the adaptive solution of threenumerical problems. The results showed a clear benefit from using DRUM, in clusters with computational or network heterogeneity as well as in non-dedicated clusters. We also showed that DRUM has a total overhead of less than 0.2% of total execution time.

Joseph E. Flaherty
Dept. Computer Science
Rensslaer Polytechnic
flaherje@cs.rpi.edu

Jamal Faik
Oracle USA Inc.
jamal.faik@oracle.com

Jim Teresco
Department of Computer Science
Williams College
terescoj@cs.williams.edu

MS31

Parallel Thermal Computation Model of a PEM Fuel Cell

In this work, the thermal management problem of proton exchange membrane fuel cell (PEMFC) is addressed. Fuel cells are electrochemical devices which efficiently convert chemical energy into DC electricity and some heat (thermal energy). PEMFCs are the most favoured fuel cell type to power electric vehicles of the future as a replacement for the internal combustion engine. However, there are many practical problems to be overcome to make this technology economically competitive. Thermal management and minimizing temperature constraints, which add complexity and cost to the system, is one of the most important issues. To understand the phenomenon, thermal computation models of a PEM fuel cell that have been proposed in the literature involve complex systems of partial differential equations and their resolution can not be done within a reasonable amount of time. Recently, a new model based on the nodal network simulation approach was proposed. In this model, the system is divided into elementary volumes that are dependent of each other according to conductance conditions. In this paper, a parallel algorithm to solve this thermal computation model is proposed together with parallel implementations using PVM and MPI.

Salah Abdelkrim, Abdellah El Moudni, Jaafar Gaber
Laboratoire SeT
Universite de Technologie de Belfort Montbeliard
abdelkrim.salah@utbm.fr, abdellah.elmoudni@utbm.fr,
Jaafar.Gaber@utbm.fr

Rachid Outbib
Laboratoire L2ES
Universite de Technologie de Belfort Montbeliard
rachid.outbib@utbm.fr

MS31

Cluster Computing for a System of Time-Dependent Reaction-Diffusion Equations on a Three-Dimensional Domain with High Resolution

The flow of calcium ions in one human heart cell can be modeled by a system of time-dependent reaction-diffusion equations. The flow is driven by sources at large numbers of discrete points in the three-dimensional domain. The use of parallel clusters allows for the resolution of these points. We present a special-purpose code that exhibits excellent scalability up to at least 32 processors on a distributed-memory cluster with high-performance interconnect.

Matthias K. Gobbert
University of Maryland, Baltimore County
Department of Mathematics and Statistics
gobbert@math.umbc.edu

MS31

Partitioning Algorithms for Parallel Applications on Heterogeneous Architectures

Existing partitioning algorithms provide limited support for load balancing simulations that are performed on heterogeneous parallel computing platforms. On such architectures, effective load balancing can only be achieved if the graph is distributed so that it properly takes into account the available resources (CPU speed, network bandwidth). We developed a graph partitioning algorithm that can address the partitioning requirements of the scientific computations, and can correctly model the architectural characteristics of emerging hardware platforms.

George Karypis
University of Minnesota / AHPCRC
karypis@cs.umn.edu

Irene Moulitsas
University of Minnesota and University of Cyprus
moulitsa@ucy.ac.cy

MS32

Large-Scale First-Principles Molecular Dynamics Simulations on the BlueGene/L Computer

We present the results of large-scale First-Principles Molecular Dynamics (FPMD) applications performed on the BlueGene/L supercomputer using a plane-wave, pseudopotential approach. Simulations were run on up to 65,536 nodes, using up to 131,072 processors, demonstrating that an efficient parallelization of plane-wave based FPMD is possible. Details of the algorithms used and their parallelization strategy will be discussed.

Francois Gygi
University of California, Davis
Lawrence Livermore National Laboratory
fgygi@ucdavis.edu

MS32

Large-Scale Parallel Molecular Dynamics Simulations of Biomolecules

Scaling simulations of biomolecules is challenging: The systems of interests have a fixed size. Typical simulations involve 10k to a few 100k atoms. Parallelizing a small amount of computation of each timestep on a large number of processors is thus challenging. This talk will review techniques used in NAMD to address this fine-grain parallelization issue, and present performance where 92k atom simulations are parallelized over several thousand processors, taking only a few milliseconds per step.

Gengbin Zheng, Eric Bohm, Laxmikant Kale
University of Illinois at Urbana-Champaign
gzheng@uiuc.edu, ebohm@cs.uiuc.edu, kale@uiuc.edu

Jim Phillips
Beckman Institute
University of Illinois at Urbana-Champaign
jim@ks.uiuc.edu

MS32

All Electron Calculations of Biomolecules Using the Fragment Molecular Orbital (FMO) Method and Massively Parallel Computers

The FMO method enables one to calculate electronic structures of very large molecules. In the method, a system is divided into small fragments and ab initio MO calculations are performed on the fragments and their dimers to obtain the energy and the properties of the entire system. The FMO method was applied to the electronic structure calculation of photosynthetic reaction center complex of Rhodopseudomonas viridis. The protein complex contains 20,581 atoms and 77,754 electrons. The single point energy calculation at the FMO-RHF/6-31G* level took 3 days using 600 2.0GHz Opetron CPUs.

Kazuo Kitaura
National Institute of Advanced Industrial Science and Technology (AIST), Japan
kazuo.kitaura@aist.go.jp

MS32

Multibillion-Atom Chemically Reactive and Nonreactive Molecular Dynamics Simulations of Materials

We present our linear-scaling embedded divide-and-conquer (EDC) algorithms: fast reactive force-field (F-ReaxFF) molecular dynamics (MD) and density functional theory (DFT) on adaptive multigrids for quantum mechanical MD. We parallelized these algorithms based on a tunable hierarchical cellular decomposition framework. On 1920 Itanium2 processors, we achieved 0.56 billion-atom F-ReaxFF and 1.4 million-atom (0.12 trillion grid points) EDC-DFT, in addition to 18.9 billion-atom space-time multiresolution MD, with parallel efficiency as high as 0.953.

Rajiv K. Kalia, Aiichiro Nakano, Priya Vashishta
University of Southern California
rkalia@usc.edu, anakano@usc.edu, priyav@usc.edu

MS33

Nonlinear Solvers in Reduced Resistive Magnetohydrodynamics

Abstract not available at time of publication.

David Keyes
Columbia University
APAM Department
david.keyes@columbia.edu

MS33

One and Two-Level Schwarz Methods in Reduced Resistive Magnetohydrodynamics.

We discuss a parallel fully implicit Newton-Krylov-Schwarz algorithm for solving a magnetohydrodynamics problem in two-dimensional space. Current density sheets become nearly singular in the process of magnetic reconnection. This behavior of the solution limits time step sizes in explicit schemes. Our approach is a fully implicit time integration using Newton-Krylov techniques with one- and two-level additive Schwarz preconditioning. We study parallel convergence of the implicit algorithm as implemented and run on an IBM BG/L supercomputer.

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

Florin Dobrian
Columbia University
dobrian@cs.odu.edu

Serguei Ovtchinnikov
University of Colorado at Boulder
Department of Computer Science
serguei.ovtchinnikov@colorado.edu

David Keyes
Columbia University
APAM Department
david.keyes@columbia.edu

MS33

Parallel Issues in Reduced Resistive Magnetohydrodynamics Using Implicit AMR

Abstract not available at time of publication.

Michael Pernice
LANL
pernice@lanl.gov

MS33

Adaptive Mesh Simulations of Tokamak Refueling

Abstract not available at time of publication.

Ravi Samtaney
Princeton Plasma Physics Laboratory
rsamtaney@pppl.gov

MS34

Finite Element Electronic Structure Code Based

on Existing Parallel AMR

Abstract not available at time of publication.

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

MS34

Efficiency of Plane-Wave Density Functional Calculations on the Cray X1E

Plane-wave density functional calculations are a broadly applicable workhorse method in nanoscience, materials science, and condensed matter physics, owing to their accuracy and applicability to large systems. Computationally, execution time is dominated by linear algebra and the evaluation of Fourier transforms. Parallel efficiency is primarily governed by the efficiency of parallel Fourier transforms, over all or a subset of processors. These methods traditionally achieve a high fraction of peak performance on scalar architectures. Here, we discuss the adaptation and efficiency of common plane wave packages to the vector-based Cray X1E, as well as the introduction of an optimized Fourier Transform kernel that can be readily integrated in plane wave density functional theory codes. Work supported by the Laboratory Directed Research and Development Program of Oak Ridge National Laboratory as well as the Scientific User Facilities Division, U. S. Department of Energy. Work performed in collaboration with T.C Schulthess, M Fahey (ORNL), J. Levesque, N Troullier (Cray).

Paul Kent
University of Cincinnati, Oak Ridge National Laboratory
kentpr@ornl.gov

MS34

Diagonalization Methods in Real Space Density Functional Theory

Density Functional Theory (DFT) is a successful technique used to determine the electronic structure of matter which is based on a number of approximations. It converts the original n-particle problem into an effective one-electron system, resulting in a coupled one-electron Schrödinger equation and a Poisson's equation. This coupling is nonlinear and rather complex. It involves a charge density ρ which can be computed from the wavefunctions ψ_i, for all occupied states. However, the wavefunctions ψ_i are the solution of the eigenvalue problem resulting from Schrödinger's equation whose coefficients depend nonlinearly on the charge density. This gives rise to a non-linear eigenvalue problem which is solved iteratively. The challenge comes from the large number of eigenfunctions to be computed for realistic systems with, say, hundreds or thousands of electrons. We will discuss a parallel implementation a finite difference approach for this problem with an emphasis on diagonalization. We will illustrate the techniques with our in-house code, called PARSEC. This code has evolved over more than a decade as features were progressively added and the diagonalization routine, which accounts for the biggest part of a typical execution time, was upgraded several times. We found that it is important to consider the problem as one of computing an invariant subspace in the non-linear context of the Kohn-Sham equations. This viewpoint leads to considerable savings as it de-emphasizes the accurate computation of individual eigenvectors and focuses instead on the subspace which they span.

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

MS34

Eigensolvers for Large Electronic Structure Calculations

The solution of the single particle Schroedinger equation that arises in electronic structure calculations often requires solving for interior eigenstates of a large Hamiltonian. The states at the top of the valence band and at the bottom of the conduction band determine the band gap that relates to important physical characteristics such as optical or transport properties. In order to avoid the explicit computation of all eigenstates, a folded spectrum method is employed to only compute the eigenstates near the band gap. In this talk, we compare the folded spectrum method with other techniques, such as conjugate gradient minimization, locally optimal block preconditioned conjugate gradient (LOBPCG) and Jacobi-Davidson.

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

Christof Voemel
Lawrence Berkeley National Lab
Scientific Computing Group, Computational Research Division
cvoemel@lbl.gov

MS35

A Distributed Simulation Framework for Scalable Local-Adaptive Multigrid

Numerical simulation using the parallel-adaptive computing paradigm is a challenging task, especially if distributed memory machines and unstructured meshes are addressed. We present a simulation platform that provides all capabilities to perform highly scalable, large-scale computations under a local-adaptive solution strategy with an optimal complexity multigrid solver as kernel. Several aspects will be discussed in more detail, e.g. dynamic load balancing and migration. For real-world examples we show the maturity of the developed concepts and their realization.

Stefan Lang
IWR,Uni Heidelberg
Germany
Stefan.Lang@iwr.uni-heidelberg.de

MS35

Composition of Components in Multiphysics Applications

Complex multiphysics applications often need to overcome several challenges to achieve accurate and efficient composition of simulation components, including compatibility between data structures, programming languages and parallel decompositions. We describe mechanisms that enable composition by employing a component-based architecture where parallelism and communication concerns are delegated to a separate component. This system enables

complex multiphysics applications that involve wide ranges of length and time scales through transitions in solution algorithms, adaptive mesh-refinement, and scalable parallelism.

Steven G. Parker
Scientific Computing and Imaging Institute
University of Utah
sparker@cs.utah.edu

MS35

Parallel Data Management and Communication in the SAMRAI Structured AMR library

Adaptive mesh refinement (AMR) is becoming an increasingly important simulation methodology for many science and engineering applications. Use of AMR introduces many complexities on parallel systems, both in programming and achieving efficient performance. We describe the data management and communication infrastructure used in SAMRAI. In particular, we will discuss software designs that facilitate communication different datatypes on complex mesh configurations, load balancing problems that have non-uniform workloads, and techniques to reduce adaptive grid overheads on large-scale parallel computer systems.

Rich Hornung
Lawrence Livermore National Laboratory
hornung@llnl.gov

Andrew M. Wissink
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
awissink@llnl.gov

MS36

Parallel Delaunay with 10 Billion Tetrahedrons

Abstract not available at time of publication.

Guy Blelloch
Carnegie Mellon University
blelloch@cs.cmu.edu

MS36

Parallel Delaunay Mesh Generation and Refinement using COTS

Delaunay refinement is a popular method for generating guaranteed quality triangular and tetrahedral meshes. It is based on inserting additional (Steiner) points to meet the element quality requirements. We address the fundamental theoretical problem of resolving the dependencies among the candidate Steiner points as well as the basic practical issue of using the available high quality sequential software in the parallel setting. The goal of this talk is to propose a clear theoretical foundation for developing practical parallel Delaunay refinement algorithms. First, we state the algorithm correctness requirements. Any parallel point insertion procedure should preserve the Delaunay refinement loop invariant, i.e. the mesh has to be conformal (simplicial) and has to satisfy the Delaunay criterion. We show under which circumstances these requirements can be violated. Second, we propose a new definition of Steiner point independence in the context of the parallel Delaunay refinement. We say that two points are Delaunay-independent if and only if they can be inserted into the mesh concurrently without violating the correctness of the algorithm. Third, we formulate and prove a criterion and a sufficient condition of Delaunay-independence. Finally, we propose a parallel Delaunay refinement algorithm which does not require the restructuring of the sequential meshing kernel, and thus allows to reuse the available sequential commercial off-the-shelf software (COTS).

Andrey Chernikov
College of William and Mary
ancher@cs.wm.edu

MS36

Parallel Hex Mesh Generation: An Overview

Abstract not available at time of publication.

Tim Tautges
Sandia National Laboratories
tjtautg@sandia.gov

MS36

Phoenix: A Computational Database System for Generating Massive Tetrahedral Meshes

High-performance parallel computers have created unprecedented new opportunities for scientists and engineers to simulate complex physical phenomena on a larger scale and at a higher resolution than heretofore possible. Unstructured meshes have, accordingly, grown larger and larger. The scalability of a mesh generator hence often imposes a limit on the size of the problem that can be solved. We propose a database method to achieve scalability. Our key idea is to build a new software system, called a computational database system, which maps mesh geometries to existing and new database structures, and execute mesh generation algorithms, such as Delaunay meshing, through a sequence of highly-optimized operations on the underlying database structures. This new approach allows us to generate massive tetrahedral meshes on workstations with limited main memory in a reasonable amount of time. Besides, a mesh thus generated is organized and indexed in a database structure that is capable of supporting efficient queries for many other purposes.

Tiankai Tu
Carnegie Mellon University
tutk@cs.cmu.edu

MS37

An Overview of Support Preconditioning

Support preconditioning is a powerful but not well known class of preconditioning techniques for sparse, SPD linear systems. In contrast to empirical techniques, support theory provides bounds on the generalized eigenvalues (condition numbers). Many support preconditioners are based on graph algorithms, and the field is rich in combinatorial problems. We give a brief overview of different types of support preconditioners. Also, we describe and compare graph-based preconditioners that have been independently developed for network optimization. Finally, we discuss open problems and some parallel computing issues.

Bruce Hendrickson
Sandia National Labs
bahendr@sandia.gov

Erik G. Boman

Sandia National Labs, NM
Discrete Math and Algorithms
egboman@sandia.gov

MS37

Efficiently Solving Linear Systems using Hierarchical Decompositions

We focus on algorithms for solving symmetric diagonally-dominant linear systems. It is known that the solution of these systems can be reduced to Laplacians of graphs. Fifteen years ago, Vaidya showed that one could carefully choose a subgraph of any given graph and use it as a preconditioner with non-trivial convergence guarantees. This thread of work culminated recently in the results of Spielman and Teng, who gave almost optimal algorithms. In a parallel thread of work, Gremban and Miller proposed the construction of preconditioners that may not be subgraphs of the original graph. They called these preconditioners support trees. The construction and analysis of support trees presented additional difficulties that were largely overcome by Maggs et al., who found a key connection with hierarchical decompositions of graphs proposed by Räcke in his seminal 2002 paper. Maggs et al. showed that Räcke's decomposition tree can be used as a support tree with convergence guarantees that exceed those of any given subtree. We extend this line of work by developing tools for analyzing support graphs. We observe that Räcke's decomposition has a collection of properties that allow the design of a sequence of support graphs which can be used in recursive iterative algorithms whose convergence guarantees always improve those of Maggs et al., and match or slightly improve Spielman's and Teng's, at least for a wide and practical class of graphs.

Yiannis Koutis
CMU
CSD
jkoutis@cs.cmu.edu

Bruce Maggs
CSD
CMU
bmm@cs.cmu.edu

R. Ravi
Tepper School of Business
CMU
ravi@cmu.edu

Maverick Woo
CSD
CMU
maverick@cs.cmu.edu

Gary Miller
Carnegie Mellon University
glmiller@cs.cmu.edu

Ojas Parekh
Emory University
ojas@mathcs.emory.edu

MS37

Finite Element Matrix Approximations and Graph-based Parallel Preconditioners

We consider the problem of constructing support graph preconditioners for matrices arising from finite element approximation of partial differential equations. It is known that the standard support graph preconditioning approach is not applicable to finite element stiffness matrices. We have developed a technique that uses element-level transformations to convert the coefficient matrix into one that is amenable to preconditioning by support graphs. Our talk will outline these transformations for typical finite element bases used for diffusion-type problems and discuss the resulting effect on the quality of the support graphs. We will describe approaches to construct support graphs that can be used as parallel preconditioners and present experimental results.

Vivek Sarin
Texas A&M University
Department of Computer Science
sarin@cs.tamu.edu

MS37

Algebraic Support Approximations for Finite Element Matrices

I will describe two techniques for algebraically approximating finite-element matrices using support-theory tools. The first technique, called diagonally-dominant element approximation, is applicable to matrices arising from scalar elliptic problems. It is similar to recent constructions by Boman-Hendrickson-Vavasis and by Sarin, except that it does not require a particular factorization of the element matrices. The second technique, called fretsaw preconditioning, is applicable to a wider class of problems, including linear elasticity. It is also purely algebraic.

Sivan A. Toledo
Tel Aviv University
School of Computer Science
stoledo@tau.ac.il

Doron Chen
Tel-Aviv University
mycroft@tau.ac.il

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

MS38

An Architecture for Reconfiguring Iterative MPI Applications in Dynamic and Heterogeneous Environments

With the proliferation of large scale dynamic execution environments such as grids, the need for providing efficient and scalable application adaptation strategies for long running parallel and distributed applications has emerged. Message passing interfaces have been initially designed with a traditional machine model in mind, which assumes homogeneous and static environments. It is inevitable that long running message-passing applications will require support for dynamic reconfiguration to maintain high performance under varying load conditions. We describe a framework that provides iterative MPI applications with reconfiguration capabilities. Our approach is based on integrating MPI applications with a middleware that supports process checkpointing, process migration, and decentralized strategies for application reconfiguration. We present our architecture for reconfiguring MPI applications and our

experimental evaluation.

Boleslaw Szymanski, Carlos Varela
Rensselaer Polytechnic Institute
not submitted, not submitted

Kaoutar El Maghraoui
Department of Computer Science
Rensselaer Polytechnic Institute
elmagk@cs.rpi.edu

MS38

Automated Dynamic Load Balancing Methods for Parallel Scientific Computing

Most load balancers only active at a certain predefined step. With highly adaptive computations, and changing computing environments, this can lead to unbalanced computation, as the next load balancing phase may not occur for some time. A more automated system will allow the actual load of the computing environment to determine when to go into the next load balancing phase. We examine how to create a cost model, and find a heuristic for this optimization problem.

Luis Gervasio
Rensselaer Polytechnic Institute
Department of Computer Science
gerval@cs.rpi.edu

Joseph E. Flaherty
Dept. Computer Science
Rensslaer Polytechnic
flaherje@cs.rpi.edu

James D. Teresco
Department of Computer Science
Williams College
terescoj@cs.williams.edu

Jamal Faik
Oracle USA Inc.
jamal.faik@oracle.com

MS38

Decentralized Data Management Framework for Distributed Environments

The tremendous growth in data requirements for scientific applications stresses the need for new data placement algorithms. In this talk we introduce distributed, adaptive, and scalable middleware that provides efficient access to data in distributed environments. At the core of our approach are dynamic data placement and location techniques that adapt replica creation to the continuously changing network connectivity and users behavior. The results show that our solution provides better data access performance with lower resource consumption rates.

Boleslaw Szymanski
Computer Science Department
Rensselaer Polytechnic Institute
szymansk@cs.rpi.edu

Houda Lamehamedi
Department of Computer Science
Rensselaer Polytechnic Institute
lamehh@cs.rpi.edu

MS38

Not Available at Time of Publication

Abstract not available at time of publication.

Andrew Lumsdaine
Open Systems Laboratory
Indiana University
lums@cs.indiana.edu

MS39

Optimal Algorithm Choice Through Statistical Modeling

Abstract not available at time of publication.

Victor Eijkhout
University of Texas at Austin
Texas Advanced Computer Center
eijkhout@cs.utk.edu

MS39

Easy Evaluation Way of ILU Preconditioning Effects for Automatically Tuned Iterative Linear Solvers

In Incomplete LU (ILU) preconditioning, orderings and a level of fill-in's affect the effect of preconditioning. The author recently proposed a simple evaluation way for the ILU preconditioning effect. The evaluation index, which has a simple relationship with the matrix norm of the remainder matrix, is easily computed without additional memory requirement. The effectiveness of the method is examined by numerical tests using coefficient matrix data from the Matrix Market and a 3-d electromagnetic field analysis.

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

MS39

ABCLibScript: A Directive to Specify Auto-tuning Facility—Its API and The Application Examples for Several Numerical Methods

A directive named ABCLibScript to specify auto-tuning processes for numerical software is presented. The main objects are performance parameters for loop unrolling depth, cache blocking size, and algorithm selection for dense and sparse kernels. We limit its function to the three specifications: Unroll, Variable, and Select. The API and examples are also shown. A prototype of the preprocessor for ABCLibScript and a visualizer to support user's description will be demonstrated in this presentation.

Takahiro Katagiri
University of California at Berkeley
The University of Electro-Communications
katagiri@is.uec.ac.jp

Kenji Kise, Hiroki Honda, Toshitsugu Yuba
The University of Electro-Communications
kis@is.uec.ac.jp, honda@is.uec.ac.jp, yuba@is.uec.ac.jp

MS39

Automatic Selection of Preconditioners for Library

of Sparse Iterative Solvers

High performance numerical libraries containing automatic selection are highly recommended. We focus on two iterative solvers that use the conjugate gradient (CG) method and the generalized minimal residual (GMRES) method. They are suitable for solving large sparse matrix equations. Our libraries select an appropriate preconditioning method for each type of matrix and target architecture in order to reduce computation time as much as possible. This automatic selection is performed at runtime.

Hisayasu Kuroda
The University of Tokyo
kuroda@cc.u-tokyo.ac.jp

MS40

Normalized Cuts Without Eigenvectors: A Multilevel Approach

Graph partitioning is widely used in scientific computing and parallel processing. Typically the goal is to partition the nodes of a graph so that the graph cut between the partitions is minimized while constraining the partitions to be (nearly) equal in size. Recently, graph partitioning (also called graph clustering) has been used in varied data clustering problems in data mining and machine learning. However, in these applications, there is no inherent need for the graph partitions to be equal in size. Hence weighted graph cuts, such as Normalized Cuts, have been used as the objectives to be minimized. Typically, spectral clustering algorithms that require eigenvector computations are used for this task. In this talk, we first show a theoretical equivalence between the objective function of minimizing weighted graph cuts, and the objective function used in the weighted kernel k-means clustering formulation. However, kernel k-means is prone to local minima and initialization problems. To avoid these problems, we have developed a multilevel approach for efficient and robust minimization of weighted graph cuts that exploits the above equivalence. Experimental results on large data sets underscore the effectiveness of our multilevel approach.

Inderjit S. Dhillon
Department of Computer Sciences
University of Texas, Austin
inderjit@cs.utexas.edu

MS40

Multilinear Algebra for Link Analysis

Linear algebra is a powerful and proven tool in web search. Typically, techniques score web pages based on the principal eigenvector (or singular vector) of a non-negative matrix that captures the hyperlink structure of the web. We propose a new methodology, based on a PARAFAC tensor decomposition, that uses multilinear algebra to elicit more information from a higher-order representation of the web graph. We present numerical results and discuss the computational challenges of scaling these techniques.

Brett W. Bader
Sandia National Laboratories
Computational Sciences Department
bwbader@sandia.gov

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

MS40

Sampling Methods for Data Applications of Linear and Multilinear Algebra

An emerging paradigm for accessing, computing with, and learning from large complex data sets involves the principled use of "sampling". In this approach, one uses only a small portion of the data, and one performs computations of interest for the full dataset by using the small portion as a surrogate. We discuss recently-developed random sampling algorithms for structure identification and extraction from data sets that may be naturally modeled within a linear-algebraic or multi-linear-algebraic framework.

Michael Mahoney
Yahoo! Research
mahoney@yahoo-inc.com

MS40

Numerical Linear Algebra for Data and Link Analysis

Modern information retrieval and data mining systems operate on large datasets and require efficient, robust and scalable algorithms. Numerical linear algebra provides a solid foundation for them. In this talk I will discuss a linear system formulation of Pagerank, a web page ranking algorithm, and Krylov subspace methods for its efficient solution. I will describe a scalable parallel implementation of PageRank and present convergence results of numerical experiments on multiple web graphs.

Leonid Zhukov
Yahoo! Research
zhukovl@yahoo-inc.com

David Gleich
Stanford
dgleich@stanford.edu

MS41

Not Available at Time of Publication

Abstract not available at time of publication.

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

MS41

Lfc: Parallel Computing with Sequential and Interactive Interface

Interactive environments such as Mathematica, Matlab, or Maple offer computational capabilities that are limited to sequential execution. We present a framework that extends these capabilities to parallel execution. Our unique design is independent of the interactive environment. We envision a possibility to switch between interactive environments while keeping the data on the parallel server. We use Lapack For Clusters package that utilizes ScaLAPACK on the server and can be accessed from Matlab or Python clients.

Piotr Luszczek
Department of Computer Science
University of Tennessee, Knoxville
luszczek@cs.utk.edu

MS41
Use of LAPACK and ScaLAPACK in MATLAB

After reviewing how MATLAB 7.0 makes use of LAPACK for matrix computation on single processor machines, we will describe how future versions of MATLAB could make use of ScaLAPACK in distributed memory, multiprocessor environments.

Cleve Moler
The MathWorks, Inc
moler@mathworks.com

MS41
Parallelizing the MRRR algorithm for ScaLAPACK

We present PDSYEVR, the ScaLAPACK version of the parallel MRRR algorithm for the symmetric eigenvalue problem. Depending on progress until the conference, we will also talk about the parallel MRRR-based Singular Value Decomposition.

Christof Voemel
Lawrence Berkeley National Lab
Scientific Computing Group, Computational Research Division
cvoemel@lbl.gov

MS42
Distributed Optimization for Complex System Design

Design of complex systems governed by computational simulations may require distributed optimization approaches for reasons of performance or tractability. In this case a conventional, "non-parallel" nonlinear programming problem formulation is possible, in principle, but a distributed formulation is preferable. In other cases, the components of a complex system are inherently nearly autonomous and require a distributed approach to design and control because non-distributed formulations are unsuitable. We investigate distributed optimization formulations and attendant algorithms for both classes of problems.

Natalia Alexandrov
NASA Langley Research Center
n.alexandrov@nasa.gov

MS42
Aspects of Optimization in Biomedical Flow Simulation

Several challenges and methods in biomedical flow device design are described. The objective often involves the unique behavior of blood as the flowing medium, necessitating, e.g., accurate modeling of cell damage. The complex constitutive behavior, in particular shear-thinning, may affect the outcome of shape optimization more than it affects direct flow analysis. Finally, target applications often involve intricate time-varying geometry, and thus, realistic solutions can be only obtained on high-performance parallel computers.

Mike Nicolai
RWTH Aachen University
nicolai@cats.rwth-aachen.de

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

MS42
Topology Optimization of Fluid Problems based on the Lattice Boltzmann Method

The geometry of a body submerged in a Navier-Stokes flow is optimized for a given performance functional. A new topology optimization approach is presented modeling the flow by the Lattice Boltzmann method (LBM). The LBM is a natural match with topology optimization as in both cases the geometry of a submerged body is defined in an "on-off" mode of computational cells of a fixed regular grid. The LBM eases parallel implementation, features scalability, and is able to deal with complex flows. The proposed design optimization method is verified by numerical studies of 2D and 3D steady-state problems.

Georg Pingen, Anton Evgrafov, Kurt Maute
University of Colorado at Boulder
georg.pingen@colorado.edu, anton.evgrafov@colorado.edu, maute@colorado.edu

MS42
Parallel Hessian Approximations in the Optimization of Systems Governed by P.D.E.

We discuss some parallel approaches to approximating the Hessian in PDE-constrained optimization. These approximations are based on decompositions that take advantage of locality of effects in space and in length-scale. These decompositions may be changed during the course of execution. We discuss techniques for determining when such a change might be advantageous, taking into account the trade-off between acceleration of the optimization and the cost of interprocessor data migration.

Robert M. Lewis
College of William and Mary
rmlewi@wm.edu

MS43
Parallel Adaptive Low Mach Number Simulation of Reacting Flows

Numerical simulation of reacting flows with comprehensive kinetics is one of the most demanding areas of computational fluid dynamics. High-fidelity modeling requires accurate fluid mechanics, detailed models for microcomponent transport and detailed chemical mechanisms. Spatial and temporal requirements based on integration of the compressible reacting Navier Stokes equations on a uniform mesh lead to prohibitive estimates for the computational cost. In this talk we describe methodology to reduce these requirements based on combining a low Mach number formulation with a local adaptive mesh refinement algorithm. The low Mach number equations, which are structurally similar to the incompressible Navier-Stokes equations, enable a larger time step while adaptive refinement reduces the number of degrees of freedom needed to represent the solution. However, integration of the low Mach number equations on adaptive grids introduces additional complexity into the discretization. We will discuss a variety of algorithmic and implementation issues that must be addressed in developing this approach for distributed-memory parallel architectures. We will present an application of the

resulting methodoloty to the simulation of turbulent premixed flames and provide comparisons between simulation and experiment.

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

MS43

VORTONICS: A Grid-Enabled Software Package for the Evolution, Identification, and Visualization of Navier-Stokes Vortices

VORTONICS is a suite of software components for studying vortical motion and reconnection under the incompressible Navier-Stokes equations in three dimensions. It includes three separate modules for the Navier-Stokes solver – a multiple-relaxation-time lattice Boltzmann code, an entropic lattice Boltzmann code, and a pseudospectral solver. It includes routines for dynamically resizing the computational lattice using Fourier resizing, wherein data on the lattice is Fourier transformed, high-k modes are added or deleted to increase or decrease the lattice size respectively, and the inverse Fourier transform is taken. It also includes routines for visualizing vortex cores using thresholding, Q criterion, Delta criterion, lambda-squared criterion, and maximal-line-integral definitions.The package is fully parallel, making use of MPI for interprocessor communication, and it includes routines for arbitrary remapping of data amongst the processors, so that one may transform from pencil to slab to block decomposition of the data, as desired. It also implements a form of computational steering, allowing the user to modify parameters and schedule tasks dynamically. We describe the VORTONICS package with particular attention to its recent grid implementation using MPICH-G2 to allow for both task decomposition and geographically distributed domain decomposition of the computational lattice.

Bruce Boghosian, Lucas Finn, Christopher Kottke
Tufts University
bruce.boghosian@tufts.edu, lucas.finn@tufts.edu, ckottke@gmail.com

MS43

Simulation of Blood Flow in Human Arterial Tree on the TeraGrid

Cardiovascular disease accounts for almost fifty percent of deaths in the western world. The formation of arterial disease such as atherosclerotic plaques is strongly related to the blood flow patterns, and is observed to occur preferentially in regions of separated and recirculating flow such as vessel branches and bifurcations. With grid computing technology we perform simulations of blood flow in the entire human arterial tree through detailed three-dimensional computations at a number of arterial bifurcations, coupled by the wave-like nature of pulse information traveling from the heart to arteries that is modeled by a reduced set of one-dimensional equations. We employ MPICH-G2 and conduct geographically-distributed coupled cross-site simulations at major TeraGrid sites in the US. The algorithms and computation results will be presented.

Leopold Grinberg
Brown University
lgrinb@dam.brown.edu

Alex Yakhot
Ben Gurion University Negev
alex_yakhot@yahoo.com

Spencer Sherwin
Imperial College London
s.sherwin@imperial.ac.uk

Suchuan Dong, George Karniadakis
Brown University
sdong@dam.brown.edu, gk@dam.brown.edu

MS43

Compute and I/O Challenges to Large-Scale Ocean Data Assimilation: Variational and Ensemble-Based Approaches

Data Assimilation (DA) in ocean science applications is a highly compute and I/O intensive process. Among the most powerful modern approaches are variational schemes using an automatically computed adjoint and sequential Monte Carlo ensemble-based non-linear (Error Subspace Statistical Estimation - ESSE) techniques. These types of methods are used for example for global scale ocean estimation, ocean prediction and parameter estimation. These two different methods stress modern computer systems in ways characteristic of a larger set of DA methods. We discuss the significant performance challenges to both regarding computation and I/O and describe how their different needs best fit in a Grid computing environment of distributed parallel platforms.

Constantinos Evangelinos
Massachusetts Institute of Technology
ce107@mit.edu

MS44

A Stable Time-Parallel and Coarseless Implicit Algorithm for Second-Order Hyperbolic Problems

Recently, it was proved that for ODEs resulting from the semi-discretization of second-order hyperbolic problems, most published frameworks for time-parallel computations are inherently unstable. The objective of this paper is to present an alternative framework for time-parallel computations that retains the desired stability properties of a preferred ODE solver. Superior performance is demonstrated over the parareal and original PITA frameworks for several linear oscillator problems including one derived from the discretization of an F-16 aircraft.

Julien Cortial, Henri Bavestrello, Climene Dastillung, Charbel Farhat
Stanford University
jcortial@stanford.edu, hbavestrello@stanford.edu, cdastillung@stanford.edu, CFarhat@stanford.edu

MS44

Time Parallel Time Integrators for Nonlinear Problems

Time parallel time integrators received substantial attention over the past years. One reason for this is the arrival of parallel computers with thousands of processors which lead to saturation of parallelism in space for space-time problems. Another reason is a variant of a time parallel integrator called the parareal algorithm, which has been applied to several problems in industry. I will present an analysis

of the parareal algorithm applied to non-linear problems.

Martin J. Gander
University of Geneva
martin.gander@math.unige.ch

MS44

Parareal in Time Method: Application to Quantum Control and Multi-Model Simulation

We first present in this talk the application of the parareal algorithm with the quantum control by laser fields: the coupling of the monotonic scheme for the optimal control and the parareal algorithm allows for full efficiency of the methodology. Another aspect of the parareal algorithm is presented by the use of simplified coarse schemes for the serial propagation. An application to chemical reactions is also given that proves also large speed up.

Yvon Maday
Université Pierre et Marie Curie
maday@ann.jussieu.fr

MS44

Space-Time Solution of Large-Scale PDE Applications

Software and algorithms are being developed to efficiently formulate and solve transient PDE problems as "steady" problems in a space-time domain. In this way, sophisticated design and analysis tools for steady problems can be brought to bear on transient and periodic problems. This new capability is being developed in the Trilinos solver framework, and allows for parallelism over both space and time domains. Numerical results from a PDE reacting flow application will be presented.

Danny Dunlavy
Sandia National Labs
dmdunla@sandia.gov

Andrew Salinger
Computational Sciences Department, Sandia National Labs
agsalin@sandia.gov

Eric Phipps
Sandia National Laboratories
etphipp@sandia.gov

MS45

Continuous Program Optimization

Our research in the context of the DARPA HPCS project aims at an infrastructure to characterize and understand the interactions between hardware and software and to affect optimizations based on these characterizations. We have developed an architecture for Continuous Program Optimization (CPO) to assist in, and automate the challenging task of performance tuning a system. CPO utilizes the data from all hardware and software layers to detect, diagnose, and eliminate performance problems. We describe the CPO architecture and several CPO agent implementations that we used do demonstrate the capabilities of this framework.

Calin Cascaval
IBM Thomas J. Watson Research Center
cascaval@us.ibm.com

MS45

Issues in Parallel Middleware Design for Many-Core Chips - Cyclops-64 Experience

In this talk, we present a case study - illustrating the problems facing the parallel middleware design of high-end parallel computing systems based on a many-core-on-a-chip architecture. We discuss several issues including execution models with fine-grain multithreading and memory hierarchy optimization. A design study derived from the experience on IBM Cyclops-64 architecture will be discussed.

Guang Gao
Computer Architecture and Parallel Systems Laboratory
Dept. of Electrical and Comp. Eng., University of Delaware
ggao@capsl.udel.edu

MS45

Support for Parallelism in Intel Compilers

Intel Compilers (F95/C++) are used to obtain optimum performance from programs when run on Intel processors. As Intel processors have increasingly incorporated various levels of parallelism for performance, the compilers have evolved to help programmers exploit this parallelism. The talk will describe the many features and their implementation in the compilers to help parallel programming.

Milind Girkar
Intel Compiler Labs
Milind.Girkar@intel.com

MS45

Transactional Memory: Architectural Support for Practical Parallel Programming

As chip-multiprocessors replace uniprocessors in all system types, practical parallel programming is becoming an urgent need. Transactional memory is a new model for shared-memory multiprocessors that relies on user-defined transactions as the basic unit of parallelism, coherence and consistency, error atomicity, and optimization. Transactional memory simplifies parallel programming by eliminating the need for manual orchestration of parallelism using locks. It also provides a flexible environment for dynamic optimizations and fault-tolerant computing. This talk will introduce the software and hardware aspects of transactional memory techniques developed at Stanford University, provide an initial evaluation, and discuss its long-term potential.

Christos Kozyrakis
Departments of Electrical Engineering and Computer Science
Stanford University
christos@ee.stanford.edu

MS45

SmartApps: Adaptive Applications for High Productivity/High Performance

The performance of current parallel and distributed systems continues to disappoint. This is due partially to the inability of applications, compilers and computer systems to adapt to their own dynamic changes. It is also due to the

current compartmentalized approach to optimization: applications, compilers, OS and hardware configurations are designed and optimized in isolation. No global optimization is attempted and the needs of each running application does not constitute the primary optimization consideration. To address these problem we propose application-centric computing, or Smart Applications (SmartApps). In the SmartApps executable, the compiler embeds most runtime system services, and a optimizing feedback loop that monitors the application's performance and adaptively reconfigures the application and the OS/system platform. At run-time, after incorporating the code's input and determining the system's resources and state, the SmartApps performs an instance specific optimization, which is more tractable than a global generic optimization between application, OS and system. The overriding philosophy of SmartApps is "measure, compare, and adapt if beneficial." SmartApps is being developed in the STAPL infrastructure. STAPL (the Standard Template Adaptive Parallel Library) is a framework for developing highly-optimizable, adaptable, and portable parallel and distributed applications. It consists of a relatively new and still evolving collection of generic parallel algorithms and distributed containers and a run-time system (RTS) through which the application and compiler interact with the OS and hardware. SmartApps are developed for several important computational science applications including a discrete-ordinates code simulating subatomic particle transport, a molecular dynamics code and a program using a novel method to simulate protein folding and ligand binding. We use current ASCI shared and distributed machines, networks of workstations that run under Linux and K42 operating systems and, as of late, the Blue Gene machine.

Nancy Amato
Texas A & M University
amato@cs.tamu.edu

Lawrence Rauchwerger
Computer Science Department
Texas A&M University
rwerger@cs.tamu.edu

MS46

GridSolve: A Seamless Bridge Between the Standard Programming Interfaces and Remote Resources

GridSolve is a client-server system that enables users to solve complex scientific problems remotely. The system allows users to access both hardware and software resources distributed across a network. GridSolve searches for computational resources on a network, chooses the best ones available, and using retry for fault-tolerance solves a problem, and returns the answers to the user. A load-balancing policy is used by the GridSolve system to ensure good performance by enabling the system to use the computational resources available as efficiently as possible. Our framework is based on the premise that distributed computations involve resources, processes, data, and users, and that secure yet flexible mechanisms for cooperation and communication between these entities is the key to metacomputing infrastructures.

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

MS46

Architecture-Aware Scientific Computing: Improving Performance and Power Characteristics

We consider techniques for improving the performance of scientific computing codes while reducing the energy consumed by the hardware. We consider several memory subsystem optimizations and their impact on performance and power of certain floating-point intensive benchmarks. Our results indicate that the right mix of architectural features and application tuning can lead to significant reduction in energy and improved performance.

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

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

Mary Irwin
The Pennsylvania State University
Department of Computer Science and Engineering
mji@cse.psu.edu

MS46

Interface of Sparse Linear Solver Library Optimized for Various Types of Architectures

Hardware-dependent tuning for optimization is an important issue for extracting power of modern parallel computers. Moreover, portability is critical, because users can access to variety types of hardware under GRID environment. Sometimes, different algorithm and data structure are required according to individual hardware. In this presentation, features of an infrastructure for efficient development of optimized and reliable scientific applications, such as finite-element method are described. Examples for real applications are also presented.

Kengo Nakajima
The University of Tokyo
Department of Earth & Planetary Science
nakajima@tokyo.rist.or.jp

MS46

EIGADEPT: An Automatic and Application-Based Eigensolver Selection

The automatic eigensolver selection engine, EIGADEPT, is composed of an intelligent engine, knowledge base, decision tree and various numerical libraries. Based on the application and matrix properties, EIGADEPT will use intelligent engine to consult with knowledge base and decision tree to derive an optimal eigensolver from the available numerical libraries connected to EIGADEPT. User also has the capability to add new library into EIGADEPT or provide his own algorithm to solve the eigenvalue problem.

Yeliang Zhang
University of Arizona
zhang@ece.arizona.edu

MS47

Developing Linear Algebra Libraries: Theory and

Practice

The Formal Linear Algebra Methodology Environment (FLAME) project at UT-Austin pursues the mechanical derivation of dense linear algebra libraries. The ultimate goal is to produce a system that will be able to take a mathematical description of a linear algebra operation and from this mechanically produce high-performance library routines together with performance and stability analyses. In this talk we report on progress towards this goal and what can be learned from the experience. This work is in collaboration with a large number of researchers at UT-Austin and elsewhere. Visit

http://www.cs.utexas.edu/users/flame/

For a complete list of participants.

Paolo Bientinesi
Department of Computer Sciences
The University of Texas at Austin
pauldj@cs.utexas.edu

MS47
The Future of LAPACK and ScaLAPACK

We are planning new releases of the widely used LAPACK and ScaLAPACK numerical linear algebra libraries. Based on an on-going user survey (www.netlib.org/lapack-dev) and research by many people, we are proposing the following improvements:

- Faster algorithms (including better numerical methods, memory hierarchy optimizations, parallelism, and automatic performance tuning to accomodate new architectures),
- More accurate algorithms (including better numerical methods, and use of extra precision)
- Expanded Functionality (including updating and downdating, new eigenproblems, etc. and putting more of LAPACK into ScaLAPACK)
- Improved ease of use (friendlier interfaces in multiple languages)

To accomplish these goals we are also relying on better software engineering techniques and contributions from collaborators at many institutions. This is joint work with Jack Dongarra.

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

MS47
Interfaces to Low-Level Kernels

A number of projects, including the ATLAS, PHiPAC, and GotoBLAS efforts, have shown that high-performance BLAS libraries can be built upon a small number of highly optimized "inner-kernels". These insight raise the question of whether the software layers, traditionally maintained by dense linear algebra libraries, need to be redefined. For example, on SMP systems the duplication of the packing of data for locality by different threads could be avoided if routines at the LAPACK library level had direct access to low level kernels rather than through the BLAS libraries. In this talk we review how the best BLAS library implementations are achieved and what opportunities there exist by exposing the natural building blocks of these implementations.

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

MS47
A Critical Look at ScaLAPACK and PLAPACK

This talk will discuss the development of parallel linear algebra libraries with a critical look at what works best in ScaLAPACK and PLAPACK. We show how to develop optimal blocking scheming, including the physical blocking used in ScaLAPACK, with the inner algorithmic blocking used in benchmarks such as HPL, with the outer algorithmic blocking used in PLAPACK. We include both a theoretical and a practical experience with different mappings other than the standard 2D-wrap Cartesian mappings such as was used in a recent Top 500 entry at NASA Ames. We make a case for when the use of dynamically re-blocking the data is useful in practice, and how some of the conventions used in ScaLAPACK actually stand in the way of the fastest performance. We discuss the practical use of recursion and alternate data structures. We make a strong case against the mindless iteration that some projects do in a misplaced effort to mislead their users into believing that the computer is actually doing useful work for them, and discuss how the future of HPC must follow another path if we are to be successful. We make a strong case for the best use of a benchmark such as LINPACK. As a proof of concept, we provide numbers on a final library which achieve superior performance to any other known library or benchmark.

Greg Henry
Intel Corporation
greg.henry@intel.com

MS48
The Optimal Placement of Sensors for Recovering the Source of a Chemical/Biological Attack

The optimal placement of sensors in a building to enable the rapid location of the source of a chemical/biological attack is an important and difficult problem. Our approach to the modeling of the problem postulates a set of possible attack scenarios against which the sensor placement must be effective. We give the details of our sensor placement model and present some numerical results. Finally, we discuss the strengths and weaknesses of our approach.

Stephen Margolis, Kevin Long
Sandia National Laboratories
margoli@sandia.gov, krlong@sandia.gov

Paul Boggs
Sandia National Lab
ptboggs@sandia.gov

MS48
Identification and Prediction of Airborne Contaminant Dispersion Via Terascale Inversion

We are interested in the localization of airborne contaminant releases in regional atmospheric transport models from sparse observations, in time scales short enough for

predictions to be useful for hazard assessment, mitigation, and evacuation procedures. In particular, our goal is rapid reconstruction — via solution of an inverse problem — of the unknown initial concentration of the airborne contaminant in a scalar convection-diffusion transport model, from limited-time spatially-discrete measurements of the contaminant concentration, and from a velocity field as predicted, for example, by a mesoscopic weather model. We employ special-purpose parallel multigrid algorithms that exploit the spectral structure of the inverse operator. Experiments on problems of localizing airborne contaminant release from sparse observations on large parallel systems suggest that ultra-high resolution data-driven inversion may be carried out sufficiently rapidly for simulation-based "real-time" hazard assessment.

Judith Hill
Sandia National Laboratories
jhill@sandia.gov

Volkan Akcelik
Carnegie Mellon University
Laboratory for Mechanics Algorithm
volkan@andrew.cmu.edu

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

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

Andrei I. Draganescu
Sandia National Laboratories
Optimization & Uncertainty Estimation
aidraga@sandia.gov

Bart G. Van Bloemen Waanders
Sandia National Laboratories
bartv@sandia.gov

MS48

Polynomial Chaos Acceleration of Bayesian Methods for Solving Inverse Problems

The Bayesian setting for inverse problems provides a foundation for inference from noisy data and uncertain models and a quantitative description of uncertainty in the inverse solution. For complex PDE-based forward models, however, the cost of likelihood evaluations may render Bayesian methods prohibitive. We address this difficulty by developing a polynomial chaos construction to propagate prior uncertainty through the forward model and efficiently evaluate Bayesian integrals, demonstrating the technique on an inverse scalar transport problem.

Youssef M. Marzouk
Sandia National Laboratories
ymarzou@sandia.gov

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

MS48

Parallel Solution of Optimal Control Problems Using an Inexact SQP Algorithm

We discuss the design, implementation and application of a parallel sequential quadratic programming (SQP) algorithm for the solution of PDE constrained optimization problems. Our algorithm is a matrix-free extension of well-known trust-region SQP methods. Each iteration requires the solution of several KKT type systems. Our SQP method dynamically adjusts the stopping tolerances for all iterative linear system solves based on the progress of the outer iteration. The stopping tolerances can be easily computed. Optimization level domain decomposition preconditioners are used to accelerate the system solves. The algorithm is implemented in C++ and separates the handling of optimization and linear algebra tasks. We present tests using the Trilinos Solver Framework for the linear algebra subtasks. Numerical results include optimal control problems governed by steady-state nonlinear elliptic and incompressible Navier-Stokes equations.

Matthias Heinkenschloss
Rice University
Dept of Comp and Applied Math
heinken@rice.edu

Denis Ridzal
Rice University
Dept. of Computational and Applied Math.
dridzal@rice.edu

MS49

Infini-T: The Infinite Thread Architecture

To ensure program correctness in multiprogramming, synchronization mechanisms are needed to enforce data dependencies and timing constraints between the various threads such as locks and transactions. These synchronization mechanisms are expensive, so programmers must deploy them frugally, making them difficult to reason about and often leading to unwanted race conditions or deadlock. We introduce the stored-processor architecture, which provides hardware support for the waiting mechanism. By facilitating direct communication between the threads, the waiting thread can efficiently suspend until it is woken up by the releasing thread.

Krste Asanovic
MIT EECS Dept
krste@mit.edu

MS49

Transactional Memory: Architectural Support for Practical Parallel Programming

As chip-multiprocessors replace uniprocessors in all system types, practical parallel programming is becoming an urgent need. Transactional memory is a new model for shared-memory multiprocessors that relies on user-defined transactions as the basic unit of parallelism, coherence and consistency, error atomicity, and optimization. Transactional memory simplifies parallel programming by eliminating the need for manual orchestration of parallelism using locks. It also provides a flexible environment for dynamic optimizations and fault-tolerant computing. This talk will introduce the software and hardware aspects of transactional memory techniques developed at Stanford University, provide an initial evaluation, and discuss its long-term

potential.

Christos Kozyrakis
Departments of Electrical Engineering and Computer Science
Stanford University
christos@ee.stanford.edu

MS49

How to Hurt Scientific Productivity

The boundary between software and hardware complexity has been increasingly shifted in the direction of the software over the past few generations of microprocessor architectures. The result is that the path towards optimization is increasingly obfuscated. This talk challenges conventional wisdom in computer architecture and investigates how hardware architects can improve system design in ways that simplify the burden on software engineering and improve scientific productivity.

David Patterson
UC Berkeley CSEE Dept
pattrsn@eecs.berkeley.edu

MS49

Critical Factors and Key Directions For Petaflops Scale Supercomputers

The rate of improvement in the effective performance of modern processor architectures is departing dramatically from historical trends. The decline of brute-force clock rate improvements opens a significant opportunity to consider more elegant approaches to system architecture. This presentation will describe the causal factors that cause machines that are fast on paper to provide sub-par performance on real scientific applications. It goes on to describe a number of architectural alternatives for petaflops-scale scientific computing platforms.

Thomas Sterling
NASA JPL / California Institute for Technolgy
LSU Center For Computation and Technology (CCT)
tron@cct.lsu.edu

MS50

Eigenvalue Solvers for Electromagnetic Fields in Cavities

We investigate the Jacobi-Davidson algorithm for computing a few of the smallest eigenvalues of a generalized eigenvalue problem resulting from the finite element discretization of the time-harmonic Maxwell equation. Multilevel preconditioners are employed to improve convergence and memory consumption of the eigensolver. We present results of very large eigenvalue problems originating from the design of resonant cavities of particle accelerators. We detail our approach for parallelizing the code by means of the Trilinos software framework.

Roman Geus
Paul Scherrer Institut
roman.geus@psi.ch

Marzio Sala
ETHZ Computational Laboratory
marzio@inf.ethz.ch

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

MS50

The Impact of the Two-Level FETI-DPH Iterative Solver on the Performance of the Inverse Shifted Lanczos Method

Domain decomposition-based linear solvers have emerged as powerful algorithms for symmetric positive definite systems. However, such methods are inapplicable to symmetric indefinite problems, $(A - sM)x = b$, arising within the shift-invert Lanczos method for the generalized symmetric eigenvalue problem. The FETI-DPH solver was designed to address such systems when s is a large scalar shift. So this talk focuses on the performance of FETI-DPH and its impact on a shift-invert Lanczos eigensolver.

Charbel Farhat, Philip Avery
Stanford University
CFarhat@stanford.edu, pavery@stanford.edu

MS50

A Rayleigh Quotient Minimization Algorithm Based on Algebraic Multigrid

Several efficient algorithms are available to solve large-scale eigenproblems. In this talk, we will review briefly some approaches and then focus on a multigrid-based algorithm. The RQMG algorithm, introduced by Mandel and McCormick, approximately minimizes the Rayleigh quotient over a sequence of geometric grids. Here we will present an algebraic extension, where we replace the geometric mesh information with the algebraic information defined by an algebraic multigrid preconditioner.

Rich Lehoucq, Ulrich L. Hetmaniuk
Sandia National Laboratories
rblehou@sandia.gov, ulhetma@sandia.gov

MS50

Performance of Parallel Projection Algorithms for Solving Nonlinear Eigenvalue Problems in Electronic Structure Calculation

One of the fundamental problems in electronic structure study is to determine wave functions associated with the minimum total energy of large atomistic systems. These functions are solutions to a nonlinear eigenvalue problem which can be solved by a number of projection algorithms. The parallel performance of these algorithms are examined and compared in this talk.

Lin-Wang Wang
Lawrence Berkeley National Lab
lwwang@lbl.gov

Juan C. Meza
Lawrence Berkeley National Laboratory
JCMeza@lbl.gov

Chao Yang
Lawrence Berkeley National Lab
CYang@lbl.gov

MS51

Runtime Support System for Partially Coupled Grid Computations

Adaptive mesh generation and refinement impose uncommon requirements on runtime support substrate. Conventional communication libraries like MPI are not well-suited for development of mesh generation codes. In this talk we will present the communication library, Clam, which aims the applications of interest, and describe our experience using this code on CoWs and on TeraGrid.

Andriy Fedorov
The College of William and Mary
Computer Science Department
fedorov@cs.wm.edu

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

MS51

Building Symbiotic Relationships Between Formal Verification and High Performance Computing

Computational simulation for scientific and engineering applications are becoming more ubiquitous as part of the engineering design cycle. The application of simulation science to complex problems often require complex models, sophisticated numerics and intricate implementations. Tremendous effort has been expended toward the development of systematic techniques for model validation and numerical method verification. As most researcher's hesitantly admit, the amount of time spent debugging intricate high performance parallel implementations of their simulations consume a large bulk of their time. In particular, many would argue that although this debugging time is necessarily, it distracts one from the science or engineering problem of interest. In this talk, we will present a new effort by the Utah Gauss Group to employ formal verification techniques to the debugging of parallel high performance computing codes using MPI. This synergistic combination of formal techniques with HPC is designed to infuse news ways of thinking about parallel code design through interaction of two normally disparate communities, with the goal of benefiting both communities.

Robert Kirby
University of Utah
kirby@sci.utah.edu

MS51

Developing and Supporting a Large Scale Parallel Application for Scientific Users - The TITAN Geophysical Mass Flow Code Experience

Over the last few years we have developed the TITAN geophysical mass flow modeling toolset. This tool is in widespread use for volcanic hazard mapping. The tool offers, parallel adaptivity as a primary mechanism for rendering solvable some of the computaitionally intractable problems in modeling hazardous mass flows over natural terrain. Optional modules also enable grid based use, dynamic data and stchastic modeling of input parameters. We will present the significant success of the development based on a common parallel infrastructure and aspects that enabled many geologists with minimal expertise in numerical methodology to use this tool successfully.

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

E Bruce Pitman, M. Jones
SUNY Buffalo
pitman@buffalo.edu, jonesm@ccr.buffalo.edu

MS51

From Physical Modeling to Scientific Understanding: An End-to-End Approach to Parallel Supercomputing

Conventional parallel scientific computing uses files as interface between simulation components such as meshing, partitioning, solving and visualizing. This approach results in time-consuming file transfers, disk I/O and data format conversions that consume large amounts of network, storage, and computing resources while contributing nothing to applications. We propose an *end-to-end* approach that replaces the cumbersome file interface with a scalable, parallel, runtime data structure, on top of which all simulation components are constructed. This framework has been implemented for octree-based finite element simulations; performance evaluation for earthquake simulations on up to 2048 Alpha processors shows good good isogranular and fixed-size scalability.

Jacobo Bielak
Carnegie Mellon University
Department of Civil & Env. Eng
bielak@cs.cmu.edu

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

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

Hongfeng Yu
University of California Davis
hfyu@ucdavis.edu

Leonardo Ramirez-Guzman
Carnegie Mellon Universtiy
lramirez@andrew.cmu.edu

David O'Hallaron, Tiankai Tu
Carnegie Mellon University
droh@cs.cmu.edu, tutk@cs.cmu.edu

MS52

PyACTS: A High-Level User Interface to The ACTS Collection

The ACTS Collection brings a robust and high performance set of software libraries to the hands of compu-

tational scientists. However, application developers often dont use tools because of the intricacy in understanding arguments and parameters that mostly relate to running efficiently in parallel computing environments. PyACTS provides a didactical user interfaces for helping users with their first application prototype and subsequent code developments. Here we look at existing functionality in PyACTS, applications and future plans.

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

Vicente Galiano, Violeta Migallon, Jose Penades
Universidad de Alicante
vgaliano@umh.es, violeta@dccia.ua.es, jpenades@dccia.ua.es

MS52

Numerical Policy Interface for Automatic Tuninig Library

Recently, parameters for performance tuning are diverse in many matrix libraries, thus automatic tuning facilities have been required. In this talk, a new framework, numerical policy, is proposed to achieve higher semantics library-user requirement. The framework is applied to Lanczos eigensolver on a SMP platform, one node of the SR11000. The result shows that the framework effectively selects a good parameter value that balances between amount of memory and computation time.

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

Hiroyuki Kidachi, Mitsuyoshi Igai
Hitachi ULSI Systems Cp., Ltd.
kidachi-hiroyuki@hitachi-ul.co.jp, igai@hitachi-ul.co.jp

MS52

Tools for Performance Discovery and Optimization

The TAU Performance System offers integrated performance instrumentation, measurement and analysis capabilities. In this minisymposium, we describe how TAU is applied to evaluate the performance of numerical kernels. Performance profiles generated by TAU are stored in a performance database. The performance database is queried for historical performance data and application-specific metadata. This enables selection of appropriate numerical algorithms to match the given problem. TAU's PerfExplorer framework is used for performance data mining and performance knowledge discovery and optimization.

Sameer Shende
Dept. of Computer and Information Science
University of Oregon
sameer@cs.uoregon.edu

MS52

OSKI: A Library of Automatically Tuned Sparse Matrix Kernels

The Optimized Sparse Kernel Interface (OSKI) is a collection of low-level C primitives that provide automatically tuned computational kernels on sparse matrices, for use by solver libraries and applications. While conventional implementations of sparse matrix-vector multiply run at 10% of peak or less, careful tuning can achieve 31% of peak and 4x speedups. However, tuning depends on the matrix and machine, and is deferred until run-time. This talk reviews the design and implementation of OSKI.

Richard Vuduc
Lawrence Livermore National Laboratory
vuduc2@mail.llnl.gov

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

Katherine Yelick
UC Berkeley
yelick@cs.berkeley.edu

MS53

A Software Framework for Parallel Agent-Based Applications

Agent-based computing can significantly improve the ability to model, design, and build complex systems in a variety of application areas. The large computational requirements of these simulations make the development of scalable software important. The presented software framework includes an agent movement module, which is crucial to spatially explicit agent models on continuous space, and an agent behavior module, which embeds parallel element searching algorithms to support the agents interaction with other agents and well as the surrounding environment.

Robert M. Hunter
Major Shared Resource Center (MSRC)
U.S. Army Engineer Research and Development Center
robert.m.hunter@erdc.usace.army.mil

Preston McAllister
Department of Computer Science and Engineering
Mississippi State University
pbm2@msstate.edu

Cary D. Butler
Information Technology Laboratory
U.S. Army Engineer Research and Development Center
cary.d.butler@erdc.usace.army.mil

Jing-Ru Cheng
Major Shared Resource Center (MSRC)
U.S. Army Engineer Research and Development Center (ERDC)
Ruth.C.Cheng@erdc.usace.army.mil

MS53

Application of a Parallelized Eulerian-Lagrangian-Agent Method (ELAM) for Ecohydraulics and Water Resource Management

We describe performance of a parallelized EulerianLagrangianagent method (ELAM), which integrates (1) a Eulerian mesh framework describing the physical, hydrodynamic, and water quality domains, (2) the Lagrangian framework describing sensory perception and movement trajectories of individuals, and (3) the agent framework

describing the behavior decisions of individuals. ELAMs are designed to mechanistically decode/forecast movement patterns of individual fish and must efficiently handle large numbers of virtual individuals, unstructured mesh elements, and time steps.

John M. Nestler, Robert M. Hunter, R. Andrew Goodwin
U.S. Army Engineer Research & Development Center
john.m.nestler@erdc.usace.army.mil,
robert.m.hunter@erdc.usace.army.mil,
andrew.r.goodwin@nwp01.usace.army.mil

Jing-Ru Cheng
Major Shared Resource Center (MSRC)
U.S. Army Engineer Research and Development Center (ERDC)
Ruth.C.Cheng@erdc.usace.army.mil

MS53

Simulating the Emergence of Hierarchical Organization in Ecological Systems

Abstract not available at time of publication.

Anthony King
Program for Ecosystem Research
Oak Ridge National Laboratory
kingaw@ornl.gov

MS53

Algorithms and Software for Agent-Based Models of Biological Systems

We present algorithms and software for agent-based simulations of biological systems. Our goal is the development of a computational framework for the modeling of the immune response at the cellular level. The framework allows for cellular interactions, with or without internal state, to be modeled with individual agents. Simple protiens can be modeled as volume concentrations. Object motion such as diffusion and chemotaxis are efficiently modeled by statistical means. Results from the use of this software framework to simulate the persisten infection of the Epstein-Barr virus and acute influenza infections of the lung are presented.

Mark Jones
Dept. of Elect. and Comp. Eng.
Virginia Tech
mtj@vt.edu

Paul E. Plassmann
Virginia Tech
plassmann@vt.edu

MS54

Parallel Adjoint-Based Methods for Electromagnetic Shape Optimization

We formulate the problem of designing the low-loss cavity for the International Linear Collider (ILC) as an electromagnetic shape optimization problem involving a Maxwell eigenvalue problem. The objective is to maximize the stored energy of a trapped mode in the end cell while maintaining a specified frequency corresponding to the accelerating mode. A continuous adjoint method is presented for computation of the design gradient of the objective and constraint. The gradients are used within a nonlinear optimization scheme to compute the optimal shape for a simplified model of the ILC. We discuss challenges involved in parallelizing the method.

Lie-Quan Lee
Advanced Computations Department
Stanford Linear Accelerator Center
liequan@slac.stanford.edu

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

Volkan Akcelik
Carnegie Mellon University
Laboratory for Mechanics Algorithm
volkan@andrew.cmu.edu

Kwok Ko
Stanford Linear Accelerator Center
kwok@slac.stanford.edu

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

MS54

PDE-Constrained Optimization Algorithms for Protein-Protein Interactions

We discuss a protein-protein interaction problem in which electrostatic interactions between the solvent and protein surfaces are modeled using a standard classical Poisson-Boltzmann (PB) formulation. The optimization parameters are the charge positions. The problem is highly-nonlinear and combinatorial in nature. We give details for a simplified version in which we fix the geometry of the two proteins and we allow a continuous variation of the position of the charge. We first present a parallel (matrix-free multigrid accelerated) solver for the forward problem and then we present the formulation and scalability results for adjoint sensitivity analysis.

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

MS54

Multi-Level Techniques for Inverse Problems Involving Incompressible Flows

We show that multi-level techniques designed for solving regularized ill-posed problems can be applied successfully to the inverse problem of reconstructing the initial state of a time-dependent incompressible flow from sparse measurements in space and time. We show that, at fine resolution, this inverse problem can be solved at a cost that is a reasonably small multiple of the cost for a forward solve.

Andrei I. Draganescu
Sandia National Laboratories
Optimization & Uncertainty Estimation
aidraga@sandia.gov

Bart G. Van Bloemen Waanders
Sandia National Laboratories
bartv@sandia.gov

MS54

The Lagged Time Method for Parallelizing Optimization with Time Dependent Constraints

In this talk we describe a method to solve optimization problems with time dependent constraints. We propose a new method we call, lagged time, where we separate the computation into two steps that can be done in parallel. This allows us to trivially parallelize our algorithm. We show performance testing which indicate that the method scales well when the number of processors increase.

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

MS55

The Bootable Cluster CD: From Student Lab to Student Cluster

The Bootable Cluster CD (BCCD) provides an instant clustering environment when booted into ram from the CDROM drives of networked workstations. This makes the BCCD ideally suited for existing Windows and Mac labs found at virtually every academic institution. A *complete* and *fully-configured* clustering environment is established *within minutes*, requires *no permanent installation*, and *leaves no trace on the host system* after rebooting, making it an ideal tool for High Performance Computing Education.

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

MS55

Teaching Practices for Introductory Parallel Computing

When introducing parallel computing concepts to students for the first time, using examples that are visual, that show speed-up, that include real-time analysis of results, and that run in a short enough period of time to fit in a standard lecture can help to emphasize key concepts. A series of publicly available examples hosted on the Computational Science Education Reference Desk of the National Science Digital Library will be presented.

David A. Joiner
Kean University
Assistant Professor
djoiner@kean.edu

MS55

Little-Fe: An Inexpensive, Portable Teaching Cluster for All Levels of Computational Science Education

As scientific research moves increasingly towards interdisciplinary collaboration, computational science education languishes on the periphery. A key cause is the lack of computing resources available to the average classroom. We have developed a portable teaching cluster, Little-Fe, which is both easy and inexpensive to replicate. Little-Fe's operating environment is the Bootable Cluster CD, which comes complete with HPC computational tools, applications and curricular material for computational science education.

Charles Peck
Cluster Computing Group, Computer Science
Earlham College
charliep@cs.earlham.edu

Thomas P. Murphy
Contra Costa College
Computer Science
tmurphy@contracosta.edu

MS55

Teaching Parallel Computing to Non-Computer Scientists

Parallel Computing can appear to be difficult to learn, and typically is presented in technical terms with substantial detail. Yet the basic concepts can be straightforwardly described in a non-technical manner, using analogies, storytelling and live demonstrations, so as to be accessible to computing novices. This talk will address techniques for teaching parallelism to diverse audiences, ranging from practicing scientists and engineers to non-programmers.

Henry Neeman
University of Oklahoma
OU Supercomputing Center for Education & Research
hneeman@ou.edu

MS56

A Distributed Packed Storage for Large Parallel Calculations

Our initial motivation was to develop for the GOCE[1] mission efficient parallel solvers that can tackle the very large dense linear least squares problems encountered in gravity field calculations. Since unlike LAPACK, there is no support for packed matrices in the standard parallel library ScaLAPACK, we implemented a packed distributed storage based on ScaLAPACK computational kernels and that exploits the symmetry or the triangular structure of a matrix. We present performance results for the Cholesky factorization and for the updating of the R factor in the QR factorization.

Serge Gratton
CERFACS
gratton@cerfacs.fr

Marc Baboulin
CERFACS, France
baboulin@cerfacs.fr

Luc Giraud
CERFACS, Toulouse, France
giraud@cerfacs.fr

MS56

BLAS 2.5, Householder Bidiagonalization, and Computer Architecture

Due to recent advances in computing matrix singular vectors, the initial reduction to bidiagonal form is typically the most computationally intensive part of a singular value

[1] Gravity field and steady-state Ocean Circulation Explorer, European Space Agency, satellite scheduled for launch in 2006

decomposition. This talk mainly explains how the order of computations can be rearranged to speed Householder bidiagonalization on cache-based computer architectures. A BLAS 2.5 operator combines several BLAS 2 matrix vector operations, increasing data locality. For example, a _GEMVT combines two matrix vector multiplications (_GEMVs) so that both can be performed in only one fetch of the matrix from RAM. For an $n \times n$ matrix, a BLAS 2.5 - BLAS 3 Householder bidiagonalization uses _GEMVT and _GEMM (matrix matrix) operations for all but $O(n^2)$ of a total of $O(n^3)$ flops. Compared to performing the same operations by a BLAS 2- BLAS 3 (_GEMV matrix vector, _GEMM) the number of fetches of data from RAM is approximately halved. The serial BLAS 2.5-BLAS 3 algorithm significantly speeds bidiagonalization on most modern architectures. Whether the BLAS 2.5-3 algorithm is useful in the distributed memory parallel case depends on cache size and communication latency. If the time to start a communication is much less than the time to refill cache memory, then parallel _GEMVT is faster than using parallel _GEMV operations. Then a BLAS 2.5-3.0 diagonalization is desirable. Automated adaptation of algorithm to architecture and extensibility to the sparse case will also be discussed.

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

J. W. Demmel
University of California, Berkeley
demmel@cs.berkeley.edu

M. Premkumar, C.T. Fulton
Florida Institute of Technology
premkumar@math.fit.edu, fulton@math.fit.edu

Sven J. Hammarling
Numerical Algorithms Group Ltd
Wilkinson House
sven@nag.co.uk

MS56

Parallel Algorithms and Software for Solving Standard and Generalized Sylvester-Type Equations

We develop parallel ScaLAPACK-style algorithms based on the Bartels-Stewart's method for solving continuous-time and discrete-time standard and generalized Sylvester-type matrix equations. The reduced triangular problems are solved using explicit matrix blocking and wavefront-like traversal of the right-hand-side matrices. Node solvers are either recursive or explicit blocked variants. We apply our solvers in condition estimation of the matrix equations, developing parallel Sep^{-1}-estimators. Our deliverables include a software package SCASY which contains general and triangular solvers and condition estimators.

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

MS56

Updating-Downdating Factorizations in ScaLAPACK

We present some parallel distributed software based on the ScaLAPACK kernels for performing the updating and downdating of various factorizations of ScaLAPACK (QR, LU and Cholesky). Example codes with performance and scalability results are given. Issues when parallelizing the sequential algorithm are described.

Julien Langou
The University of Tennessee
langou@cs.utk.edu

MS57

Programming Multi-Core Processors

Mainstream multi-core computers require mainstream parallel applications. The software industry must fundamentally change to embrace concurrency and make parallel software the norm. This presentation will explore the programming models and languages that will enable this transformation to a brave new world of ubiquitous parallel computing.

Ali-Reza Adl-Tabatabai
Programming Systems Lab
Intel Corporation
ali@gomez.sc.intel.com

MS57

Multi-Core in the Real World: The IBM Cell Processor as an Early Multi-Core Product

Abstract not available at time of publication.

Peter Hofstee
IBM
hofstee@us.ibm.com

MS57

What are Multi-Core Processors and Why Are They Inevitable?

All of the major CPU vendors are moving to multi-core products. They aren't a passing fad or the latest marketing ploy. When competing CPU vendors all move in the same direction, you know something is up. Multi-core CPUs are the only way for CPU performance to increase within a fixed power envelope as transistor density continues to climb. In this presentation, we will explain why multi-core CPUs are an inevitable consequence of semiconductor manufacturing. And we will lay out at a high level the opportunities and challenges presented by the move to multi-core CPUs.

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

MS57

Programming Multi-Core processors and SMPs with OpenFLAME

The Formal Linear Algebra Methods Environment (FLAME) project at UT-Austin pursues the systematic derivation of proven correct algorithms for dense linear

algebra operations. The FLAME/C API allows the algorithms to be conveniently coded in the C programming language. It has been used to reimplement large portions of the popular BLAS and LAPACK libraries on sequential processors. For SMP architectures, the API has been extended to use OpenMP, yielding the OpenFLAME API. Through a simple mechanism known as *work queuing* high-performance parallel implementations can be almost trivially achieved. The approach should also support multi-core systems with a medium number of cores (32-64). In this talk we present the approach and show scalability lessons that can be learned. This work is in collaboration with a number of FLAME project participants. For details, visit

`http://www.cs.utexas.edu/users/flame/`

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

MS58

Challenges Facing Future Mesh-based Codes at Large Scale: An Application-centric View

Parallel scientific codes today often make use of structured, static spatial grids partitioned among processors in a regular fashion. This practice not only simplifies the process of mesh generation and application development, but also allows the logical topology of the application to be optimally mapped onto the underlying physical topology of the parallel machine. As illustration, we present two large-scale applications typical of what can be found in national laboratories: HYCOM, a general ocean circulation model, and KRAK, a hydrodynamics code used to simulate explosive detonations. Using performance modeling techniques developed at Los Alamos National Laboratory, we examine how such codes utilize today?s computing resources. With this foundation, we then peer into the near future and surmise how the use of such applications will evolve; as codes make use of larger and more refined spatial grids, dynamic mesh adaptivity may play a critical role as parallel applications strive to make efficient use of parallel resources. At the same time, we examine recent trends in supercomputer architecture and surmise that machines of the future may become more structured in their network topology, as illustrated by IBM?s BlueGene/L machine and the Cray XT3, both of which utilize 3D torus interconnect topologies. We describe one research activity, an intelligent system software toolkit known as the Parallel Runtime Environment for Multi-computer Applications, or PREMA, that goes some way toward bridging the gap between increasingly regular parallel machine topologies and the demands of irregular and unstructured scientific codes.

Kevin Barker
Los Alamos National Laboratory
kjbarker@lanl.gov

MS58

Parallel Out-of-Core Delaunay Refinement Method using COTS

We present a solution for generating very large meshes in a small cluster of workstations environment using commercial off-the-shelf hardware and the best publicly available off-the-shelf sequential in-core Delaunay mesh generator. It is cost-effective since the total wall-clock time including wait-in-queue delays for the out-of-core methods on a small cluster is considerably shorter than that for the in-core generation of the same size mesh (using over a hundred processors). It is also high-performance since the speed of mesh generation of our out-of-core method is comparable to the speed of in-core implementation of the algorithm.

Andriy Kot
Computer Science Department
The College of William and Mary
kot@cs.wm.edu

MS58

Distributed Adaptive and Interactive Simulations Using Structured Adaptive Mesh-Refinement

Structured adaptive simulations can yield highly advantageous cost/accuracy ratios and can enable accurate solutions to realistic models of complex physical phenomenon. However, the phenomena being modeled by these simulations and their implementations are inherently dynamic and heterogeneous, and their large scale implementation presents many challenges. In this presentation I will outline the challenges of distributed SAMR and will present a computational engine that enables efficient and scalable implementations of these applications. Specifically, I will address dynamic data-management, adaptive load-balancing, and interactive monitoring and control.

Manish Parashar
Department of Electrical & Computer Engineering
Rutgers University
parashar@caip.rutgers.edu

MS58

Resource-Aware Parallel Scientific Computation for Heterogeneous and Non-Dedicated Clusters

Many modern parallel computing environments include heterogeneous resources or allow shared access to resources. We present a system called the Dynamic Resource Utilization Model to help guide dynamic load balancing in the presence of heterogeneity or non-dedicated usage of processing and communication resources. DRUM combines information from a static description of a computing environment with information from a dynamic monitoring system to produce an appropriate distribution of work among the cooperating processes.

James D. Teresco
Department of Computer Science
Williams College
terescoj@cs.williams.edu

PP0

A Parallel Multishift QZ Algorithm for Generalized Eigenvalue Problems

The QZ algorithm reduces a regular matrix pair (H, T) in Hessenberg-triangular form to a generalized Schur form (S, T). A novel parallel implementation of a small-bulge multishift QZ algorithm for solving the generalized eigenvalue problem is presented. The algorithm chases chains of small bulges instead of only one bulge in each QZ iteration, which makes it more amenable in a parallel setting. In addition, advanced deflation strategies are used, including aggressive early deflation techniques. Recent progress

will be reported including parallel performance results.

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

Björn Adlerborn, Daniel Kressner
Umeå University
adlerborn@cs.umu.se, kressner@cs.umu.se

PP0

Comparing Parallel CAMR and SAMR Hydrodynamics Simulations

We compare the relative accuracy and performance of two parallel 3D Eulerian adaptive mesh refinement (AMR) codes, one cell-wise and one patch-based, on several hydrodynamics test problems. Using software developed for analyzing the comprehensive performance behavior of large-scale parallel applications, we investigate relationships between solution time, processor counts, communication amount, message sizes, active- and ghost-cell counts, time steps per level, load balancing efficiency, flop rates, and memory access rates.

Michael Norman, James Bordner, Alexei Kritsuk
University of California, San Diego
mnorman@cosmos.ucsd.edu, jbordner@cosmos.ucsd.edu, akritsuk@cosmos.ucsd.edu

Robert Weaver
Los Alamos National Laboratory
rpw@lanl.gov

PP0

Scalable Implicit Solutions for MHD Models

We describe scalable implicit solutions for time dependent, two dimensional, stream function based, magnetohydrodynamics (MHD) models. We employ a fully implicit Newton-Krylov-Schwarz approach, the linear systems being preconditioned with one level additive Schwarz and incomplete LU on subdomains, and solved with GMRES iterations. The space discretization, based on finite differences, is second order and the time discretization, based on backward differentiation formulas is second, third or fourth order.

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

Florin Dobrian
Columbia University
dobrian@cs.odu.edu

Serguei Ovtchinnikov
University of Colorado at Boulder
Department of Computer Science
serguei.ovtchinnikov@colorado.edu

David E. Keyes
Columbia University
Department of Applied Physics & Applied Mathematics
kd2112@columbia.edu

PP0

RECSY and SCASY: Library Software for Matrix Equations

High-performance software for solving common continuous-time and discrete-time standard and generalized Sylvester-type matrix equations are presented. RECSY uses recursive blocked algorithms for solving one-sided and two-sided matrix equations on serial and SMP nodes. The recursive approach leads to an automatic variable blocking amenable for deep memory hierarchies. SCASY uses explicit matrix blocking in parallel ScaLAPACK-style algorithms for solving similar equations on distributed memory platforms. Both RECSY and SCASY provide condition estimation functionality.

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

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

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

PP0

A Case Study in Tools and Techniques for Optimization on the Power5

This presentation, in poster format, will detail the various tools and techniques that were used to optimize a Fast Poisson Solver code using one of a class of spectral methods. We cover both the tools that we used for profiling the program, including dynamic instrumentation, as well as various techniques available to optimize numerically intensive kernels in the POWER5 processor environment.

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

PP0

Performance Intercomparison of Community Atmosphere Model on High-End Computing Platforms

We analyze the performance of the Community Atmosphere Model when run at high resolution with the finite-volume dynamical core on leading high-end computing systems. This includes the vector-based Cray X1E and Earth Simulator, IBM systems including the Power-5, and the Intel Itanium 2 cluster. We consider various parallelization and communication paradigms. The throughput on a per-processor basis is highest on the Cray-X1E, which outperforms the IBM Power-3 by close to an order of magnitude.

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

Arthur A. Mirin
Lawrence Livermore Nat'l Lab.

mirin@llnl.gov

Leonid Oliker, Michael Wehner
Lawrence Berkeley National Laboratory
loliker@lbl.gov, mfwehner@lbl.gov

William Sawyer
Swiss Federal Institute, Zurich
sawyer@env.ethz.ch

David Parks
NEC Solutions America
david.parks@necsam.com

PP0

Modeling Hazardous Geophysical Mass Flows

We describe in this poster the development of a toolkit for modeling a range of hazardous geophysical mass flows. These flows range from hundreds of meters to tens of kilometers in runout length and flow durations often in the hours. Large scale parallel adaptive discontinuous Galerkin schemes are essential to enabling realistic simulations incorporating the complex representations of physics needed and stochastic computations needed for hazard analysis.

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

PP0

Measuring and Reducing Network Latency in the 2.6 Linux Kernel

Network latency continues to be a bottleneck for many computational science applications on Beowulf clusters. Using open source tools, we develop technology for making nanosecond precision measurements between timing points within the 2.6 Linux kernel. Using this technology, in conjunction with low-level and application benchmarks, we establish that there is significant network latency in the kernel and identify its origin. Lastly, we survey software-based latency reduction technologies and make recommendations based on our measurements.

Skylar Thompson, Josh McCoy, Alex Lemann, Tobias McNulty, Kevin Hunter
Cluster Computing Group
Earlham College
skylar@cs.earlham.edu, mccoyjo@cs.earlham.edu, lemanal@cs.earlham.edu, tmcnulty@ppckernel.org, hunteke@cs.earlham.edu

Charles Peck
Cluster Computing Group, Computer Science
Earlham College
charliep@cs.earlham.edu

PP0

Little-Fe - The Low-Cost Portable Cluster for Computational Science Education

Many institutions still do not have access to high performance computing platforms for computational science education. Little-Fe is a complete 8 node Beowulf style portable computational cluster. By leveraging the Bootable Cluster CD project, and its associated curriculum modules, Little-Fe makes it possible to have a powerful turn-key educational platform for under $2,500. We describe Little-Fe's hardware and software configuration, including plans for a "do it yourself" version.

Skylar Thompson, Josh McCoy, Alex Lemann, Tobias McNulty, Kevin Hunter
Cluster Computing Group
Earlham College
skylar@cs.earlham.edu, mccoyjo@cs.earlham.edu, lemanal@cs.earlham.edu, tmcnulty@ppckernel.org, hunteke@cs.earlham.edu

Charles Peck
Cluster Computing Group, Computer Science
Earlham College
charliep@cs.earlham.edu

PP0

Folding@Clusters: Harnessing Grid-Based Parallel Computing Resources for Molecular Dynamics Simulations

Folding@Clusters is an adaptive framework for harnessing grid-based cluster resources for protein folding research. It combines capability discovery, load balancing, process monitoring, and checkpoint/re-start services to provide a platform for molecular dynamics simulations on a range of grid-based parallel computing resources. We examine the design, implementation, and field-trial results for the first major release of the software, and discuss ways to improve it based on our initial production experiences.

Skylar Thompson, Josh McCoy, Alex Lemann
Cluster Computing Group
Earlham College
skylar@cs.earlham.edu, mccoyjo@cs.earlham.edu, lemanal@cs.earlham.edu

Charles Peck
Cluster Computing Group, Computer Science
Earlham College
charliep@cs.earlham.edu

Joshua Hursey
Computer Science Department
Earlham College
joshh@cs.earlham.edu

Vijay Pande
Stanford University
Pande Lab
pande@stanford.edu

PP0

The Parallel Computation of Radiative Heat Transfer Effects with Non-Grey Absorption Coefficient Models

We consider the problem of computing radiative heat transfer effects in simulations where complex, non-grey absorption models are required. These calculations are implemented within a parallel, photon Monte Carlo software framework. The implementation's parallel performance depends on the determination of a load balancing strategy that takes into account temperature and absorption coefficient differences within the computational domain. We present computational results illustrating our proposed

load balancing strategy for a variety of representative combustion problems.

Paul E. Plassmann
Virginia Tech
plassmann@vt.edu

Ivana Veljkovic
Penn State University
Department of Computer Science and Engineering
icotaki@psu.edu

Kamal Viswanath
Virginia Tech
kamalv@vt.edu

PP0

Parallel Rosenbrock Methods for Solving Option Pricing Models

In this paper we consider a meshfree radial basis function approach for the valuation of pricing options with non-smooth payoffs. By taking advantage of parallel architecture, a strongly stable and highly accurate time stepping method is developed with computational complexity comparable to the implicit Euler method implemented concurrently on each processor. This, in collusion with the radial basis function approach, provides an efficient and reliable valuation of exotic options, such as digital options and options with multiple strike prices.

David A. Voss
Western Illinois University
d-voss1@wiu.edu