IP1

Building Watson: A Brief Overview of DeepQA and the Jeopardy! Challenge

Watson, named after IBM founder Thomas J. Watson, was built by a team of IBM researchers who set out to accomplish a grand challengebuild a computing system that rivals a humans ability to answer questions posed in natural language with speed, accuracy and confidence. The quiz show Jeopardy! provided the ultimate test of this technology because the games clues involve analyzing subtle meaning, irony, riddles and other complexities of natural language in which humans excel and computers traditionally fail. Watson passed its first test on Jeopardy!, beating the show's two greatest champions in a televised exhibition match, but the real test will be in applying the underlying natural language processing and analytics technology in business and across industries. In this talk I will introduce the Jeopardy! grand challenge, present an overview of Watson and the DeepQA technology upon which Watson is built, and explore future applications of this technology.

Eric Brown
IBM Thomas J. Watson Research Center
ewbus.ibm.com

IP2

Challenges and Lessons Learned in Data-Intensive Computing

Making the best use of modern computational resources for data-intensive distributed applications requires expert knowledge of system architecture and low-level programming tools, or a productive high-level and high-performance programming framework. Even though the state-of-the-art high-level frameworks eases the programing burden, their performance are far from being optimal for many data-intensive applications. In this talk, we will present the lessons we learned developing such data-intensive applications on current hierarchical and hybrid architectures, and discuss the interplay of architecture, runtime system, algorithm, and application characteristics for developing high-performance applications.

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

IP3

Parallel N-body Algorithms on Heterogeneous Architectures

The fast multipole method (FMM) is an efficient algorithm for what is known as "N-body problems". I will present a scalable algorithm and a new implementation of the kernel-independent fast multipole method, in which both distributed memory parallelism (via MPI) and shared memory/SIMD parallelism (via GPU acceleration) are employed. I will conclude my talk by discussing the direct numerical simulation of blood flow in the Stokes regime using the FMM. I will describe simulations with 200 million red blood cells, an improvement of four orders of magnitude over previous results.

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

IP4

How to Avoid Communication in Linear Algebra and Beyond

The cost of moving data in an algorithm can surpass by several orders of magnitude the cost of performing arithmetics, and this gap has been steadily and exponentially growing over time. In this talk I will argue that this communication problem needs to be addressed by the numerical software community directly at the mathematical formulation and the algorithmic design level. This requires a paradigm shift in the way the numerical algorithms are devised, which now need to aim at keeping the number of communication instances to a minimum, while retaining their numerical efficiency. Communication avoiding algorithms provide such a novel perspective on designing algorithms that provably minimize communication in numerical linear algebra. The novel numerical schemes employed, the speedups obtained with respect to conventional algorithms, as well as their impact on applications in computational science will be also discussed.

Laura Grigori
INRIA
France
Laura.Grigori@inria.fr

IP5

Massive Parallelism for Quality Instead of Speed

For physical simulations and other technical applications, we can define "quality" as the reciprocal of the uncertainty in the answer. Uncertainty is caused by floating-point rounding errors as well as discretization error. We know that parallel processing can be applied to make spatial meshes finer and reduce discretization error; what is less obvious is that massive parallelism can be used to tighten bounds on the rounding error. Even "embarrassingly serial" problems with very few spatial variables and many time steps (like the three-body problem, say) can make good use of billions of processors to improve the quality of the calculation. We present a paradigm that fundamentally changes the approach to many classical problems in applied mathematics.

John Gustafson
Intel Labs USA
john.gustafson@intel.com.

IP6

State-of-the-Art Analysis and Perspectives of China HPC Development and Applications

In this talk, we first introduce the background of SAMSS China HPC TOP100 rank list. Then we give the total performance trend of China HPC TOP100 2011. Followed with this, the manufacturer and application area of 2011 China HPC TOP100 are analyzed briefly. Then the recent progess of high scalable application developemt onTianhe-1A are introduced. Based on public available historical data and TOP100 supercomputers peak performance data from 1993 to 2011 in China mainland, we predict the future performance trend of China HPC development.

Yunquan Zhang
Chinese Academy of Sciences, P.R. China
zyq@mail.rdcps.ac.cn; yunquan.zhang@gmail.com

IP7
Are there Exascale Algorithms?

Future computing system designs will be constrained by power density and total system energy, and will require new programming models and implementation strategies. Data movement in the memory system and interconnect will dominate computational costs, which leads to a new optimization criteria for algorithms. At the same time, the massive amount of parallel computation and performance uncertainty means the bulk synchronous computations will be increasingly inefficient. I will describe some of architectural trends and performance models for proving communication bounds, as well as emerging algorithmic ideas to avoid communication and synchnronization.

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

SP1
SIAG Junior Scientist Prize Award and Lecture: Performance-oriented Parallel Programming: Integrating Hardware, Middleware and Applications

Parallel programming is hard, optimizing parallel programming is even harder, and writing optimal parallel programs is nearly impossible. In this talk, we discuss how optimizing communication in parallel programs routinely requires to deal with low-level system details. We show portable abstractions that enable transparent optimizations but require advanced techniques in the lower layers. We conclude that scaling to larger machines demands an holistic approach to integrate hardware, middleware, and application software to develop performance-portable parallel programs.

Torsten Hoefler
University of Illinois
htor@illinois.edu

SP2
SIAG Career Scientist Prize Award and Lecture: It Seemed Like a Good Idea at the Time

Ive spent almost all of my career trying to make very fast, very parallel computing work for solving the most difficult problems science poses. In the beginning, Amdahl told us it wasn't going to work. We are still worrying about whether or not it will work at the exascale, but our success so far has shown him to be spectacularly wrong. In getting here, we have had to take on a long list of issues in machines, programming languages, and most importantly, algorithms. At different times, Ive felt it was one or the other of these areas that was the limiter to progress, so I switched focus, perhaps more frequently than was best. But in return, I got to work on many problems, do it in several settings, and work with many outstanding people. Ill survey some of that history in the talk.

Robert Schreiber
HP Labs
rob.schreiber@hp.com

CP1
Large-Scale Parallel Computations for High-Resolution Gravity Field Modelling

A determination of the Earths gravity field is formulated by the geodetic boundary value problem for the Laplace equation. To obtain its numerical solution we apply parallel versions of the boundary element method (BEM) and finite volume method (FVM) using the standard MPI subroutines. In both cases we use BiCGSTAB to solve large linear systems of equations. We present different ways how to parallelize a dense non-symmetric stiffness matrix for BEM and sparse one for FVM. Obtained numerical solutions are compared with the EGM2008 geopotential model.

Robert Cunderlik
Dept. of Mathematics, Slovak University of Technology
cunderli@svf.stuba.sk

Zuzana Faskova, Marek Macak
Dept. of Mathematics, Slovak University of Technology
Bratislava, Slovakia
faskova@math.sk, macak@math.sk

Karol Mikula
Department of Mathematics
Slovak University of Technology
karol.mikula@gmail.com

CP1
Enhanced Parallel Numerical Simulations of Carbon Capture and Storage

In the context of greenhouse gas emission into the atmosphere, CO2 capture and storage into geological formation has been considered recently as one of the mitigation option. Mathematical models are essential tools in addressing problems that arise in the context of CO2 storage in the deep subsurface. We enhance the parallel scalability of the Tough-MP (LBNL) code to consider large scale complex modeling and this talk will discuss the parallel strategies used in this context.

Fabrice Dupros
BRGM/STI
BP6009 - 45060 ORLEANS CEDEX 2
f.dupros@brgm.fr

Christophe Chiaberge
BRGM/WATER
BP6009 - 45060 ORLEANS CEDEX 2
c.chiaberge@brgm.fr

hajime yamamoto
aTaisei Corporation,344-1 Nase-cho Totsuka-ku,
Yokohama 245-
hajime.yamamoto@sakura.taisei.co.jp

Pascal Audigane
BRGM
p.audigane@brgm.fr

CP1
A Parallel High-Order Accurate Finite Element Nonlinear Stokes Ice-Sheet Model

A parallel finite element implementation on tetrahedral grids of the nonlinear three-dimensional full-Stokes model for the dynamics of ice-sheets is presented. Effective solvers

by using preconditioning techniques for the saddle-point system resulting from the discretization are discussed and implemented. We validated our finite element full-Stokes model through the use of well-known ice-sheet benchmark experiments, and the solver is shown to be efficient, robust, and scalable.

Wei Leng
Chinese Academy of Sciences
State Key Laboratory of Scientific and Engineering Computing
wleng@lsec.cc.ac.cn

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

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

Max Gunzburger
Florida State University
School for Computational Sciences
mgunzburger@fsu.edu

Stephen Price
Los Alamos National Laboratory
sprice@lanl.gov

CP1

Towards Petascale Simulation of Atmospheric Circulations with Soundproof Equations

We highlight progress with the development of a petascale implementation of a general-purpose high-resolution hydrodynamical simulation code EULAG. The new model-domain decomposition into a three dimensional processor array has been implemented to increase model performance and scalability. The results of scaling tests on IBM Bluegene/L and Cray XT4 show signi?cant improvement of the model ef?cacy compared to the original decomposition into a two dimensional processor array in the horizontal a standard in meteorological models.

Zbigniew P. Piotrowski
National Center for Atmospheric Research
piotrows@ucar.edu

Andrzej Wyszogrodzki, Piotr Smolarkiewicz
NCAR
MMM
andii@ucar.edu, smolar@ucar.edu

CP1

3D Seismic Modeling Via a Massively Parallel Structured Helmholtz Solver

We consider the modeling of seismic wave propagation on a rectangular domain via the discretization and solution of the inhomogeneous Helmholtz equation in 3D, by exploiting a parallel multifrontal sparse direct solver equipped with Hierarchically Semi-Separable (HSS) structure to reduce the computational complexity and storage. In particular, we are concerned with solving this equation on a large domain, for a large number of different forcing terms in the context of seismic problems in general, and modeling in particular. The solver is a parallel hybrid between multifrontal and HSS structure. The computational complexity is almost linear in the size of the Helmholtz matrix.

Shen Wang
Purdue University
wang273@purdue.edu

Jianlin Xia
Department of Mathematics
Purdue University
xiaj@math.purdue.edu

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

Maarten de Hoop
Center for Computational & Applied Mathematics
Purdue University
mdehoop@math.purdue.edu

CP1

A Parallel Fully Implicit Solver for Atmospheric Euler Flows

To study the mesoscale behavior of the atmosphere, a non-hydrostatic Euler model is considered. We investigate a fully implicit method whose time step size is not constrained by the CFL condition. In the solver a low dissipative method is used to discretize the equations and an inexact Newton-Krylov-Schwarz method is employed to solve the nonlinear system in each time step. Numerical experiments are provided to show that the method is scalable to thousands of processors.

Chao Yang
Software Institute, Chinese Academy of Sciences
P. R. China
yangchao@iscas.ac.cn

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

CP2

Linear-Time Verification of Correct Parallelization

Parallelization, especially the more-scalable variety using asynchronous communication, is difficult to verify in general. We propose a restricted set of primitives—a data distribution and remote invocation language—which is both sufficient to represent very aggressive, scalable asynchronous parallelism, and for which we have a linear-time algorithm (as a function of the number of primitives) which either verifies the lack of non-determinancy and non-termination or returns the offending primitive.

Christopher Anand, Michal Dobrogost, Wolfram Kahl, Jessica Pavlin
McMaster University
anandc@mcmaster.ca, michal.dobrogost@gmail.com, kahl@cas.mcmaster.ca, pavlinjl@mcmaster.ca

CP2

Efficient Optimization Algorithm for Empirical Performance Tuning of Parallel Scientific Applications

In empirical performance tuning of parallel scientific applications, finding high quality parameter configurations is a highly time consuming task. While typically parameters are tuned by hand, recent studies have shown that optimization algorithms can find high quality configurations quickly. We propose an efficient numerical optimization algorithm that exploits the characteristics of the search problem in tuning. We show the effectiveness of the proposed algorithm on a number of computational kernels from a large CFD code.

Prasanna Balaprakash, Stefan Wild
Argonne National Laboratory
pbalapra@mcs.anl.gov, wild@mcs.anl.gov

Paul D. Hovland
Argonne National Laboratory
MCS Division, 221/C236
hovland@mcs.anl.gov

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

CP2

A Hybrid-Agent Based Approach for Automated Fault Tolerance

This talk will present a hybrid approach for achieving automated fault tolerance in high-performance computing systems. In the approach, two forms of intelligence, namely intelligent computer cores and intelligent software agents are incorporated. When a core is about to fail which of these two forms of intelligence is to be given priority? A dependency heuristic used to determine the intelligence that needs to come into play and experimental studies to support the proposal will be presented.

Blesson Varghese
University of Reading, UK
blssnvarghese@gmail.com

Gerard McKee
Baze University, Nigeria
gerard.mckee@gmail.com

Vassil Alexandrov
Barcelona Supercomputing Centre, Spain
vassil.alexandrov@bsc.es

CP2

Using a Miniapp to Explore Computational Strategies for Multi-Core Machines

MiniGhost is a miniapp that mimics the MPI communication patterns of CTH, a large shock hydrodynamics code which utilizes the bulk synchronous programming model. We are using miniGhost to experiment with different methods of restructuring the code. This presentation will detail our experiments using OpenMP with MPI to adapt this programming model to machines with many cores per node such as our new ASC capability machine Cielo, a Cray XE6, in preperation for exascale machines.

Courtenay T. Vaughan, Richard Barrett
Sandia National Laboratories
ctvaugh@sandia.gov, rfbarre@sandia.gov

CP3

PHG: a Parallel Adaptive Finite Element Toolbox and Its Applications

PHG is a toolbox for developing parallel adaptive finite element programs. PHG deals with conforming tetrahedral meshes and uses bisection for adaptive local mesh refinement and MPI for message passing. In this lecture, the main algorithms in PHG will be introduced and some numerical results with up to 1 billion unknowns and using up to more than 12288 CPU cores are presented to demonstrate that PHG is robust and scalable.

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

Jie Cheng
LSEC, Chinese Academy of Sciences
chengjie@lsec.cc.ac.cn

Wei Leng
Chinese Academy of Sciences
State Key Laboratory of Scientific and Engineering Computing
wleng@lsec.cc.ac.cn

Linbo Zhang
LSEC, Chinese Academy of Sciences
zlb@lsec.cc.ac.cn

CP3

Generation of Unstructured Tetrahedral Meshes from Parallel Octree-Based Algorithm

An octree-based mesh generation method is proposed to build unstructured tetrahedral meshes from triangulated surface models. The technique recursively subdivides an octree in order to capture the features of the model. Problematic unbalanced octants are eliminated using 2:1 constraint. New refinement templates are proposed to eliminate the hanging vertices of the balanced octree resulting from the refinement process. The algorithm rely on massive parallel processing to build reasonable-quality, highly refined finite element meshes.

Igor T. Ghisi, José Camata
NACAD - High Performance Computing Center
COPPE/Federal University of Rio de Janeiro
igor.ghisi@gmail.com, camata@nacad.ufrj.br

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

CP3

Migrating the Uintah Computational Framework to CPU/GPU Architectures

Uintah is a parallel adaptive multi-physics framework used

to simulate fluid-structure interaction problems using both structured AMR, the ICE flow solver and the Material Point Method. Using hybrid parallelism (MPI/Pthreads), Uintah now scales to 196k cores on Jaguar. The focus of this work is to demonstrate and analyze the performance of Uintah on mixed CPU/GPU architectures through mapping key kernels onto GPUs, as in proposed next generation machines such as DOE's Titan.

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

Joseph Peterson
Dept of Chemistry
University of Utah
joseph.r.peterson@utah.edu

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

CP3

Experience With MPI and OpenMP in an Adaptive FEM Code

PHAML is a parallel *hp*-adaptive finite element code that was originally developed with MPI for message passing parallelism. Recently we have added OpenMP directives for shared memory parallelism on multi-core computers. In this talk we will report on our experiences using PHAML with MPI, OpenMP and hybrid MPI/OpenMP on multi-core computers and clusters of multi-core computers.

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

CP3

Efficient Parallel Matrix Assembly for Large Scale Fem

Finite element methods (FEM) for solving large scale problems are typically both time and memory consuming. While parts of FEM are perfectly suited for parallelisation, the system matrix assembly is still done more or less sequentially and, thus, prevents excellent speed-up and scalability values on massive parallel systems. To tackle this problem, we propose a bitonic sorting strategy allowing an efficient organisation of the communication. We will especially address advantages and shortcomings of this approach.

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

Rank Ernst
TUM, Faculty of Civil Engineering and Surveying
Chair for Computation in Engineering
rank@bv.tum.de

CP3

Streaming GPU-Based Triangular Surface Mesh Compression

Extremely large meshes cannot be stored in main memory, even if they are compressed or simplified. We propose a streaming, triangle-collapsing mesh compression and simplification scheme which leverages the memory and architecture of the GPU to compress extremely large triangular surface meshes. We compare our compressed and simplified meshes with the original mesh in terms of its visual quality, making it especially useful in computer graphics, where quality is based on the users perspective.

Dragos M. Nistor, Suzanne M. Shontz
The Pennsylvania State University
dmn5048@cse.psu.edu, shontz@cse.psu.edu

CP4

Parallel Two-Level Schwarz Methods for Shape Optimization of Steady Incompressible Flows

A fully coupled approach is introduced for shape optimization problems constrained by steady incompressible Navier-Stokes equations. The problem is discretized with a finite element method on unstructured moving meshes and then solved by a parallel two-level one-shot Lagrange-Newton-Krylov-Schwarz algorithm. As an application, we consider the design of an optimal artery bypass. Numerical experiments show that our algorithm performs well on supercomputers with hundreds of processors for problems with millions of unknowns.

Rongliang Chen
Department of Computer Science, University of Colorado
rongliang.chen@colorado.edu

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

CP4

Cache Optimization of a Non-Orthogonal Joint Diagonalization Method

The LUJ2D algorithm is a recently proposed numerical solution method for non-orthogonal joint diagonalization problems appearing in signal processing. The original LUJ2D algorithm attains low performance on modern microprocessors since it is dominated by cache ineffective operations. In this talk, we propose a cache efficient implementation of the LUJ2D algorithm which is dominated by matrix products.

Yusuke Hirota
Department of Computational Science,
Graduate School of System Informatics, Kobe University
hirota@stu.kobe-u.ac.jp

Yusaku Yamamoto
Kobe University
yamamoto@cs.kobe-u.ac.jp

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

CP4

A Domain Decomposed Newton Type Method for

Semiconductor Device Simulation

We discuss some recent work on domain decomposition methods for solving the nonlinear algebraic system arising from discretization of the drift-diffusion model for semiconductor device simulations. The system is highly nonlinear and becomes very difficult to solve especially when the maximum semiconductor doping concentration is large. We investigate a preconditioned Newton method, and report its robustness with respect to the physical parameters and its parallel performance on a computer with a large number of processors.

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

Xuefeng Li
Loyola University New Orleans
xli200005@yahoo.com

CP4
Scaling a Numerical Library to Emerging Compute Architectures

We present our experiences in extending the numerical library PETSc to scale with the computing resources of emerging petascale architectures. The optimizations are focused on structured grid applications, linear solvers, and preconditioning. The extended data structures facilitate aggressive vectorization, efficient memory accesses, compiler optimizations, and SIMD parallelism. The revised bottleneck-kernels make extensive use of intrinsics, multicores, and GPUs. We demonstrate the performance of our data structures and algorithms using a reservoir simulation software (PFLOTRAN).

Chekuri S. Choudary
RNET Technologies, Inc.
Dayton, OH
cchoudary@rnet-tech.com

Jeswin Godwin, Deepan Karthik
Ohio State University
jeswingodwin@gmail.com, dpnkarthik@gmail.com

Daniel Lowell, Boyana Norris
Argonne National Laboratory
dlowell@mcs.anl.gov, norris@mcs.anl.gov

Gerald Sabin
RNET Technologies, Inc.
Dayton, OH
gsabin@rnet-tech.com

P. Sadayappan
Ohio State University
sadayappan.1@osu.edu

Sravya Tirukkovalur
RNET Technologies, Inc.
Dayton, OH
stirukkovalur@rnet-tech.com

CP4
Scalability and Performances of the FEAST Eigenvalue Algorithm and Solver

The list of capabilities of the FEAST algorithm and solver recently proposed by the author, uniquely combines accuracy, robustness, high-performance and (linear) parallel scalability. An efficient parallel implementation for FEAST is here addressed at three different levels. Consequently, if enough parallel computing power is available at hand, the main computational cost of FEAST for solving the eigenvalue problem (and even for capturing millions of eigenpairs), can be ultimately reduced to solving only one linear system. Scalability results for the electronic structure problem will be presented.

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

CP4
Parallel Polynomial Evaluation

Fast polynomial evaluation (and its derivatives) is a critical component in Newton's method and in continuation methods for solving polynomial systems. In the path trackers of PHCpack, we have integrated the software library QD-2.3.9 for double double and quad double complex arithmetic to increase the accuracy of the computations. To offset the extra cost of the multiprecision arithmetic we use multiple cores. In particular, we calculate quality up factors, answering the question: if we have c cores available, how many extra digits d can we use in the working precision keeping the wall clock time fixed?

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

Genady Yoffe
Dept. of Mathematics, Statistics, and CS
University of Illinois, Chicago
gyoffe2@uic.edu

CP5
Exploiting Low-Rank Structure in Computing Matrix Powers with Applications to Preconditioning

Parallel Communication-avoiding Krylov Subspace Methods (CA-KSMs) use the matrix powers kernel to exploit sparsity of A, performing s SpMVs for the latency cost of 1 SpMV. We extend the parallel matrix powers kernel to handle A with dense but low-rank components. Our approach reduces parallel latency by a factor of $O(s)$ without asymptotically increasing computation cost. We analyze convergence for several polynomial bases, and show our approach enables preconditioned CA-KSMs with hierarchical semiseparable preconditioners.

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

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

Ming Gu

University of California, Berkeley
Mathematics Department
mgu@math.berkeley.edu

CP5
Parareal for Stochastic Parabolic Equation

We investigate a parallel domain decomposition preconditioning technique for solving a stochastic parabolic equation. The stochastic equation is decoupled into a sequence of uncoupled deterministic equations by KL expansion and double orthogonal polynomials. Based on parareal algorithm, we combines temporal and spacial discretization into one large system, which is then solved by a preconditioned iterative method. Additive Schwarz preconditioner and recycling of selective Krylov subspace are used to construct the overall solution algorithm. We report some numerical experiments obtained on a parallel supercomputer with a large number of processors.

Cui Cong
Department of Mathematics
University of Colorado at Boulder
congcui@gmail.com

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

CP5
Parareal with Adjoints

Parareal method solves time dependent ordinary differential equations parallel in time. It involves splitting the time interval into smaller subsets and solving the ODEs independently on them. We present a technique that remodels the problem of finding the intermediate solutions at each of these intervals as a non-linear optimization problem. The optimization problem is solved using gradient based methods, where the gradients are computed using adjoint models.

Vishwas Hebbur Venkata Subba Rao
Virginia Polytechnic Institute and State University
visrao@vt.edu

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

CP5
Estimating the Diagonal of Matrix Inverses in Parallel

Approximating the diagonal of the matrix inverse is of importance in areas ranging from engineering to economics. Recent algorithms on the subject have combined statistical estimation, conjugate gradients, mixed-precision arithmetic and parallelism. We present novel algorithms based on projections that aggressively pursue information sharing and non-trivial parallelism and demonstrate their effectiveness on current HPC systems.

Efstratios Gallopoulos
Department of Computer Engineering & Informatics
University of Patras, Greece
stratis@ceid.upatras.gr

Vasilis Kalantzis
CEID, School of Engineering
University of Patras, Greece
kalantzis@ceid.upatras.gr

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

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

CP5
Solving Sparse Symmetric Rank-Deficient Systems of Linear Algebraic Equations on SMP Platforms

The talk is devoted to a highly-parallel method for solving

rank-deficient symmetric sparse linear least-square problems on SMP computers. The proposed method is based on the truncated SVD approach. The algorithm computes an orthonormal basis of the kernel using FEAST Eigenvalue Solver first and then solves an augmented system of linear equations. Intel ®MKL PARDISO is used to solve resulting well-conditioned systems. Performance comparison with existing approaches is provided for the latest Intel ®multi-core processors.

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

CP5

Using Overlapping and Filtering Techniques for Highly Parallel Preconditioners

In this talk we discuss a highly parallel preconditioner for solving large sparse linear systems of equations. The preconditioner is based on nested dissection ordering and uses filtering techniques to alleviate the effect of low frequency modes on the convergence of the iterative method. We also discuss overlapping techniques to further increase the numerical efficiency of the preconditioner.

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

Laura Grigori
INRIA
France
Laura.Grigori@inria.fr

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

fezzani riadh
INRIA Saclay, France
rfezzani@lri.fr

CP6

An Analytical Framework for the Performance Prediction of Work-Stealing Dynamic Load Balancing Schemes for Ensemble Simulations

Ensemble simulations are employed to estimate the statistics of possible future states of a system, and are widely used in important applications such as climate change and biological modeling. This paper presents a new probabilistic framework to analyze the performance of dynamic load balancing algorithms (most-dividing, all-redistribution, random-polling, and neighbor-redistribution) when applied to large ensembles of stochastic biochemical simulations. The analysis can be applied to any ensemble simulations where the individual compute times vary considerably.

Tae-Hyuk (ted) Ahn
Virginia Tech
thahn@vt.edu

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

CP6

Scheduling and Mapping of Communicating Multi-Processor Tasks

This talk discusses the parallel programming model of communicating multi-processor tasks (CM-tasks), which allows both task-internal communication as well as communication between concurrently executed tasks at arbitrary points of their execution. A corresponding scheduling algorithm is proposed as well as a transformation toolset supporting the scheduling. An experimental evaluation of several application programs shows that using the CM-task model may lead to significant performance improvements compared to other parallel execution schemes.

Gudula Rünger, Joerg Duemmler
Chemnitz University of Technology
ruenger@informatik.tu-chemnitz.de,
joerg.duemmler@informatik.tu-chemnitz.de

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

CP6

Task Scheduling for Multi-Threaded/mpi Hybrid Parallelization

A concept which has arisen in recent years, is the notion of automatic decomposition of a system into its constituent data dependencies. These dependencies can be stored as a directed acyclic 'task graph', allowing us to procedurally determine which components of the system may be executed in parallel. We will show that for GPU/CPU HPC environments, under some simplifying assumptions, it is possible to obtain excellent performance without requiring hand tuning or user intervention.

Braden D. Robison
University of Utah
devin.robison@utah.edu

CP6

Considerations on Parallel Graph Coloring Algorithms

Graph coloring is a combinatorial optimization problem that appears in distributed computing to identify independent operations. Building on a distributed-memory graph coloring framework, we investigate two main directions: ordering of the vertices to take into account the structure of the graph; and an iterative improvement technique that use the information contained in a coloring to build another one. To further improve the performance on modern hierarchical machines, we discuss the use of hybrid shared-memory/distributed-memory implementations.

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

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

CP6

Evaluating Application-Level *VS.* Middleware-Level Load Balancing for Parallel Adaptive Computation

We investigate the utility of using a middleware-based load balancer to guide a simple parallel adaptive computation. Normally, such computations require application-level load balancing (managing partitioning, load balanc-

ing, and data migration). We use an actor-theater model of computation, where each actor is responsible for a subset of the computation that changes due to adaptivity. The Internet Operating System is used to manage balancing only at the actor level and remove some burden from application programmers.

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

Carlos Varela, Qingling Wang
Department of Computer Science
Rensselaer Polytechnic Institute
cvarela@cs.rpi.edu, wangq9@cs.rpi.edu

CP6

Reaction Network Partitioning and Load Balancing for Parallel Computation of Chemistry in Combustion Applications

Simulating combustion of complex hydrocarbon fuels requires the evaluation of chemical reaction terms involving thousands of species and reaction steps. The evaluation of chemical reaction terms is computationally intensive, time consuming and limits the complexity of chemical models that can be used in combustion simulations. This paper presents an approach for partitioning the reaction network that allows the computation to be parallelized into multiple concurrent computations targeted at the modern multicore and many core processors. Several partitioning strategies are explored using software such as Metis and PaToh and are evaluated based on efficiency, load balance and memory footprint. The parallelization based on reaction network partitioning is demonstrated using chemical reaction networks for gasoline, diesel and biodiesel surrogates.

Wei Wang
Joint Institute for Computational Sciences,
University of Tennessee
wangw2@ornl.gov

Ramanan Sankaran
National Centre for Computational Sciences
Oakridge National Laboratories
sankaranr@ornl.gov

Ray W. Grout
National Renewable Energy Laboratory
ray.grout@nrel.gov

CP7

Efficient Particle Swarm Methods for Optimal Space-Filling Designs on Gpu

Space-filling designs are indispensable to practical experiments and they can be generalized by solving the corresponding discrete optimization problems. Such problems are challenging as the feasible points are huge and the objective function values are irregularly sensitive. We propose variants of novel particle swarm optimization algorithms, combining with the computational power of graphic process units to find optimal designs efficiently and accurately. Numerical results are presented to demonstrate the advantages of the proposed methods.

Ray-Bing Chen
Department of Statistics
National Cheng Kung University
rbchen@stat.ncku.edu.tw

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

Dai-Ni Hsieh, Yen-Wen Shu
Department of Mathematics
National Taiwan University
dn.hsieh@gmail.com, b94201002@gmail.com

Ying Hung
Department of Statistics and Biostatistics
Rutgers University
yhung@stat.rutgers.edu

CP7

Parallel Global Optimization of Functions with Multiple Local Minima

Three optimization problems with high dimensionality and many local minima are investigated with three di?erent optimization algorithms: DIRECT, simulated annealing, Spalls SPSA algorithm, and QNSTOP, a new algorithm developed at Indiana University. Optimization is considered both globally, using a latin hypercube to aid in exploration of the feasible set, and locally, centered on the best result obtained by the global optimization. Various parallel implementations of these algorithms are explored.

David R. Easterling
Virginia Tech
dreast@vt.edu

Layne T. Watson
Virginia Polytechnic Institute and State University
Departments of Computer Science and Mathematics
ltwatson@computer.org

Brent Castle
Indiana University
bcastle@cs.indiana.edu

Michael W. Trosset
Department of Statistics, Indiana
mtrosset@indiana.edu

Michael Madigan
Virginia Tech
mimadiga@vt.edu

CP7

Massively Parallel Multi-Objective Optimization and Application

Particle accelerators are invaluable tools for research in basic and applied sciences. The successful design, commissioning, and operation of accelerator facilities is non trivial. We implemented a framework for general simulation-based multi-optimization methods automating the investigation of optimal sets of machine parameters. In order to solve the emerging, huge problems we propose a massively-parallel topology-aware scheduling approach. We employ the framework to identify optimal parameters of the 250

MeV injector currently under construction at PSI.

Yves Ineichen
ETHZ
PSI
yves.ineichen@psi.ch

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

Peter Arbenz
ETH Zurich
Chair of Computational Science
arbenz@inf.ethz.ch

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

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

CP7

Parallel Metaheuristic Optimization Framework for Population-Based Search Algorithms with Environmental and Energy Applications

Optimization problems involve finding an optimal set of parameters to minimize or maximize a given objective function subject to constraints. Population-based heuristic search methods such as evolutionary algorithms and particle swarm optimization are widely used for solving such problems especially when classical optimization techniques are inadequate. This research presents a scalable multi-population heuristic optimization framework for deployment on emergent parallel architectures and demonstrates it for applications in the environmental and energy domains.

Sarat Sreepathi, Joseph DeCarolis, Kumar G. Mahinthakumar
North Carolina State University
sarat_s@ncsu.edu, jdecarolis@ncsu.edu, gmkumar@ncsu.edu

CP8

Fast K-Selection Algorithms for Graphical Processing Units

We discuss two CUDA-based algorithms, GPUbucketSelect and GPUradixSelect, for finding the kth-largest value in a list on a GPU. Both of these algorithms are faster than sorting the list on the GPU with the highly optimized radix sorting algorithm included in the Thrust Library as `thrust::sort`. The acceleration grows dramatically with the vector length; for example, with lists of length 2^{29}, we observe a 19-times speed-up over `thrust::sort`.

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

Tolu Alabi
Department of Computer Science
Grinnell College
alabitol@grinnell.edu

Bradley Gordon
Grinnell College
gordonbr@grinnell.edu

Russel Steinbach
Department of Computer Science
Grinnell College
steinbac@grinnell.edu

CP8

Rank Computations with Parallel Random Surfers

We show that functional rankings can be expressed as random surfer processes with step-dependent transition probabilities. We call this approach multidamping and it is different from PageRank where transition probabilities are fixed. It follows functional rankings allow massively parallel, Monte Carlo type implementations. We experiment with embarrassingly parallel cluster executions, employing shared memory parallelism at the compute node - with graph sharing surfers as threads accumulating node visit frequencies - and comment on their probabilistic properties.

Giorgos Kollias, Efstratios Gallopoulos
Department of Computer Engineering & Informatics
University of Patras, Greece
gkollias@cs.purdue.edu, stratis@ceid.upatras.gr

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

CP8

Scalable SPARQL Querying with Compressed Bitmap Indexes

We present the design and implementation of FastBit-RDF, an RDF data management system for efficiently answering SPARQL pattern-matching queries on large semantic data sets. We describe a parallel index construction and query answering methodology that utilizes compressed bitmap indexes. Our approach is up to an order of magnitude faster the state-of-the-art system RDF-3X, for a variety of SPARQL queries on gigascale RDF datasets.

Kamesh Madduri
Lawrence Berkeley National Lab
madduri@cse.psu.edu

CP8

Parallel Implementation of Pca and Isomap

Manifold Learning as a way of visualizing high-dimensional, unordered data is an emerging technique. While enormity of the datasets is not an uncommon concern, example techniques like Principal Component Analysis (PCA) and Isomap require the solving of a few largest eigenvalues of dense matrices. Hence, owing to the memory constraints, most of the sequential implementations of these methods fail to scale to larger datasets. This formed the impetus for us to develop parallel implementation of SETDiR (Scalable Extensible Toolkit for Dimensionality Reduction). This framework is applied to materials science

data from Atom Probe Tomograph to extract the information helpful to understand the mechanism of Atom Probe Tomograph and to analyze and extract various physical and topological features describing the data.

Sai Kiranmayee Samudrala
Department of Mechanical Engineering, Iowa State University
Ames, Iowa -50011
saikiran@iastate.edu

Jaroslaw Zola, Srinivas Aluru, Baskar Ganapathysubramanian
Iowa State University
zola@iastate.edu, aluru@iastate.edu, baskarg@iastate.edu

CP8

Breadth First Search Implementation on the Convey HC-1ex

The Graph500 benchmark was ported to a Convey HC-1ex, a hybrid computer system with an Intel host processor and a coprocessor incorporating four reprogrammable Xilinx FPGAs. The system incorporates a unique memory system designed to sustain high bandwidth for random memory accesses. The BFS kernel was implemented as a set of parallel state machines that traverse the graph in coprocessor memory. The resulting implementation executes the benchmark at a rate of over 1 billion TEPS.

Kevin Wadleigh
Convey Computer Corporation
kwadleigh@conveycomputer.com

CP9

Parallelization of the Fast Multipole Method on Heterogeneous Architectures in Electromagnetics

The Fast Multipole Method (FMM) can dramatically speed up the solving of electromagnetic scattering problems, the corresponding complexity being $\mathcal{O}(N \log N)$ instead of $\mathcal{O}(N^2)$ (N being the number of unknowns). Unfortunately in electromagnetism, the parallel FMM is poorly scalable. To overcome this limit, our hybrid (CPUs and GPUs) implementation is based on independent tasks scheduled by StarPU, a runtime system for heterogeneous multicore architectures. We report preliminary performance results that demonstrate the relevance of our approach.

Cyril Bordage
cea.fr
cyril.bordage@gmail.com

CP9

High Performance Solution of Sparse and Dense Linear Systems with Application to Large 3D Electromagnetic Problems on a Petascale Computer

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

David Goudin, Jean-Jacques Pesque, Agnes Pujols, Muriel Sesques, Bruno Stupfel
CEA/CESTA
david.goudin@cea.fr, jean.pesque@cea.fr, agnes.pujols@cea.fr, muriel.sesques@cea.fr, bruno.stupfel@cea.fr

CP9

An Eigenvalue Solver for Extended Magnetohydrodynamics

Tokamak plasmas owe their success to the existence of a symmetry angle which provides excellent confinement of the fast moving particles. Experimentally, the appearance of disruptive, symmetry-breaking perturbations coincides with the stability boundary of the extended magnetohydrodynamic (MHD) model. We present the development of an eigenvalue solver based on the NIMROD code, a massively parallel DOE MHD simulation tool. The resulting generalized eigenproblem is solved using SLEPc, an extension of the PETSc framework.

Ben Jamroz
Tech-X
jamroz@txcorp.com

Scott Kruger
Tech-X Corporation
kruger@txcorp.com

CP9

A Highly Scalable FMM for Particle Simulations

In this talk we present our error-controlled and runtime-minimized FMM implementation for long-range interactions in particle simulations. The current code scales up to 300k BG/P cores and can handle more than three trillion particles in just over eleven minutes for an expansion up to quadrupoles. The code employs a one-sided, non-blocking parallelization scheme with a small communication overhead. We will show results and timings of low- and high-precision computations for homogeneous and clustered particle systems.

Ivo Kabadshow
Research Centre Juelich
i.kabadshow@fz-juelich.de

Holger Dachsel
Juelich Supercomputing Centre
Research Centre Juelich
h.dachsel@fz-juelich.de

CP9

A Vortex Tree Code for Extreme Scale Fluid Simulations

We present our massively parallel Barnes-Hut tree code for vortex particle methods, which is based on the recent hybrid MPI/pthread tree code PEPC. We demonstrate its capabilities on 65'536 cores of an IBM BlueGene/P system, studying vortex ring dynamics with millions of particles. Using multipole expansions of algebraic smoothing kernels and efficient remeshing by parallel sorting, this code is a versatile tool for large-scale simulations of vortex-driven

fluid flows.

Robert Speck, Rolf Krause
Institute of Computational Science
University of Lugano
robert.speck@usi.ch, rolf.krause@usi.ch

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

CP9

PEPC A Versatile Highly Scalable Parallel Barnes-Hut Treecode

We present an innovative hybrid MPI/Pthreads parallelization strategy for the tree traversal and force calculation in a parallel Barnes-Hut treecode. It allows for a native overlap of communication and computation, shows excellent scaling across all 294,912 processors of the BlueGene/P system Jugene and is prepared for upcoming HPC architectures. As an example application for our multi-disciplinary implementation PEPC, we show studies of transport properties in strongly coupled Coulomb systems consisting of up to 2,048,000,000 particles.

Mathias Winkel
Jülich Supercomputing Centre
Forschungszentrum Jülich
m.winkel@fz-juelich.de

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

CP10

Using Gpus in the Parallelization of a Simple Entropy-Based Moment Model

In transport and kinetic theory, entropy-based methods are used to derive closures for moment models in order to reduce the complexity of the problem. Such algorithms require an optimization problem to be solved numerically for every grid point in a space-time mesh. The work presented will highlight the challenges of using a GPU to parallelize the optimization solver and present a comparison of the scaling results between the GPU parallelized code and the serial code.

Charles K. Garrett
University of Texas - Arlington
charles.garrett@mavs.uta.edu

CP10

2D Gpu-Based Travel Time Computation

Commonly GPU-based seismic imaging tools achieve high performance, in particular if they are based on embarrassing parallel numerical schemes. Unfortunately, the ones-based on accurate travel time computation are data-driven, then not suitable for GPU-parallelism. The porting of this kind of popular (among geophysicist) algorithms to GPU is relevant even acknowledging that the task is difficult. Our implementation achieves up 4x speed up, and almost linear scalability for multi-GPU executions. Thus, showing feasibility and efficiency.

Mauricio Araya
BSC
araya.mauricio@repsol.com

Max Grossman
Rice University, Houston, Texas
jmaxg3@gmail.com

CP10

Parallel Sieve Processing on Vector Processor and GPU

The RSA cryptography code is based on the difficulty of factorizing long-digit composite numbers. It consists of "sieve processing", processing of 0-1 matrices, and computation of algebraic square roots. Sieve processing which is the most time consuming, is tuned.
The performance on one node of the ES2 vector processor is approximately 800 times faster than that on one core of Intel Core2, 2.3 GHz. Using GPU NVIDIA GTX580 is approximately 60 times faster than that on the same PC.

Yasunori Ushiro, Yoshinari Fukui
Earth Simulator Center
ushiro@aoni.waseda.jp, yoshinari.fukui@ieee.org

Hidehiko Hasegawa
Inst. of Library and Infomation Science
University of Tsukuba
hasegawa@slis.tsukuba.ac.jp

CP10

Gpu Accelerated Statistical Analysis of Cryptographic Functions

To be considered secure, a cryptographic function, such as a block cipher or one-way hash function, must pass statistical randomness tests. Generating the tests' input sequences requires a large number (e.g., 10^{12}) of function evaluations. The tests themselves also require a large amount of computation. However, the evaluations and tests can be executed in a massively parallel fashion and are ideally suited to run on a graphics processing unit (GPU) accelerator. This talk describes a GPU parallel statistical testing framework for cryptographic functions and gives the statistical test results for the five SHA-3 finalist hash algorithms.

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

CP10

A Parallel Monte Carlo Algorithm for Modeling Dense Stellar Systems on Hybrid Architectures

Investigating the interactions between the stars within a dense stellar system is a problem of fundamental importance in theoretical astrophysics. While simulating realistic globular clusters containing about 10^6 stars is computationally intensive, galactic nuclear star clusters with about 10^9 stars are practically unachievable. In this talk, we will present a parallel version of a Monte Carlo algorithm for simulating the evolution of such very large stellar clusters on hybrid hardware architectures. We will discuss

the numerical methods used in the algorithm, their time complexities, the strategies we used to minimize communication and how to efficiently use the computing power of GPUs in developing the parallel version. We will also talk about the performance of our implementation on a GPU cluster for various physical configurations, and discuss the scalability of the algorithm and how optimally it uses the available computational resources.

Bharath Pattabiraman
Dept. of Electrical Engineering and Computer Science
Northwestern University
bharath650@gmail.com

Stefan Umbreit
Dept. of Physics and Astronomy,
Northwestern University, Evanston, USA
s-umbreit@northwestern.edu

Wei-Keng Liao
Department of Electrical Engineering and Computer Science
Northwestern University
wkliao@ece.northwestern.edu

Frederic Rasio, Vassiliki Kalogera
Dept. of Physics and Astronomy,
Northwestern University, Evanston, USA
rasio@northwestern.edu, vicky@northwestern.edu

Gokhan Memik, Alok Choudhary
Dept. of Electrical Engineering and Computer Science
Northwestern University, Evanston, USA
memik@eecs.northwestern.edu,
choudhar@eecs.northwestern.edu

CP10

Parallel Numerical Methods for Solving 1+2 Dimensional Nonlinear Schrodinger Type Equations

The nonlinear Schrodinger equation(NLS) is of tremendous interest in both theory and applications. Various regimes of pulse propagation in optical fibers are modeled by some form of the NLS equation. In this paper we introduce parallel split-step Fourier and ADI finite difference methods for numerical simulations of the 1+ 2 dimensional NLS type equations. The parallel methods are implemented on multiprocessor systems and will be implemented on the newly acquired GPUs. Our numerical results have shown that these methods give good results and considerable speedup. The results of the GPU implementation will be reported at the conference.

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

CP11

Parallel Implementation of SSA to Bio-chemical Reaction Network

The application of stochastic simulation algorithms (SSA) in the recent past led to new developments in the dynamics of system biology. One particular example where molecular-level fluctuations play a key role is genetic networks . We applied stochastic algorithm to L1 Gene transcription model. We are currently working on the parallel implementation of the stochastic algorithm to this model to reduce the run-times in comparison to simulations run on a single processor. In this talk, we will present the results and the time comparisons.

Vani Cheruvu
Department of Scientific Computing
Florida State University, Tallahassee, FL 32306
vani.cheruvu@utoledo.edu

Devarapu Anilkumar, Seyed Roosta
Dept. of Mathematics and Computer Science
Albany State University, Albany, GA
anilkumar.devarapu@asurams.edu,
seyed.roosta@asurams.edu

CP11

Substrate Sequestration in a Multisite Phosphorylation System Produces Bi-Stability: A Numerical Approach

Cascades of coupled phosphorylation/dephosphorylation cycles,such as mitogen-activated protein kinase(MAPK) pathways, integrate external stimuli and propagate signals from plasma membrane to nucleus. A typical, three-stage cascade consists of MAPK, MAP2K and MAP3K.MAP2K is ac- tivated by MAP3K at cell membrane by an addition of a phosphate group and consequently the interior protein MAPK in the cell (near nucleolus membrane) is phosphorylated by activated MAP2K on two conserved threonine and tyrosine residues. Activated MAPK then sends some signal in nu- cleus to take the stand for the external signal. Various phosphatases undo these phosphorylations. Here we considered various mathematical models to explore the system, which involves multisite phosphorylation system with regulated substrate sequestration. Our models demonstrate that sub- strate sequestration in combination of multisite phosphorylation can produce robust switch-like and bi-stability.

Kanadpriya Basu
University of South Carolina
Department of Mathematics
basuk@mailbox.sc.edu

CP11

Scaling for Accuracy and Performance in Realistic Heart Models

Realistic simulation of the electrophysiology of the human heart requires complex high-resolution models and massively parallel computations. We present a novel hybrid OpenMP-MPI code for computational electrocardiology, capable of simulating systems with billions of mesh nodes and with excellent scalability up to 16k cores. We present the results of our thorough performance analysis of the recent hybrid implementation for explicit and implicit-explicit time integration, quantifying the advantages of hybrid execution.

Dorian Krause, Mark Potse, Thomas Dickopf, Rolf Krause
Institute of Computational Science
University of Lugano
dorian.krause@usi.ch, mark@potse.nl,
thomas.dickopf@usi.ch, rolf.krause@usi.ch

Angelo Auricchio
Fondazione Cardiocentro Ticino
angelo.auricchio@cardiocentro.org

Frits Prinzen
Department of Physiology
University of Maastricht
frits.prinzen@fys.unimaas.nl

CP11
Parallel Computational Model of Hiv Infection

We created a computational model of HIV infection to evaluate mechanisms of viral evasion from immune responses. Our efficient model uses OpenMPI and ANSI-C for a highly scalable and portable simulation. The implementation scales well across distributed memory systems, especially as the complexity of the simulation increases. Parallelism offered through a message-passing-interface and our high performance cluster greatly increased the rate that simulation data is generated. Our results may provide insights for HIV vaccine development.

Sergey Lapin
Washington State University
Department of Mathematics
slapin@math.wsu.edu

Elissa Schwartz
Washington State University
ejs@wsu.edu

L G. de Pillis
Harvey Mudd College
depillis@hmc.edu

William Bonner
Washington State University
bbonner@wsu.edu

CP11
Simulation of Electromechanics in the Heart

Cardiovascular disease is the leading cause of death in America. Computer simulation of complicated dynamics of the heart could provide valuable quantitative guidance for diagnosis and treatment of heart problems. In this poster, we present an integrated numerical model which encompasses the interaction of cardiac electrophysiology, electromechanics, and mechanoelectrical feedback. The model is solved by finite element method on a Linux cluster and the Cray XT5 supercomputer, kraken.

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

Xiaopeng Zhao, Henian Xia
University of Tennessee
xzhao9@utk.edu, xiahenian@gmail.com

CP11
Parallel Scalable Domain Decomposition Method for Blood Flow in Compliant Arteries in 3D

We develop a parallel scalable domain decomposition method for fluid-structure interaction simulation of blood flow in arteries using a monolithically coupled system consisting of incompressible Navier-Stokes equations and a linear elastic equation. We focus on the parallelization and scalability of the solution algorithm, and study the performance of the overlapping Schwarz preconditioner for solving the Jacobian systems as well as the efficiency of the Newton-Krylov-Schwarz algorithm for solving the algebraic nonlinear system on a large scale supercomputer.

Yuqi Wu
Dept. of Applied Math., University of Colorado
yuqi.wu@colorado.edu

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

CP12
Block Distributed Schur Complement Preconditioners for CFD Computations on Many-Core Systems

At the German Aerospace Center, the parallel simulation systems TAU and TRACE have been developed for the aerodynamic design of aircrafts or turbines for jet engines. For the parallel iterative solution of large, sparse equation systems within both CFD solvers, block-local preconditioners are compared with global block Distributed Schur Complement preconditioning methods for real or complex matrix problems. Performance results of preconditioned FGMRes algorithms are presented for TAU and TRACE problems on many-core systems.

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

CP12
Distributed Communication Strategies for Parallel Computational Fluid Dynamics

Nowadays, fluid flow simulations have to run in parallel in order to obtain physically significant results for large computational domains. We present a massive parallel implementation for solving the isothermal Navier-Stokes equations using a finite volume based discretised code. A specially designed data structure and addressing scheme will be presented in order to handle the inter-process communications in an efficient way. A scalability study will be performed using the Shaheen BlueGene/P installed at KAUST.

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

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

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

CP12
A Hybrid Algorithm for a Finite Element Solver

Based on Two-Level Parallelism

The current trend in hardware development is to increase the CPU core count and socket count per node on compute clusters and supercomputers. This increased parallelism comes at the cost of a more heterogeneous memory profile. It becomes necessary to develop parallelization schemes that take into account the hardware topology of such platforms. The goal of this work is to study the behaviour of a hybrid parallel implementation of an Additive-Schwarz preconditioner when solving the linear systems of equations associated with computational fluid dynamics problems. The preconditioning scheme uses a parallel subdomain solver on each subdomain. As such, it becomes possible to use a high number of MPI processes, while lowering the total number of domains. In addition to being able to control the iteration count required to solve the problem, this strategy comes with no loss of parallelism in other phases of the simulation, like the assembly of the linear system of equations.

Radu Popescu
Ecole Polytechnique Federale Lausanne
Chair of Modelling and Scientific Computing (CMCS)
radu.popescu@epfl.ch

Sivasankaran Rajamanickam
Sandia National Laboratories
srajama@sandia.gov

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

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

Michael A. Heroux
Sandia National Laboratories
Albuquerque, New Mexico, USA
maherou@sandia.gov

CP12

An Efficient Numerical Simulator for Two-Phase Flows on a GPU and Applications to Oil Reservoirs and CO2 sequestration

We consider the numerical simulation of two-phase flow problems in porous media. By ignoring the capillary pressure, two-phase flow problems are modeled by a coupled system of hyperbolic and elliptic problems. We consider heterogeneity in permeability and porosity. To solve the hyperbolic problem, we extend the Kurganov-Tadmor central scheme to handle variable porosity. Recent investigation on the use of a GPU for solving the hyperbolic problem associated with two-phase flows indicates a speed-up up to 60 in comparison to the serial calculation on a CPU [Pereira and Rahunanthan, VECPAR 2010]. In this talk, we describe a GPU parallelization of two-phase flow problems in three space dimensions and discuss the speed-up of the approach for a heterogeneous CPU-GPU system. The numerical experiments are discussed for applications in oil reservoirs and CO2 sequestration.

Felipe Pereira
Center for Fundamentals of Subsurface Flow
University of Wyoming
lpereira@uwyo.edu

Arunasalam Rahunanthan
Center for Fundamentals of Subsurface Flow
University of Wyoming, Laramie, WY 82071
rahu@uwyo.edu

CP12

Finite Element Analysis of a Load Carrying Mechanism in Tilted Pad Slider Bearing Lubrication

Finite Element Solutions are obtained for thermohydrodynamic lubrication of tilted pad slider bearings with heat conduction to the stationary pad and moving slider. The fluid film continuity, momentum and energy equations are solved in a coupled fashion along with conduction equations for the pad and slider with variable density and viscosity order to obtain the parameters responsible for load generation and also for a reduction in frictional drag. The special version of the conventional finite element methods is modified by adding a discontinuous terms order to preclude the spurious node-to-node numerical oscillations in the scheme and obtained a solution which is more suitable for complex geometry like slider bearing lubrication and also an attempt is made for parallel processing of the numerical scheme on 8-node PRAM machine.

Pentyala Srinivasa Rao
Indian School of Mines
pentyalasrinivasa@gmail.com

CP12

Comparing the Performance of Preconditioners and Linear Iterative Solvers for Unsaturated Flow Problems

This presentation will reveal results of testing the performance of preconditioners and linear iterative solvers for unsaturated flow problems using the parallel version of the PETSc library. Because the material properties of unsaturated soil vary several orders of magnitude as a function of pressure head, a linear system of ill-conditioned equations must be solved at each nonlinear iteration of a finite element solution. Test problems used range from small two-dimensional problems to large three-dimensional examples.

Fred T. Tracy, Maureen Corcoran
Engineer Research and Development Center
Waterways Experiment Statiuon
fred.t.tracy@usace.army.mil,
maureen.k.corcoran@usace.army.mil

CP12

Paladins: a Scalable Solver for the Navier-Stokes Equations.

PALADINS (Parallel ALgebraic time ADaptive solver for the Incompressible Navier-Stokes equations) is a c++ implementation of a class of high-order splitting schemes that was introduced by Gervasio, Saleri and Veneziani, and that features a hierarchical structure prone to time adaptivity. It is based on Trilinos and the finite element library LifeV (CMCS in Lousanne, MOX in Milan, Emory in Atlanta). In this talk we will present scalability results and an application to computational hemo-dynamics.

Umberto E. Villa

Dept. of Mathematics and Computer Science
Emory University
uvilla@emory.edu

Alessandro Veneziani
MathCS, Emory University, Atlanta, GA
ale@mathcs.emory.edu

CP13

Performance Parameters for Parallel Algorithms in GPU-Enhanced Environments

This work analyzes the role of graphic processing units (GPUs) in the framework of traditional parallel architectures (MIMD, SIMD, ...) to find some effective parameters for the prediction of algorithms performance in a GPU-enhanced environment. We consider a medical imaging application, namely the deconvolution of 3D Fluorescence Microscopy images, implemented in C with CUDA extension on a NVIDIA Tesla C1060: describing its design choices we intend to show how those parameters can affect actual algorithms performance.

Luisa D'Amore
Universita di Napoli, Dep. of Mathematics and
Center for Research on Parallel Computers - CNR
luisa.damore@unina.it

Valeria Mele
University of Naples Federico II
Dipartimento di Matematica e Applicazioni R. Caccioppoli
valeria.mele@unina.it

Diego Romano
ICAR - CNR
diego.romano@na.icar.cnr.it

Almerico Murli
SPACI
almerico.murli@unina.it

CP13

Parallel PDEs Based Numerical Methods for Tracking Cells in 3D+time Microscopy Images

We present parallel algorithms for filtering, segmentation and tracking cells in time lapse confocal microscopy movies. We consider 2D + time videos of Drosophila and 3D + time movies of Zebrafish. The filtering and segmentation is performed using the geodesic mean curvature flow and the generalized subjective surface algorithms starting from cell identifiers obtained by the advection-diffusion level-set center detection method. The individual cell trajectories and the cell lineage trees are obtained by means of finding an ideal path inside segmented spatio-temporal structures. The algorithms are parallelized using massively parallel achitecture and MPI. This is a common work of the groups at Department of Developmental Biology CNRS Gif-sur-Yvette, Institute Curie Paris and Department of Mathematics, Slovak University of Technology in Bratislava.

Karol Mikula
Department of Mathematics
Slovak University of Technology
karol.mikula@gmail.com

CP13

Application of Novel Symbolic Computation to MRI Velocity Inverse Problem

We will describe the problem of reconstructing velocity fields from MRI data, and use it as a case study to demonstrate the performance advantages in using our prototype symbolic code generation, which incorporates novel features which raises the level of abstraction to suit large-scale inverse imaging problems. We show how parallelization and other optimizations are encoded as rules which are applied automatically rather than manually by the programmer in the low-level implementation.

Christopher Anand, Maryam Moghadas, Jessica Pavlin
McMaster University
anandc@mcmaster.ca, maryamm@mcmaster.ca, pavlinjl@mcmaster.ca

CP13

An Efficient Integration Scheme for Voxel-Based Simulations

This contribution is about a highly efficient numerical integration scheme for PDE's within the newly developed Finite Cell Method (FCM), a high order fictitious domain approach applied to voxel-based models from biomedical applications. The integration scheme exploits a precomputation approach that can favorably apply grid computing in the setup-phase and fully exploits multi-core parallelization in the analysis step thus providing a highly responsive numerical analysis tool for user interactive simulations.

Martin Ruess
Technische Universität München
Computation in Engineering
ruess@tum.de

Zhengxiong Yang
Computation in Engineering
Technische Universität München
yang@bv.tu-muenchen.de

Alexander Düster
Technische Universität Hamburg-Harburg
Numerische Strukturanalyse
ruess@tum.de

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

CP13

Scalable Parallel Algorithms for Biological Motif Search

Finding biologically significant patterns in DNA sequences is a well-known problem in computational biology. While many research efforts have been devoted, this problem remains to be challenging, largely due to its computation and memory intensiveness. Recent advances on multi-core techniques have shed light on tackling this problem with a computing platform within reach of most researchers. In this research, we designed and implemented efficient algorithms with OpenMP for Planted Motif Search and demon-

strated scalable speedups.

Ngoc Tam Tran, Chun-Hsi Huang
University of Connecticut
ntt10001@engr.uconn.edu, huang@engr.uconn.edu

CP14

A Parallel Framework for Large Scale Particle Simulation Methods

Particle simulation methods are usually used for the large scale computing of molecular dynamics, electromagnetism, hydrodynamics, monte-carlo transport, material dislocation dynamics, and so on. Usually, these methods require careful trade-off among data structures, communication algorithms, load balancing strategies, and code optimizations. This talk gives a parallel framework for such trade-off and its integration to JASMIN infrastructure to support the peta-scale simulations while billions of particles and tens of thousands of processors are used.

Xiaolin Cao
Institute of Applied Physics and Computational Mathematics
xiaolincao@iapcm.ac.cn

Zeyao Mo
Beijing Institute of Applied Physics and Computational Mathe
zeyao_mo@iapcm.ac.cn

Aiqing Zhang, Xu Liu
Institute of Applied Physics and Computational Mathematics
zhang_aiqing@iapcm.ac.cn, liu_xu@iapcm.ac.cn

CP14

A Unified Communication Algorithm for the Patch-Based Multi-Block or Multi-Block Overlapping Structured Mesh Applications

Patch-based multi-block or multi-block overlapping structured mesh are widely used to discretize the computational domain with complex geometry especially in the three dimensional cases. However, such discretizations seriously challenge the data communication across neighboring patches of neighbor blocks. This talk presents a unified algorithm for such communication and introduces its integration to JASMIN infrastructure to support the peta-scale simulations while tens of thousands of processors are used. Performance results show its robustness.

Hong Guo
P.O.Box 8009
guo_hong@iapcm.ac.cn

Zeyao M, Aiqing Zhang
Institute of Applied Physics and Computational Mathematics
zeyao_mo@iapcm.ac.cn, zhang_aiqing@iapcm.ac.cn

CP14

Two Fast Algorithms for Boxes Operations Oriented to Patch-Based Structured Mesh Applications

Boxes intersection and boxes difference are two type of basic operations for the computing of data dependency of patches for the patch-based structured mesh applications. Two fast algorithms with optimal complexity O(NlogN) are introduced where N is the number of boxes. These algorithms have been integrated into JASMIN infrastructure to support the peta-scale simulations while millions of patches and tens of thousands of processors are used. Benchmark results and real applications show their robustness.

Xu Liu, Aiqing Zhang, Li Xiao
Institute of Applied Physics and Computational Mathematics
liu_xu@iapcm.ac.cn, zhang_aiqing@iapcm.ac.cn, xiaoli@iapcm.ac.cn

Zeyao Mo
Beijing Institute of Applied Physics and Computational Mathe
zeyao_mo@iapcm.ac.cn

CP14

JASMIN: A Parallel Infrastructure for Large Scale Scientific Computing

JASMIN is a parallel infrastructure oriented to simplify the development of parallel software for the multi-physics peta-scale simulations on multi-block or adaptive structured mesh. Patch-based data structures, efficient communication algorithms, robust load balancing strategies, scalable parallel algorithms, object-oriented parallel programming models are designed and integrated. Tens of codes have been developed using JASMIN and have scaled up to tens of thousands of processors. The design theory of JASMIN and its applications are introduced in this talk.

Zeyao Mo
Beijing Institute of Applied Physics and Computational Mathe
zeyao_mo@iapcm.ac.cn

Aiqing Zhang, Xiaolin Cao, Qingkai Liu
Institute of Applied Physics and Computational Mathematics
zhang_aiqing@iapcm.ac.cn, xiaolincao@iapcm.ac.cn, liuqk@iapcm.ac.cn

CP14

Jcogin: a Parallel Framework for Mc Transport Based on Combinatorial Geometry

JCOGIN is a parallel software framework targeted at the development of Monte-Carlo (MC) Transport applications. JCOGIN implements the optimizing geometry computing based on the combinatorial geometry representation for the pin-by-pin computing of the reactor. Similar to the JASMIN infrastructure, patch-based data structures are designed to support large-scale parallel computing, object-oriented and layered software architectures are designed to support the application development. The scalable performances of Monte-Carlo simulations using JCOGIN are demonstrated.

Baoyin Zhang, Gang Li, Yan Ma
Institute of Applied Physics and Computational Mathematics,
Beijing, 100094, P.R. China
zby@iapcm.ac.cn, li_gang@iapcm.ac.cn, ma_yan@iapcm.ac.cn

CP14

A Unified Parallel Sweeping Algorithm for Radiation Or Neutron Transport on Patch-Based Mashes

Parallel flux sweeping algorithms are widely used to iteratively solve the discrete systems arising from the radiation or neutron transport applications. While the patch-based data structures are used, many cases should be studies. This talk gives a unified version of these algorithms on patch-based structured and unstructured mesh and introduces its integrations to the JASMIN infrastructure to support th large scale simulations while thousands of processors are used. Benchmarks and real applications show their robustness.

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

Zeyao Mo
Beijing Institute of Applied Physics and Computational Mathe
zeyao_mo@iapcm.ac.cn

CP15

On Threaded Deterministic Lock-Free Cholesky and ldl^t Factorizations

We look at threaded Cholesky and LDL^T factorizations of symmetric matrices from the stand-point of determinism. We propose new DAG-based algorithms that avoid locks and achieve speedups comparable to those of nondeterminsitic solvers.

Alexander V. Andrianov
SAS Institute Inc.
alexander.andrianov@sas.com

CP15

A Class of Fast Solvers for Dense Linear Systems on Hybrid GPU-multicore Machines

We show how we can accelerate parallel computations of linear system solutions using innovative approaches to reduce the communication due to pivoting. The first technique is based on a random preconditioning of the original matrix to avoid pivoting in general or symmetric indefinite square systems. We also describe how communication-avoiding heuristics can take advantage of hybrid CPU/GPU architectures for general linear systems. For each approach, we provide performance results on current multicore-GPU parallel machines.

Marc Baboulin
INRIA/University of Paris-Sud
marc.baboulin@inria.fr

CP15

Communication-Avoiding QR: LAPACK Kernels Description, Implementation, Performance and Example of Application

We present LAPACK subroutines xTPQRT/xTPMQRT which can be used for implementing the triangle-on-top-of-triangle and triangle-on-top-of-square CAQR sequential kernels. We describe their implementation, their performance, and their uses in various settings (QR updating, CAQR and tiled algorithm). We also describe a reference ScaLAPACK implementation of (parallel distributed) CAQR and compare its performance with the standard ScaLAPACK PxGEQRF.

Rodney James
University of Colorado Denver
rodney.james@ucdenver.edu

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

CP15

Sparse Triangular Linear Solves on Gpus

Solution of sparse triangular linear systems is a common requirement in iterative approaches to solve systems of equations. Computationally efficient solutions of such problems becomes more challenging as computing facilities evolve from traditional, homogeneous supercomputers composed of CPU-only nodes to heterogeneous systems where accelerators, such as GPUs, are coupled to CPU nodes. In this project, we use NVIDIA's CUDA to develop parallel implementations of sparse triangular solves on GPUs. These algorithms have been integrated into the open-source software package PETSc (Portable, Extensible Toolkit for Scientific Computation) from Argonne National Lab (Argonne, IL, USA). Here we show new results of our algorithms for performing the expensive forward and backward substitutions steps in triangular solves.

Chetan Jhurani, Paul Mullowney
Tech-X Corporation
jhurani@txcorp.com, paulm@txcorp.com

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

CP15

Robust Memory-Aware Mappings for Parallel Multifrontal Factorizations

We study the memory scalability of the parallel multifrontal factorization of sparse matrices. We illustrate why commonly used mapping strategies (e.g. proportional mapping) cannot achieve a high memory efficiency. We propose a class of "memory-aware' algorithms that aim at maximizing performance under memory constraints. These algorithms provide both accurate memory predictions and a robust solver. We illustrate our approach with experiments performed on large matrices with the MUMPS solver.

François-Henry Rouet
Université de Toulouse, INPT-ENSEEIHT-IRIT, France
frouet@enseeiht.fr

Emmanuel Agullo
INRIA
emmanuel.agullo@inria.fr

Patrick Amestoy
ENINPT-IRIT, Université of Toulouse, France
amestoy@n7.fr or patrick.amestoy@n7.fr

Alfredo Buttari
CNRS-IRIT, France
alfredo.buttari@enseeiht.fr

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

Jean-Yves L'Excellent
INRIA-LIP-ENS Lyon
jean-yves.l.excellent@ens-lyon.fr

MS1

Code Migration Methodology for Heterogeneous Systems

Migrating legacy software to accelerator-based architectures (e.g. GPU) is a complex process that requires mastering the technological risks (e.g. loss of code portability, extensive code restructuration, debugging complexity) as well as costs. This talk presents an incremental methodology and corresponding tools that helps to implement legacy codes for heterogeneous hardware.

Francois Bodin
CAPS Enterprise, France
francois.bodin@caps-entreprise.com

MS1

Heterogeneous HPC in the Energy Industry - Opportunity or Threat?

The lions share of Shells global HPC capacity is consumed by geophysical seismic imaging and reservoir engineering fluid flow simulations for oil and gas reservoirs. Legacy algorithms and software must be replaced with fundamentally different ones that scale to 1000s of possibly heterogeneous- cores. Geophysical Reverse Time Migration is an example. Traditional techniques are now majorly memory and I/O bottlenecked. We present a novel method that uses domain-decomposition to leverage a multi-gpgpu machine.

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

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

Jon Sheiman
tba
tba@siam.org

MS1

Weather Simulation for a Hybrid-MC System

We will discuss the challanges in numerical weather prediction simulations. The stress would be on GPU considerations for the next generation weather simulations – from identifying computational bottlenecks, to performance modeling, and architecture-specific designs to improve efficiency/time to solution.

Thomas Schulthess
Swiss National Supercomputing Center
schulthess@cscs.ch

MS1

Linear Algebra for Heterogeneous HPC

A wide range of science and engineering applications depend on linear algebra (LA); these applications will not perform well unless LA perform well. We present the newest developments in numerical LA for heterogeneous architectures. Examples will be given from MAGMA - the LAPACK for HPC on heterogeneous architectures. MAGMA employs a hybridization methodology - algorithms are split into tasks and data dependencies, which are fed to a runtime system that dynamically schedules the execution.

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

Azzam Haidar
Department of Electrical Engineering and Computer Science
University of Tennessee, Knoxville
haidar@utk.edu

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

MS2

Energy Aware Performance Metrics

Energy aware algorithms are the wave of the future. The quest for the development of exaflop systems made it clear that extrapolations of current technologies, algorithmic practices and performance metrics are simply inadequate. The community reacted by introducing the FLOPS/WATT metric in order to promote energy awareness. Here we take a step forward and argue that one should instead aim to reduce the total spent energy in conjunction with minimizing the time to solution.

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

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

MS2

On the Evolution of the Green500 to Exascale

The Green500 seeks to encourage sustainable supercomputing by raising awareness to the energy efficiency of such systems. Since its launch at SC07, the list has continued to evolve to serve the high-performance computing (HPC) community.This talk will address (1)new metrics, methodologies, and workloads for measuring the energy efficiency of a HPC system and (2) trends across the list and the associated path to exascale.

Wu-chun Feng

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

MS2

Saving Energy in Sparse and Dense Linear Algebra Computations

We analyze the impact that power-saving strategies, like e.g. the application of DVFS via Linux governors or the scheduling of tasks to cores, have on the performance and energy consumption of dense and sparse linear algebra operations for multi-core processors, hardware accelerators, and clusters of computers. The study involves codes for several matrix kernels from different libraries, which exhibit different levels of concurrency and CPU/memory activity.

Pedro Alonso
Universidad Politécnica de Valencia
Spain
palonso@dsic.upv.es

Manel Dolz
Universidad Jaume I
dolzm@icc.uji.es

Francisco D. Igual
Texas Advanced Computing Center
figual@tacc.utexas.edu

Rafael Mayo, Enrique S. Quintana-Ortí, Vicente Roca
Universidad Jaume I
mayo@icc.uji.es, quintana@icc.uji.es, vroca@icc.uji.es

MS2

Predicting the Performance Impact of DVFS

Predicting performance under Dynamic Voltage Frequency Scaling (DVFS) remains an open problem. Current best practice explores available performance counters to serve as input to linear regression models that predict performance. However, the inaccuracies of these models require large-scale DVFS runtime algorithms to predict performance conservatively in order to avoid significant consequences of mispredictions. This talk will compare those models to an alternative approach based on a proposed new "Leading Loads" hardware counter.

Bronis R. de Supinski
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
bronis@llnl.gov

MS3

Utilizing Multicore and GPU Hardware for Multiphysics Simulation through Implicit High-order Finite Element Methods with Tensor Product Structure

The cost of memory, especially memory bandwidth, is becoming increasingly expensive on modern high performance computing architectures including GPUs and multicore systems. In contrast, floating point operations are relatively inexpensive when they can be vectorized. This relative cost of memory to flops will continue to become even more pronounced due to fundamental issues of power utilization, therefore it is important to rethink algorithms to effectively utilize hardware. Commonly used methods for implicit solves with finite element methods involve assembly of a sparse matrix. Unfortunately, sparse matrix kernels have an arithmetic intensity (ratio of flops to bytes of memory movement) that is orders of magnitude less than that delivered by modern hardware, causing the floating point units to be massively under-utilized. The "free flops' can be effectively utilized by higher order methods which deliver improved accuracy for the same number of degrees of freedom. Effective use of high order methods require eschewing assembled data structures for matrix storage in exchange for unassembled representations. The resulting computation reduces to small dense tensor-product operations and indepedent "physics' kernels at each quadrature point, both of which are amenable to vectorization and capable of delivering a high fraction of peak performance. To reduce the effort required to implement new physics, retain code verifiability, and experiment with different vectorization strategies and solver algorithms, we express the continuum equations in Python and use automatic differentiation, symbolic methods, and code generation techniques to create vectorized kernels for residual evaluation, Jacobian storage, Jacobian application, and adjoints for each block of the system. The performance and effectiveness of these methods is demonstrated for free-surface Stokes flows relevant to glaciology and geodynamics.

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

Aron Ahmadia
King Abdullah University of Science and Technology
aron.ahmadia@kaust.edu.sa

Matthew G. Knepley
University of Chicago
knepley@gmail.com

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

MS3

Hybrid Parallelism for Massive Scale, Fully Coupled, Fully Implicit Multiphysics Simulation

Fully coupled, fully implicit multiphysics simulation is an ideal problem to tackle using parallel processing. As the number of coupled physics is increased, the computing burden quickly expands to the point where supercomputing resources are necessary. To solve these complex problems on modern, high performance computers we've employed a hybrid parallel algorithm comprised of a shared memory threading model coupled to a distributed memory, MPI based model. The algorithm, as implemented in Idaho National Laboratory's MOOSE multiphysics framework, is presented with results demonstrating the efficacy of the approach for complex, nuclear energy related simulations performed on over 10,000 processors.

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

Cody Permann
Center for Advanced Modeling and Simulation

Idaho National Laboratory
cody.permann@inl.gov

David Andrs, John Peterson
Idaho National Laboratory
david.andrs@inl.gov, jw.peterson@inl.gov

MS3

Hybrid Parallelism for Volume Rendering on Large, Multi- and Many-core Systems

With the computing industry trending towards multi- and many-core processors, we study how a standard visualization algorithm, ray-casting volume rendering, can benefit from a hybrid parallelism approach. Hybrid parallelism provides the best of both worlds: using distributed-memory parallelism across a large numbers of nodes in- creases available FLOPs and memory, while exploiting shared-memory parallelism among the cores within each node ensures that each node performs its portion of the larger calculation as efficiently as possible. We demonstrate results from weak and strong scaling studies, at levels of concurrency ranging up to 216,000, and with datasets as large as 12.2 trillion cells. The greatest benefit from hybrid parallelism lies in the communication portion of the algorithm, the dominant cost at higher levels of concurrency. We show that reducing the number of participants with a hybrid approach significantly improves performance.

Mark Howison
Brown University
mhowison@brown.edu

E. Wes Bethel, Hank Childs
Lawrence Berkeley National Laboratory
ewbethel@lbl.gov, hchilds@lbl.gov

MS3

Next-generation Capabilities for Large-scale Scientific Visualization

As high-performance computing progresses from petascale to exascale, the underlying computer architecture is undergoing revolutionary changes. The relative costs of execution, parallelism, communication, and storage are being upended causing the basic assumptions of our workflow and our algorithms to be violated. In this talk we review the ways in which the nature of high-performance computing is changing and how we are responding to the challenges that arise for ParaView, a general-purpose large-scale scientific visualization application. Specifically, we describe the emerging capabilities of prioritized streaming, multiple forms of in situ visualization, at scale visualization, and fine grain parallel algorithms.

Kenneth Moreland, Fabian Fabian
Sandia National Laboratories
kmorel@sandia.gov, ndfabia@sandia.gov

Berk Geveci
Kitware
berk.geveci@kitware.com

Utkarsh Ayachit
Kitware Inc
utkarsh.ayachit@kitware.com

James Ahrens
Los Alamos National Laboratory
ahrens@lanl.gov

MS4

Parallel Clustered Low-rank Approximation of Social Network Graphs

Social network analysis has become a major area of research in recent times. A major problem in implementing social network analysis algorithms, such as friend or product recommendations, is the sheer size of many social networks - for example, the Facebook graph has over 700 million nodes, and even small networks may have tens of millions of nodes. One solution is dimensionality reduction using spectral or SVD analysis of the adjacency matrix of the network but these global techniques do not necessarily take into account local structures or natural community structure in the network. A more promising approach is clustered low rank approximation: instead of computing a global low-rank approximation, the adjacency matrix is first clustered, then low rank approximations of each cluster (i.e., diagonal blocks) are computed, and finally the different local approximations are stitched together. The resulting algorithm is challenging to parallelize not only because of the large size of the social network graphs but also because it requires computing with very diverse data structures ranging from extremely sparse graphs to dense matrices. In this talk, I will describe the first parallel implementation of the clustered low rank approximation algorithm for large social network graphs, and present experimental results that show that this implementation scales well on large distributed-memory machines; for example, on the LiveJournal graph, a standard data set in the social networks area with roughly 4 million nodes and 40 million edges, our implementation scales by a factor of 90 on 128 processors, and processes the graph in a little over a minute. This is joint work with Xin Sui, Keshav Pingali and Berkant Savas.

Inderjit S. Dhillon
University of Texas at Austin
inderjit@cs.utexas.edu

MS4

Parallel Community Detection in Streaming Graphs

Community detection partitions a graph into subgraphs more densely connected within the subgraph than to the rest of the graph. Streaming graphs experience frequent update actions (e. g. vertex and/or edge insertions and/or deletions) without a specified beginning or end. Tackling the current volume of graph-structured streaming data to report communities requires parallel tools. We extend our work on analyzing such massive graph data with a parallel algorithm for community detection in streaming graphs that scales to current data sizes on a massively multithreaded parallel architecture. Our approach for the initial solution is agglomerative, merging pairs of connected intermediate subgraphs to optimize different graph properties. After the graph has been changed by a batch of actions, the community structure is updated without a complete recomputation.

Henning Meyerhenke
Karlsruhe Institute of Technology
Institute of Theoretical Informatics
meyerhenke@kit.edu

Jason Riedy
Georgia Institute of Technology
School of Computational Science and Engineering
jason.riedy@cc.gatech.edu

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

MS4
Parallel Bayesian Methods for Community Detection

Communities in social networks are groups of nodes more closely connected to each other than to the rest of the network. Although there are practical community detection algorithms, they are largely ad hoc. Bayesian inference may lead to more statistically rigorous community detection and better algorithm comparison. We discuss new (non-LDA) Bayesian methods that distinguish conferences from independents in college football schedules. We describe Bayesian inference parallelization and give initial results on larger graphs.

Jonathan Berry
Sandia National Laboratories
jberry@sandia.gov

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

Jiqiang Guo, Daniel Nordman
Iowa State University
jqguo@iastate.edu, dnordman@mail.adp.iastate.edu

Cynthia Phillips
Sandia National Laboratories
caphill@sandia.gov

David Robinson
Sandia National Laboratories
drobin@sandia.gov

Alyson Wilson
IDA Science & Technology
awilson@ida.org

MS4
The Inherent Community Structure in Real-World Graphs

Finding communities in social networks has the subject of many recent research projects. In this talk, we will approach this problem from a modeling perspective, and show that these communities are actually essential defining the graph structure.

Ali Pinar, C. Seshadhri
Sandia National Labs
apinar@sandia.gov, scomand@sandia.gov

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

MS5
Accelerating Algebraic Multigrid on GPUs

Accelerating algebraic multigrid methods on massively parallel throughput-oriented processors, such as the GPU, demands algorithms with abundant *fine-grained* parallelism. Sparse matrix-vector multiplication operations dominate the performance of the cycling phase of algebraic multigrid and we use efficient GPU implementations to achieve notable speedup on a representative set of matrices. We also present novel sparse matrix operations required to construct the AMG hierarchy. The GPU sparse matrix-matrix and maximal independent operations avoid transfer operations and achieve an average of 2× speedup. Our algorithms are expressed as collections of data parallel primitives provided by the Thrust library and available as part of the Cusp sparse matrix library.

Steven Dalton
University of Illinois at Urbana-Champaign
dalton6@illinois.edu

Nathan Bell
NVIDIA Research
nbell@nvidia.com

Luke Olson
Department of Computer Science
University of Illinois at Urbana-Champaign
lukeo@illinois.edu

MS5
Reducing Communication in Algebraic Multigrid

Algebraic multgrid (AMG) solvers have shown excellent scalability on high performance computers, such as IBMs BG/L or BG/P. However, AMG's increasing communication complexity on coarser grids has shown to negatively impact AMGs performance on emerging multi-core clusters, and we expect this to be an even larger problem on future exascale machines. We describe several ways to reduce communication in AMG, including an approach, we call redundant coarse grid solve, as well as an additive AMG variant.

Hormozd Gahvari
University of Illinois at Urbana-Champaign
gahvari@illinois.edu

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

Kirk E. Jordan
IBM T.J. Watson Research
kjordan@us.ibm.com

Luke Olson
Department of Computer Science
University of Illinois at Urbana-Champaign
lukeo@illinois.edu

Martin Schulz
Lawrence Livermore National Laboratory
schulzm@llnl.gov

Ulrike Meier Yang
Center for Applied Scientific Computing

Lawrence Livermore National Laboratory
umyang@llnl.gov

MS5
Reducing Communication in Parallel AMG Utilizing a Domain Decomposition Approach

Abstract not available at time of publication.

Toby Jones
CU Boulder
jones_toby@yahoo.com

MS5
Increasing the Arithmetic Intensity of Multigrid on Many-Core Chips

The basic building blocks of a classic multigrid algorithm, which are essentially stencil computations, all have a low ratio of executed floating point operations per byte fetched from memory. This important ratio can be identified as the arithmetic intensity. Applications with a low arithmetic intensity are typically bounded by memory traffic and achieve only a small percentage of the theoretical peak performance of the underlying hardware. We propose a polynomial Chebyshev smoother, which we implement using cache-aware tiling, to increase the arithmetic intensity of a multigrid V-cycle. This tiling approach involves a trade-off between redundant computations and cache misses. Unlike common conception, we observe optimal performance for higher degrees of the smoother. We also discuss the scalability of the method on many-core chips.

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

Pieter Ghysels
Department of Mathematics and Computer Science
Universiteit Antwerpen
pieter.ghysels@ua.ac.be

Przemyslaw Klosiewicz
Universiteit Antwerpen
Department of Mathematics and Computer Science
przemyslaw.klosiewicz@ua.ac.be

MS6
Automatic Tuning Amg Library for Fluid Analysis Applications

This talk presents the online automatic tuning method of AMG library for fluid analysis based on SMAC method. In this analysis, Pressure Poisson equation needs to be solved every time step. We implemented the AMG library which determines solver parameters by measuring the parameters efficiency at each time step. In our numerical tests, auto-tuning method improved the performance of AMG solver with default setting up to 30 percent.

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

Osamu Nakamura
Sumitomo Metal Industries
nakamura-osm@sumitomometals.co.jp

MS6
Inside a GPGPU Managed Platform with an Auto-tuning JIT Compiler

A GPGPU managed platform built around an auto-tuning JIT compiler is under development. The application virtual machine supports a high level array programming language implemented as a C++ DSL similar to the PeakStream API. Auto-tuning is performed ahead-of-time and just-in-time, reducing cold start effects and increasing stability. ATI (Evergreen) and NVIDIA (Fermi) GPUs are supported. Kernels may be mixtures of images and memory buffers of different precisions, using optimal blocking and vectorization found by auto-tuning.

Christopher Jang
N/A
fastkor@gmail.com

MS6
Krylov Subspace and Incomplete Orthogonalization Auto-tuning Algorithms for GMRES on GPU Accelerated Platforms

Auto-tuning Krylov subspaces size at runtime for the GMRES(m) method may be efficient to minimize the computing time. We introduced recently first algorithms to auto-tuned at run-time the number of vectors targeted by an incomplete orthogonalizations of Krylov basis associated with the GMRES(m) method, minimizing the number of dot-products for a fixed subspace size. We present in this talk some experimental results obtained on GPUs for large sparse matrices and we compare the results with those obtained with subspace size auto-tuning methods.

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

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

Jerome Dubois
CEA-DEN
jerome.dubois@cea.fr

Laurine Decobert, Rqiya Abouchane
CNRS/LIFL
laurine.decobert@polytech-lille.net,
rqiya.abouchane@polytech-lille.net

MS6
Evaluation of Numerical Policy Function on Generalized Auto-Tuning Interface Openatlib

Matrix libraries have many parameters as inputs by the user. They include problem parameters that are difficult to set values, therefore a new approach of automatically setting them is needed. In this presentation, we will present an Auto-tuning interface named OpenATLib. OpenATLib automatically sets a numerical policy defined by users that balances among minimizing of computation time, saving of memory, and satisfying of accuracy requirement for the residual of solution, without difficult parameter settings and selection of numerical algorithms. In addition, we will show a result of performance evaluation with one node of

the T2K Open supercomputer.

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

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

Ken Naono
HITACHI Ltd.
ken.naono.aw@hitachi.com

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

Shoji Itoh
Univ. of Tokyo
itosho@cc.u-tokyo.ac.jp

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

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

MS7

Gyrokinetic Particle-in-Cell (PIC) Simulations on Multi-core and GPU Clusters

The gyrokinetic (particle in cell) PIC formalism is used for studying plasma microturbulence to harness fusion energy. There is an immense need to improve the parallel performance and scalability of GTC on emerging architectures. In this talk we will present strategies for code optimizations for different phases of computation on emerging architectures including multicore processors and nVidia Fermi GPUs. We will also show how optimizations and code transformations are influenced by the underlying architecture.

Khaled Z. Ibrahim
Lawrence Berkeley National Laboratory
Berkeley, CA, USA
kzibrahim@lbl.gov

Kamesh Mudduri
Lawrence Berkeley National Laboratory, Berkeley, USA
kmadduri@lbl.gov

Samuel Williams
Lawrence Berkeley National Laboratory
swwilliams@lbl.gov

Eun-Jin Im
School of Computer Science, Kookmin University, Seoul 136-70
eunjin.im@gmail.com

Stephane Ethier
Princeton Plasma Physics Laboratory
ethier@pppl.gov

John Shalf, Leonid Oliker
Lawrence Berkeley National Laboratory
jshalf@lbl.gov, loliker@lbl.gov

MS7

Using Dynamic Performance Modeling to Assist in Anomaly Detection

The Performance Health Monitor (PHM) is aimed at efficiently pinpointing sources of lost performance and enable applications to experience a consistent performance environment from run to run. Its goal is to explain performance issues and suggest corrective actions or pinpoint likely causes. Using performance snapshots and rapidly executed quasi-analytical performance models it provides performance expectations for the system, allowing problems to be identified even in the absence of performance anomalies or out-of-range measurements.

Darren J. Kerbyson
Pacific Northwest National Laboratory
Richland, WA, USA
darren.kerbyson@pnnl.gov

MS7

Reconciling Explicit with Implicit Parallelism

Implicit parallel programming models simplify the writing of parallel programs, by dynamically detecting and enforcing dependencies between tasks. These models are, however, incompatible with state-of-the-art optimizations used in explicitly parallel programming models, such as Cilk or OpenMP. This talk presents methods to efficiently implement a popular implicit parallel programming model, asynchronous task dataflow, as an extension of explicit parallel programming models, while retaining the performance properties of the latter.

Dimitrios S. Nikolopoulos
Queen's University of Belfast
Belfast, NI
dsniko@gmail.com

MS7

A Simple Machine Model for Refined Performance Analysis and Prediction of Loop Kernels

Simple models based on memory bandwidth and arithmetic intensity have been used successfully to understand and predict the performance of streaming loop kernels. They are insufficient, however, in more complex situations, where neither peak performance nor bandwidth are limiting factors. We use a simple machine model based on in-core code execution and cacheline transfers to arrive at more accurate performance predictions. The model is validated using applications from fluid mechanics and medical physics.

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

MS8

Adventures in Green Computing

Over the past decade, the SCAPE Laboratory has pioneered green computing in HPC through design and promotion of techniques to improve server energy efficiency

without sacrificing performance. In this talk we describe the motivation and challenges facing the Green Computing movement in HPC and beyond and our past and current efforts to build infrastructure to enable profiling, analysis, control, and optimization of the energy used by high-performance systems and applications.

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

MS8

Exascale Computing and the Electricity Bill of Science

High performance computing is getting more and more power demanding. With Exaflops computers we expect at least a power consumption of 20MW. Thus, codes being executed on these machines will cost us millions of dollars for electricity. Long running climate simulations will cost more than a million dollars and the same will hold for other complex simulations. New concepts are needed to reduce the kWh-to-solution and thus the costs of computational science.

Thomas Ludwig
University of Hamburg
ludwig@dkrz.de

MS8

Algorithmic Choices in Dense Linear Algebra and their Effect on Energy Consumption

We present power profiles of two dense linear algebra libraries: LAPACK and PLASMA. They differ substantially in their algorithms and data storage. We show results from the power profiling of the most common routines, which permits us to clearly identify the different phases of the computations. This allows us to isolate the bottlenecks in terms of energy efficiency. We also present some analysis of these numerical dense linear algebra libraries in terms of power requirements.

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

Hatem Ltaief
KAUST Supercomputing Laboratory
Thuwal, KSA
hatem.Ltaief@kaust.edu.sa

MS8

Multiresolution Simulations of Compressible Flows in Energy Efficient Architectures

While performance and energy efficiency of computing hardware are constantly growing, software development for scientific simulations is experiencing an important paradigm shift. Efficient algorithm implementations face rigid constraints about memory layouts, access patterns and FLOP/Byte ratios. We discuss the design of wavelet-based adaptive solvers for compressible flow simulations that run effectively on energy-efficient heterogeneous platforms. We report accuracy, performance and energy-related measurements of such simulations on a multi-core/multiGPU system and an Accelerated Processing Unit (APU).

Diego Rossinelli
Department of Computer Science
ETH Zurich
diegor@inf.ethz.ch

MS9

Emergent Behavior Detection in Massive Graphs

Abstract not available at time of publication.

Benjamin Miller, Nadya Bliss
MIT Lincoln Laboratory
bamiller@ll.mit, nt@ll.mit.edu

MS9

Scalable Graph Clustering and Analysis with Knowledge Discovery Toolbox

Deep analysis of large graphs is indispensable in genomics, biomedicine, financial services, marketing, and national security. Domain experts need the ability to explore graphs directly to apply their domain intuition. Knowledge Discovery Toolbox (KDT) provides a Python interface of graph abstractions and high-level operations, such as clustering and ranking, which are implemented at cluster scale. This talk covers recent advances to extend KDT to semantic graphs, whose vertices and edges have attributes.

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

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

Steve Reinhardt
Microsoft Corporation
steve.reinhardt@aya.yale.edu

MS9

Scalable Algorithms for Analysis of Massive, Streaming Graphs

Graph-structured data in social networks, finance, network security, and others not only are massive but also under continual change. These changes often are scattered across the graph. Repeating complex global analyses on massive snapshots to capture only what has changed is inefficient. We discuss analysis algorithms for streaming graph data that maintain both local and global metrics. We extract parallelism from both analysis kernel and graph data to scale performance to real-world sizes.

Jason Riedy
Georgia Institute of Technology
School of Computational Science and Engineering
jason.riedy@cc.gatech.edu

Henning Meyerhenke
Karlsruhe Institute of Technology
Institute of Theoretical Informatics
meyerhenke@kit.edu

MS9

Multiscale Approach for the Network Compression-friendly Ordering

We present a fast multiscale approach for the generalized network minimum logarithmic arrangement problem. This type of arrangement plays an important role in the network compression schemes, fast indexing, and efficient node/link access operations. The computational results show how far the existing compression-friendly ordering heuristics are from being optimal. We demonstrate significant improvement in the compression of different classes of networks.

Ilya Safro
Argonne National Laboratory
safro@mcs.anl.gov

Boris Temkin
Department of Computer Science
Weizmann Institute of Science
boris.temkin@gmail.com

MS10

Massively Parallel Algebraic Multigrid for Simulation of Subsurface Flow

We present a parallel algebraic (AMG) multigrid method for preconditioning the elliptic problem $-\nabla \cdot (K\nabla u) = f$ on bounded domains Ω. Our method is robust for highly variable or even discontinuous coefficients $K(x)$ (with a range of several orders of magnitude). Due to a greedy heuristic used for the coarsening the method has a low memory footprint and allows for scalable subsurface flow simulations with up to 150 billion unknowns on an IBM Blue Gene/P using nearly 295 thousand cores.

Markus Blatt
Interdisziplinary Center for Scientific Computing
University Heidelberg
Markus.Blatt@iwr.uni-heidelberg.de

Olaf Ippisch
Interdisziplinary Center for Scientific Computing
University of Heidelberg
olaf.ippisch@iwr.uni-heidelberg.de

Peter Bastian
Interdisciplinary Center for Scientific Computing
University of Heidelberg
peter.bastian@iwr.uni-heidelberg.de

MS10

Parallel Algebraic Multigrid for Saddle Point Problems

We present an approach to the construction of parallel AMG methods for saddle point systems of the form $K : \begin{pmatrix} U \\ P \end{pmatrix} \to \begin{pmatrix} U \\ P \end{pmatrix}$, $K = \begin{pmatrix} A & B \\ B^T & -C \end{pmatrix}$ where $A > 0$ and $C \geq 0$. We will demonstrate how to build the transfer operators P and R such that an inf-sup condition for K implies an inf-sup condition for the coarse system $K^C = RKP$ regardless of the coarse grids chosen for U and P.

Bram Metsch
University of Bonn
metsch@ins.uni-bonn.de

MS10

Hybrid Parallelisation of BoxMG on Spacetrees

Spacetree-based adaptive grid codes have proven of value for several massively parallel applications working on dynamically adaptive grids. Spacetree construction algorithms inherently yield a grid hierarchy. Hybrid algebraic-geometric approaches exploit this geometric hierarchy but combine it with the robustness of algebraic approaches. In this talk, we present a matrix-free realisation on Dendy's BoxMG in combination with two tree colourings. These colourings yield a scaling, hybrid parallelisation tested on two supercomputers at LRZ and KAUST.

Tobias Weinzierl
Technische Universitaet Muenchen
weinzier@in.tum.de

MS10

A Geometric Data Structure for Parallel Finite Elements and the Application to Multigrid Methods with Block Smoothing

We introduce the concept of distributed point objects, which allows for a very transparent and lean parallel programming interface and a very flexible support for all components of the parallel multigrid implementation. Moreover, we consider various block smoothers and a parallel direct coarse problem solver (based on nested dissection) in order to increase the robustness of the multigrid algorithm. Finally, parallel multigrid performance is demonstrated for several advanced applications including elasticity, plasticity, and Maxwell problems.

Christian Wieners
Institute for Applied and Numerical Mathematics
Karlsruhe Institute of Technology, Germany
christian.wieners@kit.edu

Daniel Maurer
Institute for Applied and Numerical Mathematics
Karlsruhe Institute of Technology,
daniel.maurer@kit.edu

MS11

Model-Guided Performance Analysis of Applications on Hybrid Architectures

Determining a scalable programming approach for hybrid distributed/shared-memory systems is a difficult problem. A common solution to this problem combines message passing and threads, but this approach can introduce additional performance components that limit scalability. We will describe a model-based performance study of a hybrid MPI/OpenMP implementation of the RandomAccess benchmark. The models identified thread contention issues and quantified how much the updates should be aggregated to increase updates by as much as 7.7X.

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

Van Bui
University of Illinois at Urbana-Champaign
vanbui1@illinois.edu

Boyana Norris

Argonne National Laboratory
norris@mcs.anl.gov

MS11

Performance Modeling for GPUs and Many-core Architectures

GPU architectures are increasingly important in the multi-core era due to their high number of parallel processors. Programming GPGPU applications is a big challenge, but understanding the performance bottlenecks of those parallel programs to improve performance is even more difficult. To provide insights into the performance bottlenecks of parallel applications on GPU architectures, in this talk, we discuss several GPU performance analytical models. These models consider memory-, thread-, instruction-level parallelism to understand performance behavior.

Hyesoon Kim, Sunpyo Hong
Georgia Institute of Technology
hyesoon@cc.gatech.edu, shong9@gatech.edu

MS11

Designing Heterogeneous Multithreaded Instruction Sets from the Programming Model Down

Historically, chip-multithreading research has largely focused on exposing basic ALU pipelining capabilities to the instruction set architecture. However, the lack of departure from traditional architecture techniques has yielded architectures that only provide weak latency hiding capabilities. This has resulted in an increase in operating system pressure to manage the thread scheduling, context switching and concurrency. This presentation shall examine an architecture framework for the tight coupling of commodity multithreading programming models to the underlying ISA.

John D. Leidel
Convey Computer Corporation
jleidel@conveycomputer.com

MS11

Scalable Software for SoC (System on a Chip) Platforms

The system on a chip (SoC) market segment is driven by rapid TTM (time to market), OS scalability, and efficiency. This requires the SW stack to be designed with TTM, scalability and efficiency as first order design constraints. In this paper we propose a layered modular architecture for SoC drivers to enable aggressive driver code reuse between OSes and platforms. This cuts SW development, validation, integration, and maintenance effort. We then discuss the implementation of such an architecture in a media driver that is highly reusable across SoCs in different market segments and operating systems.

Bratin Saha
Intel Corporation
bratin.saha@intel.com

MS12

Topology, Bandwidth and Performance: A New Approach in Linear Orderings for Application Placement in a 3D Torus

Mapping application instances onto physical processors in parallel computers affects application performance. Bin packing is complex but simpler placement algorithms can be aided by node ordering (like Hilbert curves) to bias placement decisions. Increasing scale and an anisotropic interconnect – differing speeds depending on the axis traveled – make for new challenges. Describing these we offer a new placement order balancing several tradeoffs. We demonstrate our improvements running parallel applications on a Top10 supercomputer.

Carl Albing
University of Reading and Cray Inc.
albing@cray.com

Norm Troullier
Cray Inc.
troullie@cray.com

Stephen Whalen
Cray Inc. and University of Minnesota
whalen@cray.com

Ryan Olson, Joe Glenski
Cray Inc.
ryan@cray.com, glenski@cray.com

Hugo Mills
University of Reading
h.r.mills@reading.ac.uk

MS12

Topology Aware Resource Allocation and Mapping Challenges at Exascale

Low-diameter, fast interconnection networks are going to be a pre-requisite for building exascale machines. Such networks are going to present greater challenges for application developers in the area of communication optimization. Careful routing and mapping choices will have to be made to minimize contention and optimize scaling performance of parallel applications on future machines. This talk presents recent work on topology aware mapping at the application and runtime level and possible extensions for future topologies. LLNL-ABS-495576.

Abhinav Bhatele
Lawrence Livermore National Laboratory
bhatele@llnl.gov

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

MS12

Modeling and Management Techniques for Energy Efficiency and Reliability in Multiprocessor Systems

Temperature-driven reliability and performance degradation are among the major challenges for high-performance systems. This work presents reliability-aware job scheduling and power management approaches for multiprocessor systems. Using a novel simulation framework that captures performance and thermal behavior accurately, this talk will demonstrate that techniques offering similar performance, energy, and peak temperature can differ significantly in their effects on processor lifetime. The talk also introduces run-time management techniques based on dynamic learn-

ing and forecasting of the thermal behavior.

Ayse Coskun
Boston University
Electrical and Computer Engineering Department
acoskun@bu.edu

MS12

Algorithms for Processor Allocation and Task Mapping

We present the processor allocation and task mapping problems and then present some results on processor allocation for mesh-connected systems. We describe a 1D curve-based strategy that finds allocations of comparable quality to a fully 3D algorithm MC1x1. We also propose several buddy-system strategies, the best of which finds better allocations than MC1x1 if the job sizes and mesh dimensions are powers of 2. Furthermore, these algorithms are much faster than MC1x1.

Vitus Leung
Sandia National Laboratories
vjleung@sandia.gov

David Bunde, Peter Walker
Knox College
dbunde@knox.edu, pwalker@knox.edu

MS13

Dynamic Selection of Auto-tuned Kernels to the Numerical Libraries in the DOE ACTS Collection

The Advanced CompuTational Software (ACTS) Collection is a set of computational tools and libraries developed primarily at DOE laboratories. We look for ways to improve the performance of the ACTS tools without changing the user interfaces and tool development environments. Our software dependency graph combines techniques to auto-tune numerical kernels that are later used in the implementations of ACTS functionality, making the tools scalable in multi-core systems.

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

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

MS13

Spiral: Black Belt Autotuning for Parallel Platforms

Automatically achieving performance on par with human programmers on current and emerging parallel platforms is a key challenge for the automatic performance tuning community. With the Spiral system (www.spiral.net) we have shown that it is indeed possible to achieve performance portability across a wide range of parallel platforms from embedded processors to supercomputers at or above the performance level achieved by human experts (a.k.a. Black Belt Programmers), for a restricted set of algorithms. We will discuss our experience with building the Spiral system and adapting it to the ever changing landscape of parallel platforms.

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

MS13

High-order DG Wave Propagation on GPUs: Infrastructure and Implementation

Having recently shown that high-order unstructured discontinuous Galerkin (DG) methods are a discretization method for systems of hyperbolic conservation laws that is well-matched to execution on GPUs, in this talk I will explore both core and supporting components of high-order DG solvers for their suitability for and performance on modern, massively parallel architectures. Components examined range from software infrastructure facilitating implementation to strategies for automated tuning. In concluding, I will present a selection of further design considerations and performance data.

Andreas Kloeckner
Courant Institute of the Mathematical Sciences
New York University
kloeckner@cims.nyu.edu

Timothy Warburton
Department of Computational And Applied Mathematics
Rice University
timwar@rice.edu

Jan S. Hesthaven
Brown University
Division of Applied Mathematic
Jan.Hesthaven@Brown.edu

MS13

Dynamical Variation of Eigenvalue Problems in Density-Matrix Renormalization-Group Code

The density matrix renormalization group (DMRG) has been widely employed to examine one-dimensional strongly-correlated electronic structure in high accurate manner. On the other hand, authors have parallelized the DMRG code in multi-core platforms and challenged to extend the target system to two-dimensional ones or ladder ones. The algorithm is composed of two eigenvalue problems, one of which is for large sparse matrices and the other one of which is for relatively small and block dense matrices. These eigenvalue problems dynamically vary through the renormalization process and automatic tuning may be a promising choice to improve the performance. In this paper, we will briefly summarize the algorithm and parallelization techniques and show the future tuning strategy.

Susumu Yamada
Japan Atomic Energy Agency
yamada.susumu@jaea.go.jp

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

Masahiko Machida
Japan Atomic Energy Agency
machida.masahiko@jaea.go.jp

MS14

Modeling Interactive Effects of Power and Performance

Over the past decade, the SCAPE Laboratory has pioneered green computing in HPC through design and promotion of techniques to improve server energy efficiency without sacrificing performance. In this talk we describe two of our recent efforts to model performance and power constraints. The resulting models, called power-aware speedup and iso-energy efficiency, capture key tradeoffs between power, performance and scalability in HPC systems and applications.

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

MS14

Beyond Automatic Performance Analysis

High performance computers consist today of hundreds of thousands of cores organized into a complex hierarchical structure. Programs need to be tuned carefully for that hierarchy to achieve very good performance. This presentation will present the status of Periscope, an advanced automatic performance analysis tool currently under development at Technische Universität München. It will also present the Periscope Tuning Framework (PTF), an extension of Periscope supporting automatic tuning for homogeneous and heterogeneous HPC architectures. PTF will be developed in the new EU FP7 project AutoTune which will start in October 2011.

Michael Gerndt
Technical University of Munich
Garching, Germany
gerndt@in.tum.de

MS14

A Case for More Modular and Intuitive Performance Analysis Tools

The growing scale of future machines, coupled with increasing node complexity, requires new approaches to performance analysis and optimization. Tools will have to be adjustable to particular applications and target scenarios to be able to cope with the increasing data volumes. Further, performance analysis results need to be more intuitive and tools must map their data into the domains most familiar to the user. We will present approaches that help cover these requirements along with several case studies showing their use on both BlueGene and GPU cluster architectures.

Martin Schulz, Abhinav Bhatele, Peer-Timo Bremer, Todd Gamblin
Lawrence Livermore National Laboratory
schulzm@llnl.gov, bhatele@llnl.gov, bremer5@llnl.gov, tgamblin@llnl.gov

Katherine Isaacs
University of California, Davis
isaacs6@llnl.gov

Aaditya Landge
SCI Institute - University of Utah
aaditya@sci.utah.edu

Joshua Levine
University of Utah
jlevine@sci.utah.edu

Valerio Pascucci
SCI Institute - University of Utah
pascucci@sci.utah.edu

MS14

Integrated Hardware/Software Stack for Power and Performance

The energy gap in exascale computing requires that we not move one byte further than needed, also that we have the capability to turn wires on and off adjust component frequencies at a fine-grained level and manage heterogeneous components effectively; it is common these days to hear that these problems are best solved by some particular level of the hardware/software stack (compiler, PL, runtime, the hardware itself) but the reality is graceful cooperation between all levels is required to achieve exascale efficiencies.

Allan E. Snavely
San Diego Supercomputer Center
University of California, San Diego.
allanesnavely@gmail.com

MS15

Optimising Software for Energy Efficiency

Abstract not available at time of publication.

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

MS15

Application-level Tools for Power and Energy Consumption Measurements

Abstract not available at time of publication.

Shirley Moore
Department of Electrical Engineering and Computer Science
The University of Tennessee
shirley@eecs.utk.edu

MS15

Enhancing Energy-Efficiency of Sparse Scientific Computing

We consider sparse scientific computations that typically used in the simulation of models based on partial differential equations. We discuss how we can effectively map features of such codes to features of networked multicore clusters to gain both performance and energy efficiencies. We will provide examples of energy savings and performance gains related to the effective staging of data and thread packing to manage load imbalances.

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

MS15
Mathematical Libraries and Energy Efficiency

Developing mathematical libraries that will enable energy-aware optimisations requires many building blocks to be put in place. In order to be effective there is a need to understand the requirements of the application, the capacity of the underlying system and the likely behaviour of layers of the software stack. The complexity and multiple approaches to architectures and programming models adds to what is already a difficult space to traverse. We will present some results in this area relevant to multicore and threaded environments.

Anne E. Trefethen
e-Science Core Programme
anne.trefethen@oerc.ox.ac.uk

MS16
The Coupled Physics Problem in Space Weather

These are exciting times for space weather research in Europe. There is a convergence of political support for space weather materializing in funding both from ESA and from the European Commission. In the frame of Hardware/Software co-design, it is then natural to choose space weather applications as a basis for future hardware development. The talk will be mainly centered on the modeling challenges posed by space weather numerical simulations. They are both multi-scales and multi-physics which is why they require very heavy computation resources and highly scalable algorithms adapted to the next generation of super-computer. I will especially describe the kinetic part of the simulation and the multi-levels refinement approach we use.

Giovanni Lapenta
Center for Plasma Astrophysics
K.U.Leuven
giovanni.lapenta@wis.kuleuven.be

Arnaud Beck
K.U.Leuven
arnaud.beck@wis.kuleuven.be

MS16
Exploring Software Scalability and Trade-offs in the Multi-Core Era with Fast and Accurate Simulation

Improving the scaling of scientific applications in the multi-core era requires a better understanding of hardware and software interactions. Through the use of fast and accurate computer hardware simulation, scientists and software developers can better understand application bottlenecks as well as perform detailed trade-offs at the level of detail currently not possible by simply running the applications on hardware. We review an arithmetic intensity and energy trade-off study to provide this detailed application view.

Trevor E. Carlson
Computer Systems Laboratory
Ghent University
trevor.carlson@elis.ugent.be

MS16
Communication Avoiding Strategies for the Numerical Kernels in Coupled Physics Simulations

To avoid latency penalties of global reductions for orthogonalization and normalization of vectors in each iteration of a Krylov solver, we propose the use of non-blocking or asynchronous global communication. Standard GMRES and CG are adapted to hide these latencies at the cost of some redundant computations. The resulting algorithms relax the hardware constraints to reach exascale performance on future systems. We look at stability issues and compare to s-step or communication avoiding Krylov methods.

Pieter Ghysels
Department of Mathematics and Computer Science
Universiteit Antwerpen
pieter.ghysels@ua.ac.be

Thomas Ashby
Imec, Leuven, Belgium
ashby@imec.be

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

MS16
Implementation of a Parallel Multiphysics Simulation Code within the Peano Software Framework

We describe how a multiphysics code, using a particle-in-cell algorithm, can be implemented efficiently within the Peano software framework, developed at T.U.Muenchen.

Dirk Roose, Bart Verleye
K.U.Leuven
Dept. of Computer Science
Dirk.Roose@cs.kuleuven.be, bart.verleye@cs.kuleuven.be

MS17
Petascale Simulation of Regional Seismic Wave Propagation Using AWP-ODC

We have developed a highly scalable application AWP-ODC that has achieved M8: a full dynamical simulation of a magnitude-8 earthquake on the southern San Andreas fault up to 2 Hz. M8 sustained 220 Tflop/s for 24-hours on NCCS Jaguar using 223,074 cores, a breakthrough in seismology in terms of computational size and scalability. We discuss the computational and I/O challenges in simulating the very-large scale M8, with the conclusions and an outlook on future work.

Yifeng Cui
San Diego Supercomputer Center
yfcui@sdsc.edu

Kim Olsen
San Diego State University
kbolsen@sciences.sdsu.edu

Thomas Jordan
University of Southern California
tjordan@usc.edu

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

Steven Day
Dept of Geological Sciences
San Diego State university
day@moho.sdsu.edu

Philip Maechling
University of Southern California
maechlin@usc.edu

MS17

Petascale Cosmology Simulations using ENZO

ENZO is a widely-used community code commonly used for cosmology simulations. The strategy in using multi-levels of parallelism in emerging architectures has been highly effective at O(100,000) threads. Current simulations are near the limit of what can be achieved since the cost to complete larger full-scale models is becoming prohibitive. New physics capabilities are required to address fundamental problems so it is essential to concentrate on strong scaling in order to exploit emerging exascale technology.

Robert Harkness
University of California- San Diego
San Diego Supercomputer Center
harkness@sdsc.edu

MS17

Recent Advances in Kinetic Processes In the Magnetosphere Made Possible Through Petascale Computing

Since plasma in the Earth's magnetosphere and the solar wind is collision-less, kinetic effects play a dominant role. Capturing small-scale, kinetic effects in large-scale systems such as the magnetosphere poses a major computational challenge, and capability runs on massively parallel computers are critical to our research. This work presents the latest advances in our kinetic simulations as applied to the modeling of the magnetosphere, along with analysis and visualization techniques developed for massive data sets.

Homayoun Karimabadi
University of California-San Diego
Department of Electrical and Computer Engineering, MC0407
homakar@gmail.com

MS17

Designing and Understanding Nanoelectronic Devices through Petascale Simulations

Driven by Moore's scaling law, the active components of integrated circuits and the transistors size have been drastically reduced to reach nowadays the nanometer scale. Computer-aided design tools help to conceive the next generation transistors thanks to quantum mechanical approach on electron transport. The capabilities of such tool will be presented and the importance of having an application performing at petascale level to get new insight into the physics of nanoelectronic devices will be demonstrated.

Mathieu Luisier
Purdue University
Electrical & Computer Engineering
mluisier@iis.ee.ethz.ch

MS18

Intra-node Performance Tuning of Fast n-body Methods

Abstract not available at time of publication.

Aparna Chandramowlishwaran, Richard Vuduc
Georgia Institute of Technology
aparna@cc.gatech.edu, richie@cc.gatech. edu

MS18

Free Surface Flow with Moving Particle Semi-implicit Method using GPU

The Moving Particle Semi-implicit method (MPS) is a gridless method to solve the Navier-Stokes Equations and has advantages for problems that involve large deformations and topological changes. Employing MPS, we have developed a model of surfactant transports coupled with free-surface flows and utilized the GPU to accelerate the model computation. This talk will report experiences of implementation and effective use of the NVIDIA Tesla C1060 unit and the CUDA libraries.

Hideki Fujioka
Center for Computational Science
Tulane University
fuji@tulane.edu

MS18

Parallel Implementations of Lattice Boltzmann Methods on Multicore and GPU Processors

High performance implementations of the lattice Boltzmann Method provide opportunities to study single and multiphase flow behavior in large, complex systems. Achieving optimal performance requires effective utilization of memory bandwidth at all levels of processor hierarchy. This talk will address algorithms and optimization strategies for GPU, shared memory multi-core multi-socket CPU nodes, and multiple nodes in distributed memory. Performance is analyzed for implementations developed using CUDA, C++ and MPI.

James E. McClure
Department of Environmental Sciences and Engineering
University of North Carolina at Chapel Hill
jemcclur@email.unc.edu

Cass T. Miller
University of North Carolina
Chapel Hill
casey_miller@unc.edu

Jan F. Prins
Dept. of Computer Science
University of North Carolina at Chapel Hill
prins@cs.unc.edu

MS18

CPU+GPU Hybrid Implementation of the Multipole Method for the Method of Regularized Stokeslets

We present a time-dependent Adaptive Fast Multipole Method implementation for the method of regularized stokeslets, targeting a shared-memory multi-core multi-socket CPU node augmented by multiple GPUs. Paral-

lelism is achieved through use of OpenMP tasking facilities. We describe parameter selection for balancing work between the CPU and GPUs and report on scaling as a function of problem size, processors and stokeslet distribution. We report on the efficiency of OpenMP tasking and effects of non-uniform tree depth.

Robert Overman
Department of Computer Sciences
UNC Chapel Hill
reoverma@cs.unc.edu

Ricardo Ortiz
University of North Carolina-Chapel Hill
Mathematics Department
ortiz@unc.edu

Vikram Kushwaha
University of North Carolina, Chapel Hill
kvikram@cs.unc.edu

Jan F. Prins
Dept. of Computer Science
University of North Carolina at Chapel Hill
prins@cs.unc.edu

Michael Minion
University of North Carolina-Chapel Hill
Mathematics Department
minion@unc.edu

MS19

A Highly Scalable Multigrid Solver for Structured Matrices

Structured matrices are used in a variety of applications, especially in physics. Unlike unstructered matrices that are widely used in engineering applications, this structure can be relatively easily exploited by mapping this structure to nowadays supercomputers. We present a parallel multigrid methods that maps 3D problems to torus networks, like those found in the Blue Gene supercomputers, to adress the specific demands of the increasing number of processors, techniques like stencil-collapsing are used.

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

MS19

Efficient Treatment of Varying Coefficients on Hierarchical Hybrid Grids

The Hierarchical Hybrid Grids framework is designed to close the gap between the flexibility of (linear) Finite Element's and the performance of geometric Multigrid's by using a compromise between structured and unstructured grids. A coarse input Finite Element's mesh is split into the grid primitives vertices, edges, faces, and volumes. The primitives are then refined in a structured way, resulting in semi-structured meshes. The regularity of the resulting meshes may be exploited in such a way that it is no longer necessary to explicitly assemble the global discretization matrix and thus permits an efficient matrix-free implementation. This approach allows to solve elliptic partial differential equations with a very high resolution. On such meshes, using constant coefficients for each structured region does not cause any performance drop. However, more general problems are in direct conflict with a matrix-free implementation, if a stencil per grid point has to be stored. We will discuss possibilities to calculate the stencil on-the-fly. The recalculation of each stencil can be accelerated by partial information storing. In addition, a variant for linearly varying coefficients will be suggested. Our aim is to optimize and compare these strategies on current computer architectures in a parallel setting. The talk will carefully consider the trade offs between redundant computation and memory access costs with the goal to achieve an optimal runtime behavior.

Björn Gmeiner, Ulrich Rüde
Computer Science 10 - System Simulation
Friedrich-Alexander-University Erlangen-Nuremberg, Germany
bjoern.gmeiner@informatik.uni-erlangen.de,
ulrich.ruede@informatik.uni-erlangen.de

MS19

Algebraic Multigrid on GPU Clusters

Recent developments in graphics hardware by NVidia and ATI, and associated software development tools as CUDA (and OpenACC recently) enable us to transfer numerical solver components on the recent generation of graphics processing units (GPUs). We solve systems of linear equations with sparse unstructured system matrices derived from f.e. discretizations of PDEs and we present the adaption of an algebraic multigrid solver (AMG) used as preconditioner in a conjugate gradient solver on these GPUs. We achieve an accelerations of 10 wrt. to one CPU core in various practical applications ranging from engineering to medical technology [1].
Stepping forward from one GPU to clusters of GPUs is non-trivial even with a fast interconnect between the compute nodes. Here the multigrid smoothers can be replaced by domain decomposition (DD) smoothers in order to reduce the communication. These results are compared with a simple block DD preconditioner that uses parallel AMG as solver in each block.

[1] B. Rocha, F. Campos, R. Amorim, G. Plank, R. Weber dos Santos, M. Liebmann, G. Haase, *Accelerating cardiac excitation spread simulations using graphics processing units*, Journal on Concurrency and Computation: Practice and Experience (23) pp. 708-720, 2011

Gundolf Haase
University of Graz
gundolf.haase@uni-graz.at

Aurel Neic
University of Graz, Institute for Mathematics
Graz, Austria
aurel.neic@uni-graz.at

MS19

ASIL - Advanced Solvers Integrated Library

Multigrid and domain decompostion are well known methods with optimal complexity. Both methods are widely used in parallel computing environments. In the near future, computers with 10^6 CPUs will be available. To use these computers efficiently, new scaling concepts for solvers are required. In the talk, we present scaling concepts and results of simulations of the new solver toolbox ASIL.

Gabriel Wittum

Simulation in Technology Center
University of Heidelberg, Germany
wittum@techsim.org

MS20

Automatic Code Generation and Tuning for Stencil Kernels on Modern Microarchitecture

PATUS is a code generation and auto-tuning framework for stencil computations targeted at modern multi- and many-core processors, such as multicore CPUs and graphics processing units. Its ultimate goals are to provide a means towards productivity and performance on current and future multi- and many-core platforms. The framework generates the code for a compute kernel from a specification of the stencil operation and a Strategy: a description of the parallelization and optimization methods to be applied. We leverage the auto-tuning methodology to find the optimal hardware architecture-specific and Strategy-specific parameter configuration.

Matthias Christen
University of Basel, Switzerland
m.christen@unibas.ch

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

MS20

Auto-generating Tuned High Performance Software in the Intel(R) Math Kernel Library

The Intel(R) Math Kernel Library (MKL) is well known for its high performance for HPC applications across many scientific domains. Auto-tuning is merely the beginning of our exploration, and our ultimate goal is to auto-generate tuned software that is ready to use. In this talk, we show how we are using models to predict the performance characteristics of several commonly occurring multi-variable problems in math libraries: from predicting the optimal number of threads to determining the efficiency of several subtasks in order to find at run time the optimal distribution between a host and an accelerator. Not only do we give concrete examples, but we show how we generate a decision tree that has a fixed number of leaves and use it to generate optimal code. Our goal goes far beyond finding the fastest answer; our goal is to auto-generate software that runs quick enough to allow finding the best choice at run-time. The algorithms we show are useful enough to auto-generate many different codes useful in practice, and we show how these ideas are enhancing future generations of Intel(R) MKL even more and across multiple domains from BLAS, FFTs, LAPACK, etc.. We show performance results verifying the success of these techniques.

Greg Henry
Intel Corporation
greg.henry@intel.com

MS20

ASPEN-K2: Automatic-tuning and Stabilization for the Performance of CUDA BLAS Level 2 Kernels

ASPEN-K2 is our new implementation for CUDA BLAS level 2 kernels. Most of the vendor supplied implementations for CUDA BLAS show higher performance on Level 3 kernels, especially xGEMM. On the other hand, other kernels are not tuned enough. For example, DSYMV for CUDA3.2 on a Tesla C2050 shows 18.4GFLOPS with enabling ECC (when ECC is turned off, it improves up to 25GFLOPS), though its potential of memory bandwidth is very wider (144GB/sec). Moreover, CUDA BLAS sometimes shows a saw-tooth figured performance fluctuation. In this study, we will present the mechnism of auto-tuning kernels, which selects a better parameter set to increase and stabilize the performance. Newly implemented DSYMV of ASPEN-K2 shows 34.6GFLOPS and 45.8GFLOPS on a Tesla C2050 with enabling and disabling ECC, respectively. In addition, it performs very stable.

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

MS20

BFrame: New BLAS Tuning Framework for Hybrid Supercomputers

With the advent of hybrid supercomputing systems such as Cray XK6, it poses a great challenge of extracting the performance from different computing platform together, including dense matrix kernel tuning. We address this problem through generalization of our auto-tuned BLAS framework (BFrame). BFrame utilizes meta-programming concepts to allow integration of code-generation and integration for any platform specific features seamlessly, and creates highly tuned numerical kernels for different matrix size combinations. In this talk, we will discuss our approach and the performance on Cray XK systems.

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

MS21

Soft Error Correction in Dense Linear Algebra

Many existing approaches to soft error recovery work offline, by checking the result after the computation is finished. For example, algorithm-based fault tolerance (ABFT) first computes the result, then verifies its checksum. We present an online fault tolerance technique for dense linear algebra that detects, locates, and corrects soft errors in the middle of execution. Because the unfinished corrupted computation can be terminated early, the proposed online approach significantly improves computational efficiency over offline approaches.

Teresa Davies
Colorado School of Mines
tdavies@mines.edu

MS21

Fault-tolerant Iterative Methods via Selective Reliability

Current iterative linear solvers assume reliable storage (no "bit flips') and arithmetic. Otherwise, the solver may abort or silently compute the wrong answer. Improving reliability at the system level costs energy, and this cost will become unbearable as processor counts continue to grow. Instead, if the system lets applications apply reliability *selectively*, we can develop iterations that compute the right answer despite faults. We demonstrate this for DRAM

ECC faults with our Fault-Tolerant GMRES algorithm.

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

Patrick G. Bridges, Kurt Ferreira
University of New Mexico
bridges@cs.unm.edu, kurt@cs.unm.edu

MS21

The Effects of Soft Errors on Krylov Methods

As architectures are developed with smaller components operating a lower voltages, soft errors (i.e., bit flips) become potentially worrying, and in fact may not even be detected. We investigate the sensitivity of several Krylov methods to soft errors. We apply sensitivity analysis techniques to assess the relative importance of algorithmic locations where the bit flips occur. We compare their effects on CG, FCG, GMRES, and the fault tolerant GMRES algorithm due to Hoemmen and Heroux.

Victoria Howle
Texas Tech
victoria.howle@ttu.edu

Patricia D. Hough
Sandia National Laboratories
pdhough@sandia.gov

MS21

Connections between Compressed Sensing and Error-correcting Codes

We study connections between channel coding linear programming decoding (CC-LPD) and compressed sensing linear programming decoding (CS-LPD). We show that good LDPC matrices (over the reals) are also good measurement matrices. Consequently, measurement structures based on good LDPC matrices guarantee robust CS-LPD. Using these results, we demonstrate for the first time a class of deterministic matrices that satisfy certain l_1-l_1 robustness conditions in compressed sensing, namely LDPC adjacency matrices of bipartite graphs with logarithmic girth.

Amin Khajehnejad
California Institute of Technology
amin@caltech.edu

MS22

The APGAS Programming Model for Heterogeneous Architectures

Asynchronous Partitioned Global Address Space (Asynchronous PGAS or APGAS) languages, namely X10 and Chapel, provide programmers with higher-level constructs for parallelism and locality. The language infrastructure (compiler and run-time system) provides mechanisms for supporting a program on multicore systems, multicore-accelerator hybrids and clusters of each. This talk will present examples and challenges for X10 and Chapel programs on multicore-accelerator clusters.

David E. Hudack
Ohio Supercomputer Center
dhudak@osc.edu

MS22

Bioinformatics and Life Sciences - Standards and Programming for Heterogeneous Architectures

Heterogeneous architectures hold tremendous potential to favorably impact bioinformatics and life-science applications. This is readily observed across individual technologies with accelerated performance of sequence searching with FPGAs, molecular dynamics with GPUs and large data analysis problems with multiprocessor technologies. The extensive use of open-source solutions in life-sciences and bioinformatics combined with the absence of cross-technology standards creates barriers for heterogeneous architectures. This talk will discuss challenges and potential strategies in bioinformatics and life science applications.

Eric Stahlberg
National Cancer Institute
estahlberg@gmail.com

MS22

Using OpenMP to Program Embedded Heterogeneous Systems

Efficiency improvements to power, performance, and silicon area motivate specialization of processors and accelerators. Heterogeneous processors and accelerators tend to be programmed using low-level vendor specific APIs. OpenMP language committee is working on extensions to OpenMP supporting rapid, maintainable accelerator code development, leaving performance optimizations to the compiler/runtime environment. In this presentation, we will present an overview of the current proposal for extending OpenMP to program accelerators and more general heterogeneous multicore-programming model.

Eric Stotzer
Texas Instruments
estotzer@ti.com

MS22

Developing Programming Models for Scalable Network Simulation on Multi-core Architectures

Network simulation is an area of interest for analog circuit design and neural network applications. The availability of inexpensive multi-core CPUs has inspired research in the programming models used in network simulation with the goal of developing scalable simulation tools. This presentation will include the anatomy of a network simulation code, some software migration paths inspired by homogeneous multi-core architectures, and a forward-looking perspective on the impact of heterogeneous multi-core architectures on this application.

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

MS23

Generic Finite Element Capabilities for Forest-of-octrees AMR

Adaptive mesh refinement using a forest of octrees has recently been demonstrated for a number of large-scale simulations, based on a common sequence of efficient atomic mesh operations such as refinement, 2:1 balance and partitioning. To make this approach more generally applicable, we will present interfaces to mesh generators to enable

arbitrary geometries, and discuss the generic implementation of hanging node constraints. We conclude with meshing examples for ice sheet simulation and electromagnetic scattering.

Carsten Burstedde
Institut fuer Numerische Simulation
Universitaet Bonn
burstedde@ins.uni-bonn.de

Tobin Isaac, Omar Ghattas
University of Texas at Austin
tisaac@ices.utexas.edu, omar@ices.utexas.edu

MS23

New Efficient Algorithms for Parallel AMR on Octrees

Parallel adaptive mesh refinement (AMR) can be realized in an efficient and scalable way using octrees. The two most expensive parts of refinement in terms of CPU time and communication volume are ensuring a 2:1 size balance between neighboring elements and determining element-to-element and element-to-node maps. We present new algorithms for both steps. We apply these improvements and are able to demonstrate excellent speed and scalability on 220,000 cores of the Jaguar Cray XT5 supercomputer.

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

Carsten Burstedde
Institut fuer Numerische Simulation
Universitaet Bonn
burstedde@ins.uni-bonn.de

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

MS23

Towards an Adaptive, Dynamically Load-Balanced, Massively Parallel Lattice Boltzmann Fluid Simulation

Dynamic, massively parallel, Lattice Boltzmann-based CFD applications using adaptive grid refinement or simulating free surface flows can generate severe load imbalances during execution. We will present newly developed, fully distributed data structures and algorithms that are designed to achieve high single-core performance while maintaining scalability for up to hundreds of thousands of compute nodes in simulations which are subject to strong, dynamic workload fluctuations.

Harald Köstler
Universität Erlangen-Nürnberg
harald.koestler@informatik.uni-erlangen.de

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

Christian Feichtinger
Chiar for System Simulation
University of Erlangen-Nuremberg, Germany
christian.feichtinger@informatik.uni-erlangen.de

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

MS23

Scaling Uintah's AMR to 200K Cores and Beyond?

Uintah is a parallel adaptive multi-physics framework used to simulate fluid-structure interaction problems using structured AMR, the ICE flow solver and the Material Point Method. It will be shown how the use of hybrid parallelism (MPI/Pthreads) has made it possible for Uintah to scales to 196k cores on Jaguar. We shall consider how to scale AMR beyond this scale of machine and discuss the changes that will be needed in the future.

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

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

MS24

Engineering the PFLOTRAN Subsurface Flow and Reactive Transport Code for Scalable Performance on Leadership-class Supercomputers

We describe PFLOTRAN—a code for simulation of coupled hydro-thermal-chemical processes in variably saturated, non-isothermal, porous media—and the approaches we have employed to obtain scalable performance with it on some of the largest scale supercomputers in the world. We present detailed analyses of both numerical and I/O routine performance on Jaguar, the Cray XT5 at Oak Ridge National Laboratory, and Intrepid, the IBM BlueGene/P at Argonne National Laboratory, that have guided our choices of algorithms.

Richard T. Mills
Oak Ridge National Laboratory
and The University of Tennessee, Knoxville
rtm@utk.edu

Sarat Sreepathi
North Carolina State University
sarat_s@ncsu.edu

Vamsi Sripathi
Intel Corporation
vamsi.sripathi@ncsu.edu

Kumar G. Mahinthakumar
North Carolina State University
gmkumar@ncsu.edu

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

Peter Lichtner
Division of Earth and Environmental Sciences
Los Alamos National Laboratory

lichtner@lanl.gov

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

Jitendra Kumar, Gautam Bisht
Oak Ridge National Laboratory
jkumar@climatemodeling.org, bishtg@ornl.gov

MS24

Scalability Studies of Coupled Algorithms for Flow in Elastoplastic Porous Media with a Focus on CO_2 Sequestration

Two massively parallel codes developed at Sandia National Laboratories, Aria (flow) and Adagio (solid mechanics), have been coupled to simulate the multiphase flow and deformation in porous media. The performance of the coupled code on a CO_2 sequestration application will be presented. Several aspects of the coupled simulation will be discussed including the cost of using different computational grids for each physics problem, the performance of various solvers, and the overall scalability on simulations.

Pania Newell, Daniel Turner, Mario Martinez, Joseph Bishop
Sandia National Laboratory
pnewell@sandia.gov, dzturne@sandia.gov, mjmarti@sandia.gov, jebisho@sandia.gov

MS24

A Global Jacobian Approach for Nonlinear Domain Decomposition

We describe several domain decomposition schemes for non-linear problems, where the entire coupled system is linearized first in all variables. This allows the non-linear problem to be reduced to a linearized interface system. These schemes are much simpler than previous formulations, in which non-linear subdomain problems were coupled with non-linear interface conditions, leading to nested Newton iterations and forward difference approximations. The global Jacobian approach is shown to have improved stability and scalability properties.

Benjamin Ganis
University of Texas at Austin
Center for Subsurface Modeling
bganis@ices.utexas.edu

Mika Juntunen
University of Texas at Austin, ICES / CSM
mojuntun@gmail.com

Mary F. Wheeler
Center for Subsurface Modeling
University of Texas at Austin
mfw@ices.utexas.edu

Gergina Pencheva
University of Texas at Austin
gergina@ices.utexas.edu

MS24

History Matching and Uncertainty Quantification Using Parallel Ensemble Based Algorithms

The ensemble-based data assimilation algorithms based on the Kalman filter update equation for automatic history matching problems are of increasing interest. However, the sequential implementation is computationally expensive as the methods require relatively high number of reservoir simulation runs. In this talk we present an implementation of a parallel data assimilation framework in which multiple realizations are spawned off in parallel on several partitions of a parallel machine (cluster) each of which are further sub-divided among different nodes (processors) and communication performed at data assimilation time, between the partitions before proceeding again to next assimilation step. Performance results of ensemble Kalman filter and ensemble smoother for a history matching problem are presented.

Reza Tavakoli
UT-Austin
tavakoli@ices.utexas.edu

Gergina Pencheva
University of Texas at Austin
gergina@ices.utexas.edu

Mary F. Wheeler
UT-Austin
mfw@ices.utexas.edu

MS25

Spectrum Slicing Methods for the Kohn-Sham Problem

Solving the Kohn-Sham equation is a standard procedure to determine the electronic structure of atoms, molecules, and condensed matter systems. The solution of this nonlinear eigenproblem is used to predict the spatial and energetic distribution of electronic states. However, obtaining a solution for large systems is computationally intensive because the problem scales super-linearly with the number of atoms. We demonstrate a divide and conquer method that partitions the necessary eigenvalue spectrum into slices and computes each partial spectrum on an independent group of processors.

James R. Chelikowsky
University of Texas at Austin
jrc@ices.utexas.edu

MS25

Parallel Strategy for Finite Difference Linear Scaling Density Functional Theory Calculations on Large Number of Processors

Discretizing the Kohn-Sham equations by Finite Differences on a mesh leads to a straightforward parallelization scheme by spatial domain decomposition. Large scale simulations however remain challenging due to the cubic complexity of traditional algorithms. An O(N) algorithm can be obtained by constraining the electronic wavefunctions to be localized orbitals confined in finite spatial regions. However as data structures become sparse, parallel data distribution becomes more complicated and the ratio between computation and communications becomes less favorable.

Jean-Luc Fattebert
Lawrence Livermore National Lab.

fattebert1@llnl.gov

MS25

Adaptive Local Basis Set for Kohn-Sham Density Functional Theory

We present a new discretization scheme for Kohn Sham density functional theory that constructs basis functions adaptively by solving local problems. These adaptive local basis functions capture the localized atomic structure as well as environmental effects, and the Kohn Sham orbitals are reconstructed in the discontinuous Galerkin framework. Numerical examples indicate that the new basis functions are efficient and accurate for insulating and metallic systems, and can be used for large scale calculation.

Lin Lin
Lawrence Berkeley National Laboratory
linlin@lbl.gov

Jianfeng Lu
Courant Institute of Mathematical Sciences
New York University
jianfeng@cims.nyu.edu

Lexing Ying
University of Texas
Department of Mathematics
lexing@math.utexas.edu

Weinan E
Princeton Univesity
Department of Mathematics
weinan@math.princeton.edu

MS25

Low-Order Scaling Density Functional Methods Based on Quantum Nearsightedness

We present our recent development of low-order scaling density functional (DF) methods based on quantum nearsightedness, which may extend applicability of the DF theory to large-scale systems on massively parallel computers. By making full use of quantum nearsightedness in density matrix, Green's function, and localized basis functions, it will be shown that a wide variety of low-order scaling methods are developed for diagonalization and computation of the exact exchange energy.

Taisuke Ozaki
Japan Advanced Institute of Science & Technology
t-ozaki@jaist.ac.jp

MS26

Locality-centric Optimizations of Large-scale Graph Analysis on Distributed, Multicore/multithreaded Architectures

Communication efficiency on distributed-memory machines and cache performance on cache-based SMPs are the major limiting factors for performance of large-scale graph analysis. As both are directly related to the locality behavior of memory accesses on cache-based architectures, we propose locality-centric optimizations to improve performance. Our target systems are a cluster of SMPs and a multithreaded, multicore platform. Our coordinated memory access scheduling among threads significantly improves both communication efficiency and cache performance. The optimized PGAS implementation in UPC achieved more than two orders of magnitude speedups over the initial implementation on a cluster of SMPs. For multithreaded architectures with weak caches, e.g., the Sun Niagara2, we evaluate the efficiency of memory latency hiding through multithreading. We propose efficient techniques to improve the cache performance.

Guojing Cong
IBM T J Watson Research Center
gcong@us.ibm.com

MS26

Parallel Algorithms for Matching and Coloring

Abstract not available at time of publication.

Ariful Azad
Purdue University
aazad@purdue.edu

Mahantesh Halappanavar
Pacific Northwest National Lab
mahantesh.halappanavar@pnl.gov

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

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

MS26

SMP Algorithms for Computing Spanning Trees and Connected Components

We present new SMP algorithms for computing connected components and spanning trees of large sparse graphs. The algorithms are based on the use of the disjoint-set data structure. Experiments show that the new algorithms are superior to the previous best algorithm for this problem. The algorithms are also appealing in that they are quite simple.

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

MS26

Exploring Architectural Features for Supporting Parallel Graph Algorithms

Graph algorithms, which often use pointer-based data structures, are typical irregular applications. They present unpredictable memory access patterns, control structures, and/or network transfers, require fine-grained synchronization and communication, and operate on very large data sets. All these aspects are challenging for current high performance architectures, which rely on data locality and regular computation to tolerate access latencies. In this talk we discuss the architectural features that enhance or reduce their performance while executing graph algorithms.

Antonino Tumeo, Oreste Villa, Simone Secchi

Pacific Northwest National Laboratory
Antonino.Tumeo@pnnl.gov, oreste.villa@pnnl.gov, simone.secchi@pnnl.gov

MS27

Emerging Challenges and Solutions for Manycore Scientific Computing

The transition to manycore parallel applications will be as disruptive as the transition from serial/vector codes to the first SPMD codes. Here we give an overview of important issues facing algorithms, libraries and applications developers as we migrate to manycore based systems. We present new ideas for addressing these issues and discuss how pattern-based modeling can be used to facilitate application redesign for future systems, providing a natural migration strategy from todays applications to tomorrows.

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

MS27

Hybrid Parallel Ordering Method for Parallel Geometric Multigrid Solver for Fast Finite Element Electromagnetic Field Analysis

We have developed a parallel multigrid solver for electromagnetic wave simulation on a parallel distributed memory computer. The main issue of the research is to parallelize the AFW smoother, which is one of multiplicative Schwarz method, in hybrid multi-process and thread parallel programming model. We introduce a new hybrid parallel ordering technique, in which domain decomposition and block multi-color orderings are effectively utilized. An approximately 800 million DOF problem is successfully solved in 250 s.

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

Yu Hirotani
Kyoto University
hirotani@fem.kuee.kyoto-u.ac.jp

Takeshi Mifune
Department of Electrical Engineering
Kyoto University
mifune@fem.kuee.kyoto-u.ac.jp

Toshio Murayama, Hideki Ohtani
Sony Corporation
toshio.murayama@jp.sony.com, hideki.ohtani@jp.sony.com

MS27

Programming for Mulit-peta and Exaflop Computers using Directive Based Acceleration

This talk will address a programming approach that has the potential to efficiently address future generation hybrid multi-core systems with or without accelerators. The approach consists of utilizing comment line directives to address machine specific features. The directives are OpenACC that are designed to target accelerators. The approach requires that the application developer refactor their program into an efficient form where OpenMP is used on the node and MPI between nodes and vectorization of inner looping structures. Results will be presented that illustrate the value of this approach on large scientific applications.

John M. Levesque
Cray
levesque@cray.com

MS27

Exploiting Multithreaded Tree Parallelism for Multicore Systems in a Parallel Multifrontal Solver

As the number of cores in microprocessors increases, exploiting parallelism in sparse direct solvers only through message-passing and parallel threaded BLAS has some limitations. We analyze the potential of exploiting the inherent tree parallelism of the multifrontal method through multithreaded programming. First, we use a Python simulator to model and evaluate the performance of different approaches. We then experiment the most promising ones on large-scale problems using an actual implementation in the MUMPS solver.

Patrick Amestoy
ENINPT-IRIT, Université of Toulouse, France
amestoy@n7.fr or patrick.amestoy@n7.fr

Alfredo Buttari
CNRS-IRIT, France
alfredo.buttari@enseeiht.fr

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

Jean-Yves L'Excellent
INRIA-LIP-ENS Lyon
jean-yves.l.excellent@ens-lyon.fr

Mohamed Sid-Lakhdar
ENS Lyon
mohamed.sid_lakhdar@ens-lyon.fr

MS28

Avoiding Communication for Banded Eigensolvers

The running time of an algorithm depends on both arithmetic and communication costs, and the relative costs of communication are growing over time. In this work, we present both theoretical and practical results for tridiagonalization of a symmetric band matrix: we describe an algorithm that asymptotically reduces communication compared to previous approaches, and we show that it indeed performs well in practice. We observe speedups over state-of-the-art libraries, both in sequential and shared-memory parallel cases.

Grey Ballard
UC Berkeley
ballard@cs.berkeley.edu

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

Nicholas Knight
UC Berkeley
knight@eecs.berkeley.edu

MS28

Efficient Scalable Algorithms for Hierarchically Semi-separable Matrices

Much progress has been made in developing fast algorithms for structured linear systems, such as those involving hierarchically semi-separable (HSS) matrices. Nearly linear time factorization algorithms have been developed to solve these systems. A key idea behind these algorithms is to fully exploit numerical low rankness in these structured matrices. In this talk, we present new parallel algorithms for HSS matrix operations and their use in the context of factorization-based sparse solvers and preconditioners.

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

Shen Wang
Purdue University
wang273@purdue.edu

Jianlin Xia
Department of Mathematics
Purdue University
xiaj@math.purdue.edu

Maarten de Hoop
Center for Computational & Applied Mathematics
Purdue University
mdehoop@math.purdue.edu

MS28

Extending PETSc's Composable Hierarchically Nested Linear Solvers

The changing landscape of scientific applications and high-performance computing systems requires continued innovation in mathematical algorithms and scalable software for solving sparse linear systems. We will discuss new work in the PETSc library on linear solvers comprising multiple levels of nested algorithms and data models to exploit architectural features and/or problem-specific structure. Topics will include preconditioners for Stokes problems in geodynamics, hierarchical Krylov methods in subsurface flow, and unstructured algebraic-geometric multigrid methods in gyrokinetics

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

Mark Adams
Columbia University
ma2325@columbia.edu

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

Matthew G. Knepley
University of Chicago
knepley@gmail.com

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

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

MS28

Nested Adaptive Partition Hierarchies for Fast Banded Eigensolvers

We consider computing all the eigenpairs, or singular value triplets, of a banded matrix by the fast multipole method (FMM) accelerated divide-and-conquer method. We review the divide-and-conquer method and describe several variants of FMM algorithms for acceleration. We discuss the algorithm properties in approximation, numerics and complexities in sequential and parallel computations. We then introduce algorithmic partition strategies in adaption to various parallel computing architectures.

Bo Zhang
Duke University
zhangb@cs.duke.edu

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

MS29

Software Reliability at Scale

Abstract not available at time of publication.

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

MS29

Reliable Computing on Unreliable Hardware: Correcting Soft Errors in Iterative Methods

Soft errors corrupt a computer's state without halting execution. Supercomputers are especially susceptible to soft errors because of their large number of components. Triple modular redundancy (TMR) generally protects against soft errors. However, TMR can only correct soft errors offline, after the computation is finished. It may also introduce significant overhead. In this talk, we present an efficient algorithmic fault-tolerance technique that detects and corrects soft errors online before the corrupted computation completes.

Zizhong Chen
Colorado School of Mines
zchen@mines.edu

MS29

Soft Error Resilience for Dense Matrix Factoriza-

tions

Soft errors occur silently during computation and can corrupt matrix data in numerical linear algebra codes. We show fault-tolerant algorithms that efficiently handle soft errors during one-sided dense matrix factorizations. Our algorithms are designed with special focus on performing computation in the presence of rounding error. We allow multiple errors to occur simultaneously. The proposed algorithms are evaluated in a LU factorization code on a distributed-memory system.

Peng Du
University of Tennessee - Knoxville
du@cs.utk.edu

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

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

MS29

TFOCS: Convex Optimization for Robust PCA, Error Correction, and Sparse Recovery

TFOCS is a framework for constructing efficient solvers for a variety of convex optimization problems. By design, it is well suited for solving models that arise in compressed sensing and related applications. In this talk, we will provide an overview of the software and present usage examples relevant to sparse recovery and error correction.

Emmanuel Candes
Stanford University
Departments of Mathematics and of Statistics
candes@stanford.edu

Stephen Becker
California Institute of Technology
srbecker@caltech.edu

Michael C. Grant
CVX Research
California Institute of Technology
mcg@cvxr.com

MS30

Asynchronous Execution in n-body Computations

Abstract not available at time of publication.

Aparna Chandramowlishwaran
Georgia Institute of Technology
aparna@cc.gatech.edu

MS30

DAG use in Linear Algebra Software for Contemporary Heterogeneous Multicore Architectures

In this talk, we present a generic framework for architecture aware scheduling and management of micro-tasks on many-core heterogeneous architectures. Applications we consider can be represented as a Direct Acyclic Graph of tasks with labeled edges designating data dependencies. The idea is to assigns computation threads to the cores, overlaps communications and computations and uses a dynamic scheduler based on cache awareness, data-locality and task priority.

Azzam Haidar
Department of Electrical Engineering and Computer Science
University of Tennessee, Knoxville
haidar@utk.edu

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

MS30

Template-based Generic Programming Applied to Large-scale Stabilized Finite Element Simulations

As computational algorithms, hardware, and programming languages have advanced over time, computational modeling and simulation is being leveraged to understand, analyze, predict, and design increasingly complex physical, biological, and engineered systems. Because of this complexity, significant investments must be made, both in terms of manpower and programming environments, to develop simulation capabilities capable of accurately representing the system at hand. At the same time, modern analysis approaches such as stability analysis, sensitivity analysis, optimization, and uncertainty quantification require increasingly sophisticated capabilities of those complex simulation tools. Often simulation frameworks are not designed with these kinds of analysis requirements in mind, which limits the efficiency, robustness, and accuracy of the resulting analysis. In this work, we describe an approach for building simulation code capabilities that natively support many types of analysis algorithms. This approach leverages compile-time polymorphism and generic programming through C++ templates to insulate the code developer from the need to worry about the requirements of advanced analysis, yet provides hooks within the simulation code so that these analysis techniques can be added later. The ideas presented here build on operator overloading-based automatic differentiation techniques to transform a simulation code into one that is capable of providing analytic derivatives. However we extend these ideas to compute quantities that aren't derivatives such as polynomial chaos expansions, floating point counts, and extended precision calculations. We will show example use cases including turbulent flow in a light water nuclear reactor, and stability analysis of a magnetohydrodynamics test problem. The capabilities presented have been released in the open-source Trilinos package called Phalanx and are available for download from trilinos.sandia.gov.

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

Patrick Notz
Sandia National Labs
pknotz@sandia.gov

Eric Phipps

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

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

MS30

Graph-Based Parallel Task Scheduling and Algorithm Generation for Multiphysics PDE Software

Multiphysics PDE solvers are complex in large part due to the nonlinear couplings present among various models. This complexity is exacerbated by the need to support multiple models that imply different nonlinear couplings due to different constitutive relationships, equations of state, required PDEs, etc. Graph-based design approaches, where tasks are created as nodes in a directed-acyclic graph (DAG), allow a wealth of information about the structure of the calculation to be exposed. This talk will address 1) key abstractions in PDE software that facilitate a DAG representation of the problem, 2) scheduling tasks to enable memory reuse and 3) schedulers for task-based parallelism in a multicore environment. We will also discuss hierarchical parallelization techniques for hybrid computing architectures.

James C. Sutherland, Devin Robison
University of Utah
James.Sutherland@utah.edu, devin.robison@utah.edu

MS31

Arbitrary-order Coupled Cluster without Global Arrays

The coupled cluster method is one of the most successful methods in computational chemistry. Current approaches based on fixed tiling of data such as Global Arrays must, at each step, pull tiles of data, contract, and then push tiles back to global storage, possibly reaching across the network every transaction. A new methodology is discussed, which emphasizes structured mapping of the tensors to the processor grid, and the use of systolic tensor contraction algorithms.

Devin Matthews
The University of Texas at Austin
dmatthews@utexas.edu

MS31

Usable and Flexible Foundations for High-Performance Programming Languages and Tools

Designing languages and tools that increase the productivity of scientists and engineers is difficult, due to their demanding workflows and their low tolerance for complexity and change. In this work, we demonstrate a new language extension mechanism called *active typing* that enables modular specifications of compile-time and design-time behavior directly within libraries, laying a foundation for the next-generation of languages and development tools. We demonstrate with a case study involving neurobiological circuit simulations.

Cyrus Omar
Computer Science Department
Carnegie Mellon University
cyrus@cmu.edu

MS31

Parallelizing Population-genetic Simulations

The genome is shaped by evolution through complex, interacting forces, and so understanding the interplay of these forces often falls to population-level simulation. As the scope of genomic data increases with technological advances, so too must the scope of the simulations. However, these simulations do not scale well to parallel architectures due to inherent data dependency issues. In this talk, I will explore methods and their tradeoffs for parallelizing population-level simulations.

Troy Ruths, Luay Nakhleh
Rice University
troy.ruths@rice.edu, nakhleh@rice.edu

MS31

Topology-aware Parallel Algorithms for Symmetric Tensor Contractions

Tensor contractions are the bulk of the computational work involved in many methods for electronic structure calculation. These methods often leverage symmetry within the tensors to reduce the memory footprint. In this talk, we demonstrate how to distribute symmetric tensors in a packed layout across large supercomputer partitions. Our tensor contraction library uses this distribution to map tensors in a topology-aware fashion on Blue Gene/P. This design minimizes inter-process communication and achieves excellent load balance.

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

MS32

Massively Parallel Computational Chemistry Electronic Structure Applications

Computational chemistry can now be used to reliably predict the properties of a wide range of compounds based on correlated molecular orbital theory. New basis sets coupled with effective core potentials, improved software, new correlation methods, and access to high performance, massively parallel computers make it possible to reliably calculate the energetic properties of many compounds. We will describe the software and applied mathematics issues and needs in terms of critical energy applications.

David Dixon
University of Alabama
dadixon@as.ua.edu

MS32

New Algorithms for the Acceleration of Large-Scale First-Principles Molecular Dynamics Simulations

We present recent advances in the development of algorithms for large-scale First-Principles Molecular Dynamics (FPMD). A recursive subspace bisection approach is used to accelerate the computation of the Hartree-Fock exchange energy in FPMD simulations based on hybrid

density functionals. This leads to a hierarchical decomposition of electronic orbitals in rectangular domains. The accuracy of the exchange energy is controlled by a single parameter. Applications to defects in semiconductors and simulations of liquids will be discussed.

Francois Gygi
University of California, Davis
fgygi@ucdavis.edu

MS32

Ab Initio Configuration Interaction Calculations for Nuclear Structure

The atomic nucleus is a self-bound system of strongly interacting nucleons. In Configuration Interaction calculations, the nuclear wavefunction is expanded in Slater determinants of single-nucleon wavefunctions (Configurations), and the many-body Schrodinger equation becomes a large sparse matrix problem. We use a hybrid OpenMP/MPI parallel code to solve for the lowest eigenvalues (binding energies) and eigenvectors (wavefunctions). In ab initio calculations, all nucleons of a given nucleus interact through 2-body and 3-body forces.

Pieter Maris
Iowa State University
pieter.maris@gmail.com

MS32

Linear and Nonlinear Optical Response in TDDFT

Time-dependent density-functional theory (TDDFT) is a generalization of static density-functional theory (DFT) which allows the calculation of excited-state properties. This technique can be used perturbatively in the frequency domain via the Sternheimer equation to calculate linear and nonlinear optical susceptibilities. We have implemented such calculations in periodic systems in Octopus, a open-source TDDFT code (www.tddft.org/programs/octopus), to study liquid chloroform, and benchmarked performance of various parallel linear solvers for solving the Sternheimer equation in real space.

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

MS33

Performance of a Hybrid MPI/OpenMP Version of the HERACLES Code on the Fat Node Curie System

As a grand challenge, the HERACLES code run on the PRACE supercomputer CURIE , to solve hydrodynamic Eulerian 3D problems for Interstellar physic model computing. We will present and discuss experimental results with respect to several parameters. We will analyze the scalability, the I/O management and performance optimizations for both MPI/OpenMP and MPI programming. We will present comparisons between the hybrid programming and the MPI one for large cases.

Edouard Audit
CEA, Maison de la Simulation
Gif-sur-Yvette, France
edouard.audit@cea.fr

Matthias Gonzalez
University Paris 7
mathias.gonzalez@cea.fr

Pierre Kestener
maisson de la simulation
pierre.kestener@cea.fr

Pierre-Francois lavallé
CNRS/IDRIS
pierre-francois.lavallee@idris.fr

MS33

A Library for Performance-Portable Multidimensional Array Computations on Manycore Nodes

The Trilinos Kokkos library supports performance-portable multidimensional array computations on a range of manycore devices through C++ template meta-programming. This library provides multidimensional array data structures that are mapped onto the manycore device, and parallel_for and parallel_reduce interfaces to apply computations to those arrays. Performance-portability is provided though template-based device polymorphism, and demonstrated by compiling and executing the same computational kernel and a finite element mini-application codes on CPU-multicore and NVIDIA-manycore devices.

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

MS33

An Overview of PGAS Approaches

The complexity of numerous applications poses difficulties for porting and validation on new architectures, albeit efforts for improving performance and programmability are expected not to cause holdups. Partitioned Global Address Space (PGAS) Languages offer an alternative model for code development and also porting on parallel systems. Climate applications present opportunities for assessing PGAS, e.g. for overlapping communications and computations, and we will discuss PGAS approaches for refactoring a CG-based solver in an ocean circulation model.

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

MS33

Optimizing Implicit Finite Element Applications for GPUs

MiniFE is a mini-application that approximates the performance impacting computations in an implicit finite element code, from local stiffness matrix generation and assembly to solution by iterative methods. The reference implementation was derived from CPU-oriented approaches. In this presentation we discuss changes made to MiniFE to improve its performance on GPUs.

Alan B. Williams
Sandia National Laboratory
Distributed Systems Research Department

william@sandia.gov

MS34

Holistic Analysis of Energy-performance Tradeoffs on the Shared Memory Multi-core Architectures

Multi-core processors have been poised to dominate the landscape of next generation computing. The large scalability of multi-core processor designs can lead to significant power dissipation. In this paper, we propose a combinational optimization methodology to reduce the overall energy consumption of CPUs, bus and memory hierarchy by scaling the supply voltage. Achieving the optimal voltage settings of different components, we can optimize the overall energy efficiency. We examine this methodology with different parallel algorithms for multiple hardware configurations. Experimental results show that voltage/frequency scaling on both cores and buses can achieve up to 20% in energy reduction compared with scaling on the cores only.

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

MS34

Energy-aware Mappings of Series-parallel Workflows onto Chip Multiprocessors

We discuss the problem of mapping streaming applications modeled as series-parallel-graphs onto two-dimensional tiled CMPs. The objective is to minimize the energy consumption, using DVS techniques, while maintaining a given level of performance. This mapping problem turns out to be NP-hard, but we identify simpler instances, whose optimal solution can be computed by a dynamic programming algorithm in polynomial time. Several heuristics are proposed, and we assess their performance through a set of comprehensive simulations.

Anne Benoit
LIP, ENS Lyon, France
Anne.Benoit@ens-lyon.fr

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

Paul Renaud-Goud
LIP, ENS Lyon, France
paul.renaud-goud@ens-lyon.fr

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

MS34

Speedup-aware Co-schedules for Energy Efficient Workload Management

Today's high performance computing (HPC) systems have immense parallel processing capacity. However, many applications within commonly executed workloads on such systems suffer from Amdahl's Law effects– i.e., sub-linear fixed-problem-speedup on increasing processor counts. We propose speedup-aware co-schedule schemes that seek to enhance overall HPC system energy and throughput efficiencies. We demonstrate their effectiveness at delivering overall system energy improvement by managing the tradeoffs between faster workload completion and slower execution of one or more applications.

Manu Shantharam
CSE, Pennsylvania State University, USA
shanthar@cse.psu.edu

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

MS34

Large Scale Power Aware Scientific Computing

While hardware techniques have proven sufficient to fit current systems within their power budgets, we anticipate that continued increases in processing power will require application-level power-aware techniques. Our initial experiences demonstrate that substantial energy savings are possible for programs using the MPI and OpenMP programming models. In this talk, I will detail energy saving techniques for large scale scientific workloads such as load smoothing based on dynamic voltage and frequency scaling and dynamic concurrency throttling.

Bronis R. de Supinski
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
bronis@llnl.gov

MS35

A Hybrid Data Parallel Non-Rigid Image Registration Method on a Cooperative Architecture

Long execution time of Non-rigid registration (NRR) often presents a major obstacle to its routine clinical use. We present a hybrid data parallel method to parallelize a cutting edge NRR on a cooperative architecture (GPU and multi-core) by partitioning regular data on GPU and irregular data on multi-core. The runtime of the parallel NRR, evaluated on 6 clinical cases, is reduced to less than 1 minute, therefore satisfying the real-time requirement.

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

Yixun Liu
Department of Radiology and Imaging Sciences
National Institutes of Health
yxliuwm@gmail.com

MS35

Compiler-Directed Application Mapping and Optimization on Emerging Multicores

Emerging multicore architectures exhibit a lot of disparities regarding their on-chip cache hierarchies. This motivates for cache hierarchy aware code restructuring and data optimizations. This presentation includes our recent work in this area which targets array/loop-intensive applications from the HPC domain. The techniques covered include cache hierarchy aware loop iteration distribution, neighborhood aware scheduling, inter- and intra-core data layout optimizations, and cache behavior modeling. We

will cover both linear algebraic approach and polyhedral approach.

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

MS35

On Sparse Representations for Image Pattern Classification

Generative image classification solutions commonly involve Gaussian assumptions,capitalizing on Mahalanobis difference metrics. The equivalent formulation is expressed through correlation, where optimization is conducted through a single iteration linear programming problem rather than through convex expectation maximization. The result is a statistically consistent sparse class representation. This feature space representation facilitates scalability and online training concepts. Related concepts, including hierarchical mixture models, Fisher discriminant analysis, and Bayesian decision theory, are related to the new paradigm.

Karl Ni
MIT Lincoln Laboratory
karl.ni@ll.mit.edu

Katherine Bouman
Massachusetts Institute of Technology
klbouman@mit.edu

Nadya Bliss
MIT Lincoln Laboratory
nt@ll.mit.edu

MS35

Logo Retrieval with Parallel Pattern Matching

We introduce two new and faster pattern matching algorithms for logo queries for video, which is a theme study in image-based information retrieval. One of the algorithms calculates the k-means clustering of a large point set in a high-dimensional space, the other locates the k-nearest neighbors for relatively large k values. We present experiments with established retrieval data sets on multicore computers and GPUs.

Nikos Pitsianis
Department of Electrical & Computer Engineering
Duke University
nikos.pitsianis@eng.auth.gr

Nikos Sismanis
Department of Electrical & Computer Engineering
Aristotle University
nsismani@auth.gr

MS36

Adaptive Magnetohydrodynamics Simulations with SAMRAI

Adaptive Mesh Refinement (AMR) allows for both hi-resolution and large spatial domain computations. This talk will focus on current progress in adding AMR capabilities to an existing 3D plasma model (pixie3d) using SAMRAI (Structured Adaptive Mesh Refinement Application Infrastructure). The resulting code will be able to simulate new problems that require spatial resolution not currently accessible without AMR.

Mark Berrill, Luis Chacon, Bobby Philip
Oak Ridge National Laboratory
berrillma@ornl.gov, chaconl@ornl.gov, philipb@ornl.gov

MS36

Cello: An Extremely Scalable Adaptive Mesh Refinement Framework

We present `Cello`, a highly-scalable parallel AMR framework currently under development. Its octree-like distributed data structure includes enhancements to scalability in terms of hierarchy size and depth, and supports both particle and uniform mesh data. The message-driven parallel programming system used, `CHARM++`, enables overlapping communication and computation, and provides dynamic load-balancing and checkpoint-restart capabilities. Despite being developed and tested using a specific cosmology application, `Cello` is a general AMR framework usable with other applications.

James Bordner
University of California, San Diego
jobordner@ucsd.edu

MS36

Massively Parallel Finite Element Simulations with deal.II

Implementing parallel adaptive mesh refinement in solvers for partial differential equations is complex and often results into code tailored to specific problem domains, code without the ability to scale to a large number of machines, or often both. We extended the existing finite element library deal.II to support massively parallel computations in a generic setting. We will review the difficulties, features that are currently supported, and showcase complex applications.

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

Timo Heister
Texas A&M University
Department of Mathematics
heister@math.tamu.edu

MS36

Adaptive Mesh Refinement Based Upon An Abstract Grid Interface with Application to Compressible Flow.

In this talk we present the Distributed and Unified Numerics Environment (DUNE, www.dune-project.org), a software framework for High Performance Computing. Based on an abstract grid interface adaptive mesh refinement is utilized for the simulation of compressible flow problems. In parallel applications dynamic load balancing is applied to maintain scalability of the code. Examples will demonstrate the performance of the code in parallel applications.

Robert Kloefkorn
University of Freiburg
robertk@mathematik.uni-freiburg.de

MS36

Block-Structured AMR Applications Without Distributed Meta-data

Abstract not available at time of publication.

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

Daniel Graves
Lawrence Berkekey National Laboratory
dtgraves@lbl.gov

MS37

Parallelization and Software Concepts for Memory-Efficient Dynamically Adaptive Grid Traversals

We present an inherently memory- and cache-efficient approach for dynamically adaptive simulations on recursively refined triangular grids. Grid cells are stored and processed in an order defined by the Sierpinski curve; the resulting locality properties are exploited for optimised serial implementation and parallelisation. A layered software design is used to provide an element-oriented user frontend that largely hides the algorithmic complexity. As target applications, we will present examples from porous media flow and tsunami simulation.

Michael Bader
Technische Universität München
Department of Informatics
bader@in.tum.de

Oliver Meister, Kaveh Rahnema
University of Stuttgart
Institute of Parallel and Distributed Systems
oliver.meister@ipvs.uni-stuttgart.de,
kaveh.rahnema@ipvs.uni-stuttgart.de

MS37

Shared and Distributed Memory Parallelism for Octree Construction

Abstract not available at time of publication.

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

MS37

Massively Parallel Fluid-structure Interaction Simulation of Blast and Explosions Impacting on Realistic Building Structures with a Block-structured AMR Method

We have coupled the massively parallel AMR hydro framework AMROC to the explicit solid dynamics solver DYNA3D. By utilizing a level-set-based Cartesian embedded boundary method in combination with structured AMR, efficient and scalable data exchange along the fluid-structure interface is achieved. We benchmark the parallel performance on distributed memory systems and present several large-scale simulations involving shock waves impacting and fracturing concrete structures, among them prototypical hydrogen explosions in a nuclear reactor containment building.

Ralf Deiterding, Stephen Wood
Oak Ridge National Laboratory
Computer Science and Mathematics Division
deiterdingr@ornl.gov, woodsl4@ornl.gov

MS37

A Comparison of Patch-based and Loop Level OpenMP for AMR in GeoClaw

Our goal is to parallelize GeoClaw to make use of multicore desktop machines and thereby reduce simulation times for depth-averaged geophysical flows by an order of magnitude. We have chosen OpenMP as the easiest approach to parallelism but questions of granularity still exist. We present results for both patch-based and loop-level AMR, some of the algorithmic changes that were needed for better performance, and discuss when to use which approach.

Kyle T. Mandli
University of Washington
Dept. of Applied Mathematics
kyle@ices.utexas.edu

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

Sandra May
New York University
Courant Institute
may@cims.nyu.edu

MS38

A Multiphase Model of the Thermal Evolution of Pyroclastic Density Currents Generated by a Volcanic Eruption

Volcanic eruptions create deadly pyroclastic density currents (PDCs) composed of hot gases and rocks. The deposits from PDCs produce large data sets that require comparison to large scale numerical models in order to better understand some of the physical processes involved in PDCs. Our particular focus is a multiphase Eulerian-Eulerian-Lagrangian model of PDCs. This model will provide a better understanding of the thermal evolution of a current such as ambient air entrainment and particle concentration.

Mary Benage, Josef Dufek
Georgia Institute of Technology
mary.benage@eas.gatech.edu, josef.dufek@eas.gatech.edu

MS38

Efficient Parallel Implementations of the Fast Marching Method for Solving Static Hamilton-Jacobi Equations

The fast marching method is a numerical technique for solving static Hamilton-Jacobi equations based on a Dijkstra-like ordered upwind approach. It has been used extensively in many fields including fluid dynamics, chemistry, computational geometry, and image processing. However, the serial nature of the required ordering has served as an obstacle to parallelization. We present tech-

niques to overcome this problem for narrowbanded regions and extend this to general domains with efficient load balancing by using a multilevel approach.

Jeffrey Donatelli
Department of Mathematics
Lawrence Berkeley National Laboratory / UC Berkeley
jdonatel@math.berkeley.edu

MS38

Accounting for Time Dependent Deformational Forces in Large-scale Hemodynamics Simulations

We will present a mesoscopic fluid model, the Lattice Boltzmann equation, and applying it to model bloodflow in coronary artery. By scaling our application to run on up to 294,912 cores, we can model a full heartbeat at the resolution of the red blood cells. Now, we are studying the deformational forces the heart exerts on the arterial flows across heartbeats. We will discuss the both computational optimizations and additions to our model that this required.

Amanda Peters
Harvard School of Engineering and Applied Sciences
apeters@fas.harvard.edu

Efthimios Kaxiras
Dept. of Physics and School of Engg. and Applied Sciences
Harvard University
kaxiras@physics.harvard.edu

MS38

Continuum Kinetic Method for Plasmas on GPU Clusters using OpenCL

We present a code for plasma simulation on many-core architectures such as GPU clusters. We describe the framework design, exemplary physics results, and present performance data. In order to maximize performance and power efficiency on coming systems, codes should minimize data movement. The algorithms employed are local, explicit, and benefit from predictable data access patterns as opposed to PIC methods. Presented physics results highlight the influence of hot particles in tokamaks captured with continuum model.

Noah F. Reddell
University of Washington
reddell@uw.edu

MS39

Large-Scale Parallel Simulations of CO_2 Storage in Brine Aquifers

Accurate understanding of the mechanisms that control the migration and trapping of CO_2 in subsurface is crucial to design future storage projects for long term and safe containment. Deep saline aquifers can provide vast and safe storage for CO_2 pending a proper understanding of the displacement character of CO_2-brine system at in situ conditions. Several prototype aquifer models were simulated to determine the impact of injection rates, interfacial tension variations, and history dependent relative permeabilities (hysteresis) on CO_2 migration, trapping, and injectivity.

Mojdeh Delshad
CPGE, University of Texas
mdelshad@pe.utexas.edu

Xianhui Kong
University of Texas at Austin
xianhui@ices.utexas.edu

Mary F. Wheeler
Center for Subsurface Modeling
University of Texas at Austin
mfw@ices.utexas.edu

MS39

Title Not Available at Time of Publication

Abstract not available at time of publication.

Kirk E. Jordan
IBM T.J. Watson Research
kjordan@us.ibm.com

MS39

Massive Parallel Geostatistical Inversion of Coupled Processes in Heterogeneous Porous Media

The quasi-linear geostatistical approach is used to estimate the spatial distribution of a heterogeneous hydraulic conductivity field based on point-wise measurements of observable quantities. For a high resolution 3D representation of the model domain the occurring elliptic or hyperbolic advection-dominated problems are discretized by Finite Element Schemes using fully scalable domain decomposition techniques. In order to reduce interprocess communications and to improve the scalability of the code on larger clusters another level of parallelization has been added.

Adrian Ngo
University of Heidelberg
adrian ngo ¡adrian.ngo@iwr.uni-heidelberg.de¿

Peter Bastian
Interdisciplinary Center for Scientific Computing
University of Heidelberg
peter.bastian@iwr.uni-heidelberg.de

Olaf Ippisch
Interdisziplinary Center for Scientific Computing
University of Heidelberg
olaf.ippisch@iwr.uni-heidelberg.de

Wei Li, Ronnie Schwede, Olaf A. Cirpka
Center for Applied Geoscience
University of Tuebingen
wei.li@uni-tuebingen.de, ronnie.schwede@uni-tuebingen.de, olaf.cirpka@uni-tuebingen.de

MS39

Smoothed-Aggregation Algebraic Multigrid for Porous Media Simulations

We consider smoothed aggregation algebraic multigrid (AMG) methods for Darcy flow models of porous media problems. The first is PARCEL/SDF, where we consider AMG as a subdomain solver in a mortar-based domain decomposition scheme. The second involves higher-order discontinuous Galerkin (DG) discretizations. Here we employ a specialized prolongator which allows us to move from a higher-order DG discretization to a first order nodal discretization to which standard AMG techniques can be ap-

plied.

Christopher Siefert
Sandia National Laboratories
csiefer@sandia.gov

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

Gergina Pencheva
University of Texas at Austin
gergina@ices.utexas.edu

Axel Gerstenberger
Sandia National Laboratories
agerste@sandia.gov

MS40

Insights Obtained from Optimal Sparse Matrix Partitioning

Optimal partitioning of small sparse matrices can give insights into how larger matrices should be partitioned. We present a branch-and-bound method for optimal bipartitioning of small matrices, compare it to ILP methods, and derive lessons for heuristic partitioning, for both structured and unstructured problems. We also give optimal communication volumes for about 150 out of the 200 matrices of smallest size in the Florida collection.

Rob H. Bisseling, Daan Pelt
Utrecht University
Mathematical Institute
r.h.bisseling@uu.nl, daan.pelt@gmail.com

MS40

Load-Balancing Spatially Located Computations using Rectangular Partitions

Distributing spatially located heterogeneous workloads is an important problem in parallel scientific computing. We investigate the problem of partitioning such workloads (represented as a matrix of non-negative integers) into rectangles, such that the load of the most loaded rectangle (processor) is minimized. Since finding the optimal arbitrary rectangle-based partition is an NP-hard problem, we investigate particular classes of solutions: rectilinear, jagged and hierarchical, and present new optimum and heuristic algorithms with different time/quality trade-off.

Erdeniz Bas, Erik Saule, Umit V. Catalyurek
The Ohio State University
Department of Biomedical Informatics
erdeniz@bmi.osu.edu, esaule@bmi.osu.edu, umit@bmi.osu.edu

MS40

Exploiting Geometry and Adjacencies in Mesh Partitioning

Hypergraph partitioning holds many advantages for partitioning unstructured meshes. It has been shown to reduce communication volume in mesh partitions by 10-15% relative to graph partitioning, but can take significantly more time to compute. In this talk, we show how hypergraphs can model meshes more richly than graphs. We also present a more scalable hypergraph partitioner that exploits geometric information available in meshes.

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

Nicholas Aase
University of New Mexico
nicholas.aase@gmail.com

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

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

Cedric Chevalier
CEA
France
cedric.chevalier@cea.fr

MS40

Partitioning Problems on Trees and Simple Meshes

We survey partitioning problems on tree and mesh-structured domains. Certain variants of those problems are polynomial time solvable, some other variants are NP-complete. We catalog exact algorithms for the polynomial time solvable cases and very effective, specialized heuristics for the two other class of problems. We evaluate the current state of the art partitioners with respect to aforementioned exact algorithms and heuristics. We will also try to collect applications in which the partitioning problems arise.

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

MS41

Adaptive Sampling for Improving the Results of Dynamic Network Analysis

Most network sampling methods focus on reducing the network size and aim to maintain all the important characteristics of the original network in the sample. We present adaptive sampling where the goal is to extract only those portions of the network that are relevant to the given analysis objective. In particular, we describe a parallel chordal-graph based sampling algorithm on dynamic networks, that improves the results of clustering by extracting highly connected regions of the network.

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

Kanimathi Duraisamy
University of Nebraska at Omaha
kduraisamy@unomaha.edu

MS41

Chordal Graph Preconditioners for Sparse Linear Systems

We describe a family of preconditioners for sparse linear systems that is based on computing a chordal subgraph of the system. Because these preconditioners have no-fill factorizations, they can be attractive for problems where memory is limited. We discuss both theoretical and experimental results, then planned future work including parallelization.

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

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

Sriram Srinivasan
Department of Computer Science
University of Nebraska, Omaha
sriramsrinivas@unomaha.edu

MS41

Chordal Graph Theory and Triangulation Algorithms for Solving Problems in Perfect Phylogeny

Determining whether a set of species on qualitative characters has a perfect phylogeny is a classic NP-Hard problem in computational biology. In 1975, Buneman showed that solutions are characterized by legal chordalizations of the partition intersection graph. I will discuss our recent results using minimal triangulation theory to investigate 3-state perfect phylogeny and illegal triangulations to solve the character removal problem, where one wants to minimize the characters removed to obtain a perfect phylogeny.

Rob Gysel
Department of Computer Science
U. C. Davis
rsgysel@ucdavis.edu

Dan Gusfield
UC Davis
gusfield@cs.ucdavis.edu

MS41

Weak Precedence Orders of Paths in Trees

Partial orders (and some generalizations) defined by paths in a rooted tree have a natural application to searching for the best match of a given string to a set of possible patterns. Another problem related to these orders corresponds to determining whether there is a tree in which given nodes obey constraints saying these nodes must occur consecutively on a root to leaf path; this is a form of constraint between the interval constraints solved by PQ-Trees, and the subtree constraints solved by clique trees as part of chordal graph recognition.

Jeremy Spinrad
Vanderbilt University
spin@vuse.vanderbilt.edu

Ross McConnell
Colorado State University
rmm@CS.ColoState.edu

MS42

Supporting Diverse Parallel Environments in the Trilinos Library

The prominence of multicore processors is particularly challenging for library developers. The need to support multiple parallel programming APIs can result in a significant duplication of effort. Furthermore, multicore has complicated the balance between a manageable code-base and providing users with desired capability. This talk will demonstrate approaches explored in Trilinos to utilize generic programming techniques to provide portable parallel library support while allowing user extensibility at multiple levels of granularity.

Christopher G. Baker
Oak Ridge National Laboratory
bakercg@ornl.gov

MS42

Physis: A Domain-Specific Language for FDM Code on GPU Clusters

Physis is a portable domain-specific language for stencil computations with structured grids. We extend the standard C language with a small set of custom declarative constructs that allow for portable and implicitly-parallel programming for a variety of parallel computing environments. We present performance and productivity studies of the Physis framework using the TSUBAME2.0 GPU supercomputer and demonstrate that both high productivity and high performance can be achieved via the domain-specific approach.

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

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

MS42

Coarse Grid Solvers in Parallel Multigrid Methods using OpenMP/MPI Hybrid Programming Models

In this study, we analyzed communication patterns of 3D finite-volume based simulation codes with MGCG solvers using OpenMP/MPI hybrid parallel programming models for multicore clusters. We proposed a new strategy for solving equations at the coarsest level (coarse grid solver), and evaluated that using T2K Open Supercomputer (U.Tokyo) with 8,192 cores and Cray XE6 (LBNL) with 12,288 cores. New coarse grid solver improved the scalability of the multigrid solver dramatically.

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

MS42

CnC for HPC

This talk summarizes our experiences in expressing

and tuning computations from high-performance computing (HPC) in the Concurrent Collections programming model. We discuss how CnC facilities expression of novel asynchronous-parallel algorithms for dense linear algebra, and show that with appropriate tuning, we can match or exceed vendor-optimized codes. We also report on our preliminary work to express more irregular computations in CnC, including tree-based fast multipole methods for n-body problems.

Aparna Chandramowlishwaran, Piyush Sao
Georgia Institute of Technology
aparna@cc.gatech.edu, piyush.feynman@gmail.com

Kathleen Knobe
Intel
Kath.knobe@intel.com

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

MS43

Fast Convolution for Molecular Dynamics on Parallel Computers

Molecular dynamics is a well-established technique for simulating complex systems in physics, chemistry, biology and materials science. However, fundamental hurdles remain in order for molecular dynamics to reach either spatial or temporal scale of most realistic systems. The biggest challenge is the lack of efficient and scalable algorithms for treating long-range interactions, which relies on fast convolution methods. In this talk, I will discuss some recent progress and challenges toward fast convolution algorithms with their application to molecular dynamics on emerging heterogeneous platforms.

Xiaolin Cheng
Oak Ridge National Laboratory
chengx@ornl.gov

MS43

Software Strategies for Improving Load Balance on Anton, a Special-Purpose Machine for Molecular Dynamics Simulation

Anton is a special-purpose, massively parallel machine that uses a combination of specialized hardware mechanisms and restructured software algorithms to accelerate molecular dynamics simulations by orders of magnitude compared with the previous state of the art. In this talk, I will discuss some of the software strategies developed to achieve and maintain peak performance through very long simulations, with particular attention to dynamic load balancing.

Doug J. Ierardi
D. E. Shaw Research, LLC
doug.ierardi@deshaw.com

MS43

Communication-optimal H-Matrix Multiplication

Optimal parallelizations of $\mathcal{H}$-matrix/vector and $\mathcal{H}$-matrix/$\mathcal{H}$-matrix multiplication are presented. For the former, each process requires $\mathcal{O}(\mathcal{N}\log_\in(\mathcal{N})/\sqrt{})$ work and storage, $\mathcal{O}(\mathcal{N}/\sqrt{})$ communication volume, and $\mathcal{O}(\log_\in(\sqrt{}))$ latency (over a fully connected network). Similarly, the $\mathcal{H}$-matrix/$\mathcal{H}$-matrix multiplication algorithm requires $\mathcal{O}(\mathcal{N}\log^\in_\in(\mathcal{N})/\sqrt{})$ work, $\mathcal{O}(\mathcal{N}\log_\in(\mathcal{N})/\sqrt{}$ communication volume, and either $\mathcal{O}(\log_\in(\sqrt{}))$ or $\mathcal{O}(\log_\in(\sqrt{})\log_\in(\mathcal{N}))$ latency, depending upon whether or not per-process memory usage is increased from $\mathcal{O}(\mathcal{N}\log_\in(\mathcal{N})/\sqrt{})$ to $\mathcal{O}(\mathcal{N}\log^\in_\in(\mathcal{N})/\sqrt{})$.

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

Lexing Ying
University of Texas
Department of Mathematics
lexing@math.utexas.edu

MS43

Accelerated Computation of the CS Decomposition

The CS (Cosine-Sine) matrix decomposition is a sibling of the eigenvalue and singular value decompositions. In recent years, we developed a new, two-phase method for its computation: simultaneous bidiagonalization followed by simultaneous diagonalization. We now accelerate both phases on architectures with memory hierarchies and multiple processors. Special matrix structures are crucial for efficiency and stability.

Brian D. Sutton, Kingston Kang, William Lothian, Jessica Sears
Randolph-Macon College
bsutton@rmc.edu, jiankang@go.rmc.edu, williamlothian@go.rmc.edu, jessicasears@go.rmc.edu

MS44

Large Scale Extrinsic Fragmentation Simulation

Cohesive fracture simulations usually require a fine mesh discretization in order to capture non-linear behavior near the crack tip regions. This demands a large amount of computational resources, so models of reduced geometries are used in the simulations. However, reduced models do not accurately represent the original experiments, due to material-dependent length scales, and dependency of the direction of crack propagation on the mesh refinement level. As real-scale cohesive fracture simulations may not be practical or feasible to run in a serial code executed on a single workstation, parallel processing becomes mandatory. We have investigated large-scale extrinsic fragmentation simulations. First, we have proposed a new topological data structure for distributed mesh representation. We then have identified different strategies to parallelize a serial simulation code, exploring symmetrical computations on different partitions in order to minimize network communication. Computational experiments demonstrate the efficiency and scalability of our solution. We have also explored the use of many-core architectures, such as the ones provided by modern graphics cards. A GPU-based computational framework is presented to deal with dynamic failure events simulated by means of cohesive zone elements.

Waldemar Celes
PUC-Rio, Brazil
wceles@gmail.com

Glaucio Paulino
Department of Civil and Environmental Engineering
University of Illinois
paulino@uiuc.edu

Rodrigo Espinha, Andrei Monteiro
PUC-Rio, Brazil
rodesp@tecgraf.puc-rio.br, andrei@tecgraf.puc-rio.br

MS44

Towards Massively Parallel Unstructured Mesh Drift-diffusion Semiconductor Simulations

Fully-implicit Newton-Krylov solution methods are popular for their robustness, but for massively parallel simulations this requires the scalable solution of very large sparse linear systems. Preconditioners that provide global coupling such as multigrid-based approaches are needed. Algebraic multigrid has the advantage over geometric multigrid of not requiring the generation of coarser meshes. We employ an FEM discretization on unstructured meshes and will show scaling results for two multiphysics systems: semiconductor drift-diffusion and magnetohydrodyamics (MHD).

Paul Lin
Sandia National Laboratories
ptlin@sandia.gov

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

MS44

Adjoint-Based Algorithms for Complex Aerodynamic Flows in Large-Scale Computational Environments

Recent work aimed at developing adjoint-based algorithms on unstructured grids for design optimization, error estimation, and mesh adaptation for steady and unsteady aerodynamic flows over complex aerospace configurations is presented. The difficulties associated with implementing such methodologies in large-scale parallel computing environments are discussed. Challenges include efficient load-balancing schemes, achieving optimal communication stencils, and frequent disk I/O for very large time-dependent datasets.

Eric Nielsen, Mike Park, Dana Hammond
NASA Langley
eric.j.nielsen@nasa.gov, mike.park@nasa.gov, dana.hammond@nasa.gov

MS44

Interactive Visualization from a Live Unstructured Grid CFD Simulation at 160k Cores

Scalability and time-to-solution studies have historically been focused on the size of the problem and run time. We consider a more strict definition of "solution' whereby a live data analysis (co-visualization of either the full data or in-situ data extracts) provides continuous and reconfigurable insight into massively parallel simulations. Specifically, we used the Argonne ALCF's BlueGene/P machine with 163,840 cores tightly linked through a high-speed network to 100 visualization nodes that share 200 GPUs.

Michel Rasquin
University of Colorado Boulder
michel.rasquin@colorado.edu

Patrick Marion
Kitware Inc.
pat.marion@kitware.com

Venkatram Vishwanath, Ray Loy
Argonne Leadership Computing Facility, Argonne National Lab.
venkatv@mcs.anl.gov, rloy@alcf.anl.gov

Andrew Bauer
Kitware Inc.
andy.bauer@kitware.com

Benjamin Matthews, Min Zhou, Onkar Sahni, Jing Fu, Ning Liu
Rensselaer Polytechnic Institute
matthb2@scorec.rpi.edu, zhoum@scorec.rpi.edu, sahni@scorec.rpi.edu, fuj@cs.rpi.edu, liun2@cs.rpi.edu

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

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

Mark Hereld
Argonne National Laboratory
hereld@mcs.anl.gov

Michael E. Papka
Argonne National Laboratory
University of Chicago
papka@anl.gov

Kalyan Kumaran
Argonne Leadership Computing Facility, Argonne National Lab.
kumaran@alcf.anl.gov

Berk Geveci
Kitware
berk.geveci@kitware.com

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

MS45

Parallel I/O Quo Vadis

Abstract not available at time of publication.

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

MS45
The Hdf5 Utility Toolkit

Particle-based simulations running on large high-performance computing systems can generate an enormous amount of data for post-processing and analysis. Achieving high-performance I/O for this data, effectively managing it on disk, and interfacing it with analysis and visualization tools can be challenging, especially for domain scientists who do not have I/O and data management expertise. We present the H5hut library, an implementation of several data models for particle-based simulations that encapsulates the complexity of HDF5 and is simple to use, yet does not compromise performance.

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

Mark Howison
Brown University
mhowison@brown.edu

Wes Bethel
Lawrence Berkeley National Laboratory
ewbethel@lbl.gov

Achim Gsell
PSI
achim.gsell@psi.ch

Bendikt Oswald
Paul Scherrer Institut
benedikt.oswald@psi.ch

Mr Prabhat
Lawrence Berkeley National Laboratory
prabhat@lbl.gov

MS45
Parallel HDF5 Performance - Future Plans

This session will provide a forum for HDF5 developers and users to interact. HDF5 developers will describe the current status of HDF5 and discuss future plans. Ample time will be allowed for questions and discussion, and direct users of HDF5 technologies, as well of users of software built on top of HDF5, will be encouraged to share their successes, challenges, and requests.

Quincey Koziol
HDF5 Group
koziol@hdfgroup.org

MS45
H5hut and Large Scale Particle Accelerator Simualtions

Abstract not available at time of publication.

Robert D. Ryne
LBNL
RDRyne@lbl.gov

MS46
Comparison of Unstructured Mesh Implementations in MPI, UPC and Chapel

The use of unstructured meshes is commom in certain scientific computing applications, such as computation fluid dynamics. These structures present design challenges to parallel application programmers. Irregular data access patterns, often through several levels of array indirection, along with the need to decompose data sets in a distributed environment, can lead to convoluted codes with performance problems due to poor use of the memory hierarchy. This study presents a comparison of three implementations of an unstructured mesh kernel: 1) in MPI/OpenMP 2) in UPC and 3) in Chapel. These implementations are compared in terms of readability, performance, and scalability.

Megan Cason
U.S. Army Engineer Research and Development Center
Scientific Computing Research Center
megan.l.cason@gmail.com

MS46
Performance Evaluation using Autonomous Solver Selection for Large-scale Flow Simulations

Solving linear systems is often the performance bottleneck of many large-scale scientific applications. Finding the best combination of solver, preconditioner, and input settings for a simulation requires knowledge and effort. We use machine-learning techniques to predict the best combination at each time-step during the simulation. A software package capable of plugging to application codes and selecting the best solver and preconditioner has been developed. Experimental results from solving symmetric and unsymmetric linear systems are presented.

Jing-Ru C. Cheng
Scientific Computing Research Center
U.S. Army Engineer Research and Development Center (ERDC)
Ruth.C.Cheng@usace.army.mil

Paul Eller
U.S. Army Engineer Research and Development Center
paul.r.eller@usace.army.mil

MS46
Implementing Network Analysis on GPUs

Network analysis is an active area of research with applications in various domains including social sciences, and bioinformatics. Most networks are very large and therefore parallel algorithms are essential for their efficient analysis. In recent years GPUs have become an important platform for high performance computing. In this talk we will discuss the relative merits of implementing network analysis algorithms in GPUs, with respect to implementing them on multicores, using examples from community detection algorithms.

Naresh Chinni
University of Nebraska at Omaha
nchinni@unomaha.edu

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

MS47

Optimizing Seismic Depth Imaging on a Hybrid CPU-GPU Cluster

Depth Imaging technology based on the solution of wave equation on regular grid are good candidates to take advantage of the many cores design of GPUs. Optimizing Seismic depth imaging codes on CPU-GPU cluster requires effort in porting actual codes taking care of the SIMT (single instruction multi threads), the different memory hierarchy and also the introduction an other level of data communications between host (CPU) and device (GPU). In this presentation we discuss about the experience we have in TOTAL in porting and optimizing Seismic code on CPU-GPU clusters. We'll also address the need of standard tools and programming models.

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

MS47

An Eigenvalue Solver using a Linear Algebra Framework for Multi-core and Accelerated Petascale Supercomputers

We have designed a basic linear algebra framework in order to express in a generic way classical Krylov solvers. This framework can be used etiher on massively homogeneous parallel machines, or on multicore architectures using hybrid MPI/OpenMP apraoch or on hybrid CPU-GPU clusters. We will present in this paper comparative performances of typical eigenvalue solver, namely ERAM, onto two different petascale class machines (Hopper and Curie) depending on the architecture, BLAS library and the number of threads per node used.

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

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

Jerome Dubois
CEA-DEN
jerome.dubois@cea.fr

Laurine Decobert
CNRS/LIFL
laurine.decobert@polytech-lille.net

MS47

In Search of A More Sustainable Approach to Implement Scalable Numerical Kernels

Numerical libraries do not offer a single programming abstraction to implement both high performance single thread and multi-layer parallelism. Libraries have provided easy-to-use interfaces and facilitated code reusability but without a better programming model codes will always depend on the hardware. We present two Restarted Arnoldi implementations that use SLEPc/PETSc. One uses the original libraries while the other uses SLEPc/PETSc with YML replacing MPI. We present the programming model, our implementations and results.

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

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

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

MS47

Implementation of FEM Application on GPU

In this talk, we describe the implementation of FEM application on GPU. We are focusing GPU implementation of CG solver and matrix assembling. Good performance was observed in the TeslaC2050 for our current implementations. And also we are trying to implement them on multi GPUs environment. Whole GPU implementation of FEM is considered for future work.

Satoshi Ohshima, Masae Hayashi
The University of Tokyo
Information Technology Center
ohshima@cc.u-tokyo.ac.jp, masae@cc.u-tokyo.ac.jp

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

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

MS48

Miniapps: Vehicles for CoDesign

Computing is undergoing rapid change. Many elements of architecture, system software, compilers, libraries and applications are candidates for re-design. Mini-applications (miniapps) are small programs exhibiting one performance element of a larger application. Miniapps are small, self-contained and open source. Thus, they facilitate rapid understanding, rewriting, collaboration and performance measurement. In this presentation we characterize our experience with miniapps, show how they have been used successfully so far and provide historical context for future work.

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

MS48

Utilization of Proxies in the CoCoMANS Computational Co-design Process

CoCoMANS (Computational Co-design for Multiscale Applications in the Natural Sciences) is a new computational co-design research project at Los Alamos National Laboratory. In this project we are simultaneously evolving the four corners of science, methods, software, and hardware

in an integrated computational co-design process. We are targeting multiscale applications in plasma physics, materials science, and climate. Our goal is to define, develop, and employ a computational co-design process, encompassing at one end, non-trivial computational physics problems, and at the other a range of current, emerging, and to some degree hypothetical, hardware architectures. Our co-design method is one of empirical, test- and benchmark-driven experimentation. We will develop and implement an experimental framework, along with supporting documentation and tracking capabilities, to evaluate and test multiple numerical methods, along with evolving implementations of four types of codes (or proxies): micro-benchmarks, kernels, mini-apps, and compact apps. In this talk we will discuss the evolving roles of these various levels of proxies within the CoCoMANS project.

Dana Knoll, Allen McPherson, Patrick McCormick
Los Alamos National Laboratory
nol@lanl.gov, mcpherson@lanl.gov, pat@lanl.gov

MS48

Overview of Co-Design Efforts at Los Alamos

Computational co-design is an vital part of the push towards exa-scale computing at Los Alamos National Laboratory. This session will provide an overview of multiple co-design efforts currently underway at Los Alamos. We give a detailed overview of each project along with any information on recent progress. We also describe where each project fits into the search for an overall strategy for co-design at Los Alamos. Finally, we will discuss the past and future of the Los Alamos Summer Co-Design School.

Allen McPherson
Los Alamos National Laboratory
mcpherson@lanl.gov

MS48

Using Mini-apps to Explore Multiphysics Codes at Exascale

Exascale machines will be a significant departure from traditional architectures for which multiphysics simulations codes have been developed. While the size and complexity of these codes preclude direct study of their algorithms, extraction of the most important characteristics into mini-apps provides a venue for doing so. In this presentation, we describe the methodology for using mini-apps to research algorithmic changes needed for exascale multiphysics simulation. *Work performed under the auspices of the U. S. Department of Energy by the Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Charles H. Still
Lawrence Livermore National Laboratory
still1@llnl.gov

MS49

Hybrid Solvers or Various Ways to Efficiently Combine Direct and Iterative Schemes toward Parallel Scalable Linear Solvers

Sparse hybrid solvers are a trade-off between direct and iterative methods; direct methods bring numerical robustness while iterative methods alleviate the computational complexity and memory usage. In this talk we will present and compare various approaches that combine such schemes. Experimental results arising from real-life applications and processed with PDSLin and MaPHyS hybrid solvers will illustrate our discussion.

Emmanuel Agullo
INRIA
emmanuel.agullo@inria.fr

Luc Giraud
INRIA
Toulouse, France
luc.giraud@inria.fr

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

Azzam Haidar
Department of Electrical Engineering and Computer Science
University of Tennessee, Knoxville
haidar@utk.edu

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

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

Jean Roman
INRIA
Jean.Roman@inria.fr

Ichitaro Yamazaki
UTK
ic.yamazaki@gmail.com

MS49

Fine Grain Scheduling for Sparse Solver on Many-core Architectures

The emergence of many-cores architectures introduces variations in computation costs, which makes precise cost models hard to realize. Static schedulers based on cost models, like the one used in the sparse direct solver PaStiX, are no longer adapted. We describe the dynamic scheduler developed for the super-nodal method of PaStiX to correct the imperfections of the static model. The solution presented exploit the elimination tree of the problem to keep the data locality during the execution.

Mathieu Faverge
U. of Tennessee - Knoxville
mfaverge@eecs.utk.edu

Pierre Ramet
LaBRI, Univ Bordeaux, France
ramet@labri.fr

MS49

A Parallel Black Box Multifrontal Preconditioner

that Exploits a Low-rank Structure

We present a black box preconditioner based on an approximate factorization. Our starting point is a multifrontal approach to a sparse LU factorization, which amounts to operate on a sequence of dense matrices. For problems arising in PDE applications these matrices are known to exhibit a low-rank structure. We exploit this structure in a black-box fashion to reduce memory requirements, computation and communication costs. Numerical experiments are presented to demonstrate the potentialities of the approach.

Artem Napov
Lawrence Berkeley National Lab
anapov@lbl.gov

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

Ming Gu
University of California, Berkeley
Mathematics Department
mgu@math.berkeley.edu

MS49

Performance of a Parallel Direct Solver Superlu_dist on Multicore Clusters

SuperLU_DIST is a parallel direct solver capable of solving a sparse general linear system of equations with millions of unknowns from real 3D applications on a distributed memory computer. However, most of computers are now based on multicore architectures, and the performance of SuperLU_DIST may not be ideal on such computers. In this talk, we study the performance of SuperLU_DIST on multicore clusters, and discuss techniques to enhance its performance.

Ichitaro Yamazaki
UTK
ic.yamazaki@gmail.com

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

MS50

Communication-Avoiding Nonsymmetric Eigensolver using Spectral Divide Conquer

Algorithms have two costs: arithmetic and communication. Conventional algorithms for eigenvalue problems perform asymptotically more communication than proven lower bounds require. In this paper we present parallel and sequential algorithms for the generalized nonsymmetric eigenvalue problem that use spectral divide-and-conquer and do attain these lower bounds, and analyze their convergence and communication costs.

Grey Ballard
UC Berkeley
ballard@cs.berkeley.edu

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

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

MS50

A Step Toward Scalable Eigenvalue and Singular Value Solver

In this talk, we present a new implementation in reducing a dense matrix to Tri/Bi-diagonal form, which is the preprocessing step toward solving Eigenvalue and SVD problems. The challenging trade-off between algorithmic performance and task granularity has been tackled through a grouping technique, which consists of aggregating fine-grained and memory-aware computational tasks. A dynamic runtime environment system then schedules the different tasks in an out-of-order fashion. The performance for the reduction reported is unprecedented.

Azzam Haidar
Department of Electrical Engineering and Computer Science
University of Tennessee, Knoxville
haidar@utk.edu

Hatem Ltaief
KAUST Supercomputing Laboratory
Thuwal, KSA
hatem.Ltaief@kaust.edu.sa

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

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

MS50

Steps toward a Sparse Scalable Singular Value Solver

The author has implemented an algorithm for reducing sparse matrices to small band form by block Householder operations. Random initial transformations give reliability in determining singular values. Banded algorithms from LAPACK efficiently test convergence of the larger singular values. The predominant matrix operations are with flat or tall matrices, amenable to efficient parallel computations. Some current issues: how to put the pieces together to produce an algorithm with a versatile but uncluttered interface? How to choose a good block size? How does performance compare to some other available packages?

Gary W. Howell
North Carolina State University
gary_howell@ncsu.edu

MS50

Parallel Two-Stage Reduction to Hessenberg Form

- Implementation and Scalability Issues

We recently presented a blocked two-stage algorithm for reduction to Hessenberg form targeting shared-memory systems with multicore processors. Experiments on a dual quad-core system showed good performance. In this talk, we will address the scalability of the algorithm on large-scale systems with many cores. The current bottlenecks, and some possible cures, will be discussed and evaluated. We end with an outlook on future work such as implementation on massively parallel systems with distributed memory.

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

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

MS50

Parallel Multishift QR and QZ Algorithms with Advanced Deflation Strategies – Recent Progress

Key techniques used in our novel parallel QR and QZ algorithms include multi-window bulge chain chasing and distributed aggressive early deflation (AED), which enable level-3 chasing operations and improved eigenvalue convergence. Mixed MPI-OpenMP coding techniques are utilized for DM platforms with multithreaded nodes, such as multicore processors. Some recent progress including a two-level approach for performing AED in a parallel environment are presented. Application and test benchmarks confirm the superb performance of our parallel implementations.

Bo T. Kågström, Björn Adlerborn
Umeå University
Computing Science and HPC2N
bokg@cs.umu.se, adler@cs.umu.se

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

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

Meiyue Shao
Umeå University
Computing Science and HPC2N
myshao@cs.umu.se

MS51

Tuning Applications for Multicore Architectures with Hierarchical Partitioning and Local Ordering

We consider those high performance parallel applications that partition their problem to effect load balancing and to minimize inter-process communication. Using the Zoltan dynamic partitioning and ordering library, we show the benefit of partitioning to the machine hierarchy to minimize interprocess communication time, and of locally ordering the problem in an MPI process to optimize the memory access patterns of threads.

Lee Ann Riesen
Sandia National Laboratories
lriesen@sandia.gov

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

Karen D. Devine, Sivasankaran Rajamanickam
Sandia National Laboratories
kddevin@sandia.gov, srajama@sandia.gov

MS51

An Unstructured Mesh Infrastructure for Massively Parallel Adaptive Simulation

The Flexible distributed Mesh DataBase (FMDB) is an implementation of the iMeshP interoperable unstructured mesh interface focused on supporting parallel adaptive simulation on massively parallel computers being developed as part of the Frameworks, Algorithms, and Scalable Technologies for Mathematics (FASTMath) SciDAC Institute. We will present distinctive features and on-going efforts on FMDB for massively parallel adaptive simulation support including: flexible mesh partitioning with multiple parts per process, multiple communicators, global/local mesh load balancing, and migration.

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

TIng Xie, Misbah Mubarak
Rensselaer Polytechnic Institute
txie@scorec.rpi.edu, mubarm@scorec.rpi.edu

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

MS51

ParMA: Towards Massively Parallel Partitioning of Unstructured Meshes

Parallel unstructured simulations at extreme scale require that the mesh be distributed across a large number of processors with equal workload and minimum inter-part communications. The goal of ParMA it to dynamically partition unstructured meshes directly using the existing mesh adjacency information to account for multiple criteria. Results will demonstrate the ability of ParMA to dynamically rebalance large meshes (billions of mesh regions) on large core count machines (over 100,000 cores) accounting for multiple criteria.

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

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

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

MS52

Designing Next-Generation Computational Chemistry Applications Using Extensible Software Tools

Computational chemistry applications involve many different computational motifs, including dense linear algebra, N-body solvers and irregular domain-specific kernels. The computational properties of these components differ greatly, particularly regarding communication. Application developers use different runtime systems - e.g. MPI, Global Arrays, Charm++ - to build massively-parallel computational chemistry codes, but not all runtime systems and execution models are compatible with one another, hence applications are limited in how they can compose their applications. In this talk, we describe a new high-level library inspired by Global Arrays that allows for new levels of runtime interoperability and also provides more flexible and more scalable numerical tools using the developments of the FLAME project. Emphasis will be placed on the critical role of flexible, layered dense linear algebra tools in the design of next-generation quantum many-body solvers.

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

MS52

Relayering Dense Linear Algebra Libraries

Dense linear algebra libraries have traditionally been upon the Basic Linear Algebra Subprograms (BLAS). In this talk, we discuss how strict adherence to this interface as an impenetrable layer inherently obstructs optimizations. We suggest a new layering of linear algebra functionality that overcomes this. This layering has the benefit that it greatly reduces the number of kernels that need to be customized for a given architecture while facilitating optimizations across layers. This sets the stage for a modern replacement for the BLAS being pursued as part of the FLAME project.

Lee Killough
Argonne Leadership Computing Facility
killough@alcf.anl.gov

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

MS52

Smart Distributed Memory Dense Linear Algebra Libraries

In this talk we will discuss work to mechanize the process an expert follows to implement and optimize a dense linear algebra algorithm in Elemental, a modern distributed-memory library. We will explain our success with a prototype system that breaks through the layers of Elemental to automate this implementation and parallelization process. This will motivate the development of a better-layered software stack, including LAPACK and BLAS-level functions, which will enable more automatic parallel algorithm implementation.

Bryan Marker
Department of Computer Science
UT-Austin
bamarker@gmail.com

MS52

libflame: High Performance via High-level Abstractions

An ongoing effort of the FLAME research group is to provide the scientific community with a modern, object-based dense linear algebra library that provides user-friendly, high-level abstractions while still providing the kind of high performance typically associated with low-level code. We showcase the libflame library, its APIs, and functionality while highlighting recent research developments that have been incorporated into the library.

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

MS53

Parallel Algorithms for Computational Geometry Problems in High Dimensions

We present algorithms for clustering and nearest-neighbor searches on massively parallel computer architectures. We combine parallel distance and filtering operations with locality-sensitive hashing, and a novel parallel indexing tree data structure to support exact and approximate range search operations. We outline the algorithms and provide benchmarks tests.

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

Logan Moon, Daniel Long, Bo Xiao
Georgia Institute of Technology
lmoon3@gatech.edu, dlong3@gatech.edu,
boxiao33@gatech.edu

MS53

Scalable Collocation Methods using a Combination of Krylov Basis Recycling and Adjoint Enhancement

Abstract not available at time of publication.

Michael S. Eldred
Sandia National Laboratories
Optimization and Uncertainty Quantification Dept.
mseldre@sandia.gov

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

MS53

Large-scale Parallel Hessian-based Methods for

Uncertainty Quantification in Inverse Problems

We present a parallel Langevin-accelerated MCMC method for sampling high-dimensional, expensive-to-evaluate probability densities that characterize the solution to PDE-based statistical inverse problems. The method builds on previous work in Metropolized Langevin dynamics, which uses gradient information to guide the sampling in useful directions, improving acceptance probabilities and convergence rates. We extend the Langevin idea to exploit local Hessian information, leading to what is effectively a stochastic version of Newton's method. A key issue is low rank approximation of the Hessian, exploiting the compact nature of the data misfit operator for many ill-posed inverse problems. We apply the method to the Bayesian solution of a seismic inverse problem, for which we observe several orders of magnitude faster convergence over a reference blackbox MCMC method. The method scales well to large problem sizes and numbers of cores.

Tan Bui-Thanh
The University of Texas at Austin
tanbui@ices.utexas.edu

Carsten Burstedde
Institut fuer Numerische Simulation
Universitaet Bonn
burstedde@ins.uni-bonn.de

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

James R. Martin
University of Texas at Austin
Institute for Computational Engineering and Sciences
jmartin@ices.utexas.edu

Georg Stadler
University of Texas at Austin
georgst@ices.utexas.edu

MS53

Parallel Adaptive and Robust Algorithms for the Bayesian Analysis of Mathematical Models Under Uncertainty

In recent years, Bayesian model updating techniques based on measured data have been applied to many engineering and applied science problems. At the same time, parallel computational platforms are becoming increasingly more powerful and are being used more frequently by engineering and scientific communities. Bayesian techniques usually require the evaluation of multidimensional integrals related to the posterior probability density function (PDF) of uncertain model parameters. The fact that such integrals cannot be computed analytically motivates the research of stochastic simulation methods for sampling posterior PDFs. In this talk we discuss the parallelization and enhancements of the Adaptive Multilevel Stochastic Simulation Algorithm (AMSSA, from S. H. Cheung and J. L. Beck, "New Bayesian Updating Methodology for Model Validation and Robust Predictions of a Target System based on Hierarchical Subsystem Tests', *Computer Methods in Applied Mechanics and Engineering*, accepted). We also present a variety of parallel computational results on the improved AMSSA.

Ernesto E. Prudencio
Institute for Computational Engineering and Sciences
prudenci@ices.utexas.edu

Sai Hung Cheung
School of Civil and Environmental Engineering
Nanyang Technological University, Singapore
shcheung@ntu.edu.sg

MS54

Toward Efficient Parallel in Time Methods for Partial Differential Equations

An algorithmic framework for the parallelization of numerical methods for partial differential equations in the temporal direction will be presented. In practice, temporal parallelization is only attractive if the temporal parallelization has greater parallel efficiency than (additional) spatial parallelization. Hence, the focus here is on constructing methods with good parallel efficiency. The approach presented iteratively improves the solution on each time slice by applying deferred correction sweeps to a hierarchy of discretizations at different spatial and temporal resolutions. Coarse resolution problems are formulated using a time-space analog of the full approximation scheme, and different coarsening strategies depending on the spatial discretization will be presented. The parallel efficiency and speedup for PDEs in one, two and three dimensions will be presented.

Matthew Emmett
University of North Carolina at Chapel Hill
Department of Mathematics
memmett@unc.edu

Michael Minion
University of North Carolina-Chapel Hill
Mathematics Department
minion@unc.edu

MS54

Title Not Available at Time of Publication

Abstract not available at time of publication.

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

MS54

A RIDC-DD Algorithm

We propose a parallel space–time algorithm that layers time parallelization together with a parallel elliptic solver to solve time dependent partial differential equations (PDEs). The parallel elliptic solver utilizes domain decomposition to divide a spatial grid into subdomains, and applies a classic, parallel, Schwarz iteration to find consistent solutions. The high-order parallel time integrator employed belongs to the family of Revisionist Integral Deferred Correction Methods (RIDC) introduced by Christlieb, Macdonald and Ong in 2010, which allows for the small scale parallelization of solutions to initial value problems. The two established algorithms are combined in this proposed space–time algorithm to add parallel scalability. As a proof of concept, we utilize a framework involving classical Schwarz matching conditions and a fourth order RIDC method. It will be shown that the resulting Schwarz iterations can be analyzed using standard domain decom-

position analysis, and that the four Schwarz iterations per time step can be evaluated simultaneously in parallel, after initial start-up costs.

Ronald Haynes
Memorial University of Newfoundland
St. John's, NL, Canada
rhaynes74@gmail.com

Benjamin Ong
Department of Mathematics
Michigan State University
bwo@math.msu.edu

Andrew J. Christlieb
Michigan State Univerity
Department of Mathematics
andrewch@math.msu.edu

MS54

Adaptive RIDC Algorithms

In this talk, we show that stepsize control can be incorporated within a family of parallel time integrators known as Revisionist Integral Deferred Correction (RIDC), and demonstrate that these adaptive RIDC methods perform better than constant stepsize RIDC for stiff problems. The RIDC framework framework can be layered with parallel spatial solvers to generate a parallel-space time algorithm.

Benjamin Ong
Department of Mathematics
Michigan State University
bwo@math.msu.edu

Raymond J. Spiteri
University of Saskatchewan, Canada
Department of Computer Science
spiteri@cs.usask.ca

Colin Macdonald
Oxford University
macdonald@maths.ox.ac.uk

MS55

Modified Augmented Lagrangian Preconditioners for Incompressible Flow Problems

We review recent developments in preconditioning large linear systems arising from the discretization and linearization of the incompressible Navier–Stokes equations. We focus on block triangular preconditioners based on the Augmented Lagrangian formulation of the saddle point system. Some convergence results and the determination of augmentation parameters using Fourier analysis will be discussed. Numerical experiments will be presented, including the results of a parallel implementation in the Trilinos framework. This is joint work with Zhen Wang (Oak Ridge National Laboratory).

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

MS55

Scalable Physics-based Preconditioning for 3D Extended MHD

Extended MHD (XMHD) is a very challenging hyperbolic PDE system for implicit integration techniques due to the ill-conditioning introduced by fast dispersive waves. In this talk, we will describe our physics-based preconditioning approach for 3D XMHD. The method is based on a conceptual Schur-complement decomposition, which exploits the nature of the hyperbolic couplings in XMHD to produce a block diagonally dominant PDE system, well-conditioned for multilevel techniques. Numerical experiments will demonstrate the scalability of the approach.

Luis Chacon
Oak Ridge National Laboratory
chaconl@ornl.gov

MS55

Block Preconditioners for an Exact Penalty Viscoresistive MHD Formulation

We propose a block preconditioner for an exact penalty finite element formulation of the stationary viscoresistive magnetohydrodynamics equations, implicitly enforcing the no magnetic monopoles condition. Operators arising from a block decomposition of the discrete system are analyzed from a continuous perspective to develop approximations. Commutator arguments are also employed to further simplify computation. Numerical results indicate very good behavior of the preconditioner over a range of parameters on a variety of two-dimensional test problems.

Edward G. Phillips
University of Maryland at College Park
Department of Mathematics
egphillips@math.umd.edu

Howard C. Elman
University of Maryland, College Park
elman@cs.umd.edu

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

Eric C. Cyr
Scalable Algorithms Department
Sandia National Laboratotories
eccyr@sandia.gov

MS55

Block Preconditioning of Stiff Implicit Models for Radiative Ionization in the Early Universe

Multiphysics simulations' seemingly unquenchable thirst for computing typically requires utilization of the largest parallel architectures available. Simulations of cosmological reionization during the "dark ages" of the early universe are no exception, and have pushed the limits of modern parallel architectures. In this talk I focus on scalable implicit solvers for the stiffest components of these models – radiation transport and primordial chemistry. We employ a Newton-Krylov-block Schur formulation, solving the inner system using a multigrid-preconditioned CG solver. After describing our solution approach, we present parallel scalability and simulation results.

Daniel R. Reynolds

Southern Methodist University
Mathematics Department
reynolds@smu.edu

Robert Harkness
University of California- San Diego
San Diego Supercomputer Center
harkness@sdsc.edu

Geoffrey So, Michael Norman
University of California at San Diego
gsiisg@gmail.com, mlnorman@ucsd.edu

MS56

Using Proxy Apps in the ASCR Materials Co-Design Center

Computational materials science is performed with a suite of applications that span the quantum mechanics of interatomic bonding to the continuum mechanics of engineering problems and phenomenon specific models in between. In this talk, we will review this suite and the motifs used in each of the codes with particular emphasis on how the proxy apps are used for exascale co-design, i.e., how the application 'work' translates into metrics for price, power, performance, and resiliency. Work performed under the

auspices of the U.S. DOE by LLNL under Contract DE-AC52-07NA27344.

James F. Belak
Lawrence Livermore Nat. Lab.
belak1@llnl.gov

MS56

Domain Specific Languages for Co-Design

Complex physical simulations have driven the need for exascale computing, but reaching exascale will require more power-efficient supercomputers. Heterogenous hardware offers one way to increase efficiency, but is difficult to program and lacks a unifying programming model. Abstracting problems at the level of the domain rather than hardware offers an alternative approach. In this talk we describe the design of Liszt, a domain-specific language for solving partial-differential equations on meshes. There have been many domain-specific languages and frameworks proposed for physical simulation. Liszt is unique in that it targets current and future heterogeneous platforms. We have found that designing a DSL requires a careful balance between features that allow for automatic parallelization, and features that make the language flexible.

Zach Devito
Stanford University
zdevito@stanford.edu

MS56

Early Experiences with Simulation-Based Co-Design with Proxy Applications

Making simulation codes ready for exascale platforms will require major changes in the programming model, algorithms, and runtimes. Similarly, building an exascale computer will require major changes in nearly every aspect of the hardware, from the system architecture to the packaging. To prepare for these sweeping changes, the DOE has adopted a methodology of "co-design", where both software and hardware are designed in unison. This talk shares some early experiences with using a multi-level, holistic, parallel architecture simulator (the Structural Simulation Toolkit) to practice an iterative co-design of an exascale materials code using application proxies (compact apps and skeleton apps).

Arun Rodrigues
Sandia National Laboratories
afrodri@sandia.gov

MS56

Exploring the Energy and Performance Landscape of FEM Execution

The parallel nature of FEM makes GPUs an excellent choice for accelerating FEM-based applications. However, since these applications also require data assembling, which is not readily parallelizable, it is not immediately clear whether GPU solutions will outperform CPU solutions. An in-depth study was performed to quantify energy and performance when executing the miniFE mini-application on platforms with GPU and high-performance or low-power CPUs. This talk discusses the findings and outline ways to improve these solutions.

X. Sharon Hu
Department of Computer Science and Engineering
University of Notre Dame
shu@nd.edu

Li Tang
University of Notre Dame
ltang@nd.edu

MS57

Multifrontal Sparse QR Factorization: Multicore, and GPU Work in Progress

Recent progress towards a multifrontal sparse QR factorization on a GPU will be presented. The project relies upon the hardware capabilities of recent NVIDIA GPUs, namely, the ability to run multiple different kernels simultaneously, and the event/stream model of CUDA programming. Our goal is to factorize entire subtrees of the elimination tree on the GPU, without transferring individual frontal matrices between the CPU and GPU. This transfer would be a bottleneck. Rather, we will perform the assembly operation on the GPU itself, to construct new frontal matrices from their children and from entries in the input matrix. If the single etree is too large, it will be broken in to subtrees that can be handled on the GPU, where only the root contribution block needs to be off-loaded from the GPU to the CPU.

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

MS57

Incomplete-LU Preconditioned Iterative Methods on GPUs

In this presentation we focus on parallel algorithms that are building blocks for the incomplete-LU preconditioned CG and BiCGStab iterative methods on Graphics Processing

Units (GPUs). In particular, we focus on the techniques and the associated tradeoffs used to implement sparse triangular solve and incomplete-LU factorization using the CUDA parallel programming model. Finally, numerical experiments comparing the GPU and CPU implementations are also presented.

Maxim Naumov
NVIDIA
mnaumov@nvidia.com

MS57

Preconditioners for Sparse Linear Systems on the Manycore Architectures

The amount of parallelism in the manycore architectures and good implementations available for sparse matrix-vector multiply make iterative methods for linear systems attractive in these architectures. However, extracting similar performance from constructing and applying preconditioners in the manycore architectures has remained a challenge. We address this problem with orderings and preconditioners that exploit the structure of the input matrix. We present results comparing preconditioned iterative methods in the manycore architectures with their best CPU implementations.

Sivasankaran Rajamanickam, Michael A. Heroux
Sandia National Laboratories
srajama@sandia.gov, maherou@sandia.gov

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

MS57

Performance Gains in Multifrontal for Symmetric Positive Definite Linear System on CPU-GPU Hybrid System

Multifrontal is an efficient direct method for solving large-scale linear system problems. Focusing on symmetric positive definite systems, we study how the method can be accelerated by using a graphic processing unit (NVIDIA C2070). The proposed algorithm intends to find an efficient strategy in CPU-GPU communication and CPU-GPU workload distribution to achieve the shortest running time. Analytical and numerical results will be presented to demonstrate the efficiency of the proposed algorithm.

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

Chenhan D. Yu
Department of Mathematics
National Taiwan University
b94201001@ntu.edu.tw

MS58

Multi-Dimensional Image-to-Mesh Conversion on Massively Parallel Systems

FEM simulations on four dimensional meshes have been shown to be quite accurate. In this abstract, we propose a multi-layer approach for generating four dimensional meshes from binary medical images. Our technique takes advantage of both the faster shared-memory layers when the communication cost is high, and the larger (but slower) distributed memory layers when a lot of memory allocation is required.

Panagiotis Foteinos
College of William and Mary
pfot@cs.wm.edu

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

MS58

Scalable Mesh Smoothing using Mesquite on BG/L and BG/P Platforms

We present weak scaling studies for Mesquite, a parallel mesh optimization package. For BG/L and BG/P we use local meshes with 64,000 degrees of freedom per MPI task and report nearly perfect weak scaling from eight through 64K cores (globally, over 4 billion degrees of freedom). We summarize the parallel algorithm and describe software changes needed to run at this scale. Finally we discuss mixed-mode parallelism by measuring the amount of available thread level parallelism and report preliminary results exploiting this parallelism using OpenMP + MPI.

Brian Miller, David Hysom
Lawrence Livermore National Laboratory
miller125@llnl.gov, hysom1@llnl.gov

MS58

Single-file Parallel I/O in MOAB

MOAB is a library for representing structured and unstructured mesh and associated field data. MOAB's parallel representation appears locally as a serial mesh, with additional information about shared and "ghost" or halo elements accessed through the iMesh tag and set data types. Parallel input and output must be performed to initialize a parallel mesh or dump this mesh to disk storage. In this presentation, we describe MOAB's parallel I/O capabilities. This functionality is based on input/ouput from/to a single file, based on the parallel HDF5 library. Scaling results on the IBM BG/P architecture, on up to 64k processor cores, are described.

Jason Kraftcheck
The University of Wisconsin
kraftche@cae.wisc.edu

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

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

John Magerlein
IBM Research
magerus.ibm.com

MS58

Geometry-based Parallel Mesh Generation and

Adaptation

Large scale parallel simulations deal with meshes with millions and billions of elements. As these numbers increase, it becomes necessary to generate and adapt these meshes in parallel. We overview procedures that start with a CAD model and generate these meshes in parallel. Adaptive control is supported by parallel mesh modification procedures that can perform anisotropic adaptation and maintain boundary layer meshes. These procedures are being integrated with analysis codes to support simulations running entirely in parallel.

Saurabh Tendulkar
Simmetrix Inc.
stendulkar@simmetrix.com

Mark Beall
Simmetrix, Inc.
mbeall@simmetrix.com

Rocco Nastasia
Simmetrix Inc.
nastar@simmetrix.com

MS59

Efficient I/O in Parallel Out-of-Core Multiforntal Cholesky

Solving large sparse systems of linear equations is ubiquitous in scientific computing and the size of the systems arising in various applications continues to grow. Direct methods are still the method of choice in many industrial applications due to their robustness and performance. However, as the scale of systems increases, direct linear solvers suffer from two bottlenecks: a computational bottleneck and a memory bottleneck. The computational bottleneck stems from the inherent superlinear complexity of direct solvers and the fact that CPU manufacturers are unable to improve sequential hardware performance. Therefore, linear solver must employ parallel processing in order to continue to deliver performance. The memory bottleneck stems from the fact that direct solvers require large amounts of memory due to fill-in of the factors. Although the amount of memory available in modern computers keeps increasing, the increase in memory size cannot keep pace with either the increasing number of CPUs or the memory requirements of direct solvers for growing linear system sizes. In order to address these issue, out-of-core solvers are required. Modern linear solvers for modern applications must use both techniques: parallel processing and out-of-core computations. Out-of-core computations requires a large amount of I/O, and it is important to do this I/O efficiently. The use of parallel further complicates the situation and poses unique challenges. This talk discusses these challenges and ways to overcome them. A new parallel out-of-core parallel solver is presented, alongside with numerical experiments.

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

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

MS59

CALU_PRRP: a Communication Avoiding LU Factorization Algorithm with Panel Rank Revealing Pivoting

We present the LU decomposition with panel rank revealing pivoting (LU_PRRP), an LU factorization algorithm based on a new pivoting strategy that performs Strong Rank Revealing QR on the panels to choose the pivot rows. LU_PRRP is more stable than Gaussian elimination with partial pivoting (GEPP), with a theoretical pivot growth upper bound of $(1+2b)^{\frac{n}{b}}$, where b is the size of the panel used during the factorization. For example, if the size of the panel is $b = 64$, then $(1+2b)^{n/b} = (1.079)^n \ll 2^{n-1}$, where 2^{n-1} is the pivot growth upper bound for GEPP. Our extensive numerical experiments show that the new pivoting scheme is as numerically stable as GEPP but is more resistant to pathological cases and easily beats the Kahan matrix. The LU_PRRP method only does $O(n^2)$ flops more than GEPP. We also present CALU_PRRP, a communication avoiding version of LU_PRRP that is optimal in terms of communication. Like the CALU algorithm, this algorithm is based on tournament pivoting, but is more stable than CALU in terms of worst case pivot growth.

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

Ming Gu
University of California, Berkeley
Mathematics Department
mgu@math.berkeley.edu

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

MS59

Communication-avoiding Parallel Implementation for Strassen's Algorithm

We suggest a 2.5D parallel implementations for Strassen's and Strassen-like algorithms, that attain the communication cost lower bounds. The implementations follow the finding of (Ballard, Demmel, Holtz, and Schwartz, 2011), that it is always better, communication-wise, to avoid breaking blocks created by the algorithms. To complete the implementation scheme, we spell out how to schedule the work of the P processors, gaining a speedup factor of $\Theta(P)$ in the arithmetic count, in the bandwidth cost, and in the latency cost.

Grey Ballard
UC Berkeley
ballard@cs.berkeley.edu

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

Olga Holtz
University of California, Berkeley
Technische Universitat Berlin
holtz@math.berkeley.edu

Eran Rom

IBM Research Tel-Aviv
eranr@il.ibm.com

Oded Schwartz
TU Berlin
oded.schwartz@gmail.com

MS59

A Communication Avoiding Symmetric Indefinite Factorization

We describe a communication-avoiding algorithm for factoring symmetric indefinite matrices. The algorithm works in two phases. In the first phase, the input matrix A is factored into a product $A = PLTL^TP^T$ where P is a permutation matrix, L is unit lower triangular, and T is banded. In the second phase, we factor T using one of several known algorithms. The first phase is a blocked version of Aasen's algorithm, in which we use a communication-avoiding LU algorithm to factor block columns. The algorithm is an improvement over the Shklarski-Toledo factorization which is efficient in terms of communication volume but not in terms of the number of messages.

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

Inon Peled, Alex Druinsky
Tel-Aviv University
inon.peled@gmail.com, alexdrui@post.tau.ac.il

MS60

Programming Clusters with StarSs

StarSs is a task-based programming model that aims to provide portability and flexibility to sequential codes while the performance is achieved by the dynamic exploitation of the parallelism at task level. The talk will target the programming of clusters of multicores and GPUs with StarSs. We propose MPI/SMPSs that parallelizes the code at node level and enables the overlap of communication and computation. Also, we will describe the pure StarSs programming model for clusters.

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

MS60

Essential Elements of a Superscalar Scheduling API for Numerical Libraries

Parallel Linear Algebra Software for Multicore Architectures (PLASMA) is a robust numerical library providing an array of routines for solving linear systems of equations, least square problems, singular value and eigenvalue problems. PLASMA relies completely on a superscalar task scheduler, QUARK. The objective of QUARK is to provide a proof-of-concept implementation of an API, suitable to support a numerical library. This talk highlights critical elements of such an API.

Jakub Kurzak
University of Tennessee Knoxville
kurzak@eecs.utk.edu

Asim YarKhan
Department of Electrical Engineering and Computer Science
University of Tennessee, Knoxville
yarkhan@eecs.utk.edu

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

MS60

Scalable Computing with Multiprocessor Tasks

Task-based approaches usually allow a task specification that decouples the programming from the scheduling and mapping to the execution resources of a specific hardware platform. For a distributed address space, parallel tasks offer the additional advantage of structuring the program into communication domains, which can be exploited to reduce the communication overhead. This talk gives an overview of programming approaches for parallel tasks with different characteristics.

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

Gudula Rünger
Chemnitz University of Technology
ruenger@informatik.tu-chemnitz.de

MS60

Managing Dependencies in a Task Parallel Framework

This talk will give an introduction to dependency-aware task-based programming models. By taking each task's individual dependencies into account, we can obtain fine grained synchronization and flexibility in scheduling. Also, the user effort for parallelization is low; it is enough to annotate each task's memory accesses. Focus will be on how to specify and represent dependencies. We will present our solution using data versioning, and how this conveniently allows for reordering of associative operations.

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

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

MS61

On a Class of Phase Field Models for Irradiated Materials

We give an overview of phase field models for microstructure evolution in irradiated materials. In particular, we highlight a set of important issues related to the mathematical formalism of these models, which include the thermodynamic consistency, free energy functional construction and sharp interface limits. A set of numerical examples illustrating the suitability of these models in tackling the problem of void formation in irradiated materials will be

given.

Anter A. El-Azab
Florida State University
aelazab@fsu.edu

Michael Pernice
Idaho National Laboratory
Michael.Pernice@inl.gov

MS61

An Adaptive Finite Element Moreau-Yosida-based Solver for a Non-smooth Cahn-Hilliard Problem

An a posteriori error estimator based adaptive finite element semi-smooth Newton solver for the Cahn-Hilliard model with double obstacle free energy is proposed. For the numerical solution of the optimal control problems associated with a semi-discretization in time, an algorithm combining a Moreau-Yosida regularization technique, which handles the constraints due to the obstacle potential, with a semi-smooth-Newton method is proposed. The performance of the overall algorithm is illustrated by numerical experiments.

Michael Hintermueller
Humboldt-University of Berlin
hint@math.hu-berlin.de

Michael Hinze
Universität Hamburg
Department Mathematik
michael.hinze@uni-hamburg.de

Hicham Tber
Department of Mathematics
Faculté des Sciences et Techniques Beni-Mellal
hicham.tber@gmail.com

MS61

Large-Scale Differential Variational Inequalities for Heterogeneous Materials

Recent progress on the development of scalable DVI solvers to handle the phase-field approach for mesoscale materials modeling is described. We have developed semi-smooth and reduced-space active set VI solvers in PETSc, leveraging experience by optimization community in TAO. We have achieved mesh-independent convergence rates for Allen-Cahn system using Schur complement preconditioner with algebraic multigrid, and also have validated DVI modeling approach against baseline results from the material scientists for void formation and radiation damage.

Jungho Lee
Argonne National Laboratory
Mathematics and Computer Science Division
julee@mcs.anl.gov

Shrirang Abhyankar
Illinois Institue of Technology
abhyshr@mcs.anl.gov

Mihai Anitescu, Lois Curfman McInnes, Todd Munson
Argonne National Laboratory
Mathematics and Computer Science Division
anitescu@mcs.anl.gov, curfman@mcs.anl.gov,
tmunson@mcs.anl.gov

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

Lei Wang
Argonne National Laboratory
lwang@mcs.anl.gov

Anter A. El-Azab
Florida State University
aelazab@fsu.edu

MS61

An Object-oriented Finite Element Framework for Multiphysics Phase Field Simulations

Phase field models effectively predict microstructure evolution. To facilitate the creation of new phase field models, we have created the MARMOT framework based on INLs MOOSE finite element framework. With MARMOT, 2- or 3-D phase field models can be developed with minimum coding by taking advantage of object-oriented architecture. MARMOT-based models are easily coupled to additional physics and have access to mesh and time step adaptivity. Several multiphysics examples will be shown.

Michael Tonks, Derek Gaston
Idaho National Laboratory
michael.tonks@inl.gov, derek.gaston@inl.gov

Paul Millett
Idaho National Laboratory
Nuclear Fuels and Materials
paul.millett@inl.gov

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

David Andrs
Idaho National Laboratory
david.andrs@inl.gov

MS62

Commuting Block Preconditioned Splitting with Multigrid within the Same Code Base

Implicit solution methods for multiphysics problems are either monolithic (treating the coupled problem directly) or split (solving reduced systems independently). Software for monolithic multigrid is typically more intrusive, requiring a monolithic hierarchy instead of independent hierarchies, but the required number of iterations may be smaller due to coupling on all levels. We describe generic software support for "multigrid inside splitting' and "splitting inside multigrid' and provide direct comparisons for multiphysics applications in glaciology and geodynamics.

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

Matthew G. Knepley
University of Chicago
knepley@gmail.com

Dave May
ETH Zurich
dave.may@erdw.ethz.ch

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

MS62

Block Preconditioners for Fully Implicit Atmospheric Climate Simulation in CAM-SE

We discuss the development of block preconditioners in an effort to reduce computational costs associated with fully implicit time integration of atmospheric climate models within CAM-SE. We construct a fully implicit framework based on the shallow water equations and view the subsidiary linear system as a block matrix. Preconditioners are derived based on approximate block factorization.

P. Aaron Lott
Lawrence Livermore National Laboratory
Aaron.Lott@llnl.gov

Kate Evans
Oak Ridge National Laboratory
evanskj@ornl.gov

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

MS62

Physics-based Preconditioners for Ocean Simulation

We examine physics-based preconditioners for ocean simulation based on barotropic-baroclinic splitting. Physics-based preconditioning is a highly successful approach for multiple time scale problems where an accurate simulation is desired on the dynamical time scale. Our approach is a fully implicit, fully coupled time integration of the momentum and continuity equations of ocean dynamics; thus reducing errors and increasing stability due to traditional operator splitting. The nonlinear system is solved via preconditioned Jacobian-free Newton-Krylov, where we reformulate traditional barotropic-baroclinic splitting as a preconditioner. Thus the desired solution is timestep converged with timesteps on the order of the dynamical timescale. We provide numerical examples to support the study and compare to explicit methods and methods based on operator splitting.

Chris Newman
Los Alamos National Lab
cnewman@lanl.gov

MS62

Preconditioning of Tightly Coupled Hydromechanical Systems for Reservoir Simulation

Given the increasingly urgent need to reduce atmospheric concentrations of carbon dioxide, geologic carbon sequestration has emerged as a promising technology for reducing greenhouse gas emissions at an industrial scale. In the most common configuration, captured CO_2 is injected into a deep saline aquifer that is overlain by an impermeable caprock. The injection process, however, creates large overpressures in the reservoir that can lead to significant geomechanical deformations. An understanding of the coupled hydromechanical processes taking place is therefore essential for successful operation and long-term maintenance. To this end, we consider a mixed finite element formulation of multiphase fluid flow through deformable porous media. At each timestep, the resulting discrete, nonlinear system is solved using a Newton-Krylov algorithm. A key challenge that has prevented the wide-scale adoption of tightly-coupled models, however, is the lack of robust and scalable preconditioning strategies for this class of problems. In light of this need, we discuss a block-structured approach that leads to both scalable behavior and robustness with respect to large jumps in material coefficients, a pervasive feature of natural reservoirs.

Joshua A. White
Lawrence Livermore National Laboratory
jawhite@llnl.gov

MS63

Quantitative Performance Comparison of Mantevo Miniapps and the SPEC CPU2006 Benchmark Suite on a Contemporary X86 Architecture

It has long been contended that the SPEC CPU benchmarks are not representative of the behavior of realistic scientific applications. In this talk, we present data comparing the Mantevo miniapps and SPEC benchmark suites with a diverse set of performance metrics including instruction mix, data dependence distances, memory stride and reuse distributions, spatial and temporal locality, and working set size. Preliminary data shows a significant difference in behavior of these two application suites.

Jeanine Cook, Waleed Alkohlani
New Mexico State University
jcook@nmsu.edu, wkohlani@nmsu.edu

MS63

A New Multigroup Thermal Radiation Diffusion Miniapplication

In high energy density physics simulations, multigroup radiation diffusion can dominate the total run time and memory usage. A new mini-application implements a finite-element discretization for the coupled, implicit diffusion solves. The amount of matrix data needed to solve the tens to hundreds of diffusion equations can be very large, but the related nature of the group matrices presents new opportunities for optimization not possible when optimizing a single implicit solve.[1]

Thomas A. Brunner
Lawrence Livermore National Laboratory
brunner6@llnl.gov

MS63

Automated Extraction of Skeleton Apps from Apps

Abstract not available at time of publication.

Daniel J. Quinlan
Lawrence Livermore Nat'l Lab.

[1]This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. LLNL-ABS-494432

Center for Applied Scientific Computing
dquinlan@llnl.gov

MS63
Los Alamos Summer Co-Design School

During the Summer of 2011 Los Alamos National Laboratory held its first Summer Co-Design School. This session describes the purpose and operation of the school along with the assigned problem space described below. The Boltzmann transport equation provides high fidelity simulation of a diverse range of kinetic systems. Classical methods to solve the equation are computationally and data intensive. Existing stochastic solutions to the Boltzmann equation map well to traditional large multi-core and many-node architectures but suffer performance degradations on graphics processing units (GPUs) due to heavy thread divergence. We present a a novel algorithm, Quasi-Diffusion Accelerated Monte Carlo (QDA-MC), which improves performance on heterogeneous CPU/GPU architectures.

Jeffrey A. Willert
North Carolina State University
jawiller@ncsu.edu

MS64
Scalable Solution using PETSc for Multiple Timescale Electrical Power Grid Dynamics Simulation

The solution and analysis of multiscale, multiphysics models presents one of the major hurdles for modern applied and computational mathematics. Motivated by various applications, e.g., coupling between flow and reaction variables in groundwater modeling and coupling between core and edge models in fusion simulations, we have developed capabilities in PETSc library to ease the implementation of multiphysics simulations, and can be utilized for multiscale simulations too. Our approach includes (1) simple user specification of multimodels, (2) abstractions for managing the composition of solvers, and (3) flexible solver options and efficient solution for coupled multimodels. In this talk we present our latest progress and demonstrate our approach through a multiple time-scale simulation for studying electrical power grid dynamics.

Shrirang Abhyanka
Argonne National Laboratory
abhyshr@mcs.anl.gov

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

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

Alexander Flueck
Illinois Institute of Technology
flueck@iit.edu

MS64
Threading I/O Strategy for Checkpoint/restart of a Massively Parallel Electromagnetic System

Checkpointing is an effective approach for failure recovery and post-processing purpose. However, this approach could result in heavy I/O load and may cause an I/O bottleneck on a massively parallel system. In this talk, we present our application-level checkpoint approaches using MPI-IO collective routines, application data aggregation model and a potential threading framework to do asynchronous I/O. We discuss some production performance improvement of a massively parallel electromagnetic solver (NekCEM) on the IBM BG/P at ANL on up to 64K processors.

Jing Fu
Rensselaer Polytechnic Institute
fuj@cs.rpi.edu

MiSun Min
Argonne National Laboratory
Mathematics and Computer Science Division
mmin@mcs.anl.gov

Robert Latham
Argonne National Laboratory
robl@mcs.anl.gov

Christopher Carothers
Rensselaer Polytechnic Institute
chrisc@rpi.cs.edu

MS64
Impact of Kernel-assisted MPI Communication over Scientific Applications

Collective communication is one of the most important message passing concepts, enabling parallel applications to express complex communication patterns while allowing the underlying MPI to provide efficient implementations to minimize the cost of the data movements. However, with the increase in the heterogeneity inside the nodes, more specifically the memory hierarchies, harnessing the maximum compute capabilities becomes increasingly difficult. This talk will investigate the impact of kernel-assisted MPI communication, over several scientific applications: 1) Car-Parrinello molecular dynamics(CPMD), a chemical molecular dynamics application, 2) FFTW, a Discrete Fourier Transform (DFT), and 3) ASP, a parallel all-pairs-shortest-path graph application. By focusing on the usage of Message Passing Interface (MPI), we found the communication characteristics and patterns of each application. Our experiments indicate that the quality of the MPI communication implementation plays a critical role on the overall application performance.

Teng Ma, Aurelien Bouteiller
University of Tennessee
tma@eecs.utk.edu, bouteill@eecs.utk.edu

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

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

MS64
Composable Libraries for Parallel Programming

MPI's support for parallel libraries is based on partitioning

the underlying system across ranks and time. As applications grow to encompass more components and work more dynamically, partitioning will become progressively less productive and scalable. The Charm++ parallel programming system offers an alternative in its support for asynchronous adaptive overlap of independent components. In this talk, I describe a few libraries I have worked on in Charm++, ranging across parallel I/O, dense linear algebra, and shared distributed-memory arrays. I will discuss application usage, performance results, and some future directions that this body of work points to.

Phil Miller, Laxmikant V Kale
University of Illinois at Urbana-Champaign
mille121@illinois.edu, kale@illinois.edu

PP1

Hybrid Mpi/openmp Implementation on the General Utility Lattice Program (gulp) Code

A hybrid MPI/OpenMP programming approach is implemented on a computational chemistry code, General Utility Lattice Program (GULP), with the development process and performance testing all carried out on a multi-core cluster. In this study, we observed a decent improvement in the hybrid model overall timing compared to the pure MPI code. Our performance analysis shows that the communication timing percentage of the hybrid model is much lower than that of the MPI. Ultimately, this work confirms that a scientific application on a shared memory cluster stands to gain benefit by mixing shared memory with message passing programming.

Samar A. Aseeri
Computational Scientist at thE
KAUST Supercomputing Laboratory
samar.aseeri@kaust.edu.sa

Dodi Heryadi
Computational Scientist at the KAUST
Supercomputing Laboratory
dodi.herydi@kaust.edu.sa

PP1

Txpkl: A Preconditioner Library Based on Sparsification of High-Order Stiffness Matrices

High-order finite elements are becoming increasingly common in scientific simulation codes due to their high degree of accuracy per computational cost. Unfortunately they lead to denser stiffness matrices that subsequently lead to more expensive linear solves. We present a novel finite element library that can be used to generate sparser stiffness matrices for use in the preconditioner. Parallel results for Krylov methods using these preconditioners will be examined.

Travis M. Austin
Tech-X Corporation
austin@txcorp.com

Ben Jamroz
Tech-X
jamroz@txcorp.com

Chetan Jhurani
Tech-X Corporation
jhurani@txcorp.com

PP1

Towards a Parallel High-Performance Search Engine

Text mining has long been merely an application for parallel matrix computations. We attempt to bridge the long-standing gap between large-scale text analysis and the users. To this end, we are developing a parallel high-performance text indexing service as an end-to-end application involving dimensionality reduction, document clustering and fast nearest-neighbor search. We outline algorithms and design of our prototype, report experimental results and describe future technical challenges.

Tobias Berka
Department of Computer Sciences
University of Salzburg
tberka@cosy.sbg.ac.at

Marian Vajteršic
Department of Computer Sciences, University of Salzburg
Institute of Informatics, Slovak Academy of Sciences
marian@cosy.sbg.ac.at

PP1

Scalable Mhd Simulation on Multi-Core and Gpu Clusters

continuum-level 3D simulation for magnetically confined fusion plasma results in large scale, stiff systems of time-dependent partial differential equation. They can be simplified using a convenient gauge and a stream function representation. Considering the timestep size limited by fast wave, we advance these equations in time with semi-implicit method. Discretized by finite difference along magnetic field line, this system of equations can be reformed and categorized into the following three types of 2D equations

$$\nabla_\perp u = f(u)$$
$$(\nabla_\perp + \lambda)u = f(u)$$
$$(\nabla_\perp - \lambda)u = f(u)$$

subjected to Dirichlet or Neumann boundary conditions. $\perp$ represents the operations on each 2D plane, which are perpendicular to magnetic field line. C^0 triangular finite element method on these 2D planes is employed for discretization. This results in a very large system of linear equations, and presents challenges for the nowadays state-of-art supercomputing. Solving these linear systems takes most of cpu times. The key component of achieve efficient and scalable solving is to construct an effective preconditioner for iterative linear solver. There are several linear solver and preconditoner packages available now on distributed systems. The widely used are PETSc, SuperLU, Hypre, Mumps. Each has its own advantages and disadvantages. Skillfully composing these solvers and preconditions, we can achieve scalable linear solving. Scalable Eigenvalue calculation and restart/output IO are important issues we will discuss here as well.

Jin Chen
Princeton Plasma Physics Laboratory
jchen@pppl.gov

PP1

Out-of-Core Algorithms for Dense Matrix Factor-

ization on Gpgpu

This work adapts out-of-core algorithms for dense matrix factorization for solving large problems on General Purpose Graphics Processing Unit (GPGPU). Dense matrix computations arises in diverse applications such as in modeling the response and heating of fusion plasma to radio frequency (RF) waves, modeling radiation heat transfer, boundary element method, and large scale linear least squares problems. The MAGMA library (`http://icl.cs.utk.edu/magma/`) achieves very high performance on GPGPU for dense matrix computations. However, the largest problem size is limited to the amount of local device memory on GPGPU. This work adapts out-of-core algorithms for solving large problems on GPGPU so that a matrix of size say 10 GBytes can still be factored on GPGPU with only 1 GBytes of device memory. The dense factorization using a column-panel oriented left-looking algorithm and builds upon the in-core MAGMA library. Even with the overhead for repeated data movement between the host and device memory, the software achieves over 80% performance of MAGMA.

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

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

Amy Huang, Watson Wu
Chinese University of Hong Kong
huanghang1109@hotmail.com, wuyisince927@gmail.com

PP1

Analyzing Massive Networks with GraphCT

Many open source software packages for graph analysis do not scale on multicore servers with data of moderate size. GraphCT is a portable, scalable framework for analyzing complex queries on massive graphs. Running on both the Cray XMT and commodity workstations, GraphCT is a testbed for algorithm development in areas such as parallel agglomerative clustering. We can efficiently process a series of complex analytics on a power-law graph with billions of edges.

David Ediger, Jason Riedy
Georgia Institute of Technology
School of Computational Science and Engineering
dediger@gatech.edu, jason.riedy@cc.gatech.edu

Henning Meyerhenke
Karlsruhe Institute of Technology
Institute of Theoretical Informatics
meyerhenke@kit.edu

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

PP1

An Agent-Based Simulation of a Heavy Equipment Rental Process

Heavy equipment rental houses must carefully balance the size, mix, and age of a fleet in order to maximize the profit, internal rate of return, and utilization. By modeling individual machines and aggregating the results at the end, a highly parallel, asynchronous simulation was built by assigning each machine to a GPU core. This allowed the agent-based simulation to be run across multiple GPUs on a multi-core system without the overhead of time synchronization.

Jj Lay
Thompson Machinery Commerce Corporation and
Middle Tennessee State University
jj.lay@thompsonmachinery.com

PP1

Parallel Community Detection Algorithm in Large Scale Social Networks

Social networking sites, such as YouTube, Flickr, Livejournal and Orkut, are getting popular with the increasing number of Internet users. Consequently, community detection on social network becomes more important and have been studied by many researchers. However, most of the proposed algorithms are based on a single machine which limits studying large scale social networks. In this paper, we are proposing a parallel community detection algorithm and analyze large scale social network sites.

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

PP1

Many-Core Memory Hierarchies and Parallel Graph Analysis

It is generally accepted that sparse graph analysis algorithms have poor cache performance due to low spatial locality and low to non-existent temporal locality. We captured memory traces from real-world multithreaded graph software and characterized their data access patterns through simulation of multicore coherent caches. We find a wide variety of data access patterns among common graph algorithms. We use this simulation framework to explore the performance of next-generation memory hierarchies and data structures.

Robert C. Mccoll
Georgia Institute of Technology
rmccoll3@gatech.edu

David Ediger
Georgia Institute of Technology
School of Computational Science and Engineering
dediger@gatech.edu

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

PP1

Sting: Software for Analysis of Spatio-Temporal In-

teraction Networks and Graphs

Current tools for analyzing graph-structured data and semantic networks focus on static graphs. Our STING package tackles analysis of streaming graphs like today's social networks and communication tools. STING maintains a massive graph under changes while coordinating analysis kernels to achieve analysis at real-world data rates. We show examples of local metrics like clustering coefficients and global metrics like connected components and agglomerative clustering. STING supports parallel Intel architectures as well as the Cray XMT.

Jason Riedy, David Ediger
Georgia Institute of Technology
School of Computational Science and Engineering
jason.riedy@cc.gatech.edu, dediger@gatech.edu

Henning Meyerhenke
Karlsruhe Institute of Technology
Institute of Theoretical Informatics
meyerhenke@kit.edu

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