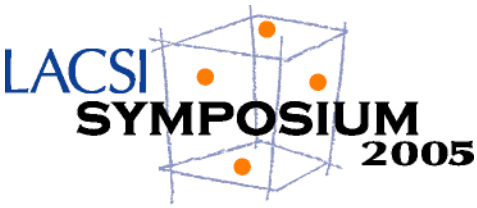

Table of Contents

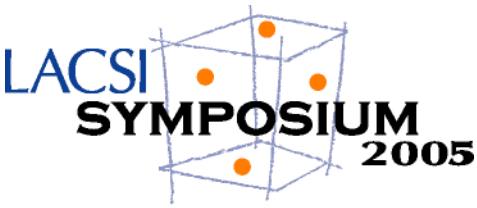

LACSI

PROCEEDINGS
Sixth LACSI Symposium
October 11-13, 2005
Santa Fe, NM USA

INTRODUCTION

The Sixth Symposium of the Los Alamos Computer Science Institute (LACSI) was held in Santa Fe, NM on October 11-13, 2005. LACSI was formed at Los Alamos National Laboratory to conduct research in computer science issues relevant to DOE's Accelerated Strategic Computing Initiative (renamed ASC). The Institute focuses on research with a longer time frame than the operational requirements of ASCI, and supports work both at LANL and at LACSI-associated universities.

The technical program for 2005 addressed the current and future state of high performance computing in several ways.

- 9 peer-reviewed technical papers were presented to the Symposium and reproduced in these proceedings.
- 11 posters were presented at the poster session.
- The keynote address, “New Architectures for a New Biology”, was delivered by David E. Shaw, D. E. Shaw Research and Development and Center for Computational Biology and Bioinformatics of Columbia University
- Two panel discussions focused on important issues for researchers and managers in high performance computing.
- Ten workshops were held focusing on technical areas of particular importance to the LACSI community.

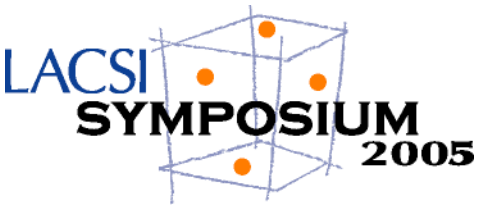

LACSI

TECHNICAL PROGRAM

The Sixth Symposium offered participants a wide variety of opportunities to discuss and learn about technical and policy issues in high-performance computing.

Tuesday, October 11: Workshop Sessions

Full Day Workshops: 9:00am – 5:30pm

High Availability & Performance Computing Workshop
Leangsuksun, Chokchai Box, Lousiana Tech University (box@latech.edu)
Author(s): Chokchai Box Leangsuksun, Louisiana Tech University; Stephen Scott, Oak Ridge National Laboratory (ORNL)

High availability & performance computing has recently become a recognized important combination to those organizations that require tremendous computing power to solve their important problems such as Energy, Climate, Fusion, Biology, & Nanotechnology. These non-trivial problems are usually characterized by massive & long running applications; therefore, Reliability, Availability, & Serviceability (RAS) management will become an increasingly paramount aspect in many computing environments. RAS management goals are to maximize uptime & therefore undoubtedly complement High End Computing (HEC) objectives by preventing performance degradation & spectrum availability. High Availability (HA) Computing has always played a critical role in commercial mission critical applications. Likewise, High Performance Computing (HPC) has equally been a significant enabler of the R&D community because of their scientific discoveries. Serviceability aims toward effective means by which corrective & preventive maintenance can be performed on a system. Higher serviceability improves availability & helps retain quality, performance, & continuity of services at expected levels. The combination of HA, Serviceability, & HPC will clearly lead to even more benefits to critical shared major HEC resource environments. This third annual WS (HAPCW2005) is a forum for the discussion of topics related to the issues affecting HAPC.

Advanced Numerical Methods for PDEs
Boyarkin, Oleg, University of Houston (boyarkin@math.uh.edu)

Participants in WS2 will discuss new developments & challenges in construction, investigation, & applications of new numerical methods & algorithms for the solution of partial differential equations relevant to LANL applications. New discretization methods for PDEs on arbitrary polyhedral meshes, their stability & accuracy, & efficient preconditioned solvers for underlying large-scale algebraic systems as well as interface reconstruction algorithms are among the major topics.

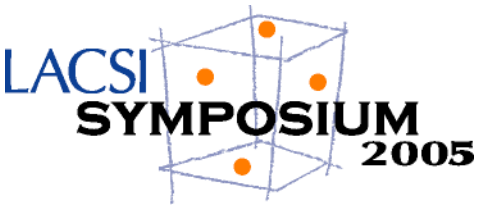

LACSI

Performance & Productivity of Extreme-Scale Parallel Systems
Hoisie, Adolfy, LANL (hoisie@lanl.gov)
Author(s): Adolfy Hoisie, PAL/LANL; Dan Reed, UNC/Institute for Renaissance Computing

In WS3 we will be concerned with the interplay across system architecture, network, applications, & system software design. The invited speakers, leaders in these fields, will cover these areas & also address the state-of-the-art in methodologies for performance analysis & optimization including benchmarking, modeling, tools development, tuning & steering, as well as metrics for productivity. We envision WS3 to be composed of 4 sessions of 3 talks each.

Models & Simulations for Large-Scale Socio-Technical Systems
Eidenbenz, Stephan J, LANL (eidenben@lanl.gov)
Author(s): James P Smith, LANL; Stephan Eidenbenz, LANL;
Gabriel Istrate, LANL; Anders Hansson, LANL; Christian Reidys, LANL

Complex socio-technical systems consist of millions of interacting physical, technological, & human /societal components. Examples of such systems include transportation systems, national commodity markets, telecommunication & computing systems including the Internet, & public healthcare systems. High-fidelity simulations capable of representing & analyzing such complex systems require the use of high performance computing platforms & tools. WS4 aims to bring together some of the leading researchers with the goal of identifying fundamental issues in designing, implementing, & using such simulations on high-performance computing architectures. Topics include the following: scalable HPC oriented design of such simulations; distributed algorithms & their implementations; & large-scale discrete event simulation systems.

High Performance Computing in Beam Physics & Astrophysics
Habib, Salman, LANL (habib@lanl.gov)
Author(s): Salman Habib, LANL; Robert Ryne, LBNL

Particle-based codes are among the most widely used high performance computing tools today, essential components of the state-of-the-art in fields, such as astrophysics & cosmology, compressible & incompressible fluid dynamics, & plasma & beam physics. Several large-scale applications are now at a threshold where they can be used as precision tools rather than as quantitative indicators of system behavior. Certain target problems in beam physics & astrophysics & cosmology have very stringent error control requirements for next-generation simulation frameworks & tools ranging from sub-percent to parts per million. Additionally, the success of major projects (e.g., the International Linear Collider) & large-scale cosmological surveys (e.g., the Joint Dark Energy Mission, Dark Energy Survey, & Large Synoptic Survey Telescope) depends on accurate & truly predictive simulations. That these projects represent a multi-billion dollar science investment further underscores the importance of high-performance simulation tools to their success.

In WS5 we will aim to bring together researchers in these fields to discuss the future challenges in high-performance simulations for beam physics & astrophysics. WS5 will enable researchers to share successful strategies that have worked in their subdisciplines & to outline the areas where more work is clearly needed. A joint strategy for attacking these problems will be a major aim of WS5.

Automatic Tuning of Whole Applications

Contact person: Kennedy, Ken, Rice University (ken@rice.edu)
Author(s): Ken Kennedy, Rice University

For many years, retargeting of applications for new architectures has been a major headache for high performance computation, requiring many person-months (or even years) of effort to retune each new architecture, & even each new model of an established architecture. Automation of this retuning process has now become a fertile area of computer science research. Most of this work is based on the strategy of using large amounts of computation time to explore a space of different variants of a loop nest, running each variant on the target architecture, & picking the best one. One example of this strategy is the Atlas system, which uses substantive amounts of computation to provide versions of a computational linear algebra kernel that are tuned in advance to different machines. If this approach can be extended more generally to components & whole programs, it would help avoid the enormous human costs involved in retargeting applications to different machines. A major research issue is how to bring the tuning time to manageable levels, given that the number of variants in a complete application can be enormous. WS6 will report on the ongoing research efforts in this area, solicit feedback from & collaboration with the application development community, & exchange ideas on future directions for this work. One specific subgoal will be to initiate an activity to develop a standard set of benchmarks for use in automatic tuning research.

Algorithm Acceleration with Reconfigurable Hardware

Contact person: Maya Gokhale (maya@lanl.gov; 505-665-9095)

Over the past 15 years, direct execution of algorithms in reconfigurable hardware has demonstrated speedup of 1-2 orders of magnitude over equivalent software. Reconfigurable Computers (RC) using Field Programmable Gate Arrays (FPGA) as processors have emerged as co-processors to augment microprocessors in workstations, clusters, & supercomputers. While RC offers remarkable opportunities for performance, research challenges abound:

- designing system architectures that balance conventional & reconfigurable processors
- developing analysis & compiler tools to automatically map
- algorithm kernels to hardware
- minimizing communications costs between hardware & software
- designing highly parallel, fine-grained computational elements for direct hardware execution
- scheduling & managing reconfigurable computing elements in large systems

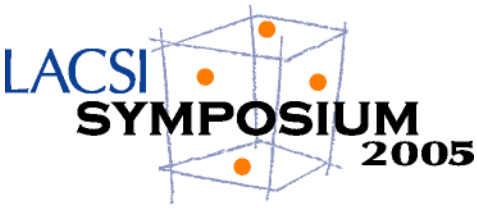

The purpose of WS7 is to discuss successes & challenges of reconfigurable supercomputing. The AM session will present introductory topics & applications; the PM session will include research topics in FPGA-based architectures, systems, tools, & future directions.

Parallel Programming with Charm++ and AMPI
Contact person: Celso L. Mendes, University of Illinois (cmendes@cs.uiuc.edu)
Author(s): Laxmikant V. Kale, University of Illinois; Celso L. Mendes, University of Illinois

Adaptive MPI (AMPI), Charm++ and the frameworks built upon them have emerged as powerful parallel programming systems in recent years. By allowing programmers to divide the computation into a large number of entities that are mapped to the available processors by an intelligent runtime system, Charm++ enables a separation of concerns between the programmers & the computing system. This approach leads to both improved programmer productivity & higher system performance. WS8 will focus on showcasing leading research in parallel processing based on Charm++ & its frameworks. Topics will include tutorial- level introduction to Charm++ & AMPI, followed by case studies of applications developed using the frameworks, as well as advances in AMPI/Charm++ technology itself. Authors & attendees will be encouraged to share their experiences & plans for the systems built upon Charm++/AMPI.

LinuxBIOS Summit
Contact person: Ron Minnich, LANL (rminnich@lanl.gov)

WS9 will include a structured set of talks & a less structured discussion period. We will explore the current status of LinuxBIOS, including presentations by vendors on how they are using or plan to use LinuxBIOS in their products. We will discuss successes as well as problems & draw lessons learned from both. We will try to determine where LinuxBIOS should be taken next & to set goals & figure out how to meet them. We plan to close by producing a consensus document on the next steps needed over the coming year.

Application Development Using Eclipse & the Parallel Tools Platform
Contact person: Watson, Gregory, LANL (gwatson@lanl.gov; 505-665-0726)

Eclipse is an extensible, open-source integrated development environment (IDE) meant to be a full-featured, commercial-quality platform for development of highly integrated software tools. Eclipse offers many features: syntax-highlighting editor, incremental code compilation, thread-aware debugger, code & class navigator, file/ project manager, interfaces to standard source control systems, & support for Java, C, C++, Fortran, & other languages. The Parallel Tools Platform (PTP) is an official Eclipse Foundation Technology Project that focuses on integrating parallel tools into the Eclipse environment for enhanced application development. PTP supports a range of architectures & runtime systems & simplifies interaction with parallel systems. WS10 will introduce Eclipse &

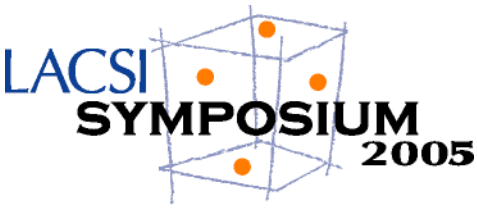

PTP, provide hands-on experience at managing & developing software, demonstrate both C/C++ & Fortran Development Toolkits, & present PTP tools. Participants will be able to use their own laptops (Linux or OS X) with supplied Eclipse & PTP software to maximize the hands-on time at software development activities.

Welcoming Reception and Poster Presentations: 6:00pm – 7:00pm

Workshops/tutorials on subjects of special interest to attendees, and the Welcome Reception and Poster Exhibit were featured. The posters were also made available over the next two days for additional inspection.

POSTER PRESENTATIONS

Using Cache Models and Empirical Search for Automatic Tuning of Applications
Contributors: Ken Kennedy, John Mellor-Crummey, Apan Qasem, (Rice University)

Parallel Space-Filling Curve Generation for Dynamic Load Balancing
Justin Luitjens, Tom Henderson, and Martin Berzins (University of Utah)

Adaptive Performance Monitoring and Profiling On Large Scale Systems
G. Todd Gamblin, Ying Zhang, Daniel A. Reed (Renaissance Computing Institute, University of North Carolina at Chapel Hill)

Support for Simultaneous Multiple Substrate Performance Monitoring
Kevin London, Shirley Moore, Daniel Terpstra, Jack Dongarra (University of Tennessee)

PathScale InfiniPath Interconnect Performance
Greg Lindahl (PathScale, Inc.)

An Initial Implementation of the Program Database Toolkit using the Open64 compiler
Oscar Hernandez (University of Houston), Sameer Shende (University of Oregon), Barbara Chapman (University of Houston)

Improving Adaptive Compilation with Truncated Execution and Loop Unrolling
Jeff Sandoval, Keith Cooper, Tim Harvey (Rice University)

Adaptive Inlining
Todd Waterman and Keith Cooper (Rice University)

Compiling for Memory Constraints on Short Vector Machines
Yuan Zhao, Ken Kennedy (Rice University)

Scout: A GPU-Accelerated Language for Visualization and Analysis
Patrick McCormick, Jeff Inman, James Ahrens (Los Alamos National Laboratory), Greg Roth, Chuck Hansen (University of Utah)

A Multi-platform Co-Array Fortran Compiler for High-Performance Computing
Yuri Dotsenko and Cristian Coarfa (Rice University)

Wednesday, October 12: Keynote and Paper Presentations

9:00am Welcoming remarks

9:30am: Keynote Address – "New Architectures for a New Biology"
David E. Shaw, D. E. Shaw Research and Development and
Center for Computational Biology and Bioinformatics, Columbia University

Abstract:
Some of the most important outstanding questions in the fields of biology, chemistry, and medicine remain unsolved as a result of our limited understanding of the structure, behavior and interaction of biologically significant molecules. The laws of physics that determine the form and function of these biomolecules are well understood. Current technology, however, does not allow us to simulate the effect of these laws with sufficient accuracy, and for a sufficient period of time, to answer many of the questions that biologists, biochemists, and biomedical researchers are most anxious to answer. This talk will describe the current state of the art in biomolecular simulation and explore the potential role of high-performance computing technologies in extending current capabilities. Efforts within our own lab to develop novel architectures and algorithms to accelerate molecular dynamics simulations by several orders of magnitude will be described, along with work by other researchers pursuing alternative approaches. If such efforts ultimately prove successful, one might imagine the emergence of an entirely new paradigm in which computational experiments take their place alongside those conducted in "wet" laboratories as central tools in the quest to understand living organisms at a molecular level, and to develop safe, effective, precisely targeted medicines capable of relieving suffering and saving human lives.

10:30am Break

11:00am – Reviewed Papers I: Systems

An Extensible Message-Oriented Offload Model for High-Performance Applications
Patricia Gilfeather, Arthur B. McCabe, University of New Mexico

Cluster Security with NVisionCC: The Forseti Distributed File Integrity Checker
Adam J. Lee, Gregory A. Koenig, William Yurcik, University of Illinois at Urbana-Champaign

Near-Real-time Availability Monitoring and Modeling for HPC/HEC Runtime Systems
Hertong Song, Chokchai Box Leangsuksun, Narasimha Raju Gottumukkala, Raja Nassar, Louisiana Tech University; Stephen L. Scott, Oak Ridge National Laboratory (ORNL); Andy Yoo, Lawrence Livermore National Laboratory (LLNL)

12:30pm – 2:00pm Lunch

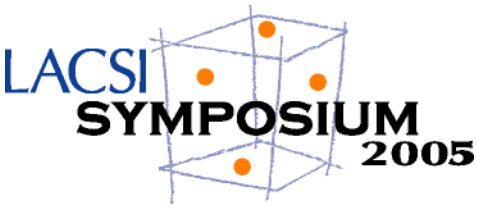

2:00pm – Reviewed Papers II: Performance

A Framework for Analyzing Linux System Overheads on HPC Applications
Sushant Sharma, Patrick G. Bridges, and Arthur B. Maccabe, University of New Mexico

Scalable Cross-Architecture Predictions of Memory Hierarchy Response for Scientific Applications
Gabriel Marin, John Mellor-Crummey, Rice University

Exploring Application Performance: a New Tool For a Static/Dynamic Approach
Lamia Djoudi, Jean-Thomas Acquaviva, Christophe Lemuet, William Jalby (LRC ITACA, CEA/DAM and Université de Versailles Saint-Quentin, France); Denis Barthou (PRiSM, Université de Versailles Saint-Quentin, France); Patrick Carribault (Bull SA, Les Clayes Sous Bois, France)

3:30pm – 4:00pm Break

4:00pm Reviewed Papers III: Algorithms and Applications

Performance Analysis, Modeling and Enhancement of Sandia's Integrated TIGER Series (ITS) Coupled Electron/Photon Monte Carlo Transport Code
Mahesh Rajan, Brian Franke, Robert Benner, Ron Kensek and Thomas Laub, Sandia National Laboratories

Using Space-filling Curves for Computation Reordering
Guohua Jin, John Mellor-Crummey, Rice University

Towards an Automatic and Application-Based Eigensolver Selection
Yeliang Zhang, Lawrence Berkeley National Laboratory; Xiaoye S. Li, Osni Marques, Lawrence Berkeley National Lab

ADJOURN

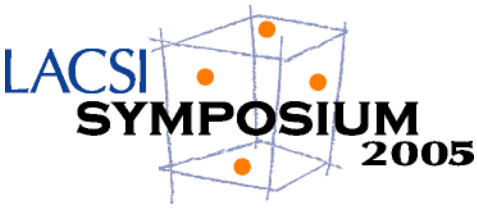

LACSI

Thursday, October 13: Panel Discussions

Two panel discussions took place focusing on important issues for researchers and managers in high performance computing and a closing reception ended the LACSI Symposium.

8:00am Breakfast

9:00am – Panel I: The Impact of ASC on Computer Science Research

Panelists: Sally McKee, Cornell University; Peter Eltgroth, Lawrence Livermore National Laboratory; Patrick Bridges, University of New Mexico; David Womble, Sandia National Laboratories; Ken Kennedy, Rice University; Rod Oldehoeft, Los Alamos National Laboratory

Representatives from each of the ASC laboratories and an academic partner present overviews of the interactions and their impact on research.

10:30am Break

11:00am – Panel II: Diverging Architectural Directions in HPC?
Moderator: Rob Fowler
Panelists: Burton Smith, Allen McPherson, Steve Poole, John Gustafson

Over the past few years, the "conventional cluster" architecture for general-purpose high-performance computing has consisted of collection of high-end, high-power microprocessors in small SMP boxes (with disks) connected by a network, either commodity or designed specifically as a cluster interconnect. These shared the HPC space with vector machines and large shared-memory systems. Recently, there has been a consensus that space, power, reliability, communication latency, and manageability are among the issues that constrain these systems. Some of the system characteristics perceived to address these problems, and on which emerging systems are based, include special purpose hardware, co-processors, processors that optimize a computation versus power function, multi-core chips, multi-threading, and vector/streaming processors. We thus appear to be entering an era in which architectural designs may be diverging. The panelists, and the audience, are invited to discuss emerging directions in high-end architectures and their implications on the user communities for these systems.

12:30pm – 2:00pm Lunch and Closing Reception

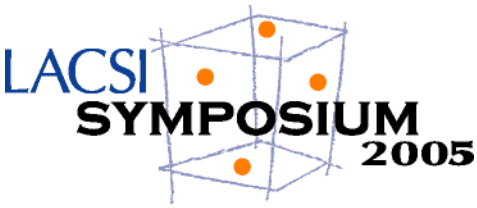

LACSI

SUPPLEMENTAL MATERIALS

Abstracts accompanying the Workshop on Algorithm Acceleration with Reconfigurable Hardware (Workshop 7)

Workshop Summary - Algorithm Acceleration with Reconfigurable Hardware
Organizers: Maya B. Gokhale (maya@lanl.gov) and Rod R. Oldehoeft (rro@lanl.gov)

The Design And Application Of Bee2 —A High-End Reconfigurable Computing System
Chen Chang (chenzh@eecs.berkeley.edu), John Wawrzynek (johnw@eecs.berkeley.edu), Robert W. Brodersen (rb@eecs.berkeley.edu)
Department of Electrical Engineering and Computer Sciences
University of California, Berkeley

High-order Finite Difference Seismic Modeling on Reconfigurable Computing Platform
Chuan He, Department of Electrical Engineering, Guan Qin, Institute for scientific computation, and Wei Zhao, Department of Computer Science, Texas A&M University, College Station, TX 77843

Terrestrial-based Radiation Upsets: A Cautionary Tale
Heather Quinn(hquinn@lanl.gov), Los Alamos National Laboratory, ISR-3 Space Data Systems, MSD440, Los Alamos, NM 87544

Trident: An FPGA Compiler Framework for Scientific Computing
Justin L.Tripp, Kristopher D. Peterson, Christine Ahrens, Jeffrey D. Poznanovic and Maya Gokhale

Implications of FPGAs for Floating-Point HPC Systems
Keith D. Underwood and K. Scott Hemmert

A Hardware-Accelerated Steady-State Power Flow Solver
Daniel G. Chavarría-Miranda and David Chassin, Pacific Northwest National Laboratory (PNNL), Richland, WA

A Reconfigurable Computing Framework for Multi-scale Cellular Image Processing
Reid Porter, Al Conti, Jan Frigo, Neal Harvey, Garret Kenyon, Maya Gokhale,
Los Alamos National Laboratory

FPGA Based High Performance Compute Platforms
Prasanna Sundararajan, Xilinx Research Labs, San Jose, CA

Hardware Accelerated Apriori Algorithm for Data Mining
Zachary K. Baker (zbaker@halcyon.usc.edu) and Viktor K. Prasanna (prasanna@ganges.usc.edu)
University of Southern California, Los Angeles, CA, USA

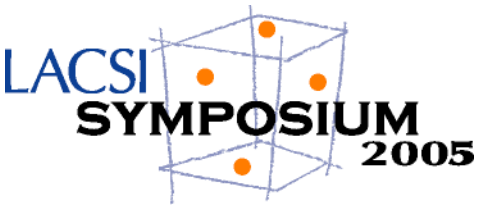

SUPPORTING CAST

The 2005 LACSI Symposium was successful because of the combined efforts of many people.

LACSI Co-Directors
Andy White, Los Alamos National Laboratory Ken Kennedy, Rice University

Organizing Committee
Rob Fowler, Program Chair, Rice University; Rod Oldehoeft, General Chair, LANL; Barney Maccabe, Local Arrangements Chair, University of New Mexico

Conference Arrangements
Deborah G. Cole, Registration Coordinator, University of New Mexico; Barbara Daniels, Meeting Coordinator, Local Arrangements, University of New Mexico

Submission and Review Management Software and Service

John Konkle, Linklings

Symposium Web Pages and LACSI CDROM

Sarah Gonzales Rice University; John Konkle, Linklings

This material is based on work supported by the Department of Energy under Contract Nos. 03891-001-99-4G, 74837-001-03 49, 86192-001-04 49, and/or 12783-001-05 49 from the Los Alamos National Laboratory.