scaling to new heigths

CUG 2006

PROGRAM

Switzerland, Lugano

May 8th - 11th 2006

cscs
Swiss National Supercomputing Centre

Welcome

Dear Friends,

The Cray User Group welcomes you to its annual technical conference, CUG 2006. CSCS is pleased to welcome you to the Italian-speaking part of Switzerland—where snow and palm trees meet. With a record 100 tutorials and technical presentations we are confident that all participants will benefit from a rich exchange of information.

Information on local attractions will be available from the local tourism office and don't hesitate to contact us at the CUG office for any assistance.

We hope you enjoy your stay in Lugano!

Yours sincerely,

Marie-Christine Sawley

CUG 2006 Local Arrangements Chair

Bordering on both Switzerland and Italy, Lake Lugano represents one of the greatest attractions to the area.

Our Theme

The theme "Scaling to new heights" was inspired by the purchase of our Cray XT3 system in the summer of 2005 and the natural environment of Switzerland. When visiting CUG 2005 in Albuquerque we had just received the first cabinet of the system. We are pleased to say that the system has been in general production since January 2006 and it is fulfilling the expectation and hopes that were communicated with our theme. Already during the previous pilot usage phase, users where very impressed by the new dimensions the system allowed them to work in. The system is rightfully named Horizon—it's acquisition having enabled us to look towards a whole new range of applications. But best of all it has enabled our users to scale to new heights and challenge problems that were previously out of reach.

Small and hardy as the chamois (the wild mountain goat), CSCS continues to scale to new heights in search of new horizons.

Welcome

On Behalf of the CUG Program Committee

Welcome to the 48th Cray User Group meeting, CUG 2006—Scaling to New Heights. This is truly a unique opportunity for you to exchange problem-solving information and enjoy professional interactions with your fellow Cray Inc. high performance computing system users.

The CUG Program Committee has assembled a diverse and impressive array of presentations in General and Parallel Technical Sessions, Tutorials, and Special Interest Group (SIG) meetings. Furthermore, Cray Inc. has committed to having many technical experts on-hand throughout the entire program.

- General Sessions will provide you with the latest corporate and technical information from Cray Inc. executives, as well as general interest technical presentations.
- Parallel Technical Sessions will give you the opportunity to focus on the specific knowledge domains of the SIG's.
- Tutorials are a great opportunity to update your technical skills with the help of selected technical experts from Cray Inc. and/or other CUG sites.
- SIG Meetings include discussion of the technical areas for each Cray platform (X1, XT3, and XD1). These meetings are where we work together to maintain and improve our CUG technical program.

In addition to these prepared presentations, tutorials, and SIG meetings, there are many opportunities for informal discussions.

- Birds of a Feather (BoF) Sessions, the most dynamic aspect of a CUG conference, are scheduled as needed. You are welcome to organize a BoF session. Notices of BoF sessions will be posted on the Message Board.
- The CUG Night Out, the Cray Reception, and luncheons are among the many occasions you will have to exchange information with your colleagues from other CUG sites and to engage the insight and expertise of representatives from Cray Inc.

This year we have two excellent Keynote Speakers.

- **Manuel C. Peitsch—Global Head of Systems Biology, Novartis Institutes for BioMedical Research**. Dr. Peitsch directs a department spanning experimental sciences (Proteomics), Computational Systems Biology, Computational Knowledge Management and Text Mining as well as Information Sciences. Prior to his current position, Dr. Peitsch was Global Head of Informatics and Knowledge Management where he was responsible for Research IT, Scientific Computing and Knowledge Management. Dr. Peitsch has published numerous papers, book chapters, technical reports and patents and made his most important contributions in the fields of life science informatics (bioinformatics, protein structure modeling, and knowledge management) and cell death research. He is the recipient of several honors and awards. In addition to his work at Novartis, he serves on the Research Council of the Swiss National Science Foundation, the Foundation Council of the Swiss Institute of Bioinformatics, and the Steering Boards of the Swiss Supercomputing Centre and the SwissBioGRID. In 1997, Dr. Peitsch co-founded the startup company Geneva Bioinformatics and the Research site GlaxoWellcome Experimental Research in Geneva. In 1998 he co-founded the Swiss Institute of Bioinformatics and later played a key role in extending this Institute to Basel. In 2003 he co-founded the SwissBioGRID. Mr. Peitsch received his Ph.D. in Biochemistry from the University of Lausanne in Switzerland and spent his post-doctoral years at the National Cancer Institute of the NIH and at the University of Lausanne. Since 2002, he has been a Professor for Bioinformatics with the University of Basel.

- **Marie-Christine Sawley—CEO Swiss National Supercomputing Centre**. Dr. Sawley completed her doctoral studies in physics at EPFL (Ecole Polytechnique Federale de Lausanne-The Swiss Federal Institute of Technology Lausanne) in 1985 in the Plasma Physics Institute. After a post-doc at the University of Sydney, she directed her career towards HPC management and was appointed head of the User Group at the Scientific Computing Centre (SIC) of EPFL when the first supercomputers were installed in Switzerland (1988). Over the past 15 years Dr. Sawley has been an active member in a number of selection committees for HPC systems purchase and development projects such as the PATP (PSC Parallel Applications Technology Program) collaboration with Cray, the Swiss Tx series with Compaq, and the establishment of the Vital IT centre with HP and Intel. From 1995 to 1998 Dr. Sawley was President of the Speedup forum and since 1997 she has been a committee member for the SOS workshops (Sandia National Labs, Oak Ridge National Laboratory, and Switzerland). She was appointed CEO of the Swiss National Supercomputing Centre, CSCS, in July 2003.

CUG 2006 promises to be a rewarding, educational, and entertaining conference.

I look forward to meeting you at the May conference in Lugano, Switzerland.

David J. Gigrich
CUG Vice-President and Program Chair
The Boeing Company

CRAY USER GROUP INCORPORATED

Contents

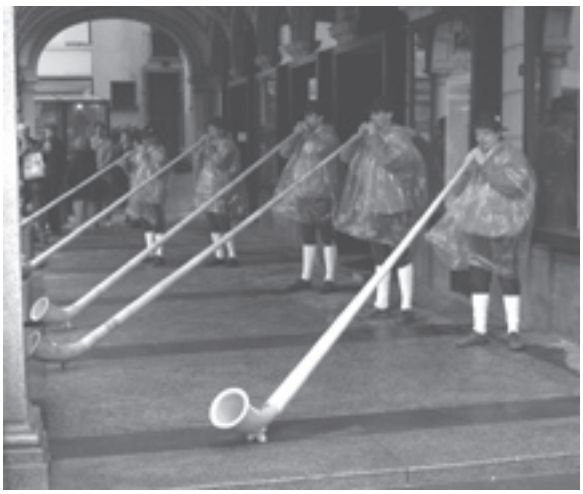

Traditional alphorn players in downtown Lugano.

Interactive Sessions

We have transformed the Special Interest Group (SIG) meetings into "Interactive Sessions" to better reflect the needs of our members. A set of Interactive Sessions will be held for each Cray product (X1, XT3, and XD1). These meetings are open to all interested CUG attendees and will take place on Monday, Tuesday, and Wednesday. They provide a forum to become acquainted with other CUG sites that share your interests and to discuss future direction, important issues, and areas of special interest. Liaisons from Cray Inc. will be available for each Interactive Session to help address questions and to foster communication between the membership and Cray Inc. You are encouraged to attend any of these open meetings.

Program Committee

The Program Committee needs your ideas! We encourage you to share your suggestions of how to improve the conference with the CUG Program Chair, CUG Directors, and/or any of the SIG Chairs. The CUG Conference Program is the net result of the combined efforts of all interested CUG sites. You can make a unique contribution to the next CUG conference!

Cell Phone and Pager Policy

Please turn off your cell phone and/or pager during conference sessions.

Smoking Policy

There is no smoking allowed at the Conference.

Changes to the Program

We always anticipate there will be some changes to the Program schedule. Check the Message Board on site at the conference each day for schedule changes. Thank you for your patience and cooperation.

Pocket Schedule

The Pocket Schedule provides the schedule and other information in a handy pocket size and replaces the program schedule layout used in past programs.

Conference Evaluation and Cray Inc. Service Survey

The CUG 2006 Evaluation and the Cray Inc. Service Survey will be available online starting Wednesday at 8:00 AM until Thursday at 5:00 PM. Please go to www.cug.org/4-organization/eval_survey/ and take time to complete them both. Once you've completed them, please stop by the CUG office. We have a gift for you.

MONDAY

1A Tutorials 8:00 Room B1

Portable Performance with Vectorization, *Mark Fahey and James B. White III (ORNL)*

Getting high performance is more than just finding the right compiler options. Modifying a code to tune for a specific system, such as a vector system, can slow the code down on other systems and make it more difficult to maintain. This tutorial uses examples from real applications to show how to program for portable performance while preserving maintainability.

1B Tutorials 8:00 Room B2

Portals Programming on the XT3, *Ron Brightwell and Rolf Riesen (SNLA); Arthur B. Maccabe, University of New Mexico*

This tutorial will describe the Portals network programming interface as it is implemented on the Cray XT3. We will describe the fundamental building blocks of Portals and how they can be combined to support upper-layer network transport protocols and semantics. We will also cover performance and usability issues that are specific to the XT3 system.

1C Tutorials 8:00 and 10:20 Room B3

Lustre Tutorial for Cray XD1 and Cray XT3, *Rick Slick, Cray Inc. and Jeff Denworth, CFS*

In this tutorial, we will give an overview of Lustre, including architecture and common vocabulary, and describe differences between Cray XD1 and Cray XT3 Lustre implementations. We will discuss Lustre configuration guidelines and samples for both Cray XD1 and Cray XT3. We will also discuss Lustre administration, including initialization, monitoring, starting, and stopping. We will present the common user level commands with special attention to the lsf command and striping. The current state and future of Lustre will be given by CFS.

2A Tutorials 10:20 Room B1

Performance Measurement, Tuning, and Optimization on the Cray XT3, *Luiz DeRose and John Levesque, Cray Inc.*

In this tutorial we will discuss techniques and tools for application performance tuning on the Cray XT3. We will briefly review the system and architecture, focusing on aspects that are important to understand in order to realize performance; we will discuss compiler optimization flags and numerical libraries; we will present the Cray performance measurement and analysis tools; and will conclude with optimization techniques.

2B Tutorials 10:20 Room B2

PBS Tutorial, *Michael Karo, Cray Inc.*

We will have a PBS tutorial this year based on responses from last year's CUG. Joining us this year will be representatives from Altair to assist in the tutorial (at the suggestion of the CUG group last year). We will discuss PBS in terms of the X1/X1E, the XT3, and the XD1 systems. We would expect a fair amount of interaction with the audience for this tutorial.

LUNCH 12:10 Atrio/Foyer, Ground Floor

3 Opening-General Session 1:30 Room B1

1:30 Welcome, *Barbara Horner-Miller (ARSC), CUG President and Marie-Christine Sawley (CSCS), Swiss National Supercomputing Centre, CEO*

Keynote Address: High Performance Computing in Drug Discovery, *Dr. Manuel C. Peitsch—Global Head of Systems Biology, Novartis Institutes for BioMedical Research*

2:20 Leadership Computing at the NCCS, *Arthur Bland, Ann Baker, R. Shane Canon, Ricky Kendall, Jeffrey Nichols, and Julia White (ORNL)*

We will describe the Leadership Computing Facility (LCF) at the National Center for Computational Sciences (NCCS), both the current facility and our aggressive plans for upgrade. Project allocations were announced in January, and we will introduce some of the groundbreaking science enabled by the LCF. We will also describe the unique service model of the NCCS, designed for capability computing and breakthrough science.

4A Development of Large Scale Systems/Applications 2:55 Room B1

Implementing the Operational Weather Forecast Model of MeteoSwiss on a Cray XT3, *Angelo Mangili, Emanuele Zala, Jean-Marie Bettems, Mauro Ballabio, and Neil Stringfellow (CSCS)*

MeteoSwiss uses the computer facilities of CSCS to carry out the twice daily weather forecasts for Switzerland using a specialized alpine model called aLMo based on the LM code from Deutscher Wetterdienst. This paper describes the implementation of the aLMo suite on the Cray XT3 at CSCS and outlines some of the issues faced in porting and configuring a complex operational suite on a new architecture. Performance figures for the suite are given, and an outline is presented of some of the challenges to be faced in carrying out the next generation of high resolution forecasting.

4A Development of Large Scale Systems/Applications 3:30 Room B1

An Investigation of Application Fit to Architecture, *Courtenay Vaughan, Mahesh Rajan, and Robert Benner (SNLA)*

There are several differences between large commodity clusters and more custom HPC machines like Red Storm/XT3. We use standard benchmarks and Sandia applications to compare and contrast these effects on performance on large numbers of processors. This will allow us to better match the application space requirements of the user community with installed and future HPC systems.

4A Development of Large Scale Systems/Applications 4:00 Room B1

Performance Analysis and Modeling of Discrete Simulation Monte Carlo (DSMC) Code for Transient Microscale Gas Flows, *Mahesh Rajan, Michail Gallis, and John Torczynski (SNLA)*

Sandia's DSMC code is currently used to simulate physical devices like MEMS, Neutron Generators, and other low density noncontinuum flows. Scalability of the code to thousands of processors is measured. We will discuss development of performance models and their use to predict performance of DSMC on the Red Storm/XT3.

4B Chapel/Black Widow 2:55 Room B2

Chapel: Cray Cascade's High Productivity Language, *Mary Beth Hribar, Steven Deitz, and Bradford L. Chamberlain, Cray Inc.*

In 2002, Cray joined DARPA's High Productivity Computing Systems (HPCS) program, with the goal of improving user productivity on High-End Computing systems for the year 2010. As part of Cray's research efforts in this program, we have been developing a new parallel language named Chapel, designed to: support a global view of parallel programming with the ability to tune for locality; support general parallel programming including codes with data-parallelism, task-parallelism, and nested parallelism; help narrow the gulf between mainstream and parallel languages. We will introduce the motivations and foundations for Chapel, describe several core language concepts, and show some sample computations written in Chapel.

4B Chapel/Black Widow 3:40 Room B2

The Cray Programming Environment for Black Widow, *Luiz DeRose, Terry Greyzck, Mary Beth Hribar, Brian Johnson, Bill Long, and Mark Pagel, Cray Inc.*

In this paper we will present the Cray programming environment for Black Widow, the next generation of Cray vector systems. The paper will cover programming models, scientific libraries, and tools. In addition, the paper will describe the new features from the Fortran standard that will be available on the Black Widow compiler.

4C System Operations-Schedulers & Monitoring Tools 2:55 Room B3

Application Level Placement Scheduler (ALPS), *Michael Karo, Richard Lagerstrom, Marlys Kohnke, and Carl Albing, Cray Inc.*

The Application Level Placement Scheduler (ALPS) is being introduced with the vector follow-on product to the X1/X1E. It is planned for use with future Cray products. ALPS is the replacement product to psched. It will manage scheduling and placement for distributed memory applications and support predictable application performance. It will integrate with system configuration and workload management components with an emphasis on extensibility, scalability, and maintainability.

4C System Operations-Schedulers & Monitoring Tools 3:30 Room B3

Moab Workload Manager on Cray XT3, *Michael Jackson, Don Maxwell, and Scott Jackson, Cluster Resources Inc.*

Cluster Resources Inc. has ported their Moab Workload Manager to work with Cray XT3 systems (both PBS Pro, runtime, and other resource manager environments) and the National Center for Computational Sciences (NCCS) is now using Moab to schedule production work. We will discuss the Moab port, how NCCS is using it, as well as how it helps provide added control over XT3 systems in terms of advanced policies, fine-grained scheduling, diagnostics, visualization of resources and improved resource utilization. We will also review plans and possibilities for Moab technologies on other Cray products.

4C System Operations-Schedulers & Monitoring Tools 4:00 Room B3

CrayViz: A Tool for Visualizing Job Status and Routing in 3D on the Cray XT3, *John Biddiscombe and Neil Stringfellow (CSCS)*

When initially installed, the Cray XT3 at CSCS showed node failures which appeared to occur in a pattern. The tools available for the analysis of the failures were limited to 2D printouts of node status and were difficult to interpret given the 3-dimensional connectivity between the processors. A visualization tool capable of displaying the status of compute nodes and the routes between them was developed to assist in the diagnosis of the failures. It has the potential to assist in the analysis of the working system and the development of improved allocation and job scheduling algorithms.

5A X1/E-Interactive Session 4:30 Room B1

Cray X1/E Special Interest Group Meeting (SIG), *Mark Fahey (ORNL), Chair*

Systems, *Jim Glidewell (BOEING) and Brad Blasing (NCS-MINN), Focus Chairs*

Users, *Mark Fahey (ORNL) and Rolf Rabenseifner (HLRS), Focus Chairs*

The purpose of this Interactive Session is to provide a forum for open discussions on the Cray X1/E system and its environment. The "Systems" area focuses on issues related to Operations, Systems & Integration, and Operating Systems. The "Users" area covers issues related to Applications, Programming Environments, and User Services. Cray Inc. liaisons will be on hand to provide assistance and address questions. This session is open to everyone that has an interest in the Cray X1/E and wishes to participate.

MONDAY (CONTINUED)

5B Interactive Sessions 4:30 Room B2

Reserved for interactive session.

5C BoF 4:30 Room B3

Cray Service Today and Tomorrow, *Charlie Clark and Don Grooms, Cray Inc.*

Cray will present its current support model and reporting tools used for XT3, X1, and other legacy products. The discussion will then explore possibilities in which Cray could provide service differently in the future, and discuss the potential changes required to both its products and to its service offerings to accommodate these changes. The session will start with a short presentation then move to Q&A.

TUESDAY

6 General Session 8:20 Room B1

8:20 Keynote Speaker Introduction, *Ladina Gilly (CSCS), Local Arrangements Coordinator*

Keynote Address, *Dr. Marie-Christine Sawley (CSCS), Swiss National Supercomputing Centre, CEO*

8:50 Cray Corporate Update, *Pete Ungaro, CEO and President, Cray Inc.*

9:15 Adaptive Supercomputing Vision, *Jan Silverman, Senior Vice President, Corporate Strategy and Business Development, Cray Inc.*

9:30 Cray Product Roadmap, *Steve Scott, Chief Technology Officer, Cray Inc.*

9:55 Cray General Software Direction, *Paul Krueger, Vice President Software, Cray Inc.*

7A Physics 10:40 Room B1

Quantum Mechanical Simulation of Nanocomposite Magnets on Cray XT3, *Yang Wang (PITTSCC); G.M. Stocks, D.M.C. Nicholson, A. Rusanu, and M. Eisenbach (ORNL)*

In this presentation, we demonstrate our capability of performing the quantum mechanical simulation of nanocomposite magnets using a Cray XT3 supercomputing system and the Locally Self-consistent Multiple Scattering (LSMS) method, a linear scaling *ab initio* method capable of treating tens of thousands of atoms. The simulation is intended to study the physical properties of magnetic nanocomposites made of FePt and Fe nanoparticles. We will discuss our results on the electronic and magnetic structure and the magnetic exchange coupling between nanoparticles embedded in metallic matrices.

7A Physics 11:20 Room B1

Astrophysical Particle Simulations and Reconfigurable Computing on the Cray XD1, *Rainer Spurzem and Andreas Ernst, Astronomisches Rechen-Institut; Tsuyoshi Hamada and Naohito Nakasato, RIKEN Institute*

We have developed and tested FPGA-based pipelines to compute gravitational forces between particles in astrophysical N-body simulations, and implemented them on the Cray XD1 at NIC Jülich. First results regarding the implementation, the integration in the standard code, and the performance are given. Perspectives are discussed for the future use of FPGA as well as other special purpose hardware (GRAPE) for this and other astrophysical high-performance computing applications.

7A Physics 11:45 Room B1

Exploring Novel Quantum Phases on Cray X1E, *Evgeni Burovski, Nikolay Prokofev, and Boris Svistunov, University of Massachusetts; Matthias Troyer, ETH Zürich*

At temperatures close to the absolute zero (-273°) some metals go superconductive, i.e., lose the electrical resistance [Nobel Prizes 1913, 1972]. As a related phenomenon, some liquids go superfluid at sufficiently low temperature [Nobel Prizes 1978, 1996, 2001]. These phenomena are well understood in the two limiting cases, namely the Bardeen-Cooper-Schrieffer (BCS) limit, and the Bose-Eistein condensation (BEC) limit. Using the Cray X1E aphonic of the Oak Ridge National Lab, we performed Quantum Monte Carlo simulations to study the universal features of the crossover between these two extremes and, for the first time, obtained an accurate and quantitative description. The QMC code is highly parallelizable, and ideally suited to vector machines, with the most computationally expensive procedure being the vector outer product. The vector processors of the X1E were indispensable for the project, which would otherwise be impossible to complete.

7B Mass Storage & Grid Computing 10:40 Room B2

Storage Architectures for Cray-based Supercomputing Environments, *Matthew O'Keefe, Cray Inc.*

HPC data centers process, produce, store and archive large amounts of data. In this paper, existing storage architectures and technologies, including file systems, HSM, and large tape archives will be reviewed in light of current experience. The biggest storage problem facing Cray users may not be attaining peak IO and storage processing speeds, but instead may be sharing that data out to the rest of the data center from the Cray supercomputer. We will discuss technical problems in making this work efficiently and suggest possible solutions.

7B Mass Storage & Grid Computing
11:20 Room B2

Current Progress of the GRID Project at KMA, *Hee-Sik Kim, Tae-Hun Kim, Cray Korea; Ha-Young Oh and Dongil Lee (KMA)*

KMA is one of the pioneer groups to implement GRID projects of KOREA in meteorology. Cray has supported the establishment of a grid infrastructure on the largest Cray X1E. We will introduce the effort and describe current progress.

7B Mass Storage & Grid Computing
11:45 Room B2

PDIO Interface, Functionality and Performance Enhancements, *Nathan Stone, D. Balog, B. Gill, B. Johanson, J. Marsteller, P. Nowoczynski, R. Reddy, J.R. Scott, J. Sommerfield, K. Vargo, and C. Vizino (PITTSCC)*

PSC's Advanced Systems Group has created Portals Direct I/O ("PDIO"), a special-purpose middleware infrastructure for sending data from compute processor memory on Cray XT3 compute nodes to remote agents anywhere on the WAN in real-time. The prototype provided a means for aggregation of outgoing data through multiple load-balanced PDIO daemons, end-to-end parallel data streams through XT3 SIO nodes, and a bandwidth feedback mechanism for stability and robustness. However, it was limited by the special-purpose nature of its initial design: adopters had to modify their application source code and each write was bound to an independent remote file. This version was used by one research group, demonstrated live at several conferences and shown to deliver bandwidths of up to 800 Mbits/sec.

PSC is advancing a beta release with a number of interface, functionality and performance enhancements. First, in order to make PDIO a general-purpose solution, developers are re-implementing the API to make it transparent. Users will invoke it through the standard C API, obviating the need for any changes to application source code. Second, developers already have prototypes for more scalable and robust portals IPC and support for broader file semantics. This includes a more general dynamic allocation mechanism to allow the PDIO daemons to be a truly shared resource for all users on the XT3 system. Finally, together with an updated WAN communication protocol, the goal of these enhancements is to continue to deliver performance limited only by the WAN connectivity and file systems of the remote host.

7C System Operations 10:40 Room B3

XT3 Operational Enhancements, *Chad Vizino, Nathan Stone, and J. Ray Scott (PITTSCC)*

The Pittsburgh Supercomputing Center has developed a set of operational enhancements to the supplied XT3 environment. These enhancements facilitate the operation of the machine by allowing it to be run efficiently and by providing timely and relevant information upon failure. Custom scheduling, job-specific console log gathering, event handling, and graphical monitoring of the machine will be reviewed and discussed in depth.

7C System Operations 11:20 Room B3

A Preliminary Report on Red Storm RAS Performance, *Robert A. Ballance and Jon Stearley (SNLA)*

The Cray XT3 provides a solid infrastructure for implementing and measuring reliability, availability, and serviceability (RAS). A formal model interpreting RAS information in the context of Red Storm was presented at CUG 2005. This paper presents an implementation of that model—including measurements, their implementation, and lessons learned.

7C System Operations 11:45 Room B3

Extending the Connectivity of XT3 System Nodes, *Davide Tacchella (CSCS)*

XT3 system service nodes provide a limited connectivity to the outside world. The boot node can be accessed only from the control workstation and the sdb is connected only to the systar network. CSCS has extended the connectivity of both nodes on its XT3. These extensions allow us to back up the file system of the boot node at run time, and to query the system database without being logged on the machine. The talk will present the chosen solutions for extending the node connectivity and the experiences made with them.

LUNCH 12:10 Atrio/Foyer, Ground Floor

8A Performance/Evaluations
1:30 Room B1

Performance Evaluations of User Applications on NCCS's Cray XT3 and X1E, *Arnold Tharrington (ORNL)*

The National Center of Computational Sciences has granted select projects the use of its Cray XT3 and X1E for fiscal year 2006. We present highlights of the projects and their performance evaluations.

8A Performance/Evaluations
2:00 Room B1

Evaluation of the Cray XT3 at ORNL: A Status Report, *Richard Barrett, Jeffrey Vetter, Sadaf Alam, Tom Dunigan, Mark Fahey, Bronson Messer, Richard Mills, Philip Roth, and Patrick Worley (ORNL)*

Over the last year, ORNL has completed installation of its 25 TF Cray XT3, named Jaguar, and has started running leadership-scale applications on this new platform. In this paper, we describe the current state of Jaguar and provide an overview of the capabilities of this system on numerous strategic DOE application areas that include astrophysics, fusion, climate, materials, groundwater flow/transport, and combustion.

TUESDAY (CONTINUED)

8A Performance/Evaluations
2:30 Room B1

Performance of Applications on the Cray XT3, *Deborah Weisser, Nick Nystrom, Shawn Brown, Jeff Gardner, Dave O'Neal, John Urbanic, Junwoo Lim, R. Reddy, Rich Raymond, Yang Wang, and Joel Welling (PITTSCC)*

Application performance on the XT3 has increased as our detailed understanding of factors affecting performance has improved. Through copious acquisition and careful examination of data obtained through CrayPat, PAPI, and Apprentice2, we have identified (and subsequently eliminated) performance impediments in key applications. This paper presents widely applicable strategies we have employed to improve performance of these applications.

8B BlackWidow Overview 1:30 Room B2

BlackWidow Hardware System Overview, *Brick Stephenson, Cray Inc.*

The BlackWidow system is a second-generation scalable vector processor closely coupled to an AMD Opteron based Cray XT3 system. Major hardware improvements over the Cray X1 system will be discussed including: single thread scalar performance enhancement features, high performance fat tree interconnect, high speed I/O bridge to the Cray XT3 system, and the high density packaging used on the BlackWidow system. Prototype hardware is starting to arrive in Chippewa Falls; a complete status of the system checkout progress will be given.

8B BlackWidow Overview 2:00 Room B2

BlackWidow Software Overview, *Don Mason, Cray Inc.*

The paper overviews the BlackWidow software. The OS for BlackWidow is Linux based, different than previous vector architectures. OS and IO software, including scheduling and administration is common to both BlackWidow and the scalar product line. The compiler and library software is an evolution of software developed for the Cray X1.

8B BlackWidow Overview 2:30 Room B2

Use of Common Technologies Between XT3 and BlackWidow, *Jim Harrell and Jonathan Sparks, Cray Inc.*

This talk will present how Cray uses common technologies between XT3 and BlackWidow. How past proven technologies and new emergent technologies help in administration and configuration of large scale systems.

8C FPGAs 1:30 Room B3

Using an FPGA Designer System to Program FPGAs in the XD1, *Karen Haines, Adam Deller, Steven Tingay, Aidan Hotan, and John Dickey (WASP)*

While FPGAs are configurable, configuring these devices remains a formidable task for any application. This work will investigate the application of an FPGA-based designer to develop an FPGA-based software correlator on the Cray XD1 suitable for the correlation of data from radio telescopes.

8C FPGAs 2:00 Room B3

FPGA-accelerated Finite-Difference Time-Domain Simulation on the Cray XD1 Using Impulse C, *Peter Messmer, David Smithe, and Paul Schoessow, Tech-X Corporation; Ralph Bodenner, Impulse Accelerated Technologies*

The Finite-Difference Time-Domain (FDTD) algorithm is a well-established tool for modeling transient electromagnetic phenomena. Using a spatially staggered grid, the algorithm alternately advances the electric and magnetic fields, according to Faraday's and Ampere's laws. These simulations often contain millions of grid cells and run for thousands of timesteps, requiring highly efficient grid update algorithms and high-performance hardware. Here we report on the implementation of the FDTD algorithm on the application accelerator of a Cray XD1. The FPGA is programmed using the Impulse C tool suite. These tools translate C code into VHDL and therefore enable the domain scientist to develop FPGA-enhanced applications. Different optimization strategies, ranging from algorithmic changes to exploitation of the high degree of parallelism on the FPGA, will be presented.

8C FPGAs 2:30 Room B3

Benchmarking of Cray Systems at CCLRC Daresbury Laboratory, *Miles Deegan, Martyn Guest, Richard Wain, Igor Kozin, Christine Kitchen, and Mike Ashworth (DARESBURYLAB)*

We present an overview of our experience with the XD1 at CCLRC Daresbury Laboratory. Data for a number of synthetic and real applications will be discussed along with data gathered on other x86-64 based systems, e.g., Pathscale's Infinipath. In addition we will cover RAS and general usability issues we have encountered. Time permitting we will go on to discuss our initial efforts to exploit the Virtex IV FPGA as a co-processor on the XD1, along with details of benchmarking exercises carried out on XT3 and X1E systems at collaborating sites.

9A Architecture-Catamount 3:20 Room B1

Overview of the XT3 Systems Experience at CSCS, *Richard Alexander (CSCS),*

In July 2005 CSCS took delivery of a 1100+ node XT3 from Cray. This system completed acceptance successfully on 22 December 2005. This paper describes most aspects of the work needed to bring this machine from loading dock to production status. We will also briefly mention the uncompleted goals, including batch scheduling, scalability of systems administration tools, and deployment of the Lustre file system.

9A Architecture-Catamount 4:00 Room B1

Catamount Software Architecture with Dual Core Extensions, *Ron Brightwell, Suzanne M. Kelly, and John P. VanDyke (SNLA)*

Catamount is the lightweight kernel operating system running on the compute nodes of Cray XT3 systems. It is designed to be a low overhead operating system for a parallel computing environment. Functionality is limited to the minimum set needed to run a scientific computation. The design choices and implementations will be presented. This paper is a reprise of the CUG 2005 paper, but includes a discussion of how dual core support was added to the software in the fall/winter of 2005.

9B Tuning with Tools 3:20 Room B2

Performance Tuning and Optimization with CrayPat and Cray Apprentice2, *Luiz DeRose, Bill Homer, Dean Johnson, and Steve Kaufmann, Cray Inc.*

In this paper we will present the Cray performance measurement and analysis infrastructure, which consists of the CrayPat Performance Collector and the Cray Apprentice2 Performance Analyzer. These performance measurement and analysis tools are available on all Cray platforms and provide an intuitive and easy to use interface for performance tuning of scientific applications. The paper will describe, with examples of use, the main features that are available for understanding the performance behavior of applications, as well as to help identify sources of performance bottlenecks.

9B Tuning with Tools 4:00 Room B2

Performance and Memory Evaluation Using TAU, *Sameer Shende, Allen D. Malony and Alan Morris, University of Oregon; Pete Beckman, Argonne National Laboratory*

The TAU performance system is an integrated performance instrumentation, measurement, and analysis toolkit offering support for profiling and tracing modes of measurement. This paper introduces memory introspection capabilities of TAU featured on the Cray XT3 Catamount compute node kernel. One technique that TAU employs examines the memory headroom, or the amount of heap memory available, at routine entry, and correlates it to the program's callstack as an atomic event.

9C Simulation & Research Opportunities 3:20 Room B3

Supercomputer System Design Through Simulation, *Rolf Riesen (SNLA)*

Changes to an existing supercomputer such as the Cray XT3, or planning and designing the next generation, are very complex tasks. Even small changes in network topology, network interface hardware, or system software can have a large impact on application performance and scalability. We are working on a simulator that will be able to answer questions about system performance and behavior before actually changing or improving hardware or software components. This paper gives an overview of our project, its approach, and some early results.

9C Simulation & Research Opportunities 4:00 Room B3

Western Australian Supercomputer Program—Scientific and Collaborative Opportunities, *Karen Haines (WASP)*

This talk will discuss research areas supported by WASP. Applications use the XD1 and/or XT3 computers. I will also discuss funded visiting researcher opportunities within the WASP.

10A XT3-Interactive Session 4:30 Room B1

Cray XT3 Special Interest Group Meeting (SIG), *Robert A. Ballance (SNLA), Chair*

Systems, *Robert A. Ballance and Keith Underwood (SNLA), Focus Chairs*

Users, *Thomas Oppe (CSC-VBURG) and James Kasdorf (PITTSCC), Focus Chairs*

The purpose of this Interactive Session is to provide a forum for open discussions on the Cray XT3 system and its environment. The "Systems" area covers issues related to Operations, Systems & Integration, and Operating Systems. The "Users" area deals with issues related to Applications, Programming Environments, and User Services. Cray Inc. liaisons will be on hand to provide assistance and address questions. This session is open to everyone that has an interest in the Cray XT3 and wishes to participate.

10B BoF 4:30 Room B2

Open for BoF

10C BoF 4:30 Room B3

OpenFPGA—The Role of Reconfigurable Technology in HPC, *Eric Stahlberg (OSC)*

Advances in reconfigurable computing technology, such as in the Cray XD1, have reached a performance level where they can rival and exceed the performance of general purpose processors in the right situations. As present general purpose processors undergo transformation to increase throughput without increasing clock speeds, the opportunity for broader use of reconfigurable technology arises. OpenFPGA has emerged as an international effort aiming to ease the software transition required for applications to take advantage of reconfigurable technology. This BOF session will provide an opportunity to share viewpoints and opinions regarding the future incorporation of reconfigurable technology in computing, challenges faced, and approaches that can be taken within the context of an community wide effort to make progress.

WEDNESDAY

11A Performance 8:20 Room B1

FV-CAM Performance on the XT3 and X1E, *Patrick Worley (ORNL)*

The Community Atmospheric Model (CAM) is currently transitioning from a spectral Eulerian solver for the atmospheric dynamics to a finite volume-based solver that has the conservation properties required for atmospheric chemistry. We analyze the performance of CAM using the new dynamics on the Cray XT3 and Cray X1E systems at Oak Ridge National Laboratory, describing the performance sensitivities of the two systems to the numerous tuning options available in CAM. We then compare performance on the Cray systems with that on a number of other systems, including the Earth Simulator, an IBM POWER4 cluster, and an IBM POWER5 cluster.

11A Performance 9:00 Room B1

HPCC Update and Analysis, *Jeff Kuehn (ORNL); Nathan Wichmann, Cray Inc.*

The last year has seen significant updates in the programming environment and operating systems on the Cray X1E and Cray XT3 as well as the much anticipated release of version 1.0 of HPCC Benchmark. This paper will provide an update and analysis of the HPCC Benchmark Results for Cray XT3 and X1E as well as a comparison against last years' results.

11A Performance 9:25: Room B1

Comparing Optimizations of GTC for the Cray X1E and XT3, *James B. White III (ORNL); Nathan Wichmann, Cray Inc.; Stephane Ethier, Princeton Plasma Physics Laboratory*

We describe a variety of optimizations of GTC, a 3D gyrokinetic particle-in-cell code, for the Cray X1E. Using performance data from multiple points in the optimization process, we compare the relative importance of different optimizations on the Cray X1E and their effects on XT3 performance. We also describe the performance difference between small pages and large pages on the XT3, a difference that is particularly significant for GTC.

11A Performance 9:50 Room B1

Performance Comparison of Cray X1E, XT3, and NEC SX8, *Hongzhang Shan and Erich Strohmaier, Lawrence Berkeley National Laboratory*

In this paper, we focus on the performance comparison of Cray X1E, XT3, and NEC SX8 using synthetic benchmarks and SciDAC applications. Further, we will analyze the relationship between performance differences and hardware features.

11B UPC 8:20 Room B2

Using Unified Parallel C to Enable New Types of CFD Applications on the Cray X1E, *Andrew Johnson (NCS-MINN)*

We are using parallel global address-space (PGAS) programming models, supported at the hardware-level on the Cray X1E, that allow for the development of more complex parallel algorithms and methods and are enabling us to create new types of computational fluid dynamics codes that have advanced capabilities. These capabilities are the ability to solve complex fluid-structure interaction applications and those involving moving mechanical components or changing domain shapes. This is a result of coupling automatic mesh generation tools and techniques with parallel flow solver technology. Several complex CFD applications are currently being simulated and studied using this method on the Cray X1E and will be presented in this paper.

11B UPC 9:00 Room B2

Evaluation of UPC on the Cray X1E, *Richard Barrett (ORNL); Tarek El-Ghazawi and Yiyi Yao, George Washington University; Jeffey Vetter (ORNL)*

Last year we presented our experiences with UPC on the Cray X1. This year we update our results, in consideration of the upgrade to X1E of the machine at Oak Ridge National Laboratory as well as modifications to the UPC implementation of the NAS Parallel Benchmarks.

11B UPC 9:25 Room B2

Performance of WRF Using UPC, *Hee-Sik Kim and Jong-Gwan Do, Cray Korea Inc.*

The RSL_LITE communication layer of WRF was modified, which can do UPC and/or SHMEM. Its performance comparison with MPI version will be discussed.

11B UPC 9:50 Room B2

Parallel Position-Specific Iterated Smith-Waterman Algorithm Implementation, *Maciej Cytowski, Witold Rudnicki, Rafal Maszkowski, Lukasz Bolikowski, Maciej Dobrzynski, and Maria Fronczak (WARSAWU)*

We have implemented a parallel version of Position-Specific Iterated Smith-Waterman algorithm using UPC. The core Smith-Waterman implementation is vectorized for Cray X1E, but it is possible to use an FPGA core instead. The quality of results and performance are discussed.

11B UPC (paper only; no presentation)

Atomic Memory Operations, *Phillip Merkey (MTU)*

The Atomic Memory Operations provided by the X1 are the foundation for lock-free data structures and synchronization techniques that help users make the most of the UPC programming model. This paper focuses on applications that either could not be written without AMOs, or were significantly improved by the use of AMOs on the X1. We will cover basic data structures like shared queues and heaps and show how their use improves scheduling and load balancing. We will give examples of graph algorithms that can be intuitively parallelized only if AMOs are available.

11C Performance-XD1/FPGAs
8:20 Room B3

Early Experiences with the Naval Research Laboratory XD1, *Wendell Anderson, Marco Lanzagorta, Robert Rosenberg, and Jeanie Osburn (NRL)*

The Naval Research Laboratory has recently obtained a two-cabinet XD1 that has 288 dual core Opteron 275 dual core CPUs and 144 Vertex 2 FPGAs. This paper will examine our experiences with running scientific applications that take advantage of these two new technologies.

11C Performance-XD1/FPGAs
9:00 Room B3

Molecular Dynamics Acceleration with Reconfigurable Hardware on Cray XD1: A System Level Approach, *Brice Tsakam-Sotch, XLBiosim*

Molecular dynamics simulations are most realistic with explicit solvent molecules. More than 90% of the computing time is consumed in the evaluation of forcefields for simulation runs that last months on standard computers. Cray XD1 systems allow designing improved solutions for these problems by the mean of reconfigurable hardware acceleration and high performance buses. We present the performance analysis of the molecular dynamics software application gromacs and the resulting optimizations at the system level, including hardware acceleration of forcefield evaluation and optimization of communications.

11C Performance-XD1/FPGAs
9:25 Room B3

High Level Synthesis of Scientific Algorithms for the Cray XD1 System, *Paolo Palazzari, Giorgio Brusco and Paolo Verrecchia, Ylichron Srl; Alessandro Marongiu and Vittorio Rosato, Ylichron Srl and ENEA; Paolo Migliori, ENEA; Philip LoCascio (ORNL)*

The rapid increase in integrated circuit densities, with the slowing in the increase in commodity CPU clock rates, has made the investigation of hybrid architectures using FPGA technologies to provide some specialized accelerated computational capability, such as on the Cray XD1 system. In order to take full advantage of the FPGA enabled nodes of the Cray XD1 system, avoiding the traditional problematic development times typical of custom digital circuit design, High Level Synthesis (HLS) methodologies appear to be a promising research avenue. An HLS allows one to obtain a synthesizable VHDL description, in a near automated way starting from a specific algorithm given in some High Level language (such as C, FORTAN etc.). In this presentation we describe the HLS methodology developed by Ylichron S.r.l. [a spin-off company of the Italian Agency for the New Technologies, the Energy and the Environment (ENEA)] in the framework of the HARWEST project and used by ENEA researchers to implement core algorithms utilizing the FPGA module of the XD1 system. A case study based on the ENEA-ORNL collaboration in the context of the implementation of some illustrative biological and molecular modeling kernels will be presented to describe the whole design flow, from the initial algorithm up to a working parallel prototype which employs both Opteron CPUS and FPGA resources.

11C Performance-XD1/FPGAs
9:50 Room B3

Compiling Software Code to FPGA-based Application Accelerator Processors in the XD1, *Doug Johnson, Chris Sullivan, and Stephen Chappell, Celoxica*

We present a programming environment that enables developers to quickly map and implement compute-intensive software functions as hardware accelerated components in an FPGA. As part of the XD1 application acceleration subsystem, these reconfigurable FPGA devices provide parallel processing performance that can deliver superlinear speed up for targeted applications. We will demonstrate how supercomputer developers can quickly harness this potential using their familiar design methodologies, languages and techniques.

12A User Services 10:30 Room B1

Case Study: Providing Customer Support in the High Performance Computing Environment, *Barbara Jennings (SNLA)*

This study presents the experience gained through the design and implementation of a tool to deliver HPC customer support. Utilized by developers, consultants, and customers, the results provide insights to developing useful on-line tools to assist this caliber of user to complete their work. This presentation will showcase the actual product demonstrating the model that was discussed at the CUG 2005 for delivering collaboration, learning, information and knowledge.

12A User Services 10:55 Room B1

Beyond Books: New Directions for Cray Customer Documentation, *Nola Van Vugt and Lynda Feng, Cray Inc.*

Last year you told us what you liked and what you didn't like about our customer documentation. We listened, we responded. Come see a demonstration of the new online help system that's part XT3 SMW GUI and see what else we are doing to make our documentation more accessible and up-to-date.

12B Optimizing Performance (Compilers & Tools) 10:30 Room B2

Optimizing Application Performance on Cray Systems with PGI Compilers and Tools, *Douglas Miles, Brent Leback, and David Norton, The Portland Group*

PGI Fortran, C and C++ compilers and tools are available on most Cray XT3 and Cray XD1 systems. Optimizing performance of the AMD Opteron processors in these systems often depends on maximizing SSE vectorization, ensuring alignment of vectors, and minimizing the number of cycles the processors are stalled waiting on data from main memory. The PGI compilers support a number of directives and options that allow the programmer to control and guide optimizations including vectorization, parallelization, function inlining, memory prefetching, interprocedural optimization, and others. In this paper we provide detailed examples of the use of several of these features as a means for extracting maximum single-node performance from these Cray systems using PGI compilers and tools.

WEDNESDAY (CONTINUED)

12C Performance—MTA 10:30 Room B3

Characterizing Applications on the MTA-2 Multithreading Architecture, *Richard Barrett, Jeffrey Vetter, Sadaf Alam, Collin McCurdy, and Philip Roth (ORNL)*

ORNL is currently evaluating various strategic applications on the MTA-2 architecture in order to better understand massively parallel threading as an architectural choice beyond scalar and vector architectures. In this paper, we describe our initial experiences with several applications including molecular dynamics, sparse matrix vector kernels, finite difference methods, satisfiability, and discrete event simulation.

13 General Session 11:20 Room B1

11:20 CUG Business/Elections, *Barbara Horner-Miller (ARSC), CUG President, Chair*

11:40 CUG and the Cray CEO: AKA, "100 on 1," *Pete Ungaro, CEO and President, Cray Inc.*

Cray Inc. will provide the Cray User Group (CUG) with an opportunity to go one-on-one with its CEO and President, Pete Ungaro. This will give attendees a chance to reflect on what they have heard at the conference and have a direct line to the CEO without any concern of upsetting others at Cray. The expectation is that CUG members will share their feelings and concerns with Cray Inc. openly, without having to temper their feedback. Pete will offer his thoughts and comments, and hopes that this session will provide Cray Inc. with valuable insight into the Cray User Group.

12:25 Election Results, *Kevin Wohlever (OSC), CUG Treasurer*

LUNCH 12:30 Atrio/Foyer, Ground Floor

14A XT3 Performance Analysis 1:45 Room B1

Enabling Computational Science on BigBen, *Nick Nystrom, Shawn Brown, Jeff Gardner, Roberto Gomez, Junwoo Lim, David O'Neal, R. Reddy, John Urbanic, Yang Wang, Deborah Weisser, and Joel Welling (PITTSCC)*

BigBen, PSC's 2090-processor Cray XT3, entered production on October 1, 2005 and now hosts a wide range of scientific inquiry. The XT3's high interconnect bandwidth and Catamount operating system support highly scalable applications, with the Linux-based SIO nodes providing valuable flexibility. In this talk and its associated paper, we survey applications, performance, utilization, and TeraGrid interoperability, and we present status updates on current PSC XT3 applications initiatives.

14A XT3 Performance Analysis 2:10 Room B1

Boundary Element–Finite Element Method for Solid-Fluid Coupling on the XT3, *Klaus Regenauer-Lieb (WASP)*

An Implicit Finite Element (FEM) solver is employed for the solid and an Implicit Boundary Elements (BEM) code is used for calculating only the drag caused from the surrounding fluid. The BEM approach is applied along horizontal planes (assumption validated by 3-D flow model). This application is implemented on the XT3 and results are presented.

14A XT3 Performance Analysis 2:35 Room B1

Analysis of an Application on Red Storm, *Courtenay Vaughan and Sue Goudy (SNLA)*

CTH is a widely used shock hydrodynamics code developed at Sandia. We will investigate scaling on Red Storm to 10,000 processors and will use those results to compare with an execution time model of the code.

14B Co-Array Fortran 1:45 Room B2

Scaling to New Heights with Co-Array Fortran, *Jef Dawson (NCS-MINN)*

Supercomputing users continually seek to solve more demanding problems by increasing the scalability of their applications. These efforts are often severely hindered by programming complexity and inter-processor communication costs. Co-Array Fortran (CAF) offers innate advantages over previous parallel programming models, both in programmer productivity and performance. This presentation will discuss the fundamental advantages of CAF, the importance of underlying hardware support, and several code examples showing how CAF is being effectively used to improve scalability.

14B Co-Array Fortran 2:10 Room B2

Porting Experience of Open64-based Co-Array Fortran Compiler on XT3, *Akira Nishida (UTOKYO)*

Co-Array Fortran (CAF) is one of the most hopeful technologies for seamless parallel computing on multicore distributed memory machines. In this presentation, we report our experience of porting an open source CAF compiler for Cray XT3.

14B Co-Array Fortran 2:35 Room B2

Co-Array Fortran Experiences Solving PDE Using Finite Differencing Schemes, *Richard Barrett (ORNL)*

A broad range of physical phenomena in science and engineering can be described mathematically using partial differential equations. Determining the solution of these equations on computers is commonly accomplished by mapping the continuous equation to a discrete representation and applying a finite differencing scheme. In this presentation we report on our experiences with several different implementations using Co-Array Fortran, with experiments run on the Cray X1E at ORNL.

14C System Support/Management
1:45 Room B3

Job-based Accounting for UNICOS/mp, *Jim Glidewell (BOEING)*

With the absence of CSA (Cray System Accounting) from UNICOS/mp, the ability of sites to provide job-based accounting is limited. We have developed tools for utilizing UNICOS process accounting to provide both end-of-job and end-of-day reporting based on jobid. We will provide a look at our strategies for doing so, a detailed look at the contents of the UNICOS/mp process record, and some suggestions for how sites might make use of this data.

14C System Support/Management
2:10 Room B3

Resource Allocation and Tracking System (RATS) Deployment on the Cray X1E, XT3, and XD1 Platforms, *Robert Whitten, Jr. and Tom Barron (ORNL); David L. Hulse, Computer Associates International, Inc.; Phillip E. Pfeiffer, East Tennessee State University; Stephen L. Scott (ORNL)*

The Resource Allocation and Tracking System (RATS) is a suite of software components designed to provide resource management capabilities for high performance computing environments. RATS is currently supporting resource allocation management on the Cray X1E, XT3 and XD1 platforms at Oak Ridge National Laboratory. In this paper, the design and implementation of RATS is given with respect to the Cray platforms with an emphasis on the flexibility of the design. Future directions will also be explored.

14C System Support/Management
2:35 Room B3

Red Storm Systems Management: Topics on Extending Current Capabilities, *Robert A. Ballance (SNLA); Dave Wallace, Cray Inc.*

This paper will discuss how by using an extensible System Description Language, current management capabilities can be extended and/or supplemented. We will discuss two implementations on Red Storm hardware and how these concepts can be easily extended to address other management issues.

15A Performance via Eldorado & MTA 3:20 Room B1

Eldorado Programming Environment, *Simon Kahan, Cray Inc.*

Cray's Eldorado programming model supports extremely concurrent computation on unpartitioned, shared data: hundreds of thousands of hardware streams may simultaneously execute instructions affecting arbitrary locations within large, dynamic data structures. While this flexible model enables programmers to use the most efficient parallel algorithms available, to do so induces demands on the system software that bottleneck other large parallel systems. This talk presents the Eldorado programming model, identifying and elaborating on some specific shared-memory challenges and how Eldorado uniquely addresses them.

15A Performance via Eldorado & MTA 4:00 Room B1

A Micro-kernel Benchmark Suite for Multithreaded Architectures, *Fabrizio Petrini and Allan Snavely (SDSC); John Feo, Cray Inc.*

The micro-benchmarks used to assess the performance of conventional parallel architectures, such as computing clusters, measure network bandwidth and latency and other network communication characteristics such as message startup and throughput. These benchmarks are not appropriate for shared-memory, multithreaded architectures such as the Cray MTA2 and Eldorado that have no cache and do not support message passing. In this paper we define a more appropriate suite of benchmarks for these machines. We discuss how the suite of kernels may be different for the MTA2 (an UMA machine) and Eldorado (a NUMA machine). Finally, we use the kernels to measure the performance of the MTA2 and predict the performance of Eldorado.

15B Libraries 3:20 and 4:00 Room B2

Cray and AMD Scientific Libraries, *Chip Freitag, AMD; Mary Beth Hribar, Bracy Elton and Adrian Tate, Cray Inc.*

Cray provides optimized scientific libraries to support the fast numerical computations that Cray's customers require. For the Cray X1 and X1E, LibSci is the library package that has been tuned to make the best use of the multistreamed vector processor based system. For the Cray XT3 and Cray XD1, AMD's Core Math Library (ACML) and Cray XT3/XD1 LibSci together provide the tuned scientific library routines for these Opteron based systems. This talk will summarize the current and planned features and optimizations for these libraries. And, we will present library plans for future Cray systems.

15C Operations 3:20 Room B3

Providing a Shared Computing Resource for Advancing Science and Engineering Using the XD1 at Rice University, *Jan E. Odegard, Kim B. Andrews, Franco Bladilo, Kiran Thyagaraja, Roger Moye, and Randy Crawford (RICEU)*

In this paper we will discuss the recent procurement and deployment of one of the largest XD1 systems in the world. The system, deployed by the Computer and Information Technology Institute at Rice University, is a major addition to the HPC infrastructure and will serve as the primary HPC resource supporting cutting edge research in engineering, bioX, science, and social science. At this point we envision that the majority of the paper will focus on hardware and software procurement, benchmarking, and deployment. The paper will also provide a view into how the system will be managed as well as discussions of performance, application porting, and tuning as well as discuss operation and support.

WEDNESDAY (CONTINUED)

15C Operations 4:00 Room B3

Safeguarding the XT3, *Katherine Vargo (PITTSCC)*

Service nodes and the management server provide critical services for the XT3 and its users. Using open-source software, these critical services can be monitored and safeguarded from common Linux attacks. Patch management, configuration verification, integrity checking, and benchmarking software provide a solid foundation for insulating services from assaults and enable system administrators to increase system availability and reliability.

16A XD1 Interactive Sessions
4:30 Room B1

Cray XD1 Special Interest Group Meeting (SIG), *Liam Forbes (ARSC), Chair*

Systems, *Liam Forbes (ARSC), Focus Chair*

Users, *Peter Cebull (INL), Focus Chair*

The purpose of this Interactive Session is to provide a forum for open discussions on the Cray XD1 system and its environment. The "Systems" area covers issues related to Operations, Systems & Integration, and Operating Systems. The "Users" area deals with issues related to Applications, Programming Environments, and User Services. Cray Inc. liaisons will be on hand to provide assistance and address questions. This session is open to everyone that has an interest in the Cray XD1 and wishes to participate.

16B Interactive Sessions 4:30 Room B2

Reserved for interactive session.

16C BoF 4:30 Room B3

Batch Schedulers, *Richard Alexander (CSCS)*

The objective of this BoF will be to discuss the various batch schedulers. It will not discuss batch systems per se, but rather, the scheduling philosophies and implementations of each of the following sites:

A) CSCS with a TCL based scheduler for PBS
B) PSC with original work on external PBS scheduling
C) ORNL with MOAB, the commercial version of MAUI, for PBS
D) ERDC with LSF

Bordering on both Switzerland and Italy Lake Lugano represents one of the greatest attractions to the area.

THURSDAY

17A FPGAs & Supercomputers
8:20 Room B1

Turning FPGAs Into Supercomputers—Debunking the Myths About FPGA-based Software Acceleration, *Anders Dellson, Mitrionics AB*

When considering whether to use field programmable gate arrays (FPGAs) as co-processors to accelerate software algorithms, it is easy to get the impression that this is an option limited to a small elite group of highly skilled and heavily funded hardware savvy researchers. In this presentation, we will show that this is no longer true and that the 10x to 30x performance benefit and the low power consumption of FPGA-based software acceleration are now readily available to all researchers and software developers. As an example, we will show an actual application that has been accelerated using the FPGAs in the Cray XD1 and walk through the steps needed for the implementation using the Mitrion Platform.

17A FPGAs & Supercomputers
9:00 Room B1

Experiences with High-Level Programming of FPGAs on Cray XD1, *Thomas Steinke, Konrad Zuse-Zentrum für Informationstechnik Berlin; Thorsten Schuett and Alexander Reinefeld, Zuse-Institut Berlin*

We recently started to explore the potential of FPGAs for some application kernels that are important for our user community at ZIB. We give an overview of our efforts and present results achieved so far on a Cray XD1 using the Mitrion-C programming environment.

17A FPGAs & Supercomputers
9:25 Room B1

Experiences Harnessing Cray XD1 FPGAs and Comparisons to other FPGA High Performance Computing (HPC) Systems, *Olaf Storaasli, Melissa C. Smith, and Sadaf R. Alam (ORNL)*

The Future Technologies Group at Oak Ridge National Laboratory (ORNL) is exploring advanced HPC architectures including FPGA-based systems. This paper describes experiences with applications and tools (including Mitrion C) on ORNL's Cray XD1 and comparisons with results from other FPGA-based systems (SRC, Nallatech, Digilent) and tools (VIVA, Carte) at ORNL.

17A FPGAs & Supercomputers
9:50 Room B1

Status Report of the OpenFPGA Initiative: Efforts in FPGA Application Standardization, *Eric Stahlberg and Kevin Wohlever (OSC)*; *Dave Strenski, Cray Inc.*

The OpenFPGA initiative began officially in early 2005. From early formative discussions, the effort has matured to become a community resource to foster and advance the adoption of FPGA technology for high-level applications. Participation includes vendors, application providers, and application users across academic, commercial and govern-

mental organizations. With widespread and international participation by hundreds of participants, an active discussion list on FPGA related topics, and emerging working groups, the OpenFPGA effort is building a strong foundation for broad incorporation of FPGA accelerated high-level applications. This presentation will cover insight from early discussions, current organizational overview, future directions towards standardization, and information for becoming part of the OpenFPGA effort.

17B Parallel Programming 8:20 Room B2

Hybrid Programming Fun: Making Bzip2 Parallel with MPICH2 and pthreads on the XD1, *Charles Wright (ASA)*

The author shares his programming experience and performance analysis in making a parallel file compression program for the XD1.

17B Parallel Programming 9:00 Room B2

Parallel Performance Analysis on Cray Systems, *Craig Lucas and Kevin Roy (MCC)*

Solving serial performance problems is very different to solving the performance problems on parallel codes. In many cases the behaviour of an application can depend on the number of processors used. Some parts of the code scale well and others don't; when we look at parallel codes it is these nonscaling parts that are of most interest. Identifying the components and when they become a problem is not often an easy task. In this talk we present some software that interfaces to existing Cray performance analysis tools and presents clear and easy to view information.

17B Parallel Programming 9:25 Room B2

An Evaluation of Eigensolver Performance on Cray Systems, *Adrian Tate and John G. Lewis, Cray Inc.; Jason Slemons, University of Washington*

Solutions to the Eigenvalue equation are the computational essence of many large-scale applications, especially in engineering and physical chemistry. Eigensolver performance is expected to vary according to the physical properties of the underlying data, but when run on parallel systems, performance is also heavily dependent on the subtleties of interconnect network and the message passing implementation of the particular system. We will compare the performance of various eigensolvers across each of Cray's products. We will show which systems are better suited to certain methods and show how the performance of some real applications can be maximized by careful selection of method. We will also describe what research has been carried out, both inside and outside of Cray, to improve performance of parallel eigensolvers.

17B Parallel Programming 9:50 Room B2

Symmetric Pivoting in ScaLAPACK, *Craig Lucas (MCC)*

While the ScaLAPACK parallel numerical linear algebra library supports symmetric matrices, there is no symmetric (or "complete") pivoting in its algorithms. Here we look at an implementation of the pivoted Cholesky factorization of semi-definite matrices, which uses symmetric pivoting and is built on existing PBLAS and BLACS routines. We examine the performance of the code which is written in ScaLAPACK style and therefore uses block cyclic data distribution.

17C Mass Storage 8:20 and 9:00 Room B3

Lustre File System Plans & Performance on Cray Systems, *Charlie Carroll and Branislav Radovanovic, Cray Inc.*

Lustre is an open source, high-performance, distributed file system from Cluster File Systems, Inc. designed to address performance, availability, and scalability issues. A number of high-performance parallel Lustre file systems for supercomputers and computer clusters are available today. This talk will discuss Cray's history with Lustre, the current state of Lustre, and Cray's Lustre plans for the future. We will also provide recent performance results on a Cray XT3 supercomputer, showing both scalability and achieved performance.

17C Mass Storage 9:25 Room B3

XT3 File I/O Offloading, *Paul Nowoczynski, Jason Sommerfield, Doug Balog, and J. Ray Scott (PITTSCC)*

The offloading of Lustre I/O traffic onto remote storage servers may provide several possible advantages in the areas of performance and increased storage manageability. Foremost will be the exploration of end-to-end service node throughput via the Lustre Routing Interface between Seastar and both Infiniband and 10 Gigabit Ethernet. This paper will explore the pros and cons of such an effort as well as describe the configuration and setup procedures.

17C Mass Storage 9:50 Room B3

A Center Wide File System Using Lustre, *Shane Canon and H. Sarp Oral (ORNL)*

The National Leadership Computing Facility at Oak Ridge National Laboratory is currently deploying a Lustre based center wide file system. This file system will span multiple architectures and must meet demanding user requirements. The objectives for this file system will be presented along with an overview of the architecture. A discussion of issues that have been encountered during the deployment, as well as current performance numbers, will also be provided. The presentation will conclude with future plans and goals.

18A Operating Systems—UNICOS 10:40 Room B1

Recent Trends in Operating Systems and their Applicability to HPC, *Rolf Riesen and Ron Brightwell (SNLA); Patrick Bridges and Arthur Maccabe, University of New Mexico*

In this paper we consider recent trends in operating systems and discuss their applicability to high performance computing systems. In particular, we will consider the relationship between lightweight kernels, hypervisors, microkernels, modular kernels, and approaches to building systems with a single system image. We will also give a brief overview of the approaches being considered in the DoE FASTOS program.

THURSDAY (CONTINUED)

18A Operating Systems—UNICOS
11:20 Room B1

Compute Node OS for XT3, *Jim Harrell, Cray Inc.*

Cray is working on different kernels for the compute node operating system. This talk will describe the rationale, requirements, and progress.

18A Operating Systems—UNICOS
11:45 Room B1

XT3 Status and Plans, *Charlie Carroll and David Wallace, Cray Inc.*

This paper will discuss the current status of XT3 software and development plans for 2006.

18B Libraries—MPI 10:40 Room B2

Message Passing Toolkit (MPT) Software on XT3, *Howard Pritchard, Doug Gilmore, and Mark Pagel, Cray Inc.*

This talk will focus on using the Message Passing Toolkit (MPT) for XT3. An overview of the MPI and SHMEM implementations on XT3 will be given. Techniques for improving the performance of message passing and SHMEM applications on XT3 will also be discussed. In addition, new features in the MPT package for XT3 1.3 and 1.4 releases will be presented.

18B Libraries—MPI 11:20 Room B2

Open MPI on the XT3, *Brian Barrett, Jeff Squyres, and Andrew Lumsdaine, Indiana University; Ron Brightwell (SNLA)*

The Open MPI implementation provides a high performance MPI-2 implementation for a wide variety of platforms. Open MPI has recently been ported to the Cray XT3 platform. This paper discusses the challenges of porting and describes important implementation decisions. A comparison of performance results between Open MPI and the Cray supported implementation of MPICH2 are also presented.

Together with the camelias that are celebrated at the annual Camelia Festival in Locarno, magnolias announce the beginning of the warm season and all its outdoor activities.

18B Libraries—MPI 11:45 Room B2

What if MPI Collectives Were Instantaneous?, *Rolf Riesen and Courtenay Vaughan (SNLA)*

MPI collectives, such as broadcasts or reduce operations, play an important role in the performance of many applications. How much would these applications benefit, if we could improve collectives performance by using better data distribution and collection algorithms, or move some functionality into the Seastar firmware, closer to the network? We answer these questions by simulating a Cray XT3 system with hypothetical, instantaneous MPI collectives.

18C Eldorado/MTA-2 10:40 Room B3

Graph Software Development and Performance on the MTA-2 and Eldorado, *Jonathan Berry and Bruce Hendrickson (SNLA)*

We will discuss our experiences in designing and using a software infrastructure for processing semantic graphs on massively multithreaded computers. We have developed implementations of several algorithms for connected components and subgraph isomorphism, and we will discuss their performance on the existing Cray MTA-2, and their predicted performance on the upcoming Cray Eldorado. We will also describe ways in which the underlying architecture and programming model have informed algorithm design and coding paradigms.

18C Eldorado/MTA-2 11:20 Room B3

Evaluation of Active Graph Applications on the Cray Eldorado Architecture, *Jay Brockman, Matthias Scheutz, Shannon Kuntz, and Peter Kogge, University of Notre Dame; Gary Block and Mark James, NASA JPL; John Feo, Cray Inc.*

In this paper, we discuss an approach to such problems called an "active graph," where each node in the graph is tied to a distinct thread and where the flow of data over the edges is expressed by producer/consumer exchanges between threads. We will show how several cognitive applications, including a neural network and a production system, may be expressed as active graph problems and provide results on how active graph problems scale on both the MTA-2 and Eldorado architectures.

Ticino is known the world over for its famous architects.

18C Eldorado/MTA-2 11:45 Room B3

Scalability of Graph Algorithms on Eldorado, *Keith Underwood, Megan Vance, Jonathan Berry, and Bruce Hendrickson (SNLA)*

The Eldorado platform is a successor to the Cray MTA-2 that is being built within the Cray XT3 infrastructure. The XT3 infrastructure brings with it a number of limitations in terms of network bisection bandwidth and random access memory access that are not present in the MTA-2. While the MTA-2 is an exceptional performer on graph algorithms, the new limitations being introduced by the Eldorado platform could have negative implications for scaling. This paper analyzes Eldorado to explore the domain of applications requirements to enable scaling. By mapping current graph algorithms into this domain, we find that, while Eldorado will not scale quite as well as the MTA-2, it will scale sufficiently to offer orders of magnitude better graph performance than any other platform that is available.

LUNCH 12:10 Atrio/Foyer, Ground Floor

19A Sizing: Page, Cache, & Meshes 1:30 Room B1

The Effect of Page Size and TLB Entries on XT3 Application Performance, *Neil Stringfellow (CSCS)*

The AMD Opteron processor allows for two page sizes to be selected for memory allocation, with a small page size of 4 Kilobytes being used for most standard Linux systems, and a larger 2 Megabyte page size which was selected for the Catamount lightweight kernel. Although the larger page size appears more attractive for HPC applications running under a lightweight kernel, the increase in page size comes at a cost in that there are very few entries in the TLB available to store references to these pages. This paper reports on work which highlighted problems with the use of the large page size and small number of TLB entries and shows some of the performance improvements which became possible with the introduction of a small page option to the yod job launcher.

19A Sizing: Page, Cache, & Meshes 2:00 Room B1

Performance Characteristics of Cache-Insensitive Implementation Strategies for Hyperbolic Equations on Opteron Based Super Computers, *David Hensinger and Chris Luchini (SNLA)*

For scientific computation moving data from RAM into processor cache incurs significant latency, and comes at the expense of useful calculation. Cache-insensitive strategies attempt to amortize the costs of movement of high latency data by carrying out as many operations as possible using in-cache data. For explicit methods applied to hyperbolic equations, cache insensitive algorithm advance regions of the solution domain through multiple time steps. Success of these strategies requires the performance advantage associated with computing in-cache to offset the overhead associated with managing multiple solution states. The performance characteristics of several cache insensitive implementation strategies were tested. The implications for multiprocessor simulations were also examined.

19A Sizing: Page, Cache, & Meshes 2:30 Room B1

Investigations on Scaling Performance and Mesh Convergence with Sandia's ASC FUEGO Code for Fire Model Predictions of Heat Flux, *Mahesh Rajan and Amalia Black (SNLA)*

Performance characteristics for coupled fire/thermal response prediction simulations are investigated using coarse, medium and fine finite-element meshes and running the simulation on the Red Storm/XT3. These simulations have leveraged computationally demanding mesh convergence studies to obtain detailed timings of the various phases of the computation and will be helpful in performance tuning.

19B XD1 Applications 1:30 Room B2

Alef Formal Verification and Planning System, *Samuel Luckenbill, James R. Ezick, Donald D. Nguyen, Peter Szilagyi, and Richard A. Lethin, Reservoir Labs, Inc.*

We have developed a parallel SAT solver using algorithms and heuristics that improve performance over existing approaches. We have also developed compiler technologies to transform problems from several problem domains to our solver. Our system runs on the Cray XD1 machine and particularly makes use of the high performance communication hardware and interface library.

19B XD1 Applications 2:00 Room B2

XD1 Implementation of a SMART Coprocessor for Fuzzy Matching in Bioinformatics Applications, *Eric Stahlberg and Ben Smith (OSC)*

Efficiently matching or locating small nucleotide sequences on large genomes is a critical step in the Serial Analysis of Gene Expression method. This challenge is made increasingly difficult as a result of experimental assignment errors introduced in the match and target sequence data. Early versions of SAGESpy were developed incorporating the Cray Bioinformatics Libraries to enable large-scale pattern matching of this type. This presentation describes the design of an FPGA-based scalable fuzzy DNA sequence matching algorithm implemented specifically for the XD1. Results of this implementation will be compared to performance on earlier Cray SV1 and current Cray X1 systems.

19B XD1 Applications 2:30 Room B2

Simulating Alzheimer on the XD1, *Jan H. Meinke and Ulrich H. E. Hansmann (KFA)*

Misfolding and aggregation of proteins are the causes for Alzheimer, BSE, and other neurodegenerative diseases. We simulate the aggregation of a 7 amino acid long fragment of the protein Abeta, responsible for the formation of toxic fibrils in Alzheimer. We use parallel tempering, an advanced Monte Carlo algorithm that scales nearly optimally.

THURSDAY (CONTINUED)

19C Benchmarking/Comparison
1:30 Room B3

Performance Comparison of Cray X1 and Cray Opteron Cluster with Other Leading Platforms Using HPCC and IMB Benchmarks, *Rolf Rabenseifner (HLRS); Subhash Saini and Robert Ciotti (NAS); Brian T. N. Gunney, Thomas E. Spelce, Alice Koniges, and Don Dossa, Lawrence Livermore National Laboratory; Panagiotis Adamidis and Sunil R. Tiyyagura (HLRS); Matthias Mueller, Dresden University of Technology; Rod Fatoohi, San Jose State University*

The HPC Challenge (HPCC) benchmark suite and the Intel MPI Benchmark (IMB) are used to compare and evaluate the combined performance of processor, memory subsystem and interconnect fabric of Cray X1 and Cray Opteron Cluster together with SGI Altix BX2, Dell Xeon cluster, IBM Blue/Gene and NEC SX-8. These six systems use six different networks (Cray X1 network, Myrinet, SGI NUMAlink4, InfiniBand, IBM Blue/Gene 3D torus, and NEC IXS).

19C Benchmarking/Comparison
2:00 Room B3

Performance Analysis of Cray X1 and Cray Opteron Cluster, *Subhash Saini (NAS); Rod Fatoohi, San Jose State University; Johnny Chang and Robert Ciotti (NAS)*

We study the performance of two Cray systems (Cray X1 and Cray Opteron Cluster) and compare their performance with an SGI Altix 3700 system. All three systems are located at NASA Ames Research Center. We mainly focus on network performance using different number of communicating processors and communication patterns—such as point-to-point communication, collective communication, and dense communication patterns. Our results show the impact of the network bandwidth and topology on the overall performance of each system.

19C Benchmarking/Comparison
2:30 Room B3

An Accelerated Implementation of Portals on the Cray SeaStar, *Ron Brightwell, Trammell Hudson, Kevin Pedretti, and Keith D. Underwood (SNLA)*

This paper describes an accelerated implementation of the Portals data movement layer on the Cray SeaStar used in the XT3 platform. The current supported implementation of Portals is interrupt-driven and does not take full advantage of the embedded processor on the SeaStar. The accelerated implementation offloads a significant portion of the network protocol stack to the SeaStar, providing significant improvements in network performance. This paper will describe this new implementation and show results for several network micro-benchmarks as well as applications.

The Piazza della Riforma is the hub of Lugano's shopping centre that also hosts the weekly market and on Sunday becomes a favorite spot for enjoying sun, espresso and people-watching.

20 General Session 3:10 Room B2

3:10 The AWE HPC Benchmark 2005, *Ron Bell, S. Hudson, and N. Munday (AWE)*

Near the end of 2005, the UK's Atomic Weapons Establishment (AWE) placed an order for a Cray XT3 system with 3936 dual core nodes (over 40 Tflops peak) to replace its existing HPC system. This paper describes the design of the benchmark used during the competitive procurement preceding this order and presents details of the evaluation process and the results. These include the more than 2-times speedup obtained by Cray by tuning the source code of the most important application.

3:50 Next CUG (Seattle), *Hank Heeb (BOEING)*

4:00 LAC Appreciation and Adjournment, *CUG President*

Architectural characteristics such as this narrow passage convey a feeling of the proximity to Italy.

Local Arrangements

How to Contact Us

Until May 15, 2006, and after the conference

Marie-Christine Sawley

Local Arrangements Chair

CEO, CSCS—Swiss National Supercomputing Centre
Galleria 2
Via Cantonale
6928 Manno
Switzerland

(41-91) 610 82 01 Fax: (41-91) 610 83 08

Ladina Gilly

Local Arrangements Coordinator

Event Manager, CSCS—Swiss National Supercomputing Centre
Galleria 2
Via Cantonale
6928 Manno
Switzerland
(41-91) 610 82 67 Fax: (41-91) 610 83 08

During the conference, May 8–11, 2006

CUG 2006 Conference
Palazzo dei Congressi
Piazza Indipendenza 4
6900 Lugano
(41-91) 911 04 25 Fax: (41-91) 911 04 24

Conference Registration

CUG Office
3224 Southwood Drive
Philomath, Oregon 97370 USA
(1-541) 929-3776 Fax: (1-509) 272-7239
bobwinget@toolboxcm.com

Attendance and Registration

Who May Attend?

CUG meetings are not open sessions, according to the organization's corporate bylaws and guidelines. Admittance to any meeting session requires the appropriate identification badge. Generally speaking, employees of a CUG Member site (usually a facility or company using a Cray Inc. computer and identified by its CUG site code), and users of computing services provided by a CUG member site, may attend a CUG meeting as a registered Installation Participant from that CUG Member site.

Additionally, CUG bylaws specify that a "CUG Visitor" is "Any individual who is not an Installation Representative and who is invited to attend a specific function of the Corporation ... [with] the prior approval of the Board of Directors." This generally applies to an occasional invited speaker who brings important perspectives to our meeting in support of a strong technical program. Finally, CUG bylaws state that, "Meetings shall not be used for marketing or other commercial purposes."

Conference Registration Location/Hours

Conference registration is located on the first floor of the Palazzo dei Congressi.

Badges and registration materials are available during these scheduled hours. All attendees must wear badges during CUG Conference activities.

Office Hours

Sunday .....3:00 p.m.–6:00 p.m.

Monday–Wednesday .....7:30–a.m.–5:30 p.m.

Thursday .....7:30 a.m.–3:00 p.m.

Conference Registration Fees

Your registration fee includes

- Admission to all program sessions, meetings, and tutorials
- Lunch and refreshments
- CUG Night Out

Proceedings

All conference attendees will have access to the conference Proceedings on our secure web server and will receive them on CD-ROM distributed approximately two months after the conference.

Cancellations

Conference registration cancellations must be received by the CUG Office before April 29, 2006. All registration fees will be refunded (less a $25 handling fee) if the cancellation is received by this date.

Conference Location

The conference will be held at the Palazzo dei Congressi in Lugano, Switzerland.

Hotel Booking

Hotel bookings are managed through the Lugano Tourism office. Given the size and layout of hotels in Lugano the booking system has been modified for this conference. The Tourism office will try to accommodate your requests as far as possible. We thank you for your understanding.

Special Assistance

For special assistance during the conference come to the Conference Office.

Dining Services

Breakfast, Lunch, and Refreshments

Breakfast is included in the price of all hotel rooms. Lunch is included in the conference registration fee and will be served Monday through Thursday in the ground floor Foyer of the Palazzo dei Congress.

Refreshments will be provided in the first floor Foyer during the morning and afternoon breaks.

Dinner on Wednesday

Dinner Wednesday evening is provided at the CUG Night Out for all attendees who have registered for the full conference. Tickets for attendees' guests are $40 and are available in the Conference Office. This special event is organized by the Local Arrangements Committee.

Dining Out

A wide variety of restaurants are available in downtown Lugano. Local instinct is to go for the full ones as that is usually a sign of good food.

On-Site Facilities

Messages

During the conference (May 8-11), telephone messages will be taken at the on-site CUG conference office at (41-91) 911 04 25.

Faxes

A fax machine will be available at the conference office. The fax number during the conference (May 8-11) will be (41-91) 911 04 24. Notices for incoming fax messages will be posted on the conference notice board.

Wireless Access, E-mail, and Personal Computers

There is wireless access for your laptop computer throughout the conference venue for which you must provide an 802.11b/g compatible card. Please make sure your wireless access has not been disabled by your company/institution as we will not be able to modify this on site.

The e-mail room will be equipped with 8 personal computers. SSH will be available on all systems. Two printers will be available on the LAN. One computer and printer are designated for the use of speakers for last minute changes to presentations. You may also bring your laptop computer and connect to the ethernet hub. The ethernet hub accepts RJ-45 modular jacks.

E-mail room hours are:

Monday–Wednesday ....7:00 a.m.–9:00 p.m.

Thursday .....7:00 a.m.–4:00 p.m.

Speaker Preparation

We invite speakers to use the personal computer and printer in the e-mail room for last minute changes. Video projectors will be provided in all session rooms. Please find time before sessions and during breaks to test your laptop connection to the video projectors. Should you require any assistance please come to the conference office.

Business Services and Photocopying

A photocopying machine will be available on site. However, if you plan to distribute copies of your presentation, please bring sufficient copies with you. If you need to make copies while you are at the conference, you may go to COLOR X S.A. in via Frasca 1, Lugano (Phone: 091 923 58 91).

Voltage

Switzerland operates on a 240V system; those traveling from abroad may need an adaptor. Check the Voltage Valet at voltagevalet.com/idx.html for information on requirements and to find adapters.

Travel Information

Lugano can be reached via the airports of Lugano (connect through Zurich or Geneva), Zurich (3 hour train ride to Lugano) or Milano Malpensa that has a regular shuttle bus connection to Lugano (1 hour ride). Make sure you book the shuttle in advance.

Ground Transportation

Useful links:

- Swiss Rail—www.sbb.ch (English site available)
- Malpensa Shuttle—www.starbus.info
- Connection from Geneva —www.darwinairline.com

We strongly advise you to use public transport as it is an efficient and easy solution to the very congested traffic situation in Lugano.

Currency

Although Euros are accepted in some places, Switzerland works with Swiss Francs. It is therefore advisable to have the local currency with you. You can exchange foreign currencies at any of the banks around the conference venue or in town.

Get an update on the current exchange rate at http://www.oanda.com/convert/classic.

Weather

May is spring season for Lugano. The weather can vary quite a bit during this season but you should expect daytime temperatures between 10° and 20°C. Check http://www.meteoswiss.ch/web/en/weather.html for a current update on the weather situation.

Tourist Information

Tourist information is available from the CUG 2006 Local Arrangements pages on cug.org. Brochures and additional information will be available on site at the conference.

Social Events

Cray Inc. Reception

Cray Inc. will host a reception Tuesday evening at 7:00 PM at Ristorante Parco Ciani (on site at the Palazzo dei congressi). All conference participants and their guests are invited to attend.

CRAY THE SUPERCOMPUTER COMPANY

The Cray User Group is especially grateful to Cray Inc. for their ongoing support as we celebrate 28 years of CUG conferences and workshops.

CUG Night Out

On Wednesday, all attendees are invited to enjoy a night out organized by the Local Arrangements Committee. The CUG Night Out will be held at one of Bellinzona's medieval castles, which are part of the UNESCO world heritage. The bus ride to the venue will take approximately 20 minutes.

During the night you will have the opportunity to taste some local specialties ranging from traditional dried meats, fresh water fish, cheese, and vegetables to the well known Risotto and meat dishes and last but not least—the local wines.

Buon appetito e buon divertimento!

The three medieval castles and the fortifications of the town of Bellinzona have been named a world heritage site by UNESCO.

AMD

Altair Engineering

DataDirect NETWORKS

We thank Cray partners AMD and Altair for their sponsorship of the CUG night out and DataDirect Networks for their sponsorship.of the coffee breaks.

Contacts

CUG Board of Directors

President
Barbara Horner-Miller
Arctic Region Supercomputing Center
909 Koyukuk Dr, Suite 105
P.O. Box 6020
Fairbanks, AK 99775-6020 USA
(1-907) 450-8629 Fax: (1-907) 450-8601
horner@arsc.edu

Vice President
David Gigrich
The Boeing Company
P.O. Box 3707
MS 67-EE
Seattle, WA 98124-2207 USA
(1-425) 237-3003 Fax: (1-425) 234-0408
david.j.gigrich@boeing.com

Treasurer
Kevin Wohlever
Ohio Supercomputer Center–Springfield
One South Limestone St.
Suite 310
Springfield, OH 45502 USA
(1-937) 322-2743 Fax: (1-937) 322-7869
kevin@osc.edu

Secretary
Rolf Rabenseifner
High Performance Computing Center Stuttgart
Nobelstr. 19
D-70550 Stuttgart Germany
(49-711) 685-5530 Fax: (49-711) 685-5832
rabenseifner@hlrs.de

Director—Americas
James B. White III
Oak Ridge National Laboratory
ORNL MS-6008
1 Bethel Valley Road
Oak Ridge, TN 37831-6008 USA
(1-865) 241-2103 Fax: (1-865) 241-2850
trey@ornl.gov

Director—Europe
Mike Pettipher
University of Manchester
Kilburn Building
Oxford Road
Manchester, Lancashire M13 9PL United Kingdom
(44-161) 275 6095 Fax: (44-161) 275 6800
Mike.Pettipher@manchester.ac.uk

Director-at-Large
John Noe
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-0807 USA
(1-505) 844-5592
jpnoe@sandia.gov

Local Arrangements Chair

Marie-Christine Sawley
CEO, CSCS–Swiss National Supercomputing Centre
Galleria 2
Via Cantonale
6928 Manno Switzerland
(41-91) 610 82 01 Fax: (41-91) 610 83 08

Advisory Council Secretary and CUG.log Editor

Angela Harris
Oak Ridge National Laboratory
One Bethel Valley Road
PO Box 2008, MS-6164
Oak Ridge, TN 37831-6164 USA
(1-865) 241-2863 Fax: (1-865) 241-2850
harrisad@ornl.gov

Regional CUG Representatives

German CUG Representative
Uwe Wolfram
Amt fuer Militaerkunde
Am Nippenkreuz 19
D-53179 Bonn Germany
(49-228) 350176280 Fax: (49-228) 856110
rz@amk-wfb.de

Cray Inc. Representatives to CUG

Vito Bongiorno, Business Liaison
Cray Inc.
1340 Mendota Heights Road
Mendota Heights, MN 55120 USA
(1-651) 605-8935 Fax: (1-651) 605-9123
vb@cray.com

Steve Scott, Technical Liaison
Cray Inc.
1050 Lowater Road
P.O. Box 6000
Chippewa Falls, WI 54729 USA
(1-715) 726-4757
sscott@cray.com

Ellen Roder, Program Liaison
Cray Inc.
1340 Mendota Heights Road
Mendota Heights, MN 55120 USA
(1-651) 605-9058 Fax: (1-651) 605-9123
ellen@cray.com

Linda Yetzer, Arrangements Liaison
Cray Inc.
1340 Mendota Heights Road
Mendota Heights, MN 55120 USA
(1-651) 605-8961 Fax: (1-651) 605-9123
yetzerll@cray.com

Contacts

Special Interest Groups

Special Interest Groups (SIGs) are divided into two areas (Users and Systems). The Users SIG deals with issues related to Applications, Programming Environments, and User Services. The Systems SIG covers issues related to Operations, Systems & Integration, and Operating Systems. The contacts for CUG Special Interest Groups, by platform, are shown in the table below. The CUG SIG Chairs, Cray Inc. Representatives to CUG, Local Arrangements Chair, and CUG Office, are members of the Program Lead Committee and are the leaders who develop CUG programs in support of the Program Chair.

XD1 SIG, Chair: Liam Forbes (ARSC)

Users

Chair: Peter Cebull (INL)

Cray Inc. SIG Liaisons

Applications: Kristyn Maschhoff
Programming Environments: Amar Shan
User Services: Nola Van Vugt

Systems

Chair: Liam Forbes (ARSC)

Cray Inc. SIG Liaisons

Systems & Integration: Amar Shan
Operating Systems: Amar Shan
Operations: Charlie Clark

XT3 SIG, Chair: James Kasdorf (PITTSCC)

Users

Chair: Thomas Oppe (CSC-VBURG)

Deputy Chair: James Kasdorf (PITTSCC)

Cray Inc. SIG Liaisons

Applications: Kristyn Maschhoff
Programming Environments: Luiz DeRose
User Services: Nola Van Vugt

Systems

Chair: Bob Ballance (SNLA)

Deputy Chair: Keith Underwood (SNLA)

Cray Inc. SIG Liaisons

Systems & Integration: Jim Harrell
Operating Systems: Mike McCardle
Operations: Charlie Clark

X1/E SIG, Chair: Mark Fahey (ORNL)

Users

Chair: Mark Fahey (ORNL)

Deputy Chair: Rolf Rabenseifner (HLRS)

Cray Inc. SIG Liaisons

Applications: Kristyn Maschhoff
Programming Environments: Mary Beth Hribar
User Services: Nola Van Vugt

Systems

Chair: : Jim Glidewell (BOEING)

Deputy Chair: Brad Blasing (NCS-MINN)

Cray Inc. SIG Liaisons

Systems & Integration: Steve Johnson
Operating Systems: Wayne Margotto
Operations: Charlie Clark

SIG Chairs and Cray Inc. SIG Liaisons

Robert Ballance
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-1110 USA
(1-505) 284-9361 Fax: (1-505) 844-2067
raballa@sandia.gov

Bradford Blasing
Network Computing Services, Inc.
1200 Washington Ave. South
Minneapolis, MN 55415 USA
(1-612) 337-3525 Fax: (1-612) 337-3400
brad@networkcs.com

Peter Cebull
Idaho National Laboratory
P.O. Box 1625
MS 3605
Idaho Falls, ID 83415-3605 USA
(1-208) 526-1909 Fax: (1-208) 526-4017
Peter.Cebull@inl.gov

Charles Clark
Cray Inc.
1050 Lowater Road
Chippewa Falls, WI 54729 USA
(1-715) 726-4020 Fax: (1-715) 726-4991
cmc@cray.com

Luiz DeRose
Cray Inc.
1340 Mendota Heights Road
Mendota Heights, MN 55120 USA
(1-651) 605-8925 Fax: (1-651) 605-9001
ldr@cray.com

Mark Fahey
Oak Ridge National Laboratory
P.O. Box 2008, MS 6008
Oak Ridge, TN 37831-6008 USA
(1-865) 574-5622 Fax: (1-865) 241-2850
faheymr@ornl.gov

SIG Chairs and Cray Inc. SIG Liaisons (continued)

Luiz DeRose
Cray Inc.
1340 Mendota Heights Road
Mendota Heights, MN 55120 USA
(1-651) 605-8925 Fax: (1-651) 605-9001
ldr@cray.com

Mark Fahey
Oak Ridge National Laboratory
P.O. Box 2008, MS 6008
Oak Ridge, TN 37831-6008 USA
(1-865) 574-5622 Fax: (1-865) 241-2850
faheymr@ornl.gov

Liam Forbes
Arctic Region Supercomputing Center
P.O. Box 756020
909 Koyukuk Drive
Fairbanks, AK 99775-6020 USA
(1-907) 450-8618
lforbes@arsc.edu

Jim Glidewell
The Boeing Company
PO Box 3707 M/C 7J-04
Seattle, WA 98124-2207 USA
(1-425) 865-1516 Fax: (1-425) 865-2221
james.glidewell@boeing.com

Jim Harrell
Cray Inc.
1345 Northland Dr.
Mendota Heights, MN 55121 USA
(1-651) 605-9131 Fax: (1-651) 683-7115

Mary Beth Hribar
Cray Inc.
411 First Avenue South, Suite 600
Seattle, WA 98104 USA
(1-206) 701-2072 Fax: (1-206) 701-2500
marybeth@cray.com

Steve Johnson
Cray Inc.
1050 Lowater Road
Chippewa Falls, WI 54729 USA
(1-715) 726-4816 Fax: (1-715) 726-4739
sjj@cray.com

James Kasdorf
Pittsburgh Supercomputing Center
4400 Fifth Avenue
MI-417
Pittsburgh, PA 15213 USA
(1-412) 268-6802 Fax: (1-412) 268-1490
kasdorf@psc.edu

Wayne Margotto
Cray Inc.
1340 Mendota Heights Rd.
Mendota Heights, MN 55120 USA
(1-651) 605-9062
wtm@cray.com

Kristyn Maschhoff
Cray Inc.
411 First Avenue South, Suite 600
Seattle, WA 98104 USA
(1-206) 701-2041 Fax: (1-206) 701-2500
kristyn@cray.com

Mike McCardle
Cray Inc.
411 First Avenue South, Suite 600
Seattle, WA 98104 USA
(1-206) 701-2286 Fax: (1-206) 701-2500
mikem@cray.com

Thomas Oppe
Computer Sciences Corporation–Vicksburg
ERDC MSRC
3909 Halls Ferry Road, ATTN: CEERD-IH-C
Vicksburg, MS 39180-6199 USA
(1-601) 634-2797 Fax: (1-601) 634-2324
Thomas.C.Oppe@erdc.usace.army.mil

Rolf Rabenseifner
High Performance Computing Center Stuttgart
Nobelstr. 19
D-70550 Stuttgart Germany
(49-711) 685 5530 Fax: (49-711) 6787 626
rabenseifner@hlrs.de

Amar Shan
Cray Canada
4621 Canada Way, Ste 301
Burnaby, BC V5G 4X8 Canada
(1-604) 484-2253 Fax: (1-604) 484-2221
shan@cray.com

Keith Underwood
Sandia National Laboratories
P.O. Box 5800
Mailstop 1110
Albuquerque, NM 87185-1110 USA
(1-505) 284-9283
kdunder@sandia.gov

Nola Van Vugt
Cray Inc.
411 First Avenue S., Suite 600
Seattle, WA 98104 USA
(1-206) 701-2239
nolav@cray.com

CUG Office

Bob and Karen Winget
Toolbox Conference Management
3224 Southwood Drive
Philomath, OR 97370 USA
(1-541) 929-3776 Fax: (1-509) 272-7239
bobwinget@toolboxcm.com kwinget@fpes.com

Meeting Room Map

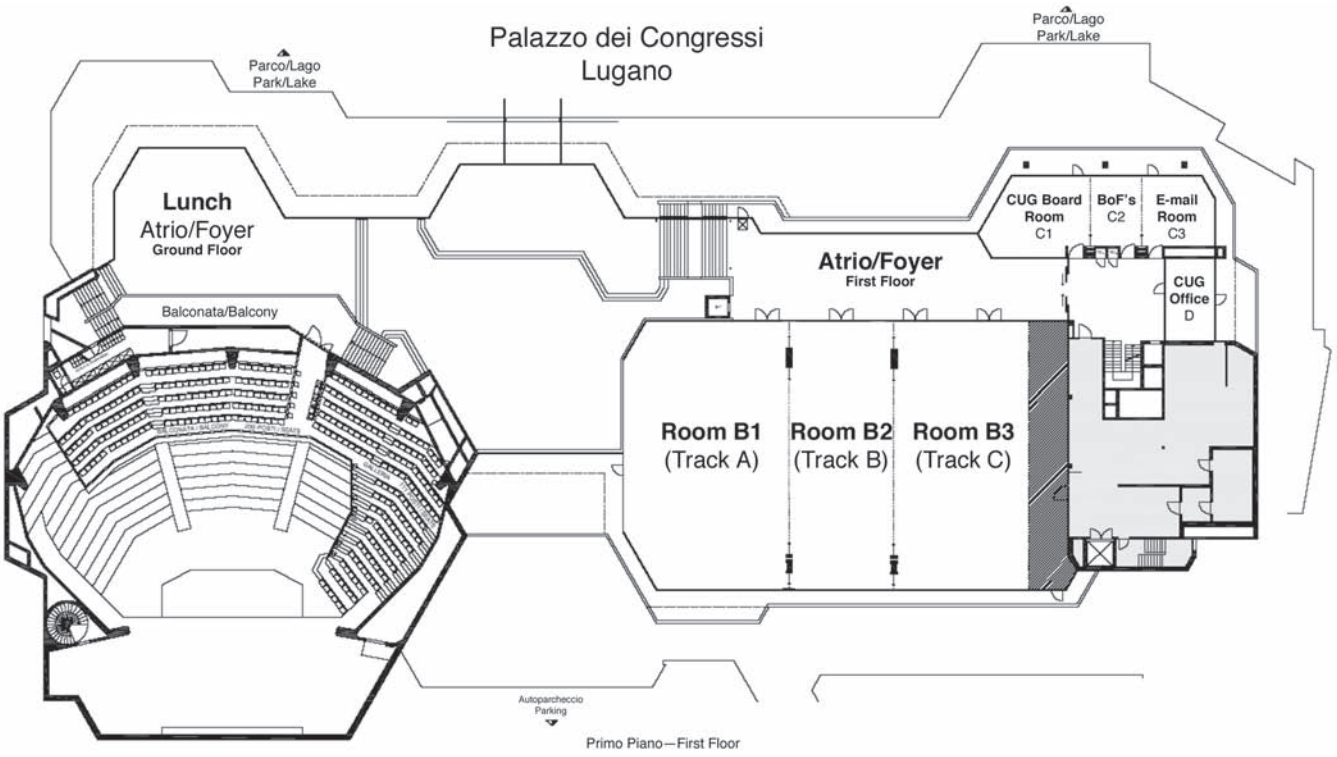

CUG 2007—New Frontiers

Dear Friends,

The Cray User Group invites you to its annual technical conference, CUG 2007, in Seattle, Washington, USA from April 30–May 4, 2007 at the Red Lion Hotel. The Boeing Company is pleased to host this event.

Our theme for this event is "New Frontiers" which reflects upon how the many improvements in High Performance Computing have significantly facilitated advances in Technology and Engineering. These advances have enabled the creation of better designs at lower cost. This is especially critical within manufacturing environments such as The Boeing Company that need to maintain product excellence while still staying competitive regarding cost. Faster processing and increased capabilities have been a catalyst in the development of our most advanced aircraft, the 787 Dreamliner, and have also vastly improved other product lines.

Here at Boeing we depend heavily on our Cray X1 to run critical engineering codes in support of The Boeing Commercial Airplane Company. We first installed our X1, more affectionately known as "Stonehenge," in fall 2003. At that time it was a single cabinet 16 node machine with 32Meg of memory per node. We upgraded this system to 32 nodes in late 2005.

Engineers using High Performance Computing have embarked upon a new frontier in aircraft development by pioneering designs based on computational fluid dynamics to obtain maximum speed and fuel economy, by creating innovations in aerodynamics, and in utilizing materials to be more environmentally friendly and efficient.

In order to take advantage of the lower fee we encourage you to register early. You will find all the necessary information about the venue, the hotel, and Seattle and its surroundings using the links from the cug.org web site. Registration will be open by December 1, 2006.

We look forward to welcoming you to Seattle, the home of The Boeing Commercial Airplane Company and Cray Inc.

Sincerely,

Henry Heeb
CUG 2007 Local Arrangements Chair
Manager BCA High Performance Computing
The Boeing compnay

Online and printed program design and production by Toolbox Conference Management (toolboxcm.com)