Dear Friends,

On behalf of the staff and faculty of the Arctic Region Supercomputing Center (ARSC) at the University of Alaska Fairbanks, we welcome you to CUG 2011 in Fairbanks, Alaska. This is the second time that ARSC has hosted CUG and we hope that this experience will be even better than it was sixteen years ago, in the Fall of 1995. Our theme, Golden Nuggets of Discovery, reflects the gold rush history of the far north as well as the many valuable discoveries made possible through High Performance Computing on Cray systems.

ARSC's rich history with Cray began with a CRAY Y-MP in 1993. Our continuous relationship has been sustained with leading installations of the T3D, T3E, SX6, X1 and XE6 architectures. Throughout, we have enjoyed the benefits of CUG membership that we have passed on to our users. Serving the CUG organization itself has brought its own rewards to many ARSC individuals.

We hope that this week will bring each of you new nuggets of discovery whether through new insights for supporting users of Cray systems or new ideas for performing your own research on those systems. This week is a great time to start dialogs and collaborations towards the goal of learning from each other.

We welcome you to Interior Alaska and hope that you will find time while you are here to enjoy some of its many charms. ARSC is very pleased you have come to the Golden Heart City.

Sincerely,

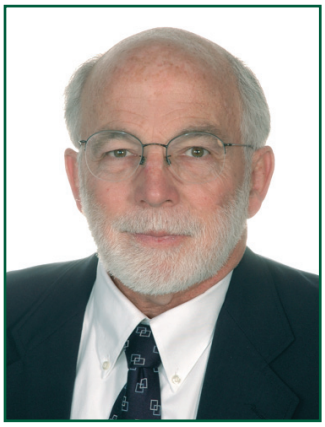

Image courtesy of ARSC

Dr. Frank Williams,
ARSC Director

Image courtesy of ARSC

Barbara Horner-Miller,
Local Arrangements Chair

Program Notes

Contents

Image courtesy of FCVB

The milepost at Golden Heart Plaza can help you figure out how far Fairbanks is from your home.

Special Interest Group (SIG) Meetings

The Special Interest Groups (SIGs) hold interactive sessions that are open to all interested CUG attendees. A separate interactive session will be held again this year for each of the three standing groups. Please check the schedule for day and time of the meetings. These meetings provide a forum to become acquainted with people from other CUG sites who share your interests, to discuss the future direction of the SIG and other important issues, and to hear talks of special interest. Liaisons from Cray Inc. will attend each SIG to help address questions and to foster communications between the SIG and Cray. You are encouraged to attend any of these open meetings.

Program Suggestions

The Program Committee needs your ideas! We encourage you to share your suggestions on how to improve the conference with the CUG Vice-President, 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 participants. You can make a unique contribution to the next CUG conference with your suggestions. Help establish the future direction of CUG and let us know how to help you make the most of your Cray system!

Conference Evaluation

The conference evaluation will be available online starting Thursday, May 26 at 8:00 AM until Friday, June 3 at 5:00 PM. Please go to www.cug.org/4-organization/eval_survey/ to complete the evaluation. If you complete your survey during the conference, please stop by the conference office and pick up a small thank you gift.

Cell Phone and Pager Policy

Please silence your cell phone and/or pager during conference sessions. There will be ample time to check your messages outside of the conference sessions. Please extend this courtesy to our presenters and your fellow attendees.

Changes to the Program

We always anticipate there will be some changes to the Program schedule. Check the Message Board outside the conference each day for schedule changes. Changes should be announced at the first morning session as well. Thank you for your patience and cooperation.

DAVID NEWMAN

David Newman is a professor of Physics at the University of Alaska Fairbanks as well as Director of the Center for Complex Systems Studies also at UAF. Before coming to UAF in 1998, he was selected as a Wigner Fellow at Oak Ridge National Laboratory in 1993 and worked there as a research scientist from 1993-1998. In 1997, he was awarded the Presidential Early Career Award for Scientists and Engineers. Before getting his PhD at the University of Wisconsin in 1993, Professor Newman served in the U.S. Peace Corps teaching in Kenya for two years.

His research work has covered plasma physics, fluid turbulence, and nonlinear dynamics as well as modeling the dynamics of complex systems such as the power transmission grid and communications networks.

Image courtesy of ARSC

Image courtesy of FCVB

The Trans-Alaska Oil Pipeline passes through Fairbanks on its route from Prudhoe Bay to Valdez. Visitors can view the pipeline on the outskirts of Fox, Alaska.

MONDAY

1A Tutorial

8:00 Introduction to Programming for GPU Accelerators, *Jeff Larkin, Cray Inc.*

At NVIDIA's GTC '10, Cray announced the future support for NVIDIA GPUs in XE6 systems. At SC10, GPUs were one of the hottest topics of the conference. This tutorial will teach the basics of GPU computing in preparation for gpu-accelerated Cray XE6 systems. I will establish a baseline knowledge in GPU architecture. This will be followed by a discussion of the currently available options for GPU programming, including CUDA C, CUDA Fortran, OpenCL, and compiler directives. The tutorial will also include a demonstration of GPU performance analysis and basic optimization techniques.

1B Tutorial

8:00 A Deep Dive on New Features of the Cray Programming Environment, *Luiz DeRose and Heidi Poxon, Cray Inc.*

This tutorial is intended for users that are interested in learning in more depth some of the more recent features in the Programming Environment for the Cray XT and Cray XE systems. The tutorial will cover, with examples, features in the Cray Compiling Environment (CCE), the Cray Performance Measurement and Analysis Tools (CPMAT), the Cray Debugging Supporting Tools (CDST), and the Cray Scientific and Math Libraries (CSML).

1C Tutorial

8:00 An Overview of the Chapel Programming Language, *Sung-Eun Choi and Bradford Chamberlain, Cray Inc.*

Chapel is a new parallel programming language being developed under the DARPA High Productivity Computing Systems program (HPCS). In this tutorial, we will present an introduction to Chapel, from context and motivation to a detailed description of Chapel via example computations.

2 Opening General Session

11:00 CUG Welcome, *Nicholas P. Cardo, CUG President, National Energy Research Scientific Computing Center*

11:10 Welcome, *Frank Williams, Director, Arctic Region Supercomputing Center, University of Alaska Fairbanks*

11:15 The Dynamics of Turbulent Transport, *David Newman, University of Alaska Fairbanks*

One of the great open problems in physics is a fundamental understanding of turbulence and turbulent transport. The issues involved in this have underpinned a number of "Grand Challenge" scale problems from climate to plasmas. One of the difficulties in studying turbulent transport is the extreme multi-scale nature of the problem to be solved. In the past few years a great deal of progress has been made on understanding turbulent transport in the presence of flows using computational methods. This new understanding includes classes of non-diffusive transport that have implications for all areas that model turbulent transport as a diffusive process. To continue to improve our ability to address the multi-scale nature of the problem, the Parareal technique (parallelization in time) has been applied to turbulence models. The implication of these advances on transport modeling and large-scale turbulence simulations will be discussed as will the application of the Parareal technique to turbulence. This work has been performed with the support of DOE-OFES grants and computational support from ARSC.

3A

1:00 The Cray Programming Environment: Current Status and Future Directions, *Luiz DeRose, Cray Inc.*

The Cray Programming Environment has been designed to address issues of scale and complexity of high-end HPC systems. Its main goal is to hide the complexity of the system, such that applications can achieve the highest possible performance from the hardware. In this talk I will present the recent activities and future directions of the Cray Programming Environment, including an overview of the Cray PE plans to support heterogeneous CPU/GPU systems.

1:45 Cray Scientific Libraries: Overview, Performance Evaluation and Advanced Usage, *Adrian Tate, Cray Inc.*

The Cray Scientific Libraries enable highly efficient usage of Cray systems with the minimum programmer effort. The standard and near-standard scientific libraries for dense linear algebra, sparse linear algebra and FFTs are all provided and tuned extensively for AMD processors, the Cray network, or both. A comprehensive set of custom tools are also provided that allow simpler usage, higher degrees of control or better performance than the standardized libraries. Some parts of LibSci use auto-tuning and adaptation, the tailoring of numerical kernels to the calling problem at run-time for increased performance. This talk will describe the technical innovations required to implement the optimizations and features in LibSci, provide a performance evaluation of the currently released libraries and detail upcoming feature and performance optimizations.

3B

1:00 Metrics and Best Practices for Host-based Access Control to Ensure System Integrity and Availability, *Urpo Kaila, Marco Passerini, and Joni Virtanen, CSC Scientific Computing Ltd.*

Open access in academic research computing exposes the servers to many kinds of brute-force attacks and vulnerability exploits. The systems administrator has a delicate task to similarly ensure system integrity by proper access controls and by applying security patches but also to enable service availability and ease of use. This paper will present an analysis of aggregated log metrics for access history and service up-time processed with tools such as Nagios and Splunk in conjunction with a set of cases of vulnerabilities, intrusions and faults. The paper will also compare and suggest improved best practices to be shared between sites.

1:45 High Performance Network Intrusion Detection in the HPC Environment, *Jim Mellander, National Energy Research Scientific Computing Center*

Today's high-performance, high-bandwidth systems require innovative approaches in order to provide effective Network Intrusion Detection. The National Energy Research Scientific Computing Center utilizes the advanced Intrusion Detection System, Bro, to detect and/or deflect network attacks that may threaten the open computing environment. We describe several novel approaches to security issues, such as intrusion detection on encrypted streams, and clustering of IDS nodes to accommodate high-bandwidth network requirements, and future challenges in this area.

3C

1:00 Overview of the Cray XMT2, *Andrew Kopser and Dennis Vollrath, Cray Inc.*

The Cray XMT2 system will be the newest offering in Cray's line of scalable multithreaded computers. The Cray XMT2 is based on the latest Cray Threadstorm processor. This paper will describe the Cray XMT2 with particular emphasis on the new architectural features provided by this processor.

1:45 Multi-TB file IO for Semantic Databases on the XMT, *Rolland Waters and Eric Lund, Cray Inc.*

Cray has developed an I/O library for the XMT that allows applications to store and restore multi-terabyte files rapidly and accurately. The newest version of this library features dynamically scaled transfers which maximize the speed of large transfers and minimize the setup time of small ones. A new daemon facilitates these features without any user intervention. Based on the elements of this design, Cray expects to be offering new approaches to hybrid computational environments using the XMT.

4A

3:00 Software's One-Sided Challenge, *Larry Kaplan, Edwin Froese, Forest Godfrey, Igor Gorodetsky, Chris Johns, Matt Kelly, and Brent Shields, Cray Inc.*

All Cray machines of recent times have scaled really well for MPI applications. The XE is no exception. However, introducing the PGAS programming model with direct user access to the network interface hardware required many changes and unprecedented cooperation between a lot of different functional development groups.

This paper describes how a team was able to deal with a number of challenging design and implementation issues, some over a very short period of time, in developing, testing, and releasing the software to support the rollout of the new system.

3:30 Application Performance Evaluation Studies of Multi-Core Nodes and the Gemini Network, *Mike Ashworth, Xiaohu Guo, Stephen Pickles, Martin Plummer, Andrew Porter, Andrew Sunderland, and Ilian Todorov, EPCC*

The UK National HPC Service (HECToR) has recently been upgraded from Cray XT4 with quad-core nodes to a Cray XT6 with 24-core nodes and then to a Cray XE6 with the Gemini network. We examine the performance implications for a range of applications including detailed performance profiling studies to examine the effect of the increase in the number of cores and the change in architecture from XT series Seastar to XE6 Gemini.

4:00 Application Characteristics and Performance on a Cray XE6, *Courtenay Vaughan, Sandia National Laboratories*

In this paper, we will explore the performance of two applications on a Cray XE6 and their performance improvement from previous machines, including the XT5 and the XT6. These two applications show different scaling effects as we go from machine to machine and we will explore the differences in the applications to explain these differences. We will use profiling and other tools to better understand resource contention within and between nodes and how that changes with the evolution of the machines with changes in processors and network.

4:30 A Performance Comparison Framework for Numerical Libraries on Cray XT5 Systems, *Bilel Hadri and Haihang You, National Institute for Computational Sciences*

Kraken, a Cray XT5 operated by the National Institute for Computational Sciences (NICS) enables scientific discoveries of researchers nationwide by providing leading-edge computational resources. Numerical libraries are frequently used by applications that are linked to different vendor libraries such as Libsci, ACML and MKL. Choosing the most efficient library for a given application is essential for achieving good performance. The performance comparison framework designed at NICS will help researchers determine the fastest library choices for their application.

4B

3:00 Cray's Security Update Process (or Why Can't We Have It Right Now?!), *Wendy Palm, Cray Inc.*

"Security vulnerability" is a frightening

and annoying fact of life in the computer world. There is a lot of confusion regarding Cray's process for monitoring and evaluating vulnerabilities as well as providing updates to our customers. This paper will describe the current process of periodic security updates and the fast-paced exciting "security scramble" we all know and don't love so much.

3:30 Overview of Node Health Checker, *Jason Sollom, Cray Inc.*

As the size of Cray systems increases, the meantime to failure of an individual compute node decreases. As a result, Cray's Node Health Checker (NHC) is playing an ever increasing role in ensuring job completion and system administrator sanity. NHC, marketed as NodeKARE, is a system management tool that runs a series of built-in and system administrator defined health tests on compute nodes. NHC automatically sequesters unhealthy nodes, dumps them for future debugging, and returns them to the pool of available nodes with a reboot. This increases the availability of nodes while simultaneously decreasing administrator intervention. Sequestering unhealthy nodes prevents them from causing job failure and saves tens of thousands of hours of compute time. This paper provides an overview of NHC.

4:00 Real-Time System Log Monitoring/ Analytics Framework, *Raghul Gunasekaran, Byung H. Park, David Dillow, Sarp Oral, Galen Shipman, and Al Geist, Oak Ridge National Laboratory*

Analyzing system logs provides useful insights for identifying system/application anomalies and helps in better usage of system resources. Nevertheless, it is simply not practical to scan through the raw log messages on a regular basis for large-scale systems. First, the sheer volume of unstructured log messages affects the readability, and secondly correlating the log messages to system events is a daunting task. These factors limit large-scale system logs primarily for generating alerts on known system events, and post-mortem diagnosis for identifying previously unknown system events that impacted the systems performance. In this paper, we describe a log monitoring framework that enables prompt analysis of system events in real-time. Our web-based framework provides a summarized view of console, netwatch, consumer, and apsched logs in real-time. The logs are parsed and processed to generate views of applications, message types, individual/group of compute nodes, and in sections of the compute platform. Also from past application runs we build a statistical profile of user/application characteristics with respect to known system events, recoverable/non-recoverable error messages and resources utilized. The web-based tool is developed for Jaguar XT5 at the Oak Ridge Leadership Computing facility.

4:30 User Application Monitoring through Assessment of Abnormal Behavior Recorded in RAS Logs, *Byung H. Park, Raghul Gunasekaran, Thomas Naughton, David Dillow, Al Geist, and Galen Shipman, Oak Ridge National Laboratory*

Abnormal status of an application is typically detected by "hard" evidence, e.g., out of memory, segmentation fault. However, such information only provides clues for the notification of abnormal termination of the application; lost are any implications as to the application's termination with respect to the particular context of the platform. Restated, the generic exception the application reports is devoid of the overall system context that is captured elsewhere in the system, e.g., RAS logs. In this paper we present an "activity entropy" based application monitoring framework that extracts both facts (events) and context with regard to applications from RAS logs, and maps them into entropy

scores that represent degrees of "unusualness" for applications. The paper describes our results from applying the framework to the Cray "Jaguar" system at Oak Ridge National Laboratory, and discusses how it identified applications running abnormally and implications based on the type of abnormality.

4C

3:00 Transitioning Applications from the Franklin XT4 System with 4 Cores Per Node to the Hopper XE6 System with 24 Cores Per Node, *Katie Antypas, Yun (Helen) He and Harvey Wasserman, National Energy Research Scientific Computing Center*

As NERSC users move from the Franklin XT4 system with 4 cores per node to the Hopper XE6 system with 24 cores per node, they have had to adapt to a lower amount of memory per core and on-node I/O performance which does not scale up linearly with the number of cores per node. This paper will discuss the practical implications of running on a system with 24 cores per node, exploring advanced aprun and memory affinity options for typical NERSC applications as well as strategies to improve I/O performance out of a node.

3:30 Building an Electronic Knowledge Base to Aide in Support of Jaguar, *Robert Whitten and Ashley Barker, Oak Ridge National Laboratory*

Supporting the world's largest XT5 can be a daunting task. By designing, developing and implementing an online knowledge base, the Oak Ridge Leadership Computing Facility (OLCF) has provided a valuable user support resources to its users. This paper will describe the process and results of this development effort.

4:00 Memphis on an XT5, *Collin McCurdy, Jeffrey Vetter, and Patrick Worley, Oak Ridge National Laboratory*

Memphis is a tool that makes use of Instruction Based Sampling (IBS) hardware counters, available in recent AMD processors, to help pinpoint the sources of memory system performance problems. This presentation will describe our experiences porting Memphis to a test XT5 system at ORNL, including modifications required by Compute Node Linux to the kernel module that interfaces with IBS, and low impact modifications to the batch queue that enable the module's use at runtime. The presentation will also include results from running production applications, such as CAM/HOMME, through Memphis on the XT5, highlighting performance problems pointed out by the tool.

4:30 Using Platform LSF Workload Scheduler with CLE, *Mehdi Bozzo-Rey, Platform Computing*

The Platform LSF 8.0 Workload Scheduler is now available on the Cray Linux environment. In addition to offering advanced scheduling capabilities, Platform LSF supports Cluster Compatibility Mode (CCM). In this session, we will present an overview of the Platform LSF and Cray Linux environment integration and give several examples to illustrate the scheduling capabilities of LSF on Cray XT/XE and XTm/XEm systems.

5A Interactive Session

5:15 Open Discussion with the CUG Board

5B Interactive Session

5:15 Programming Environments, Applications and Documentation SIG

TUESDAY

6 General Session

8:30 Introduction, *David Hancock, CUG Vice-President and Program Chair, Indiana University*

Cray Corporate Update, *Peter Ungaro, President and CEO, Cray Inc.*

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

Cray R&D Highlights and Challenges, *Peg Williams, Senior Vice President of R&D, Cray Inc.*

Cray Software Update, *Peter Young, Vice President of Software Engineering, Cray Inc.*

7 General Session

10:30 Titan: ORNL's New System for Scientific Computing, *Arthur (Buddy) Bland, Jim Rogers, Galen Shipman, Oak Ridge National Laboratory*

ORNL is planning to install a 10-20 petaflops computer system over the next 18 months that will be the next generation system for scientific computing for the U.S. Department of Energy. While there will be many similarities to the existing Jaguar system, there will also be architectural differences. In this paper, we discuss the accelerator-based architecture of Titan and the reasons for our decision to go in this direction. We also discuss our choice of the file systems to support the system.

11:00 1 on 100 (or More), *Peter Ungaro, President and CEO, Cray Inc.*

8A

1:00 Cray's Lustre Model and Road-Map, *Cory Spitz, Cray Inc.*

Since 2003, Cray, our customers and the wider HPC community have developed Lustre as a key technology component for our success. In order to ensure that Lustre will continue to grow and develop Cray has played a founding role, with other leaders in the HPC community, in launching OpenSFS and have joined the two other open Lustre consortia, HPCFS and EOFSCS. Cray plans to incorporate new Lustre features, produced through the efforts of these consortia and their member companies, into its products. This paper will lay out the support model and new software release details for Cray's use of Lustre in CLE and esFS in 2011, 2012, and beyond.

1:45 Update on Lustre, *Brent Gorda, Whamcloud*

Whamcloud is a startup focused on high-end filesystems: Lustre on Linux for HPC. The talk will introduce the company, its offerings and discuss the roadmap for continuing development in a vendor neutral manner for the open source technology.

8B

1:00 Authoring User-Defined Domain Maps in Chapel, *Bradford Chamberlain, Sung-Eun Choi, and David Iten, Cray Inc.*

One of the most promising features of Cray's parallel Chapel programming language is its support for 'user-defined domain maps' which permit advanced users to specify their own implementation for a parallel, distributed array that supports high-level global array operations. In choosing to write a domain map the user has control over high-level decisions like how data and iterations are divided among the target nodes of the machine as well as finer-grained decisions like the memory layout used to store the array's indices and values. In this paper, we give an overview of Chapel's user-defined domain map strategy and provide a summary of the developer's interface used to specify them.

1:45 Automation-Assisted Debugging on the Cray with TotalView, *Chris Gottbrath, Rogue Wave Software*

A little bit of automation can go a long way towards streamlining and simplifying the process of debugging scientific applications. This talk will demonstrate using a new TotalView feature, C++View, to transform complex data structures and automatically perform validity checks within them. C++View is an element of TotalView's extensive scripting framework, which also includes a type transformation facility, a fully programmable TCL-based CLI, a C and Fortran

expression evaluation system, and the scripting tools MemScript and TVScript.

8C

1:00 Optimizing Nuclear Physics Codes on the XT5, *Rebecca Hartman-Baker and Hai Ah Nam, Oak Ridge National Laboratory*

Scientists studying the structure and behavior of the atomic nucleus require immense high-performance computing resources to gain scientific insights. Several nuclear physics codes are capable of scaling to more than 100,000 cores on Oak Ridge National Laboratory's petaflop Cray XT5 system, Jaguar. In this paper, we present our work on optimizing codes in the nuclear physics domain.

1:45 Performance of Density Functional Theory Codes on Cray XE6, *Zhengji Zhao and Nicholas Wright, National Energy Research Scientific Computing Center*

Around 1/3 of the computing cycles are consumed by the materials science and chemistry users in each allocation year at NERSC, and ~75% of them run various Density Functional Theory (DFT) codes, among which the majority are pure MPI codes. In this paper, we select a few representative codes and discuss their performance on the Cray XE6, especially the performance impact from the multicore architecture in comparison with that on Cray XT4. We also explore how OpenMP and/or the multi-threaded BLAS library help to address the reduced per-core memory on Cray XE6 and to improve the parallel performance of the codes.

9A

3:00 Cray OS Road Map, *Charlie Carroll, Cray Inc.*

This paper will discuss Cray's operating system road map. This includes the compute node OS, the service node OS, the network stack, file systems, and administrative tools. Coming changes will be previewed, and themes of future releases will be discussed.

3:45 XE System Reliability and Resiliency: Observations and Impact to Operations, *Steve Johnson, Cray Inc.*

In 2010, Cray Inc. introduced the XE6 product with new technology in the software, interconnect, blades and cabinets. This paper will discuss the reliability trends that have been observed since the introduction of the product and our observations of the effectiveness of the network resiliency improvements incorporated into the system. We will also discuss how these improvements are impacting reliability and availability metrics and Cray support activities.

9B

3:00 Benchmark Performance of Different Compilers on a Cray XE6, *Michael Stewart and Yun (Helen) He, National Energy Research Scientific Computing Center*

There are four different supported compilers on NERSC's recently acquired XE6, and our users often request guidance from us in determining which compiler is best for a particular application. In this paper, we will describe the comparative performance of different compilers on several MPI and Hybrid MPI/OpenMP benchmarks with different characteristics. For each compiler and benchmark, we will establish the best set of optimization arguments to the compiler.

3:30 Targeting AVX-Enabled Processors Using PGI Compilers and Tools, *Brent Leback, The Portland Group*

AMD and Intel will release new microprocessors in 2011 based on the extended AVX architecture. In this paper we will show examples of compiler code generation and new library and tools capabilities that support these new processors. Performance data comparing the

new platform vs. previous generations will also be included.

4:00 The Chapel Tasking Layer Over Qthreads, *Kyle Wheeler, Sandia National Laboratories, Bradford Chamberlain, Cray Inc., and Richard Murphy, Sandia National Laboratories*

The Chapel compiler provides an abstraction of compute tasks that allows for the use of external libraries to provide the task management functionality. This paper describes the experiences and insights learned in porting Chapel to use the qthread lightweight threading library for task management.

9C

3:00 Grand-Scale WRF Testing on the Cray XT5 and XE6, *Don Morton, Oralee Nudson, and Don Bahls, Arctic Region Supercomputing Center, and Peter Johnsen, Cray Inc.*

The Arctic Region Supercomputing Center (ARSC) continues to push the Weather Research and Forecasting (WRF) model in ambitious directions. With the help of Cray, Inc. and WRF developers, a model size of more than one billion grid points was tested on a real-world weather scenario over the North Pacific and Arctic, providing 1-kilometer horizontal resolution over the entire region. With research and operations groups increasingly interested in horizontal resolutions of 100 meters or less, and fine-scale vertical resolutions near the surface, it becomes imperative to begin testing the full WRF environment (pre-processing, model execution and post-processing) on domains consisting of billions of grid points, executed on tens to hundreds of thousands of cores. In this year's paper and presentation, we extend our work of previous years by attempting to run real-world weather simulations domains of billions of grid points, employing various optimization schemes suggested by Cray and WRF developers. The platforms used for this work include the Cray XT5 (kraken) and the Cray XE6 (chugach).

3:30 Acceleration of Porous Media Simulations on the Cray XE6 Platform, *Nicholas Wright, National Energy Research Scientific Computing Center, George Pau and Michael Lijewski, Lawrence Berkeley National Laboratory, and Kirsten Fagnan, National Energy Research Scientific Computing Center*

Simulating carbon sequestration and reacting groundwater flow is important because each of these problems involve processes that can not be sufficiently simulated in the laboratory. Investigation of these issues requires resolution of spatial scales on the order of meters within domains that are on the order of tens of kilometers, which necessitates the use of adaptive mesh refinement. Also, stiff chemical reactions and large acoustic wave speeds limit the size of the time step, but simulations must be able to predict results for 10-15 years in the future, making this a challenging multi-scale physics problem. This paper examines the memory requirements of the porous media code used for these simulations and discusses improving performance through the use of a hybrid (OpenMP+MPI) programming on the Cray XE6 platform.

4:00 A Study of Scalability Performance for Hybrid Mode Computation and Asynchronous MPI Transpose Operation in DSTAR3D, *Lucian Anton and Ning Li, NAG Ltd., and Kai Luo, Southampton University*

A necessary condition for good scalability of parallel computation at large core counts is to minimize the data communication and overlap it with computation whenever this is possible. Along these ideas we study the parallel performance of the code DSTAR3D that simulates reactive turbulent flows using direct numerical simulation. The studied algorithm

uses a two-dimensional domain decomposition with OpenMP threads inside each local domain for numerical intensive kernels and asynchronous MPI for the needed transpose operation of subdomains. This new algorithm allows DSTAR3D to use around 10,000 cores with excellent scalability, a significant improvement from hundreds of cores used by the initial algorithm based on one-dimensional decomposition.

10A Birds of a Feather

4:45 The Future of Dynamic Linking for the Cray XT/XE Programming Environment

10B Birds of a Feather

4:45 Node Health Checker Discussion

6:00 *Cray Reception*

Buses will depart from the front door of the Westmark at 5:45, 6:00, and 6:15 pm to take attendees and their guests to Pioneer Park for the Cray Reception. See Social Events, p. 27 for details.

WEDNESDAY

11 General Session

8:15 CUG Business: AC Introductions, Election Presentations, Voting, *Chair: Nicholas P. Cardo, CUG President, National Energy Research Scientific Computing Center*

9:00 Greenland Ice Sheet Flow Computations: Scaling-Up to High Spatial Resolution and Fast Time-Scale Boundary Processes, *Ed Bueler, University of Alaska Fairbanks*

Scientists need physics-based models which connect warming in the polar regions to the behavior of ice sheets—especially their sea level contribution—but understanding of ice flow dynamics remains limited. Models that connect hard-to-observe local processes to global flow consequences are not as mature as for other climate components. Fortunately, ice sheet modeling is growing up. Modelers are converging on effective tools to make high-resolution multi-physics simulations on supercomputers actually useful. I'll address these challenges for a Greenland ice sheet model: Does 1 km ice-sheet-wide resolution resolve fast flow in 5 km wide fjords? Are long-distance stress transmissions in floating or well-lubricated ice, and the fast time-scale processes at its boundary, modeled well-enough to capture observed ice sheet changes? (Can we solve such huge linear systems at every timestep?) What is the physical limit on the speed of flowing ice?

9:45 CUG Business: Election Results

12A

10:30 I/O Congestion Avoidance via Routing and Object Placement, *David Dillow, Galen Shipman, and Sarp Oral, Oak Ridge National Laboratory, and Zhe Zhang, IBM T. J. Watson Research Center*

As storage systems get larger to meet the demands of petascale systems, careful planning must be applied to avoid congestion points and extract the maximum performance. In addition, the large size of the data sets generated by such systems make it desirable for all compute resources in a center to have common access to this data without needing to copy it to each machine. This paper describes a method of placing I/O close to the storage nodes to minimize contention on Cray's SeaStar2+ network, and extends it to a routed Lustre configuration to gain the same benefits when running against a center-wide file system. Our experiments show performance improvements for both direct attached and routed file systems.

11:15 Tips and Tricks for Diagnosing Lustre Problems on Cray Systems, *Cory Spitz and Ann Koehler, Cray Inc.*

As a distributed parallel file system, Lustre is prone to many failure modes. The manner

in which it breaks can make diagnosis and serviceability difficult. Cray deploys Lustre file systems at extreme scales, which compounds the difficulties. This paper discusses tips and tricks for diagnosing and correcting Lustre problems for both CLE and esFS installations. It will cover common failure scenarios including node crashes, deadlocks, hardware faults, communication failures, scaling problems, performance issues, and routing problems. Lustre issues specific to Cray Gemini networks are addressed as well.

12B

10:30 The NCRC Grid Scheduling Environment, *Frank Indiviglio, National Climate-Computing Research Center, and Don Maxwell, Oak Ridge National Laboratory*

In support of the NCRC, a joint computing center between NOAA and ORNL, a grid-based scheduling infrastructure was designed to allow geographically separate computing resources to be used as production resources in climate and weather research workflows. These workflows require job coordination between the two centers in order to provide a complete workflow of data staging, computation, post-analysis and archival. This paper details the design, implementation and initial production phase of the infrastructure and lessons learned from the process.

11:15 Case Studies from the OLCF Center for Application Acceleration Readiness: The Importance of Realizing Hierarchical Parallelism in the Hybrid Multicore Era, *Bronson Messer, Ricky Kendall, Richard Graham, and Oscar Hernandez, Oak Ridge National Laboratory, and John Levesque, Cray Inc.*

We will present several case studies - including kernel identification and work performed to date - for our suite of Early Science applications targeted for Titan. This work has, we believe, made clear a productive path forward for application developers to maximize effective use of hybrid architectures, where the realization of hierarchical parallelism - from distributed-memory to SMP-like to vector-like - is the essential ingredient.

12C

10:30 The NERSC- Cray Center of Excellence: Performance Optimization for the Multicore Era, *Nicholas Wright, Hongzhang Shan, Filip Blagoievic, Harvey Wasserman, Tony Drummond, and John Shalf, National Energy Research Scientific Computing Center, Karl Fuerlinger, UC Berkeley, Katherine Yelick, National Energy Research Scientific Computing Center, Stephane Ethier, Princeton Plasma Physics Laboratory, Marcus Wagner, Nathan Wichmann, Sarah Anderson, and Mike Aamodt, Cray Inc.*

We compare performance of several NERSC benchmarks on three Cray platforms, Franklin (XT4), Jaguar (XT5) and Hopper (XE6). We also report our work on evaluating the hybrid MPI-OpenMP programming model for several of these benchmarks. By using detailed timing breakdowns, we measure the contributions to the total runtime of the applications from computation, communication, and from runtime overhead, such as that due to OpenMP regions; and we discuss their effect on the performance differences observed. Finally, we report preliminary results of our PGAS (UPC and CAF) performance measurements.

11:15 Providing Runtime Clock Synchronization with Minimal Node-to-Node Time Deviation on XT4s and XT5s, *Terry Jones and Gregory Koenig, Oak Ridge National Laboratory*

We present a new high precision clock synchronization algorithm designed for large XT4 and XT5 leadership-class machines. The algorithm, which is designed to support OS noise

reduction through co-scheduling, is suitable for usage cases requiring low overhead and minimal time deviation between nodes. Unlike most high-precision algorithms that reach their precision in a post-mortem analysis after the application has completed, the new ORNL-developed algorithm rapidly provides precise results during runtime. Previous to our work, the leading high-precision clock synchronization algorithms that made results available during runtime relied on probabilistic schemes that are not guaranteed to result in an answer.

13A

1:00 Determining the Health of Lustre Filesystems at Scale, *David Dillow, Jason Hill, Dustin Leverman, and Scott Koch, Oak Ridge National Laboratory*

Monitoring the components of a Lustre file system is crucial to meeting mission requirements as the scale and complexity of the installation grows. Determining the health and performance of the file system becomes non-trivial, and the complexity increases faster than the size of the installation. This paper discusses the ongoing work at the Oak Ridge Leadership Computing Facility to monitor the health of its center-wide Lustre file systems.

1:30 DVS, GPFS and External Lustre at NERSC - How It's Working on Hopper, *Tina Butler, Rei Lee, and Gregory Butler, National Energy Research Scientific Computing Center*

Providing flexible, reliable and high performance access to user data is an ongoing problem for HPC centers. This paper will discuss how NERSC has partitioned its storage resources between local and global filesystems, and the configuration, functionality, operation and performance of these several different parallel filesystems in use on NERSC's Cray XE6, Hopper.

2:00 DVS and DSL Implementation at NCAR, *Irfan Elahi and Junseong Heo, National Center for Atmospheric Research*

The presentation will detail NCAR's experience with setting up DVS (Data Virtualization Service) and DSL (Dynamic Shared Libraries) on Lynx, the NCAR's Cray XT5m System.

Driven by local users' need at NCAR, we implemented GPFS shared files systems from external servers to the XT5m system using DVS. Setting up four service nodes as DVS servers for stripe-parallel mode, we were able to mount several GPFS file systems and an external Lustre, making all of them available on compute nodes as well as on service nodes.

We also implemented DSL support on our Cray XT5m. Lynx allows dynamically linked applications with shared libraries to run on the compute nodes. This also allows us to define and use our own DSOs (Dynamic Shared Objects).

13B

1:00 The Programmability and Performance of Cray's Programming Environment for Accelerators, *Jeffrey Poznanovic, Sadaf Alam and Gilles Fourestey, CSCS Swiss National Supercomputing Center, and Adrian Tate, Cray Inc.*

Although GPU accelerators are rapidly being adopted within the HPC community due to their high floating point performance, the related programming languages, such as CUDA and OpenCL, require significant application code modifications to achieve desirable performance. Cray has addressed this issue with the prototype Programming Environment for Accelerators that seeks to allow efficient usage of hybrid GPU/CPU systems using a directive-based programming model and accelerated scientific libraries. Directive-based programming has the potential to drastically improve programmability and portability by abstracting away the complex

underlying hardware characteristics and providing an interface to specify parallelism to the compiler. Scientific libraries that are designed to make highly efficient use of the accelerators through standardized interfaces remove the burden of memory management from the programmer and hence increase user productivity. Our paper evaluates the performance and programmability of Cray's Programming Environment for Accelerators. Initially, we ported multiple PGI Accelerator applications into the accelerator version of OpenMP. Then, we modified existing multi-core OpenMP codes to use the accelerator-based OpenMP directives. We also subjected the Cray Scientific Library for accelerators to a series of benchmark tests, and attempted to show how the Cray PE as a whole can be used to accelerate a real scientific application. Based on these experiences, this paper reports our performance results, evaluates the effort required to obtain these results and presents the insights gained from the porting process.

1:30 Porting a Particle Transport Code to GPGPU Using Hybrid MPI / OpenMP / Cuda Programming Models, *Paul Graham, EPCC, Dave Barrett, AWE PLC, and Ron Bell, EPCC*

This paper will present the work undertaken through collaboration with EPCC, required to port AWE's benchmark code, Chimaera, to NVIDIA Fermi GPU using Double Precision arithmetic and hybrid MPI / OpenMP / Cuda programming. This paper will describe these algorithmic changes and performance results obtained from both the original evaluation project and from the final ported application code using Cuda and OpenMP.

2:00 Porting Solvers to a Multi-Processor, Multi-Core, Multi-gpu Platform Using the PGI Accelerator Model and PGI CUDA Fortran, *Brent Leback, The Portland Group*

In this paper we present work and results from a port of a standard solver library to a heterogeneous computing platform, using a minimally-intrusive set of accelerator directives and language extensions. Results comparing alternative coding solutions are given.

13C

1:00 PBS Plug-Ins: a Run-time Environment for Agility and Innovation, *Bill Nitzberg, Altair*

Nitzberg's definition of HPC is "computing that requires pushing the limits of today's technology just beyond where it works well". HPC combines leading-edge hardware, operating systems, networks, enterprise utilities, applications, and more. The key to making HPC work is agility. PBS plug-ins allow you to quickly and easily integrate, extend, and customize PBS Professional to meet the unique requirements and constantly-evolving demands in your enterprise. Plus, by providing a well-defined, modular platform, standardizing on Python "hooks", and embedding it everywhere, our run-time environment enables sharing and reusing innovative plug-ins throughout the PBS community. This presentation provides a technical look at PBS plug-ins, including real examples (yes, code) to get you started.

1:30 Prospects for Truly Asynchronous Communication with Pure MPI and Hybrid MPI/OpenMP on Current Supercomputing Platforms, *Georg Hager, Erlangen Regional Computing Center, Rainer Keller, High Performance Computing Center Stuttgart, Thomas Zeiser and Johannes Habich, Erlangen Regional Computing Center, Thomas Schoenemeyer, CSCS Swiss National Supercomputing Centre, and Gerhard Wellein, Erlangen Regional Computing Center*

We investigate the ability of MPI implementations to perform truly asynchronous communication with nonblocking point-to-point calls on current highly parallel systems,

including the Cray XT and XE series. For cases where no automatic overlap of communication with computation is available, we demonstrate several different ways of establishing explicitly asynchronous communication by variants of functional decomposition using OpenMP threads or tasks, implement these methods in application codes, and show the resulting performance benefits. The impact of node topology and the possible use of simultaneous multithreading (SMT) is studied in detail.

2:00 A uGNI-Based MPICH2 Nemesis Network Module for Cray XE Computer Systems, *Howard Pritchard and Igor Gorodetsky, Cray Inc.*

Recent versions of MPICH2 have featured Nemesis - a scalable, high-performance, multi-network communication subsystem. Nemesis provides a framework for developing Network Modules (Netmods) for interfacing the Nemesis subsystem to various high-speed network protocols. Cray has developed a User-Level Generic Network Interface (uGNI) for interfacing MPI implementations to the internal high-speed network of Cray XE and follow-on compute systems. This paper describes the design of a uGNI Netmod for the MPICH2 nemesis subsystem. Performance data on the Cray XE will be presented. Planned future enhancements to the uGNI MPICH2 Netmod will also be discussed.

14A

3:00 Petascale Capability Computing, *Doug Doerfler and Sudip Dosanjh, Sandia National Laboratories, John Morrison, and Manuel Vigil, Los Alamos National Laboratory*

The LANL/Sandia Alliance for Computing at the Extreme Scale (ACES) is partnering with Cray to deploy a Petascale capability system, Cielo, for the Department of Energy's Advanced Simulation and Computing (ASC) program. Many targeted national security applications are extremely large and will require a significant fraction of the cores on Cielo to execute. Cielo is one of three Cray Petascale systems, has 6,704 nodes and uses a 3-D Torus topology and Cray's Gemini interconnect. Each node is composed of two 8 core AMD Magny Cours processors with 16 GB of memory each, resulting in a total of 107,264 cores and 160 TB of memory. It is the first large Cray system that employs a Panasas filesystem. An upgrade in early 2011 will add 0.34 Petaflops to Cielo. This presentation describes Cielo's architecture, the user environment and preliminary performance results. We discuss some of our system integration efforts and challenges, as well as our experience with initial use of the system for simulations. An ASC R&D partnership with Cray to design an interconnect for the 2014 timeframe is also described.

3:45 The Hopper System: How the Largest XE6 in the World Went from Requirements to Reality, *Jonathan Carter and Tina Butler, National Energy Research Scientific Computing Center*

This paper will discuss the entire process of acquiring and deploying Hopper from the first vendor market surveys to providing 3.8 million hours of production cycles per day for NERSC users. Installing the latest system at NERSC has been both a logistical and technical adventure. Balancing compute requirements with power, cooling, and space limitations drove the initial choice and configuration of the XE6, and a number of first-of-a-kind features implemented in collaboration with Cray have resulted in a high performance, usable, and reliable system.

14B

3:00 Storage Adventures at Petascale and Beyond, *Keith Miller, DataDirect Networks*

The first Petaflop system was put into

production in 2008. By 2011, several systems in the 10's of Petaflops will be deployed. The race to Exascale has begun and the face of storage will change dramatically as we push new boundaries of I/O and archive scalability. Very large compute systems are creating unprecedented challenges for the storage systems that support them. This talk will explore experiences and concepts which DDN has recently gained through deploying many of the world's fastest HPC file systems, including:

- Maximizing cluster performance over the aggregate of the production lifespan, not just during benchmarking
- Exploiting NAND-based storage devices for intelligent storage tiering, buffering and automatic I/O path acceleration
- Data protection technologies which eliminate silent data corruption across scalable storage pools

3:45 Deployment and Implementation of a Workflow-Oriented Data Storage and Transfer Framework, *Douglas Fuller, National Climate-Computing Research Center*

Many emerging compute environments increasingly handle complex task workflows utilizing geographically separate resources. These resources may include data sources, compute platforms, pre- or post-processing platforms, and archive locations. The National Climate-Computing Research Center's Gaea deployment includes multiple Lustre filesystems and workflow-integrated data management subsystems. This paper details Gaea's storage system architecture, deployment, implementation and early experiences in support of its workflow software.

14C

3:00 Performance of the Time-Dependent Close-Coupling Approach to Electron-Impact Ionization on the Cray XE6, *James Colgan, Los Alamos National Laboratory, Michael Pindzola, Auburn University, Katie Antypas, Woo-Sun Yang, and Yun (Helen) He, National Energy Research Scientific Computing Center*

We report on time-dependent close-coupling calculations of the electron-impact ionization of small atoms and molecules. Such calculations are required to accurately predict the angular distributions and energy sharings of the two outgoing electrons after ionization by an incident electron. Our calculations treat the long-range electron-electron interaction without approximation, resulting in computationally intensive problems. We have performed calculations for electron-impact ionization of helium and molecular hydrogen on the NERSC facilities Franklin (Cray XT4) and Hopper II (Cray XE6). We report on the scaling properties of our codes on these platforms, and discuss some performance issues we have encountered. This work is supported in part by grants from the U.S. Department of Energy and the U.S. National Science Foundation. The Los Alamos National Laboratory is operated by Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under Contract No. DE-AC5206NA25396.

3:45 Cielo Full-System Simulations of Multi-Beam Laser-Plasma Interaction in NIF Experiments, *Steven Langer, Bert Still, Denise Hinkel, Ed Williams, and Todd Gamblin, Lawrence Livermore National Laboratory*

pF3D simulates laser-plasma interactions in experiments at the National Ignition Facility (NIF), the home of the world's most powerful laser. Simulations of recent NIF experiments require 100 billion zones and have been run on Cielo, a Cray XE6 with a Panasas parallel file system located at Los Alamos National Laboratory. This paper compares the techniques used to obtain good I/O and message passing scaling in ~100,000 processor runs on Cielo and

on BlueGene/P systems with Lustre and GPFS parallel file systems.

15A Interactive Session

4:15 Systems Support SIG

15B Birds of a Feather

4:15 Cray GPU Programming Tools

15C Birds of a Feather

4:15 XTreme BoF – Better Coordination of Our Security Efforts to More Effectively Support Our Sites

6:00 *CUG Night Out at the Antique Auto Museum*

Buses will depart from the front door of the Westmark at 6:00 and 6:30 for the Antique Auto Museum for the CUG Night Out. See Social Events, p. 27 for details.

THURSDAY

16A

8:30 Application-Driven Acceptance of Cielo, an XE6 Petascale Capability Platform, *Douglas Doerfler and Mahesh Rajan, Sandia National Laboratories, Cindy Nuss, Cray Inc., and Cornell Wright and Thomas Spelce, Los Alamos National Laboratory*

Cielo is one of the first instantiations of Cray's new XE6 architecture and will provide capability computing for the NNSA's Advanced Simulation and Computing (ASC) Campaign. A primary acceptance criteria for the initial phase of Cielo was to demonstrate a six times (6x) performance improvement for a suite of ASC codes relative to its predecessor, the ASC Purple platform. This paper describes the 6x performance acceptance criteria and discusses the applications and the results. Performance up to tens of thousands of cores are presented with analysis to relate the architectural characteristics of the XE6 that enabled the platform to exceed the acceptance criteria.

9:00 Discovering the Petascale User Experience in Scheduling Diverse Scientific Applications, *Troy Baer, R. Glenn Brook, Lonnie D. Crosby, and Matt Ezell, National Institute for Computational Sciences*

Newly emerging petascale computational resources are popular for both capability and capacity computing. However, these varied job classes have widely different resource requirements that make scheduling challenging. Beyond machine utilization, the scheduling of computational resources should provide reasonable throughput for all classes of jobs. This work will examine the user impact of scheduling various job classes on a petascale, shared-computing resource with a diverse workload, including scheduling policies and user behavior.

9:30 Cosmic Microwave Background Data Analysis at the Petascale and Beyond, *Julian Borrill, Christopher Cantalupo, Theodore Kisner and Radek Stompor, National Energy Research Scientific Computing Center*

The analysis of Cosmic Microwave Background (CMB) data is an ongoing high performance computing challenge. For more than a decade now the size of CMB data sets has tracked Moore's Law, and we expect this to continue for at least the next 15 years. In this talk we will review the work done to date to follow this scaling, and discuss the steps we are taking to continue to do so to the petascale and beyond.

16B

8:30 Large-Scale Performance Analysis of PFLOTRAN with Scalasca, *Brian Wylie, Juelich Supercomputing Centre*

The PFLOTRAN code for reactive multiphase flow and transport has featured prominently in US Department of Energy SciDAC and INCITE programs, where its execution performance with up to 128k processor cores on Cray XT and

IBM BG/P systems has been analyzed. Although the complexities of PFLOTRAN executions employing PETSc, LAPACK, BLAS, HDF5 and MPI libraries at large scale were challenging, the open-source Scalasca [www.scalasca.org] toolset was able to provide valuable insight into a variety of performance-limiting aspects.

9:00 Evolution of the Cray Performance Measurement and Analysis Tools, *Heidi Poxon, Cray Inc.*

The goal of the Cray Performance Measurement and Analysis Tools is to help the user identify important and meaningful information from potentially massive data sets by providing hints around problem areas instead of just reporting raw data. Analysis of data that addresses multiple dimensions of scalability including millions of lines of code, lots of processes or threads and long running applications is needed. The Cray toolset supports these dimensions by collecting information at process and thread levels, and providing features such as load imbalance analysis, derived metrics based on hardware events, and optimal MPI rank placement strategies. This paper focuses on recent additions to the performance tools to enhance the analysis experience and support new architectures such as hybrid X86 and GPU systems. Work presented includes support for applications using PGAS programming models, loop work estimates that help identify parallel or accelerator loop candidates, and statistics around accelerated loops.

9:30 Debugging at Petascale and Beyond, *Chris January, David Lecomber, and Mark O'Connor, Allinea Software*

Debugging at scale is now a reality - with Allinea DDT 3.0 achieving whole-machine interactive debugging of Petaflop Cray systems at high speed in production usage. This paper explores the need and opportunities for debugging at the target application scale, and shows how Allinea's fast and scalable debugging architecture enables new possibilities that are simplifying the task of debugging at a time when system architectures are becoming more complex.

16C

8:30 A Pragmatic Approach to Improving the Large-Scale Parallel I/O Performance of Scientific Applications, *Lonnie D. Crosby, Glenn Brook, Mikhail Sekachev, and Kwai Wong, National Institute for Computational Sciences*

I/O performance in scientific applications is an often neglected area of concern during performance optimizations. However, various scientific applications have been identified which benefit from I/O improvements due to the volume of data or number of compute processes utilized. This work will detail the I/O patterns and data layouts of real scientific applications, discuss their impacts, and demonstrate pragmatic approaches to improve I/O performance.

9:00 The Design of an Auto-Tuning I/O Framework on Cray XT5 Systems, *Haihang You and Qing Liu, National Institute for Computational Sciences, and Zhiqiang Li, University of Tennessee*

The Cray XT5 is equipped with Lustre, a parallel file system. To utilize I/O effectively is essential for an application to scale up. We have developed a mathematical model based on queuing theory and built an experimental I/O auto-tuning infrastructure for XT5 system.

9:30 Cheetah: A Scalable Hierarchical Collective Operation Framework, *Richard Graham, Pavel Shamis, Joshua Ladd, and Manjunath Gorentla Venkata, Oak Ridge National Laboratory*

The performance and scalability of collective operations play a key role in the performance and

scalability of many scientific applications. Within the Open MPI code base we have developed a general purpose hierarchical collective operations framework called Cheetah, and applied it at large scale on the Oak Ridge Leadership Computing Facility's Jaguar platform, obtaining better performance and scalability than the native MPI implementation, in measurement taken up to order 49K process count. This talk discusses Cheetah's design and implementation, as well as the results of large-scale benchmark data.

17A

10:30 Increasing Petascale Resources and Utilization through Bi-Modal Scheduling Policies, *Patricia Kovatch, Phil Andrews, Victor Hazelwood, Troy Baer, Matt Ezell, Ryan Braby, Glenn Brook, Justin Whitt, Tabitha Samuel, and Lonnie D. Crosby, National Institute for Computational Sciences*

In late 2009, the National Institute for Computational Sciences placed in production the world's fastest academic supercomputer (third overall), a Cray XT5 named Kraken, with almost 100,000 compute cores and a peak speed in excess of one Petaflop. Delivering over 50% of the total cycles available to the National Science Foundation users via the TeraGrid, Kraken has two missions that have historically proven difficult to simultaneously reconcile: providing the maximum number of total cycles to the community, while enabling full machine runs for "hero" users. Historically, this has been attempted by allowing schedulers to choose the correct time for the beginning of large jobs, with a concomitant reduction in utilization. At NICS, we used the results of a previous theoretical investigation to adopt a different approach, where the "clearing out" of the system is forced on a weekly basis, followed by consecutive full machine runs. As our previous simulation results suggested, this led to a significant improvement in utilization, to over 90%. The difference in utilization between the traditional and adopted scheduling policies was the equivalent of a 300+ Teraflop supercomputer, or several million dollars of compute time per year.

11:00 Producing Weather Forecasts on Time in Denmark Using PBS Professional, *Thomas Lorenzen and Thor Olason, Danish Meteorological Institute, and Frithjov Iversen and Paola Palazzi, Cray Inc.*

Running a mix of research jobs together with jobs that shall run in a timely manner to a predefined schedule adds an extra layer of complexity to job scheduling. More traditional approaches for ensuring resource availability include suspend/resume or checkpoint/restart, which are either not available on XT systems or have practical limitations. This presentation outlines the use of advance reservation scheduling on an XT system doing Numerical Weather Prediction (NWP) at the Danish Meteorological Institute (DMI). A basic and simple strategy for exploiting the advance reservation scheduling feature of PBS Professional to ensure resource availability at predefined timeslots is presented along with experiences made over the past years.

11:30 Deploying SLURM on XT, XE, and Future Cray Systems, *Gerrit Renker, Neil Stringfellow, Morris Jette, Danny Auble, and Sadaf Alam, CSCS Swiss National Supercomputing Centre*

We describe porting the open-source SLURM resource manager to the Cray BASIL/ALPS interface; and report on experiences of using it on our main 20-cabinet Cray XT5 production platform, as well as several development systems of a heterogeneous multi-cluster environment. Since some of these systems are GPU-based, we also discuss issues in extending the existing interface to future systems (Cray XE6 with GPU accelerators), in order to take advantage of

SLURMs cutting-edge GPU support.

17B

10:30 Shared Libraries on a Capability Class Computer, *Suzanne Kelly, Ruth Klundt, and James Laros, Sandia National Laboratories*

Popularity of dynamically linked executables continues to grow within the scientific computing community. The system software implementation of shared libraries is non-trivial and has significant implications on application scalability. This presentation will first provide some background on the Linux implementation of shared libraries, which was not designed for distributed HPC platforms. This introductory information will be used to identify the scalability issues for massively parallel systems such as the Cray XT/XE product lines. Lastly, the presentation will describe the considerations and lesson learned in file system placement of the shared libraries on Cielo, a Cray XE6 system with over 100,000 cores. Scaling results and comparisons will be included.

11:00 Scalability of Paraview's Coprocessing Capability, *Nathan Fabian, Sandia National Laboratories*

For exceedingly large high performance computing runs, writing all data to disk is unmanageably slow. It becomes necessary to have analysis and visualization communicate with the simulation in memory instead of through disk, retaining access to all available data for analysis. The open source visualization and analysis tool, Paraview, has recently added a coprocessing API allowing it to be linked into simulation codes. We will demonstrate scalability of Paraview coprocessing on up to 64,000 cores on the new NNSA platform, Cielo.

11:30 A Programming Environment for Heterogeneous Multi-Core Computer Systems, *Richard Graham, Pavel Shamis, Oscar Hernandez, Christos Kartsaklis, Tiffany Mintz, and Chung Hsing Hsu, Oak Ridge National Laboratory*

As part of the OLCF-3 project, the Oak Ridge National Laboratory is working with several vendors and engaged in research to develop a Programming Environment for mixed CPU/Accelerator based ultra-scale computer systems. The environment provides a toolset to port or develop CPU/Accelerator systems while reducing development time to improve the performance and portability of the codes while minimizing sources of errors. Our toolset consists of compilers for high-level Accelerator directives, libraries, performance tools and a debugger with synergistic interfaces among them. In this paper we show how these tools work together and how they support the different stages of the program development/porting cycle. Our paper will describe how we successfully used the tools to port DOE codes to a CPU/Accelerator system.

17C

10:30 Topology, Bandwidth and Performance: A New Approach in Linear Orderings for Application Placement in a 3D Torus, *Carl Albing, Norm Troullier, Stephen Whalen, and Ryan Olson, Cray Inc.*

Application performance was improved in Cray XT supercomputers using node ordering with a simple, one-dimensional allocation strategy but improvements were dependent on sizes and shapes of systems and applications so a new "thicker" ordering was developed to improve bisection bandwidth of placed applications, both large and small. The Cray XE line provides yet another challenge – differing speeds depending on the axis travelled which result in different bi-section bandwidth characteristics. To meet this challenge an enhanced ordering was developed which varies with system size, benefiting a wide range of

applications, all without user input. This paper summarizes the approach to placement that the Cray Application Level Placement Scheduler (ALPS) now offers based on the underlying node topology, the reasons for this approach, and the variations that sites can choose to optimize for their specific machines.

11:00 Data Systems Modernization (DSM) Project: Development, Deployment, and Direction, *Robert Whitten, Oak Ridge National Laboratory*

The Data Systems Modernization (DSM) project was undertaken to consolidate and update the current information systems of the Oak Ridge Leadership Computing Facility (OLCF). The project combined the Resource Allocation and Tracking System (RATS), the New Account Creation System (NACS) and open-source process management and business intelligence software to streamline the data processing systems of the OLCF. This paper will discuss the development, deployment and future directions of this ongoing project.

18A

1:00 Information Environment in JAIST, *Teruo Matsuzawa, Japan Advanced Institute of Science and Technology*

The Center for Information Science supports users world-class research and educational environment by providing a high-speed advanced information environment. A high-speed, high-availability network provides the foundation for the high performance file servers, massively parallel computers, and various servers that have enabled JAIST since its foundation to continuously provide users a convenient information environment. I would like to introduce the information environment in JAIST and my research of bio-fluid mechanics.

1:30 Cray Goes Bright for HPC Services, *Matthijs van Leeuwen, Bright Computing*

Bright Cluster Manager provides complete, end-to-end cluster management in one integrated solution: deployment, provisioning, monitoring, and management. Its intuitive GUI provides complete system visibility and ease of use for multiple clusters simultaneously; it's powerful cluster shell enables automated tasks and intervention. Bright scales from desk-side to TOP500 installations. Cray Custom Engineering has pioneered the use of Bright Cluster Manager® for external HPC systems: large-scale Lustre file systems, login servers, data movers, pre- and post-processing servers. Cray has also leveraged Bright to create additional services. This presentation is an overview of Bright Cluster Manager and its capabilities.

2:00 Scheduling Multi-Petaflop Systems and New Technologies, *Gary Brown, Adaptive Computing*

New multi-petaflop HPC systems present greater scheduling challenges to their effective utilization. New technologies, such as GPGPU, likewise present scheduling difficulties. The Moab scheduler has recent enhancements that address scheduling larger quantities of nodes, cores, and new resource types. This presentation introduces the latest Moab enhancements that make effective resource scheduling easier and more efficient.

18B

1:00 Future Proofing WL-LSMS: Preparing for First Principles Thermodynamics Calculations on Accelerator and Multicore Architectures, *Markus Eisenbach, Oak Ridge National Laboratory*

The WL-LSMS code has a very good track record for scaling on massively parallel architectures and achieves a performance of approx. 1.8 PF on the current Jaguar system at ORNL. Yet the code architecture assumes a distributed memory with a single thread of

execution per MPI rank, which is not a good fit for multicore nodes and the emerging accelerator based architectures. This talk will present the ongoing work to restructure the WL-LSMS code to take advantage of these new architectures and continue to work efficiently during the next decade.

1:30 Parallel Finite Element Earthquake Rupture Simulations on Quad- and Hex-Core Cray XT Systems, *Xingfu Wu, Oak Ridge National Laboratory and Texas A&M University, Benchu Duan and Valerie Taylor, Texas A&M University*

In this paper, we illustrate an element-based partitioning scheme for explicit finite element methods, and based on the partitioning scheme, we discuss how efficiently to use hybrid MPI/OpenMP to parallelize a sequential finite element earthquake rupture simulation code in order to not only achieve multiple levels of parallelism of the code but also to reduce the communication overhead of MPI within a multicore node by taking advantage of the shared address space and on-chip high inter-core bandwidth and low inter-core latency. We evaluate the hybrid MPI/OpenMP finite element earthquake rupture simulations on quad- and hex-core Cray XT 4/5 systems from Oak Ridge National Laboratory using the Southern California Earthquake Center (SCEC) benchmark TPV 210, which is to test convergence of TPV 10 problem with increasingly higher spatial resolutions (smaller element sizes). The benchmark solves dynamic rupture propagation along a 60° dipping normal fault (30 km x 15 km) and wave propagation in a homogeneous three-dimensional half space, and the initial stress on the fault linearly increases with depth. Our experimental results indicate that the parallel finite element earthquake rupture simulation obtains the accurate output results and has good scalability on these Cray XT systems.

2:00 Performance Characterization and Implications for Magnetic Fusion Co-Design Applications, *Praveen Narayanan, Alice Koniges, Leonid Oliker, and Robert Preissl, National Energy Research Scientific Computing Center, Samuel Williams, National Climate-Computing Research Center, Nicholas Wright, National Energy Research Scientific Computing Center, Stephane Ethier and Weixing, Princeton Plasma Physics Laboratory, Maxim Umansky and Xueqiao Xu, Lawrence Livermore National Laboratory, and Jeff Candy, General Atomics*

Co-design in high performance computing is a process that tightly couples applications and computer hardware architecture aimed at developing designs for the exascale level. In order to effectively perform the co-design process, it is important to study the current performance of the targeted applications. In this paper we present detailed benchmarking results for a set of magnetic fusion applications with a wide variety of underlying mathematical models including a particle code, and grid-based codes requiring both implicit and explicit numerical solvers. The analysis focuses on profiling these codes in terms of critical performance characteristics, which include such metrics as scalability, memory/network bandwidth limitations, communication versus computation and I/O versus computation. We compare and describe the available tools for performing this sort of study. The magnetic fusion codes represent a suite of applications that were selected as part of the co-design effort. Results are given for Cray XT4 and XE6 platforms.

19 General Session

2:45 LAC Appreciation

3:00 Next CUG-Stuttgart, *Neriman Emre, High Performance Computing Center Stuttgart (HLRS)*

3:15 Adjourn

Local Arrangements

How to Contact Us

After the conference

Barbara Horner-Miller (ARSC) Local Arrangements Chair

Arctic Region Supercomputing Center
P.O. Box 756020
Fairbanks, AK 99775-6020 USA
(1-907) 378-2907 Fax: (1-907) 450-8601
cug2011@cug.org

During the conference, May 23-26, 2011

CUG 2011 Conference Office
Westmark Hotel, Chena Room
813 Noble Street
Fairbanks, AK 99701-4977 USA
(1-907) 378-2907 Fax: (1-907) 451-7478
cug2011@cug.org

Conference Registration

James Rogers
Oak Ridge National Laboratory
1 Bethel Valley Road
Oak Ridge, TN 37831-6008 USA
(1-865) 576-2978 Fax: (1-865) 241-9578
jrogers@ornl.gov

Attendance and Registration

Who May Attend?

Based on CUG's corporate bylaws and guidelines, CUG meetings are not open sessions. Admittance to all meeting sessions 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.

In addition, 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 occasional invited speakers who bring 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 Office and Registration

Registration locations:
Sunday: Hall outside of the Gold Room
Monday: Minto Room
Tuesday – Thursday: Chena Room (Conference Office)

Registration hours:

Sunday	3:00 PM - 5:00 PM
Monday	7:00AM - 5:30 PM
Tuesday - Wednesday	8:00 AM - 5:30 PM
Thursday	8:00 AM - 3:00 PM

All attendees must wear badges during CUG Conference activities.

Conference Registration Fees

Your registration fee includes

- Admission to all program sessions, meetings, and tutorials
- Morning and afternoon breaks, and lunch Monday through Thursday
- CUG Night Out on Wednesday night

Proceedings

All conference attendees will have access to the Conference Proceedings on our secure web server. We will post a downloadable disk image of all the papers and slides and the complete program with papers and slides linked to it.

Cancellations

Conference registration cancellations must be received by the CUG Office before April 22, 2010. All registration fees will be refunded (less a $25 handling fee) if the cancellation was received by this date.

Special Assistance

For special assistance during the conference come to the Conference Office located in the Chena Room.

Dining Services

Lunch and Refreshments

Lunches and breaks are included in the conference registration fee and will be served in the East Gold Room.

Wednesday Night Out

Wednesday evening, the CUG Night Out, will include drinks and heavy hors d'oeuvres. Tickets for attendee's guests are available in the Conference Office. (See Social Events, p. 27 for details.)

Dining Out

There are many options for dining in Fairbanks during the conference. Please see the handout provided in your conference bag or stop by the Conference Office for suggestions.

On-Site Facilities

Messages

Call the hotel at (1-907) 456-7722 to leave telephone messages for attendees. Ask that the message be either delivered to your room or to the conference office during the conference (May 23-26).

Faxes

Please send faxes to your hotel for room delivery. The numbers are (1-907) 451-7478 for the Westmark Hotel and (1-907) 459-2720 for the Regency.

Wireless Access, E-mail, and Power

There is wireless internet access throughout the Westmark. A limited number of wired connections (ethernet) will be available in the E-mail room. Each conference meeting room will have power strips servicing the first few rows of seating. If you need power, sit near the front.

Image courtesy of FCVB

Panning for gold was once a way of life in Interior Alaska.

E-mail and Personal Computers

The E-mail room, available throughout the conference, is equipped with a limited number of personal computers with standard software and connected to the internet.

Speaker Preparation

We invite speakers to use the personal computers in the E-mail room for last minute changes. We will provide video projectors in all technical and general session rooms. Please find time before sessions and during breaks to test your laptop connection to the video projectors.

Local Arrangements

Image courtesy of FCVB

Visitors enjoy a sunny stroll through the Golden Heart Plaza in downtown Fairbanks.

Please come to the Conference Office if you need assistance.

Voltage

The United States operates on a 120 V. Those traveling from outside the United States will need an adapter.

Travel Information

Fairbanks Airport

A taxi between the airport and the hotel will cost approximately $20 and should take about 15 minutes.
Airlink, (1-907) 452-3337, provides shuttle service to local hotels for a minimum charge of $10.50 for 1-3 passengers, $3.50 for each additional passenger, and $1.00 per bag. Attendees staying at the Regency Hotel, (1-907) 459-2700, can call for complimentary shuttle service.

Ground Transportation

Taxis are available in Fairbanks although you might find the fare high. If you will do much auto travel, you may find it worthwhile to rent a car. There are several national car rental companies in Fairbanks as well as a few local ones. Both conference hotels have free day and overnight parking.
Some local attractions have bus service from the hotel to their venue. More information on these attractions and on the local bus service is provided in your registration bag.

Currency

Your best option to exchange foreign currency is at the airport in Seattle or another gateway city; there is not a foreign exchange at the Fairbanks airport. Foreign currency can be exchanged at banks near the hotel. Cash machines (ATMs) are available in the airport outside of security control and in the conference hotels.

Weather

May is a beautiful time of year in Fairbanks. The average high in Fairbanks in May is 60° F (16° C) and the average low is 36° F (2° C). Average precipitation is .6 inches (15.2mm). Expect well over 19 hours of daylight upon your arrival with a gain of 6+ minutes daily.

Tourist Information and Attractions

Tourist information is available from the CUG 2011 Local Arrangements pages at cug.org. Brochures and additional information are available at the hotel and the Conference Office. Check out the Fairbanks Visitor Bureau website for Fairbanks tourist suggestions, www.explorefairbanks.com

Cray Social, Alaska Salmon Bake, Tuesday at 6 PM

Image courtesy of FCVB

The Alaska Salmon Bake is a popular spot for visitors.

Cray Inc. will host a social on Tuesday evening at 6:00 PM at the Alaska Salmon Bake in Pioneer Park. All conference participants and their guests are invited to attend. The all-you-can-eat buffet features salmon, halibut, Bering Sea cod, crab, prime rib, side dishes and dessert. The local band, Celtic Confusion, will provide music. An indoor pavilion has been reserved for exclusive use by CUG. Take time to walk around the park, take a train ride, visit the old cabins or visit the Pioneer Air Museum. Pioneer Park, previously called Alaskaland, was the site of the 1995 Fall CUG Conference.

For those interested in seeing the 8:15 PM Golden Heart Revue at the Palace Theatre in Pioneer Park, reserve your seat by calling (1-907) 452-7274; say you're with CUG to get a special $10 price. Walk-ins can get that great rate if seats are still available.

CUG buses will leave the front of the Westmark Hotel at 5:45 PM and 6:00 PM. Blue Salmon Bake buses (not CUG-chartered) will leave from the Salmon Bake entrance periodically throughout the evening, including after the Palace Revue Show, to take you back to the hotels. Please feel free to hop on one whenever you are ready to go.

CUG Night Out, Fountainhead Antique Auto Museum, Wednesday at 6 PM

Image courtesy of FCVB

The CUG Board gets ready for a test drive at the Antique Auto Museum.

The CUG Night Out on Wednesday, open to all registered attendees, will be held at the Fountainhead Antique Auto Museum located on the grounds of the Wedgewood Resort. In addition to more than 70 antique cars, there is an historic photographic display depicting old Alaska. Numerous vintage clothing displays are located throughout the museum. This adds up to a very interesting venue for our night out.

The 17-piece Fairbanks Community Jazz Band will provide music for dancing and listening. A buffet of heavy hors d'oeuvres including reindeer sausage, salmon and halibut is provided.

There will be at least one photo opportunity available at the museum. Dress up in vintage clothing and sit in a vintage car while your friends snap your picture. Next to the museum is the Wedgewood Wildlife Sanctuary with 1.5 miles of nature trails going through a boreal forest and a model wetland. Should you want to step outside for some fresh air during the evening, take a quiet walk with the distinct possibility of seeing a variety of birds, small mammals, and possibly moose.

Buses will leave the Westmark Hotel at 5:45, 6:00 and 6:15 PM. They will depart the museum to return to the hotel at 9:00, 9:30 and 10:00 PM. If you wish to drive and need directions, pick them up in the Conference Office.

This event is open to all CUG attendees. Guest tickets may be purchased from the CUG Office for $50.

Contacts

CUG Board of Directors

President

Nicholas P. Cardo
National Energy Research Scientific Computing Center
1 Cyclotron Road, Bldg. 943r0256
Berkeley, CA 94720 USA
(1-510) 486-4765 Fax: (1-510) 486-4316
cardo@nersc.gov

Vice President

David Hancock
Indiana University
2711 East 10th Street,
Bloomington, IN 47408 USA
(1-812) 855-4021 Fax: (1-812) 855-8299
dyhancoc@indiana.edu

Treasurer

James Rogers
Oak Ridge National Laboratory
1 Bethel Valley Road
Oak Ridge, TN 37831-6008 USA
(1-865) 576-2978 Fax: (1-865) 241-9578
jrogers@ornl.gov

Secretary

Liam Forbes
Arctic Region Supercomputing Center
909 Koyukuk Drive, Suite 105
P.O. Box 756020
Fairbanks, AK 99775-6020 USA
(1-907) 450-8618 Fax: (1-907) 450-8601
lforbes@arsc.edu

Past President

David Gigrich
The Boeing Company
P.O. Box 3707
MC 7R-47
Seattle, WA 98124-2207 USA
(1-425) 865-6342 Fax: (1-425) 865-6844
david.j.gigrich@boeing.com

Director-at-Large

Mike Ashworth
STFC Daresbury Laboratory
Daresbury Science & Innovation Campus
Warrington, Cheshire WA4 4AD United Kingdom
(44-01925) 603178 Fax: (44-01925) 603634
m.ashworth@dl.ac.uk

Director-at-Large

Liz Sim
EPCC
The University of Edinburgh
James Clerk Maxwell Building
Mayfield Road
Edinburgh EH9 3JZ, United Kingdom
(44-0131) 6505029 Fax: (44-0131) 6506555
esim@epcc.ed.ac.uk

Director-at-Large

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

Local Arrangements Chair

Barbara Horner-Miller

Arctic Region Supercomputing Center
909 Koyukuk Drive, Suite 105
Fairbanks, AK 99775 USA
(1-907) 450-8606 Fax: (1-907) 450-8601
cug2011@cug.org

Cray Advisor to the CUG Board

Christy Adkinson

Cray Inc.
901 Fifth Avenue, Suite 1000
Seattle, WA 98164 USA
(1-206) 701-2118 Fax: (1-206) 701-2500
cadkinso@cray.com

Advisory Council Secretary and CUG.log Editor

Claire Hepwood

Atomic Weapons Establishment
E2.1, AWE Aldermaston
Reading, R40 4PR United Kingdom
(44-0118) 982 6886 Fax: (44-0118) 9857343
Claire.hepwood@awe.co.uk

SIG Chairs and Cray Inc. SIG Liaisons

Program Environments, Applications and Documentation

Yun (Helen) He, Chair

National Energy Research Scientific Computing Center
1 Cyclotron Road, MS 943
Berkeley, CA 94720-8043 USA
(1-510) 486-5180 Fax: (1-510) 486-4316
yhe@lbl.gov

Rolf Rabenseifner, Deputy Chair

High Performance Computing Center Stuttgart
Nobelstr. 19
D-70550 Stuttgart Germany
(49-71) 685 65530
rabenseifner@hlrs.de

Neil Stringfellow, Deputy Chair

CSCS Swiss National Supercomputing Centre
Galleria 2
Via Cantonale
Manno 6928 Switzerland
(41-91) 610 82 11 Fax: (41-91) 610 82 09
nstring@cscs.ch

Jef Dawson, Cray Liaison

Cray Inc.
380 Jackson Street, Suite 210
Saint Paul, MN 55101 USA
(1-651) 605-9036 Fax: (1-951) 605-9123
jef@cray.com

Luiz DeRose, Cray Liaison

Cray Inc.
380 Jackson Street, Suite 210
Saint Paul, MN 55101 USA
(1-651) 605-8925 Fax: (1-651) 605-9001
ldr@cray.com

Wayne Kugel, Cray Liaison

Cray Inc.
901 Fifth Avenue, Suite 1000
Seattle, WA 98164 USA
(1-651) 605-8963 Fax: (1-651) 605-9088
wjkugel@cray.com

Systems Support SIG

Joni Virtanen, Chair

CSC - IT Center for Science Ltd.
P.O. Box 405
FI-02101 Espoo Finland
(358-9) 457 2229
Joni.Virtanen@csc.fi

Charlie Carroll, Cray Liaison

Cray Inc.
380 Jackson Street, Suite 210
Saint Paul, MN 55101 USA
(1-651) 605-8987 Fax: (1-651) 605-9001
charliec@cray.com

Steve Johnson, Cray Liaison

Cray Inc.
901 Fifth Avenue, Suite 1000
Seattle, WA 98164 USA
(1-715) 726-4816 Fax: (1-715) 726-4739
sjj@cray.com

XTreme Systems

Jim Craw, Chair

National Energy Research Scientific Computing Center
1 Cyclotron Road
NERSC - 943-0217
Berkeley, CA 94720-8043 USA
(1-510) 486-5920 Fax: (1-510) 486-4316
craw@nersc.gov

Lee Whatley, Deputy

Navy DSRC
Lockheed Martin
1002 Balch Boulevard
Stennis Space Center, MS 39522 USA
(1-228) 688-4999
lwhatley.ctr@navo.hpc.mil

Vito Bongiorno, Cray Liaison

Cray Inc.
380 Jackson Street, Suite 210
Saint Paul, MN 55101 USA
(1-651) 605-8935 Fax: (1-651) 605-9123
vb@cray.com

CUG 2011 SPONSORS

The LAC would like to thank all of the people who have assisted in making this conference possible. We would especially like to thank our sponsors for their support.

Barbara Horner-Miller,
Local Arrangements Chair

Diamond

Platinum

Gold

Silver

Bronze

The Cray User Group is especially grateful to Cray Inc. for their ongoing support as we celebrate 33 years of CUG conferences and workshops.

Dear Friends and Colleagues,

The High Performance Computing Center Stuttgart (HLRS) at the University of Stuttgart, Germany invites you to attend CUG 2012. The conference theme "Greengineering the Future" reflects the fact that Stuttgart is the capitol of Baden-Württemberg which is the place where most engineering innovations worldwide have been created. Making these innovations turn the world into a greener one is the key challenge of the future. The conference will be held at the center of Stuttgart at the Maritim Hotel from April 29 – May 3 2012.

HLRS was created in 1995 as a spin-off of the computing center of the University of Stuttgart. HLRS was Germany's first national HPC Center providing researchers from all over Germany with access to leading edge technology in HPC. Right from the start HLRS created a subsidiary together with Porsche and T-Systems to provide supercomputers to industry. Most recently we have established the Automotive Simulation Center Stuttgart (ASCS) pooling researchers, automotive companies, hardware vendors and ISVs for precompetitive research in automotive simulation. Today HLRS is the leading European center for computational engineering in research and industry. As a main partner in the European HPC framework PRACE it provides users all over Europe with the most recent HPC technology. Starting in 2011 HLRS will operate a 110,000 core Cray XE6 which will be upgraded in 2013 to an expected 500,000 core system.

CUG 2012 will be the forum for Cray users and developers to exchange their expertise, problems, and solutions. Participants will have the opportunity to share ideas and experiences, and meet key members of the Cray team. Social events will allow communicating in a relaxed atmosphere. A visit to the Mercedes museum gives participants a view of 125 years of automobile development. From cars to stars!

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 and Stuttgart. Registration will be open by December 1, 2011.

We look forward to welcoming you to Stuttgart, the automobile center of Europe, where Daimler, Porsche, and Bosch made history and are about to launch their new green technology initiatives.

Sincerely,

Prof. Dr.-Ing. Dr. h.c. Michael Resch
HLRS Director
resch@hlrs.de

Ms. Neriman Emre
Local Arrangements Coordinator
emre@hlrs.de

Meeting Room Map

Westmark Fairbanks Hotel and Conference Center

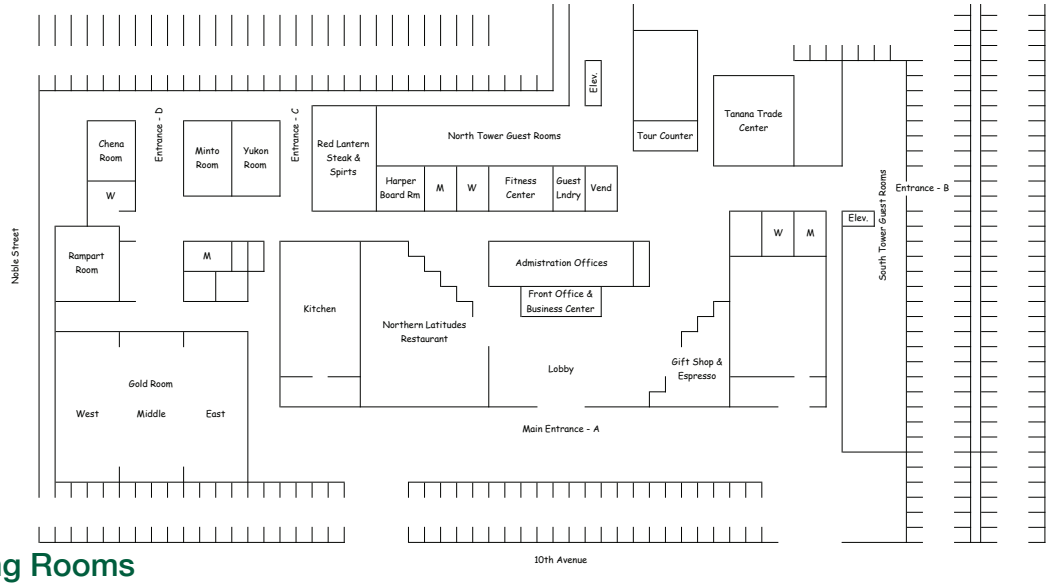

CUG 2011 Meeting Rooms

Event/Facility	Room
General Sessions	Gold West
Track A	Gold West
Track B	Yukon
Track C	Rampart
Conference Office	Chena
Email/Speaker Prep	Minto
Lunch & Breaks	Gold East
Board Room	Harper

CALL FOR PAPERS

You and your colleagues from CUG sites around the world hold the key to sustaining this conference as a source of information and technical interchange. Please begin now to consider presenting at next year's conference in Stuttgart. A year passes very quickly and our speaking requests have been growing consistently of late. Please share your experiences and help us work cooperatively toward solutions to problems encountered in the exciting venue of High Performance Computing.

As with this conference, Technical Presentations are invited in the following categories:

Applications and Programming Environments
- Programming Environment
- Compilers
- Libraries
- Tools
- 3rd Party Applications
- User Code Optimization

User Services
- Account Administration
- Consulting
- Documentation
- Training

Systems Support
- Operations
- Environmental Monitoring
- Facilities and Site Preparation
- Tuning and OS Optimizations
- System Operations
- Architecture
- Mass Storage
- Networking
- User Support

Take note of the presentations given this year, and discuss possible contributions with your colleagues. Seek advice from the Special Interest Group Chair or Deputy Chair. Work on your idea and prepare a draft to discuss at your home organization. Look for the official Call For Papers reminders you will receive in the fall and be ready to submit your Abstract to www.cug.org following the suggestions for publications listed therein. Help us make CUG in 2012 the best yet. Hope to see you in Stuttgart!