CUG 2010

SIMULATION COMES OF AGE

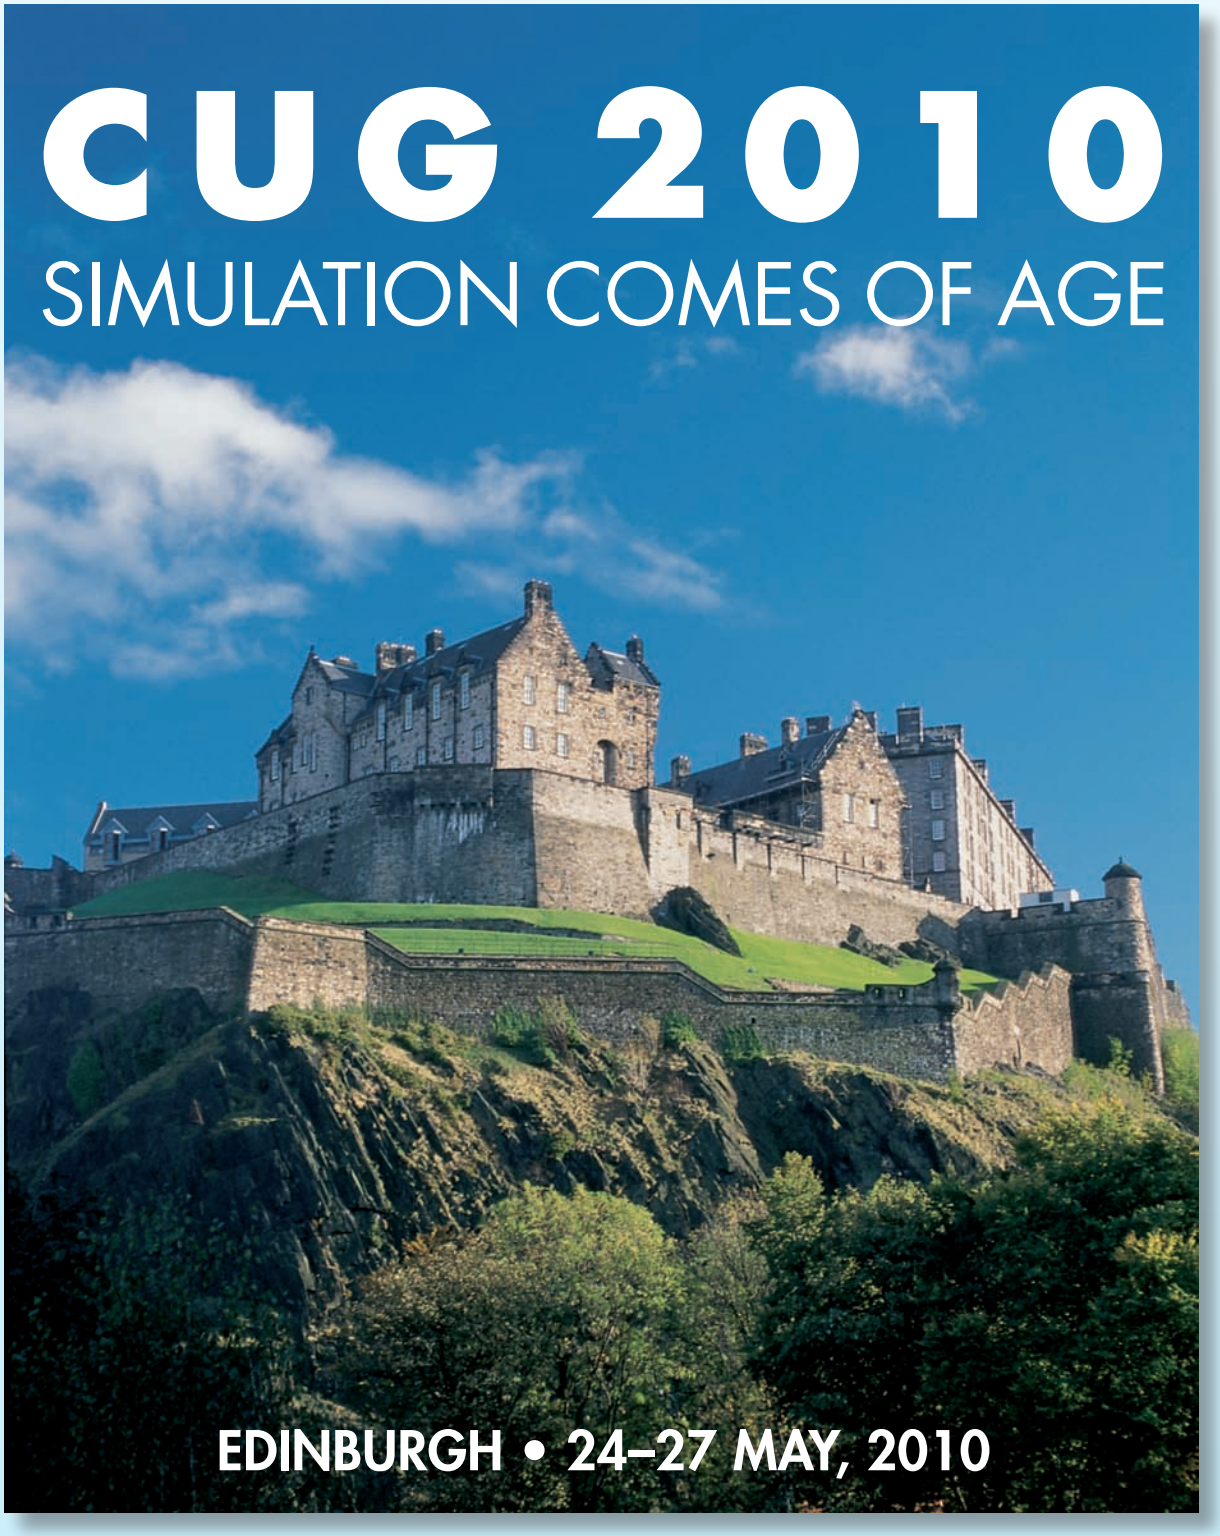

EDINBURGH • 24–27 MAY, 2010

Edinburgh Castle.

CONFERENCE PROGRAMME

Hosted by EPCC, the supercomputing centre at The University of Edinburgh
www.epcc.ed.ac.uk

epcc

Welcome

Old College, University of Edinburgh

Dear Friends,

EPCC at the University of Edinburgh warmly welcomes you to CUG 2010.

EPCC, a leading European centre of expertise in advanced research, technology transfer and the provision of supercomputer services to academia and business, is proud to host HECToR: a Cray system which acts as the UK's main national supercomputing service and is used to facilitate world leading research in a wide range of disciplines. This CUG meeting is happening at an exciting time as we are currently upgrading to XT6 technology: the first such system worldwide.

Our theme for CUG 2010 is

Simulation Comes of Age

This reflects that we have now reached another key point in the history of scientific research. Simulation has become established as the third fundamental research technique, alongside theory and experiment, enabling us to take our understanding and ability to a new level. Edinburgh is proud to have been one of the leading centres in this new methodology over the last 30 years. Simulation can be used to tackle problems that are too big, small, distant, quick or slow for experiment and too complex to solve analytically. This not only permits a deepening of our scientific understanding, but also contributes to important developments in areas that affect our everyday lives such as health, the environment, energy and industry.

We are delighted that you have joined us in Edinburgh for what we expect to be a thoroughly enjoyable CUG 2010.

Sincerely,

Professor Arthur Trew, EPCC Director

Dr. Alan Gray, Local Arrangements Chair

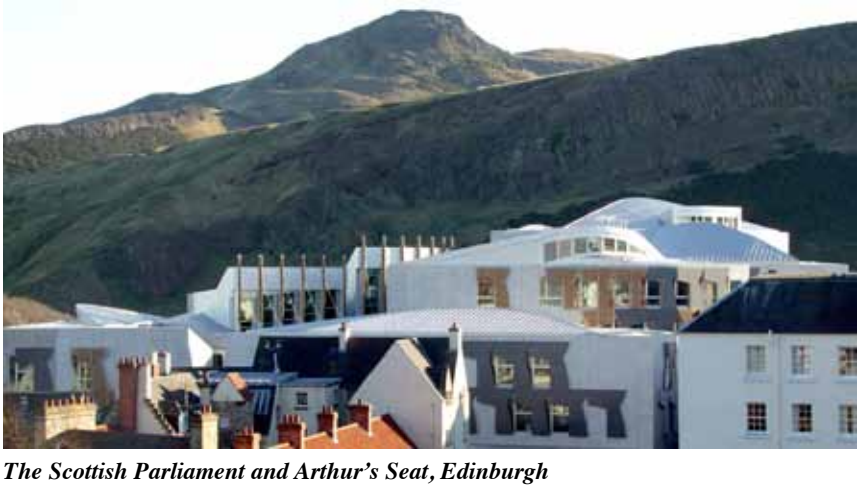

The Scottish Parliament and Arthur's Seat, Edinburgh

Welcome

On Behalf of the CUG Program Committee

Welcome to the 52nd Cray User Group meeting, CUG 2010 **Simulation Comes of Age**. This technical conference provides a unique opportunity for you to exchange problem-solving information and enjoy professional interactions with your Cray Inc. high performance computing system colleagues.

The CUG Program Committee has assembled a diverse and impressive array of detail-packed presentations in General and Parallel Technical Sessions, Tutorials, and Special Interest Group (SIG) meetings. The technical program includes many talks encompassing all aspects of the system from managing the system to taking full advantage of the architecture. A number of talks will highlight the improvements and benefits to the scientific community. Furthermore, Cray Inc. is committed to having many technical experts on site throughout the entire program. In addition, we expect many of Cray's vendor partners to participate in the technical program and to be available for informal discussions.

- General Sessions provide you with the latest corporate and technical information from Cray Inc. executives, as well as topics of general interest to this audience.
- Parallel Technical Sessions give you the opportunity to focus on the specific knowledge domains of the SIGs. The presentations in the technical sessions have been reviewed and selected for you by the Program Committee.
- Tutorials are a great opportunity for you to update your technical skills with the help of selected experts from Cray Inc. and/or other CUG sites.
- SIG Meetings include discussion of the technical issues surrounding Legacy Systems, Applications and Programming Environment, User Support and System Support. These meetings are where we work together to maintain and improve our CUG technical program.

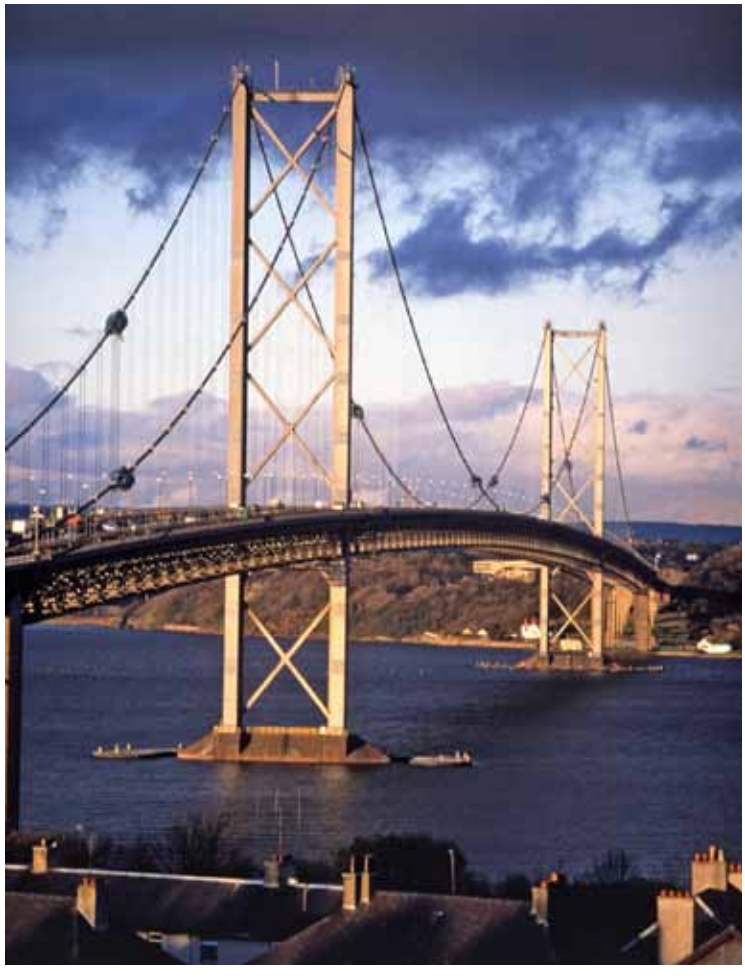

Forth Road Bridge

In addition to these prepared presentations and workshops, 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 Social, luncheons, and breaks 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. and their technology partners.

We are excited to have Professor Richard Catlow from the University College London as our keynote speaker. Professor Catlow will present a selection of research highlights from the HPC Materials Chemistry Consortium.

This year's program includes two additional General Session talks:

- *Jaguar—The World's Most Powerful Computer System*, **Buddy Bland**, Director of the Leadership Computing Facility project, Oak Ridge National Laboratory (ORNL)
- *Virtual Palaeontology: Gait Reconstruction of Extinct Vertebrates Using High Performance Computing,* **Dr. Bill Sellars**, The University of Manchester

CUG 2010 will be a rewarding, educational and memorable experience for all attendees. I look forward to seeing you in Edinburgh, United Kingdom, at CUG 2010, **Simulation Comes of Age**.

Nicholas P. Cardo

CUG Vice-President and Program Chair

National Energy Research
Scientific Computing Center

CRAY USER GROUP INCORPORATED

Program Notes

Contents

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!

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 SIG subject area. 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.

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. Please 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

The CUG 2010 Evaluation 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 it. Once you've completed it, please stop by the CUG office. We have a gift for you.

Edinburgh Castle viewed from the Old Town

Keynote Speaker Bio

Richard Catlow

Professor Richard Catlow has worked for over thirty years in the field of computational and experimental studies of complex inorganic materials. His group has pioneered a wide range of applications of computational techniques in solid state chemistry to systems and problems including microporous and oxide catalysts, ionic conductors, electronic ceramics and silicate minerals. This applications programme has been supported by technique and code development, including recent work on embedded cluster methodologies for application to the study of catalytic reactions. The computational work has been firmly linked with experimental studies, using both neutron scattering and synchrotron radiation techniques, where the Royal Institute group has also made notable contributions to development as well as application studies. Professor Catlow's research has led to over 800 publications, and in 2004 he was elected to the Fellowship of the Royal Society for "pioneering the development and application of computer modeling in solid state and materials chemistry." He is currently Dean of the Mathematical and Physical Sciences Faculty at University College London. Professor Catlow is the Principal Investigator for the "HPC Materials Chemistry Consortium" project on the national UK "HECToR" supercomputing service, which is based on Cray hardware. This project, which is one of the largest in terms of resource allocation, involves an extensive collaboration and spans a wide range of application studies aimed at modelling and predicting the structures, properties and reactivities of functional materials.

Victoria Street, Old Town, Edinburgh

Edinburgh Castle

For the location of Tutorials, General Sessions, Tracks A-C, BoFs, Lunch and Breaks, see the Meeting Room Map on the inside back cover.

MONDAY

1A Tutorial

8:00 Cray Debugging Support Tools for Petascale Applications, ***Bob Moench, Cray Inc.***

As HPC systems have gotten ever larger, the amount of information associated with a debugging failing parallel application has grown beyond what the beleaguered applications developer has the time, resources, and wherewithal to analyze. With the release of the Cray Debugging Support package, Cray introduces several innovative methods of attacking this vexing problem. FTD (Fast Track Debugging) achieves debugging at fully optimized speeds. STAT (Stack Trace Analysis Tool) facilitates the evaluation and study of hung applications. ATP (Abnormal Termination Processing) captures a STAT-like view of applications that have taken a fatal trap. And Guard, the Cray comparative debugger, delivers an automated search for the location of program errors by comparing a working version of an application against a failing version. This paper describes and explores each of the above technologies.

1B Tutorial

8:00 Scaling Applications on Cray XT Systems, ***Luiz DeRose, John Levesque, and Bob Moench, Cray Inc.***

In this tutorial we will present tools and techniques for application performance tuning on the Cray XT system, with focus on multi-core processors. Attendees will learn about the Cray XT architecture and its programming environment. They will have an initial understanding of potential causes of application performance bottlenecks, and how to identify some of these bottlenecks using the Cray Performance tools. In addition, attendees will learn advanced techniques to deal with scaling problems and how to access the on-line documentation for user help. Attendees will also have some exposure to the Cray debugging support tools, which provide innovative techniques to debug applications at scale.

1C Tutorial

8:00 An Overview of the Chapel Programming Language and Implementation, ***Steve Deitz, Cray Inc.***

Chapel is a new parallel programming language under development at Cray Inc. as part of the DARPA High Productivity Computing Systems (HPCS) program. Chapel has been designed to improve the productivity of parallel programmers working on large-scale supercomputers as well as small-scale, multicore computers and workstations. It aims to vastly improve programmability over current parallel programming models while supporting performance and portability at least as good as today's technologies. In this tutorial, we will present an introduction to Chapel, from context and motivation to a detailed description of Chapel via many example computations. This tutorial will focus on writing Chapel programs for both multi-core and distributed-memory computers. We will explore the optimizations added to the Chapel implementation this past year that helped with the most recent Chapel HPCC entry.

2 Opening General Session

11:00 CUG Welcome, ***David Gigrich, CUG President, The Boeing Company (BOEING)***

11:10 Welcome, ***Arthur Trew, Director, EPCC, University of Edinburgh***

11:15 Keynote Address: Massively Parallel Molecular Modeling of Heterogeneous Catalysis, ***Professor Richard Catlow, Dean, Mathematical and Physical Sciences Faculty, University College London***

We describe the development and application of a massively parallel version of the ChemShell code for the study of catalysis. This package combines classical potential modeling and density functional theory and allows us to study the influence of metal oxide surfaces on the chemistry of adsorbed molecules. We have developed a number of massively parallel schemes which compile parallel versions of the component codes with a task farming approach to divide the problem further and improve efficiency on very large numbers of processors. The work has been performed as part of the dCSE support for the HECToR service.

12:00 Lunch

3A

1:00 Tools, Tips and Tricks for Managing Cray XT Systems, ***Kurt Carlson, Arctic Region Supercomputing Center (ARSC)***

Managing large complex systems requires processes beyond what is taught in vendor training classes. Many sites must manage multiple systems from different vendors. This paper covers a collection of techniques to enhance the usability, reliability and security of Cray XT systems. A broad range of activities, from complex tasks like security, integrity and environmental checks of the Cray Linux Environment, to relatively simple things like making 'rpm -qa' available to users will be discussed. Some techniques will be XT specific, such as monitoring L0/L1 environment, but others will be generic, such as security tools adapted from other systems and re-spun as necessary for the XT Cray Linux Environment.

1:30 RAVEN: RAS Data Analysis Through Visually Enhanced Navigation, ***Al Geist, Guruprasad Kora, and Byung-Hoon Park, Oak Ridge National Laboratory (ORNL) and Junseong Heo, National Institute for Computational Sciences (NICS)***

Supercomputer RAS data contain various signatures regarding system status, thus are routinely examined to detect and diagnose faults. However, due to voluminous sizes of logs generated during faulty situations, a comprehensive investigation that requires comparisons of different types of RAS logs over both spatial and temporal dimensions is often beyond the capacity of human operators, which leaves a cursory look to be the only feasible option. As an effort to better embrace informative but huge supercomputer RAS data in a fault diagnosis/detection process, we present a GUI tool called RAVEN that visually overlays various types of RAS logs on a physical system map where correlations between different fault types can be easily observed in terms of their quantities and locations at a given time. RAVEN also provides an intuitive fault navigation mechanism that helps examine logs by clustering them to their common locations,

types, or user applications. By tracing down notable fault patterns reflected on the map and their clustered logs, and superimposing user application data, RAVEN, which has been adopted at National Institute of Computational Science (NICS) at the University of Tennessee, identified root causes of several system failures logged in Kraken XT5.

2:00 Correlating Log Messages for System Diagnostics, ***Al Geist, Raghul Gunasekaran, Byung Park, and Galen Shipman, Oak Ridge National Laboratory (ORNL)***

In large-scale computing systems the sheer volume of log generated has challenged the interpretation of log messages for debugging and monitoring purposes. For a non-trivial event, the Jaguar XT5 at the Oak Ridge Leadership Computing Facility with more than eighteen thousand compute nodes would generate a few hundred thousand log entries in less than a minute. Determining the root cause of such events requires analyzing and understanding these log messages. Most often, these log messages are best understood when they are interpreted collectively rather than being read as individual messages. In this paper, we present our approach to interpreting log messages by identifying commonalities and grouping them into clusters. Given a set of log messages within a time interval, we parse and group the messages based on source, target, and/or error type, and correlate the messages with hardware and application information. We monitor the XT5's console, netwatch and sys log and show how such grouping of log messages help in detecting system events. By intelligent grouping and correlation of events from multiple sources we are able to provide system administrators with meaningful information in a concise format for root cause analysis.

3B

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, which consists of state of the art compiler, tools, and libraries, supporting a wide range of programming models.

1:30 Dynamic Shared Libraries and Virtual Cluster Environment, ***Jason Schildt, Cray Inc.***

Cray is expanding system functionality to support Dynamic Shared Libraries (DSL) on compute nodes, and the ability to run a wide range of packaged ISV applications on compute nodes. Built upon Data Virtualization Service (DVS), a more standard Linux runtime environment is distributed across the system by the DSL capability via DVS Server nodes to the Compute Node clients. The CLE Virtual Cluster Environment (VCE) adds a further layer of functionality, by supporting natively installed and executed ISV applications. This three component solution allows customers to meet a wide range of runtime environment demands with limited impact and complexity while increasing productivity.

2:00 DMAPP—An API for One-sided Program Models on Baker Systems, ***Monika ten Bruggencate, Cray Inc.***

Baker Systems and follow-on systems will deliver a network with advanced remote memory access capabilities. A new API (DMAPP) has been developed to expose these capabilities to one-sided program models. This paper presents the DMAPP API as well as some preliminary performance data.

3C

1:00 The NEMO Ocean Modelling Code: A Case Study, ***Fiona Reid, EPCC (EPCC)***

We present a case study of a popular ocean modelling code, NEMO, on the Cray XT4 HECToR system. HECToR is the UK's high-end computing resource for academic users. Two different versions of NEMO have been investigated. The performance and scaling of the code has been evaluated and optimised by investigating the choice of grid dimensions, by examining the use of land versus ocean grid cells and also by checking for memory bandwidth problems. The code was profiled and the time spent carrying out file input/output was identified to be a potential bottleneck. We present a solution to this problem which gives a significant saving in terms of runtime and disk space usage.

1:30 Use of the Cray XT5 Architecture to Push the Limits of WRF Beyond One Billion Gridpoints, ***Don Morton and Oralee Nudson, Arctic Region Supercomputing Center (ARSC)***

The Arctic Region Supercomputing Center (ARSC) Weather Research and Forecasting (WRF) model benchmark suite continues to push software and available hardware limits by successfully running a 1km resolution case study composed of more than one billion grid points. Simulations of this caliber are important for providing detailed weather forecasts over the rugged Alaska terrain and are intended for benchmarking on systems with tens of thousands of cores. In pursuing these large scale simulations, we have incurred numerical, software and hardware limitations that have required us to use various parallel I/O schemes and to explore different PBS "aprun" options. In this paper we will discuss issues encountered while gradually expanding the problem sizes in which WRF can operate and our solutions in running high resolution and/or large-scale WRF simulations on the Cray XT5 architecture.

2:00 Optimising and Configuring the Weather Research and Forecast Model on the Cray XT, ***Mike Ashworth and Andrew Porter, STFC Daresbury Laboratory***

The Weather Research and Forecast (WRF) Model is a well-established and widely used application. Designed and written to be highly scalable, the code has a large number of configuration options at both compile- and run-time. We report the results of an investigation into the effect of these options on the performance of WRF on a Cray XT4 with a typical scientific use-case. Covering areas such as MPI/OpenMP comparison, cache usage and I/O performance, we discuss the implications for both regular WRF users and the authors of other application codes.

4A

3:00 Automatic Library Tracking Database, ***Mark Fahey, Bilel Hadri and Nicholas Jones, National Institute for Computational Sciences (NICS)***

The National Institute for Computational Sciences and the National Center for Computational Sciences (both located at Oak Ridge National Laboratory) have been working on an automatic library tracking database whose purpose is to track which libraries are used on their Cray XT5 Supercomputers. The database stores the libraries that are used at link time and it records which executable is run during a batch job. With this data, many operationally important questions can be answered like which libraries are most frequently used and who is using deprecated libraries or applications. The infrastructure design and reporting mechanisms will be presented with production data to this point.

3:30 A Pedagogical Approach to User Assistance, ***Robert Whitten, Jr., Oak Ridge National Laboratory (ORNL)***

This presentation will focus on a pedagogical approach to providing user assistance. By making user education the central theme in training, outreach, and user assistance activities, a set of competencies can be developed that encompasses the knowledge required for productive use of leadership-class computing resources such as the Cray XT5 Jaguar system.

4:00 PRACE Application Enabling Work at EPCC, ***Xu Guo and Joachim Hein, EPCC (EPCC)***

The Partnership for Advanced Computing in Europe (PRACE) created the prerequisites for a pan-European HPC service, consisting of several tier-0 centres. PRACE's aim has now moved to the implementation of this service. The now completed work looked into all aspects of the pan-European service, including the contractual and organisational issues, the system management, application enabling and future computer technologies. This talk discusses the work done by EPCC on the application codes HELIUM (from Queen's University Belfast, UK) and NAMD (from University of Illinois at Urbana Champaign, US) with a particular focus on the work carried out for the Prace prototype Louhi, which is Cray XT5 at CSC in Finland. We will also include a performance comparison with non-Cray systems available to PRACE.

4:30 Towards a European Training Network in Computational Science, ***Pekka Manninen and Ari Turunen, CSC-Scientific Computing Ltd. (CSC)***

The implementation phase of The Partnership for Advanced Computing in Europe (PRACE) project will develop and maintain a European training network in the field of computational science. Its key ingredients are solid contacts between the partner organisations and European research centres, as well as establishing new links to universities. In this talk, I will review the completed training-related activities of the preparatory phase of PRACE as well as plans for the implementation phase.

4B

3:00 Application Acceleration on Current and Future Cray Platforms, ***Jihan Kim, Alice Koniges, Robert Preissl, and John Shalf, National Energy Research Scientific Computing Center (NERSC) and David Eder, Aaron Fisher, Nathan Masters, and Mlaker Velimir, Lawrence Livermore National Laboratory***

Application codes in a variety of areas are being updated for performance on the latest architectures. We describe current bottlenecks and performance improvement areas for applications including plasma physics, chemistry related to carbon capture and sequestration, and material science.

3:30 General Purpose Timing Library (GPTL): A Tool for Characterizing Performance of Parallel and Serial Applications, ***James Rosinski, Oak Ridge National Laboratory (ORNL)***

GPTL is an open source profiling library that reports a variety of performance statistics. Target codes may be parallel via threads and/or MPI. The code regions to be profiled can be hand-specified by the user, or GPTL can define them automatically at function-level granularity if the target application is built with an appropriate compiler flag. Output is presented in a hierarchical fashion that preserves parent-child relationships of the profiled regions. If the PAPI library is available, GPTL utilizes it to gather hardware performance counter data. GPTL built with PAPI support is installed on the jaguar machine at ORNL.

4:00 Five Powerful Chapel Idioms, ***Brad Chamberlain, Sung-Eun Choi, Steve Deitz, David Iten, and Lee Prokowich, Cray Inc.***

The Chapel parallel programming language, under development at Cray Inc., has the potential to deliver high performance to more programmers with less effort than current practices provide. This is especially the case with the many-core architectures that are already becoming more and more prevalent. This paper presents five reasons why: 1. Chapel supports easy-to-use asynchronous and synchronous remote tasks, 2. Chapel supports local and remote transactions, 3. Chapel supports simple data-parallel abstractions when applicable, 4. Chapel supports user-defined data distributions, and 5. Chapel supports arbitrarily nested parallelism.

4:30 Overview and Performance Evaluation of Cray LibSci Products, ***Adrian Tate, Cray Inc.***

This talk serves as both an introduction to the Cray scientific library suite and as a tutorial on obtaining advanced performance with applications that utilize scientific libraries. The talk will include a thorough and frank performance evaluation of all scientific library products on Cray XT systems, including dense kernels on single core and multiple cores, dense linear solvers and eigensolvers in serial and parallel, serial and distributed Fourier Transforms and Sparse kernels within sparse iterative solvers. The emphasis will be on usage and how to increase performance by using different algorithms or libraries, better configurations, or advanced controls of the scientific libraries.

4C

3:00, 3:30, 4:00, and 4:30 XTreme SIG, ***Chair, James Craw, National Energy Research Scientific Computing Center (NERSC)***

This group works very closely with Cray, under Non-Disclosure Agreements, to provide valuable input into the Cray XT system development cycle. For this reason, these are "closed door" sessions.

5A Interactive Session

5:15 Open Discussion with the CUG Board, ***Chair: David Gigrich, CUG President, The Boeing Company (BOEING)***

The CUG Board members will be present to hold open discussions with member institutions. This is an excellent opportunity to learn more about decisions that have been made or to share your concerns with the Board.

5B Interactive Session

5:15 Cray Knowledge Management

5C Birds of a Feather

5:15 Dynamic Shared Libraries and Virtual Cluster Environment, ***Chair: Jason Schildt, Cray Inc.***

As a follow-up to the Dynamic Shared Libraries and Virtual Cluster Environment paper/presentation, Jason Schildt will chair a technical BOF covering deployment, configuration and functionality focused questions from the participants.

6:00 Drinks Reception

A short complimentary drinks reception to welcome CUG 2010 attendees to Edinburgh. Elliot's Bar (within the conference hotel, next to the reception desk).

Tuesday

6 General Session

8:30 Introduction, ***Nicholas P. Cardo, Program Chair, National Energy Research Supercomputing Center (NERSC)***

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 Jaguar-The World's Most Powerful Computer System, ***Arthur Bland, Oak Ridge National Laboratory (ORNL)***

At the SC'09 conference in November 2009, Jaguar was crowned as the world's fastest computer by the web site www.Top500.org. In this paper, we will describe Jaguar, present results from a number of benchmarks and applications, and talk about future computing in the Oak Ridge Leadership Computing Facility.

11:00 1 on 100 (or more), ***Peter Ungaro, President and CEO, Cray Inc. (No other Cray personnel and no Cray Vendor Partners please.)***

12:00 Lunch

8A

1:00 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.

1:30 XT System Reliability: Metrics, Trends, and Actions, ***Steve Johnson, Cray Inc.***

In 2009, the XT product family saw a significant improvement in overall reliability as measured by Cray's support organization. This paper will discuss the reliability trends that have been observed and the main reasons for the improvements. We will also discuss the tools used to collect the field data, the metrics generated by Cray to evaluate XT product reliability and the actions taken as a result of this analysis.

2:00 ALPS, Topology, and Performance, ***Carl Albing, Cray Inc.***

Application performance can be improved or reduced depending on the compactness of the set of nodes on which an application is placed (as demonstrated convincingly by PSC at a recent CUG). This paper describes the approach to placements that ALPS now uses based on the underlying node topology, the reasons for this approach, and the variations that sites can use to optimize for their specific machine and workload.

8B

1:00 Interactions Between Application Communication and I/O Traffic on the Cray XT High Speed Network, ***R. Glenn Brook and Lonnie D. Crosby, National Institute for Computational Sciences (NICS)***

The massive size of modern leadership computing resources often leads to the discovery of application performance bottlenecks not seen at smaller scales. Many of these performance bottlenecks originate within individual applications; however, recent application testing on a Cray XT5 indicates that an application's I/O pattern can negatively impact the communication performance of another application via interaction over the shared high speed network (HSN). This study seeks to identify and to quantify such interactions on the HSN of Kraken, the Cray XT5 operated by the National Institute for Computational Sciences (NICS).

1:30 Optimisation of the I/O for Distributed Data Molecular Dynamics Applications, ***Ian Bush, Numerical Algorithms Group and William Smith and Ilian Todorov, STFC Daresbury Laboratory Laboratory***

With the increase in size of HPC facilities it is not only the parallel performance of applications that is preventing greater exploitation, in many cases it is the I/O which is the bottleneck. This is especially the case for distributed data algorithms. In this paper we will discuss how the I/O in the distributed data molecular dynamics application DL_POLY_3 has been optimised. In particular we shall show that extensive data redistribution specifically to allow best use of the I/O subsystem can result in a code that scales to many more processors, despite the large increase in communications required.

2:00 Analyzing Multicore Characteristics for a Suite of Applications on an XT5 System, ***Douglas Doerfler and Courtenay Vaughan, Sandia National Laboratories (SNLA)***

In this paper, we will explore the performance of applications important to Sandia on an XT5 system with dual socket AMD 6 core Istanbul nodes. We will explore scaling as a function of the number of cores used on each node and determine the effective core utilization as core count increases. We will then analyze these results using profiling to better understand resource contention within and between nodes.

8C

1:00 Performance Monitoring Tools for Large Scale Systems, ***David Dillow, Jason Hill, Don Maxwell, Ross Miller, Sarp Oral, Galen Shipman, and Feiyi Wang, Oak Ridge National Laboratory (ORNL)***

Operating computing systems, file systems, and associated networks at unprecedented scale offer unique challenges for fault monitoring, performance monitoring and problem diagnosis. Conventional system monitoring tools are insufficient to process the increasingly large and diverse volume of performance and status log data produced by the world's largest systems. In addition to the large data volume, the wide variety of systems employed by the largest computing facilities present diverse information from multiple sources, further complicating analysis efforts. At leadership scale, new tool development is required to acquire, condense, correlate, and present status and performance data to systems staff for timely evaluation. This paper details a set of system monitoring tools developed by the authors and utilized by systems staff at Oak Ridge National Laboratory's Leadership Computing Facility, including the Cray XT5 Jaguar. These tools include utilities to correlate I/O performance and event data with specific systems, resources, and jobs. Where possible, existing utilities are incorporated to reduce development effort and increase community participation. Future work may include additional integration among tools and implementation of fault-prediction tools.

1:30 Hierarchy Aware Blocking and Nonblocking Collective Communications—The Effects of Shared Memory and Torus Topologies in the Cray XT5 Environment, ***Richard Graham and Joshua Ladd, Oak Ridge National Laboratory (ORNL)***

MPI Collective operations tend to play a large role in limiting the scalability of high-performance scientific simulation codes. As such, developing methods for improving the scalability of these operations is critical to improving the scalability of such applications. Using infrastructure recently developed in the contest of the FASTOS program we will study the performance of blocking collective operations, as well as those of the recently added MPI nonblocking collective operations taking into account both shared memory and network topologies.

2:00 Multi-core Programming Paradigms and MPI Message Rates—A Growing Concern?, ***Scott Hemmert, Sandia National Laboratories (SNLA)***

The continued growth in per-node core count in high performance computing platforms has lead the community to investigate alternatives to an MPI-everywhere programming environment. A hybrid programming environment, in which MPI is used for coarse-grained, inter-node parallelism and a threaded environment (pthreads, OpenMP, etc.) is used for fine-grained, intra-node parallelism presents an appealing target for future applications. At the same time, memory and network bandwidth both continue to grow at a significantly slower pace than processor performance. This trend, combined with increased parallelism due to larger machine sizes, will drive applications away from the bandwidth-limited BSP model to one with a higher number of smaller messages, which avoids unnecessary memory-to-memory copies inside a single node. The increase in small message transfers requires a higher message rate from a single node. Current network designs rely on a number of tasks on a single node injecting messages into the network in order to achieve optimal message rates. This paper quantifies the impact of local process count on node-level message rate for Cray XT5 hardware. The results are an important metric in designing both MPI implementations and applications for the hybrid programming future.

9A

3:00 Franklin Job Completion Analysis, ***Yun (Helen) He, Hwa-Chun Wendy Lin, and Woo-Sun Yang, National Energy Research Scientific Computing Center (NERSC)***

The NERSC Cray XT4 machine Franklin has been in production for 3000+ users since October 2007, where about 1800 jobs were run each day. There has been an on-going effort to better understand how well these jobs run, whether failed jobs are due to application errors or system issues, and to further reduce system related job failures. In this paper, we will talk about the progress we made in tracking job completion status, in identifying job failure root cause, and in expediting resolution of job failures, such as hung jobs, that are caused by system issues. In addition, we will present some Cray software design enhancements we requested to help us track application progress and identify errors.

3:30 Collecting Application-Level Job Completion Statistics, ***Matthew Ezell, National Institute for Computational Sciences (NICS)***

Job failures are common on large high performance computing systems, but logging, analyzing, and understanding the low-level error messages can be difficult on Cray XT systems. This paper describes a set of tools to log and analyze applications in real-time as they run on the system. By obtaining more information about typical error scenarios, system administrators can work to resolve the underlying issues and educate users.

4:00 Using Quality of Service for Scheduling on Cray XT Systems, ***Troy Baer, National Institute for Computational Sciences (NICS)***

The University of Tennessee's National Institute for Computational Sciences (NICS) operates two Cray XT systems for the U.S. National Science Foundation (NSF): Kraken, an 88-cabinet XT5 system, and Athena, a 48-cabinet XT4 system. Access to Kraken is allocated through the NSF's Teragrid allocations process, while Athena is currently being dedicated to individual projects on a quarterly basis; as a result, the two systems have somewhat different scheduling goals. However, user projects on both systems have sometimes required the use of quality of service (QoS) levels for scheduling of certain sets of jobs. We will present case studies of three situations where QoS levels were used to fulfill specific requirements: two on Kraken in fully allocated production service, and one on Athena while dedicated to an individual project. These case studies will include lessons learned about impact on other users and unintended side effects.

9B

3:00 FutureGrid: Design and Implementation of a National Grid Test-Bed, ***David Hancock and Gregor von Laszewski, Indiana University (INDIANAU)***

Indiana University is leading the creation of a grid test-bed for the National Science Foundation with nine partner institutions. FutureGrid is a high performance grid test-bed that will allow scientists to work collaboratively to develop and test novel approaches to parallel, grid, and cloud computing.

3:30 Scalable Performance Analysis of Large-scale Parallel Applications on Cray XT Systems with Scalasca, ***Brian Wylie, Juelich Supercomputing Centre***

The open-source Scalasca toolset [www.scalasca.org] supports integrated runtime summarization and automated trace analysis on a diverse range of HPC computer systems. An HPC-Europa2 visit to EPCC in 2009 resulted in significantly enhanced support for Cray XT systems, particularly the auxiliary programming environments and hybrid OpenMP/MPI. Combined with its previously demonstrated extreme scalability and portable performance analyses comparison capabilities, Scalasca has been used to analyse and tune numerous key applications (and benchmarks) on Cray XT and other PRACE prototype systems, from which experience with a representative selection is reviewed.

4:00 Evaluation of Productivity and Performance Characteristics of CCE CAF and UPC Compilers, ***Sadaf Alam, Matthew Cordery, William Sawyer, Tim Stitt, and Neil Stringfellow, CSCS-Swiss National Supercomputing Centre (CSCS)***

The Co-Array Fortran (CAF) and Unified Parallel C (UPC) functional compilers available with the Cray Compiler Environment (CCE) on the Cray XT5 platform offer an integrated framework for code development and execution for Partitioned Global Address Space (PGAS) programming paradigm together with message-passing MPI and shared-memory OpenMP programming models. Using micro-benchmarks, conformance test cases and micro-kernels of representative scientific calculations, we attempt to evaluate the following characteristics of the CCE PGAS compilers: (1) usability of the framework for code development and execution; (2) completeness and integrity of code generation; (3) efficiency of the generated code, particularly usage of the communication layer (GASNet on SeaStar2); and (4) tools availability for performance measurement and diagnostics. Our initial results show that the current version of compiler provides a highly productive code development environment for CAF or UPC code development on our target Cray XT5 platform. At the same time however, we observe that the code transformation and generation processes are unable to aggregate remote memory access for simple access patterns causing significant slowdown. We will compare and contrast code generation with two multi-platform PGAS compilers: Berkley UPC environment that uses the Intrepid UPC compiler and the g95 CAF compiler extensions. In the full paper version, we would also include comparative results using the Rice CAF 2.0 compiler, if it becomes available in due time.

9C

3:00, 3:30, and 4:00 XTreme SIG, ***Chair, James Craw, National Energy Research Scientific Computing Center (NERSC)***

This group works very closely with Cray, under Non-Disclosure Agreements, to provide valuable input into the Cray XT system development cycle. For this reason, these are "closed door" sessions.

10A Interactive Session

4:45 User Services SIG, ***Chair: Jim Glidewell, The Boeing Company (BOEING)***

The purpose of this Interactive Session is to provide an open forum for discussion of any issues related to the services or processes relating to User Services, whether supplied by Cray or provided internally by operating sites. This is an opportunity to help each other and learn from the collective experience. This session is open to everyone wanting to learn about the area of User Services.

10B Birds of a Feather

4:45 LSI Storage Best Practices for Deploying and Maintaining Large Scale Parallel File Systems Environments, ***Gava Didier, LSI***

This BoF session will focus on the challenges and rewards of deploying and implementing a parallel file system to improve cluster performance. The discussion will focus on the impact that deployment and ongoing support has in terms of system performance, system availability and pain for support personnel. Discussion will highlight experiences with different file systems, different types of platform approaches, and balancing vendor support vs. in-house support. We will discuss best practices and ask for audience participation to try to refine those best practices to help users understand how to leverage parallel file systems successfully.

10C Birds of a Feather

4:45 A Quick-start and HOWTO Guide for CMS Tools, ***Chair: Jason Schildt, Cray Inc.***

The CMS team will lead an open discussion for customers interested in deploying and using CMS Tools in the CLE-2.2 / SMW-4.0 releases. Topics will include tool overview and configuration, followed by CMS Tool capabilities and use cases.

6:00 Cray Inc. Social

The National Archives of Scotland

(See Social Events, p. 20 for details)

Wednesday

11 General Session

8:30 CUG Business: AC Introductions, Election Presentations, Voting, ***Chair: David Gigrich, CUG President, The Boeing Company (BOEING)***

9:00 Invited Talk: Virtual Paleontology: Gait Reconstruction of Extinct Vertebrates Using High Performance Computing, ***Bill Sellers, Manchester University***

Gait reconstruction of extinct animals requires the integration of paleontological information obtained from fossils with biological knowledge of the anatomy, physiology and biomechanics of extant animals. Computer simulation provides a methodology for combining multimodal information to produce concrete predictions that can be evaluated and used to assess the likelihood of competing locomotor hypotheses. However, with the advent of much faster supercomputers, such simulations can also explore a wider range of possibilities, allowing the generation of gait hypotheses de novo.

9:45 CUG Business: Election Results, ***Chair: James Kasdorf, CUG Secretary, Pittsburgh Supercomputing Center (PITTSCC)***

12A

10:30 Petascale Debugging, ***Chris January, David Lecomber, and Mark O'Connor, Allinea Software***

The need for debugging at scale is well known—yet machine sizes have raced ahead of the levels reachable by debuggers for many years. This paper outlines major development of Allinea's DDT debugging tool to introduce production-grade petascale debugging on the Oak Ridge Jaguar XT5 system. The resulting scalable architecture is raising the bar of usability and performance in a debugger by multiple orders of magnitude—and has already achieved record 225,000 core debugging at ORNL.

11:00 Improving the Productivity of Scalable Application Development with TotalView, ***Chris Gottbrath, TotalView Technologies***

Scientists and engineers who set out to solve grand computing challenges need TotalView at their side. The TotalView debugger provides a powerful and scalable tool for analyzing, diagnosing, debugging and troubleshooting a wide variety of different problems that might come up in the process of such achievements. These teams, and teams of scientists pursuing a wide range of computationally complex problems on Cray XT systems are frequently diverse and geographically distributed. These groups work collaboratively on complex applications in a computational environment that they access through a batch resource management system. This talk will explore the productivity challenges faced by scientists and engineers in this environment -- highlighting both long standing (but perhaps unfamiliar) and recently introduced capabilities that TotalView users on Cray can take advantage of to boost their productivity. The list of capabilities will include the CLI, subset attach, Remote Display Client, TVScript, MemoryScape's reporting, and ReplayEngine.

11:30 MRNet: A Scalable Infrastructure for Development of Parallel Tools and Applications, ***Luiz DeRose, Cray Inc.; Barton Miller, University of Wisconsin; and Philip Roth, Oak Ridge National Laboratory (ORNL)***

MRNet is a customizable, high-throughput communication software system for parallel tools and applications. It reduces the cost of these tools' activities by incorporating a tree-based overlay network (TBON) of processes between the tool's front-end and back-ends. MRNet was recently ported and released for Cray XT systems. In this talk we describe the main features that make MRNet well-suited as a general facility for building scalable parallel tools. We present our experiences with MRNet and examples of its use.

12B

10:30 Imperative Recovery for Lustre, ***Nic Henke, Chris Horn, and Cory Spitz, Cray Inc.***

Recovery times for Lustre failover are mainly a function of the overriding bulk data timeout because clients must timeout to a server twice before initiating contact with its backup. As a result, failover completion times exceeding ten minutes are common. During failover and recovery, all IO operations stall and the long duration can lead to job timeouts, poor system utilization, and increased administrator load. To improve overall failover times we are implementing Imperative Recovery, the framework by which Lustre can initiate and finish failover without waiting for long timeouts. Imperative Recovery directs clients to switch server connections based on automatic processing of node health data. With these changes and Version Based Recovery, it is possible to begin recovery very fast, reducing overall failover times to a few minutes. This paper discusses Imperative Recovery from a system perspective and characterizes the speedup achieved.

11:00 File System Monitoring as a Window Into User I/O Requirements, ***Katie Antypas and Andrew Uselton, National Energy Research Scientific Computing Center (NERSC), Daniela Ushizima, Computational Research Division (CRD), Lawrence Berkeley National Lab and Jefferey Sukharev, University of California, Davis***

The effective management of HPC I/O resources requires an understanding of user requirements, so the National Energy Research Scientific Computing center (NERSC) annually surveys its project leads for their anticipated needs. With the advent of detailed monitoring on the Lustre prarallel file system of the Franklin Cray XT it becomes possible to compare actual experience with the expectations presented in the surveys. A correlation of the Lustre Monitoring Tool (LMT) data with job log statistics reveals I/O behavior on a per-project basis. This feedback for both the users and the center enhances NERSC's ability to manage and provision Franklin's I/O subsytem as well as to plan for future I/O requirements.

11:30 High Performance Computing Driven Software Development for Next-Generation Modeling of the World's Oceans, ***Mike Ashworth, Xiaohu Guo, and Andrew Sunderland, STFC Daresbury Laboratory and Gerard Gorman, Stephan Kramer, and Matthew Piggott, Imperial College London***

The Imperial College Ocean Model (ICOM) is an open-source next generation ocean model build upon finite element methods and anisotropic unstructured adaptive meshing. Since 2009, a project has been funded by EPSRC to optimize the ICOM for the UK national HPC service, Hector. Extensive use of profiling tools such as CrayPAT and Vampiz has been made in order to understand performance issues of the code on the Cray XT4. Of particular interest is the scalability of the sparse linear solvers and the algebraic multigrid preconditioners required to solve the system of equations. Scalability of model I/O have been examined and we have implemented a parallel I/O strategy in the code for the Lustre filesystem.

12C

10:30 Things to Consider When Developing and Deploying a Petascale System to Ensure Optimum Scheduling of Resources, ***Trev Harmon, Adaptive Computing***

As Petascale systems are establishing themselves as the standard for high-end Top 500 machines, there are many important considerations in the design and efficient operation to be resolved. In this paper, Adaptive Computing will discuss how scalability issues are being addressed by some of these leading sites and outstanding issues that are on the horizon for future systems.

11:00 High Performance Computing with Clouds: Past, Present, and Future, ***Bill Nitzberg, Altair Engineering, Inc.***

Cloud Computing has its roots in technologies spanning the past 30 years: in network operating systems, distributed systems, metacomputing, clusters, and Grids. The promise of Clouds—-seamless access to computing power and information resources, on-demand-—is delivering substantial benefits in real production settings for datacenter applications. Although, private Clouds are having success in HPC, key characteristics that make public Clouds a compelling solution for the datacenter are orders-of-magnitude larger in HPC. Whether this gap can be crossed in the coming years is an open question; Clouds for HPC are at their infancy.

11:30 Advanced Job Scheduling Features for Cray Systems with Platform LSF, ***William Lu, Platform Computing***

On large Cray systems where all simulation jobs are running through workload management, visibility of the system and jobs are critical for users and administrators to troubleshoot problems. Features in Platform LSF such as scheduling performance, resource reservation and job level data display help simulation users and system administrators easily overcome this challenge. Benchmark data will show how Platform LSF outperforms other workload schedulers. We will also discuss additional technologies from Platform including Platform MPI and its integration with Platform LSF.

12:00 Lunch

13A

1:00 External Services on the Cray XT5 System Hopper, ***Katie Antypas, Tina Butler, and Jonathan Carter, National Energy Research Scientific Computing Center (NERSC)***

Cray External Service offerings such as login nodes, data mover nodes, and file systems which are external to the main XT system, provide an opportunity to make Cray XT High Performance Computing resources more robust and accessible to end users. This paper will discuss our experiences using external services on Hopper, a Cray XT5 system at the National Energy Research Scientific Computing (NERSC) Center. It will describe the motivation for externalizing services, early design decisions, security issues, implementation challenges and production feedback from NERSC users.

1:30 Lessons Learned in Deploying the World's Largest Scale Lustre File System, ***David Dillow, Jason Hill, Dustin Leverman, Don Maxwell, Ross Miller, Sarp Oral, Galen Shipman, James Simmons, and Feiyi Wang, Oak Ridge National Laboratory (ORNL)***

The Spider parallel file system at Oak Ridge National Laboratory's Leadership Computing Facility (OLCF) is the world's largest scale Lustre file system. It has nearly 27,000 file system clients, 10 PB of capacity, and over 240 GB/s of demonstrated I/O bandwidth. In full-scale production for over 6 months, Spider provides a high performance parallel I/O environment to a diverse portfolio of computational resources. These range from the high end, multi-Petaflop Jaguar XT5, the mid-range, 260 Teraflop Jaguar XT4, to the low end, with numerous systems supporting development, visualization, and data analytics. Throughout this period we have had a number of critical design points reinforced while learning a number of lessons on designing, deploying, managing, and using a system of this scale. This paper details our operational experience with the Spider file system, focusing on observed reliability (including MTTI and MTTF), manageability, and system performance under a diverse workload.

2:00 Using I/O Servers to Improve Performance on Cray XT Technology, ***Thomas Edwards and Kevin Roy, Cray Inc.***

Amdhal's Law proposes that parallel codes are combinations of parallel and serial tasks. In many cases these tasks are inherently parallel and can be decomposed and performed asynchronously. Each task operates on a dedicated subset of processors with highly scalable tasks operating on very large numbers of processors and less scalable tasks (like IO) operating on a smaller number. By moving to this Multiple Instruction Multiple Data paradigm codes can achieve greater parallel efficiency and scale further. This paper specifically addresses the implementation and experiences of adapting several codes important to HECToR to offload writing output data onto a set of dedicated server processors.

13B

1:00 Performance on AMD Opteron™ 6000 Series platforms (Socket G34), *David Cownie, AMD*

We will cover the feature set of the new 12-core AMD Opteron™ 6100 Series processor, show key industry performance benchmark results, and share insight into achieving optimal performance with the new processors.

1:30 DataDirect™ Networks' Storage Fusion Architecture, *Josh Goldenhar and Francesco Torricelli, DataDirect Networks, Inc.*

High performance storage systems usually fall into one of two categories: those with high IOPS capability or those with high throughput capability. In the world of supercomputing where the focus is usually on massive scale, the preference (and need) has traditionally favored storage systems with high throughput capabilities. The move to multi-core processors coupled with ever increasing number of nodes in supercomputing clusters has fundamentally changed the data patterns scalable storage must handle. Traditional storage systems cannot handle the need to capable of both high IOPS and high throughput. This paper presents a new storage architecture that can adapt to modern compute environments and the unique data storage challenges they present. Additionally, it outlines how the architecture allows for embedding clustered file systems directly into the storage resulting in reductions in complexity and latency.

2:00 Resiliency Features in the Next Generation Cray Gemini Network, *Forest Godfrey, Cray Inc.*

As system sizes scale to ever increasing numbers of nodes and network links, network failures become an increasingly important problem to address. With its next generation high speed network (code named Gemini), Cray will introduce a number of new resiliency features in this area. These features, including network link failover, are discussed in this paper as well as a comparison to other, more familiar, network technologies such as Ethernet and Infiniband.

13C

1:00 Mixed Mode Computation in CASINO, *Dario Alfe, University College of London and Lucian Anton, Numerical Algorithms Group*

CASINO is a quantum Monte Carlo code that solves many particle Schroedinger equations with the help of configurations of random walkers. This method is suitable for parallel computation because it has a very good computation/communication ratio. The standard parallel algorithm increases the computation speed by distributing equally the configurations among the available processors. For a computation with P processing elements the computation time for Nc configurations is proportional with Nc*tc/P, where tc is the average time taken for one configuration step. On petascale computers one can have more processing elements than configurations and besides that for models with more that 1000 electrons tc increases significantly. We present a mixed mode implementation of CASINO that takes advantage of the architectures with large numbers of multicore processors to improve computation speed by using multiple OpenMP threads for the computation of each configuration step.

1:30 Combining Open MP and MPI within GLOMAP Mode to Take Advantage of Multiple Core Processors: An Example of Legacy Software Keeping Pace with Hardware Developments, *Graham Mann, University of Leeds and Mark Richardson, Numerical Algorithms Group*

The MPI version of GLOMAP MODE is being used in production runs for research into atmospheric science. The memory requirement prohibits use of high resolution scenarios so 32 MPI tasks is the usual decomposition. One way to attempt higher resolution simulations is to under-populate the nodes, making more memory available per MPI task. Although this is wasteful of resources, it does provide a shorter time per existing simulation. The NAG Ltd DCSE service has examined the code and introduced Open MP so that the otherwise "idle" cores can contribute to the MPI task. This improves the performance so that the additional cost of a simulation is reduced.

2:00 A Hybrid MPI/Openmp Code Employing a High-Order Compact Scheme for the Simulation of Hypersonic Aerodynamics, *David Emerson and Vincenzo Fico, STFC Daresbury Laboratory and Jason Reese, University of Strathclyde*

High-order compact schemes are excellent candidates for Direct Numerical Simulation and Large Eddy Simulation of flow fields. We have devised a high-order compact scheme suitable for the simulation of hypersonic flows, to exploit both shared and distributed memory paradigms. Our hybrid application, employing both MPI and OpenMP standards, has been tested on HECToR.

14A

3:00 CUDA Fortran 2003, *Brent Leback, Douglas Miles, and Michael Wolfe, The Portland Group*

In the past year, The Portland Group has brought to market a low-level, explicit, Fortran GPU programming language, a higher-level, implicit, directive-based GPU programming model and implementation, and object-oriented features from the Fortran 2003 standard. Together, these provide a rich environment for programming today and tomorrow's many-core systems. In this paper we will present some of the latest features available in the PGI Fortran compiler from these three areas, and explain how they can be combined to access the performance of CPUs and GPUs while keeping application developers hidden from many of the messy details.

3:30 HPC at NCAR: Past, Present and Future, *Tom Engel, National Center for Atmospheric Research (NCAR)*

The history of high-performance computing at NCAR is reviewed from Control Data Corporation's 3600 through the current IBM p575 cluster, but with special recognition of NCAR's relationship with Seymour Cray and Cray Research, Inc. The recent acquisition of a Cray XT5m is discussed, along with the rationale for that acquisition. NCAR's plans for the new NCAR-Wyoming Supercomputing Center in Cheyenne, Wyoming, and the current status of that construction project, are also described.

14B

3:00 Overview of the Current and Future Cray CX Product Family, ***Ian Miller, Cray Inc.***

Please join us for a detailed product briefing of an exciting new product from Cray. This new product will be a significant enhancement to the Cray portfolio, and will expand the range of capabilities and programming models available to our customers and prospects.

3:30 Validating File System Permissions on Multi-OS Systems, ***Robert Ballance, Sandia National Laboratories (SNLA)***

The Cray XT series of HPC computers presents the system security officer and system administrator with a range of operating systems (Linux, CNL, CVN), job launch (shell/exec, ALPS, yod) and file systems (UFS, NFS, Lustre, LibSysIO, DVS). Available open-source packages do not span this range of requirements. As the system integrator, Cray provides the fundamentals for validating that file system permissions are correctly enforced. However, due to Sandia's security requirements, we were forced to develop a software tool for checking POSIX permission handling across multiple combinations of OS's and file systems. This paper presents the architecture and design of a novel Lisp-based POSIX file system validation tool that uses multi-methods and object-oriented programming to validate tester-specified combinations of access patterns.

14C

3:00 The Graph 500, ***Jim Ang, Brian Barrett, Bruce Henrickson, Richard Murphy, and Arun Rodrigues, Sandia National Laboratories (SNLA)***

New large-scale informatics applications require radically different architectures from those optimizing for 3D Physics. The 3D physics community is represented in the Top 500 list by a LINPACK as a single, simple, dense algebra benchmark. Informally, the Cray XMT performs significantly better than other known architectures on large-scale graph problems, which is a core informatics application kernel. The Graph 500 list, to be introduced at Supercomputing 2010, will formalize a single, unified graph benchmark for the informatics community to rally around and to precipitate innovation in the informatics space. This paper will discuss the need for this kind of benchmark, the benchmark itself, an initial set of results on a small subset of platforms (including XMT), and why those platforms are fundamentally different from other classes of supercomputer.

3:30 A Comparison of Shared Memory Parallel Programming Models, ***Jace Mogil and David Hadin, Pacific Northwest National Laboratory (PNNL)***

The dominant parallel programming models for shared memory computers, Pthreads and OpenMP, are both "thread-centric" in that they are based on explicit management of tasks and enforce data dependencies through task management. By comparison, the Cray XMT programming model is data-centric where the primary concern of the programmer is managing data dependencies, allowing threads to progress in a data flow fashion. The XMT implements this programming model by associating tag bits with each word of memory, affording efficient fine grained synchronization of data independent of the number of processors or how tasks are scheduled. When task management is implicit and synchronization is abundant, efficient, and easy to use, programmers have viable alternatives to traditional thread-centric algorithms. In this paper we compare the amount of available parallelism in a variety of different algorithms and data structures when synchronization does not need to be rationed, as well as identify opportunities for platform and performance portability of the data-centric programming model on multi-core processors.

15A Interactive Session

4:15 Applications & Programming Environments SIG, ***Chair: Rolf Rabenseifner, High Performance Computing Center Stuttgart (HLRS)***

The Applications and Programming Environment SIG welcomes attendees with a focus on compilers and programming environments. Topics include compilers, scientific libraries, programming environments and the Message Passing Toolkit. SIG business will be conducted followed by open discussions with other attendees as well as representatives from Cray. All attendees are welcome to participate in this meeting.

15B Interactive Session

4:15 Systems Support SIG, ***Chair: Joni Virtanen, CSC - IT Center for Science Ltd. (CSC)***

The purpose of this Interactive Session is to provide an open forum for discussion of any and all issues related to the installation, integration, maintenance and operation of major computing resources. Operations managers are familiar with the day to day challenges presented by complex computing systems and the constant goal is to improve all aspects of performance, reliability, ease of use, and maintainability, while providing secure reliable service to a demanding set of users. Listen to what others have tried in this area and contribute your ideas for improvement. Especially helpful are hints at what not to do. This session is open to everyone with an interest in Systems Support.

15C Birds of a Feather

4:15 New Product BoF, ***Ian Miller, Cray Inc.***

5:30 CUG Night Out at The Hub

See Social Events, p. 20 for details.

Thursday

16A

8:30 Analyzing the Effect of Different Programming Models Upon Performance and Memory Usage on Cray XT5 Platforms, ***Alice Koniges, John Shalf, Hongzhang Shan, and Nick Wright, National Energy Research Scientific Computing Center (NERSC); Haoqiang Jin, NAS Systems Division (NAS); and Seung-Jai Min, Lawrence Berkeley National Laboratory***

Harnessing the power of multicore platforms is challenging due to the additional levels of parallelism present. In this paper, we will examine the effect of the choice of programming model upon performance and overall memory usage. We will study how to make efficient use of the memory system and explore the advantages and disadvantage of MPI, OpenMP, and UPC on the Cray XT5 multicore platforms for several synthetic and application benchmarks.

9:00 MPI Queue Characteristics of Large-scale Applications, ***Richard L. Graham and Rainer Keller, Oak Ridge National Laboratory (ORNL)***

Applications running at scale have varying communication characteristics. By employing the PERUSE introspection interface of Open MPI, this paper evaluates several large-scale simulations running production-level input data-sets on the jaguar installation at ORNL. Maximum number of queued messages, average duration of unexpected receives and late sender and receiver information as a function of job size is being presented.

9:30 Performance Analysis of Pure MPI Versus MPI+OpenMP for Jacobi Iteration and a 3D FFT on the Cray XT5, ***Glenn Luecke and Olga Weiss, Iowa State University***

Today many high performance computers are collections of shared memory compute nodes with each compute node having one or more multi-core processors. When writing parallel programs for these machines, one can use pure MPI or various hybrid approaches using MPI and OpenMP. Since OpenMP threads are lighter weight than MPI processes, one would expect that hybrid approaches will achieve better performance and scalability than pure MPI. In practice this is not always the case. This paper investigates the performance and scalability of pure MPI versus hybrid MPI+OpenMP for Jacobi iteration and a 3D FFT on the Cray XT5.

16B

8:30 Improving the Performance of COSMO-CLM, ***Mathew Cordery, Will Sawyer, and Ulrich Schaettler, CSCS-Swiss National Supercomputing Centre (CSCS)***

The COSMO-Model, originally developed by Deutscher Wetterdienst, is a non-hydrostatic regional atmospheric model which can be used for numerical weather prediction and climate simulations and is now in use by a number of weather services for operational forecasting (e.g. MeteoSwiss). One current software engineering goal is to improve its scaling characteristics on multicore architectures by making it a hybrid MPI-OpenMP code. We will present hybridization strategies for different components of the model, show some first performance results, and discuss the impact on further development of the model.

9:00 Thermodynamics of Magnetic Systems from First Principles: WL-LSMS, ***Gregory Brown, Florida State University; Markus Eisenbach and Donald Nicholson, Oak Ridge National Laboratory (ORNL); Jeff Larkin, Cray Inc.; Thomas Schulthess, CSCS-Swiss National Supercomputing Centre (CSCS), and Chengang Zhou, JPMorgan Chase & Co.***

We describe a method to combine classical thermodynamic Monte Carlo calculations (the Wang-Landau method) with a first principles electronic structure calculation, specifically our locally self-consistent multiple scattering (LSMS) code. The combined code shows superb scaling behavior on massively parallel computers and is able to calculate the transition temperature of Fe without external parameters. The code was the recipient of the 2009 Gordon-Bell prize for peak performance.

9:30 Improving the Performance of CP2K on the Cray XT, ***Iain Bethune, EPCC (EPCC)***

CP2K is a freely available and increasingly popular Density Functional Theory code for the simulation of a wide range of systems. It is heavily used on many Cray XT systems, including 'HECToR' in the UK and 'Monte Rosa' in Switzerland. We describe performance optimisations made to the code in several key areas, including 3D Fourier Transforms, and present the implementation of a load balancing scheme for multi-grids. These result in performance gains of around 30% on 256 cores (for a generally representative benchmark) and up to 300% on 1024 cores (for non-homogeneous systems). Early results from the implementation of hybrid MPI/OpenMP parallelism in the code are also presented.

16C

8:30, 9:00, and 9:30 XTreme SIG, ***Chair, James Craw, National Energy Research Scientific Computing Center (NERSC)***

This group works very closely with Cray, under Non-Disclosure Agreements, to provide valuable input into the Cray XT system development cycle. For this reason, these are "closed door" sessions.

17A

10:30 The Evolution of a Petascale Application: Work on CHIMERA, ***John Blondin, North Carolina State University; Stephen Bruenn, Florida Atlantic University; Raph Hix, Bronson Messer, and Anthony Mezzacappa, Oak Ridge National Laboratory (ORNL)***

CHIMERA is a multi-dimensional radiation hydrodynamics code designed to study core-collapse supernovae. We will review several recent enhancements to CHIMERA designed to better exploit features of the CRAY XT architecture, as well as some forward-looking work to take advantage of the next generation of Cray supercomputers.

11:00 What is a 200,000 CPUs Petaflop Computer Good For (a Theoretical Chemist Perspective)?, ***Edoardo Apra and Vinod Tipparaju, Oak Ridge National Laboratory (ORNL); Ryan Olson, Cray Inc.***

We describe the efforts undertaken to efficiently parallelize the computational chemistry code NWChem on the Cray XT hardware using the Global Arrays/ARMCI middleware. We show how we can now use 200K+ processors to address complex scientific problems.

11:30 An Alliance for Computing at the Extreme Scale, ***James Ang and Sudip Dosanjh, Sandia National Laboratories (SNLA); Ken Koch and John Morrison, Los Alamos National Laboratory***

Los Alamos and Sandia National Laboratories have formed a new high performance computing center, the Alliance for Computing at the Extreme Scale (ACES). The two labs will jointly architect, develop, procure and operate capability systems for DOE's Advanced Simulation and Computing Program. This presentation will discuss (1) a petascale production capability system, Cielo, that will be deployed in late 2010, (2) a technology roadmap for exascale computing and (3) a new partnership with Cray on advanced interconnect technologies.

17B

10:30 Automatic Iterative Optimization of Parallel Applications, ***Olli-Pekka Lehto and Sebastian von Alfthan, CSC-Scientific Computing Ltd. (CSC)***

Manual software optimization is effective, but also time-consuming, and can thus benefit from complementary automatic optimization schemes. This paper describes a novel cross-platform framework that is able to optimize parallel applications by tuning three sets of parameters: compiler options, environment variables and internal program parameters. The optimization is carried out using a genetic algorithm, where the trial simulations may be run in parallel allowing the optimization algorithm to scale. The performance of this framework is assessed on a Cray XT5 by optimizing both real world applications, as well as well-known synthetic benchmarks such as the High-Performance Linpack (HPL) benchmark. The results show that our optimization framework increases the performance of the test cases significantly.

11:00 2DECOMP&FFT—A Highly Scalable 2D Decomposition Library and FFT Interface, ***Sylvain Laizet and Ning Li, Numerical Algorithms Group***

As part of a HECToR distributed CSE support project, a general-purpose 2D decomposition (also known as 'pencil' or 'drawer' decomposition) communication library has been developed. This Fortran library provides a powerful and flexible framework to build applications based on 3D Cartesian data structures and spatially implicit numerical schemes (such as compact finite difference method or spectral method). The library also supports shared-memory architecture which becomes increasingly popular. A user-friendly FFT interface has been built on top of the communication library to perform distributed multi-dimensional FFTs. Both the decomposition library and the FFT interface scale well to tens of thousands of cores on Cray XT systems. The library has been applied to Incompact3D, a CFD application performing large-scale Direct Numerical Simulations of turbulence, enabling exciting scientific studies to be conducted.

11:30 A Scalable Boundary Adjusting High-Resolution Technique for Turbulent Flows, ***Sergey Karabasov, University of Cambridge and Phil Ridley, Numerical Algorithms Group***

To accurately resolve turbulent flow structures, high-fidelity simulations require the use of millions of grid points. The Compact Accurately Boundary Adjusting High-Resolution Technique (CABARET) is capable of producing accurate results with at least 10 times more efficiency than conventional schemes. CABARET is based on a local second-order finite difference scheme which lends itself extremely well to large scale distributed systems. For Reynolds numbers of 10^4 the method gives rapid convergence without requiring additional preconditioning for Mach numbers as low as 0.05. In this paper we shall discuss the implementation and performance of the CABARET method on the HECToR XT4/6 system. We shall describe the development and optimization of an irregular parallel decomposition for the hexahedral numerical grid structure. Scalability of the code will be discussed in relation to i) the effectiveness of the load balancing for grids generated from the partitioning method ii) compiler performance and iii) efficient use of MPI and memory utilisation.

17C

10:30, 11:00, and 11:30 XTreme SIG, ***Chair, James Craw, National Energy Research Scientific Computing Center (NERSC)***

This group works very closely with Cray, under Non-Disclosure Agreements, to provide valuable input into the Cray XT system development cycle. For this reason, these are "closed door" sessions.

12:00 Lunch

18A

1:00 Parallelism in System Tools, ***Kenneth Matney, Sr. and Galen Shipman, Oak Ridge National Laboratory (ORNL)***

The Cray XT, when employed in conjunction with the Lustre filesystem, provides the ability to generate huge amounts of data in the form of many files. This is accommodated by satisfying the requests of multiple Lustre clients in parallel. In contrast, a single service node (Lustre client) cannot provide timely management for such datasets. Consequently, as the dataset enters the 10+ TB range and/or hundreds of thousands of files, using traditional UNIX tools like cp, tar, or find . exec ... ; to manage these datasets causes the impact to user productivity to become substantial. For example, it would take about 12 hours to copy a 10 TB dataset from the service node via cp if dedicated resources were employed. In general, it is not practical to schedule dedicated resources for a data copy and, as a result, a typical duty factor of 4X is incurred. This means that, in practice, it would take 48 hours to perform a serial copy of a 10 TB dataset. Over the next three to four years, datasets are likely to grow by a factor of 4X. At that point, the simple copy of a dataset may be expected to take over a week and represents significant impediment to the investigation of science. In this paper, we introduce the Lustre User Toolkit for Cray XT, developed at the Oak Ridge National Laboratory Leadership Computing Facility (OLCF) and demonstrate that, by optimizing and parallelizing system tools, an order of magnitude performance increase or more can be achieved, thereby reducing or eliminating the bottleneck. The conclusion is self-evident: parallelism in system tools is vital to managing large datasets.

1:30 Regression Testing on Petaflop Computational Resources, ***Troy Baer, Lonnie Crosby, and Mike McCarty, National Institute for Computational Sciences (NICS)***

As the complexity of supercomputers increases, it is becoming more difficult to measure how system performance changes over time. Routine system checks performed after scheduled maintenance or emergency downtime give administrators an instantaneous glimpse of system performance; however, rigorous testing, such as that performed for machine acceptance, provides more in-depth information on system performance. Both routine and rigorous testing are necessary to fully characterize system performance, and a mechanism to store and compare previous results is needed to determine the change in system performance over time. A regression testing framework has been developed at the National Institute for Computational Sciences (NICS) which provides a mechanism to measure the change in system performance over time. These performance

results can also be correlated to system events such as downtimes, system upgrades, or any other documented system change. We will describe the design and implementation of the regression testing framework, including the development of test suites, interfaces to the batch system, and the extraction of performance data. The import of extracted data into a relational database for long- term storage, report generation, and real- time analysis will also be discussed.

2:00 Reducing Application Runtime Variability on Jaguar XT5, ***Jeff Becklehimer and Jeff Larkin, Cray Inc.; Dave Dillow, Don Maxwell, Ross Miller, Sarp Oral, Galen Shipman, and Feiyi Wang, Oak Ridge National Laboratory (ORNL)***

Operating system (OS) noise is defined as interference generated by the OS that prevents the compute core from performing useful work. Compute node kernel daemons, network interfaces, and other OS related services are major sources of such interference. This interference on individual compute cores can vary in duration and frequency and can cause de-synchronization (jitter) in collective communication tasks and thus results in variable (degraded) overall parallel application performance. This behavior is more observable in large-scale applications using certain types of collective communication primitives, such as MPI_Allreduce. This paper presents our efforts towards reducing the overall effect of OS noise on our large-scale parallel applications. Our tests were performed on the quad-core Jaguar, the Cray XT5 at the Oak Ridge National Laboratory Leadership Computing Facility (OLCF). At the time of these tests, Jaguar was a 1.4 PFLOPS supercomputer with 144,000 compute cores and 8 cores per node. The technique we used was to aggregate and merge all OS noise sources onto a single compute core for each node. The scientific application was then run on the remaining seven cores in each node. Our results show that we were able to improve the MPI_Allreduce performance by two orders of magnitude and to boost the Parallel Ocean Program (POP) performance over 30% using this technique.

18B

1:00 Running Hadoop on a Cray XT System, ***Shane Canon, Lavanya Ramakrishnan, and John Shalf, National Energy Research Scientific Computing Center (NERSC) and Keith Jackson, Lawrence Berkeley National Lab***

Hadoop is an open source implementation of the MapReduce programming model popularized by Google. Hadoop has been heavily adopted in the Web 2.0 community and is now making inroads in the scientific and research communities. The flexibility of the MapReduce programming model combined with the power of the Cray XT can impact the size and nature of scientific explorations. In this paper we will explain the motivations for using Hadoop and describe the steps to deploy the framework on a Cray XT. We will examine some of the configuration options and their impact on performance. We will compare the performance of several applications running in Hadoop on the Cray system with the performance on standard Hadoop deployments on clusters and Cloud systems. We will conclude with some assessment on the feasibility and efficacy of running Hadoop on HPC systems and future work.

1:30 XGC1: Performance on the 8-core and 12-core Cray XT5 Systems at ORNL, ***Mark Adams, Columbia University; C-S Chang and Seung-Hoe Ku, New York University; Eduardo D'Azevedo, Collin McCurdy, and Patrick Worley, Oak Ridge National Laboratory (ORNL)***

The XGC1 code is used to model multiscale tokamak plasma turbulence dynamics in realistic edge geometry. In June 2009, XGC1 demonstrated nearly linear weak and strong scaling out to 150,000 cores on a Cray XT5 with 8-core nodes when solving problems of relevance to running experiments on the ITER tokamak. Here we compare performance, and discuss further performance optimizations, when running XGC1 on an XT5 with 12-core nodes on up to 224,000 cores.

2:00 Multi-Core Aware Performance Optimization of Halo Exchanges in Ocean Simulations, ***Stephen Pickles, STFC Daresbury Laboratory***

The advent of multi-core brings new opportunities for performance optimization in MPI codes. For example, the cost of performing a halo exchange in a finite-difference simulation can be reduced by choosing a partition into sub-domains that takes full advantage of the faster shared-memory mechanisms available for MPI communication between tasks on the same node. We have implemented these ideas in the Proudman Oceanographic Laboratory Coastal-Ocean Modelling System, and find that multi-core aware optimizations can offer significant performance benefit, especially on hex-core systems.

18C

1:00, 1:30, and 2:00 XTreme SIG, ***Chair, James Craw, National Energy Research Scientific Computing Center (NERSC)***

This group works very closely with Cray, under Non-Disclosure Agreements, to provide valuable input into the Cray XT system development cycle. For this reason, these are "closed door" sessions.

19 General Session

2:45 LAC Appreciation

3:00 Next CUG–Fairbanks, *Barbara Horner-Miller, Arctic Region Supercomputing Center (ARSC)*

3:15 Adjourn

Local Arrangements

How to Contact Us

Until May 21, 2010, and after the conference

Alan Gray (EPCC)
Local Arrangements Chair
EPCC
2404 James Clerk Maxwell Building
Mayfield Road
Edinburgh EH9 3JZ UK
(44-131) 650 5201 Fax: (44-131) 650 6555
cug2010@epcc.ed.ac.uk

During the conference, May 21–27, 2010
CUG 2010 Conference
Apex Waterloo Place Hotel
23-27 Waterloo Place
Edinburgh EH1 3BH UK
(44-131) 523 1819

Conference Registration
CUG Office
2601 NE Jack London, #104
Corvallis, Oregon 97330 USA
(1-541) 758-9000 Fax: (1-971) 285-9052
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.

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 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 Office and Registration

Registration on Sunday and Monday is in the hotel lobby. Tuesday–Thursday registration is in the CUG Office.

Badges and registration materials are available:

Sunday.....*3:00 p.m.–5:00 p.m.*
Monday.....*7:0 a.m.–5:30 p.m.*
Tuesday–Wednesday.....*8:00 a.m.–5:30 p.m.*
Thursday.....*8:00 a.m.–3:00 p.m.*

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 including dinner that evening

Proceedings

All conference attendees will have access to the conference Proceedings on our secure web server where 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.

Dining Services

Lunch and Refreshments

The conference registration fee includes lunch and refreshment breaks for the duration of the conference which will all be served in Elliot's Restaurant.

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 available in the Conference Office. (See Social Events, p. 20 for details)

Dining Out

There are many options for dining in Edinburgh near the conference venue. A comprehensive guide, including reviews, is available at http://www.list.co.uk/food-and-drink/

On-Site Facilities

Messages

Call the hotel at (44-131) 523 1819 to leave telephone messages for attendees and ask that the message be either delivered to your room or to the conference office during the conference (May 24–27).

Faxes

There are no fax facilities within the conference, but if required the Conference Office staff will be able to help with the sending or receiving of faxes using external University of Edinburgh facilities.

Local Arrangements

Wireless Access, E-mail, and Power

There is wireless internet access throughout the hotel facility for computers equipped with an 802.11b/g-compatible card. 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 want to plug in, sit near the front!

E-mail and Personal Computers

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

Speaker Preparation

We invite speakers to use the personal computers and the printer 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. Please come to the Conference Office if you need assistance.

Voltage

The United Kindom operates on a 230V system and the sockets are UK specific with 3 pins. Those traveling from outside the UK will need an adaptor. Check the Voltage Valet at voltagevalet.com/idx.html for information on requirements and to find adapters.

Travel Information

Edinburgh Airport

A taxi to the airport should cost about £15–20 one way plus tip and should take approximately 20 minutes, depending on the time of day and traffic.

Airport Shuttle Services

Airlink operates an express shuttle service between the airport and Edinburgh city centre. A single journey costs £3.50, a return journey (round trip) is £6. The bus terminates at Edinburgh Waverley station, which is a very short walk to the Apex Waterloo Place Hotel.

Ground Transportation

Edinburgh has excellent integrated transport – the city's bus services are operated by two companies: Lothian Buses and First. Lothian Buses provide the main service in the city centre. A single journey costs £1.20 and an unlimited travel day pass costs £3. Please note that you need to have the exact fare as bus drivers are not able to give change!

Car hire is not recommended in Edinburgh, and the conference hotel does not have its own car park, but it does offer a discount at a nearby facility: for more info see http://www.apexhotels.co.uk/hotels/edinburgh-waterloo-place/edinburgh-map/#car

Currency

Foreign currency can be exchanged at banks near the hotel. Cash machines (ATMs) are available at various points in the airport terminal building, before and after security control. All cash dispensers will accept valid UK bank and building society cards (e.g., Link affiliated) without surcharge. They also accept major credit cards. Go to the Universal Currency Converter™ web site located at www.xe.net/ucc/ for help with foreign exchange rates.

Weather

May is a beautiful time of year in Edinburgh. May and June are generally the best bet for dry, sunny weather with average temperatures between 6-16° C (61-43° F), but you can expect rain at any time.

Tourist Information and Attractions

Tourist information is available from the CUG 2010 Local Arrangements pages on cug.org. Brochures and additional information are available at the hotel.

Edinburgh is famous for its spectacular city landscape and world-class attractions:

- Edinburgh Castle (http://www.edinburghcastle.gov.uk/)
- Palace of Holyroodhouse (http://www.royalcollection.org.uk/default.asp?action=article&ID=36)
- The Royal Mile (http://www.edinburgh-royalmile.com)
- The Royal Yacht Britannia (http://www.royalyachtbritannia.co.uk/)
- Scott's Monument (http://www.edinburgh.gov.uk/internet/Leisure/Museums_and_galleries/Monuments/Scott_monument/CEC_the_scott_monument)
- National Galleries of Scotland (http://www.nationalgalleries.org/) & National Museums of Scotland (http://www.nms.ac.uk/)
- Royal Botanic Gardens (http://www.rbge.org.uk/http://www.rbge.org.uk/)

The Edinburgh Pass (http://www.edinburgh.org/pass) gives free entry to top attractions, including the Camera Obscura, Our Dynamic Earth or Edinburgh Zoo.

For spectacular 360° views over Edinburgh climb to the top of Arthur's Seat (http://en.wikipedia.org/wiki/Arthur%27s_Seat,_Edinburgh), an extinct volcano in Holyrood Park in the heart of the city!

Apex Hotel, Edinburgh

Social Events

Drinks Reception

A short complimentary drinks reception to welcome CUG 2010 attendees to Edinburgh. Elliot's Bar (within the conference hotel, next to the reception desk). Monday 6:00 PM.

Cray Social

Cray Inc. will host a social on Tuesday evening at 6:00 PM at the The National Archives of Scotland, Adam Dome and Adam Dome Gallery, 2 Princes Street, Edinburgh. All conference participants and their guests are invited to attend.

CUG Night Out

The CUG Night Out on the Wednesday, provided for all attendees who have registered for the full conference, includes a 1 hour open-top bus tour of Edinburgh and Dinner at the impressive Hub building.

The Hub, on the Royal Mile close to Edinburgh Castle, was built in the middle of the 19th century as the Assembly Hall for the Church of Scotland. In 1995 the building was acquired by the Edinburgh International Festival and has since then been used for events such as weddings, banquets and conferences.

Tour buses will leave the conference venue at 6pm and will arrive at the Hub at 7pm. Note that there is no transport after the event: participants should enjoy the 15 minute walk back to the hotel which is not only very scenic but passes several of the many pubs in Edinburgh (those requiring special assistance should contact the conference office). Tickets for attendees' guests are available to purchase in the Conference Office. This special event is organized by the Local Arrangements Committee.

6pm Open-top bus tour with live guide - leaving from outside the Apex Hotel.

7pm Arrival at The Hub for dinner

Contacts

Special Interest Groups

Special Interest Groups (SIGs) are organized by subject areas rather than by platform. The Applications and Programming Environment SIG deals naturally with Applications and Programming Environments issues. The Systems Support SIG covers topics related to Operations, Systems & Integration, and Operating Systems. The User Support SIG deals with issues related to User Services. This year, a new XTreme SIG works very closely with Cray, under Non-Disclosure Agreements, to provide valuable input into the Cray XT system development cycle.

User Services SIG, Chair: Jim Glidewell (BOEING)

Deputy Chair: Helen He (NERSC)

Cray Inc. SIG Liaison
Nola Van Vugt

Applications and Programming Environments SIG, Chair: Rolf Rabenseifner (HLRS)

Deputy Chair: Neil Stringfellow (CSCS)

Cray Inc. SIG Liaisons
Applications: Jef Dawson
Programming Environments: Luiz DeRose

Systems Support SIG, Chair: Joni Virtanen (CSC)

Deputy Chair: Jim Glidewell (BOEING)

Cray Inc. SIG Liaisons
Systems & Integration: Jim Harrell
Operating Systems: Charlie Carroll
Operations: Charlie Clark

XTreme SIG, Chair: Jim Craw (NERSC)

Deputy Chair: Neil Stringfellow (CSCS)

Cray Inc. SIG Liaison
Vito Bongiorno

Contacts

CUG Board of Directors

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

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

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
James Kasdorf
Pittsburgh Supercomputing Center
300 South Craig Street #304
Pittsburgh, PA 15213 USA
(1-412) 268-4960
kasdorf@psc.edu

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
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

Director-at-Large
Claire Hepwood
Atomic Weapons Establishment
E2.1, AWE Aldermaston
Reading, R40 4PR United Kingdom
(44-118) 982 6886
Claire.hepwood@awe.co.uk

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

Local Arrangements Chair

Alan Gray
EPCC
2404 James Clerk Maxwell Building
Mayfield Road
Edinburgh EH9 3JZ United Kingdom
(44-131) 650 5201 Fax: (44-131) 650 6555
cug2010@epcc.ed.ac.uk

Advisory Council Secretary and CUG.log Editor

Claire Hepwood
Atomic Weapons Establishment
E2.1, AWE Aldermaston
Reading, R40 4PR United Kingdom
(44-118) 982 6886
Claire.hepwood@awe.co.uk

Cray Inc. Representatives to CUG

Vito Bongiorno, Business Liaison
Cray Inc.
380 Jackson Street, Suite 210
Saint Paul, MN 55101–2907 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 Fax: (1-715) 726-4739
sscott@cray.com

Ellen Roder, Program Liaison
Cray Inc.
380 Jackson Street, Suite 210
Saint Paul, MN 55101–2907 USA
(1-651) 605-9058 Fax: (1-651) 605-9001
ellen@cray.com

Linda Yetzer, Arrangements Liaison
Cray Inc.
380 Jackson Street, Suite 210
Saint Paul, MN 55101–2907 USA
(1-651) 605-8961 Fax: (1-651) 605-9123
yetzerll@cray.com

Contacts

SIG Chairs and Cray Inc. SIG Liaisons

Vito Bongiorno
Cray Inc.
380 Jackson Street, Suite 210
Saint Paul, MN 55101–2907 USA
(1-651) 605-8935 Fax: (1-651) 605-9123
vb@cray.com

Charlie Carroll
Cray Inc.
380 Jackson Street, Suite 210
Saint Paul, MN 55101–2907 USA
(1-651) 605-8987 Fax: (1-651) 605-9001
charliec@cray.com

Charlie Clark
Cray Inc.
1050 Lowater Road
P.O. Box 6000
Chippewa Falls, WI 54729 USA
(1-715) 726-4020 Fax: (1-715) 726-4739
cmc@cray.com

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

Jef Dawson
Cray Inc.
380 Jackson Street, Suite 210
Saint Paul, MN 55101–2907 USA
(1-651) 605-9036 Fax: (1-951) 605-9123
jef@cray.com

Luiz DeRose
Cray Inc.
380 Jackson Street, Suite 210
Saint Paul, MN 55101–2907 USA
(1-651) 605-8925 Fax: (1-651) 605-9001
ldr@cray.com

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

Jim Harrell
Cray Inc.
380 Jackson Street, Suite 210
Saint Paul, MN 55101–2907 USA
(1-651) 605-9131 Fax: (1-651) 605-9001
ejh@cray.com

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

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

Neil Stringfellow
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

Nola Van Vugt
Cray Inc.
901 Fifth Avenue, Suite 1000
Seattle, WA 98164 USA
(1-206) 701-2239 Fax: (1-206) 701-2267
nolav@cray.com

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

Duddingston Loch, Edinburgh

CUG Office

Bob and Karen Winget
Toolbox Conference Management
2601 NE Jack London, #104
Corvallis, OR 97330 USA
(1-541) 758-9000 Fax: (1-971) 285-9052
bobwinget@toolboxcm.com kwinget@fpes.com

Our Sponsors

CUG 2010 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.

- Cray Inc. for sponsoring the Cray Social
- AMD for sponsoring our CUG Night Out
- DataDirect Networks and Platform computing for their sponsorship of the lunches
- Allinea, Altair, LSI, PGI, and TotalView for their sponsorship of the coffee breaks

Finally, I would like to thank Cray Inc. for their assistance and support throughout the entire planning process as well as the CUG Board for their guidance and patience.

Alan Gray, Local Arrangements Chair

Our Host

epcc

Diamond

CRAY
THE SUPERCOMPUTER COMPANY

The Cray User Group is especially grateful to Cray Inc. for their ongoing support as we celebrate 32 years of CUG conferences and workshops.

Platinum

AMD

Gold

DataDirect NETWORKS
Performance. Capacity. Innovation.

Platform™

Silver

allinea
SCALE TO NEW HEIGHTS

Altair

LSI

PGI®

TOTALVIEW
TECHNOLOGIES

Bronze

Adaptive COMPUTING

ORACLE

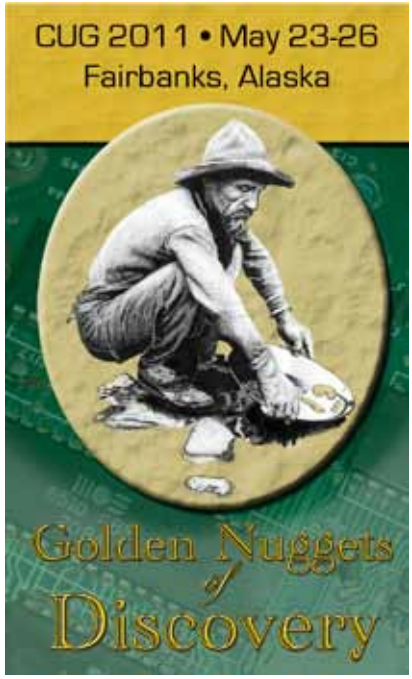

Dear Friends,

The Arctic Region Supercomputing Center (ARSC) at the University of Alaska Fairbanks (UAF) invites you to attend CUG 2011. The conference theme, Golden Nuggets of Discovery, reflects upon the discovery of gold that led to the founding of Fairbanks and the scientific discoveries made possible through high performance computing. ARSC will host the meeting at the downtown Westmark Hotel.

ARSC was opened in 1993. The center, located at two geographical locations, supports high performance computational research in science and engineering with emphasis on high latitudes and the Arctic. Computations are performed by scientists within the High Performance Computing Modernization Program of the DoD for which ARSC serves as a DoD Supercomputing Resource Center and by researchers, faculty and students at the University of Alaska Fairbanks.. UAF is America's Arctic University and home to more than a dozen research centers and institutes that focus on the Arctic.

ARSC's newest system, Chugach—named after the mountain range located in Southcentral Alaska—is a Cray XE7 comprising 11,520 cores. This system will be installed at ARSC's location within the National Petascale Computing Facility that is part of NCSA in Illinois and operated remotely from Fairbanks.

Residents of the Golden Heart City are eager to share their city and local attractions, including gold mining adventures, museums, canoeing, kayaking, hot springs powered by geothermal energy, wildlife, music, theatre and more. Day trips for fishing, visiting Mount MCKinley and the Denali National Park, sightseeing on the Alaska Railroad and flight-seeing above the Arctic Circle are accessible from Fairbanks. In late May we have 21 hours of daylight to play in!

We hope that you will join us in Fairbanks for what we expect to be a thoroughly educational, enlightening, and enjoyable CUG 2011.

Sincerely,

Barbara Horner-Miller
Local Arrangements Chair
horner@arsc.edu

Cynthia Steiner
Local Arrangements Coordinator
steiner@arsc.edu

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 Fairbanks. 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 2011 the best yet. Hope to see you in Fairbanks!

Meeting Room Map

Apex Waterloo Hotel

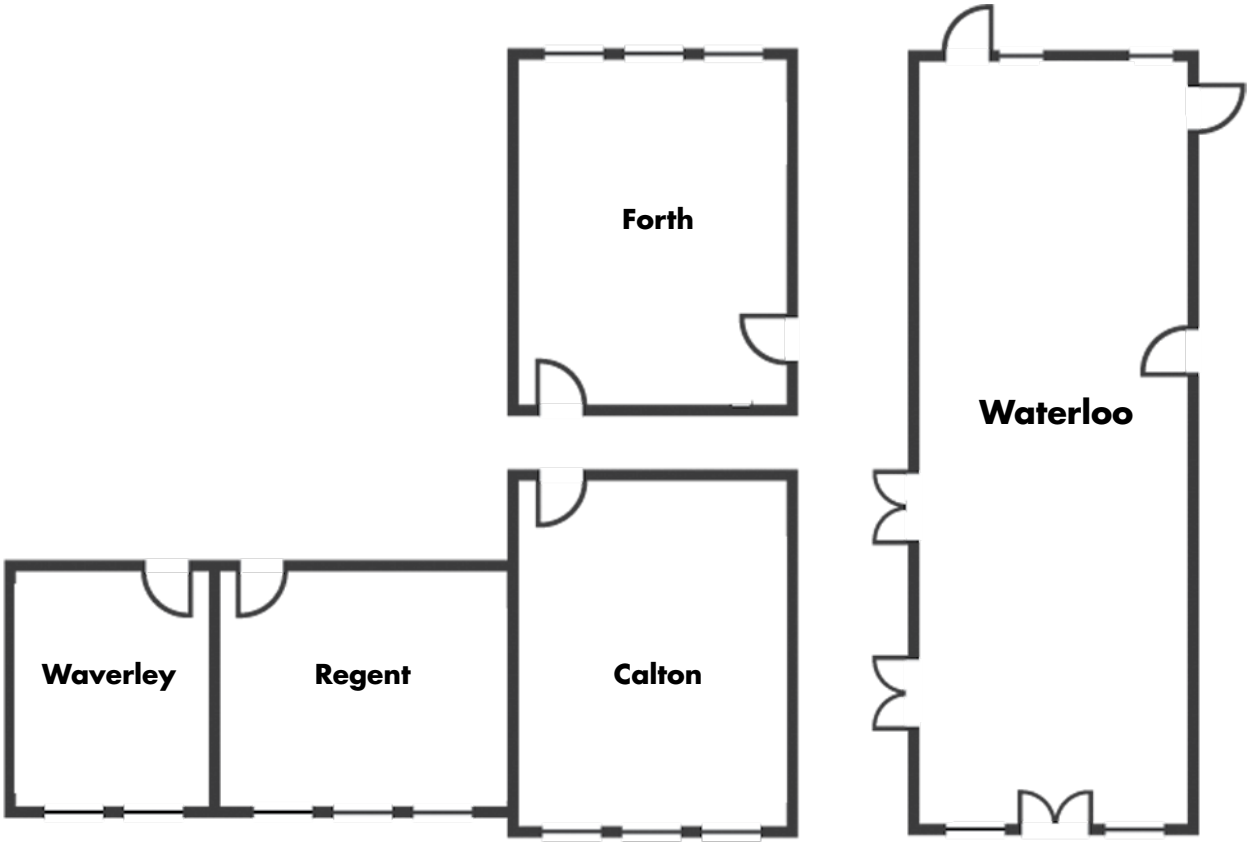

CUG 2010 Meeting Rooms

Event/Facility	Room
General Sessions	Waterloo*
Track A	Waterloo
Track B	Forth
Track C	Calton
CUG Conference Office	Waverley
E-mail/Speaker Prep	Across from Waverley
Conference Registration	
Sunday–Monday	Hotel Lobby
Tuesday–Thursday	Waverley
Lunch and Breaks	Elliot's Restaurant
Board Room	Regent

* Note: The Waterloo room is on the floor below the other rooms.

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

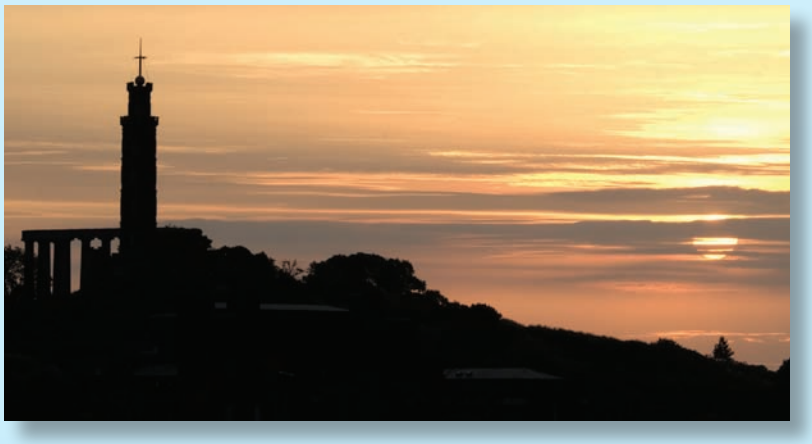

CUG2010
EDINBURGH
SIMULATION
COMES OF AGE