



**PDC Center for
High Performance Computing**

no:2
2012

PDC Newsletter

SNIC: Opening the Gate for Swedish HPC Centres to PRACE by Lilit Axner, PDC, page 4

DEISA Evolves into PRACE by Lilit Axner, PDC, page 6

The Scientific Case for the European High-End HPC Infrastructure by Lilit Axner, PDC, page 7

Application Scaling Activities within PRACE by Lilit Axner, PDC, page 8

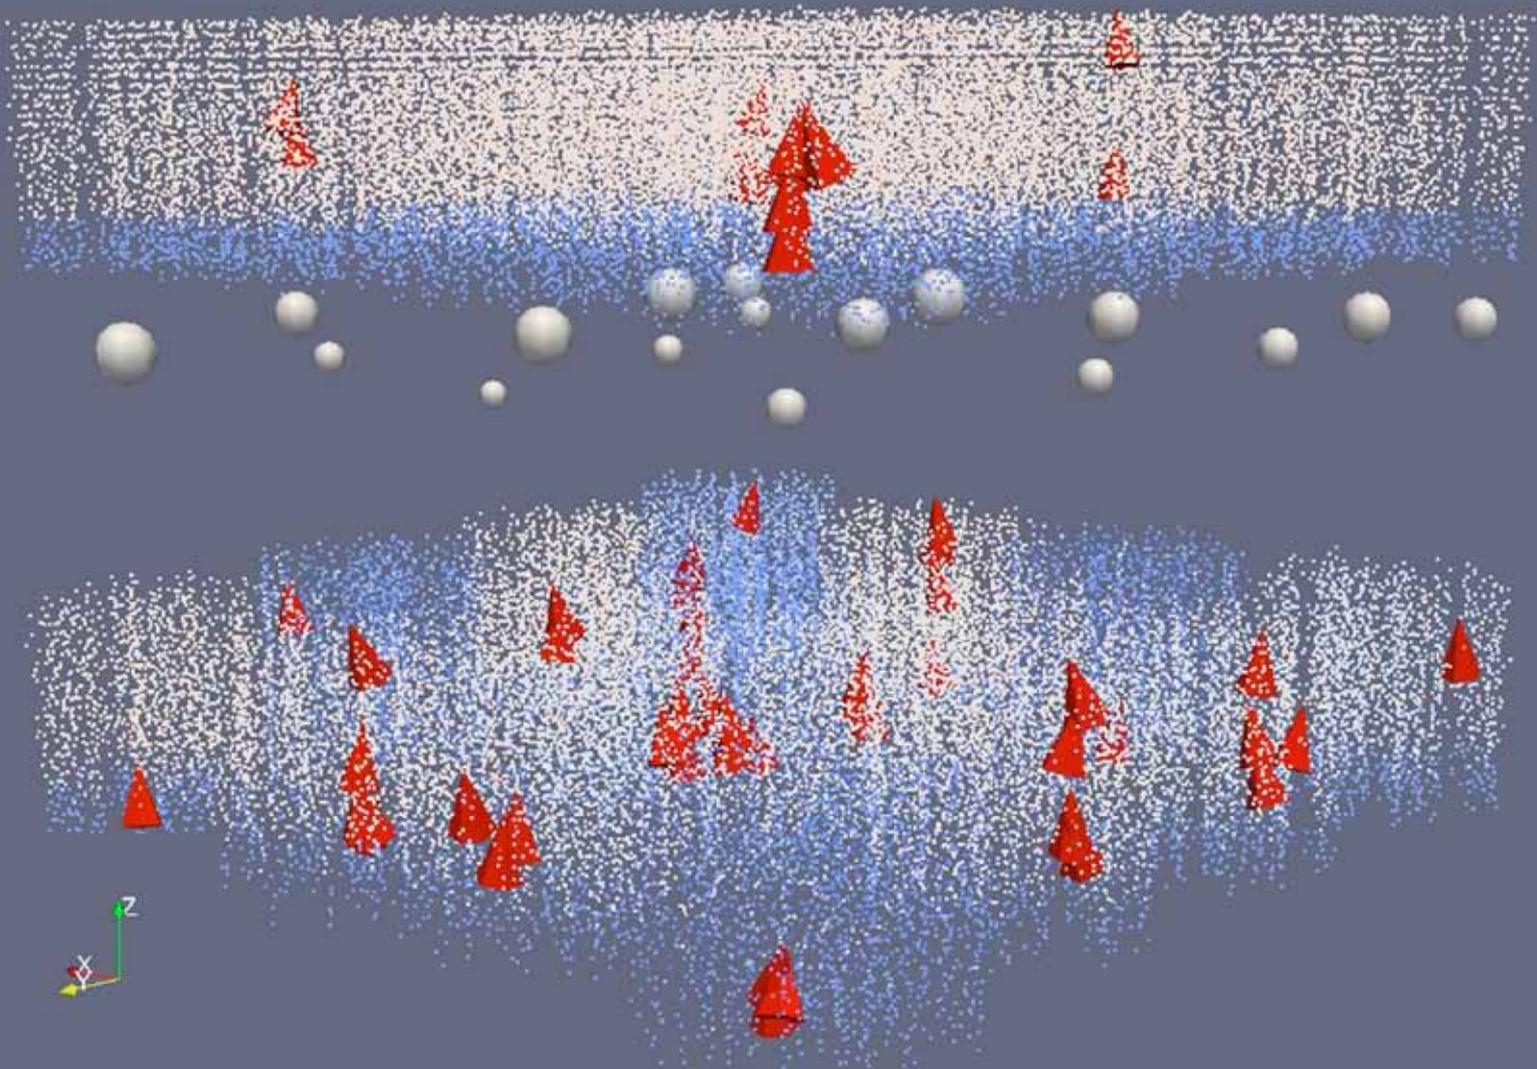
PDC's and PRACE's Commitment to Energy-Efficient Computation

by Gert Svensson and Gilbert Netzer, PDC, page 11

Zorn - New GPU Cluster Available for Research by Michael Schliephake, PDC, page 13

CRESTA Collaboration Meeting at KTH by Stefano Markidis, PDC/HPCViz, page 14

Introduction to High Performance Computing by Stefano Markidis, PDC/HPCViz, page 15



HOW TO CONTACT PDC

Visiting address:
Teknikringen 14, "Plan 4",
KTH, Stockholm

Mail: PDC, KTH,
SE-10044, Stockholm, Sweden

E-mail: support@pdc.kth.se
www: <http://www.pdc.kth.se>
Phone: +4687907800
Fax: +468247784

PUBLISHED BY PDC AT CSC, KTH

The center operates leading-edge, high-performance computers as easily accessible national resources. These resources are primarily available for Swedish academic research and education. PDC, hosted by KTH, is one of the six centers in the Swedish National Infrastructure for Computing (SNIC).

Editor:

Erwin Laure

Editorial staff:

Genet Edmondson

Layout:

Emma Bayne & Maria Malmqvist

Graphic Design:

Maria Malmqvist

E-mail: pdc-newsletter@pdc.kth.se

ISSN 1401-9671

Cover: Neural activity of two simulated neocortical areas that are modular in terms of hypercolumns with internal minicolumns. The individual dots are neurons, which become larger when firing. The red pyramids are active pyramidal neurons and the rising spheres are the aggregated spiking activity in projections between minicolumns of each area.



ERWIN LAURE,
DIRECTOR PDC

In September 2012, the third implementation phase of the PRACE project (PRACE 3IP) started with a kick-off meeting in Paris. Collaboration within Europe has always been high on the SNIC agenda to ensure that Swedish scientists have access to the best resources, wherever they might be located. Consequently, Sweden has been actively engaged in the PRACE projects from their beginnings in 2007. Thanks to the contributions of the SNIC centres, Swedish scientists have been able to benefit significantly from these European efforts. In this edition of the PDC newsletter, we look back at the past five years of PRACE and highlight the achievements and impact of PRACE on the Swedish scientific computing landscape. PRACE has not only provided resources at a level that would not have been possible in Sweden alone, but has also given valuable impetus in improving the scalability and efficiency of many applications that are important for Swedish scientists. And, as new supercomputers are using more and more energy, the research work on energy-efficient hardware that has been initiated by PRACE will be of great value in the future. The Swedish PRACE efforts were being coordinated by PDC's Lilit Axner, who went on maternity leave in September. Michaela Barth (formerly Lechner), who many of you will know already from her involvement in the European Grid Infrastructure (EGI), is taking over and will ensure a successful and productive continuation of the Swedish PRACE efforts.

Over the summer, the new SNIC organization has taken shape and the new director, Jacko Koster, together with the SNIC board, are working on the new strategic directions for SNIC. This is a user-driven process that includes surveying users and updating the scientific case. Many of you will probably be involved in this process over the coming months. I am confident that with the new SNIC organization we will be able to continue providing the best possible e-Infrastructure resources for Swedish scientists.

It is not only SNIC that is changing its organization; KTH has also undergone some changes with the establishment of a new department at the School for Computer Science and Communications (CSC) that is dedicated to

High Performance Computing and Visualization (HPCViz). This new department, led by PDC's director, Erwin Laure, combines the HPC research activities at PDC, the modelling and algorithmic work at the Computational Technology lab, and the visualization research and infrastructure of the KTH Visualization Studio (VIC). This new department will add momentum to the HPC activities at KTH and provide excellent synergies with the infrastructure work at PDC.

In August, PDC ran its flagship training event, the PDC Summer School "Introduction to High Performance Computing", where over 50 students from Sweden and abroad were given a kick-start into HPC by leading Swedish and international HPC experts.

PDC's international activities are also continuing and a highlight was the hosting of the CRESTA exascale project all-hands meeting in September. Improving the scalability and efficiency of HPC software is a key activity at PDC, which is achieved through the work of our SNIC application experts, and our involvement in SeRC and international projects (such as ScalaLife and CRESTA). Increasingly, the importance of software is also being recognized at the EC level, and HPC software is expected to be a priority area in the upcoming Horizon 2020 program.

The PDC staff has also seen some changes over summer. We enlarged our application expert team by hiring David Silverstein as an application expert in neuroinformatics, and a bioinformatics expert will be appointed during the autumn. With these valuable additions, PDC is able to provide expert support in the areas of bioinformatics, computational chemistry, fluid dynamics, molecular dynamics, and neuroinformatics. We also reinforced our team of system administrators with two new members: Carl Johan Håkansson and a part-time student, Johannes Bulin. Erik Järleberg, another part-time student, has also joined our first line support team recently. You will find short introductions to all of them in this issue. In June we had to bid "farewell" to Elisabet Molin, who has been working in PDC support for many years and leading the support group since 2011. Elisabet is in pursuit of new adventures elsewhere, and we wish her all the best for her future. PDC's support group is now led by Jonathan Vincent, one of PDC's best parallelization experts.

With this, I wish you a productive autumn on the PDC resources, and look forward to exciting HPC developments in 2013.

Erwin Laure, Director PDC

In This Issue:

• PDC EDITORIAL

page 2

• SNIC: OPENING THE GATE FOR SWEDISH HPC CENTRES TO PRACE

page 4

• DEISA EVOLVES INTO PRACE

page 6

• THE SCIENTIFIC CASE FOR THE EUROPEAN HIGH-END HPC INFRASTRUCTURE

page 7

• STAFF FOCUS

page 4, 5, 14

• APPLICATION SCALING ACTIVITIES WITHIN PRACE

page 8

• PDC'S AND PRACE'S COMMITMENT TO ENERGY-EFFICIENT COMPUTATION

page 11

• 2011 PRACE EVENTS WITH SNIC-PDC

page 12

• ZORN - NEW GPU CLUSTER AVAILABLE FOR RESEARCH

page 13

• CRESTA COLLABORATION MEETING AT KTH

page 14

• INTRODUCTION TO HIGH PERFORMANCE COMPUTING - PDC SUMMER SCHOOL 2012

page 15

• PDC RELATED EVENTS & SOURCES

page 16

Staff Focus



JOHANNES BULIN

Johannes Bulin joined PDC at the beginning of June this year. His current job is to get the storage software iRODS running on the PDC systems – when he is not busy studying scientific computing at KTH.

Before Johannes came to Sweden last year, he studied applied mathematics in Trier in the western part of Germany. When he is not working – and has no homework to do – Johannes enjoys cycling, reading and cooking. He also tries to improve his Swedish, but needs some more training!

Staff Focus



DAVID SILVERSTEIN

David Silverstein joined PDC in September 2012 as an applications expert in neuroinformatics. He received an M.Sc. in Complex Adaptive Systems from Chalmers University of Technology,



PRACE Partnership for Advanced Computing in Europe

SNIC: Opening the Gate for Swedish HPC Centres to PRACE

by Lilit Axner, PDC

Since SNIC was formed, the meta-centre has taken an active role in European projects in high performance and grid computing. During autumn 2006, SNIC participated (together with ten other European countries) in the High Performance Computing in Europe Taskforce (HET). Its mandate was to draft a strategy for the European HPC ecosystem. The HET strategy, presented in January 2007, formed the basis for the PRACE initiative, which is the only e-Infrastructure activity included in the ESFRI roadmap. The aim of PRACE is to give European researchers access to a rich portfolio of application software and computer systems, including the most powerful systems on a global scale. In April 2007, 16 European partners, including SNIC, signed the PRACE Memorandum of Understanding (MoU) with the goal of realizing the HPC initiative described in the ESFRI roadmap, namely facilitating world-class scientific research.

During 2007, the PRACE partners initiated the EU FP7 PRACE project, which is currently preparing the creation of the persistent PRACE pan-European HPC service. The PRACE Preparatory Phase ran over 24 months, and finished at the end of 2009. It had a total budget of 20 million EUR. The Swedish Research Council's Committee for Research Infrastructures (VR/KFI) made additional funding available to SNIC during 2007: 6.6 million SEK were contributed to a PRACE prototype system. This extra funding, together with the funding for staff allocated from the baseline SNIC budget, meant that SNIC occupied a prominent position in the project, especially in relation to Sweden's small population on a European scale. Thus SNIC was in an excellent position to cater for the increased needs for capability computing within Swedish research groups. Consequently SNIC entered the PRACE implementation phase at the end of 2009 as a high-profile partner. The SNIC prototype, which is a highly energy-efficient HPC system built from modern but standardized components, is currently installed at the SNIC centre, PDC, at KTH.

During 2008, the PRACE project analysed close to 70 scientific applications, many of which are heavily used by leading Swedish research groups. Based on the analysis of these applications, PRACE identified a comprehensive set of suitable architectures and selected a set of related prototypes. Most of these prototypes are located at major HPC centres in the larger EU

countries, including the PRACE prototype systems at the SNIC centre PDC. In 2008 PRACE also organized a Summer School in Stockholm, and a Winter School in Athens to teach programming for these prototype architectures. These courses included hands-on experience using the prototypes.

PRACE has set up a lightweight application procedure to enable researchers to port their applications to the prototypes and then evaluate the performance of the applications on the prototypes. The prototypes are available to all researchers within the PRACE countries. SNIC has also allocated funding for advanced user support to assist Swedish users in this process.

The Distributed European Infrastructure for Supercomputing Applications (DEISA) project was a predecessor to PRACE that deployed and operated a distributed supercomputing environment (consisting of national resources) with a European scope. The DEISA project was completed in April 2011. The SNIC centre PDC was an active partner in the DEISA project from 2009 onwards.

SNIC is also one of the twelve partners in STRATOS, the PRACE advisory group on Strategic Technologies. This group was created as the initial step in a continuous process to watch for promising hardware and software components to be used in future multi-petaflops systems, and to foster collaboration with industry in order to increase HPC expertise in Europe.

As a result of these activities, in 2010 the PRACE Research Infrastructure (RI) was established as an international non-profit association with a seat in Brussels, known as the Partnership for Advanced Computing in Europe AISBL. It has 21 member countries whose representative organizations are creating a pan-European supercomputing infrastructure that provides access to computing and data management resources and services for large-scale scientific and engineering applications at the highest performance level. These member countries are Austria, Bulgaria, Cyprus, the Czech Republic, Finland, France, Germany, Greece, Hungary, Ireland, Italy, the Netherlands, Norway, Poland, Portugal, Serbia, Spain, Sweden, Switzerland, Turkey and the UK.

In 2007, SNIC signed an official document that included KTH/PDC and the National Supercomputer Centre in Linköping in conjunction with Linköping University as third parties within the PRACE project. Later, in 2010, Umeå University also joined the PRACE project during the first implementation phase. On the 1st of September 2011, the second implementation phase of the PRACE project started with SNIC continuing to be an active partner. Since the start of the second phase, three more Swedish centres have become third parties in the project: Uppsala University, Lund University and Chalmers University of Technology. Lilit Axner, KTH/PDC, has been coordinating the efforts of these six SNIC centres within the PRACE research infrastructure. In September 2012, PDC's Michaela Barth took over this task when Lilit's maternity leave started.

and also holds dual Bachelors degrees in mathematics and cognitive science (specializing in neuroscience) from the University of California San Diego. David is currently working to complete a Ph.D. at KTH in the Department of Computational Biology, modelling cortical and subcortical brain circuits and behaviour using parallel computation. David's background includes a lot of industrial experience in the areas of computer science, bioinformatics, finance and software, where he was particularly interested in applying AI and database techniques for software solutions.

Staff Focus



ERIK JÄRLEBERG

Erik Järleberg recently finished his B.Sc. in computer science at KTH, and is currently pursuing the last year of his M.Sc. studies in discrete mathematics, also at KTH. Previously Erik lived for a year in Dublin, working with technical support for Hewlett-Packard's server range. Last summer he worked at the Centre for Autonomous Systems, KTH, on a research project developing algorithms for spatial predictions on floor plan abstractions for robot environments. At PDC, Erik is currently involved with support, where he helps researchers to use PDC's resources effectively. In his spare time, Erik likes to read literature, listen to music and spend time with friends.

Successful Swedish PRACE Applications

TIER-0 PROJECT ACCESS

- *Project name: REFIT - Rotation effects on flow instabilities and turbulence*
Project leader: Prof. Arne Johansson, Department of Mechanics, KTH
- *Project name: Direct numerical simulation of reaction fronts in partially premixed charge compression ignition combustion: structures, dynamics*
Project leader: Prof. Xue-Song Bai, Department of Energy Science, Lund University

PREPARATORY ACCESS

- *Project name: Visualization of output from Large-Scale Brain Simulations*
Principal Investigators: Prof. Anders Lansner, Simon Benjaminsson and David Silverstein, Department of Computational Biology, KTH. Research field: Medicine and Life Sciences
- *Project name: Grand Challenging Simulations in Materials Science. Project leader: Prof. Sven Öberg, Department of Engineering Sciences and Mathematics, Luleå University of Technology. Research field: Chemistry and Materials Sciences*
- *Project name: Automated Network Topology Identification and Topology Aware MPI Collectives. Principal Investigator: Dr. Chandan Basu, Dr. Soon-Heum Ko, Johan Raber, National Supercomputer Center, Linköping University. Research field: Mathematics and Computer Science*

DECI7 TIER-1 ACCESS

- *Project Name: DiSMuN - Diffusion and Spectroscopical Properties of Multicomponent Nitrides. Principal Investigator: Prof. Igor Abrikosov. Research area: Materials Science*
- *Project Name: SPIESM - Seasonal Prediction Improvement with an Earth System Model. Principal Investigators: Dr. Colin Johns and Prof. Francisco Doblas-*

DEISA Evolves into PRACE

by Lilit Axner, PDC

DEISA, the Distributed European Infrastructure for Supercomputing Applications, was a consortium of leading national Supercomputing centres that aimed to foster pan-European world-leading computational science research. DEISA deployed – and very successfully operated – a persistent, production quality, distributed supercomputing environment of continental scope. It set out to deliver a turnkey operational solution for a future European HPC ecosystem and realized that aim with great success. DEISA and its successor, DEISA2, were funded by the European Commission in FP6 and FP7 (www.deisa.eu).

DEISA started on the 1st of May 2004 and was completed on the 1st of May 2011. These seven years of activity were divided into two phases: DEISA and DEISA2. DEISA involved a consortium of 11 principal partners (BSC - Spain, CINECA - Italy, CSC - Finland, ECMWF - UK, EPCC - Scotland, FZJ, HLRS, LRZ and RZG - Germany, IDRIS-CNRS - France, and SARA - The Netherlands) and four associate partners (KTH/PDC - Sweden, CEA - France, CSCS - Switzerland and JSCC RAS - Russia).

The DEISA Extreme Computing Initiative (DECI) was launched in May 2005 by the DEISA Consortium as a way to enhance its impact on science and technology. The main purpose of this initiative was to foster a number of “grand challenge” applications in all areas of science and technology. These ground-breaking applications dealt with complex, demanding and innovative simulations that would not have been possible without the DEISA infrastructure, and that benefited from the exceptional resources provided by the consortium.

Since joining DEISA, SNIC has provided a dedicated 10 Gbit/s network to DEISA. This has given DECI users an opportunity to have a fast connection to supercomputers. From 2009 onwards, SNIC-KTH/PDC has committed five million CPU hours on the Ekman machine for the DECI6 application call. This last DECI6 call within the DEISA project was opened in December 2009, and the number of proposals the DEISA consortium received was record breaking: 122 applications were submitted from 30 countries – 22 of which were in Europe. The total computing time requested was over half a billion hours, which caused an oversubscription by a factor of ten. Three Swedish projects were among the 56 applications that were accepted for access to the pan-European supercomputing infrastructure. These projects were:

- **Braincor** (www.deisa.eu/science/deci/projects2010-2011/BRAIN-COR), Principal investigator: Prof. Anders Lansner, Head of the Department of Computational Biology at KTH.
- **Sive** (www.deisa.eu/science/deci/projects2010-2011/SIVE), Principal investigator: Prof. Erik Lindahl, Professor of Computational Structural Biology at KTH.
- **Wallpart** (www.deisa.eu/science/deci/projects2010-2011/WALL-

PART), Principal investigators: Dr. Luca Brandt and Dr. Philipp Schlatter, Department of Fluid Mechanics, Linné FLOW Centre, KTH.

Braincor received its access to the IBM Blue Gene at Jülich Supercomputer Center in Germany, whereas Sive and Wallpart received their time allocations on the CRAY XT5 machine at EPCC in Scotland. All three projects successfully used these resources during the period October 2010 - August 2011 and produced excellent results.

In exchange, five EU projects (from Cyprus, Finland, France, Italy and Scotland) used the resources on the Ekman cluster at PDC for their simulations during the same period. All five projects were completed successfully. This DECI6 call was an excellent example of successful collaboration between Sweden and other EU partners within the DEISA consortium.

DECI CALLS ACTIVE IN PRACE

After the completion of the DEISA project, PRACE decided to continue launching DECI calls due to the excellent outcome of the previous DECI calls. SNIC continues to provide the dedicated 10 Gbit/s network, and this year KTH/PDC again committed twelve million CPU hours on the supercomputer Lindgren for these calls. The DECI calls are made twice per year. The pilot DECI7 call was open from May – June 2011 and the applicants that were accepted were given access to resources starting from November 2011. The next DECI8 call was opened in January 2012, followed by DECI9 in April 2012. In November 2012, the DECI10 call was opened. Lindgren at PDC is one of the major providers of computing time for these calls for access to European Tier-1 systems.

The Scientific Case for the European High-End HPC Infrastructure

by Lilit Axner, PDC

Late in 2011, PRACE started a process to update the Scientific Case for the European high-end High-Performance Computing (HPC) infrastructure.

The original Scientific Case for high-end scientific computing in Europe was produced between August 2005 and April 2006 by an international scientific panel, consisting of representatives from Finland, France, Germany, Italy, the Netherlands, Spain and the United Kingdom. The origins of this endeavour lay in the recognition of the strategic role of HPC at the leading edge, so-called “Leadership-Class Supercomputing”, for European science and economy. At the heart of the HPCEUR initiative was the conviction that isolated European countries would not be able to provide their researchers and engineers with resources that would be competitive on the world stage.

In focusing on the requirements for the successful exploitation of

Reyes. Research field: Earth Sciences and Environment

• *Project Name: MUSIC. Principal Investigator: Dr. Mikael Djurfeldt. Research field: Computational Neuroscience*

• *Project Name: SIVE-2. Principal Investigator: Prof. Erik Lindahl. Research area: Biosciences: molecular dynamics simulation of viral entry*

DECI8 TIER-1 ACCESS

• *Project Name: PIPETURB - Large scale simulation of turbulent pipe flow. Principal Investigator: Dr. Philipp Schlatter. Research area: Engineering, Fluid Dynamics*

• *Project Name: CANONS - Comprehensive Ab initio studies of Nitride and Oxide fuels and Nuclear Structural materials. Principal Investigator: Dr. Pär Olsson. Research field: Materials Science*

• *Project Name: MBIOMARK - Multifunctional biomarkers for electron paramagnetic resonance imaging. Principal investigator: Dr. Zilvinas Rinkevicius. Research area: Materials Science*

• *Project Name: PLANETESIM - Towards an initial mass function of planetesimals. Principal Investigator: Dr. Anders Johansen. Research area: Astro Sciences*

DECI9 TIER-1 ACCESS

• *Project Name: CoStAFuM - Computational Studies of Advanced Functional Materials. Principal Investigator: Prof. Olle Eriksson. Research area: Materials Science*

• *Project Name: DifVib - Diffusion in multicomponent nitrides and vibrational thermodynamics from first-principles. Principal Investigator: Prof. Igor Abrikosov. Research area: Materials Science*

• *Project Name: HydFoEn - Hydrogen storage materials for energy applications. Principal Investigator: Prof. Rajeev Ahuja. Research area: Materials Science*

Leadership-Class Supercomputing – namely, providing in the order of 1 petaflops peak performance around 2009 – the panel emphasized that the exploitation of such a resource demanded an associated computational infrastructure.

Five years after the publication of the initial Scientific Case, the HPC landscape in Europe has changed significantly: the PRACE Research Infrastructure is providing Tier-0 HPC services to Europe; a global effort has been launched towards achieving exascale HPC by the end of this decade; and the importance of HPC in solving socio-economic challenges and maintaining Europe's competitiveness has become even more evident. Therefore, PRACE started an initiative to update the Scientific Case to define the current and expected future needs of the scientific communities.

During the first six months of 2012, five scientific panels, each consisting of about 15 researchers, have been working on the updating process. The panel members were chosen from the panels that contributed to the 2007 Scientific Case or that were active in the European Exascale Software Initiative (EESI), and from the PRACE Scientific Steering Committee (SSC).

The work done by the five panels essentially resulted in an entirely new Scientific Case, rather than an update of the previous document. The panels covered the following domain areas:

- Weather, Climate and Solid Earth Sciences,
- Astrophysics, High-Energy Physics and Plasma Physics,
- Materials Science, Chemistry and Nanoscience,
- Life Sciences and Medicine, and
- Engineering Science and Industrial Applications.

Two SNIC partners, PDC and NSC, have been closely involved in updating the Scientific Case for the European high-end HPC infrastructure. The two centres recommended eminent scientists from Sweden who make heavy use of HPC in their research as potential panel members. Of those recommended, Prof. C. Jones (SMHI), Prof. Erik Lindahl (KTH) and Dr. Philipp Schlatter (KTH) were selected as representatives from Sweden in the International Scientific Panel.

The updated Scientific Case is now available online at: <http://www.prace-ri.eu/PRACE-The-Scientific-Case-for-HPC>.

Application Scaling Activities within PRACE

by Lilith Axner, PDC



The costs for setting up and operating supercomputers, such as the PRACE Tier-0 systems, are enormous. Due to these high costs, it is important that code running on such systems makes efficient use of the resources.

To increase the usage efficiency of its resources PRACE has recently produced a range of optimization and petascaling guides for various European high performance computer systems. The main goals behind creating these guides were to investigate architecture-specific factors influencing performance, to collect best practices on how to achieve good performance on the systems, and to disseminate this knowledge to users. The topics covered in these best practice guides include:

- optimal porting of applications (for example, the choice of numerical libraries and compiler options),
- architecture-specific optimisation and petascaling techniques,
- optimal system environment (for example, tuneable system parameters, job placement and optimised system libraries),
- debuggers,
- performance analysis tools, and
- programming environments.

As a result of these activities a range of best practice guides were produced and published at www.prace-ri.eu/Best-Practice-Guides.

PDC and HPC2N in Umeå made major contributions to the work on the Cray best practice guide. The experts at these two centres carried out detailed investigations on the Cray XE6 system Lindgren at PDC (www.pdc.kth.se/resources/computers/lindgren). The results were published in the Cray XE guide, along

with the results obtained by other PRACE partners on several Cray systems such as Hermit, a Tier-0 system at HLRS in Stuttgart, Germany (www.hlrs.de/systems/platforms/cray-xe6-hermit), and HECToR, a Tier-1 system at EPCC, in Edinburgh, Scotland (www.hector.ac.uk).

These best practice guides are very useful and have already received widespread praise from users of these systems. The guides contain many hints and examples that can be used to optimize runs on these systems drastically.

PETASCALING WITH COMMUNITIES

PRACE is also directly working with research communities to develop relationships with the communities responsible for important European applications, and to provide petascaling expertise to ensure that the code for their key applications could effectively exploit the Tier-0 systems.

Eleven key applications that are of value to the main scientific communities at European level were selected for petascaling in this activity. The applications were:

- GROMACS and DL-POLY from Molecular Dynamics,
 - Quantum Espresso and CP2K from Materials Science,
 - GPAW from Nanoscience,
 - DALTON from Computational Chemistry,
 - EUTERPE from Plasma Physics,
 - the EC-Earth 3 suite from Meteo-climatology,
 - SPECFEM3D from Earth Sciences (earthquakes),
- and
- OpenFOAM and Code_Saturne from Engineering and CFD.

This petascaling work involved the owners of the application codes, and the scientific communities interested in the codes, collaborating with PRACE experts.

PDC was one of the PRACE partners in this work and, together with several other EU partners, proposed working on GROMACS and DALTON – the molecular dynamics and computational chemistry community codes.

The PDC work on GROMACS was carried out in collaboration with PRACE partners from CINECA SuperComputing Applications and Innovation Department in Italy, and the National Center for Supercom-

Apply for Parallel Programming Support through SNIC

SNIC now has a number of application experts in Computational Science and Parallel Programming who are available to help with your code development projects. For example, these experts can help with debugging, profiling, and tracing, as well as with simpler optimization, parallelization or code porting tasks. Up to ten full workdays of assistance may be allocated to each request. The requests are handled as the experts have time available.

Any researcher affiliated with a Swedish tertiary institution (including Ph.D. students) can submit a request for Parallel Programming support.

It may be possible to allocate help to projects that require more assistance. If you have a project that might require more than ten workdays of assistance, please submit a request, so an application expert can assess how much assistance would be required. We can then discuss the options for assigning an application expert to the project for a longer term.

Please submit requests to application-support@snic.se. For more information, see <http://docs.snic.se/wiki/Support>.

puting Applications in Bulgaria, as well as with the GROMACS code developers, Dr. Berk Hess and Prof. Erik Lindahl from the Department of Theoretical Physics, Royal Institute of Technology (KTH), Stockholm, Sweden. The work on this project also involved collaboration with the ScalaLife EU project, which deals with Scalable Life Science Software including GROMACS.

In this work we carried out extensive studies of the performance of several versions of GROMACS on a diverse set of computer architectures and with varying runtime conditions and program options. The aim was to determine the optimal set of conditions for a given architecture and input, without making changes to the code itself. The second set of studies investigated the hybrid MPI/OpenMP version of GROMACS as released by the ScalaLife project (version 4.6) where OpenMP

threads are used for the PME (Particle Mesh Ewald) calculation.

The results demonstrated that on systems with very high numbers of cores, such as those of the Blue Gene range, the all-to-all communication scheme of the PME calculation will limit parallel scaling quite quickly. This does not mean that users should not attempt to optimize the run conditions on Blue Gene-type architectures, but rather that a priori it is difficult to gauge which set of conditions would work best. The investigations of the hybrid-OpenMP/MPI mode showed a significant improvement of scaling on large systems with a large number of cores, typically with over 200,000 particles on more than 500 cores. On systems with very fast interconnects, such as the Cray XE6 Lindgren at PDC, the hybrid mode did not offer much advantage. These investigations demonstrated that increasing communications is the main reason for the performance deterioration of the MD simulation packages on large numbers of cores, which is typical for the case of petascale computations.

The work on DALTON was carried out in collaboration with PRACE partners from The Centre for Theoretical and Computational Chemistry (CTCC), Norway, as well as the University Center for Information Technology (USIT) at the University of Oslo in Norway, STFC Daresbury Laboratory in the UK and the Barcelona Supercomputing Center in Spain. The developers of the DALTON code in Norway and Sweden, including Dr. Zilvinas Rinkevicius and Prof. Olav Vahtras from the School of Biotechnology, Division of Theoretical Chemistry & Biology at KTH, were also closely involved in this collaboration, along with the ScalaLife project.

The aims of the work with the DALTON code were as follows:

- 1) Analyse the current status of the DALTON code and identify bottlenecks. Implement several performance improvements in DALTON QM/MM (Quantum mechanics/Molecular mechanics) and make a first attempt at hybrid parallelization.

- 2) Incorporate MPI integral components into LSDALTON, and improve the optimization and scalability, along with the interface of matrix operations to PBLAS and the ScaLAPACK numerical library routines.

- 3) Interface the DALTON and LSDALTON QM codes to the ChemShell QM/MM package and bench-

mark the QM/MM calculations using this approach.

- 4) Analyse the impact of the DALTON QM/MM system components with Dimemas.

As a result of this work on the DALTON QM/MM system, the modification of the master-worker design pattern to a team of masters made it possible to calculate several contributions at the same time and increased the scalability of the application accordingly.

For the LSDALTON code, a good scalability up to 2,048 cores for pure DFT calculations was demonstrated. This was achieved through the development of MPI integral components, and utilizing a newly-developed task partitioning scheme. The interface to the PBLAS and ScaLAPACK numerical library routines also contributed significantly. Thus the main code framework and strategies needed to exploit parallelism at all levels of the code were established.

Both the LSDALTON and DALTON QM codes were also successfully interfaced to the ChemShell QM/MM package with both binary and directly-linked interfaces. Through ChemShell, the DALTON codes have access to a range of MM approaches and supporting functionality that make flexible QM/MM modelling possible on PRACE systems. The LSDALTON interface also supports shared MPI communicators for advanced task-farming parallelization techniques. Benchmark QM/MM calculations using ChemShell/LSDALTON achieved good scaling up to at least 1,024 cores.

This successful work on both the GROMACS and DALTON applications is discussed in more detail in white papers that are available on the PRACE web site at: www.prace-ri.eu/white-papers.



PRACE Scientific Seminar 2011 (see page 12)

PDC's and PRACE's Commitment to Energy Efficient Computation

by Gert Svensson and Gilbert Netzer, PDC

Energy efficiency is an important goal for PRACE projects, which is not surprising when we are talking about the largest, and hence hottest, supercomputers in Europe. PDC, together with its PRACE partners, is taking a dual approach to this problem. On the one hand, new technologies can help to reduce the overall power consumption of supercomputers, and, on the other hand, the heat produced by supercomputers can be reused in an environmentally friendly way.

PROTOTYPES WITH ENERGY EFFICIENT TECHNOLOGIES

The HPC industry is racing to achieve even higher performance measured in exaflops (that is, 10^{18} or a million million floating point operations per second) by the end of the decade. However one particular roadblock in this race is that even today's warehouse-sized supercomputer installations (where the fastest provides a modest 0.016 exaflops) require around eight million watts of electrical power. This burns deep holes in the pockets of the institutes funding these machines. Therefore a major concern for the industry, and for hardware prototyping activities, is to squeeze more calculations out of the same amount of electrical energy, thus keeping the power demand under control.

Sweden and in particular PDC/KTH have been very active in this respect with a series of three prototype systems. Two of these prototypes are currently active, and the third has been selected by the PRACE project for the latest wave of prototype deployments.

The first of these prototypes at PDC was installed during the PRACE preparatory project in the summer of 2009. It consisted of an energy-optimized standard x86 cluster, built out of off-the-shelf components supplied by Supermicro. The system used six-core AMD Opteron 8425HE high efficiency processors in a diskless four socket blade configuration interconnected by a scalable QDR Infiniband network. While the system could not beat the contemporary leading IBM Blue Gene/L line of special purpose supercomputers in the all-important Linpack benchmark (measuring raw floating point

performance versus energy consumption in FLOPS/W), it was able to deliver higher application energy efficiency using the GROMACS molecular dynamics simulation software than the much more expensive IBM system. Today, similar stripped-down standard approaches are used by large cloud data centres and Cray supercomputers like PDC's Lindgren.

The second PRACE prototype project at PDC attempts to reach even higher energy efficiency by taking advantage of an alternative computer architecture, called digital signal processors (DSPs), that is not traditionally used in HPC. These special purpose chips are usually employed in embedded and telecommunications applications where they, for instance, provide the computational muscles for medical ultrasound devices, or transcode voice and video signals in telecommunication networks. The second PDC prototype studied the TMS320C6678 processor made by Texas Instruments. This processor packs eight DSP cores (capable of delivering 25 gigaflops) into a single package while consuming only 12 watts of electrical power. Preliminary measurements of 2 gigaflops/W place the chip at about the same energy efficiency as the leading Blue Gene/Q supercomputer line with 2.1 gigaflops/W.

As a continuation of this effort, PDC put forward a proposal in cooperation with the Irish Centre for High-End Computing (ICHEC) to investigate a promising processor architecture manufactured by the small Irish start-up company Movidius Ltd. The chip in question is primarily made for video processing in future generation smartphones. It only consumes about 0.35 watts of power, yet projections indicate that it should be able to deliver up to 50 gigaflops and reach a computational efficiency of up to 140 gigaflops/W. This would make it possible to achieve about 2.8 exaflops from the power budget of 20 MW that is currently envisioned for future exascale systems, and still allow for the interconnect network and system losses. This just leaves us with a small problem to solve: how to make the necessary twenty million chips work together!

THE PDC HEAT REUSE PROJECT

A couple of years ago, PDC started a project to install a new Cray system. The new Cray would nearly double the power consumption at PDC from around 800 kW to 1300 kW, and consequently produce a great deal of

additional heat. We wanted to find a way to make use of the extra heat from the new Cray in a manner that contributed positively to the environment.

After considering various possibilities, we decided to use the excess heat from the new Cray to heat the Chemistry building, which is relatively close to the PDC computer room. The air cooled version of the Cray system takes in cold air (at less than 16 °C) from under a raised floor and passes that air through the racks of computer components. The air is heated to temperatures of 35–45 °C, on its way to being passed out at the top of the racks. By placing standard industry air-water heat exchangers above the Cray, we could heat water which could then be sent to the Chemistry building for heating purposes.

This system for heating the Chemistry building with waste heat from the Cray became fully operational during the winter of 2011-12. It has been effective in providing environmentally friendly cooling for the new Cray that also offsets part of its operational costs. The project has been reported in the PRACE White Papers, as well as at the Cray User Group meeting 2012. More details can also be found in the previous issue of the PDC Newsletter.

REFERENCES

Green500 List - www.green500.org

Povel - www.pdc.kth.se/resources/computers/povel

PRACE Prototypes & White Papers -

www.prace-ri.eu

www.prace-ri.eu/white-papers

Texas Instruments TMS320C6678 -

www.ti.com/product/tms320c6678

Movidius Ltd. - www.movidius.com



PDC Summer School 2012 lecture (see page 15)

2011 PRACE Events with SNIC-PDC

THE 3RD PRACE INDUSTRY SEMINAR “EUROPE GOES HPC: INDUSTRIAL COMPETITIVENESS”

The aim of this seminar was to bridge the gap between industrial HPC users and researchers from academia. The event was held on 28–29 March 2011 at the Waterfront Congress Centre in Stockholm, Sweden, and was organised by SNIC, with the PDC Center for High Performance Computing at KTH serving as the local organiser. There were over 80 participants from 22 countries, representing 39 companies and 13 PRACE partners. For more information, see

www.prace-ri.eu/Presentations-from-the-3rd-PRACE.

PRACE SCIENTIFIC SEMINAR – HPC BOOSTS SCIENCE AT KTH

PRACE organized a three day scientific seminar “HPC Boosts Science” in Stockholm in February 2011. The aim of this seminar was to tackle different aspects of HPC within multiple scientific disciplines. The seminar was held at the Royal Institute of Technology (KTH) and was organized by SNIC-PDC with the help of the CSC-IT Center for Science Ltd. in Helsinki, Finland. For more details see

www.prace-ri.eu/PRACE-Scientific-Seminar-HPC.

PRACE SUMMER SCHOOL 2011

The 2011 PRACE annual summer school was held from 29 August – 1 September 2011 in Espoo, Finland. The school was hosted by the CSC – IT Center for Science Ltd, Finland, who co-organized the event with SNIC-PDC. During the four day event, the participants learned advanced parallel programming skills, which are necessary for getting the most out of the largest (Tier-0) supercomputers that the PRACE Research Infrastructure makes available to European scientists and engineers. Further information is available at

www.prace-ri.eu/PRACE-Summer-School-2011-29-August

Zorn - New GPU Cluster Available for Research

by Michael Schliephake, PDC

PDC's newest cluster utilises graphics processing units (GPUs) and has therefore been dubbed after the Swedish painter Anders Zorn. The system is surprisingly small but nevertheless has impressive capabilities: Zorn reaches a peak performance of more than 45 teraflops. Its 12 compute nodes are packed with 40 GPUs and connected through an Infiniband QDR (Quad Data Rate) network. The local disk capacity for everyday usage amounts to 10 TB.

Systems with comparable computational power first started to appear on the TOP500 list about a decade ago. They were big installations at that time and required several hundred kilowatts of electrical power to run – in comparison, Zorn fits into a single rack and uses less than 20 kilowatts. Such an exciting development over just a few years has been made possible by using graphics processing units for numerical simulations. This started about ten years ago with enthusiasts who used the programmable components of graphics cards through clumsy interfaces for calculations (rather than using the cards for their original aim of producing images on computer screens). Manufacturers soon recognised the potential of this idea, and have since been driving development to make GPUs equally well-suited both for the display of graphics and for performing numerical simulations. The use of GPUs for numerical simulations is a well-established practise nowadays. Three of the ten largest supercomputers in the world use this technology to boost their performance, although this approach is not limited to just high-end supercomputers. These days many scientists have a desktop computer equipped with a GPU that provides permanent and cheap access to processing capabilities that, only a few years ago, would have required support from a university computer centre.

GPU computing at PDC started in 2010 with an initial small installation. It was supported by a grant from the Knut and Alice Wallenberg Foundation that was awarded to one of KTH's research groups involved in the development of advanced simulation programs. SNIC also picked up the trend and started a large-scale evaluation of GPU technology. That was an opportu-

nity for PDC to triple the computational power of the initial system. We currently provide Zorn to all Swedish researchers. Its size allows jobs to be run at production-level, and thus delivers a valuable contribution to ongoing research projects.

A handful of projects have already started on Zorn: users from both KTH and Linköping University are active on the system. The projects they are working on include simulations in the field of molecular dynamics and hydrodynamic calculations in astrophysics, the development of applications in computational chemistry as well as for the general solution of partial differential equations, and last, but not least, the development of software tools to simplify the process of programming highly parallel and heterogeneous computer systems.

We hope this promising start will motivate even more users to try the advantages of GPU computing for themselves. PDC supports this in various ways. We offer, for example, a compact two-day course that provides an introduction to programming GPUs. Graduate students can start their exploration of GPU programming through a course that combines an overview of GPU technology with a programming project. The project work provides each student with the opportunity to work on a problem related to their research field, and to exchange experience with fellow students and the teacher.

For individual researchers or research groups, PDC also offers personal consulting tailored to the specific needs of the research in question. Such a consultation might, for example, involve a discussion about general approaches to GPU programming and best practices, or analysis and consulting to solve specific programming issues.

For further information about GPU programming or to arrange a consultation, contact Michael Schliephake (michs@kth.se). Michael is the Zorn system manager, and he is also responsible for providing advanced support for GPU computing.



Four compute nodes of the GPU cluster Zorn

Staff Focus



CARL JOHAN HÅKANSSON

Carl Johan Håkansson started working as a system administrator at PDC in August this year. He holds a B.A. in philosophy from Stockholm University, and has also studied computer science at Uppsala University and at KTH. For most of his career, Carl Johan has worked as a computer consultant with an emphasis on databases and Linux servers. He has also worked for several years in the media software industry, mainly with storage software for the television broadcasting industry. As a system administrator at PDC, Carl Johan will initially be working with data storage and is currently involved with the iRODS project.

Carl Johan is a local boy who has lived in the Stockholm area for most of his life. Apart from having a keen interest in computers, he likes playing around with martial arts, and is currently spending a lot of his spare time exploring the southern Chinese self-defence system Wing Chun.

CRESTA Collaboration Meeting at KTH

by Stefano Markidis, PDC/HPCViz



KTH is one of the main partners in the FP7 European Commission project about Collaborative Research into Exascale Systemware, Tools and Applications – otherwise known as CRESTA. The goal of this three-year project is to study scalable software – from compilers to run-time systems, from debuggers to performance-monitoring tools, from numerical libraries to pre- and post-processing tools – for the next generation of exascale supercomputers. A large part of the CRESTA work is based on the concept of “co-design”, that is, the study of how advances in programming models, performance-monitoring tools, debuggers, and algorithms influence scientific applications (and vice-versa) on the road to the exascale era.

PDC and the new KTH High Performance Computing and Visualization department (HPCViz) hosted the first annual collaboration meeting of the CRESTA project at the KTH library on the 11th and 12th of September this year. The goals of the meeting were first to review the achievements of the initial year of the project, second to assess the CRESTA co-design activities to date, and third to outline the steps to be undertaken in the future by the CRESTA consortium.

The meeting opened with a keynote speech by Dr. Mark Parsons, CRESTA coordinator and project manager from the Edinburgh Parallel Computing Center (EPCC), who congratulated the project for the successes of the first year and presented the new challenges waiting ahead for CRESTA. In particular, Dr. Parsons showed that the only way to meet these challenges will be with new “disruptive” solutions. He invited CRESTA members to “think the unthinkable” in the quest to find solutions to the new problems that we will face in the exascale era.

The majority of the presentations were dedicated to co-design work within CRESTA, showing the current status of various activities and outlining the directions to be taken in future. (Part of this work will be presented at the co-design workshop – which is being co-organized by the KTH HPCViz department – at the next Supercomputing meeting in Salt Lake City in November 2012.) All the CRESTA scientific applications (GROMACS, Nek5000, IFS, HemeLB, OpenFOAM, and Elmfire) were presented and new programming techniques were also explained – for example, the OpenACC compiler directives for programming Graphics Processing Units (GPUs).



Introduction to High Performance Computing - PDC Summer School 2012

by Stefano Markidis, PDC/HPCViz

PDC, in collaboration with the KTH School of Computer Science and Communication (CSC), held the 2012 summer school – “Introduction to High Performance Computing” – at the main KTH campus on 20–31 September. The course is endorsed by the KTH Computational Science and Engineering School (KCSE), and is supported primarily by the National Graduate School in Scientific Computing (NGSC) and this year also by Google.

The summer school was attended by a total of 55 students who were selected from over a hundred applicants. Most of the students were from KTH and Stockholm University. Other Swedish and European universities and Peking University were also represented, and there were also attendees from various businesses. The participants came from a wide range of scientific backgrounds: computational biology, neuroinformatics, materials science, plasma physics, and numerical analysis, to name but a few.

The school lasted two weeks and comprised a series of lectures about traditional programming techniques for parallel computers (such as OpenMP and MPI), and about computer architectures and networks for High Performance Computing platforms. Part of the course focused on optimizing scientific codes in order to improve their performance. The architecture of Graphics Processing Units (GPUs) was also introduced, along

with OpenCL programming. The final topics that were discussed were the next generation of exascale supercomputers, and the novel programming techniques that are emerging for these supercomputers.

The summer school students attended lectures, given by international speakers, in the mornings. In the afternoons, the students were given the opportunity to experience parallel computing, optimization techniques and GPU programming in the lab. They were able to work on Lindgren (a massively parallel supercomputer) and Zorn (a GPU cluster at KTH). During the lab sessions, PDC experts and CSC teaching assistants helped the students to complete the given assignments.

As part of the course, the summer school students are working on projects that require them to use parallel computing and the programming techniques they learned during the summer school. PDC and CSC experts are assisting the students in this work.

The summer school was enlivened by various social activities, including a picnic on the first day of the school, which gave the students a chance to meet and get to know each other. There were also visits to the PDC supercomputer hall, and a final dinner.

The feedback from the students reported that the summer school was very successful. In particular, the students appreciated the high quality of the lectures, and found the laboratory assignments very useful for broadening their understanding and practical skills in the realms of supercomputing.

PDC Related Events (Sponsored/Associated)

SNIC Cloud Computing Course at Chalmers

October 30, 2012, Chalmers University of Technology, Gothenburg, Sweden

PDC has organized a beginners' course on cloud computing for SNIC users and Chalmers University researchers. The course is being run by Zeeshan Ali Shah and is part of the SNIC Cloud Project, helping SNIC users to start using cloud computing.

For more information, see:

www.pdc.kth.se/resources/computers/swecloud/Chalmers_cloud-computing-course-101

Preparing Applications for Exascale Through Co-design Workshop

November 16, 2012, Supercomputing12 (SC12) Conference, Salt Lake City, Utah, U.S.A.

Stefano Markidis of PDC/HPCViz is co-organizing a workshop to be held as part of Supercomputing12 in Salt Lake City – the International Conference for High Performance Computing, Networking, Storage and Analysis.

For more information, see:

sc12.supercomputing.org/schedule/event_detail.php?evid=wksp122



PDC Summer School 2012 visits to PDC computer hall (see page 15)

Sources

We can recommend the following sources for other interesting HPC opportunities and events:

CERN

cerncourier.com/cws/events
cdsweb.cern.ch/collection/Conferences?ln=en

EGI

www.egi.eu/about/events

HPC UNIVERSITY

www.hpcuniv.org/events/current

HPCWIRE

www.hpcwire.com/events

LINUX JOURNAL

www.linuxjournal.com/events

NETLIB

www.netlib.org/confdb

PRACE

www.prace-project.eu/prototype-access

www.prace-project.eu/hpc-training-events

www.prace-project.eu/news

SNIC

www.snic.vr.se/news-events

US DEPARTMENT OF ENERGY

hpc.science.doe.gov

XSEDE

www.xsede.org/conferences-and-events