NSC News

2011:1

National Supercomputer Centre at Linköpings universitet

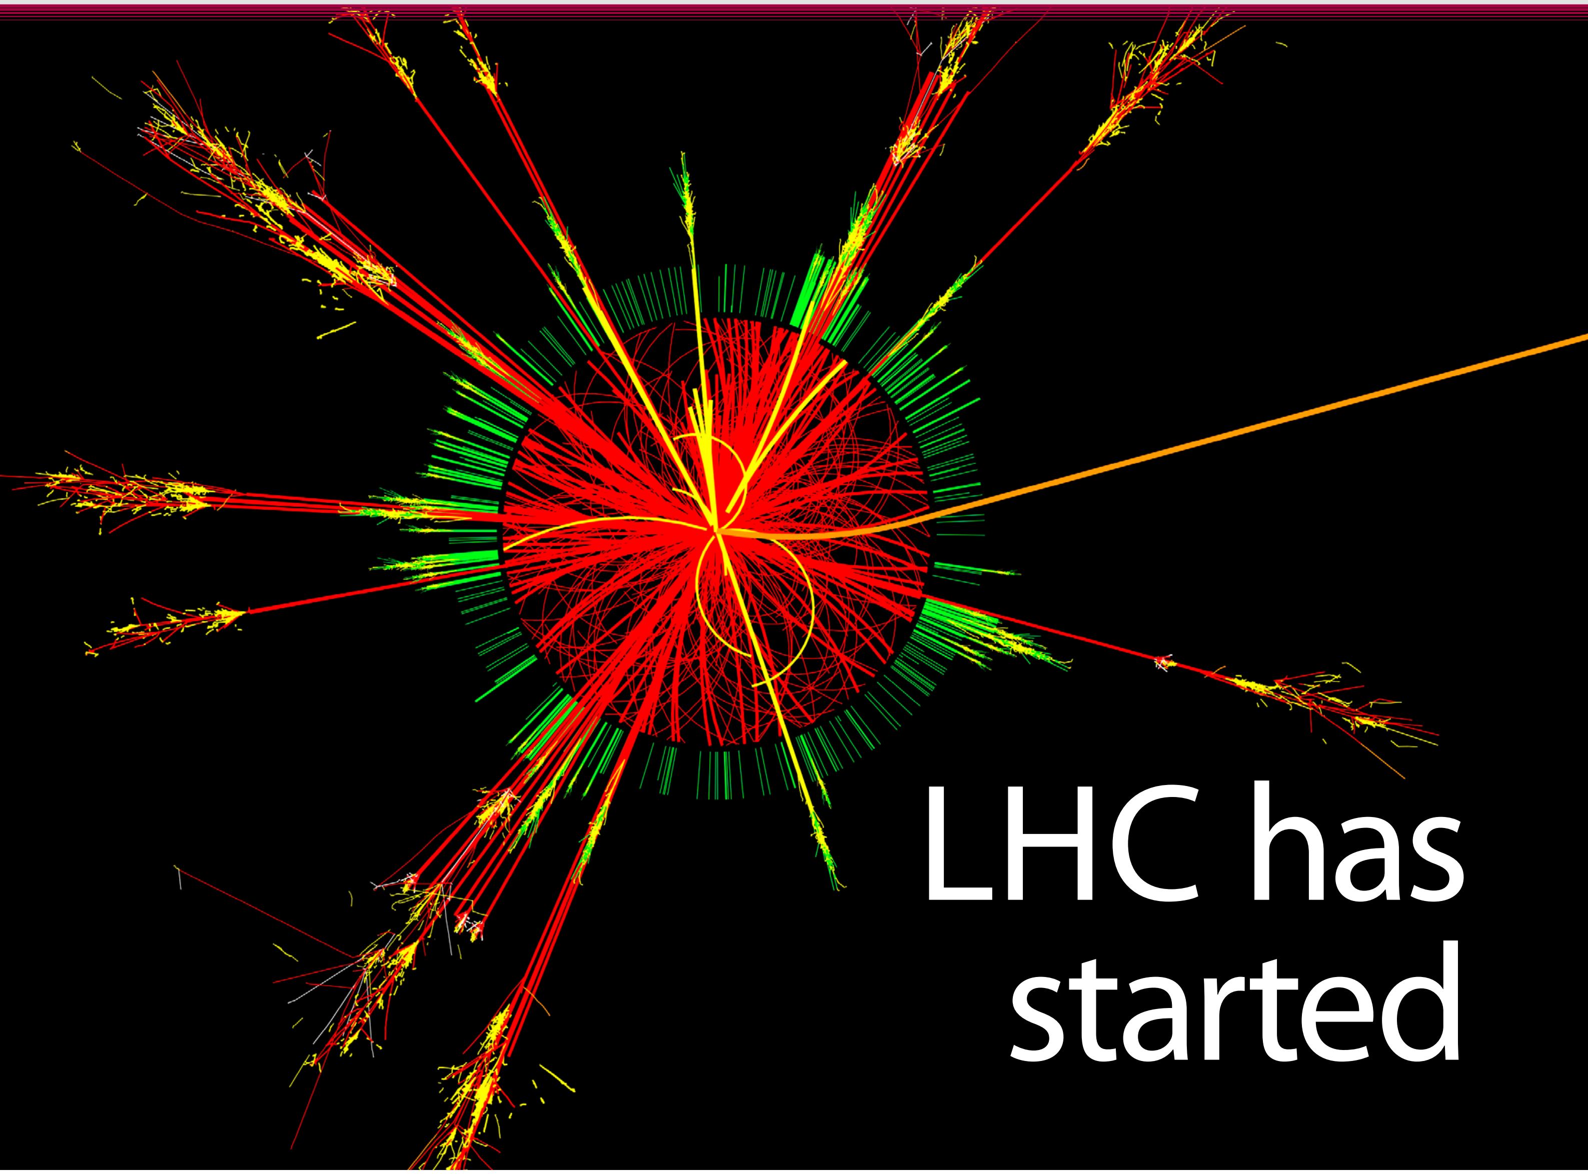

The great discovery journey with LHC has started

The large hadron collider, LHC, at CERN has just started up after a service period earlier this year. Professor Tord Ekelöf at Uppsala University and chair of the Swedish LHC consortium explains how it can be used to answer fundamental scientific question and how the World LHC Computing Grid, WLCG, will be used to analyze the huge amount of generated data.

Read further on page 3

Successful training event at NSC

NSC hosted a very popular Intel HPC training event 9 – 11 February. In fact it was so popular that the number of participants was more than doubled compared to what was originally planned for.

Read further on page 5

National application experts

A number of application experts has now been hired at several SNIC centres. Initiated by NSC and PDC through SeRC (Swedish e-Science Research Centre), a national application experts meeting was held on 2–3 December in Linköping, which was a very successful event, allowing for fruitful discussions between application experts from all SNIC centres. Among the noticable results from this meeting, I would like to mention the new Application Knowledge Base, initially suggested by SeRC and now getting support also from all other centres. SNIC encourages this initiative, and the Knowledge Base is hosted at NSC under the name http://snicdocs.nsc.liu.se. Here, you will find condensed and important information on software running on SNIC systems. Another important development is the setup of a national common queue for support questions to application experts, available at E-mail "application-support@nsc.liu.se". In order to coordinate all these efforts, Torben Rasmussen at NSC has been assigned the role as national application expert coordinator.

At NSC, we have since 1 January 2011 Joel Hedlund as application expert in bioinformatics, a relatively young scientific discipline that is increasingly using more and more SNIC resources as more and more biomedicial data become available and need interpretations and additional computational analyses. Furthermore, NSC is currently recruiting a full-time application expert in materials sciences in addition to the combined researcher/application expert within the MATTER consortium that also is under recruitment.

I am also happy to welcome our new computational scientist, Soon-Heum Ko, recrutied from USA, who started at NSC in March 2011. We now have two computational scientists at NSC, providing advanced skills in code optimisation issues, which will be increasingly important as the parallellism in modern clusters is steadily increasing. Furthermore, I am happy to welcome our software developer Henrik Wiberg, who adds important knowledge in the field of creating user-friendly, "tailor-made" and efficient solutions for our users.

Thus, we are very happy that NSC will be able to help our users in more and more areas associated with HPC – computational efficiency, large-scale storage solutions, and application expertise.

BENGT PERSSON, NSC DIRECTOR

Apply for resources

The dead-line this spring for large scale allocations on SNIC resources is April 18 at 15:00. The Swedish National Allocations Committee, SNAC, accepts applications for large scale allocations on Swedish National Infrastructure for Computing, SNIC, resources twice a year. One application round is in the spring and one in the autumn with allocations valid from first of July and first of January respectively. Large scale applications are for allocations above the maximum limit for medium scale allocations which currently is 80 kilo-core-hours/month, per resource. In this application round 1200 kh/m will be allocated on Neolith and 250 kh/m on Kappa. Instructions and an application form is available on: http://www.snic.vr.se/apply-for-resources/large-scale-applications.

Medium scale applications on the other hand can be submitted throughout the year. Applications are handled locally on the six SNIC centers and usually within two weeks. Instructions and links to application forms are available on: http://www.snic.vr.se/apply-for-resources/medium-scale-applications.

PETER MÜNGER

New staff member

Hi there! I am a fresh recruit to NSC as an applications expert in bioinformatics. Born in Happy Hudiksvall halfway up the eastern coast I came to Linköping in 1998 to study, and now, one M.Sc. in Engineering Biology and one Ph.d. in Bioinformatics later, I'm still here. My research has mainly been in analysis of biological sequences (proteins and genes) with a focus on protein families and hidden Markov models, and I used as well as developed bioinformatic software for HPC and Grid systems.

JOEL HEDLUND

The great discovery journey with LHC

In November 2009, the first proton-proton collisions were recorded at CERN's large hadron collider LHC in Geneva. The collision energy was then 0.9 TeV (1 TeV = 10^{12} electron volts) and the rate of proton-proton collisions a few per second. Since then the energy has been increased to 7 TeV and the rate of proton-proton collisions to about 10 million per second. This has been sufficient to observe at the LHC during the summer and autumn of 2010 all the elementary particles discovered in earlier particle physics experiments since the 1970's: the charm, bottom and top quarks and their bound states with other quarks, the W^+, W^- and Z^0 bosons and the gluon, all describe with high precision by the Standard Model theory of particle physics. Towards the end of 2010 LHC was run with lead nuclei instead of protons and new observations could be made of very high-energy, collective quark-gluon states.

During January and February 2011 LHC has undergone service and is now about to be started up again in March. The plan is to run LHC almost continuously at 7 TeV collision energy until the end of 2012. The expectation is that it will be possible to raise the rate of collisions gradually during the spring and summer this year to about 100 million per second. LHC has been built with the aim to detect new elementary particles of large mass. To produce particles of large mass requires a lot of energy ($E = mc^2$) and the high energy of the LHC is one of the two factors that will make possible the discovery of new heavy particles – i.e. particles that are heavier than those that have already been discovered at the accelerators of lower collision energy used before LHC, in particular LEP at CERN and the Tevatron at Fermilab in Chicago.

But the production of an elementary particle of large mass is demanding not only by the fact that it requires a lot of energy – it also is an extremely rare process such that many billions of proton-proton collisions are required for the production of a single particle of that type. That is why the planned increase in the rate of collision at the LHC is crucial for the possibility to discover new high mass particles.

Which are then the hypothetical high mass particles that could possibly be detected at LHC? According to the Standard Model there is only one more particle to be discovered – the Higgs boson – which is a heavy, electrically neutral particle without spin, which, with the LHC collision frequencies planned for 2011, could be discovered even after only a few months of operation of the LHC according to the Standard model. But why then look for other large mass particles if the successful Standard Model predicts that there should be no more particles to discover except the Higgs boson? The answer is that although the Standard Model well describes all the particles and interactions so far observed in accelerator experiments, there are other fundamental questions in physics which the Standard Model cannot answer, e.g. why there is only matter and no anti-matter in the world around us and which is the nature of the particles that make up the Dark Matter in Universe. To also explain these phenomena a more general theory is required. A large variety of candidates for such a more general theory have been put forward.

Many of the new candidate theories assume the existence of a new symmetry in nature, supersymmetry, which is a symmetry between the particles that are the smallest constituents of matter (quarks and leptons) and the particles that convey the interactions between these constituents (photons, gluons and the weak interaction bosons). The supersymmetric theories appear in different forms but all have in common that they predict the existence of a complete set of heavier

A view of the Large Hadron Collider (LHC) tunnel at CERN in Geneva

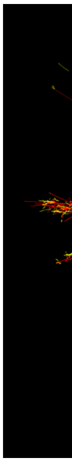

A high ene

New staff member

Hi, I am Soon-Heum Ko, who joined NSC as a computational scientist in March 2011. Prior to NSC, I worked as a postdoctoral researcher at a Center for Computation and Technology in US, on the development of a hybrid computational fluid dynamics (CFD) - molecular dynamics (MD) toolkit and the implementation of CFD codes on the Cactus framework. I was also involved in a Korean e-Science project for designing/developing a CFD research environment. These experiences will be helpful in conducting application code porting, performance optimization and other relevant works at NSC. I am looking forward to getting along with everybody at NSC!

SOON-HEUM KO

has started

supersymmetric particles that are partners to all the “ordinary” elementary particles we currently know of. According to many of these theories, with the collision frequencies planned for 2011 it is likely that supersymmetric particles will be produced and detected at the LHC even during the current year. Furthermore, supersymmetric theories predict the existence of, not only one, but several Higgs bosons, both electrically neutral and electrically charged ones. The discovery of supersymmetric particles and several, both neutral and charged, Higgs bosons would be a convincing proof for that nature is supersymmetric.

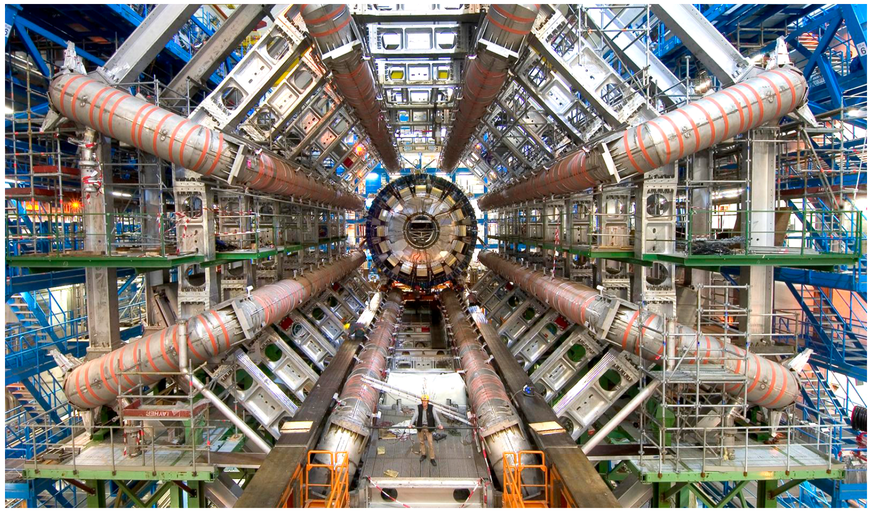

The ATLAS experiment under construction

How are these measurements then made in practice in the experiments at LHC? Enormous detector systems have been built in large caverns situated at the points around the underground 27 km long LHC accelerator ring where the two counter-circulating proton beams are made to collide. These detector systems, large as apartment buildings, contain a huge amount of particle detectors. The detectors register the 100's of particles that are created in, and projected out from, each proton-proton collision. The amount of detector data recorded for each collision is about one Megabyte. To record one Megabyte of data from each of 100 million collisions occurring every second would require 10^{14} bytes per second to be recorded – a totally impossible task.

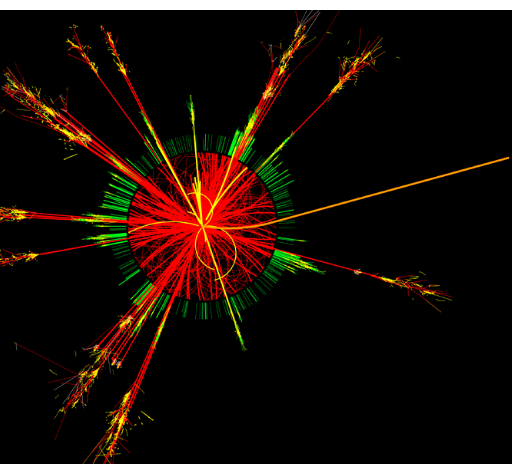

ergy proton-proton collision as detected in the ATLAS experiment at LHC

However, we have already noted that the production of a particle of very large mass is a very rare event. So it is the data from just a minute fraction of all collisions that are of interest. The collisions of interest contain a particle of very large mass which implies that its measurement data has very special features. The heavy particle is unstable and immediately decays into other more stable particles, each of which has a much smaller mass and is stable enough to be recorded by the particle detectors. The decaying particle's mass is converted into kinetic energy and, if the mass is very large, each of the stable low-mass particles from its decay thereby obtains a large kinetic energy. A proton-proton collision featuring such very energetic stable low-mass particles can be distinguished from the very large number of ordinary proton-proton collisions, as the low-mass particles produced in the latter type of collisions in general have a much lower kinetic energy. The presence of very energetic, low-mass stable particles in a collision can be rapidly detected by the use of very fast processors, called trigger processors. These processors are linked to the data stream from the particle detectors and used in real-time for the selection of the small number of collisions in which there are very high-energy, low-mass stable particles. In this way, 100 out of the 100 million collisions per second are selected in real time and recorded every second. Thereby, the amount of raw detector-data that needs to be recorded and stored is reduced to approximately 10^{8} bytes (100 Megabytes) per second or about 10^{15} bytes (1 Petabyte) per year.

The storage and analysis of this enormous flow of data poses a huge challenge in terms of disk storage space and data processing power. To meet this challenge the LHC research community started already in the late 1990s to develop the Computing Grid concept that had

New staff member

I am a software developer that has joined NSC to take part in the IS-ENES project to develop services for storing and retrieving climate data. I come from Linköping where I also studied computer science and engineering. Now me and my family lives in a house in Ljungsbro. My background as a software developer ranges from real-time systems for train security (AdTranz, Stockholm 3 years), Work flow management systems (Idalnfront, Linköping 2 years) and Warehouse management systems (Swisslog, Boxholm 6 years). My basic interest in computer science lies in software design and program quality.

HENRIK WIBERG

just been proposed at that time. This development has now led to a properly functioning World LHC Computing Grid (WLCG) with a disk storage space which currently is of the order of 20'000'000 Gigabytes (=20 Petabytes) and processing power corresponding to about 70 000 cores (a laptop usually contains one or two cores). Sweden contributes through SweGrid with the order of 4% of the total storage and processing capacity of the WLCG, a high fraction if compared with the fraction of Sweden's contributions in other global contexts. One explanation for Sweden's significant contribution to WLCG is the effective coordination of Swedish national data centers, of which NSC is one, to create SweGrid. This has required both a well functioning management structure for the coordination of the national computer and network resources, provided by SNIC, and the creation of an efficient grid system-software (middleware) ARC, which has been provided, and currently is developed and maintained, by the Swedish and Nordic High Energy Physics community. A further requirement has been the Nordic cooperation within the Nordic Data Grid Facility (NDGF) that is responsible for the main WLCG node (Tier1) that the Nordic data centers make up together. The author of this article regularly does shift work at CERN in the ATLAS experiment control room as surveyor of the WLCG operation, and can during these shifts directly observe how well the Swedish and Nordic data centers perform within the WLCG.

TORD EKELÖF

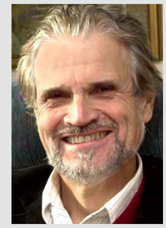

Professor Tord Ekelöf holds the Chair of Elementary Particle Physics at Uppsala University. In 2003 he initiated the creation of SweGrid. He currently works with the ATLAS experiment at the CERN LHC collider and chairs the Swedish LHC Consortium. He has earlier worked with the DELPHI experiment at the CERN LEP collider for which he, in particular, developed the Ring Imaging Cherenkov (RICH) detector. His main physics interest at LHC, as earlier also at LEP, is the search for the charged Higgs boson. He is also active in the development of new high energy particle accelerators and currently chairs the NorduCLIC and Nordic TIARA consortia.

Intel HPC training at NSC

On February 9 – 11, NSC hosted an Intel HPC training event held by the three Intel experts, Michael Klemm, Heinz Bast, and Per Hammarlund. The event was held at Quality Hotel Ekoxen in Linköping with roughly 65 participants from all over Sweden, as well as from a few places in both Denmark and Norway.

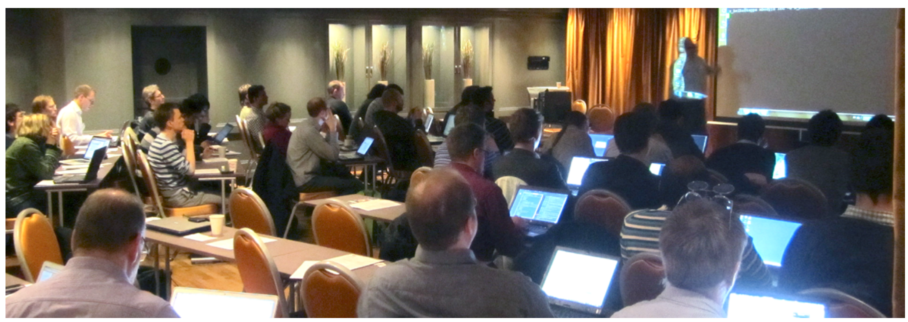

The training covered the latest Intel processor architecture and the various Intel tools relevant for HPC. The instructors, for example, gave an in-dept coverage of the Intel compilers and covered several other Intel tools, such as the Math Kernel Library, Intel MPI, the memory and thread checking tool, and performance analysis tools for both threaded and MPI parallel programming models.

All participants were offered extended Intel software trials for use on, for example, their personal laptop or by way of a temporary account on one of the NSC resources.

Overall the event received very good evaluations, but several participants did remark that the program was perhaps a little too intense and more or longer breaks would have been appreciated. In fact, we intentionally scheduled an intense program to cover as much of the relevant topics in this relatively short event, so feedback regarding the intense program was rather expected. Several participants also noted that a few hands-on sessions would have been very valuable.

Initially we did actually plan for a smaller event with a few hands-on sessions for approximately 30 participants, but during the registration phase for the event, we clearly saw that the event was in high demand among our users. Hence, we decided to change both the lecture room and the format of the event to accommodate as many prospective participants as possible. As we finally landed on a total of 65 participants, all plans of hands-on sessions had to be abandoned. Hopefully we will have a chance to revisit the original plans in a future NSC event with a narrower focus on one or a few of the topics.

NSC would like to thank Intel and especially the three instructors for a very educational and successful event.

TORBEN RASMUSSEN

Michael Klemm is part of the Software and Services Group, Developer Relations Division. His focus is on High Performance and Throughput Computing. Michael obtained an M.Sc. in Computer Science in 2003 and received a Doctor of Engineering degree (Dr.-Ing.) in Computer Science from the Friedrich-Alexander-University Erlangen-Nuremberg, Germany. His research focus was on compilers and runtime optimizations for distributed systems. Michael's areas of interest include compiler construction, design of programming languages, parallel programming, and performance analysis and tuning. Michael is an Intel delegate at the OpenMP Language Committee.

After studying Computer Science in Bonn, Germany, **Heinz Bast** worked in the context of the German HPC initiative SUPRENUM in a team developing SUPERB — the very first semi-automatic system to parallelize programs for MIMD-parallel execution. Later he maintained the Fortran compiler of this project. After joining Intel 1993, he spent 6 years with Intel's Super Computer System Division ('Intel iPSC and Intel Paragon'). He then moved to the Developer Relation Division to support all kind of independent software developers in the gaming, enterprise, server and HPC domain. Since 2006 he is working in Intel's Developer Product Division on software tools focusing on the compiler and performance tools.

Per Hammarlund received a PhD from the Royal Institute of Technology (KTH) in Stockholm, Sweden, in 1996. Per started studying computer science and engineering at KTH in 1985 (D85). Per joined Intel in 1997. He worked on the Willamette processor (Pentium 4), especially Hyper Threading and the memory system. He was one of the lead architects on the Nehalem processor (Core i7), responsible for the server segment. Per is the chief architect for an upcoming client processor. Big interests he has driven at Intel are improvements in performance modeling and micro architecture power efficiency. He holds 40 patents.

UPCOMING EVENTS

Advanced School on High Performance and Grid Computing
April 11 – 22, 2011, Trieste, Italy.
http://cdsagenda5.ictp.trieste.it/full_display.php?ida=a10135

IPDPS 2011; 25th IEEE International Parallel & Distributed Processing Symposium
May 16 – 20, 2011, Anchorage, Alaska, USA.
http://www.ipdps.org

CCGrid11; 11th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing
May 23 – 26, 2011, Newport Beach, CA, USA.
http://www.ics.uci.edu/~ccgrid11/

ICCS 2011: International Conference on Computational Science
June 1 – 3, 2011, Tsukuba, Japan.
http://www.iccs-meeting.org

The 26th NORDUnet Conference
June 7 – 9, 2011, Reykjavik, Iceland.
http://www.nordu.net/conference/ndn2011web/welcome.html

HPCS 2011; The 2011 International Conference on High Performance Computing & Simulation
July 4 – 8, 2011, Istanbul, Turkey.
http://hpcs11.cisedu.info/

ParCo2011; International Conference on Parallel Computing
August 30 – September 2, 2011, Ghent, Belgium.
http://parco2011.elis.ugent.be

Euro-Par 2011; Aspects of Parallel Computing and Distributing Computing
August 29 – September 2, 2011, Bordeaux, France.
http://europar2011.bordeaux.inria.fr

PSTI 2011; Second International Workshop on Parallel Software Tools and Tool Infrastructures
September 13 – 16, 2011, Taipei, Taiwan.
http://www.psti-workshop.org

ICPP2011; 40th International Conference on Parallel Processing
September 13 – 16, 2011, Taipei, Taiwan.
http://icpp2011.org

PaCT-2011; Parallel Computing Technologies-2011
September 19 – 23, 2011, Kazan, Russia.
http://ssd.sscc.ru/conference/pact2011

IEEE Cluster 2011
September 26 – 30, 2011, Austin, Texas, USA.
http://www.cluster2011.org

PACT; 20th International Conference on Parallel Architectures and Compilation Techniques
October 8 – 12, 2011, Galveston Island, Texas, USA.
http://www.pactconf.org

LISA'11; 25th Large Installation System Admininstration Conference
December 4 – 9, 2011, Boston, MA, USA.
http://www.usenix.org/event/lisa11

SC11; International Conference for High Performance Computing, Networkng, Storage and Analysis
November 12 – 18, 2011, Seattle, WA, USA.
http://sc11.supercomputing.org

CACS 2011; The 2011 2nd International Congress on Computer Applications and Computational Science
15 – 17 November 2011, Bali, Indonesia.
http://irast.net/conferences/CACS/2011

HiPC 2011; 18th IEEE International Conference on High Performance Computing
December 18 – 21, 2011, Bengaluru (Bangalore), India.
http://www.hipc.org.

Linköpings universitet

National Supercomputer Centre, Linköpings universitet, 581 83 Linköping, Sweden
Tel: +46 (0) 13-28 26 18, fax: +46 (0) 13-28 25 35, e-mail: info@nsc.liu.se
www.nsc.liu.se

Editor: Peter Münger To subscribe, for a paper version and/or PDF-file, send an email to info@nsc.liu.se