

## e-Science

### SeRC – the Swedish e-Science Research Centre

Dan Henningson at KTH is director of the Swedish e-Science Research Centre, SeRC, formed by the universities KTH, KI and SU in Stockholm and LiU in Linköping.

Read further on page 3

## Intel HPC training at NSC

NSC hosts an extensive 3-day HPC training class given by Intel experts. The class covers the latest Intel processor architecture, relevant Intel HPC tools, such as compilers, MKL, performance and profiling tools etc., application tuning and relevant parallel programming models.

9 -11 February 2011 in Linköping

<http://www.nsc.liu.se/intel-hpc-training>



## Successful Nordic collaborations

*In November, NDGF (Nordic DataGrid Facility) arranged a Strategic Workshop in Copenhagen. The timing was selected with respect to that NDGF will during 2011 transfer from the current second phase (2006 – 2010) into its third phase. A number of scientific areas were discussed for which a Nordic*

*collaboration would be beneficial. NDGF will in the future, besides successfully acting as the Nordic Tier-1 for CERN, also focus on coordination of Nordic efforts in EGI (European Grid Initiative) and support to other large projects on the Nordic level, such as research infrastructures, where a Nordic collaboration would provide added value. Another successful example of Nordic collaborations within NDGF is the development of large-scale distributed storage. Swestore – the Swedish national storage – utilises this technique, and Swestore is now increasingly used in a number of data-excessive sciences, providing easy access to data from all SNIC centres and also via web sites if desired.*

*Early December, SeRC (Swedish e-Science Research Centre) arranged a national application experts meeting, which was hosted by NSC. The meeting brought together application experts from all SNIC centres representing several scientific areas. The meeting was very successful, giving the experts opportunity to get to know each other better, which will lay the grounds for providing even better SNIC user support in the future by the excellent application experts. More about this meeting in a separate article in this issue.*

*SNIC was evaluated during autumn, and the report will soon be public. The review panel has done a thorough work interviewing a large number of people from SNIC centres, user communities, research infrastructure providers, e-science projects etc, and I am sure that the report will contain a lot of useful information that will be valuable for the future development of SNIC.*

*Finally, I would like to thank all NSC personnel and all our users at universities all over Sweden including our partners SMHI and Saab for a year full of events and positive development on the HPC side, and I am looking forward to an equally successful 2011.*

*Merry Christmas and Happy New Year 2011*

BENGT PERSSON, NSC DIRECTOR

## SC10 and HP-CAST 15 in New Orleans

This year the annual Supercomputing Conference (SC) visited New Orleans, Louisiana. SC is a fairly big event with around 10000 attendees from academia, government and business. This influx of (SC)nerds is never contained by the conference center but usually spills over into restaurants, hotels, river boats, etc. leaving its mark on the city.

Attending SC10 is tough. The conference (including everything that happens around it) doesn't stop for things such as week-ends or nights. Trying to take it all in can easily be overwhelming with 4000+ booths on the show floor, interesting talks, scheduled meetings, people to talk to, ... all this happening at once (it seems).

Technology-wise I think it's fair to say that this was a year of software (and of course, a year to deploy and stabilize last years bleeding edge hardware). Among the news: many new GPU related products, debugging at scales previously unheard of and PGAS support in the new Intel compilers (CAF). From a hardware perspective expect SC11 to be more exciting (new CPUs, faster networks, ...).

As is the tradition a new top500 list was released (36th ed.). For the first time China took the number one spot with Tianhe-1A a GPU powered beast that reached over 2.5 PF running linpack. Five systems from SNIC centers could be found on the list (two each from KTH/PDC and NSC, and one from HPC2N).

PETER KJELLSTRÖM



## Staff highlight

My first experience of high performance computing (HPC) was as a master student working on my thesis at CERN. There I developed and used a Monte Carlo code to calculate synchrotron radiation in order to optimize the design of the Large Electron Positron accelerator (LEP) i.e. the predecessor of the Large Hadron Collider (LHC).

After a couple of years with Feynman diagrams I returned to HPC. Following a brief stay at Bergen, Norway, I stayed a year at the Supercomputations Research Institute (SCRI) at Florida State University Tallahassee, USA. At the time they had the first G/8 supercomputer, capable of 10.8 GFLOPS world record in 1989. It was a vector system with liquid CPUs, run ETA's EOS operating system and was

# SeRC – the Swedish e-Science Research

SeRC – the Swedish e-Science Research Centre – is formed by the universities in Stockholm and Linköping (KTH, KI, SU and LiU) around the two largest high-performance computing centres in Sweden (NSC at LiU and PDC at KTH). This project will bring together a core of nationally leading IT research teams (tool makers) and leading scientists in selected strategic application areas (tool users) and will focus on three main objectives:

1. Formation of e-Science communities that connect application groups with relevant core e-Science groups and computer experts at PDC and NSC.
2. Research in core e-Science methods such as distributed resources, database technology, numerical analysis, visualization and interaction, mathematical modeling, and parallel algorithms – all focusing on problems critical for several e-Science communities.
3. Much closer collaboration between PDC and NSC, and a substantial increase in advanced support staff, which will turn the centers into comprehensive e-Science enablers.

SeRC will take a national responsibility in the e-Science area in terms of hosting a majority of the Swedish e-Science infrastructure through PDC and NSC (e.g., almost 80% of the total cores and capacity for academic research in Sweden). To ensure commercialization and non-academic use are dealt with at the highest decision making level, SeRC will have an advisory board that includes experts from the private and public sectors alongside international science and e-Science experts. SeRC will also feature industry representatives taking an active role in the e-Science communities. Thus

SeRC, as a Swedish e-Science Research Center, will constitute a leading visionary e-Science node with a national scope and strong international ties. The four partner universities are together committed to strengthening and shaping the emerging e-Science landscape in Sweden and to give research in this field clear priority in their strategic plans.

During 2010, SeRC has started seven e-Science communities in the areas of bioinformatics, complex diseases, electron structure, flow, particle simulation, visualization and numerical analysis. In 2011 additional communities within the areas of climate, databases, distributed resources, parallel algorithms will be set up.

The SeRC management group consists of Dan Henningson (director), Olof Runborg, and Anna Delin from KTH, in addition to Erwin Laure, PDC director, as adjunct member. From Linköping University Anders Ynnerman (co-director) and Bengt Persson are involved. Stockholms University is represented by Juni Palmgren and Erik Lindahl. Karolinska Institutet is represented by Bengt Persson and Juni Palmgren.

SeRC will support existing communities and the catalyzation of new communities that connect application-oriented groups with relevant core e-Science groups in the SeRC research platform and the SeRC services and infrastructure platform. Each of the supported communities will thus comprise computer experts, e-Science method developers and scientists from application areas who jointly run e-Science projects. There will not be a fixed set of communities, but they will be created dynamically as the research environments evolve. It is

crucially important that communities do not become isolated subcenters, but that they are effectively linked together. A non-exhaustive list of e-Science communities is given below:

**Fluids.** The Fluids community will include application groups from fluid mechanics, geophysical flow, aeronautics, and astrophysics, some of which already collaborate in the FLOW and BBCC centers. There is also a strong connection to the strategic areas of transport (laminar wings) and energy (wind energy). The development of efficient and accurate methods for the simulation of turbulent flow is a core activity. Major breakthroughs are expected via numerical simulations of turbulence, which is at present often only attainable in experimental studies.

**Climate and the Environment.** The climate modeling community will involve scientists from BBCC, SMHI, and FLOW. BBCC and SMHI are partners in MISTRA-SWECIA and use climate models as the research tool. On the international scene they are associated with the EC Earth project, aimed at the development of an earth system model for climate prediction. The e-Science components in climate modeling include the development of efficient coupling techniques between climate model components and numerical techniques for space-time discretization.

**Bioinformatics and Sequence Databases.** This community will build on the constellation of SBC (SU+KTH), CBR (SU), the possible new high-throughput genomics core facility in SciLifeLab (KTH+SU+KI), and additional participants from LiU. Bioinformatics is a strong strategic field in Sweden, with

as a graduate student at the IBM center in Computer Computational University (FSU) in delivered ETA10- which was the world and Nitrogen cooled programmed in

FORTRAN or assembler. Unfortunately the mean time between failures, mostly OS, was very short and internal check pointing was absolutely necessary in all production codes. I used to check-point every 10 minutes!

In 1997 I became one of eight HPDR-lecturers, appointed by the Council for High Performance Computing (HPDR), with a mission to promote the use of HPC. Based at Linköpings universitet I worked close to NSC and in 1999 I started as associate scientific

director on part time at NSC. My main task is to be a link between NSC and Swedish scientists. Part of that is to represent NSC in SNAC and handle applications for HPC resources at NSC. To predict future user needs and assist in decisions on new NSC systems is another task. And I am the editor for NSC News, NSC progress reports and SNIC progress reports. It has been a fantastic journey to follow how NSC has grown, in number of users, systems, system performance, storage, staff, etc, since I started.


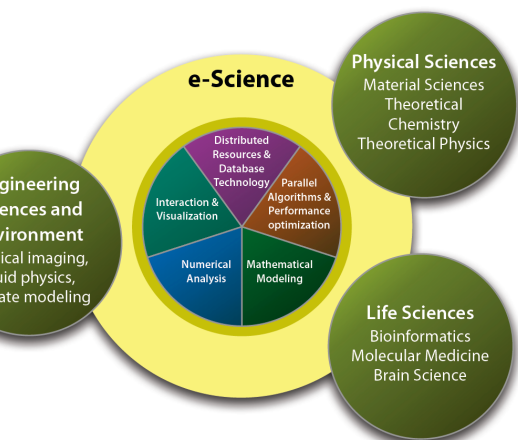
PETER MÜNGER

# h Centre

world-leading research on membrane protein structure prediction, maintenance of high-profile sequence databases (Pfam, InParanoid), and the Stockholm-Linköping groups have been driving the formation of the Swedish BILS node of the ELIXIR infrastructure for bioinformatics in Europe.


**Complex Diseases.** For the first time in medical history, the technical prerequisites exist for merging cascades of molecular data from biological samples, with images of the brain or other body parts, with longitudinal information on lifestyle and health, and with clinical data from hospital databases and from disease registers. High-throughput molecular technology allows mining a large number of genes, transcripts, metabolites or proteins in order to detect association with clinical outcomes. Adding the interplay with environmental and lifestyle factors allows new insight into the complex multi-factorial mechanisms underlying human health.

**Particle Simulation.** The core of this e-Science community is particle-based




Visualisation

modeling, using e.g., time-dependent molecular dynamics or Monte Carlo simulation to study extremely complex systems with methods based in statistical mechanics. Some of the currently most important areas are the field of simulations of biological macromolecules and materials modeling (with strong groups both at SU and KI), for instance to design new drug molecules to interact with specific proteins or understand complex dynamics on the molecular level. Advancing the state of the art of these modeling methods will also be crucial for several other areas such as condensed matter and nanophysics.



Molecular modeling



Climate

**Electron structure.** This community will focus on first-principle calculations based on, for example, density functional theory and Hartree – Fock methods, but also on multiscale simulation techniques. Groups involved in the community will be, for example, theoretical physics, theoretical chemistry, and materials science. More powerful models will have strong impact on applications such as materials science, nanoscience, molecular biosciences, energy, security, and cancer (e.g., two-photon sensitizers for photodynamic therapy), etc.

**Medical visualization.** The human visual sense is superior to today's computers in terms of perceiving content in images. Using vision, high bandwidth can be created between digital data representations and the user. It is this human capability that visualization builds upon by generating images representing the content of large and complex data sets. In image science the goal is to translate complex spatio-temporal patterns into forms that can be understood by humans, through a series of processing steps to extract descriptions of the objects embedded in the image or volume dataset. To achieve this goal there is a need to develop a language describing derived image content, understandable by visualization pipelines

## SLURM User Group 2010

SLURM is an open-source resource manager, developed and supported by Lawrence Livermore National Laboratory. It also include code contributions from sites all over the world, including NSC. At NSC all recent clusters run SLURM.

In October a user group meeting was held in Paris, France. The event was organized and sponsored by the French Atomic and Alternative Energies Commission, CEA.

Presentations covered experiences of SLURM from various sites, benchmark results, and new features being worked on. Also a road map for version 2.2 and 2.3 was presented and a Q&A session was held with the developers.

The meeting was an excellent opportunity to meet systems administrators as well as the developers. Several suggestions made by NSC during the meeting have been incorporated into SLURM 2.2 which is scheduled for release in December.

allowing them to act upon the metadata in the data stream. This will turn the visualization and Image processing system into an increasingly autonomous companion to the e-Scientist.

**Numerical analysis.** Development of numerical algorithms is critical in successful e-Science based research, often matching or outperforming improvements of hardware in terms of speed gains, in particular in areas at the forefront of research which have only recently become amenable to computer simulations. The numerical analysis group will work together with several other e-Science communities in their projects and also lead its own projects. The work will focus on general software and theory for first principle computational mathematical modeling; multiscale, multi-physics and stochastic problems, as well as application fields like turbulent and multiphase flow and high frequency wave propagation.

## Conclusion

SeRC will encompass many strong research environments in the Stockholm and Linköping regions, including ten strategic centers of excellence in areas where e-Science methods are of paramount importance. World leading scientific software with thousands of licensed users have been developed or developed in part within the consortium. One of the defining characteristics of leading e-Science is the vertical integration of cutting edge e-Infrastructure, advanced algorithms, and important scientific applications. The Stockholm–Linköping region has been highly successful in a number of such cases.

Through SeRC, the already existing close collaboration between PDC and NSC will be further developed which will have a large impact on the Swedish development in e-Science and on the dissemination of computational infrastructure to large new groups of researchers.

As e-Science is an emerging discipline there is a great need for competence building in both academia and society; this will be a major focus for SeRC that will be addressed at a national level. The comprehensive training and education programmes that SeRC is committed to delivering will underpin success in the long term by establishing e-Science infrastructures in broad scientific areas.

For more information, see <http://www.e-science.se/>

DAN HENNINGSON

Dan Henningson is a Professor in Fluid Mechanics at KTH, the head of the Department of Mechanics and the Director of the Swedish e-Science Research Centre (SeRC). His research group has made seminal contributions to the field of Computational Fluid Dynamics and performed some of the largest direct numerical simulations of turbulent flow to date.



## Application experts meeting in Linköping on December 2 – 3, 2010

On December 2 – 3, NSC hosted the first of hopefully many application expert meetings. There were a total of 18 participants and all six HPC centers in Sweden were represented. The purpose of the meeting was two-fold. First of all it was an opportunity to meet and get to know fellow application experts at other centers and learn about their background and expertise. It was also a first attempt to find and initiate concrete initiatives that application experts can

collaborate on for the ultimate benefit of all HPC users.

Among the things that we decided to initiate is the construction of a "knowledge base" for widely used scientific programs and tools. We also decided to make some behind the scenes changes to the handling of support cases that are application specific, rather than HPC system specific. The intention is that application specific support cases should, to the

extend possible, always be assigned to a person with some expertise on the application in question regardless of which HPC center the support case was submitted to and regardless of which center the expert is associated with.

The next application expert meeting has tentatively been scheduled for late April or early May, 2011.

TORBEN RASMUSSEN



# UPCOMING EVENTS

**Intel HPC Training at NSC processor architectures, compilers, libraries, debugging, tuning, MPI**  
February 9 – 11, 2011, Linköping, Sweden.  
<http://www.nsc.liu.se/intel-hpc-training>

**PPoPP 2011; 16th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming**  
February 12 – 16, 2011, San Antonio, TX, USA.  
<http://www.pppop.org>

**FAST'11; 9th USENIX Conference on File and Storage Technologies**  
February 15 – 18, 2011, San Jose, CA, USA.  
<http://www.usenix.org/events/fast11>

**Advanced School on High Performance and Grid Computing**  
April 11 – 22, 2011, Trieste, Italy.  
[http://cadsagenda5.ictp.trieste.it/full\\_display.php?id=a10135](http://cadsagenda5.ictp.trieste.it/full_display.php?id=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>

**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>

**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.sssc.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 Administration Conference**  
December 4 – 9, 2011, Boston, MA, USA.  
<http://www.usenix.org/event/lisa11>

**SC11; International Conference for High Performance Computing, Networking, Storage and Analysis**  
November 12 – 18, 2011, Seattle, WA, USA.  
<http://sc11.supercomputing.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](mailto:info@nsc.liu.se)  
[www.nsc.liu.se](http://www.nsc.liu.se)