



no.2 2014

PDC Center for High Performance Computing

PDC Newsletter

The Arrival of Beskow - Gert Svensson, PDC, page 4

PDC Business Development Update - Leif Nordlund, PDC, page 7

EUDAT's B2SHARE - Carl Johan Håkansson, PDC, page 8

PRACE Experts Help Optimize Swedish Scientific Software - Lilit Axner, PDC, page 9

Milner Marches On - Mikael Djurfeldt, PDC/INCF, page 10

PDC Summer School 2014 - Stefano Markidis, HPCViz, page 12

Update on Cloud Activities at PDC - Åke Edlund, PDC & HPCViz, page 13

KTH e-Innovation Center (KeIC) at PDC - Lilit Axner, PDC, page 15

Allinea DDT Course - Jonathan Vincent, PDC, and Patrick Wohlschlegel, Allinea, page 17

Introduction to PDC Systems Course - Henric Zazzi, PDC, page 17

NeIC News - Michaela Barth, PDC, page 18

PDC Experts Contribute to the PRACE SHAPE Programme - Lilit Axner, PDC, page 18

PRACE Best Practice Guide for Xeon Phi Processors - Michael Schliephake, PDC, page 20





Erwin Laure
Director PDC and HPCViz

**The PDC Newsletter is
published by PDC
at CSC, KTH.**

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

Editor: Erwin Laure

Editorial Staff: Genet Edmondson

Layout: Genet Edmondson

Graphic Design: Maria Malmqvist

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

ISSN 1401-9671

Cover

Beskow - PDC's new Cray XC40 supercomputer honours Swedish children's author and illustrator Elsa Beskow. Milner is in the back row.

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: +46 8 790 7800

Fax: +46 8 247 784

Editorial

On Wednesday the 22nd of October, PDC received delivery of its new flagship supercomputer – a 2 petaflops Cray XC40 system replacing our current 305 teraflops system, Lindgren. With Lindgren reaching the ripe old age (in supercomputer terms!) of four and a half years at the end of 2014, its lifetime is coming to an end, and consequently the Swedish National Infrastructure for Computing (SNIC) and KTH decided to invest in a new flagship resource at PDC. In a public tendering process, Cray submitted the winning bid offering a system with a theoretical peak performance of close to 2 petaflops, which is more than six times faster than Lindgren and which makes the new system the most powerful supercomputer in the Nordic countries and also the first petascale system in Sweden.

With this new flagship resource, computational research in Sweden is receiving a superb boost. Being available to Swedish researchers via SNIC, the new system will serve the most demanding applications from areas such as climate modelling, fluid dynamics, neuroscience, plasma physics, materials science and molecular simulation. The system will also be available to European researchers via the PRACE infrastructure.

Following an old PDC tradition, the new system is named after a Swedish writer – this time our choice was Elsa Beskow, who was both a famous author and illustrator of children's books. Some of Beskow's illustrations have been selected for the panel design on the system's cabinets (which you can see on the front cover).

During November the system underwent a rigorous testing procedure and the first pilot users were able to start using the system in the beginning of December. As of the 1st of January 2015, the system will become available to all Swedish researchers via SNIC. This means that the system will be in full production use only four months after the contract was signed – a record time!

Beskow will be officially inaugurated in a high profile ceremony taking place on the 27th of January 2015 at KTH, which will also include a scientific workshop where some of the scientific research that will benefit from being able to run on Beskow will be presented.

The total budget for Beskow, including four years of operations as well as supporting systems like storage and pre- and post-processing facilities, amounts to 170 million SEK and the funding to cover this has been contributed by SNIC, KTH and industrial collaborators. You can read more about Beskow's installation at PDC in this edition of the newsletter.

But Beskow is certainly not the only activity that has kept us busy over the past couple of months: our Neuroinformatics system “Milner” has been providing excellent services to the neuroscience research community, while our newly created business development unit has started to explore how PDC’s services can be more beneficial to Swedish industry, and we have developed the concept of the KTH e-Innovation Center, KeIC. Working with small and medium-sized enterprises (SMEs) was also piloted within the PRACE-based SME HPC Adoption Programme in Europe (known as SHAPE).

Together with PRACE, we worked on optimizing scientific software and published a best practice guide for Intel’s new MIC platform. Our work within EUDAT, particularly the B2SHARE service, as well as our cloud activities, have continued and the latter are now being reinforced with a NeIC-sponsored project named Glenna.

There were also other, non-technical, things happening related to NeIC and PDC’s Michaela Barth, who was coordinating NeIC’s “Generic Area” will soon start an extended leave of absence for a very good reason: she is expecting a child in early December. During her absence, PDC’s Dejan Vitlacil will take over her NeIC tasks.

And finally, PDC continued to offer a wide range of education and training events. So, it has been a busy autumn and Beskow will keep us fully occupied in the New Year too. But now it is time to break and wish all our users “Happy Holidays” as we look forward to great new scientific discoveries on Beskow and the other PDC services in 2015.

Erwin Laure, Director PDC and HPCViz

Facts about Beskow

- Cray XC40 with 1,676 compute nodes
- Each node has two 2.3 GHz Intel Xeon Haswell E5-2698v3 CPUs with 16 cores
- Total number of cores: 53,632
- Peak performance: 1,973 petaflops
- Linpack performance: 1,397 petaflops
- Total memory: 104.72 TB
(64 GB per node)
- Interconnect: Cray Aries
(Dragonfly topology)
- Lustre file system: 5 PB
- Rank in **Top500**: 32

In This Issue

Editorial

Erwin Laure 2

Staff Focus

Sara Berenji 8

The Arrival of Beskow

Gert Svensson 4

PDC Business Development Update

Leif Nordlund 7

EUDAT's B2SHARE

Carl Johan Håkansson 8

PRACE Experts Help Optimize Swedish Scientific Software

Lilit Axner 9

Milner Marches On

Mikael Djurfeldt 10

PDC Summer School 2014

Stefano Markidis 12

Update on Cloud Activities at PDC

Åke Edlund 13

KeIC at PDC

Lilit Axner 16

Efficient Debugging with Allinea DDT Course

Jonathan Vincent & Patrick Wohlschlegel 17

Introduction to PDC Systems Course

Henric Zazzi 17

NeIC News

Michaela Barth 18

PDC Experts Contribute to the PRACE SHAPE Programme

Lilit Axner 18

PRACE Best Practise Guide for Xeon Phi Processors

Michael Schliephake 20

PDC-Related Events

..... 20

HPC Sources

..... 20

The Arrival of Beskow

Gert Svensson, PDC

Procurement

The background work for procuring Beskow was actually initiated in 2013 when PDC started discussions with SNIC about replacing our current main system, Lindgren. We surveyed the market to find out what would be available in 2014. This included visits to the Supercomputing 2013 conference (SC13) in Denver and the International Supercomputing Conference (ISC'13) in Leipzig where we held non-disclosure discussions with all the major supercomputer vendors. We also surveyed the major users of Lindgren and asked what they needed in our next system. By the end of 2013 we had a pretty good idea what would be required and what capacity would be available for specific levels of funding. We also had a promise from the president of KTH that substantial co-funding would be provided if we got a positive decision from SNIC.

In the beginning of 2014 we worked with the procurement department at KTH to formulate the precise statement of requirements. We decided to ask for the largest possible system for a Total Cost of Ownership (TCO), including maintenance and estimated power and cooling costs, over four years. The capacity of the system was estimated with application benchmarks and datasets from different areas, which we selected together with major users of Lindgren and application experts. This resulted

in a formal invitation to tender that was divided into three lots. Lot 1 covered the main high-performance computing (HPC) system, lot 2 addressed the pre- and post-processing systems, and lot 3 was about storage for the parallel file system. The invitation allowed vendors to bid on one or several lots.

For the main system we required a system that, as a minimum, had the capacity to run the Linpack benchmark tests with a speed of at least 1.2 petaflops. We also required that the vendor supplying the system should have sold or manufactured major parts for a system that was in the 35th place or higher of the [Top500](#) list at the time. In assessing the submissions from potential vendors, significant weight was put on the application benchmarks representing typical cases and workloads from PDC users, with some weight being given to the synthetic benchmarks such as HPL and Stream.

On the 9th of May 2014, SNIC decided to fund the new system. Then on the 14th of May we received confirmation that KTH would contribute 60 million SEK, so we were able to send out the Invitation to tender the next day with a due date of the 10th of July. Cray submitted the winning bid where they offered two alternatives: a system with either 1,676 or 1,740 nodes each containing two of the new 24- or 16-core Haswell processors respectively, which would give a peak-performance of 1,737, respectively 1,973, petaflops. We soon found that Cray's alternative with the



Left: Beskow is water-cooled and requires a substantial amount of water. Here new pipes are being installed for Beskow under the old Cray, Lindgren.



Above: The basement of PDC was used as a workshop to produce the necessary cooling pipes.



Above: The power distribution extension was a time-critical component for the installation of the new Cray. Here it is arriving at PDC just two weeks ahead of the installation of Beskow.

Below: The new power distribution unit is being connected to the existing one. Luckily one of the electricians is not so tall.



Below: Only a couple of days before Beskow arrives everything is hectic. Each of Beskow's nine cabinets requires dual 100-ampere cables and connections to the cooling water.



Below: Three large trucks carrying Beskow arrived at KTH early in the morning of the 22nd of October.



16-core CPUs was more flexible, so we decided to accept this proposal.

Contract discussions with Cray started on the 21st of August and KTH signed the contract with Cray on the 8th of September. In our initial requirements, we had talked about an installation between January and March 2015 but, during the contract discussions, Cray offered to install the system in October 2015. As we had to do substantial work to prepare the physical infrastructure in the PDC computer hall prior to the installation, we were initially uncertain that we could be ready in time to do the installation in October. However a careful review of the time plan showed that it would indeed be possible, though highly hectic! To benefit our users, we agreed to the early installation. For the storage system, we chose a 5 PB system from DDN.

Infrastructure Preparations

Now the task at hand was to prepare the infrastructure for the new Cray system in two months, instead of five! The new system needed around 700 kW in power and cooling. Each of the nine cabinets in the system required dual 100-ampere three-phase connections and their own water connections. In principal we knew we had the capacity to satisfy these requirements however the connections were not in the right place. The first thing we noted was that we needed an additional power distribution unit but the delivery time for such equipment is normally at least six weeks. So our first action was to specify and order the power distribution unit. To be able to make the necessary changes quickly and using the same team we had

worked with during previous installations, we contracted the KTH landlord, Akademiska Hus, to perform the infrastructure upgrade for us. We asked Akademiska Hus to use the same consultants and contractors as when we installed Lindgren so they would be familiar with our computer hall and thus be able to complete the work rapidly.

After deciding where the new system would be placed in the computer hall, the next step was to move the equipment in that area, some of which was currently unused and some of which was actually in use.

From the beginning, we were concerned that we might have to partly shut down Lindgren when we started the new system. We were well aware that this would not be popular with the researchers using Lindgren, so we made a considerable effort to avoid this. On the power side of the problem, we had enough power available but did not have enough power distribution units of the right capacity. To purchase a power distribution unit for just a short period of time was too costly and there was nowhere to place one either. The solution we came up with was to rent one of the portable power distribution units that are used at construction sites and place it in the computer room while temporarily removing parts of one APC enclosure and using its power feed.

With Lindgren, we had developed and installed a system to capture the heat from the hot air exhaust of the computer. The new computer was a water-cooled system from the start, and used the same principle to capture the heat internally. Thus the cooling for Beskow could use the same

Below: The first crate is going into PDC. The person to the right is working for the transport company and is documenting that everything is undamaged and handled in the correct way.



Below: The first cabinet is being installed in the computer room. Cray uses lifting devices which are specially made to fit the Cray cabinets.



Below: The cables for the internal network are being installed. Copper cables are to connect pairs of cabinets and fibre connections are then used to connect the pairs to each other.



water circuit as Lindgren but the capacity of that was far from sufficient to cool both systems. However we had spare cooling capacity from the unused APC enclosures. This circuit also had the necessary back-up cooling but the temperature of it was considerable lower than what we required for the Cray systems. The problem with the lower temperature was that it could lead to condensation in the coolers... and electronic equipment and condensation are not a good mix. We decided to temporarily decrease the humidity in the computer room to decrease the dew point. Lindgren is actually cooled by four heat exchangers but due to space and cost restrictions, it was only possible to connect two out of the four heat exchangers to the cooling circuit from the APC enclosure. By utilizing the cooler water in that circuit, it was possible to have two of the heat exchangers cool the air going out of the room to a temperature considerably lower than the normal 20 °C and thus we could turn off the other two heat exchangers and just let the warm air out in the computer hall. We hoped that the cool air and the warm air would mix fast enough not to disturb any other equipment. It almost worked like that, and with the help of some guides made of plywood, we were able to obtain quite a good temperature distribution.

The last weeks before the arrival of the hardware and the installation team from Cray were really hectic. The power distribution unit arrived after a two-day delay on a truck from southern Sweden.

Although many things were prepared and planned long in advance, a lot of things happened during that last week. Thanks to heroic efforts by the teams of contractors and Johnny Eriksson from PDC who led the construction team in the computer hall, everything was in order the day before the installation.

Arrival and Installation

Three large trucks arrived early in the morning at the same time as a large installation team from Cray. This team was highly efficient and by the evening of the first day, all the cabinets had been placed in the correct positions in the computer hall. The next day, all the power and water connections were completed. Over the next few days, all the interconnecting cables running between the nodes and all the fibre cables connecting the far away cabinets were attached. After a week, all the hardware was in place and tested by running rigorous tests, so the installation team went home. Next the installation of the software and the acceptance tests started. At the time of writing, the first part of the acceptance tests (where the performance is demonstrated by running the benchmarks from the Invitation to tender) is almost complete. The next phase of acceptance testing will last one month during which test users will run their research programmes on the system while we monitor the availability of the system to check that it performs satisfactorily.

Getting Time on Beskow

The plan is that Beskow will be fully operational on the 1st of January 2015. The first round of large time allocations on Beskow was opened in September and is currently being evaluated by the Swedish National Allocation Committee (SNAC). It is expected that the next round of large time allocations for the second half of 2015 will be opened at end of March 2015. Medium-size time allocations for Beskow will be opened in January 2015. For more information, please see:

<http://www.snic.vr.se/apply-for-resources>.

PDC Business Development Update

Leif Nordlund, PDC

During the summer and autumn of 2014, the Business Development Unit at PDC has been focused on three projects: the KTH e-Innovation Center or KeIC (which you can read about in the article by Lilit Axner on page 16), KTH strategic partner mapping with regards to high-performance computing (HPC), and infrastructure development for industry needs. In the latter project, a number of areas have been defined where we are discussing how to best meet the legal and external requirements for data security and ease-of-use for industrial users.

Strategic Partner Mapping

KTH has a number of strategic partners from Swedish industry, where PDC, with its HPC capabilities is an important facilitator for joint research projects. The use of simulations and HPC computational resources within these companies varies depending on the industry and the company profile. Together with the KTH Business Liaison department and the Swedish e-Science Research Centre (SeRC), we aim to assist future e-science projects relating to Swedish industrial research and innovation. Mapping the computational needs of companies to the different areas of science research at

KTH and PDC will make it easier for companies to find and use the resources that KTH and PDC can offer.

KTH PDC Collaboration

To further strengthen the capabilities of PDC, we are also forming collaborative partnerships with complementary organisations. The Swedish Institute of Computer Science (SICS Swedish ICT) is one of PDC's close partners and we are discussing how to best collaborate on data analytics and cloud platforms. We have been monitoring how our European counterparts are developing their collaboration with industry, and this summer we visited the "Industry Innovation through HPC" track at the International Supercomputing Conference (ISC'14) in Leipzig to learn more about this. A lot of work is being done in the manufacturing and automotive sectors where the leading independent software vendors (ISVs) are streamlining the process from innovation to production within large industries. It is clear that the benefits for small and medium-sized enterprises (SMEs) of using HPC within their design and development processes can be huge – for example, producing better quality products, enjoying a higher return on investment (ROI), avoiding product failure early in the design process, and shortening time to market. However we still face the challenge of introducing this particular category of user to HPC throughout Europe; at present every European country is looking into how to do this effectively.

On the 2nd and 3rd of October, a representative from the PDC Business Unit gave a lecture at the 5th International Industrial Supercomputing Workshop in Tokyo where leading global supercomputing centres shared their strategic plans, success stories, and technical issues related to industrial supercomputing. Later this autumn we are planning a visit to the new Digital Manufacturing and Design Innovation Institute for the USA in Chicago. This new centre plans to serve as a demonstration showcase for the latest innovations in advanced digital manufacturing, research and development, and to facilitate learning and collaboration among its industrial partners.

Overall it has been a rather busy phase for PDC's industry-related development work this autumn.

Staff Focus



Sarah Berenji

Sarah joined PDC at the beginning of 2014 as a Python developer in the EUDAT B2Share project. Now her main responsibility is developing REST API features for B2Share. After a while, Sarah also started working in the Support group to help researchers use the PDC supercomputing resources. Before that, Sarah worked in Iran's operating system research centre for seven years as a researcher. During that time, she was working on the design and implementation of a vulnerability detection tool for C developers; it was a static tool based on compiler-related and syntax-oriented detection approaches. This led her to start working on the area of security in compilers and designing a new JVM. After leaving there, she became a MySQL administrator in a telecom company for about two years.

In her spare time, Sarah likes to read books, walk, or watch movies. Skiing is her favourite sport and she does like travelling around the world.

Free AVL Fire License

If you are interested in running AVL FIRE at PDC, please request a free license from support@pdc.kth.se.



B2SHARE

EUDAT's B2SHARE Store and Share Research Data

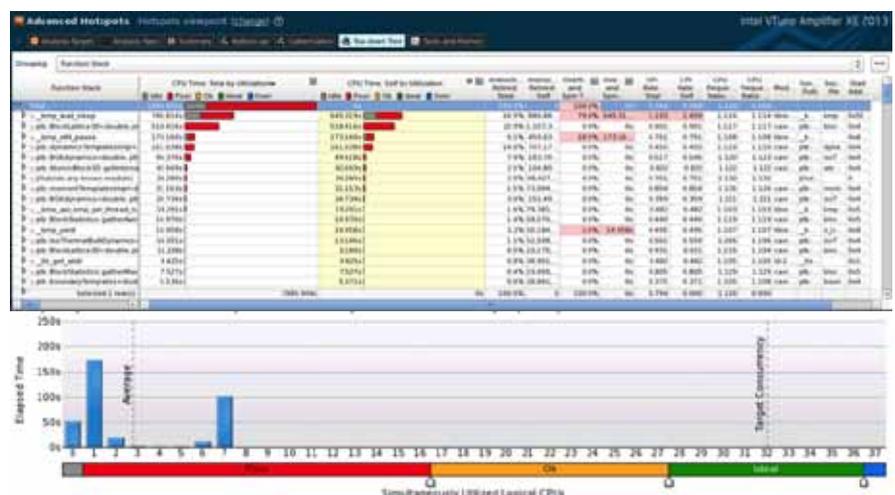
Carl Johan Håkansson, PDC

As a part of EUDAT's development of a common European e-infrastructure for research data, PDC is leading the development of B2SHARE – EUDAT's service for storing and sharing research data, particularly small-scale data.

The B2SHARE service was deployed in pre-production in January this year, and version 1.0 was released as the first stable version in time for EUDAT's third conference in late September. In the beta state, the B2SHARE service already supported persistent identifiers for tracing research data, plus integration with both the EUDAT metadata catalogue and the search engine B2FIND. Further developments are currently underway: among the improvements that EUDAT is working on are a REST API for data and metadata retrieval and deposition, B2ACCESS - EUDAT's service for federated authentication and authorisation, and integrating B2SHARE with EUDAT's other services B2SAFE (safe storage and replication) and B2DROP (personal cloud for data synchronisation and exchange).

Collaboration with the research communities has continued, and in the near future we will also see some more communities from the Nordic countries entering B2SHARE as a result of NEIC's project, B2SHARE Nordic.

The full version of B2SHARE is available at <https://b2share.eudat.eu>. If you would just like to try it out, a training instance is available here: <https://trng-b2share.eudat.eu>. For more information about EUDAT's services in general, see <http://www.eudat.eu>.



Above: Hotspot (top) and concurrency (bottom) analysis of the original SLILAB code

PRACE Experts Help Optimize Swedish Scientific Software

Lilit Axner, PDC

DECI (Distributed European Computing Initiative) is a European single-project high performance computing (HPC) access scheme which is currently supported by PRACE. Under the PRACE projects, DECI has been integrated into the PRACE ecosystem of HPC facilities and services. In order to provide seamless access to users, PRACE took over the organisation, operation and dissemination of DECI scientific results. You can read more about this at: <http://www.prace-ri.eu/deci-12-call>.

DECI enables European researchers to obtain access to the most powerful national (Tier-1) computing resources in Europe regardless of the researcher's country of origin or employment. Proposals for DECI access must deal with complex, demanding, innovative simulations that would not be possible without Tier-1 access. In addition to offering access to computing resources, application-enabling and porting, assistance from experts at the leading European HPC centres is also offered (on request and depending on availability) by DECI to enable projects to be run on the most appropriate Tier-1 platforms in PRACE.

Researchers both from academia and industry are eligible to submit proposals to DECI, as long as the project leader is undertaking non-proprietary research in a European country (which, for these purposes, includes European Union (EU) countries, EU candidate countries, EU associated countries and PRACE member countries). Project leaders will typically be employed in academic or industrial research organizations. Individual HPC centres may have further restrictions on who is eligible to use their machines, for example, due to US export rules.

One of the most recent successful DECI projects from Sweden was called DNSTF - Direct numerical simulation of finite-size fibres in turbulent flow. The goal of this project was to adapt a numerical model of the dynamic behaviour of finite-size fibres

in high Reynolds number turbulent flow. The chosen approach was to start from direct numerical simulations for turbulent flow and to resolve the flow around individual fibres. The turbulent flow was modelled with an entropy lattice Boltzmann method and the interaction between the fibres and carrier fluid was modelled using an external boundary force method (EBF). The project was led by Prof. Gustav Amberg and Dr. Minh Do-Quang from the KTH Department of Mechanics. A code named SLILAB (which was developed at the same department) was used in the project. The project received 8,437,500 DECI standard core hours on the Tier-1 machine at the Edinburgh Parallel Computing Centre (EPCC), UK, for a period of one year from May 2013 to May 2014 (which is equivalent to 6,750,000 core hours on a CRAY XE6 with 12 cores at 2.1 GHz).

The project also requested three months of application expert help from PRACE to assist in the OpenMP parallelization of SLILAB. The help was provided by the application expert Dr. Evghenii Gaburov from the PRACE partner SURFsara (the Dutch national HPC and e-science support centre) in collaboration with Dr. Lilit Axner from PDC. The work was done on parallelization of the SLILAB code with OpenMP for a shared-memory execution model when focusing on multiphase flow simulations, such as fibre suspensions in turbulent channel flows. The addition of OpenMP parallelization made it possible to minimize the number of MPI ranks in favour of a single-node parallelism, thereby mitigating MPI imbalance. With OpenMP parallelism in place, the code performance was also analysed on Intel XeonPhi. The figures on the previous pages show analyses of the original SLILAB code.

As a result of the work done in this project, the OpenMP version of the code was optimized substantially and will be used as the basis for further improvements that will be implemented in the next DECI round (as the DNSTF group was given approval to continue the work within the recent DECI12 round of applications). You can read more about this project at <http://www.prace-project.eu/IMG/pdf/wp178.pdf>.

Milner Marches On

Mikael Djurfeldt, PDC/INCF

Our Cray system Milner, which is named in recognition of significant work in neuropsychology that was done by [Brenda Milner](#) and her then husband Peter in the 1950s, is part of a high-performance computing platform for research in neuroinformatics and was funded by an infrastructure grant from the Swedish Research Council (VR). It is a Cray XC30 system, based on 2.5 GHz Intel Ivy Bridge 10-core processors and Cray's Aries interconnect. The aggregate peak performance is equal to 48 teraflops and the aggregate compute memory is equal to 3.75 TB. The system also includes a Lustre file system with a usable capacity of more than 150 TB.

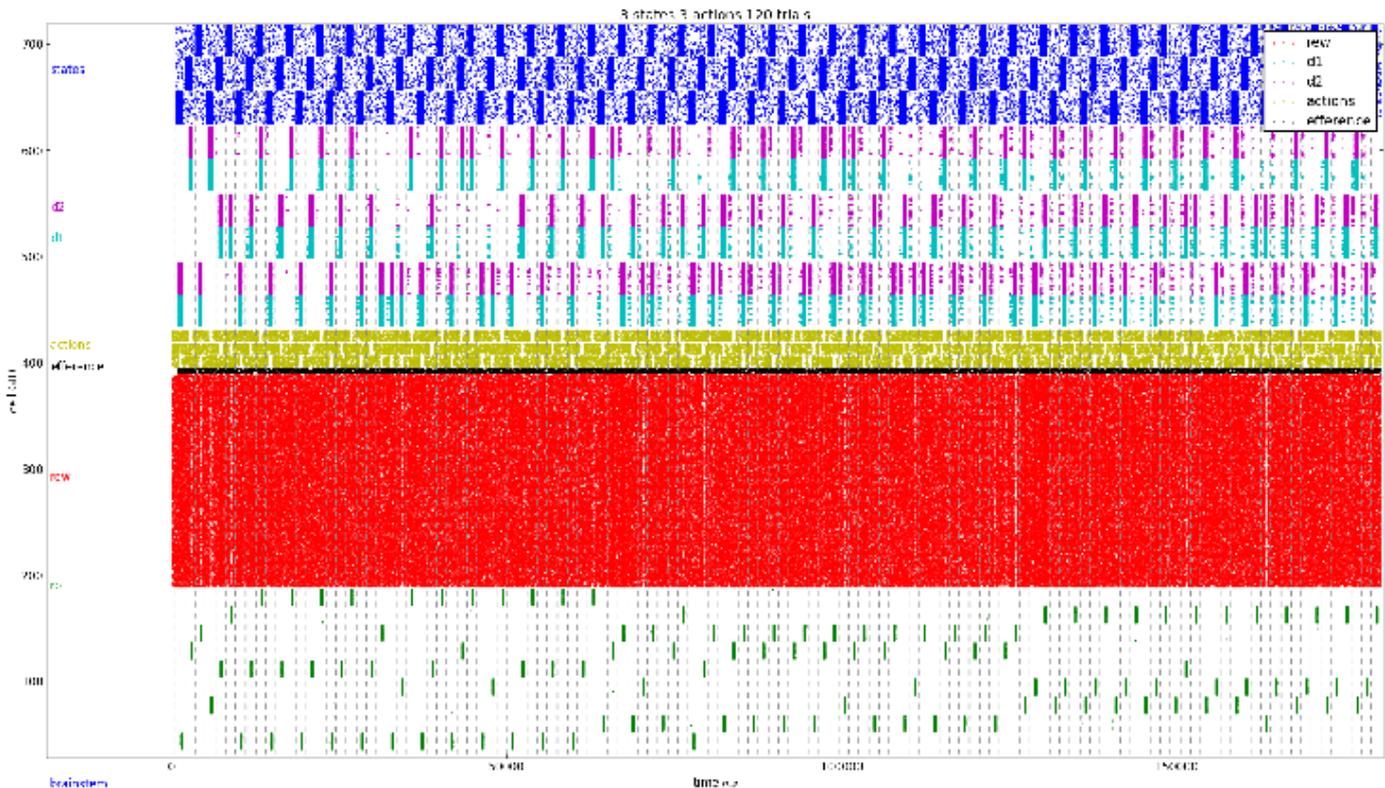
During the initial testing phase, the system was mainly used by members of the Department of Computational Biology (CB) in the School of Computer Science and Communication at the KTH Royal Institute of Technology. Other Swedish researchers in neuroinformatics are also invited to use the resource – if you are interested in doing

so, please apply for access to Milner here: <https://www.pdc.kth.se/support/accounts/user>. Efforts are underway to contact and inform potential users. In addition, neuroinformatics researchers based in other countries can obtain access to Milner through the International Neuroinformatics Coordination Facility (INCF) – anyone who is interested in doing this should contact Mikael Djurfeldt (djurfeldt@incf.org).

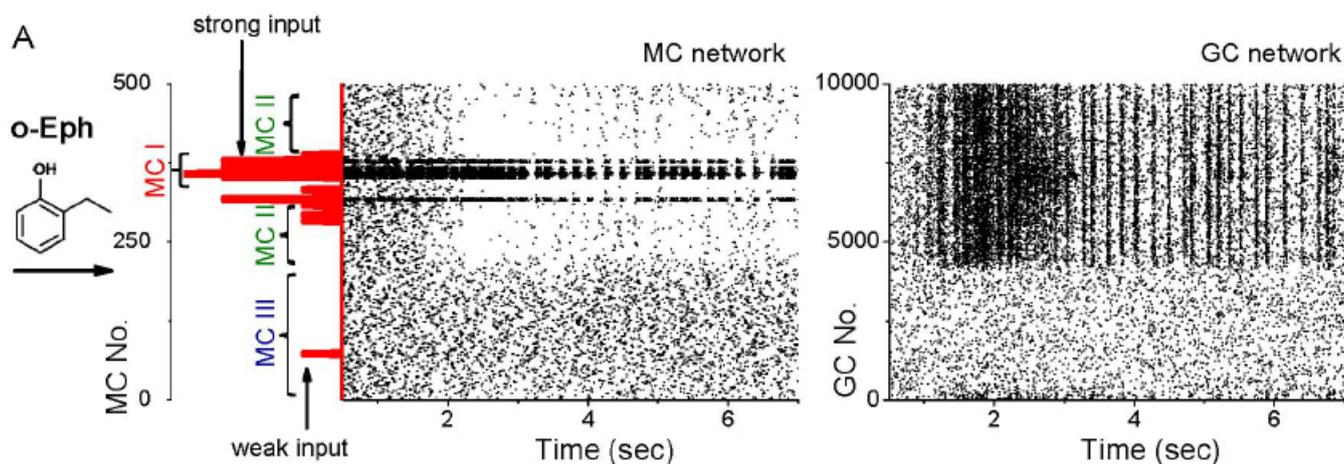
We are pleased to report that the Milner system has been working satisfactorily and the computational neuroscientists who have tested the machine have appreciated the specialized queue policy, which gives them access to large parts of the machine for shorter amounts of time.

Current work on Milner can be broadly divided into three categories:

1. abstract brain network models based on the Bayesian Confidence Propagation Neural Network (BCPNN) formalism, developed by Lansner and Ekeberg,
2. biophysical brain network models, and
3. neuroinformatics tools.



Above: A raster plot of activity in different structures in Berthet's model of action selection in an example simulation of 180 seconds of biological time with three states and three actions



Above: Different activity states of a simulated olfactory bulb following the application of an odour

Abstract Brain Network Models

This work uses the BCPNN model and has the dual, overarching aim of increasing our understanding of brain function and developing new computing technology. The figure on the previous page shows one example of work in this category. It is a raster plot of spikes from a simplified network model of action selection in the basal ganglia developed by Pierre Berthet, from Prof. Anders Lansner's group at CB. Bernhard Kaplan, also of the Lansner Lab, CB, has been using Milner to simulate learning in a closed perception-action loop in the oculomotor system, using a reward-based paradigm. In addition he is investigating self-organization of lateral connectivity in visual areas involved in motion-based perception. Phil Tully, who is also in the Lansner group, CB, is studying how sequential activity patterns can be learnt in a BCPNN network.

Biophysical Brain Network Models

This work uses spiking neuronal network models with varying degrees of biophysical detail, again with the purpose of studying brain function. Florian Fiebig, from the Lansner Lab, is working on a hypothesis of long-term memory formation based on the Markram-Tsodyks model of synaptic plasticity. David Silverstein, who is a doctoral student at the KTH Department of Computational Biology and who has been an application expert at PDC, is working with Martin Rehn (Lansner Lab) on a biophysical model of the primary visual cortical

area, which includes feature detection and motion. David is also working on the storage capacity of attractor memories in a biophysical model of the association cortex. Yuguo Yu and Michael Hines from Yale University were granted access to Milner through INCF. They are simulating a large-scale model of the olfactory bulb (see figure above).

Neuroinformatics Tools

In the area of Neuroinformatics tools, Pawel Herman, CB, is using Milner in his work on an innovative tool where brain-inspired connectionist systems are used for the analysis of large electroencephalography (EEG) and magnetoencephalography (MEG) data sets. Ekaterina Brocke, CB, and Mikael Djurfeldt, PDC, are using Milner to develop a tool, MUSIC, for connecting neuroscientific simulation tools into co-simulations (where multiple tools exchange data at run-time).



Above: David Silverstein and others at Allinea DDT course

PDC Summer School 2014

Stefano Markidis, HPCViz

The PDC “Introduction to High-Performance Computing” Summer School runs every year for two weeks at the end of summer. It is held at the main campus of the KTH Royal Institute of Technology in Stockholm, and this year it went from 18-29 August.

There were 56 students at the school this year with Swedish attendees from KTH, Linköping University, Stockholm University, Lund University, Uppsala University and the Karolinska Institute. Participants came from as far afield as India and Brazil, with Turkey, France, Belgium, Latvia, Spain, Austria, Germany, Kazakhstan, the UK, Italy and Switzerland all being represented!

The course includes explanatory lectures combined with practical lab sessions so students can put what they learn into practise. If you would like to learn more about programming supercomputers, please attend the school next summer! Details will be available on the Events section of the PDC website in March next year <https://www.pdc.kth.se/events>

The course will be held in the last two weeks of August.



Update on Cloud Activities at PDC

Åke Edlund, PDC and HPCViz

SNIC Cloud Infrastructure Project

The SNIC Cloud Infrastructure project is approaching the end of its first year. During 2014, we have been gathering user requirements from cloud user projects at PDC and three other Swedish high performance computing centres, HPC2N, C3SE and UPPMAX. In parallel with this, the SNIC Cloud Infrastructure project has been evaluating candidate cloud middleware for the prototype (at HPC2N). The number of users evaluating the cloud environments has increased (both terms of the number of people using the cloud and in the number of areas for which the cloud is being used), which is resulting in valuable new input on what is needed from a production environment. At PDC, one particularly interesting project is being run on the cloud by Arne Elofsson's group at SciLifeLab. From the PDC Cloud they can create, and administer, their software pipeline for consensus prediction of membrane protein topology or ranking protein models by assessing their quality. The pipeline consists of a web front-end and several bioinformatics software tools, such as blast, modhmm, gnuplot, and topology predictor amongst others. The advantage for this project

in using the cloud comes from the uneven usage they experience with their TOPCONS/PCONS pipeline - ranging from a couple of runs per week to around 1,000 runs per week. With the cloud's ability to dynamically spawn new instances, we will be able to handle the changes in load for the calculations and therefore use the hardware optimally. As a first step, TOPCONS is being used to determine the feasibility of the project, but the main target of the research is PCONS/PCONSC, with increased hardware requirements.

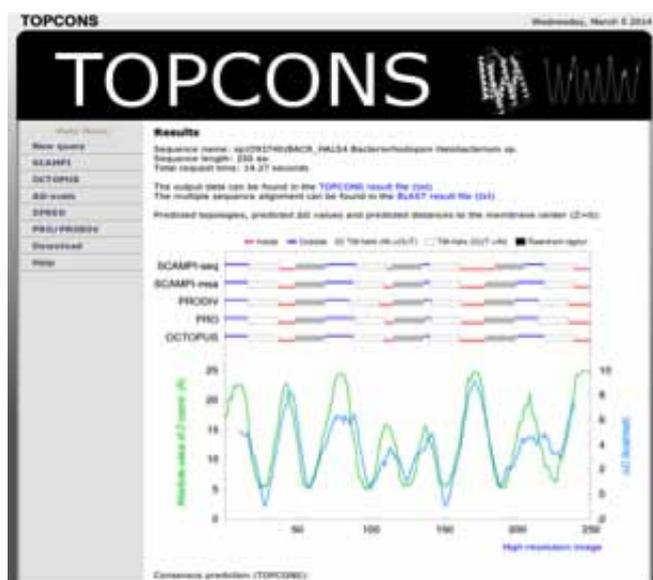
The requirements from the user groups guide the SNIC Cloud Infrastructure project when it comes to technical decisions. Research projects that join the infrastructure project either run directly on the SNIC Cloud, or use their own cloud resources in combination with the SNIC Cloud. Cloud resources that are added to the SNIC infrastructure need to conform to the technical specifications decided within the SNIC Cloud project. Completed projects will provide information regarding their resource use (including, for example, the number and types of virtual machines used, the type and format of metadata that was used, as well as details of the application scripts, package recipes, and project documentation that were used) thereby ensuring that the services can be reused for later projects

PDC is running the Galaxy platform on the SNIC cloud. Current user groups that are involved in the project cover the following research areas:

- forest remote sensing and planning (SLU, Umeå),
- bioinformatics (SciLifeLab, BILS, UPPNEX), Systems and synthetic biology (SYSBIO, Chalmers), and
- protein prediction pipeline (SciLifeLab, as mentioned earlier).

In addition we are working with groups from these areas:

- neuroscience (KI),
- proteomics – using Galaxy (SciLifeLab),
- fusion (Chalmers),
- traffic analysis (KTH),
- variant databases (SciLifeLab Uppsala),



Above: TOPCONS at PDC Cloud

- Chipster (KI, UPPMAX),
- CellProfiler image analysis (UPPMAX),
- MassSpectrometry platform (SciLifeLab Uppsala), and
- INCF federated data development (SciLifeLab, Uppsala).

The long-term goal with the SNIC Cloud Infrastructure project is to create a sustainable and generic SNIC Cloud infrastructure (Infrastructure as a Service, IaaS) to form a basis for future SNIC Cloud projects and to provide the necessary structure for running project specific Platform as a Service (PaaS) services. The SNIC Cloud infrastructure will provide an elastic resource pool, where additional resources (in the form of virtual machine hosts) can dynamically join and leave the core resource.

The project started on the 1st of January 2014 and runs in two 12-month phases, aiming at a scale up of the cloud infrastructure in 2016. The partners in the project are from PDC, HPC2N, C3SE and UPPMAX, and the project is being coordinated by Åke Edlund from HPCViz while the LifeScience project contact is Henric Zazzi from PDC.

NeIC Nordic Cloud Project “Glenna”

The goal of the “Glenna” project is to share knowledge and establish best practices for managing cloud services, as well as to create a Nordic federated cloud service, driven by the needs of Nordic researchers. The Nordic Cloud aims to do the following:

- enable Nordic researchers to run their experiments and manage their data on any Nordic Cloud resource,
- make it easy to access the Nordic Cloud through the KALMAR2 federated trust service,
- share technologies to improve the quality and security of cloud services,
- share user experiences on a Nordic level to improve quality and to increase the available set of services,
- share cloud administrative work and thus im-

- prove the service availability to the users,
- make it possible for researchers to share data (and thus facilitate new research) and increase the overall availability and security of data (thereby avoiding data loss),
- create a resource-sharing solution to simplify the usage and sharing of Nordic Cloud resources, and
- establish a means of billing and accounting within the Nordic Cloud to create a fair sharing of resources and funding. (This will also simplify matters in terms of accounting for external funding of research.)

The main project contacts are Dan Still from CSC (who is the project leader) and Michaela Barth and Dejan Vitlacil (from NeIC). The SNIC contact people in relation to the project are Åke Edlund from HPCViz (who is a member of the Glenna steering board and who distributes the work within the SNIC Cloud) and Björn Torkelsson from HPC2N (who is the main SNIC Cloud technical contact and who attends the technical meetings).

This two-year project started on the 1st of September this year, with partners from all Nordic countries. The SNIC Cloud participation in the project adds up to a total of 30 person months.

PDC Staff - June 2014

KTH Courtyard



KTH e-Innovation Center (KeIC) at PDC

Lilit Axner, PDC

Academic supercomputing centres are increasingly being recognized as a valuable resource not only for academia but also for industry. The use of advanced computing technologies as provided by such centres helps industry to shorten the production cycle, produce better products, and stay competitive in the global race. Thereby, supercomputing is directly contributing to the economic growth of a region.

Increasingly, the commercial success of companies and their capacity for innovation is dependent on their ability to perform large-scale simulations and to analyse large volumes of data on large-scale computing and storage systems. This allows for increased innovation cycles, better products targeting their market segment, reduced prototyping costs, and overall reduced time to market. “To out-compute is to out-compete” is the motto. This is particularly true for a number of industry segments of particular importance to the Stockholm region, like design and manufacturing, the automotive and aerospace industries, the pharmaceutical industry, life sciences and the energy sector, but increasingly also in the services sector (including the financial and insurance sector as well as the digital media industry).

Key shortcomings are still however the difficulties companies face in having computing and storage capacities of sufficient scale at their disposal and having the necessary expertise to use those infrastructures. In addition, many companies are not yet aware of the full potential such infrastructures can provide to their business. Commercial cloud computing offerings like those provided by companies such as Amazon and Google are only a partial answer to the problem, as their infrastructure offerings are often not suitable, prohibitively expensive, and involve as yet unsolved trust and privacy issues. In addition, the lack of necessary expertise problem remains open.

To overcome these problems, many countries, such as Scotland, Germany, the Netherlands and

the USA, are putting special measures in place to focus on supercomputing efforts for industries.

For instance, Scottish businesses in the energy, life science and finance sectors are set to benefit from a unique initiative that will enable them to channel the UK’s world-class supercomputing power and expertise. Scottish Enterprise and EPCC – the supercomputing centre at The University of Edinburgh – have joined forces to create Supercomputing Scotland. This 3-year, £1.2 million programme will enable Scottish companies of all sizes to use the UK’s most powerful supercomputer to solve key business challenges such as product and service design.

In Germany, the company Sicos BW (Simulation, COmputing, Storage) was founded in 2011 at the Karlsruhe Institute of Technology (KIT) and the University of Stuttgart. Their mission is to facilitate simple access to the best resources in high-performance computing (HPC) and large data facilities for companies and institutions and to support their usage of these systems. To accomplish this task, a close partnership between research and industry has been established.

SURFsara, the Dutch national government-financed supercomputing centre, and Vancis, an advanced ICT-services provider in the Netherlands, signed a technology transfer contract in 2013 to jointly establish projects to share their existing and future knowledge. This non-exclusive partnership with Vancis allows SURFsara to contribute with specific know-how and expertise to the growing needs of industry, small and medium-sized enterprises (SMEs) and the (semi-)public sector on HPC and big data services.

The same trend can be seen in other highly developed countries as well. Countries like the USA, China, Russia and Japan are putting infrastructures into place to more widely apply supercomputing across many different industries and segments.

In the USA the federal government has recently made a significant investment at the University of Illinois at Urbana-Champaign for the establishment of the “Digital Lab for

Manufacturing”. It is an applied research institute that will develop digital manufacturing technologies and commercialize these technologies with key industries. The lab has been launched with a contribution of \$70 million from the U.S. Department of Defense and more than \$250 million in additional funding will come from industrial, academic, government and community partners.

The Stockholm region is one of the most dynamic regions in innovation with an ever-growing knowledge-intensive labour market both in industry and academia. Research and innovation are at the core of it. In this context the KTH Royal Institute of Technology (KTH) plays an indispensable role that stretches from the practical education of future technicians to the most advanced education of researchers with strong connections to large industries and SMEs as well as the public sector. Strategic collaboration with large industries, such as ABB, Ericsson, Sandvik, Scania, and Vattenfall, as well as inter-university collaboration with places such as the Karolinska Institute and Stockholm University, has a long history at KTH.

A regional competence centre, providing large-scale computing and storage services together with the necessary expertise, will be a pivotal factor in the competitiveness of the Stockholm region as it will provide the regional economy with the required resources, as well as training and educating a highly skilled workforce, and facilitating the transfer of technology from academia to industry, both in relation to computing and the research relevant for the commercial sector.

Recently KTH-PDC applied for a pre-study grant from the Swedish Agency for Economic and Regional Growth. The agency has a government mandate to manage and distribute funding from the European Regional Development Fund in such a way as to support projects that promote growth and jobs. The goal of the KTH-PDC pre-study application was to examine the needs and prerequisites for establishing a KTH e-Innovation Center (KeIC) that would form a regional

competence centre for strategic collaboration with large industries and SMEs by providing HPC services.

The application was successful and the KeIC project received funding to conduct the pre-study by the end of December 2014. The application received support from several companies and organizations such as Saab, Scania, Sandvik, Ericsson, the SICS Swedish Institute of Computer Science and the Swedish Defence Research Agency (FOI). In total there were 22 pre-study applications of which 21 were successful.

Currently Lilit Axner and Leif Nordlund, the coordinators of the PDC Business Unit, together with the PDC director, Erwin Laure, are carrying out the KeIC pre-study in close collaboration with several other projects that also received funding from the Swedish Agency for Economic and Regional Growth, as well as with large companies and SMEs in the Stockholm region. The final outcome of the pre-study will be a solid business plan that is intended to serve as the basis of the main KeIC project to come.

If you have any questions about KeIC or are interested in the future activities of the centre, please contact Lilit Axner (lilit@pdc.kth.se), Leif Nordlund (leifnor@pdc.kth.se), or Erwin Laure (erwinl@pdc.kth.se).



Above: Participants and presenter from the Efficient Debugging with Allinea DDT course outside PDC building at KTH

Efficient Debugging with Allinea DDT Course

Jonathan Vincent, PDC and
Patrick Wohlschlegel, Allinea

On the 25th of September, PDC hosted a course on debugging with Allinea DDT. About fifteen people attended the course.

Allinea DDT is a software development tool to assist people writing code. Its purpose is to speed up the problem resolution process as much as possible. The tool is produced by a UK-based company called Allinea, which specializes in software tools to help the high performance computing (HPC) community move science forward. Allinea also produces a development environment called "Allinea Unified environment" (to improve the performance and sturdiness of HPC applications) and a performance analysis tool called "Allinea Performance Reports" (to improve the efficiency of clusters). These tools are widely spread amongst the bigger scientific institutes all over the world.

Allinea DDT is useful for resolving bugs quickly. It identifies the location of problems in the applications, and provides the user with:

- mechanisms to reveal hidden and sneaky problems,
- diagnostics to understand the root causes of potential issues, and
- intuitive tools to retrieve context information from the application and present them to the end-user.

The training course reviewed the core mechanisms of Allinea DDT, went through how and when to use the tool's capabilities to resolve specific problems, and then time was dedicated to using the tool on the attendees' real life examples. This course was tailored for HPC developers - ranging from scientists to application and parallelism experts. It was useful for anyone who feels that they are spending too much time resolving problems with HPC applications written in C/C++ or Fortran.



Above: Researchers at the Efficient Debugging with Allinea DDT course work on debugging their own code.

Introduction to PDC Systems Course

Henric Zazzi, PDC

PDC runs a range of courses to help people learn to use the PDC computer systems, and to help experienced users make more effective use of the PDC resources. On the 16th of September, PDC's application experts Henric Zazzi and Radovan Bast ran an "Introduction to PDC Systems" course. Ten people from KTH attended the course - some of them had background knowledge about high performance computing (HPC) prior to the course, whereas others had no previous experience of using HPC systems. Many of the participants had already started using PDC's resources, and attended the course as they wanted to know more.

For anyone who missed the course, it will be run again in February next year. Information about the course is available here:

<https://www.pdc.kth.se/events/event-repository/introduction-to-pdc-systems-february-2015>

If you would like to read up in advance, you can find the first part of the course presentation here:

<https://www.pdc.kth.se/support/Presentations-about-using-PDC-resources/introduction-to-the-pdc-environment-part-1>

and the second part here:

<https://www.pdc.kth.se/support/Presentations-about-using-PDC-resources/introduction-to-the-pdc-environment-part-2>

NeIC News

NeIC Generic Area Still Coordinated from PDC

Michaela Barth, PDC

Changes are afoot in the NeIC executive team thanks to vigorous efforts on the part of Harald Barth who rendered the NeIC Generic Area coordinator, Michaela Barth, temporarily unfit for duty.

Dejan Vitlacil from PDC joined the Nordic e-Infrastructure Collaboration (NeIC) on the 1st of October. In a smooth transition, he took over from the previous NeIC Generic Area coordinator, Michaela Barth, who went on maternity leave on the 1st of November. This way the coordination of the NeIC Generic Area is still being handled from the Stockholm area, which is geographically well positioned within the Nordic countries.

The Generic Area Coordinator is part of the NeIC Executive Team and thus works to help the Nordic e-infrastructure evolve towards a high quality solution that benefits Nordic scientific researchers and improves their supporting infrastructures.

The biggest project recently started in the NeIC Generic area is the Glenna Nordic Cloud project, with Glenna being an Icelandic name meaning "Opening in the Clouds". The project has an ambitious agenda and the participation of all five Nordic countries in the project shows its potential.

Another interesting domain with opportunities in the Generic Area falls within data and storage management. NeIC hopes that Dejan with his extensive experience (including management, configuration and development activities for storage and data infrastructures at the institutional, national and international levels) will drive this important question forward during his time as Generic Area coordinator. Dejan will work part-time for NeIC while also continuing with some of his previous responsibilities at PDC.



PDC Experts Contribute to the PRACE SHAPE Programme

Lilit Axner, PDC

SHAPE, the SME HPC Adoption Programme in Europe is a pan-European, PRACE-based programme supporting the adoption of high-performance computing (HPC) by small and medium-sized enterprises (SMEs). The programme aims to raise awareness and equip European SMEs with the expertise necessary to take advantage of the innovation possibilities opened up by HPC, thus increasing their competitiveness. SHAPE (see <http://www.prace-project.eu/shape>) helps European SMEs overcome barriers to using HPC, such as the cost of operation, lack of knowledge and a lack of resources. The programme also facilitates the process of defining a workable solution based on HPC and defining an appropriate business model.

SHAPE Pilot was a trial programme launched to prove the viability and value of the main programme. Its objective was to refine the details of the initiative and prepare its launch in a fully operational form. The pilot (see <http://www.prace-ri.eu/shape-pilot-call>) aimed to work with a few selected SMEs and introduce HPC-based tools and techniques into their business, operational, or production environments. It was open to all European SMEs that had an interesting idea that could be implemented using HPC.

The pilot started in June 2013 with an open Call for Applications to all European SMEs that had a business project that could be implemented using HPC. The call had a brilliant response, involving 14 SMEs, from 7 different countries (Bulgaria, France, Germany, Ireland, Italy, Spain and the UK), that spanned various industrial domains. The PRACE SHAPE experts worked with the ten SMEs that were selected from the fourteen that applied to

develop suitable solutions. Amongst other things, the experts provided the SMEs with knowledge that allowed them to make informed decisions on the HPC solutions that were selected and helped them plan future actions, while taking into account the support of independent service providers at a later stage. (For more details, see <http://www.prace-ri.eu/shape-pilot-selection>.)

The ten companies that were selected were Albatern (UK), ENTARES Engineering / NEXIO and Audionamix (France), Juan Yacht Design (Spain), NSilico (Ireland), OPTIMA Pharma (Germany), MONOTRICAT, AMET, LAPCOS and THESAN (Italy).

MONOTRICAT was one of the four Italian small companies that were selected. MONOTRICAT designs an innovative boat hull, characterized by hydrodynamic efficiency; the aim of the proposed project was to use computational fluid dynamics (CFD) on HPC methodologies applied to MONOTRICAT's hull in order to optimize their research and development workflow.

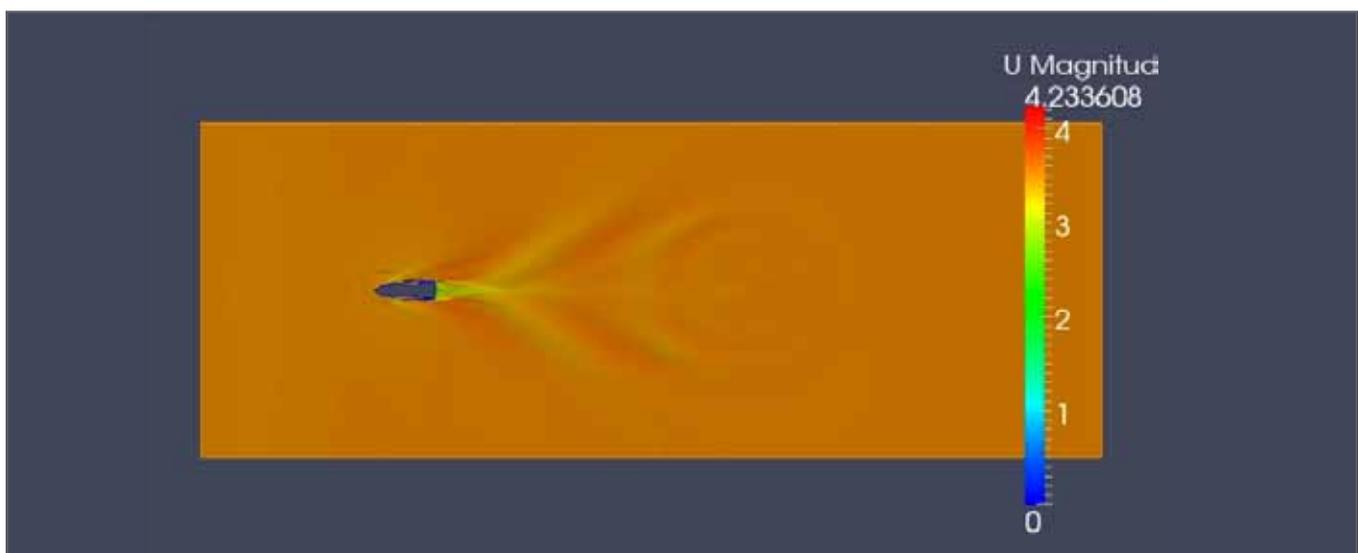
PDC's CFD application expert, Dr. Jing Gong, was the PRACE expert assigned to help implement the proposed MONOTRICAT project. Jing Gong, in collaboration with CINECA as a contact point for MONOTRICAT, has successfully implemented the optimization of the hull design using OpenFOAM software. In the project a whole hull simulation of a ship was conducted, (including mesh generation of

complex geometries, efficient solvers with various turbulent modelling and optimized parameters as well as visualization), along with a performance analysis of the simulation. This project made it possible for MONOTRICAT to take advantage of HPC-enabled simulation tools and potentially replace the traditional methods for refining hull designs.

All the simulations were conducted on the PRACE Tier-1 system Lindgren at PDC. OpenFOAM v2.2.0 built with the GNU compiler v4.6.3 was used in the project.

The results of the project verified that OpenFOAM is an accurate and reliable CFD tool for MONOTRICAT hull resistance simulations. The numerical results agree well with the experimental measurements. The scalability tests illustrated that a linear speed-up can be obtained using between 6 and 72 cores and that an efficiency of 72.9% can be achieved for up to 144 cores. The accurate numerical results and the scalability tests are very encouraging as they make it possible for MONOTRICAT to focus on CFD simulations in future instead of basing their development process on traditional experiments – where, for example, one builds a physical boat and tests its performance in the water.

More about these results and the results of other 9 SHAPE pilot projects can be found at <http://www.prace-ri.eu/shape-prototypes>.



Above: The velocity field on the air/water interface surface with $\alpha = 0.5$.

PRACE Best Practice Guide for Xeon Phi Processors

Michael Schliephake, PDC

The Partnership for Advanced Computing (PRACE) is not only known for its ambition to provide European researchers with access to world class computing facilities. PRACE also provides concrete help to prepare applications for their use on large-scale machines and offers extensive training and education material focused on technology and software development for parallel computing.

Under the auspices of PRACE, a series of best practice guides has been developed to support users of high-performance computing systems. Topics for these best practice guides include optimal porting of applications, architecture-specific optimisation and petascaling techniques, optimal system environments, debuggers, performance analysis tools and programming environments.

PDC took part in the development of the best practice guide for the recent Intel Xeon Phi processors with contributions from staff members Michaela Barth and Michael Schliephake. This document provides information about Intel's Many Integrated Cores architecture (MIC) and programming models for the Intel Xeon Phi coprocessor. It covers a wide range of topics related to programming models, as well as giving information about tools and strategies for analysing and improving application performance that enable programmers to achieve good performance with their applications.

You can find all the PRACE best practice guides, including the Intel Xeon Phi guide, online here:

<http://www.prace-ri.eu/best-practice-guides>.

PDC-Related Events

Beskow Inauguration

27 January 2015, KTH Main Campus

Please join us for the inauguration of our new supercomputer system Beskow. For details, see <http://agenda.albanova.se/conferenceDisplay.py?confId=4684>.

Introduction to PDC Systems

9 February 2015, PDC, KTH

For more details about this introductory course, please see <https://www.pdc.kth.se/events/event-repository/introduction-to-pdc-systems-february-2015>.

PDC Summer School 2015

17-28 August 2015, KTH Main Campus, Stockholm

Details about the PDC Summer School will be available here in March 2015:
<http://www.pdc.kth.se/education/summer-school>

HPC Sources

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

CERN

<http://cerncourier.com/cws/events>
<http://cds.cern.ch/collection/Forthcoming%20Events?ln=en>

EGI

<http://www.egi.eu/about/events>

HPC University

<http://www.hpcuniv.org/events/current>

HPCwire

<http://www.hpcwire.com/events>

Linux Journal

<http://www.linuxjournal.com/events>

PRACE

<http://www.prace-ri.eu/HPC-access>
<http://www.training.prace-ri.eu>
<http://www.prace-ri.eu/events>
<http://www.prace-ri.eu/news>

SNIC

<http://www.snic.vr.se/news-events>
<http://docs.snic.se/wiki/Training>

XSEDE

<https://www.xsede.org/conferences-and-events>

Save the date!

NeIC 2015 Conference

5-8 May 2015, Finland

<http://neic.nordforsk.org/neic-2015-conference>

Please join us in Finland next spring and enjoy getting in touch with your Nordic colleagues while we all benefit from the dedicated workshops at the NeIC 2015 Conference!