



**no.1** 2016

**PDC Center for High Performance Computing**

# PDC Newsletter

**EASC2016 in Stockholm** - Erwin Laure, Stefano Markidis and Apostolos Vasileiadis, PDC/CST, page 3

**SeRC: The First Five Years and Future Plans**

- Olivia Eriksson, Erwin Laure, Erik Lindahl, Dan Henningson and Anders Ynnerman, SeRC, page 4

**BioExcel's New Webinar Series** - Rossen Apostolov, PDC, page 12

**Allinea Performance and Debugging Tools Workshop** - Cristian Cira, PDC, page 13

**Are You Using the PRACE CodeVault?** - Michaela Barth, PDC, page 13

**SNIC All Hands Meeting** - Gert Svensson, PDC, page 14


**New Type of PRACE Preparatory Access Available Soon!** - Michaela Barth, PDC, page 15

**VR-Cube Retrospective Seminar** - Gert Svensson, PDC, page 16

**PDC Pub 2016** - Thor Wikfeldt, PDC, page 17

**Our Ongoing Swedish SHAPE Projects** - Michaela Barth, PDC, page 18

**DECI (Tier-1) Access** - Michaela Barth, PDC, page 20





**Erwin Laure**  
**Director of PDC, and CST**

**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:** [support@pdc.kth.se](mailto:support@pdc.kth.se)

ISSN 1401-9671

**Cover**

SeRC research on 3D vortical structures in a torus: The curvature is changing the turbulence characteristics leading to laminarisation on the inside and increased turbulence outside.

**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](mailto:support@pdc.kth.se)  
**www:** <http://www.pdc.kth.se>  
**Phone:** +46 8 790 7800  
**Fax:** +46 8 247 784

**Editorial**

The Swedish e-Science Research Centre (SeRC) has had a tremendous impact on computational research in Sweden and many of PDC's users are SeRC researchers. PDC, together with the National Supercomputer Centre (NSC), constitute the SeRC core computational infrastructure. In this newsletter we discuss some of SeRC's key successes and present some of the strategic areas that SeRC will address in the future. One of these is a Scientific Computing Lab that will be established at PDC over summer - this will help researchers who use PDC's resources in their collaborations with PDC's application experts.

PDC has been busy organizing events in recent months. Before Christmas we held the first Swedish National Infrastructure for Computing (SNIC) all-hands meeting where the SNIC employees met to discuss improving SNIC's services. April found us running the 4th Exascale Applications and Software Conference (EASC) – a high-profile international event where the challenges associated with exascale computing are discussed. The traditional PDC pub in May provided an informal setting for users and PDC staff to meet, and earlier on in January we had hosted a workshop in collaboration with Allinea that focused on their performance and debugging tools. More recently the BioExcel project started a sequence of webinars, targeting biomolecular researchers. Nostalgic sentiments were stirred in March when we participated in a seminar looking back at the Virtual Reality Cube that PDC operated between 1998 and 2008.

There is also quite a lot of news on the international front as PRACE is introducing new types of preparatory access mechanisms, as well as a code vault. PDC is also active in two PRACE industrial SHAPE projects and we are delighted that quite a few Swedish researchers have been granted additional computing time through the DECI Tier-1 access.

This spring the Swedish Research Council (VR) started a process with the Swedish universities to discuss the future long-term financing of SNIC, a process that will extend well beyond summer. Another more local process that will keep us busy over the summer is the establishment of the SeRC Scientific Computing Lab (which is mentioned in the SeRC article in this newsletter). During the summer we will be re-building parts of PDC to create an open interaction space where PDC experts and users can collaborate. You will find more about both of these activities in the next edition of our newsletter.

But, for now, I wish you a relaxing summer and I'm looking forward to all the great research that will be done on PDC's systems in autumn.

*Erwin Laure – Director PDC, and CST*

# EASC2016 in Stockholm

Erwin Laure, Stefano Markidis & Apostolos Vasileiadis, PDC/CST

The 4th International Exascale Applications and Software Conference (EASC2016) was held in Stockholm from the 25th to the 29th of April 2016. The conference was organized by PDC with support from the Swedish e-Science Research Centre (SeRC), the KTH School of Computer Science and Communication, Cray, NVIDIA and the EPIGRAM project. Over 80 attendees gathered to discuss application and software challenges in the upcoming Exascale era.

The program featured 37 contributed presentations, selected by an international programme committee, as well as four keynote presentations by high-profile international experts. Satoshi Matsuoka from the Tokyo Institute of Technology discussed what impact the foreseeable end of Moore's law will have and concluded that we are about to change from a FLOPS-oriented to a bytes-oriented regime. Karlheinz Meier from Heidelberg University gave a different perspective on computing, discussing brain-inspired neuromorphic approaches. John Shalf from the Lawrence Berkley National Laboratory gave a sneak preview of the Exascale era and the programming systems that we will be using in the future. Finally, Erland Källén from the European Centre for Medium-Range Weather Forecasts explained how Exascale systems can help producing ever more detailed and precise weather forecasts.

The conference also featured a dedicated session about the recently started High Performance Computing Centres of Excellence (CoE), as well as a workshop on European Exascale projects. A panel discussion explored what role CoEs and other European Exascale efforts will have and what efforts are going to be needed to smooth the transition from Peta- to Exascale.

Videos of the keynote presentations, as well as slides from many of the other presentations, are available at the conference website: <http://www.easc2016.com> and on EPIGRAM's YouTube channel at <https://goo.gl/VfNiNu>.

In 2017, EASC will be back to Edinburgh, where this successful conference series started.



Above: EASC2016, Stockholm

## In This Issue

### Editorial

Erwin Laure ..... 2

### Staff Focus

Dana Akhmetova ..... 8

Juris Vencels ..... 9

Ragnar Sundblad ..... 15

Thor Wikfeldt ..... 16

### EASC2016 in Stockholm

Erwin Laure, Stefano Markidis & Apostolos Vasileiadis ..... 3

### SeRC: The First Five Years and Future Plans

Olivia Eriksson, Erwin Laure, Erik Lindahl, Dan Henningson & Anders Ynnerman ..... 4

### BioExcel's New Webinar Series

Rossen Apostolov ..... 12

### Allinea Performance and Debugging Tools Workshop

Cristian Cira ..... 13

### Are You Using the PRACE CodeVault?

Michaela Barth ..... 13

### SNIC All Hands Meeting

Gert Svensson ..... 14

### New Type of PRACE Preparatory Access Available Soon!

Michaela Barth ..... 15

### VR-Cube Retrospective Seminar

Gert Svensson ..... 16

### PDC Pub 2016

Thor Wikfeldt ..... 17

### Our Ongoing Swedish SHAPE Projects

Michaela Barth ..... 18

### DECI (Tier-1) Access

Michaela Barth ..... 20

PDC-Related Events ..... 20

HPC Sources ..... 20

# SeRC

## Swedish e-Science Research Centre

### SeRC: The First Five Years and Future Plans

Olivia Eriksson, Erwin Laure, Erik Lindahl, Dan Henningson and Anders Ynnerman, SeRC

#### What is SeRC?

The Swedish e-Science Research Centre (SeRC) is a scientific environment within the strategic research area (SRA) of e-Science, and is funded by the Swedish Government. SeRC is based on a collaboration between four Swedish universities: the KTH Royal Institute of Technology (KTH), Stockholm University (SU), the Karolinska Institute (KI) and Linköping University (LiU). SeRC's mission statement is to develop state-of-the-art e-Science tools and provide e-infrastructure support to existing and emerging e-Science research communities to help bring about scientific breakthroughs in Sweden.

SeRC was founded in 2010 as the result of the SRA initiative launched by the Swedish Government Bill on Research Policy in 2008, where a total of 24 different strategic research areas were defined – one of which was e-Science. Initially SeRC was granted funding for five years. During those first five years, SeRC has built an organisation for e-Science research, which has been highly successful – this was reflected in the excellent grades that SeRC received when the SRAs in Sweden were evaluated in 2015. The next phase of SeRC will partly focus on activities relating to emerging technologies (such as exascale and big data) while also consolidating SeRC's ongoing efforts in working towards a long-lasting e-Science environment in Sweden.

#### e-Science and SeRC

We are right in the middle of a revolution where computers are redefining the process

of scientific research as we know it. While this trend started in the natural and engineering sciences, it has evolved to become a pervasive component of virtually all scientific fields: modern aeroplane design depends more on computational fluid dynamics than on wind tunnels, bioinformatics has turned biology into a quantitative subject, and we are seeing entirely new data-driven research, for instance in medicine.

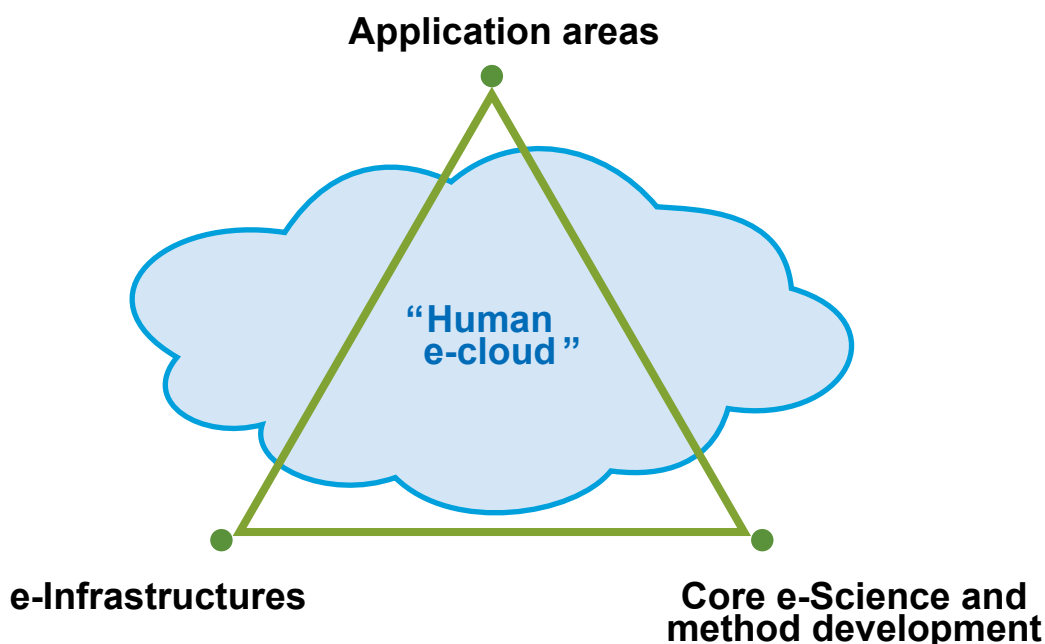
Many of the challenges that are being faced in this new e-Science are similar across different scientific disciplines: all these research directions rely on having access to large computational resources and the ability to work with and archive huge amounts of data, and they also need new algorithms (for performing both computations and large-scale analysis of data). Consequently there is also a demand now for new types of e-Science experts, known as application experts, who are highly skilled in a particular applied research area and also in e-infrastructures or method development.

The concept of e-Science is, in its most basic form, simply the notion of using digital information and processing that information to gain new scientific insights. This means that a lot of e-Science research can be conducted using the standard hardware and software that is widely available. However there are also research problems that require additional resources, which is where SeRC comes in.

#### SeRC's Scope, Model and Mission

SeRC's scope is to address the e-Science problems that need solutions which go beyond the standard ones that are currently available, and that require researchers to have access to a state-of-the-art e-infrastructure as well as new e-Science methodology and technology. SeRC is thereby pushing the development of new e-Science tools forward. In this way SeRC paves the path for new high-end e-Science that will eventually mature and benefit not only the groups involved in SeRC, but also the whole academic base of e-Science users.

*Below: The work in SeRC relies on three pillars of existing excellence in application areas, method development, and e-infrastructures, connected through a "human e-cloud".*



To address these kinds of problems, SeRC is based on a three-pillar model (shown above) in which existing excellence in application areas, core e-Science and method development, and computational infrastructures interact to find new e-Science solutions. In particular, this is done through collaborations and application experts active in at least two of these pillar areas. For the long-term, SeRC is working on developing a "human e-cloud" in order to better address challenges that are common to different application areas, and also to provide clear career paths for this new kind of researcher with dual expertise (that is, the application experts).

### **Pillar 1: Application Areas**

One of SeRC's central goals is to have direct impact in application areas of strategic relevance for our partner universities, and we also want to focus on research topics where SeRC has world-leading research groups. These SeRC application areas include fields such as computational fluid dynamics, materials science, bioinformatics, cancer epidemiology and screening, neuroinformatics, molecular simulation, and climate modelling.

### **Pillar 2: e-Infrastructures**

SeRC wants to facilitate the development of a world-class computational infrastructure at PDC, at the National Supercomputer Centre at Linköping University, and at the visualisation centres affiliated with SeRC (which are based in Norrköping and at KTH in Stockholm). The development of this infrastructure includes competitive hardware investments in collaboration with the Swedish National Infrastructure for Computing (SNIC), although it is equally important to recruit advanced infrastructure experts to develop key simulation codes further.

### **Pillar 3: Core e-Science and Method Development**

An integral part of SeRC is leading-edge research in core e-Science and method development, for instance in fields such as numerical analysis, visualisation, parallelization, acceleration, and data engineering. SeRC has several core research groups that provide both the latest research results for e-Science tools as well as a broad competence base in e-Science tool usage and methodology.

## “The Human e-Cloud”

For SeRC, the central challenge is to get these three pillars to strengthen and support each other through collaborations. Infrastructures and core e-Science do not exist in a vacuum, but need to engage directly in research on important applications. In particular SeRC focuses on supporting application areas where we see an opportunity to translate advances in fundamental e-Science into research impact.

Our primary tool to achieve this is the recruitment of cross-disciplinary e-Science experts (application experts) that work between these areas – we call them the “human e-cloud” in the figure on the previous page. These experts combine a deep knowledge of an application area with expertise either in method development or large-scale computational infrastructures. By integrating the application experts into the applied research environments (while they work close to the method development groups or the infrastructure), it provides opportunities for close collaborations between application experts and applied researchers, as well as between application experts facing similar challenges (for instance in parallelization, data management or method development) in different application areas.

## Organisation

SeRC consists of ten research communities, within which the e-Science research takes place, together with the SeRC e-Science infrastructure. Some of the SeRC research communities focus on core e-Science methodology, like numerical analysis, parallel software and data engineering and visualization, whereas others focus on applying the methodology in a particular area, like bioinformatics, neuroinformatics, cancer epidemiology and screening, molecular simulation, electronic structure, computational fluid dynamics or climate research. The infrastructure on the other hand consists of:

- two supercomputer centres – the PDC Center for High Performance Computing at KTH (PDC) and the National Supercomputing Centre at LiU (NSC), both of which are part of SNIC,

- two visualisation centres – the Visualization Center C (VC), located in Norrköping, and the Visualisation Studio at KTH (VIC), and
- a number of application experts who are employed at the supercomputer centres, and who work closely with the research communities.

SeRC co-finances about 30 projects and 15 faculty positions, including a number of larger Multidisciplinary Collaboration Programmes that were initiated in order to increase the level of collaboration between researchers in the areas of applications, methodology and infrastructure.

In December 2013, to consolidate the organisation, SeRC was turned into a research centre with KTH as the host, and we recruited an external chair of the board, Professor Morten Dæhlen, who is Dean of Mathematics and Natural Sciences at the University of Oslo and a prominent international e-Science leader. More information about the SeRC organisation can be found at <http://www.e-science.se/org>.

## SeRC Faculty

Since the inception of SeRC, we have recruited a substantial number of highly skilled assistant professors. In addition, we have selected a number of e-Science investigators who were either hired recently or who recently joined SeRC with support from other funding. Together, these researchers form the SeRC faculty.



*Above: SeRC Faculty, Management Group and Steering Board at a strategy meeting at Ljungbergsgården together with Perlan organisation consultants, September 2015*

The main goal of the SeRC faculty is to facilitate interdisciplinary exchange between e-Science communities in application areas and core e-Science and method development, and to provide SeRC with a network of e-Science researchers that straddle the various research communities.

### **Collaborations and Scientific Output**

SeRC hosts a large number of outstanding scientists who are pushing the state-of-the-art, in application areas as well as method development, with several high-impact results. However, SeRC's mission is to go a step beyond merely providing some marginal funding and instead build a scientific environment that promotes excellence within research disciplines through multi-disciplinary collaboration. SeRC makes it possible to initiate new, ambitious, and risky collaborative projects that have the potential to redefine research fields. In particular, the centre has acted as a catalyst in the formation of several strong research environments that work on the application, method-development, and infrastructure aspects of computational research (for example, see the figure to the right). These multidisciplinary collaborative projects for instance make it possible to engage computational experts in long-term development efforts that had previously been very difficult to fund.

The new initiatives have had tremendous scientific impact: one of the SeRC software projects produced what became Sweden's single most cited scientific paper during 2014, teams that did not even know each other prior to SeRC have merged, and the collaborations between infrastructures and researchers have led to a range of new high-profile international research projects being funded, not to mention the award of several prestigious European Research Council grants. The publication output by SeRC research projects has almost doubled from 2010 to 2014, and the publications of the SeRC researchers have an average field-normalized citation rate, which is a measure of scientific impact, of 1.92 for the period 2010-2014, which is almost two times the world average.

*Below: Example of SeRC collaboration*

*The "e-Science for Cancer Prevention and Control" program has an e-Science team which includes epidemiologists, bio-informaticians, image analysts and computer scientists who are developing, implementing and evaluating methods for using large amounts of biomedical data for risk-stratification in cancer screening and diagnosis. You can read more about it here: <http://www.e-science.se/ecpc>.*



However, instead of being content with these achievements, SeRC will keep raising the bar by continuously evaluating our collaborations with a view to further improving the centre's effectiveness. SeRC's goal is not to maintain the status quo, but to promote discoveries and new applications by helping and encouraging talented researchers to team up and leave their comfort zone. Collaboration, multidisciplinary, change, and impact are the genes of SeRC.

### **Computing Resources**

e-Science is critically dependent on efficient large-scale computing and storage resources. Within SeRC, these resources are provided by the leading Swedish high performance computing centres, PDC and NSC, which offer more than two thirds of the Swedish national compute capacity. A strategic partnership between these centres has been established. This allows for better coordination and alignment of their services with the needs of SeRC researchers.

Since 2015, SeRC researchers have had access to petascale computing resources (see next page) through these centres. But it is noteworthy that PDC and NSC not only provide hardware, they also provide expertise in many application areas in terms of computing experts who work closely with

## Staff Focus



Dana Akhmetova

Dana Akhmetova gained her first degree at the Department of Computer Science at Lomonosov Moscow State University in 2007. She started post-graduate studies at PDC and is now a third year Ph.D. student at the Department of Computational Science and Technology at the KTH Royal Institute of Technology under the supervision of Professor Örjan Ekeberg.

During 2015 Dana was a Ph.D. intern at the Pacific Northwest National Laboratory in the United States, where she was focusing on task-based programming models and analysing the connection between the granularity of application tasks and the scheduling overhead in many-core shared memory systems.

Currently Dana is working on two large, EC-funded projects called INTERTWinE and AllScale. The INTERTWinE project addresses the problem of programming model design and implementation for Exascale computing. The AllScale project is proposing an environment for the effective development of highly scalable, resilient and performance-portable parallel applications for future Exascale systems.

When she is not working, Dana enjoys hiking and reading books.

certain scientific domains, including bioinformatics, climate research, computational fluid dynamics, materials sciences, molecular dynamics, and neuroinformatics.

### Petascale: Beskow and Triolith

In January 2015, the SeRC infrastructure received a significant boost with the inauguration of Beskow at PDC – a 2-petaflop Cray XC40 system which is more powerful than all the other SNIC systems together. With Beskow, SeRC researchers for the first time had a world-class petascale system at their disposal – it was much needed for their research and also as a stepping stone towards using the larger European-level HPC infrastructure. Beskow was designed to cater for the needs of SeRC and important SeRC applications were used as benchmarks to guide the selection process.

In autumn 2017, NSC's flagship system, Triolith, will be replaced with another petascale system, ensuring that SeRC researchers and other e-Science researchers in Sweden continue to have access to the latest hardware technologies.

### The Future and e-Science Agenda

During its first five years, SeRC has had tremendous success in establishing e-Science in Sweden and fostering the e-Science paradigm in many research areas. e-Science is now well established and is a priority area at the SeRC partner universities. The next phase of SeRC (2015–2019) will, on the one hand, further consolidate these efforts by working towards a long-lasting e-Science environment, and, on the other hand, increase SeRC activities relating to emerging technologies, such as exascale and big data.

The recent report Swedish Science Cases for e-Infrastructure (<http://www.e-science.se/sscei>) is an important document when it comes to shaping the future of e-Science in Sweden. This document discusses strategic e-Science development in Sweden and makes a number of recommendations that will help to guide future SeRC developments. Key findings, from a SeRC perspective, include the following.

- “Development of methods, tools and software within core disciplines is necessary to make breakthroughs.”
- “Advanced and long-term user support and human infrastructures are keys to e-Science adoption.”
- “The simulation paradigm dominates the current Swedish needs for e-infrastructure. A complementary and more data-centric aspect of e-Science should be promoted.”
- “e-Social science and e-humanities are potentially very large users, but need active support, like other research communities new to e-Science.”
- “e-Science methods and tools are in increasing demand and will be instrumental in increasing interaction between tool makers and tool users.”



## Staff Focus



Juris Vencels

Juris Vencels started his university studies in Latvia where he obtained a B.Sc. degree in physics from the University of Latvia. In 2013 he started studying electrophysics at KTH and very soon became involved in the computational plasma physics research group led by Stefano Markidis. Juris enjoyed working at PDC as a research engineer during the summer of 2014 as people at PDC are very friendly and helpful. In 2015 Juris spent ten months at the Los Alamos National Lab working on numerical methods for plasma physics and, in March of 2016, he successfully graduated from KTH with an M.Sc. in Electrophysics.

For now Juris is working as a freelancer and getting involved in many crowdsourcing projects, such as big data mining for General Fusion's experimental fusion device. He is excited about the booming start-up culture in Northern Europe, and in the future he sees himself working on industrial problems, developing technologies for renewable energy and contributing to a new company.

During his spare time, Juris participates in different sports, and he also enjoys travelling, getting to know different people and cultures, and photography.

Responding to these findings entails work at several levels ranging from infrastructure and tool development to policy definitions. In its strategy process, SeRC has defined an agenda for the next five years. This agenda addresses many of the items from the report and describes how SeRC is in a unique position to spearhead this work.

Later sections give an overview of the new SeRC initiatives that have been defined to help SeRC evolve as an organisation, and to pave the way for the next stage of e-Science development and deployment in Sweden.

### The Impact of e-Science

The impact of e-Science can be seen in all areas of academia. The use of basic e-Science in new domains can generate fundamentally new and important results. At the same time, access to advanced e-Science tools and infrastructures enables the leading research groups that are spear-heading Swedish e-Science to compete at the international forefront.

In view of this, SeRC intends to intensify the collaboration with relevant Swedish groups and universities that are currently not part of SeRC and aims for a national scope in domains like bioinformatics, numerical analysis and materials science. By expanding such collaborations, SeRC will develop into a truly national organisation. This will also allow SeRC to take responsibility for further developing the Swedish e-infrastructure landscape together with SNIC.


The impact of the broad adoption of e-Science methodology cannot be overestimated. Through extensive collaborations with industry and governmental agencies, SeRC has a direct impact in society and contributes to societal development and commercial competitiveness. SeRC has several industrial representatives on its advisory board, and we are delighted with their interest and engagement in learning how computing can change what they do, for instance at our annual meetings. One of the primary indirect mechanisms for SeRC to have an effect is through knowledge transfer in the form of recent Ph.D. graduates, who have unique experiences of e-Science and form a network consisting of the next generation of e-scientists.

The single largest overall success of SeRC is that it has helped application researchers and computational experts to find each other and speak with one voice. Large-scale data analysis and simulation have been introduced in areas where they have never been used before, the strongest computational groups in the country suddenly collaborate internationally instead of competing nationally, and we will keep redefining these boundaries by identifying areas where the combination of computational and application research excellence makes a field ripe for disruptive changes to the science.

## SeRC Initiative: Multidisciplinary Collaboration Programmes

In the first phase of SeRC, larger, multidisciplinary constellations have been prototyped through a multidisciplinary collaboration programme (MCP) promoting multidisciplinary approaches to strategic research areas with high potential for gains. The e-Science for Cancer Prevention and Control (eCPC) was SeRC's first MCP, which was followed by Visualisation In e-Science Applications (see figure below) and SESSI - SeRC Exascale Simulation Software Initiative (see figure on the opposite page). SeRC will now build on these successes and further develop the MCPs by consolidating the SeRC efforts into larger constellations.

A key requirement of MCPs is the existence of external third-party support contributing to the programme goals. These multiple funding streams are one component to ensure the sustainability of SeRC efforts. For 2015–2019, three additional MCPs have been started: Data-Driven Computational Materials Design, FAST Climate Science and Brain-IT.



*Above: An illustration from the SeRC Multidisciplinary Program 'Visualisation in e-Science Applications' showing a visualisation of a protein, combining hybrid data from different sources using a hybrid visualisation of geometric protein representations and volumetric data*

*The improved A-buffer algorithm is capable of rendering the entire scene with a performance increase of 5.9 times compared to prevalent techniques.*

## SeRC Initiative: Scientific Computing Lab

SeRC groups are engaged in the development of computational methods and infrastructures, which are necessary ingredients for world-class computer-based science research. With architectural changes, which are being driven partly by the challenge of building exascale systems, these efforts need to be reinforced and strengthened. Developments – such as massively increased core-counts, large vector units, deep memory hierarchies, and the like – require significant efforts in the development of efficient, scalable methods and implementations.

Such work is planned in the multidisciplinary collaboration programs referred to earlier. These constitute ideal cases for the SeRC vision of integrating researchers (from the three pillars that were described earlier) with each other in such a way that collaboration between application area, e-infrastructure and methods researchers is being facilitated by the application experts in the Human e-Cloud.

Following international examples, for instance from the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign, we intend to prototype a scientific computing lab, where researchers working on applications, methods, and infrastructure can effectively collaborate in such a manner. A physical interaction space that will be created at PDC is a key component in this endeavour. With this lab, SeRC will build the nucleus of the Human e-Cloud identified as a key ingredient above.

## SeRC Initiative: Data-driven Science

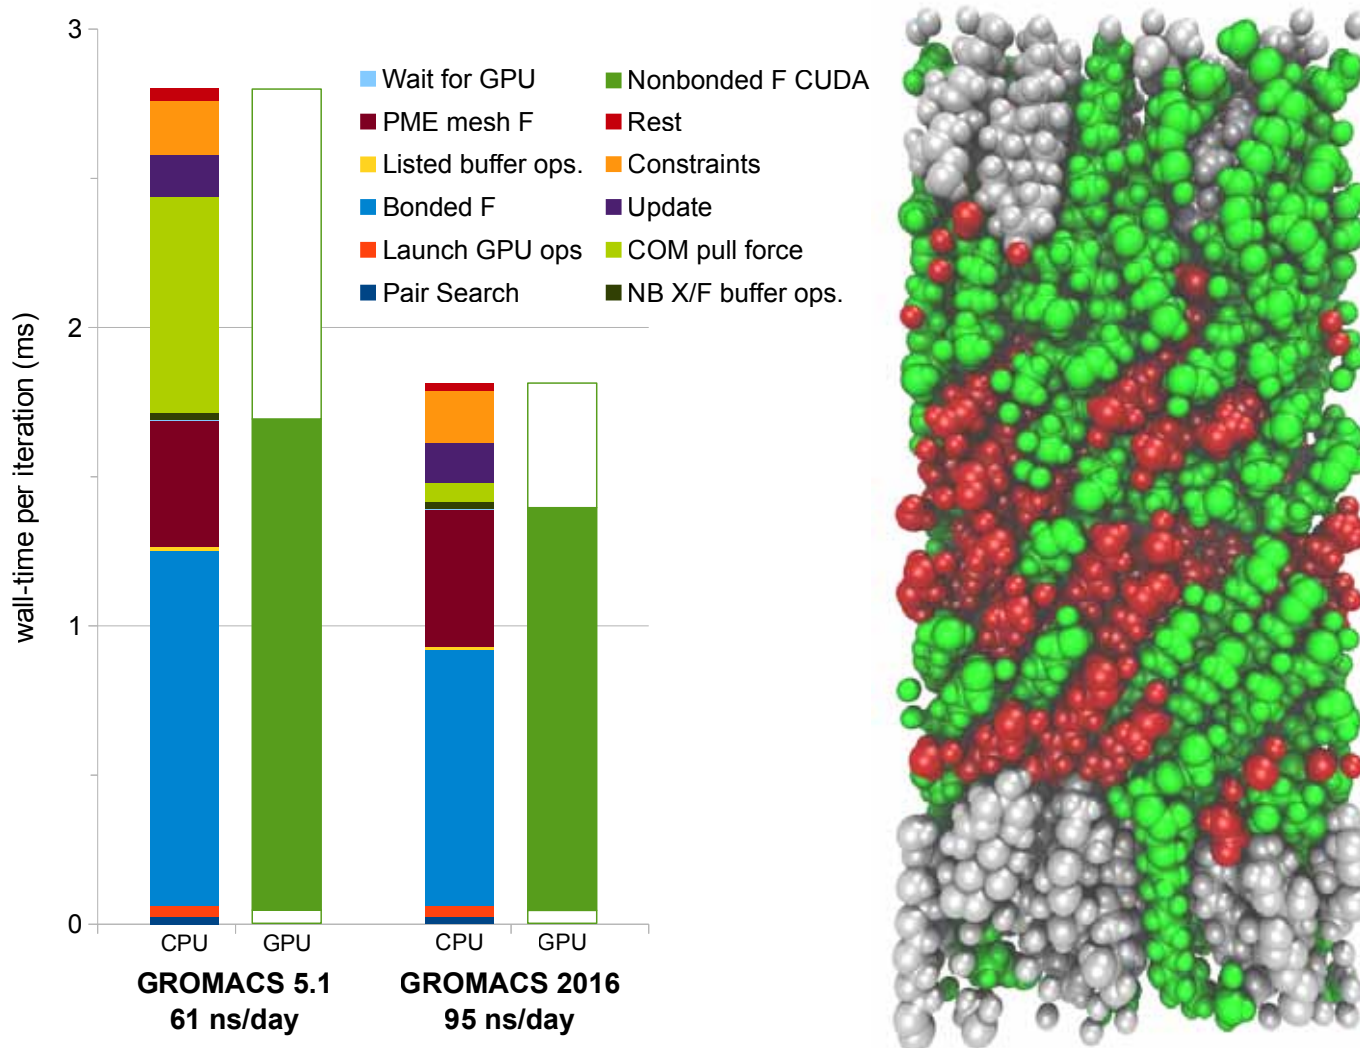
The data deluge that is arising from new scientific instruments, large simulations, and data on the internet, together with new technologies for handling and analysing this data, has significant impact on e-Science methodology. SeRC is therefore embracing big data, with an initial emphasis on the life sciences. This includes research regarding storage, management, integration, visualisation, and analysis of big data.

Below: This illustration from the SeSSI Multidisciplinary collaboration program shows results from computational modelling of the skin barrier (the lipid matrix of the stratum corneum) using molecular dynamics simulations, which gives us further insight into the largest organ of the human body and will also further clinical experiments. Thanks to the highly optimized heterogeneous parallelization in the code GROMACS, complex computational studies can be carried out quickly and efficiently. The image on the right shows the modelled molecular system with ceramide molecules in green, cholesterol molecules in white, and fatty acids in red.

Recent work in SeSSI (that focused on improving SIMD, GPU, and thread parallelization) resulted in speeding up the calculations by up to 50%. The figures on the left show the execution time breakdown of the CPU and GPU tasks in GROMACS versions 5.1 and 2016; improved performance of multiple tasks leads to an increase in simulation throughput from 61 ns/day to 95 ns/day.

The simulations were performed on a workstation equipped with a Core i7-5960X CPU and a GeForce TITAN X GPU.

Source: <https://doi.org/10.6084/m9.figshare.3423395.v1> (Magnus Lundborg and Szilárd Páll) – image reused under CC-BY



In recent years, new platforms and tools have appeared that can efficiently and cost-effectively store and process up to petabytes of data. This allows many organisations to handle data that was previously considered to be too expensive to store and manage. This data-driven science trend is transforming scientific research, by making data-driven discovery and prediction possible. The requirements for data analytics on a large scale include data management on public and

private clouds, data-parallel algorithms, and proficiency with a complex ecosystem of tools and platforms such as Hadoop. Currently these are not widely available within the existing SeRC research communities, so one of the goals for SeRC is to help transfer knowledge about the use of methods and tools from data science to facilitate the use of new research methodologies in other SeRC communities.



## BioExcel's New Webinar Series

Rossen Apostolov, PDC

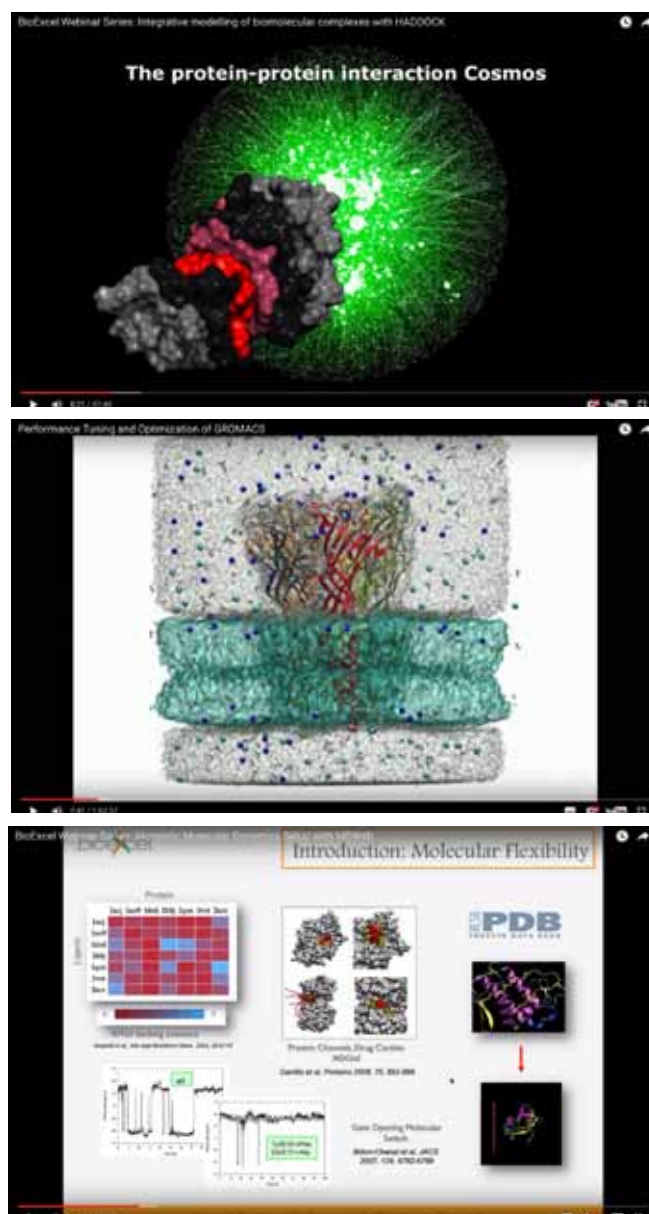
The BioExcel Center of Excellence (<http://bioexcel.eu>) has launched a new series of educational webinars for computational biomolecular research. The series will cover broad topics related to the latest developments with the major software packages, their application to modelling and simulation, best practices for performance tuning and efficient usage on high performance computing (HPC) and novel architectures, introductory tutorials for novel users and much more.

The series began at the end of April with Alexandre Bonvin from Utrecht University talking about **computational prediction of quaternary structure** of biomolecular macromolecules, which is of paramount importance for the fundamental understanding of cellular processes and drug design. One way of increasing the accuracy of modelling methods that are used to predict the structure of biomolecular complexes is to include as much experimental or predictive information as possible in the process. **HADDOCK** (<http://www.bonvinlab.org/software/haddock2.2/>) is one of the most widely used tools for integrating the information derived from biochemical, biophysical or bioinformatics methods to enhance sampling, scoring, or both.

In the second webinar series, Mark Abraham from KTH addressed the very important topic of **squeezing every cycle out of your hardware with GROMACS**, one of the most highly tuned tools for Molecular Dynamics (MD) simulations and one of the most heavily used codes at PDC and other HPC centres world-wide. Developers are making exciting progress with the automation of the GROMACS code, but finding the optimal balance for a particular embedded-solvated-multi-protein-membrane system on a multi-node HPC machine

with accelerators can be a challenge. A bit of theory though can easily give you extra power!

The series will continue with a talk for novice MD users and another on free energy calculations. All the talks are recorded and available on the BioExcel website (<http://bioexcel.eu/webinars>) and through the **BioExcel YouTube channel** - make sure to watch them and share them with your colleagues! You can also subscribe to the BioExcel newsletter (through BioExcel's website) to receive information about these and other useful upcoming BioExcel events.




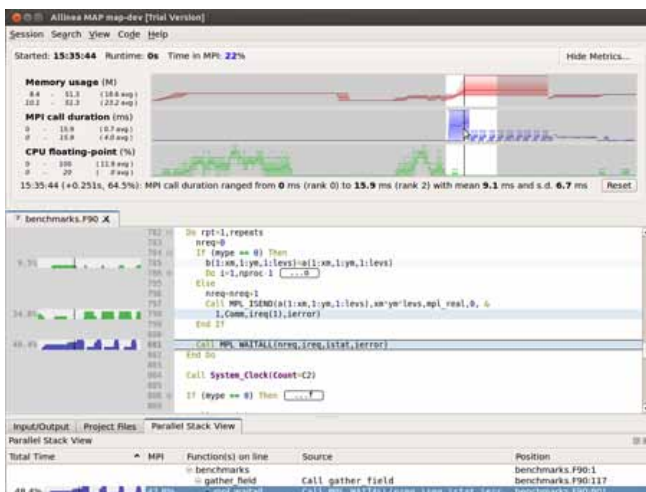
Above: BioExcel's webinars on Integrative modelling of biomolecular complexes with HADDOCK (top), Performance Tuning and Optimization of GROMACS (middle), and Atomistic Molecular Dynamics Setup with MDWeb (bottom)

# Allinea Performance and Debugging Tools Workshop

Cristian Cira, PDC

The high performance computing software specialists Allinea ran a technical workshop at PDC on the 25th of January 2016. The course gave an overview of Allinea's tools for debugging (DDT) and performance analysis (MAP). The first part of the day focused on their flagship product, Allinea DDT, which is a widely used tool for debugging parallel code. DDT can point to a part of the code and indicate the processing element (PE) responsible for erroneous behaviour. The tutorial during the course consisted of examples ranging from identifying straightforward crashes to MPI deadlocks and memory leaks.

Among the interesting new features that were presented during the workshop, it is worth mentioning the version-controlled detection mechanism and the Allinea Forge graphical user interface (GUI). This user interface has matured into a common interface for both DDT, and Allinea's profiling tools (like MAP). MAP gathers samples of performance data from running applications and reports the results in graphical form in Allinea Forge. This common interface facilitates a very smooth transition between debugging and performance analysis, offering an intuitive and direct way to navigate and to identify bogus code, as well as bottlenecks or hotspots. If you missed this workshop by Allinea, we recommend that you keep an eye on the PDC mailing list or website for future courses.



## Are You Using the PRACE CodeVault?

Michaela Barth, PDC

If not, maybe you should be!

The PRACE CodeVault went public at the beginning of this year (on the 29th of January) and a lot of PRACE effort had already gone into it by then.

So, what is the CodeVault? It is an open repository containing various high performance computing (HPC) code samples. The project aims to support self-learning of HPC programming and is intended to be used as an open platform for the HPC community to share example code snippets, proof-of-concept codes and so forth.

The CodeVault contains training material from PRACE partners, as well as example codes of common HPC kernels (such as dense and sparse linear algebra, spectral and N-body methods, structured and unstructured grids, Monte Carlo methods and parallel I/O). The code samples are published as open-source and can be used both for educational purposes and as parts of real application suites (as permitted by the particular licences).

The PRACE CodeVault is being hosted at <https://gitlab.com/PRACE-4IP/CodeVault> and it is open with anonymous read access.

If you would like to contribute code to the CodeVault, please see the instructions in the repository.

*Left: Allinea Forge interface showing code and performance characteristics side by side*

# SNIC All Hands Meeting

Gert Svensson, PDC

PDC was the local organizer of the first SNIC All Hands meeting which was held at SÅstaholm, north of Stockholm, on the 9th and 10th of December 2015. SNIC is the acronym for the Swedish National Infrastructure for Computing, the body that organizes the academic use of high performance computing (HPC) in Sweden. Gert Svensson from PDC chaired the organization committee for the meeting with Anna Jänis, Head of Centre at the National Supercomputer Centre, Linköping, and Ann-Charlotte Sonnhammer from SNIC as members. The purpose of the All Hands meeting was to share new information about SNIC and the SNIC activities and projects, from both the technical and organizational perspectives, with all the staff at the six HPC centres in Sweden that are funded by SNIC.

A large part of the meeting was devoted to discussions about different ways to make SNIC and the interactions between SNIC and the SNIC staff more efficient and to improve the collaboration between the SNIC centres. To facilitate these discussions, Lena Mc Evenue and Sven Schelin from Gällöfsta Perlan Ledarskap were brought in to function as discussion leaders.

The program started with a thorough presentation of the visions and activities of SNIC by the SNIC Director, Jacko Koster. During the subsequent discussion, we divided the participants into groups with a maximum of seven members. Each group presented their ideas for improvement on a hand written poster at the end of the session. Numerous ideas for improvement were raised in many different areas. One common recurring theme was the need for a longer term perspective in the planning at SNIC. However it is not so easy to rectify this as SNIC is dependent on decisions by the Swedish Research Council (VR) and these are made on relatively short term bases.

During the meeting, which lasted from lunchtime on the 8th of December to lunchtime on the 9th, a number of projects and activities were presented by participants from the different SNIC

centres and many issues were discussed in groups and documented for SNIC to follow up on.

As usual, SÅstaholm provided a charming environment with a relaxed atmosphere and a traditional Swedish Christmas dinner (with some not so traditional elements too), which was excellent. During the evening, a number of people were spotted relaxing in the bar and at the disco dressed in costumes from the time when SÅstaholm was a retirement home for Swedish actors and actresses.



Above: The Director of SNIC, Jacko Koster, discusses collaboration between the SNIC centres.



Above: Lena Mc Evenue and Sven Schelin summarize a discussion about improving SNIC.



Above: Around 80 people from all the SNIC Centres took an active part in the discussions.



## New Type of PRACE Preparatory Access Available Soon!

Michaela Barth, PDC

A new type of PRACE Preparatory Access (PA) will be introduced soon. It is called Type D “Tier-1 for Tier-0” and will provide researchers with initial access to a Tier-1 system plus access for a demonstration run on a Tier-0 system. Building on the popularity of PA Type C, Type D access will also provide researchers with enabling support from PRACE experts in Task 7.1. The goal of PA Type D is to provide a more effective and explicit migration path for researchers to move from using national level (Tier-1) high performance computing resources to using European level (Tier-0) resources. In addition, new computing resources were made available to researchers in the most recent PRACE PA call (which closed on the 13th of June).

Keep an eye out for calls for the new type of preparatory access, and for project access, at: <http://www.prace-ri.eu/call-announcements>. Speaking of project access, we are pleased to report that Sweden was successful in the 12th PRACE Project Access Call with a group based at the Linné Flow Centre, Swedish e-Science Research Centre, and lead by Philipp Schlatter being awarded 31,000,000 core hours on the MareNostrum system based at the Barcelona Supercomputing Center (BSC) in Spain for their project *Numerical experiments in a “virtual wind tunnel”: LES of the flow around a wing section at high Re*.



Above: The PRACE Sweden Face-to-face meeting was held at PDC on the 24th of May 2016. The main topics that were discussed were the current PRACE 2.0 business model draft (and its implications for Swedish researchers), along with presenting the PRACE EUDAT connection and the ongoing planning for the PRACE Spring School in April 2017 in conjunction with BioExcel.

## Staff Focus



Ragnar Sundblad

Ragnar Sundblad was borrowed from the CSC Systems Group to work with upgrading the long-term TSM storage system in the busy autumn of 2014, and ended up also working a bit with some of the many things that had to be done around the new systems, Beskow and Tegner, and the new Klemming file system.

Ragnar was moved over to PDC in January 2016 and now works permanently in the PDC Systems group, primarily with the long-term storage, the Tegner cluster, and as a contact person for the XENON Dark Matter project, but also with some of all the other things the PDC Systems group does.

Ragnar has been working with running computer systems at KTH for many years, and has also been involved in KTH internal and external projects, the last one building a new distributed frequency, phase and time distribution system for the Swedish Post and Telecom Authority.

He is more commonly known as “ragge”, which may or not be easier for a non Swedish-speaking person to pronounce. But the pronunciation, or denomination in general, means nothing - it is all about the intention!

## Staff Focus



Thor Wikfeldt

Thor Wikfeldt joined PDC in December 2015 as an application expert in molecular dynamics (MD). From the age of 5, Thor grew up in Iceland, but in 2002 he moved to Sweden to study physics at Stockholm University, and in 2011 he finished his Ph.D. in chemical physics. His research up to that point had focused primarily on trying to understand – from a molecular level – the many anomalous but life-supporting properties of ordinary liquid water, and in particular supercooled water, using a multitude of molecular simulation techniques. However, despite his profound understanding of its molecular properties, water still tastes the same to him.

After his Ph.D., Thor moved to London for a postdoctoral year at the University College London, where he employed path integral MD simulations to investigate nuclear quantum effects in various systems. A three-year stint as an independent researcher followed at the University of Iceland, which allowed Thor to further broaden his research experience.

In his free time Thor likes to spend time with his family, go out for a run in the forest or play a game of chess (against a human opponent!).

## VR-Cube Retrospective Seminar


Gert Svensson, PDC

The Virtual Reality Cube or VR-cube at PDC was a room to display computer generated stereoscopic 3D images. The VR-cube was a development of the original Cave design that was created in the Electronic Visualization Laboratory at the University of Illinois in Chicago by Carolina Cruz-Neira, Dan Sandin, and Tom DeFanti. The KTH VR-cube gave a realistic experience of actually being immersed in a virtual world. Our VR-cube was the first facility in the world to have a projection on all six surfaces of the cube, and the VR-cube also had a unique 3D sound capability.

The VR-cube at KTH was inaugurated in 1998 by the King of Sweden and soon became a major attraction for visitors to KTH. The cube was used both for scientific visualization and by artists and designers. For ten years until 2008 the cave was used in a number of projects: some were projects of a scientific nature and others were more artistic. The VR-cube was shut down 2008 as the technology was becoming old and the cost of operating the gigantic installation had become too high.

Recently Gert Svensson (from PDC) and Leif Dahlberg (from KTH Media Technology and Interaction Design) arranged a retrospective seminar about the VR-cube and its use by artists. The seminar was held on the 3rd of March 2016, and aimed to summarize the ten years of experience of the scientists and engineers collaborating with artists on projects involving the VR-cube. The seminar also had a secondary purpose, which was to look forward and discuss how science and art could collaborate in the future. Leif Dahlberg is responsible for creating a Centre for Art, Technology and Design, and also a doctoral program in these areas.

During the seminar Gert Svensson and Hans Hauska from KTH talked about the technical background of the VR-cube and about the experience of working with artists. The artists Teresa Wennberg, Birgitta Nord Nylund and Tomas Colbengtson (Bino and Cool) described their experience as artists working with this technology and together with engineers. The seminar ended with an open discussion on how to promote collaboration between science and art.



Above: Teresa Wennberg displays her VR-world "Brainsongs - Welcome to my brain" during the seminar.



## PDC Pub 2016

Thor Wikfeldt, PDC

The annual “PDC Pub and Open House” was held on the 3rd of May and was particularly well attended this year. The event started at 14:30 with drinks, snacks, some mingling and a welcome address by our Deputy Directory, Gert Svensson. This year’s Pub featured an inspiring mini-symposium in which a handful of young researchers who are using PDC’s high performance computing (HPC) resources gave fifteen-minute presentations on their scientific work. The symposium was a great success and we intend to continue this as a tradition in coming years.

Five speakers at the postgraduate and postdoctoral levels gave very interesting talks about their research - impressive work indeed! Joakim Halldin Stenlid went first and talked about his research based on electronic structure calculations of surface reactivity. Two talks followed on numerical simulations of airflow: first Elektra Kleusberg presented her research on wind turbines using direct numerical and large eddy simulations, and then Prabal Singh Negi talked about simulations of a pitching airfoil at high Reynolds numbers. After that Petter Johansson discussed his results from large-scale classical molecular dynamics (MD) simulations on water droplet spreading on surfaces, and finally Henrik Öberg talked about the fragment molecular orbital theory method and its application to enzyme catalysis.

All of the talks were followed by a number of questions from the audience, and it was apparent that ideas were exchanged between people working in rather disparate fields of research. Roughly 25 people attended the symposium and, although the meeting room was somewhat overcrowded, it was perfect for the occasion. PDC thanks all those who participated, and in particular all the speakers for their outstanding contributions!

After the symposium Gert started his much appreciated small-group guided tours of the PDC computer hall where the Beskow, Tegner and Milner systems are hosted, along with disk

*Below: Elektra electrifying the audience*



and tape storage devices. At the same time, the rest of the visitors and PDC staff enjoyed more snacks, sandwiches and drinks and participated in the pub quiz which featured 13 extremely tricky questions relating to HPC, PDC, and programming languages with references to computers in Hollywood blockbusters. It was an even race with four participants in shared second place, scoring six correct answers (which is statistically higher than a random distribution!), but the coveted prize of a luxurious lunch for two at the Syster O Bror restaurant went to Philipp Schlatter with eight correct answers (which is most likely more than the majority of staff at PDC would score). Congratulations to Philipp! The four silver medallists were awarded exclusive prizes in the form of elegant, yet practical, PDC coffee mugs. The Pub ended at 17:30, although some PDC staff stayed around longer and brought out something stronger from secret places at PDC...



*Above: Early stage system administrator training  
(Note: Bonus points for any readers who can decipher the hieroglyphics on the t-shirt!)*

## Our Ongoing Swedish SHAPE Projects:

### Towards a Greener and Quieter Future in Vehicular Transport

Michaela Barth, PDC

The PRACE (Partnership for Advanced Computing in Europe) Research Infrastructure runs a program to help European small to medium-sized enterprises (SMEs) take advantage of the innovation possibilities opened up by high performance computing (HPC) techniques and technology. This SME HPC Adoption Programme in Europe is known as SHAPE. The third SHAPE Call for Proposals was held earlier this year and is supporting eight projects from European SMEs, two of which are Swedish.

The two successful Swedish companies are receiving two person months (PMs) worth of support each from Task 1 in PRACE work package 7


Applications Enabling and Support, which is called Enabling Applications Codes for PRACE Systems. Support is also being provided from within the Swedish PRACE partner, that is, by the Swedish National Infrastructure for Computing (SNIC) with PDC acting in the executive and coordinating role.

Interestingly enough, in this round of the SHAPE calls, the Swedish companies that were awarded projects – Airinnova and Creo Dynamics – are both focusing on aerodynamic flow field simulations using accurate, life-like Computational Fluid Dynamics (CFD) simulations that will soon provide all of us with greener and much quieter methods of transport.



### Engineering Solutions and Innovations

Airinnova is a start-up company focusing on providing advanced computational technology for cutting-edge aircraft preliminary design, computational aerodynamics, and multi-disciplinary optimization.



*Above: Aircraft colour-coded by pressure coefficient plus starting vortex and wake development for the CRM configuration using RANS analysis with CFD code SU2, lift coefficient  $C_L=0.5$ , Mach number = 0.85, chord Reynolds number =  $5 \times 10^6$ , and Spalart-Allmaras turbulence model*

*Correctly predicted vortices and wake developments are relevant for correct drag prediction.*

The company is developing computational solutions for aerodynamic shape optimization, which is an important task in aircraft design. The goal is to design a lighter, greener, and quieter aeroplane by reducing drag especially at high speeds. Aerodynamic shape optimization for reduced drag requires a large number of CFD solutions, and computational power is a limiting factor. Current ideas for surrogate modelling are being developed to improve computational efficiency.

The aerodynamic shape optimization problem is highly non-linear. Gradient-based methods are comparatively efficient, but are myopic and deliver solutions – if at all – that are close to the baseline shape. Therefore, other methods (such as genetic algorithms or other non-local schemes) are also being investigated, albeit that they involve CFD analysis for many more candidate shapes. Accurate CFD runs (for example, the Reynolds-average Navier-Stokes equations or RANS solver with turbulence models) must be parallelized to reduce the run-time.

Within its SHAPE project Airinnova will focus on carrying out high-fidelity RANS aerodynamic shape optimization for the NASA common research model (CRM) aircraft based on the gradient-based optimization algorithms, using different CFD solvers and methods while getting essential benchmarking help in doing the performance analysis using the proposed CFD solvers.



Creo Dynamics, which is based in Linköping, is an engineering company with core competence in fluid mechanics, acoustics and structural dynamics. Creo participates frequently in national and international research programs where the focus is often on the development of new emerging technologies for the automotive or aerospace industry.

They have been awarded a SHAPE contract for their project proposal “Large scale aero-acoustic applications using open-source CFD”.


Together with researchers and application experts from the KTH Royal Institute of Technology and PDC, Creo engineers will work on the development and tuning of simulation processes based entirely on open-source software. The aim is to demonstrate an efficient and robust CFD workflow for real-life applications deployed in parallel at large scale.

From the side of Creo Dynamics, the project will be coordinated by Torbjörn Larsson, who was one of the main driving forces for introducing CFD in Formula One racing during the early 2000s and who was head of the CFD departments at BMW Sauber F1 and Scuderia Ferrari F1 for over a decade.

Torbjörn explains that “The chosen test case will be demanding, challenging and of high industrial relevance. A prime focus is on the implementation of efficient strategies for computer-aided design (CAD) handling and meshing in distributed mode on large parallel clusters. The PDC expertise should help us in identifying critical bottlenecks and performance deficits in the processes, and give guidance and assistance in finding improved solutions for a fine tuning of the overall methodology”. Results from the project will be made public to showcase the technology and encourage new users to adopt and invest in HPC and open-source software.

The envisaged CFD methodology is expected to deliver simulation accuracy in line with current industry best practices at a considerably lower cost.

PDC will help these two industrial business partners to assess how the use of HPC can increase their competitiveness in this exciting field of innovative business even more.



Above: For more information about the project, please visit CREO at <http://www.creodynamics.com>.



## DECI (Tier-1) Access

Michaela Barth, PDC

The latest PRACE Call for Proposals for HPC Compute Resources using Tier-1 systems (DECI 13) saw five successful Swedish projects.

Project & no. of hours	System the project will be run on	Principal investigator	Research area
MMIC 13,727,000	Salomon @ VSB Technical University of Ostrava	Natalia Skorodumova, KTH	Materials Science
PIPESUB 20,000,000	SiSu @ CSC – IT Center for Science Ltd.	Philipp Schlatter, KTH	Engineering
GraSiC 7,000,000	Archer @ EPCC (Edinburgh Parallel Computing Centre)	Nuala Caffrey, Linköping University	Materials Science
NTCPROJ 18,000,000	Archer @ EPCC (Edinburgh Parallel Computing Centre)	Ralf Döscher, Swedish Meteorological and Hydrological Institute	Earth Sciences
CHARTERED 16,800,000 + 6,720,000 + 84,000	Salomon @ VSB Technical University of Ostrava	Biplab Sanyal, Uppsala University	Materials Science

In total these projects have been awarded 82,331,000 PRACE standard CPU hours, which are equivalent to 20,582,750 CPU hours on Beskow. In addition, a Norwegian/Swedish proposal (FOPZOX, which is lead by Clas Persson from the University of Oslo) is being seen as an internal proposal and will get an additional 13,351,040 PRACE standard CPU hours.

Please note that the call for DECI 14 may be delayed since the DECI Allocation Assignment Committee needs to confirm that there are sufficient reviewers available for the increasing number of proposals – this is in order to avoid the kinds of delays that were experienced with the approval process for the DECI 13 call. This, and other matters relating to the next DECI call, should be discussed at the PRACE Council meeting on the 22nd of June, so further information about the call will be available here later on: <http://www.prace-ri.eu>.

## PDC-Related Events

### PDC Summer School 2016

15-26 August 2016, KTH, Stockholm

Further information about the summer school is available here:

<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.hpcuniversity.org/events/current/>

### HPCwire

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

### Linux Journal

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

### NeIC

<http://neic.nordforsk.org>

### 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>

### SeSE

<http://sese.nu>

### SNIC

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

### XSEDE

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

### NeIC Conference News

Join us in Umeå in 2017!

The 2017 NeIC Conference will run from the 29th of May to the 1st of June in Umeå, Sweden. Further details will be available from:

<http://neic.nordforsk.org>.