no.1 2020

PDC Center for High Performance Computing

PDC Newsletter

The PDC Newsletter is published by PDC at EECS, 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 EECS, KTH, is one of the six centres in the Swedish National Infrastructure for Computing (SNIC).

Editor: Patrick Norman
Editorial Staff: Genet Edmondson
Layout: Genet Edmondson
Graphic Design: Maria Malmqvist
E-mail: support@pdc.kth.se

ISSN 1401-9671

Cover

The cover image is the logo of the new pre-exascale pan-European computer, LUMI, that will be based in Finland and available for research from early next year. LUMI stands for "Large Unified Modern Infrastructure" and the word "lumi" means snow in Finnish.

How to Contact PDC

Visiting address:
Teknikringen 14, "Plan 4",
KTH, Stockholm
Mail: PDC, KTH,
SE-100 44 Stockholm, Sweden
E-mail: support@pdc.kth.se
www: http://www.pdc.kth.se
Phone: +46 8 790 7800

Editorial

After some ten years of service as director of PDC, Erwin Laure has decided to leave his post and continue his career as director of the Max Planck Computing and Data Facility in Germany. PDC and the Swedish high-performance computing (HPC) community gratefully acknowledge his service and wish to express our very best wishes in regard to his future career. It is difficult and perhaps even unfair of me as a newcomer to summarize Erwin's achievements in a few sentences, and there are certainly people from within the PDC organization that would be better suited to do so. But, as a long-term user of the Swedish HPC infrastructure and a previous director of the National Supercomputer Centre (NSC), I feel able to at least provide a biased view of Erwin's work at PDC. As a user and program developer, I have been challenged by his profiling PDC as the leading Swedish HPC centre for large-scale parallel computing which has served as an impetus for us to rid ourselves of code legacy issues and start over with a new software project under the name of VeloxChem as presented in the PDC Newsletter No. 2, 2019 (https://www.pdc.kth.se/publications/pdc-newsletter-articles/2019-no-2/veloxchem-quantum-molecular-modelling-in-hpc-environments-1.945304). As a peer director of NSC, I felt that Erwin embodied Swedish HPC on the European arena and his engagement in European projects made PDC stand out as the most international among the Swedish supercomputer centres. With the Swedish participation in the EuroHPC project and the LUMI system, I believe that Swedish users will increasingly benefit from European HPC services in the coming future and that Erwin has made an important contribution to prepare us (HPC users and centres alike) for this change. His article titled "EuroHPC – Europe's Path to Exascale" in the PDC Newsletter No. 1, 2019 (https://www.pdc.kth.se/publications/pdc-newsletter-articles/2019-no-1/eurohpc-europe-s-path-to-exascale-1.911735), is a worthwhile read in this regard.

While the process of finding a long-term replacement for Erwin goes on, I have been asked to function as interim director of PDC. My stake in HPC stems from research in the field of theoretical chemistry where we conduct method and program development for the purpose of simulating light–matter interactions in complex molecular systems. I came to the KTH Royal Institute of Technology in 2016 and became the new head of the Department of Theoretical Chemistry and Biology (TCB), which is now part of the School of Engineering Sciences in Chemistry, Biotechnology and Health (CBH). My HPC background includes serving as the director of NSC for four years and being a member of the PDC board for seven years, which are experiences that have prepared me with a fair insight into the operations of PDC and the ability to bridge the gap in between directorships. Equally importantly, this background has also put me in contact with several individuals of the highly skilled staff at PDC and, upon arriving, I at once felt warmly welcomed and everyone has given me fantastic support so as to make the best of the situation. Apart from issues regarding the standard operations of a supercomputer centre, the single most important task for me to attend to is the replacement of the Beskow cluster. This is an important milestone not only for PDC but for Swedish HPC as a whole and it is fortunate that the procurement project is led by a person with vast experience, namely Gert Svensson. So, in regard to points that PDC can largely control, I am comfortably guided by experienced co-workers, but there are also points that are beyond our control – such as the effects of the global pandemic on hardware production and delivery, exchange rates, and so forth – that are worrisome. The only guarantee I can provide is that all the staff members at PDC will do their utmost to bring a science-enabling system into operation (in two stages) to serve you. Stage 1 (a CPU-based cluster) is scheduled to accept SNAC allocations by July 2021 and Stage 2 (with a higher capacity) should be installed within nine months thereafter.

As impressive as we hope the capacity of PDC's new main cluster resource will turn out, it will not reach the extreme performance levels of the resources planned by the EuroHPC Joint Undertaking. In the cover article, Anni Jakobsson and Pekka Manninen describe the LUMI system which is the Finnish-based pre-exascale EuroHPC system scheduled for operation in 2021. Sweden has a share percentage of the LUMI system and the process for Swedish researchers to get access to it is described.

In light of the global pandemic, many scientists are devoting their efforts and expertise to improve on the difficult situation and researchers in e-sciences are no exception in this sense. In this newsletter you can read about related activities of researchers from the BioExcel and SeRC centres.

Despite the upheavals locally and globally, PDC has continued with its education and outreach programme by taking our courses, seminars and other events, like the PDC cafes, online where possible, although we were able to hold a PDC-SeRC seminar on climate modelling and a system introduction course before restrictions came into force. Regrettably, this year's PDC Summer School has needed to be postponed until 2021 and we are keeping our fingers crossed that it will be possible to have face-to-face meetings and discussions with a new batch of bright students at that time.

Patrick Norman, Interim Director PDC

In This Issue

Staff Focus

Jonathan Vincent

Jonathan Vincent is now rejoining PDC after working at NVIDIA for four years. He will be working as a researcher improving the graphics processing unit (GPU) implementation of the NEK5000 code. During his time at NVIDIA, Jon was mainly involved with writing and optimizing CUDA code for scientific applications, most recently GROMACS.

In his previous time at PDC he was an application expert and head of user support. Prior to that, Jon worked with meteorological codes as a computational scientist in Reading, after post-doctoral studies in Uppsala, Santiago de Compostella, Leiden and Gothenburg. Jon's qualifications include a Ph.D. in semiconductor physics and a B.Sc. in theoretical physics, both from the University of Exeter, UK. Outside of work, Jon enjoys watching sport, playing computer games and reading.

Above: PDC-SeRC Seminar on next-generation climate modelling by Professor Thorsten Mauritsen, SeRC room, PDC, 19 February 2020

Sweden, Get Ready for LUMI!

Anni Jakobsson & Pekka Manninen, CSC - IT Center for Science

The European High-Performance Computing Joint Undertaking (EuroHPC JU) is pooling European resources to develop top-of-the-range exascale supercomputers for processing big data. One of the pan-European pre-exascale supercomputers will be located at the data centre belonging to the Finnish IT Center for Science (CSC) in Kajaani, Finland. This new system is called LUMI, which comes from Large Unified Modern Infrastructure.

The new supercomputer will be hosted by the LUMI Consortium, which consists of nine countries: Finland, Belgium, the Czech Republic, Denmark, Estonia, Norway, Poland, Sweden and Switzerland. The plan is for the LUMI system to be in operation by early 2021.

Preparations at CSC's data centre in Kajaani are in full swing as 2021 will be here in no time. Researchers using computational methods should also start thinking about how to best exploit LUMI's capabilities. However, to do that, you will need to know what kind of supercomputer will LUMI be. Let's go through some of the details of LUMI's architecture.

LUMI: One of the Most Competitive Systems in the World

When LUMI goes into operation at the beginning of 2021, it will be one of the most competitive supercomputers in the world. The design philosophy for LUMI was to create a platform that would make it possible to use Artificial Intelligence (AI) techniques (especially deep learning) and traditional large-scale simulations combined with high-performance data analytics to solve a single research problem.

The theoretical peak performance of LUMI will be over 200 petaflops (2×10^{17} floating point operations per second) – that will make the new system approximately ten times faster than the Piz Daint supercomputer, which is the fastest in Europe at the moment. The performance of the applications on LUMI will likewise be faster than those on Piz Daint by a factor of about 10. LUMI will achieve such a high level of performance thanks to having a large number of nodes with accelerators (otherwise known as Graphics Processing Units or GPUs).

In addition to the GPUs, LUMI will have a number of supporting compute and storage resources, which will maximize the overall value: the system will be complemented by a partition consisting only of Central Processing Units (CPUs), Infrastructure as a service (IaaS) cloud services and a storage solution for large objects.

LUMI will have a highly capable parallel file system: it will consist of over 60 petabytes of storage with a sizeable flash layer for input/output (I/O). The flash layer is expected to have a volume of

Below: LUMI architecture

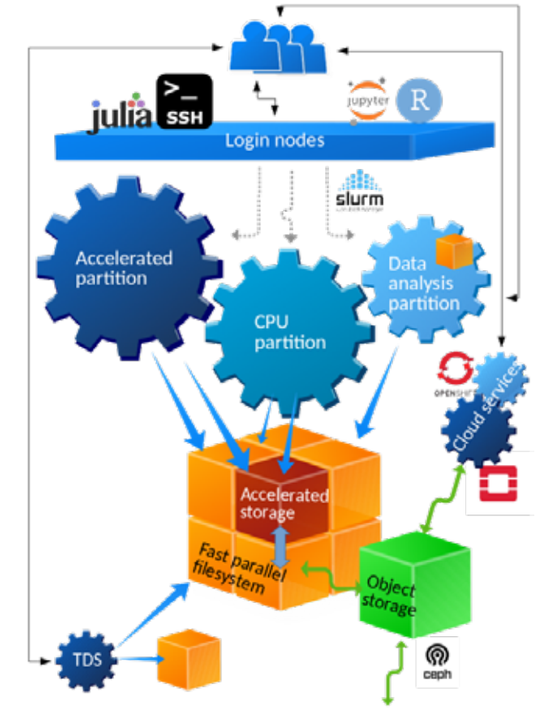

around 5 PB and will provide a bandwidth of more than 1 terabyte/s and an extreme input/output capability per second. The object storage, which is anticipated to be 30 PB in size, will be there to facilitate convenient management of data.

LUMI: A Leadership-Class AI Platform

LUMI's massive computational capability, which is based on GPUs and extreme connectivity between nodes, makes LUMI a leadership-class platform for AI-based research. The I/O capabilities of LUMI make it possible to use, and perform computations on, big datasets. As LUMI will have both GPU and CPU processors, it will be possible to mix CPUs and GPUs in the same workflow, even for the same application, thus making it possible to combine AI methods and simulations.

Easy-to-Use, Rich Stack of Pre-Installed Software for Research

Now that you understand the overall architecture of LUMI, we will look at the software that is likely to be available on LUMI.

In addition to offering traditional command-line interfaces, the intention is that LUMI will support high-level interfaces. Applications such as Jupyter Notebooks, RStudio and Matlab will be seamlessly integrated into the back-end of LUMI, so that ultimately LUMI will be an extension to researchers' laptops. LUMI will have a rich stack of pre-installed software, which will include code developed by research communities and also commercial applications. Furthermore, LUMI will feature a collection of reference datasets that will be readily available and curated (Datasets as a Service). The system will also provide capabilities for the interactive visualization of results during the execution of simulations.

The plan is that LUMI will be in operation next year, that is, early in 2021. Research communities should, therefore, start thinking now about how to utilize the significant increase in computing capacity that LUMI will provide, and should start to prepare software for LUMI's architecture. Software that currently uses only conventional processors can be modernized to employ GPUs, or alternatively, research groups may prefer to start using software that already makes use of GPUs with similar functionality.

Eco-Friendly Data Centre

Reducing CO_2 emissions is a globally critical goal, and the location of the EuroHPC machines has a huge impact on that, as supercomputers consume plenty of electricity. Supercomputers of LUMI's size would generate 50,000 tons of CO_2 emissions annually if they were powered by fossil-fuel-based electricity. However, the CSC data centre in Kajaani uses renewable and CO_2-neutral electricity.

Furthermore, LUMI will use warm-water cooling, which enables its waste heat to be utilized in the district heating network of Kajaani, and thus can replace heat that would otherwise be produced by fossil fuels. The amount of waste heat from LUMI that will be used in Kajaani's district heating network will be equivalent to up to 20 percent of Kajaani's energy needs in terms of the area's district heating. The reuse of this waste heat

Staff Focus

Tianzong Wang

Tianzong Wang is a first-year master's student specializing in Machine Learning at KTH. Before he moved to Sweden, Tianzong completed his bachelor's degree in statistics and computer science at the University of Hong Kong, China, and worked as an actuary in a Dutch insurance firm based in Hong Kong. He is currently working on computer music generation projects using deep neural networks and machine learning techniques. At the moment, Tianzong is working part-time in first-line support at PDC, taking care of account management and addressing support requests. For him, PDC is not only a job but also a place where he is learning about Swedish culture. He celebrated Saint Lucy's Day, Shrove Tuesday (The Swedish name literally means "Fat Tuesday" and the celebrations feature delicious buns filled with almond paste and cardamon-flavoured cream!) and the success of a colleague's Ph.D. defence with the team and was very happy to be a part of it.

Tianzong loves nature and, during his free time, he and enjoys working out activities such as hiking and playing basketball. He also likes algorithmic trading and is the founder of a small trading firm.

will reduce the annual CO_2 footprint of Kajaani by 13,500 tons: an amount equivalent to that of 4,000 passenger cars!

An interesting point to note about the name LUMI is that, although it is the acronym for "Large Unified Modern Infrastructure", as a literal word, "lumi" means snow in Finnish. So "snow" will be heating Kajaani from next year onwards.

Who Can Access LUMI?

Now you may be wondering who will be eligible to use LUMI and how researchers can apply for time on the system.

Half of the LUMI resources belong to the EuroHPC Joint Undertaking, and the other half of the resources belong to the participating countries, that is, the LUMI Consortium countries. Each consortium country will be given access to a share of the resources with the size of the share being based on that country's contribution to the LUMI funding. For Sweden, the share will be 4% of the LUMI resources annually. The shares of each of the countries will be allocated according to local considerations and policies within each country – so LUMI's resources will be seen and handled as extensions to national supercomputing resources.

The share of the LUMI resources belonging to the EuroHPC JU will be allocated by a peer-review process (comparable with applications for PRACE Tier-0 access). In particular, up to 20% of the EuroHPC JU resources will be available to researchers in the industrial sector, from both large companies and small to medium-sized enterprises (SMEs).

So, what does this mean for a Swedish researcher? A researcher affiliated with either a Swedish research institute or a company that has its headquarters in Sweden will be able to apply for LUMI resources via Sweden's national share and, in addition, via EuroHPC's peer-review process. The means that Swedish-based researchers have access to 4% of the total LUMI resources and can additionally apply for the EuroHPC resources via a technical and scientific peer-review process.

By partnering with European research groups and/or companies, the EuroHPC share of the LUMI resources will be available to non-European researchers as well. Note that the principal investigator (PI) for any project applying to use LUMI will need to be based in the EU or an associated country.

The peer-review and application processes for LUMI are still being negotiated so further details will be announced later nearer the inauguration of LUMI.

Note that the resources of LUMI will be allocated for the projects in terms of three different pools: GPU-hours, CPU-hours, and storage-hours. All users of the system will have access to the whole system as per batch job policies; in other words, none of the hardware will be dedicated or reserved specifically for any of the partners.

What Will Be Solved with LUMI?

The creativity of the research communities in the coming years will determine the kinds of problems that will be solved with LUMI. We foresee that, for example, LUMI will help remarkably in the following kinds of research areas.

- More precise climate models and the interconnection of different existing climate models: How will living conditions change as the climate is warming?
- The sequencing and analyzing of full genomes, combined with data analysis and correlations to clinical data, will shed more light on diseases and particularly hereditary diseases, thus helping to identify causes of illness and determine personalized treatment and medicine.
- Artificial intelligence (deep learning): analyzing large data sets (either from simulations or measurements) and re-analyzing such data, for example, in atmospheric science, environmental science, climate modelling, materials science and linguistics
- Self-driving cars and vessels: the study and development of algorithms related to these modes of transport taking advantage of LUMI's previously unprecedented computing power to make this type of transport feasible
- Social sciences: large-scale data set analytics from social networks and the modelling of different phenomena

What are you going to discover with LUMI? To get ready for LUMI, you are welcome to subscribe to the LUMI newsletter and follow the latest news about LUMI at https://www.lumi-supercomputer.eu, or you can follow LUMI on social media:

Twitter: @LUMIhpc
LinkedIn: LUMI supercomputer
YouTube: https://www.youtube.com/channel/UCb31KOJ6WquosRpIRi_k8Mw
#lumisupercomputer #lumieurohpc

bioexcel

Would you like some assistance with your COVID-19 research?

Any academic or industry researchers who are working on modelling and simulating systems related to SARS-CoV-2 are invited to get in touch with BioExcel at

info@bioexcel.eu

to discuss how we can assist your research. We can provide expert support in the areas of

- biomolecular integrative modelling,
- molecular dynamics simulations,
- free energy calculations, and
- docking.

In particular, we welcome researchers using the applications that BioExcel supports explicitly:

- HADDOCK,
- GROMACS,
- PMX, and
- CP2K (hybrid-QM/MM).

In addition, all our partner supercomputing centres

- the PDC Center for High Performance Computing (Sweden),
- Barcelona Supercomputing Center (Spain),
- EPCC (UK) and
- Jülich Supercomputing Centre (Germany)

are offering priority access to HPC infrastructures for COVID-19 research.

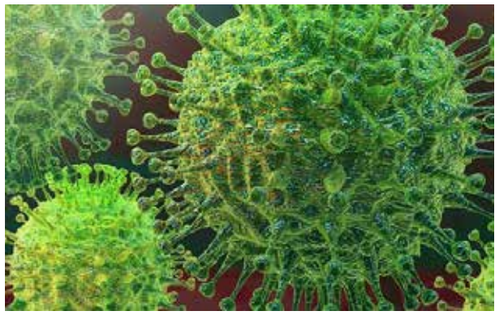

BioExcel Supports Range of COVID-19 Research Activities

Rossen Apostolov, PDC

The BioExcel Centre of Excellence for Computational Biomolecular Research provides support to academic and industrial researchers in the areas of molecular dynamics simulations, biomolecular integrative modelling, free energy calculations and docking. In recent months BioExcel has launched a series of activities to support research on the SARS-CoV-2 virus to assist with the global fight against COVID-19.

The ongoing pandemic associated with the virus has already affected most of the world in a dramatic way. It endangers the lives of many people, and puts a tremendous strain on medical systems, while the second-order effects on the economy and daily life have been equally devastating. Governments and the private sector are undertaking massive concerted efforts to tackle the crisis. Unquestionably, the discovery of a vaccine or cure will be of paramount importance when it comes to controlling the spread of the illness in the long-term.

To this end, BioExcel experts are partnering with numerous international initiatives to lend all of our advanced software applications and expertise to the search for a vaccine or cure and BioExcel's partners are actively working on a range of modelling and simulation projects related to SARS-CoV-2.

Sharing data for COVID-19 applications has become more vital than ever due to the urgent need to identify compounds (known as “leads”) which have pharmacological or biological activity that means they could be suitable for developing therapies, diagnostics, and vaccines for COVID-19. To maximize the impact of molecular simulation methods in this crisis, BioExcel members have signed the international community letter which aims to connect researchers and improve communication between simulation, experimental and clinical data investigators.

COVID-19 Molecular Structure and Therapeutics Hub

The BioExcel Centre of Excellence is partnering with the Molecular Sciences Software Institute (MolSSI) to set up a new community-driven data repository and curation service for molecular structures, models, therapeutics, and simulations related to computational research associated with therapeutic opportunities for COVID-19.

https://covid.bioexcel.eu

https://covid.molssi.org

The data can be used as a reliable starting point for further studies by research groups worldwide. Researchers are invited to join this endeavour by contributing any data (for example, small molecules, models, simulation data, or new target compounds) which may be useful to the wider research communities. Researchers can also join the data curation teams to review other groups’ submissions. For more information, see:

https://covid.molssi.org/contributing.

Nostrum Biodiscovery (NBD) is focusing on screening its proprietary virtual library, ChemistriX, as well as other open libraries (like Zinc), into the 3C main protease of SARS-CoV-2; blocking this enzyme would mean the virus could not replicate efficiently. For this reason, NBD is using a hierarchical docking approach combining Glide (a docking tool from Schrödinger) with PELE, NBD’s proprietary Monte Carlo code which is a computational protocol that has shown remarkable results in international blind competitions. In addition, this approach is highly parallel and can take advantage of the supercomputer facilities at the Barcelona Supercomputing Center (BSC). The work is being done in collaboration with the Electronic and Atomic Protein Modelling Group led by Victor Guallar at BSC, and contributes altruistically to the EXSCALATE4CoronaVirus consortium. Within the same consortium, the KTH Royal Institute of Technology is providing GROMACS and consultancy expertise to meet the needs of researchers who have to tailor molecular dynamics simulations for large-scale executions.

The Institute of Neuroscience and Medicine (INM-9) at Jülich Research Center (JSC) is using local supercomputing resources to identify ligands (molecules) that bind to proteins contributing to the coronavirus-host interactions, which are key to viral pathogenesis. The team (together with CINECA, Bologna, Italy) is currently performing virtual ligand screening experiments (in parallel on different proteins) on the JSC high-performance computing (HPC) systems. The resulting libraries of molecules will then undergo in-vitro testing.

Researchers from the University of Jyväskylä (JYU), Finland, are performing large-scale molecular dynamics simulations of the spike protein from the SARS-CoV-2 virus binding to the human ACE2 receptor as a first step to understand the interactions that control host cell recognition in lung tissue. (The spike protein is the main protein that the virus uses to invade human cells. It binds to another protein, known as a receptor, which is part of the human cell membrane, then the viral membrane fuses with the human cell membrane, allowing the genome of the virus to enter human cells and begin infection.) The group is also investigating the spike protein complexed with aptamer candidates to systematically search for oligonucleotides that could selectively bind to the spike protein and thus prevent it from binding to a human cell. (Aptamers are oligonucleotide or peptide molecules that bind to a specific target molecule.) The large-scale simulations are being performed using custom workflows, which are running on HPC resources at CSC in Finland. Moreover, the workflows and protocol used in this project will be readily available if future outbreaks of other infections occur. The project involves a team of specialists in chemistry, molecular modelling, and infectious diseases as well as in aptamer selection and modification. Collaborators include researchers from Moscow State University and the Russian Academy of Sciences, who have a strong track record in developing DNA aptamers for diagnostic and therapeutic purposes.

In a parallel project, JYU researchers (including Petri Pihko, an expert in organic synthesis, and Varpu Marjomäki, a virology specialist) are working on inhibiting the SARS-CoV-2 RdRP (which is the enzyme that assists the virus to replicate its RNA) with nucleotide analogues. The RdRP structure was recently resolved to a fine degree and, since then, the group has been performing simulations to investigate the viability of various potential approaches that might stop the virus from replicating.

Below: Docked complex between DNA aptamer and the RBD of the SARS-CoV2 spike protein

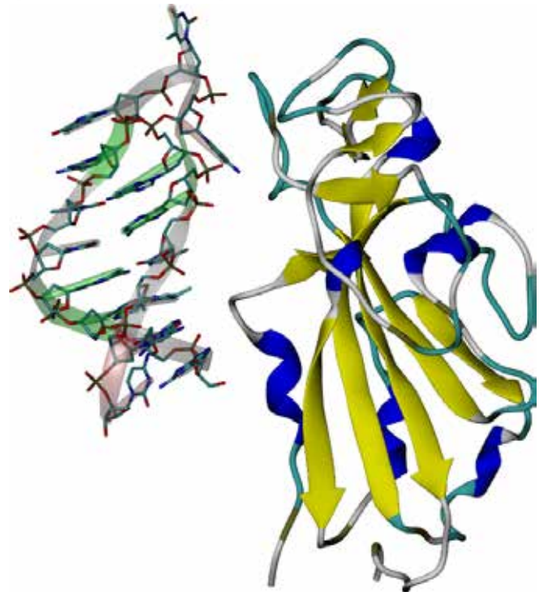

A team at the Max Planck Institute for Biophysical Chemistry, Göttingen, has started to simulate the complex of the SARS-CoV-2 coronavirus spike protein bound to the human ACE2 receptor using a combination of GROMACS and PMX. Their goal is to design a derivative of ACE2 with enhanced affinity for the SARS-CoV-2 spike protein for either diagnostic or therapeutic purposes.

The Worldwide e-Infrastructure for NMR and Structural Biology (WeNMR) portal, developed and run by BioExcel's partners at Utrecht University, has seen an increase in registrations over recent weeks with many users indicating they intend to use it for COVID-19 projects. The HADDOCK WeNMR team is already involved in several collaborations ranging from drug screening against the protease to modelling COVID-19-related protein-protein

interactions. For this purpose, together with EGI/EOSC experts, the team is looking both into expanding the processing capacity of the HADDOCK portals and providing customized solutions to support researchers. These might take the form of dedicated virtual clusters with a HADDOCK front end running on EOSC cloud resources, and customized virtual machines with ready-to-run local HADDOCK installations for experienced users wanting to use the software at the command line.

The Institute for Research in Biomedicine Barcelona (IRB Barcelona) is working on understanding how the virus has evolved by comparing its structure/genome to other corona-viruses, including the earlier SARS-CoV. Their work, which includes MD simulations, as well as bioinformatic analysis, aims to reveal how the virus evolves and how the hosts are selected by identifying key structural determinants and pathogenic mutations that have occurred going from one viral strain to another - which is also of utmost importance for future pandemic surveillance. The information from these studies will be used to identify virus inhibition opportunities by means of in silico drug screening, starting from known and commercially available drugs first. This work is being undertaken in collaboration with Roberto Burioni's laboratory for viral research in San Raffaele, Italy.

Save the date: Nordic-RSE2020
1-2 December 2020, Stockholm

The first Nordic-RSE conference will bring together those who develop software for research purposes and contribute to building a research software engineer (RSE) community in the Nordic area. The program will contain invited talks, lightning talks, poster sessions and workshops on best practices for creating research software, as well as other popular RSE-related topics. For further information, see:

https://nordic-rse.org/conference

https://coderefinery.org/blog/2020/04/24/nordic-rse-conference

SNIC Accepted as a PRACE Training Centre

Lilit Axner, PDC

Between 2010 and 2020, PRACE established ten PRACE Training Centres (PTCs) in the Czech Republic, Finland, France, Germany, Greece, Ireland, Italy, the Netherlands, Spain and the United Kingdom to make PRACE training widely available across Europe. These centres collectively provide a coordinated programme of training courses about high-performance computing (HPC) and computational science for European researchers.

The demand for PRACE training courses has increased over time. Therefore, PRACE decided to increase the number of training centres. In October 2019, PRACE opened a call to establish additional PTCs. The new training centres would be required to implement a collective training programme jointly coordinated by PRACE, where each PTC would focus on subject areas that would be of particular relevance or interest to industry and/or the research communities in their local area and in the surrounding region. Each training centre would be expected to deliver two or three training events annually, totalling up to nine training days per year.

Usually, the PRACE PTC courses need to fulfil the following requirements.

- Every PTC course is open to all researchers, free of charge (for participation).
- Course material is mostly in English and is made available on the PRACE Training Portal where feasible.
- Courses are delivered verbally in English to ensure accessibility, although verbal instruction may be given in a local language as long as there is sufficient justification to do so.
- Each PTC course is set up on the PRACE website so there is consistency in terms of providing course details, and collecting information (for example, for registrations or evaluations).
- Each PTC course is usually 1-3 days in duration, which means a PTC programme should not consist of a series of half-day courses.

SNIC with three of its centres – PDC at the KTH Royal Institute of Technology, the National Supercomputer Centre (NSC) at Linköping University (LiU) and the High Performance Computing Centre North (HPC2N) at Umea University (UmU) – applied to the PRACE call to become a PRACE training centre. In January 2020, PRACE made its decision and four more countries now have PTCs:

- Austria (VSC Research Center, Vienna University of Technology),
- Belgium (the Flemish Supercomputer Center, VSC, & the CÉCI consortium of Belgian HPC centres),
- Slovenia (HPC Center Slovenia), and
- Sweden (SNIC with PDC, NSC and UmU).

The SNIC PTC has now prepared its preliminary training plan for the spring 2020 - spring 2021 period. Please note that due to COVID-19 restrictions, the training courses are currently being run online. In the future, we hope to hold these training courses both face-to-face and online, taking into account the success of recent and current online training. Be aware that this list is preliminary and the dates for the autumn 2020 - spring 2021 courses may change.

2020	
27 May	Writing Parallel Applications Using MPI (SNIC-KTH)
27 May	Scientific Computing in R (SNIC-UmU)
10 June	Working effectively with HPC systems (SNIC-LiU)
2-4 Sept	BioExcel GROMACS workshop (SNIC-KTH)
September	OpenMP for accelerators (may be postponed due to COVID-19) (SNIC-KTH)
October	Heterogeneous computing with GPUs and performance modelling: understanding the modern high core count CPUs and GPUs (SNIC-UmU)
19-20 Oct	VASP workshop (SNIC-LiU)
Oct/Nov	CodeRefinery workshop (SNIC-KTH)
2021	
Feb/March	Introduction to Artificial Intelligence and Deep Learning (SNIC-KTH)
March/April	CodeRefinery workshop (SNIC-KTH)
20-22 April	Introduction to R in an HPC environment (SNIC-KTH)
April/May	NEK5000 workshop (SNIC-KTH)

The confirmed courses are announced on the web pages of each SNIC centre as well as on the PRACE training events web page. More information is available on the following pages.

- https://events.prace-ri.eu/category/2
- https://www.pdc.kth.se/about/events
- https://www.hpc2n.umu.se/events/courses
- https://www.nsc.liu.se/support/Events

Staff Focus

Muhammed Abdullah Al Ahad

Ahad is a Master's student in Computer Science at KTH. This is his second master's degree; his first was in Scientific Computing (also from KTH). Ahad is studying Computational Fluid Dynamics, high-performance computing, and numerical analysis amongst other courses. During his first master's, Ahad had a summer job at the Swedish company ABB where he developed an application for fluid simulation through an impinging jet for the cooling process on a Run Out Table (ROT). Ahad has also worked with the INTERTWinE project on developing the Global Address Space Programming Interface (GASPI).

Currently Ahad is working at PDC in first-line support and is also involved in the BioExcel project working on creating GROMACS containers with all their dependencies.

During his free time, Ahad likes to watch movies, play cricket and travel with his family.

Above: Introduction to PDC Systems Course, 5 February 2020

How to Acknowledge the Use of PDC Resources in Your Research

Gilbert Netzer, PDC

If you utilise any of PDC's resources (for example, computer systems, data storage or assistance from application experts) in your research, we ask that you acknowledge PDC in an appropriate way.

Here are the recommended acknowledgement formats from various organisations that grant researchers access to PDC's resources. The following information is also available on PDC's website in the Research section in case you need to check it later: www.pdc.kth.se/research/acknowledgement-of-using-pdc-resources-for-research-1.918148.

Access to PDC Resources through SNIC

> *"The [computations/data handling/SIMILAR] were/was enabled by resources provided by the Swedish National Infrastructure for Computing (SNIC) at the PDC Center for High Performance Computing, KTH Royal Institute of Technology, partially funded by the Swedish Research Council through grant agreement no. 2016-07213"*

If applicable, also add an acknowledgement for application support using either of the following.

> *"[NAME] at PDC is acknowledged for assistance concerning technical and implementational aspects [OR SIMILAR] in making the code run on the [OR SIMILAR] PDC resources."*

> *"We thank [NAME] at PDC for [his/her] assistance with [describe tasks such as porting, optimization, etc.], which was made possible through application support provided by SNIC."*

The above text is based on the recommendations at www.snic.se/allocations/apply4access (which were current as of 25 May 2020).

Access to PDC Resources through PRACE

> *"We acknowledge PRACE for awarding us access to [Beskow/Tegner/...] hosted by the PDC Center for High Performance Computing, KTH Royal Institute of Technology, Sweden."*

The above text is based on the recommendations at https://prace-ri.eu/hpc-access/project-access/project-access-information-for-awarded-projects and https://prace-ri.eu/hpc-access/preparatory-access/preparatory-access-information-for-awarded-projects (current as of 25 May 2020).

Access to PDC Resources Granted Locally by PDC

> *"We thank the PDC Center for High Performance Computing, KTH Royal Institute of Technology, Sweden, for providing access to the [computing resources/application support/...] used in this research."*

PDC-SeRC Seminars

Thor Wikfeldt, PDC

The closure of the KTH campus due to the COVID-19 pandemic, unfortunately, ended our PDC-SeRC seminar series prematurely this semester, but we did have an excellent seminar on the 19th of February by Professor Thorsten Mauritsen from the Department of Meteorology (MISU) at Stockholm University. Thorsten told a SeRC room filled with listeners about next-generation climate modelling with the ICON code. Climate science seems to touch a nerve with almost everyone as, after the talk, the audience was brimming with questions!

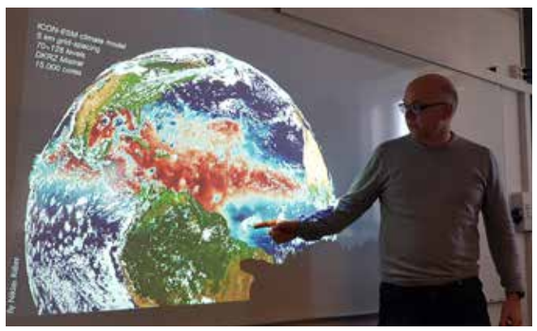

Above: PDC-SeRC seminar, 19 February 2020

New Major HPC System for PDC

Gert Svensson, PDC

As reported in the previous PDC Newsletter, PDC received a substantial grant from the Swedish National Infrastructure for Computing (SNIC) to install a new general-purpose high-performance computing (HPC) system for academic research. The process of procuring this system (which will replace PDC's current flagship system, Beskow) is now well underway. The plan is that the new system will be installed in two steps: the first part of the system (phase one) is expected to be delivered early in 2021 and the second part (phase two) is planned to be in place in early 2022 at the latest.

Background

The new supercomputer system at PDC is intended for a wide range of academic research use. It will be able to execute highly parallel jobs using a large number of nodes, as well as jobs using a single node or a small number of nodes, in an efficient way. The new system will have a partition using only central processing units (CPUs), and another partition which will be equipped with graphics processing unit (GPU) accelerators or high-speed CPUs with some properties similar to accelerators. This second partition will also be suited to Artificial Intelligence/Machine Learning (AI/ML) workloads, especially in combination with HPC simulations. In addition, the system will include a fast Lustre storage subsystem.

The capacity of the CPU-module of the new system will replace that of the following SNIC systems: Beskow and Tegner at PDC, Aurora at Lund University and Hebbe at Chalmers University of Technology. All of these systems will be retired from SNIC duty when the new system is in operation at PDC.

Sweden is a part of the LUMI pre-exascale consortium, so researchers who plan to run their codes on the LUMI system (or other pre-exascale or exascale systems) will be able to use the new system at PDC as a stepping stone. Development and testing of new codes could partly be done at PDC while the LUMI system would be available to execute extensive simulations exceeding the capabilities of the SNIC systems. Both computer systems are likely to have a substantial GPU-partition. However, due to the different time schedules for procuring the two systems, the type and manufacturer of each of the systems may differ.

The budget from SNIC for the new system at PDC is 129 million SEK in Total Cost of Ownership (TCO) over five years. This includes the purchase price, plus the costs for the installation, maintenance, power and cooling. In addition to that, SNIC will provide funding of 41 million SEK for system administration and running expenses for the computer hall for five years.

For many years the KTH Royal Institute of Technology and PDC have been involved in an extensive research collaboration with Scania. In the PDC part of the project, HPC simulations are used to improve the efficiency of vehicle designs. In 2017 the capacity of the Beskow system was extended significantly, thanks to the Scania collaboration. Future funding arising from this collaboration has the potential to increase the size of the new system at PDC substantially, which would mean that larger simulations could also be performed by academic SNIC researchers.

The Procurement Process

As mentioned previously, the procurement process is in full swing. At the time of writing, the initial invitation had been published and submitted to a range of potential vendors. In the invitation, we have roughly described the type of system we intend to purchase, together with some requirements that any companies submitting tenders would need to satisfy.

Once we know which vendors are interested in bidding for the new system, we provide those companies with detailed specifications and benchmarks (programs the vendors need to run on the system that they propose selling to us). The benchmarks represent a typical workload for the system and have been carefully selected with the help of a scientific reference group. As we

mentioned earlier, the system will be installed in two phases, and the same company will be chosen to provide the hardware for both phases (so there is only one company to deal with for maintenance and trouble-shooting).

Details of Phase One

The first phase of the new system will consist of a CPU module and a disk module. The CPU module should reach a total High-Performance Linpack (HPL) benchmark performance of at least 2 PFLOPS. The nodes will have Intel x86-64 compatible CPUs and will probably have two CPUs per node. There will be a range of nodes with the same architecture but with differing amounts of memory. Most of the nodes will have 256 GB of memory and will be known as thin nodes, but some nodes will have more memory as shown in the table.

Name of nodes	Memory	Number of nodes
large	512 GB	20
huge	1 TB	8
giant	2 TB	2

Our market analysis indicates that the number of cores per CPU is increasing and we expect in the order of 70-100 cores per node.

The disk module will support a Lustre high-speed file system, similar to Klemming, but with a capacity of at least 7 PB. The metadata disk of the disk module will be completely based on solid-state drives (SSD), which will increase the speed for metadata operations (like creating or deleting files). Phase one of the system is expected to be delivered in the first quarter of 2021.

Details of Phase Two

Phase two of PDC's new system will consist of a further module that is expected to be GPU-based (although it might be based on CPU technology that provides the same level of performance as GPUs). We are allowing for up to four GPUs per CPU. Each CPU should have a memory of 128 GB per GPU attached to the CPU, and each GPU should have at least 32 GB of memory. With the current budget, we expect to reach an HPL performance of at least 5 PFLOPS for this module. The phase two module should be delivered before the end of the first quarter of 2022.

Preparing for the Future

It is important to note that the number of cores per node in the CPU-module of the new system will be significantly higher than for Beskow, Tegner and similar systems. This may require changes to existing applications to incorporate the necessary parallelism so that they run efficiently on the hardware of the first phase. To take full advantage of the high capacity of the accelerators in the second phase, more extensive changes to the software may be required.

Researchers who have access to the source code might be able to undertake this potentially complex task themselves. In other cases, for example with commercial software, we recommend contacting the developers or vendor of the application. While this may sound intimidating, we expect that similar changes will be required to prepare for future exascale systems that will use comparable accelerator technologies. To make the transition easier, PDC plans to offer workshops and assistance to help to convert research codes. For further details of the workshops, please join the PDC announcements mailing list, watch our Events page or follow our Facebook page. You are also welcome to ask questions about this or other HPC-related research matters at the PDC cafes.

Storage Allocations Introduced for Klemming

Mattias Claesson, PDC

The Swedish National Infrastructure for Computing (SNIC) has recently introduced a uniform system for storage allocations at all SNIC sites. This will make it easier to keep track of what the storage resources are used for, as required by the Swedish Research Council (VR). This applies, for example, to both storage at individual SNIC centres and to Swestore. The allocations are handled in the SNIC User and Project Repository (SUPR) in a manner very similar to allocations for computing time. For projects only needing small amounts of storage at PDC, there will be a default storage allocation that can be accepted directly within the application for computing time. Projects with larger storage needs will have to submit one or more separate storage applications in the corresponding round of calls. Coupled with this is a more unified way of handling project directories, including naming and access control.

For PDC, this means moving away from personal user directories in the Klemming file system to project directories – which are owned by the Principal Investigator (PI) of each project – for the majority of the data. One major difference with this approach is that the PI of the storage allocation will have full access to all the data in the project directory. That means PIs will be able to change or delete any data stored there, if they need to. Individual users will still be allowed to store a very limited amount of data in their personal directories in Klemming, as a complement to the storage in their normal home directories in AFS. More information about these changes can be found at https://www.snic.se/allocations/storage and will soon be available on the PDC web site.

Introduction to PDC Systems Course

Thor Wikfeldt, PDC

On the 5th of February the biannual Introduction to PDC Systems course took place in room 304 at Teknikringen 14 – the home of PDC. This time the course was attended by around 15 people who were eager to learn the basics of using high-performance computing (HPC) resources for their research. The material for these courses has been roughly the same for a couple of years and can be found at https://pdc-support.github.io/hpc-intro. It is composed of brief presentations of theory, plus type-along sessions and exercises where participants learn how to use the file systems at PDC, how to compile parallel programs and how to run jobs through the SLURM scheduler. If you want to attend a future course, make sure to subscribe to the PDC-announce mailing list at https://www.pdc.kth.se/contact/joining-pdc-mailing-lists-1.736925, watch our Events page or follow PDC's Facebook page!

Below: Introduction to PDC Systems Course, 5 February 2020

NordicHPC

Thor Wikfeldt, PDC

NordicHPC (https://nordichpc.github.io) is a loose collaboration of Nordic computing facility staff and friends with a special focus on usability and reimagining high-performance computing (HPC) for the modern age. Most members fall somewhere between being application experts or system administrators of clusters – they are largely the people who are the front line supporting research and who deal with all of the daily mess.

The goals of NordicHPC are to

- share problems and suitable solutions across HPC and scientific computing centres and across countries, and collaboratively maintain these solutions,
- improve the usability of HPC and large research computing systems for the modern world, and
- meet once or twice per year in an informal setting to share recipes and issues with the hope of improving the user experience and also the "sysadmin experience".

NordicHPC has a GitHub organization at https://github.com/NordicHPC that is meant for sharing and collaborating on clever solutions or hacks to everyday computing and administration tasks. Are you sitting on a nice tool that needs a home? Consider sharing it via NordicHPC! It can give your work more visibility, and shared maintenance can produce better software for all of us. See the repository information at https://nordichpc.github.io/repo for more details.

COVID-19-Related Research at SeRC

Olivia Eriksson, SeRC

Researchers from the Swedish e-Science Research Centre (SeRC) are involved in a number of research projects directed towards the new coronavirus, SARS-CoV-2, and the disease that results from infection by that virus, COVID-19. Many of these projects use tools that have been developed earlier within the centre and build on knowledge from previous projects.

Within the Data Science MCP (Multi-disciplinary Collaboration Program), large-scale modelling of the human-SARS-CoV interactome is being performed to find out how the SARS-CoV-2 proteins interact with the different human proteins. This project involves data science, machine learning and physical simulations. Another project within the Data Science MCP performs studies of structural modelling and docking of proteins and protein complexes related to this new coronavirus.

The SeRC Exascale Simulation Software Initiative (SeSSI) MCP is involved in molecular dynamics simulations of virus proteins within the FOLDING@HOME project and they contribute to pharmaceutical research by the new EU consortium (EXSCALATE4COV). SeSSI is also involved in fluid dynamics simulations of complex fluids related to COVID-19, such as simulation of mucus droplet formation in alveoli and assessing the sensitivity of simulations to problem parameters (including geometry, resolution, fidelity, setup, and mask age).

The eScience for Cancer Prevention and Control (eCPC) MCP is comparing different national and international COVID-19 epidemiological models in collaboration with epidemiologists, biostatisticians and modellers from other universities. Some SeRC groups are also applying Bayesian methodology to epidemiologic models using, for example, Swedish data in order to make informed forecasts including uncertainty quantification.

EuroHPC & Sweden

Lilit Axner, PDC

The European High-Performance Computing Joint Undertaking (https://eurohpc-ju.europa.eu) or EuroHPC JU, is a one billion EUR joint initiative between the EU and various European countries to develop a world-class supercomputing ecosystem in Europe. The members of the joint undertaking are the following:

- the European Union, which is represented by the European Commission;
- Austria, Belgium, Bulgaria, Croatia, Cyprus, the Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Iceland, Ireland, Italy, Latvia, Lithuania, Luxembourg, Montenegro, the Netherlands, North Macedonia, Norway, Poland, Portugal, Romania, Slovakia, Slovenia, Spain, Sweden, Switzerland and Turkey; and
- the European Technology Platform for High-Performance Computing (ETP4HPC) Association plus the Big Data Value Association (BDVA).

The EuroHPC JU enables the EU and participating countries to coordinate their efforts and share resources with the objective of deploying a world-class supercomputing infrastructure and a competitive innovation ecosystem in super-computing technologies, applications and skills in Europe.

The Joint Undertaking intends to pool EU and national resources in high-performance computing with the aims of:

- acquiring and providing a world-class petascale and pre-exascale supercomputing and data infrastructure for Europe's academic, industrial and public users, matching their demanding application requirements by 2021 (This would be widely available to users from the public and private sectors, to be used primarily for research purposes.), and
- supporting an ambitious research and innovation agenda to develop and maintain a world-class high-performance computing ecosystem (exascale and beyond) in the EU covering all scientific and industrial value chain segments including low-power processor and middleware technologies, algorithms and code design, applications and systems, services and engineering, interconnections, know-how and skills for the next-generation supercomputing era.

To accomplish the first goal, the EuroHPC JU has been permitted to buy and install three pre-exascale machines by 2021. The systems will be located in Finland, Spain and Italy but, even though the machines will be situated in those particular countries, they will each be hosted by a consortium of several countries.

The system that will be based in Finland is called LUMI and is being hosted by a consortium of the following countries: Finland, Belgium, the Czech Republic, Denmark, Estonia, Norway, Poland, Sweden, and Switzerland. This means that a part of LUMI will be available to Swedish researchers, both in academia and industry, in a manner similar to that of Swedish national HPC resources. (For more details, see this article earlier in the newsletter.)

In addition to the three pre-exascale systems, the EuroHPC JU will be buying five petascale machines that will be sited in the Czech Republic, Bulgaria, Luxemburg, Portugal and Slovenia.

These machines will be interconnected with the existing national supercomputers and will be made available throughout Europe, to public and private users for developing leading scientific and industrial applications.

To reach its second goal, the EuroHPC JU has launched calls for proposals to fund research and innovation activities that will help Europe to remain globally competitive in the field of supercomputing. One of the aims of these calls is to support existing HPC competence centres or the creation of up to one national HPC competence centre in each of the EU member states and EuroHPC JU participating states.

Many countries jointly applied for this call. Out of these, applications from 31 countries have been accepted and one of these successful applications is from Sweden – its goal is to establish a national HPC competence centre in Sweden.

The national HPC competence centres should provide HPC services to industry (including to SMEs), academia and public administration organisations, as well as delivering tailored/modular solutions for a wide variety of users, in such a way as to ease and foster the transition towards the wider uptake of HPC in Europe. These centres should be the focal points coordinating all national initiatives, facilitating access for national stakeholders to European HPC competence and opportunities in different industrial sectors and domains. These national HPC competence centres are expected to open in September 2020.

PDC Cafes

Thor Wikfeldt, PDC

This semester we have been continuing with roughly monthly PDC Cafe events where we invite our users to join us for an informal meeting during which questions can be answered and problems solved. After the February Cafe, we moved these events online since everyone is working from home. Using the breakout room feature in Zoom we think that online Cafes can become just as useful as the corresponding in-person events – users and application experts first meet in a main room where a general round of introductions takes place and users get to describe their question or problem, then users are paired up with application experts to discuss particular problems in breakout rooms.

Starting from the 19th of May, the PDC Cafe will be merged with a corresponding outreach activity organized by our friends at the National Supercomputer Centre (NSC)! Joining forces with NSC for such online events will both strengthen ties between Sweden's two largest high-performance computing (HPC) centres, and make it easier for users who often use multiple clusters for their research to get help from the right HPC expert.

Above: PDC Cafe, 22 February 2020

PDC & PRACE Helping COVID-19 Research

Lilit Axner, PDC

The whole world is fighting against COVID-19 and so is PRACE, the European high-performance computing (HPC) infrastructure, together with its 26 members, representing European Union member states and associated countries. Fighting COVID-19 involves extensive research in areas like bioinformatics, chemistry and molecular modelling to find a cure as well as finding ways to eradicate the virus. However, to be successful, these research areas need steady ongoing support in terms of computing resources, such as HPC systems. We also know that COVID-19 is having an immense impact on the world economy and thus drastic measures will be required to boost the economy globally. Economical models and predictions that use HPC in their calculations may help the world economy to recover.

PRACE is welcoming project proposals requesting computing resources to contribute to mitigating the impact of the COVID-19 pandemic. This Call for Proposals follows a fast-track review process to provide swift feedback to the applicants. PRACE has established a scientific committee to discuss the merits of the proposals; it is supported by external scientific experts, and a technical team.

Already three projects – two from Spain and one from Germany – are taking advantage of this opportunity to use their calculation models to study, for example, the transmission dynamics of the SARS-CoV-2 virus or the molecular and dynamical SARS-CoV-2/hACE2 recognition and inhibition mechanisms. (ACE2 is an enzyme that is attached to the outer surface of certain cells, such as in the lungs, and which often enables corona-viruses to enter those cells. The human version of the enzyme is known as hACE2.) These projects have access to the seven largest HPC systems in Europe.

However, the support from PRACE does not end with that. Most of the PRACE partners have expressed their interest in going even further and opening up their Tier-1 systems for COVID-19 research as well. Access possibilities through fast track DECI calls, where PDC has also committed up to 5 million core hours to COVID-19 international research, are likely to be open for applications soon. This commitment from PDC via PRACE is in addition to PDC having already announced that Beskow and Tegner are available for Swedish COVID-19 research.

ARM HPC Hands-On Workshop

Xavier Aguilar, PDC

On the 18th of February, PDC held the ARM HPC Hands-on Workshop. This workshop hosted by PDC was part of a series of seminars and workshops offered by system vendors in the procurement to replace Beskow and Tegner. In such workshops, users are presented with exciting new hardware platforms that may become the new supercomputer hosted at PDC.

In this case, the workshop introduced the ARM HPC hardware and software ecosystems, and provided training for current generation hardware and instruction sets (Marvell ThunderX2 and NEON), as well as future generation ones, such as A64FX and Scalable Vector Extension (SVE).

Through a set of pre-prepared tutorial examples, participants were guided through an introduction to the compiler and mathematics libraries, and on to the compilation and emulation of SVE vector instructions. Additionally, participants received help to port and compile their own code examples on real ARM hardware, such as Marvell ThunderX2 processors.

Above: ARM HPC Hands-On Workshop, 18 February 2020

Code-Refinery Update

Thor Wikfeldt, PDC

Because of the ongoing COVID-19 pandemic, all the planned in-person CodeRefinery workshops have been cancelled this spring. Fortunately, however, this has not meant that the project has needed to go into hibernation. On the contrary, the CodeRefinery team has been using the time to convert lessons into an online format and to develop an online training programme – in fact, two online workshops have already been held during the spring. Although neither teaching nor attending an online workshop offers quite the same experience as attending in-person events, the team of instructors was pleasantly surprised about how well the workshop went and feedback from attendees was very positive. Online workshops will also provide a means to scale up and reach a larger audience than is possible with only in-person events.

Would you like to attend a future online CodeRefinery workshop? If so, fill in the "notify me" form at https://coderefinery.org/workshops/upcoming/#notify-me to be notified of upcoming events. Have you already participated in one or more workshops, and now want to join the CodeRefinery project in spreading better software development practices to the world, starting perhaps with your colleagues? Then register instead as a helper for upcoming workshops in the "notify me" form at https://coderefinery.org/workshops/upcoming/#notify-me. Bear in mind that there is no better way to consolidate your knowledge than teaching it to others!

Further details about CodeRefinery going online can be found in the blog post at https://coderefinery.org/blog/2020/04/24/online-workshops-update which came out in the latest CodeRefinery newsletter. To sign up for the newsletter, visit https://coderefinery.org/#stay-informed-and-connected.

Programming Future HPC Machines

Gilbert Netzer, PDC

To deliver high-performance end-user applications, software developers have to become increasingly aware of the micro-architecture innovations used in modern high-performance computing (HPC) hardware – this also applies to the replacement for Beskow, PDC's current main system. So, in addition to a workshop about the ARM HPC ecosystem, PDC hosted two webinars on programming NEC's and Intel's upcoming HPC solutions.

On the 14th of February, NEC presented their latest product in a long line of vector computers, the SX Aurora TSUBASA. The accompanying compiler, which is capable of understanding both FORTRAN and C++, attempts to extract suitable vector operations from the nested loop structures often found in scientific applications. As demonstrated in examples, programmers can get feedback in the form of detailed vectorization reports and run-time profiling information that helps to restructure the source code into recognizable idioms while avoiding stumbling blocks like unsuitable function calls.

A webinar on the 24th of March covered how to program Intel's entry into the discrete graphics processing unit (GPU) market, which is known as the Xe architecture. A cornerstone of Intel's strategy is a new software stack called oneAPI that offers a new programming model, data parallel C++, which is an extended version of ISO C++ based on the Khronos SYCL and community efforts.

In addition to working with GPUs, oneAPI is able to target Intel's CPUs, Field Programmable Gate Arrays (FPGAs) and other accelerator technologies. The package also includes further programming languages like OpenCL, as well as support for OpenMP target offload directives in both C++ and FORTRAN. Attendees could try out a beta version of this new software product in hands-on exercises using Intel's DevCloud.

In Memoriam

Johnny Eriksson

25 August 1957 - 9 January 2020

Johnny came to PDC in 2010 to work with the infrastructure extensions that were needed to install the new Cray system at the time, Lindgren. As a jack-of-all-trades (computer-related and real-life hands-on things), his expertise has been invaluable in keeping the PDC infrastructure in tip-top condition since then, and also in the installation of (and later upgrades to) PDC's most recent system, Beskow. Johnny's kindness and willingness to help, along with his sense of humour, are fondly remembered by his colleagues.

PDC-Related Events

PDC Summer School 2020 - postponed till 2021!

last two weeks of August 2021, KTH, Stockholm

Due to the COVID-19 pandemic, this year's PDC Summer School has been postponed till the summer of 2021. Details of the exact dates for the summer school will be available here: https://www.pdc.kth.se/about/events early next year.

HPC Sources

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

BioExcel
https://bioexcel.eu/news-and-events/events

CERN
https://home.cern/scientists/events/computing

EGI
https://www.egi.eu/category/events

HPC University
http://www.hpcuniversity.org/events/current

HPCwire
http://www.hpcwire.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.se/news-events
http://docs.snic.se/wiki/Training

XSEDE
https://www.xsede.org

SeRC Meeting Video

The **SeRC Annual Meeting** was held online this year due to the COVID-19 restrictions. For those who were unable to attend, a **recording of the main parts is available** at: **https://e-science.se/2020/05/serc-annual-emeeting-video**.