



no.1 2022

PDC Center for High Performance Computing

PDC Newsletter

Dardel Second Phase on the Way - Gert Svensson, PDC, page 4

eChem: Computational Chemistry from Laptop to HPC - Thomas Fransson, KTH, et al., page 6

Dardel Cloud & Dardel Inauguration - Dirk Pleiter, PDC, page 8 & page 11

EuroCC Lithuania Visits Stockholm - Lilit Axner, ENCCS, page 9

LUMI Inauguration and Applying to Use LUMI - Peter Larsson & Dirk Pleiter, PDC, page 10

CodeRefinery Update - Johan Hellsvik, PDC, page 10

SNIC Transition News - Lars Nordström, SNIC, page 12

VeloxChem Workshop - Roberto Di Remigio Eikás, ENCCS, page 12

Neko: Portable Framework for High-Fidelity CFD - Niclas Jansson, PDC, et al., page 13

Update on BioExcel and PerMed CoEs - Alessandra Villa & Rossen Apostolov, PDC, page 16


PRACE: Change of Pace - Michaela Barth, PDC, page 18

Boosting AI/ML Research on Dardel - Xavier Aguilar, PDC, page 22

GROMACS for Heterogeneous Exascale HPC Systems - Szilárd Páll, PDC, et al., page 22

Introduction to PDC Systems Workshop - Xin Li, PDC, page 28





Dirk Pleiter
Director, PDC

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: Dirk Pleiter

Editorial staff: Genet Edmondson

Graphic Design: Maria Malmqvist

E-mail: support@pdc.kth.se

ISSN 1401-9671

Cover

The second phase of the Dardel system at PDC will be installed soon. The extra cabinets will be put in a new row and the door panels will be rearranged somewhat so the updated arrangement will feature two more paintings by Nils Dardel: "Waterfall" and "Crime of Passion".

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

The upcoming summer holiday period is an opportunity to reflect on a busy period where there have been a lot of challenges and also many good results, various examples of which are covered in this newsletter.

The most prominent challenge for PDC is the continuing work on getting Dardel in place in its final configuration. The year started with the inauguration of the first phase (*Dardel Inauguration*). The event was attended by over a hundred people, which shows the strong interest in the system. During the initial phase, a partition comprising CPU-only nodes and a large-scale parallel file system were deployed. These rapidly became a major workhorse for many research projects, including large-scale projects using many nodes as well as small-scale projects using smaller numbers of cores. Some good news is that this compute partition and the capacity of the parallel file system will be significantly extended towards the end of the year. The new hardware, which is funded by SNIC, was ordered in May.

Before the new extension is installed, the second phase of Dardel will be put in place (*Dardel Second Phase on the Way*). It is based on innovative new hardware, namely AMD's MI250X graphics processing units (GPUs). Their impressive performance has only recently been disclosed publicly: each GPU features a performance of almost 50 TFLOPS, and there will be four such GPUs per node. To put this into perspective: it would only take two of Dardel's 56 new GPU nodes to outperform the whole Lindgren system, which PDC operated ten years ago. This new technology has made it possible to realise the world's first exaflop system, Frontier, at the Oak Ridge National Lab, USA. This system, together with its smaller sibling LUMI, made it to the top of the recent Top500 and Green500 lists. The Top500 list orders the systems with the highest performance in terms of throughput for floating-point operations during the execution of the LINPACK benchmark programme, while the Green500 list orders the systems with the best power efficiency.

This power efficiency only translates into energy efficiency for real-life applications if the hardware is exploited efficiently by the software. Therefore, work on applications and getting them ready for the GPU-phase of Dardel and for LUMI is of utmost importance. In this context, PDC can report on two highlights: a new computational fluid dynamics code (*Neko*) and porting the molecular dynamics code GROMACS (*GROMACS for Heterogeneous Exascale HPC Systems*), both of which demonstrate the usability of Dardel's GPU partition for high-performance computing (HPC) simulations. However, this new partition will also be very interesting for machine learning workflows. We are preparing for the necessary software to be available as soon as the hardware becomes available (*AI Research on Dardel*).

HPC applications are not only supposed to make the best possible use of the underlying hardware resources but to be user-friendly and versatile. A best-practice example is the VeloxChem software (*VeloxChem Workshop*). While this software is being prepared for Dardel's powerful GPUs, it is also having its features improved on a smaller scale as it becomes a central element of a new electronic book called eChem (*eChem: Computational Chemistry from Laptop to HPC*). This makes the application particularly suitable for learning and exploring quantum chemistry problems interactively. This is just one example of PDC's efforts in supporting training and education in the context of HPC. We are also very happy to be part of the third phase of CodeRefinery, a very successful training program that is partially funded by the Nordic e-Infrastructure Collaboration, NeIC, (*CodeRefinery Update*) and, of course, we are continuing our regular introductory training courses to help researchers to get started on Dardel (*Introduction to PDC Systems Workshop*).

The Dardel HPC system at PDC will soon be augmented by the Dardel Cloud, which will provide a virtualised computing infrastructure as well as a cloud storage system (*Dardel Cloud*). With this resource, PDC will expand its service portfolio to serve more researchers and make data sharing easier.

While PDC's prime focus is on enabling computational research in Sweden, we continue doing this in a European context that is evolving rapidly. The PRACE implementation projects have, in the past, offered many opportunities to work with European partners and to contribute to the European HPC ecosystem. While PRACE as an organisation continues to exist, the implementation projects are coming to an end, and thus the role of PRACE is changing (*PRACE: Change of Pace*). At the same time, the EuroHPC Joint Undertaking activities are increasing momentum. This includes the procurement of pre-exascale systems like LUMI, where we encourage Swedish researchers to apply for computing resources (*LUMI Inauguration*), as well as developing the European network of competence centres (*EuroCC Lithuania visits Stockholm*). Also, many of PDC's efforts towards the support of application domains remain European (*BioExcel & PerMed Updates*).

It is not just the European HPC landscape that is changing, significant changes are on the way in the Swedish landscape too. After twenty years, SNIC will be replaced by a new organisation (*SNIC Transition News*). Providing services for Swedish research based on high-end HPC resources will remain a national effort, but the way this is organised is changing. Decisions on how this will change are still in the making and will be reported on in the next edition of this newsletter.

The summer months will hopefully provide the necessary recreation opportunities to address this and other challenges!

Dirk Pleiter, Director PDC

In This Issue

Editorial	
<i>Dirk Pleiter</i>	2
Staff Focus	
<i>Ruth Hammond</i>	4
<i>Peter Larsson</i>	17
<i>Mustafa Arif</i>	23
Dardel Second Phase	
<i>Gert Svensson</i>	4
eChem: From Laptop to HPC	
<i>Thomas Fransson et al.</i>	6
Dardel Cloud	
<i>Dirk Pleiter</i>	8
EuroCC Lithuania Visit	
<i>Lilit Axner</i>	9
LUMI Inauguration	
<i>Peter Larsson & Dirk Pleiter</i>	10
CodeRefinery Update	
<i>Johan Hellsvik</i>	10
Dardel Inauguration	
<i>Dirk Pleiter</i>	11
SNIC Transition News	
<i>Lars Nordström</i>	12
VeloxChem Workshop	
<i>Roberto Di Remigio Eikås</i>	12
Neko	
<i>Niclas Jansson et al.</i>	13
BioExcel & PerMed Updates	
<i>A. Villa & R. Apostolov</i>	16
PRACE: Change of Pace	
<i>Michaela Barth</i>	18
AI/ML Research on Dardel	
<i>Xavier Aguilar</i>	22
GROMACS for Heterogeneous Exascale HPC Systems	
<i>S. Páll & A. Alekseenko</i>	22
Introduction to PDC Workshop	
<i>Xin Li</i>	28
PDC-Related Events	28
HPC Sources	28

Staff Focus



Ruth Hammond

Ruth Hammond joined PDC last semester as an exchange student under the Apple Scholar fellowship program. She is pursuing a B.Sc. in mechanical engineering with a computer science minor from Rose-Hulman Institute of Technology, USA. Ruth's previous research experience includes performance testing on the Summit and Frontier systems at Oak Ridge National Laboratory, algorithm processing at a U.S. Department of Defense laboratory, and applying machine learning to surgical robotics at Johnson & Johnson. After her time in Sweden, Ruth will return to the U.S. to complete an internship with IBM Research where she will apply AI and high-performance computing (HPC) to accelerate the discovery of carbon capture materials in the fight against climate change.

While studying HPC, deep learning, and machine learning at KTH, Ruth worked in first-line support at PDC. She thoroughly enjoyed her time at both and looks forward to applying to the master's programme at KTH next academic year.

In her spare time, Ruth holds leadership positions in non-profit organisations
 ...continued on page 17

Dardel Second Phase on the Way

Gert Svensson, PDC

Current Status of Dardel

The first phase of the new Swedish National Infrastructure for Computing (SNIC) flagship system at PDC (which is an HPE Cray EX system called Dardel) was inaugurated at the start of this year. All the researchers who had been using PDC's previous systems, Beskow and Tegner, have been transferred to Dardel (except for those from Scania) and have been successfully using the system for their scientific and technical simulations. Researchers from Scania will continue using Beskow until later this year when the Scania partition on Dardel is fully operational. The number of cores (or processing units) per CPU increased dramatically from 16 in Beskow to 64 in Dardel. The memory bandwidth was also increased, though not as much as the number of cores. Some researchers who are using Dardel have needed to adapt their code to the changes, and the resulting performance is usually good

Training for Using Dardel

Together with HPE and AMD, which is the manufacturer of the Dardel CPUs and graphics processing units (GPUs), PDC has arranged programming classes and introductory courses on how to use the system, as well as a GPU Hackathon week where experts from HPE and AMD helped to adapt locally developed code to run well on the AMD Instinct™ MI250X GPUs that will feature in the second phase of Dardel. There will be more courses and hackathons like these coming up, so keep an eye on the PDC Events page: <https://www.pdc.kth.se/about/events>.



Above: To fit the extra cabinets in the PDC computer hall, a new row will be created in front of the row of cabinets for the first phase of Dardel. The new row for the second phase will consist of two compute cabinets on the left end and a narrower cooling distribution cabinet (which controls the water for cooling the system) on the right. In terms of the images on the door panels, two more paintings by Nils Dardel will be added and some of the panels will be rearranged, as shown above. The paintings that will be added are known as "Waterfall" and "Crime of Passion" and the full view of "Waterfall" can be seen on the cover page.

Technology Readiness Level

HPE used a number of innovative techniques to make the Dardel system highly scalable and resilient to failures. This makes the system quite complex to operate from a system administration point of view, particularly as some of the new technology has not yet reached the technology readiness level needed to facilitate simple and smooth operation of the system. Some of the software tools that are expected for the day-to-day operation of the system are also not ready yet. In addition, although PDC has worked closely with the HPE team to schedule the updates needed to alleviate these issues, the maximum amount of downtime that was guaranteed by HPE has been exceeded due to unforeseen problems. PDC apologises for any inconvenience this has caused for the researchers using Dardel; work is underway with HPE to achieve more stable operation.

GPU Nodes for Second Phase of Dardel

In the previous PDC Newsletter, it was mentioned that the GPUs for the second phase of Dardel were planned to arrive during the spring this year. However, due to factors outside PDC's control, their arrival has been delayed till the summer. Adding the GPUs to Dardel could result in a significant amount of downtime, so PDC and HPE have spent considerable time investigating the best way to start using the GPUs without disrupting the ongoing operation of the first phase of the system. The first phase of Dardel uses a Slingshot 10 interconnect (rated at 100 Gb/s); it is the network that enables the compute nodes to exchange data at high speeds. However, the GPU cards for the second phase use Slingshot 11 (a 200 Gb/s version of the current interconnect). To resolve this in such a way that both phases would

use Slingshot 11 would involve changing all the network interface cards in the first phase and hence having the whole system down for a considerable amount of time. To avoid this, a separate system will be built for the GPU nodes; it will work with Slingshot 11 and be controlled using a cluster manager known as the HPE Performance Cluster Manager (HPCM). (In contrast, the first phase of Dardel uses Shasta technology.) This solution should provide the same level of functionality for those using the GPUs and will have the advantage of not affecting the ongoing stable operation of the existing Dardel compute nodes. The GPU nodes will access the current Lustre file system through a couple of Lustre router nodes. The plan is for the second phase of the system to be installed during the summer. PDC plans to join the CPU and GPU systems when all of the second phase has been installed and the whole system is operating reliably. That could happen towards the end of this year or early next year.

Upcoming Expansion of Dardel

The Swedish Science Council (VR) and SNIC have decided to invest more in Dardel than was initially planned (to mitigate the risk of a future lack of adequate capacity for Swedish research). This funding, along with the compensation from HPE for the late delivery of the first phase of Dardel, will be used to extend the system with an extra 468 CPU nodes and an increase of 50% in the disk capacity. At this stage, it is expected the additional hardware will be delivered in September or October this year. These extra nodes, plus the GPU phase, mean that Dardel will be increased with another two compute cabinets and one cooling distribution cabinet (for controlling the water that cools the system).

Type of nodes	Memory	Number of CPU nodes					Number of GPU nodes
		SNIC initial	Industry/business	SU Dept. of Astronomy	SNIC extra	Total	
Thin	256 GB	488	36		212	736	
Large	512 GB	20	236	12	248	516	56
Huge	1 TB	8				8	
Giant	2 TB	2			8	10	
Total		518	272	12	468	1270	56

Above: This table shows the number and types of nodes that will be in the Dardel system after both phases and all the currently-planned extensions are installed. The coloured boxes indicate the nodes that have already been installed.

eChem: Computational Chemistry from Laptop to HPC


Thomas Fransson, Mickael Delcey, Iulia Brumboiu, Xin Li, Manuel Hodecker, & Patrick Norman, TCB

Higher education in natural sciences has a long and honourable tradition with course curricula and textbook literature that are mature and well established. This is certainly the case at undergraduate and master's degree levels, and it is also often so at the level of doctoral studies. This is highly advantageous since it provides a common ground for communication between students and researchers from different parts of the world in their respective scientific sub-disciplines. We would argue, however, that to reach deep levels of understanding, it is beneficial to blend in components of exploration by students themselves and, being computational scientists, we advocate the use of an interactive framework such as Jupyter Notebook. With these modern, flexible software packages and frameworks, we are more capable than ever of providing platforms for interactive and easily adaptable computational

workflows, allowing a deep exploration of scientific concepts and theories. In this way, students can create their own algorithms and solutions for the physical processes at hand, reaching an understanding in line with Feynman's quip: *"What I cannot create, I do not understand."*

At the Department of Theoretical Chemistry and Biology (TCB) at the KTH Royal Institute of Technology, we are developing an electronic book suitable for such explorations titled "eChem: Computational Chemistry from Laptop to HPC" (<https://kthpanor.github.io/echem>). This e-book is published as a web page and collects a repository of Jupyter notebooks developed with the dual purpose of explaining and exploring the underlying theory behind computational chemistry in a highly interactive manner, as well as providing tutorials and workflows for solving specific chemistry-related questions. The theory parts are intertwined with explicit numerical examples, as well as suggested reading and exercises. With this learning environment, we aim to attract the interest of a wide audience, ranging from graduate students just getting into the field to experienced users and developers who want to see how we would address specific questions.

Below: Participants and instructors at the eChem workshop, 2-4 May 2022, KTH, Stockholm



Central to our e-book project is the VeloxChem software [1] (<https://veloxchem.org>). In 2018, the VeloxChem program started its development at TCB, and it has since positioned itself as the leading European quantum chemical software for contemporary and future supercomputer architectures. The project owes its success to the giant leap of starting from scratch and introducing a modern, sustainable, strictly object-oriented, code design that features a hybrid Python/C++ language, hybrid MPI/OpenMP parallelism, and auto-generated and highly optimised source-code kernels. In accordance with requests from users, data structures in the C++ layer are exposed as NumPy arrays by means of the header-only pybind11 library and these are thereby made available for exploration and integration in modules written for educational purposes. The eChem book and teaching material leverage this feature of the VeloxChem software design.

During May this year, we held a first workshop based on the eChem book. The workshop was carried out over three days, and we had 13 students with very different backgrounds and skills. While



Above: eChem workshop class, May 2022

some were well versed in software development, others had no prior experience in coding. The participants turned out to be very receptive to the strategy of combining lectures with explicit numerical examples, and in the future, we plan to give the workshop in the format of a five-day event to allow for it to be even more interactive. The topics covered were very varied, including tutorials designed to explore the intricacies of different electronic structure methods, as well as tutorials focused on the application of implemented routines to probe particular chemistry problems.

For the outcome of a workshop like this to be successful, it is essential that participants do not lose time and interest due to software installation issues. As trivial as this may sound, it becomes a challenge when everyone brings his or her own laptop to the event, which could be running any one of the three main operating systems that in turn have been customised according to personal preferences. On this occasion, to resolve that challenge, we used Conda; it is an open-source package and environment management system that works on any of the three main operating systems, and it can be used to install, run, and update packages. Before the event, we guided the participants to perform a Conda installation of a predefined and tested Conda environment. To make it feasible to study more realistic chemistry problems, however, it will be important for future workshops to offer a more powerful computer backend. In collaboration with experts at PDC, we will look into alternatives such as JupyterHub to let the Dardel cluster serve in this capacity and thereby truly live up to the title "Computational Chemistry from Laptop to HPC".

References

1. Z. Rinkevicius, X. Li, O. Vahtras, K. Ahmadzadeh, M. Brand, M. Ringholm, N. H. List, M. Scheurer, M. Scott, A. Dreuw, P. Norman, VeloxChem: A Python-driven Density-functional Theory Program for Spectroscopy Simulations in High-performance Computing Environments. *WIREs Comput. Mol. Sci.* 2020, 10 (5), e1457. <https://doi.org/10.1002/wcms.1457>.



Computational
Chemistry from Laptop
to HPC

Search this book...

Getting started

TUTORIALS AND WORKFLOWS

UV/vis spectroscopies

X-ray spectroscopies

Vibrational spectroscopies

ELECTRONIC STRUCTURE THEORY

General aspects

Wave function theory

Density functional theory

POTENTIAL ENERGY SURFACES

General aspects


Molecular structure optimization

Transition-state theory

Computational Chemistry from Laptop to HPC

A notebook exploration of quantum chemistry

For the things we have to learn before we can do them, we learn by doing them. -Aristotle



Next
Getting started >

Above: Cover page of the eChem book project (<https://kthpanor.github.io/echem>)

Dardel Cloud

Dirk Pleiter, PDC

In the past, the primary role of high-performance computing (HPC) centres has been to operate and provide access to one or more supercomputers. However, as the complexity and capabilities of HPC systems increase, and as our utilisation of them becomes more demanding, the role of HPC centres is changing to adapt to these ongoing developments. In the future, these centres will transform into service providers where the services they provide are based on different types of underlying computing and storage resources, rather than a single supercomputer. While the resources on offer will continue to include HPC systems, adding on-premise instances of private clouds will become the default. (For anyone unfamiliar with the terminology, a “private cloud” refers to a set of computational and data storage resources that are based on the same technologies as those used by big public cloud providers, such as Amazon Web Services (AWS). AWS provides on-demand computing processing capacity, data storage and software on a pay-as-you-go basis to companies, government organisations and individuals through a global network of “farms”, each containing many computers. In contrast, a private cloud uses a different business model, and the infrastructure is located on the premises of the organisation that runs the cloud.)

With these architectural changes on the way, HPC centres are responding to changing user needs, as well as those of researchers in emerging new science and engineering domains who need HPC resources for their research. One trend to highlight encompasses the efforts of HPC centres towards establishing domain-specific platform services. An example of this is the brain research community which is establishing a pan-European research infrastructure called EBRAINS. The services offered by this platform need to be deployed in a flexible environment, which is also stable. This is best realised through the use of virtual servers. Amongst other things, EBRAINS will allow end users to run HPC workflows through open portal services. Another trend is the emergence

of workflows that extend beyond a single HPC centre. Cloud instances allow the deployment of domain-specific services that effectively extend a data centre and make it possible, for example, to connect HPC model simulations or data processing to external data sources.

PDC is in the process of deploying such a private cloud instance, which will be known as the Dardel Cloud. It will feature more than a thousand physical CPU cores that can be offered to users as an even larger number of virtual CPUs. This not only adds a significant amount of compute capability, but also enables researchers and other services providers to deploy their services in a robust environment. These could range from Jupyter Notebook services for training to database services through to data analytics workflow services.

Another part of the Dardel Cloud system will be a large-scale storage system based on a technology called Ceph. It comes with interfaces that are widely used by the big public cloud providers and are designed for supporting the sharing of data, unlike the technology used for Dardel's parallel file system, Klemming. One reason to use the Ceph technology is that it supports authentication mechanisms which make it easy to integrate with providers of virtual identities, and therefore it is possible to avoid having to create local accounts for new users. For example, SUNET (the Swedish University Computer Network which provides high-speed internet access to academic institutions in Sweden) offers a service known as SWAMID where any Swedish researcher can obtain such a virtual identity. While the Klemming-based storage for Dardel is optimised for performance, the Ceph storage will be optimised for capacity that can be easily extended. PDC is thus laying the foundations for coping with the increasing demands for flexible data management as well as storage capacity.

Any researchers who are interested in obtaining early access to the Ceph system are warmly encouraged to contact PDC. Information about contacting PDC can be found at https://www.pdc.kth.se/support/documents/contact/contact_support.html.

ENCCS EuroCC National Competence Centre Sweden EuroCC Lithuania Visits Stockholm

Lilit Axner, ENCCS

As part of the collaboration between the EuroCC National Competence Centres in Lithuania and Sweden (which are known respectively as NCC Lithuania and ENCCS), the director of the Lithuanian centre, Mindaugas Mačernis, visited ENCCS in Stockholm between the 25th and 27th of January 2022 to become acquainted with the Swedish high-performance computing (HPC) ecosystem. This mentoring activity was jointly financed by CASTIEL (the Coordination and Support Action project for National Competence Centres on a European Level) and by EuroCC.

The ENCCS director, Lilit Axner, introduced Mindaugas Mačernis to the ENCCS team and together they had a full day of interesting discussions during which Mindaugas learned how ENCCS has successfully built up the centre in Sweden in just one and half years, including gathering a large number of researchers who are interested in taking advantage of the services provided by the centre to improve their research practices and outcomes. Lilit Axner discussed the activities and best practices that ENCCS had shared so far with its industrial and public sector users, as well as the importance of being agile and user-oriented when providing HPC and artificial intelligence (AI) services.



Above: (from left to right) Kjartan Thor Wikfeldt (ENCCS), Apostolos Vasileiadis (ENCCS), Mindaugas Mačernis (EuroCC Lithuania), Lilit Axner (ENCCS), Dirk Pleiter (PDC), and Ragnar Sundblad (PDC) at PDC

Below: (from left to right) Lilit Axner (ENCCS), Mindaugas Mačernis (EuroCC Lithuania) and Björn Alling (NSC) at NSC



The ENCCS training coordinator, Thor Wikfeldt, explained how ENCCS develops training material and how it tunes the training portfolio towards the needs of the Swedish users. They also talked about the periodic surveys ENCCS carries out with its users. Similarly, the ENCCS dissemination coordinator, Apostolos Vasileiadis, discussed the efficient outreach techniques that ENCCS uses to connect and communicate with the researchers who utilise its training and other services.

On the second day, Mindaugas visited the PDC Center for High Performance Computing at the KTH Royal Institute of Technology where the PDC director, Dirk Pleiter, together with PDC staff members, introduced Mindaugas to PDC's activities and described PDC's engagement with local researchers, as well as PDC's activities within several international projects. Mindaugas also visited the PDC computer room to see the new Swedish general-purpose system, Dardel, which is an HPE Cray EX supercomputer.

On the last day of the visit to Sweden, Mindaugas and Lilit visited the National Supercomputer Centre (NSC) at Linköping University. There the NSC director, Björn Alling, and NSC's technical director, Niclas Andersson, presented NSC's HPC services and the NSC staff, as well as discussing the details of how everyday work is organised at NSC. After the presentations, the group visited the NSC computer rooms to see the AI system, Berzelius, (which is funded by the Wallenberg AI, Autonomous Systems and Software Program) and the Swedish National Infrastructure for Computing (SNIC) system, Tetralith, both of which are hosted at NSC.



LUMI Inauguration & Applying to Use LUMI

Peter Larsson & Dirk Pleiter, PDC


The LUMI pan-European pre-exascale system was inaugurated on the 13th of June 2022 and is expected to reach full capacity with all hardware in place in the autumn. The supercomputer is located at the CSC – IT Center for Science in Kajaani, Finland, and is funded jointly by the EuroHPC Joint Undertaking (EuroHPC JU) and the LUMI consortium countries (which include Sweden).

LUMI is an HPE Cray EX supercomputer with several partitions: a large graphics processing unit (GPU) partition, a smaller partition that only contains CPUs, plus an auxiliary partition for data analytics with large memory nodes and some GPUs for data visualisation. Half of these LUMI resources belong to the LUMI consortium countries and are divided up based on the member countries' contributions to the funding. The resources in this pool are allocated through national organisations. The Swedish Research Council (VR) has contributed around 3.5% of the funding for LUMI so a corresponding share of the system is reserved for Swedish research. The other half of the LUMI resources constitute a EuroHPC quota which is available to researchers from all over Europe with the allocations being determined by the EuroHPC JU.

Researchers who are affiliated with an academic research organisation in Sweden may apply to use LUMI through the Swedish National Infrastructure for Computing (SNIC). Companies that are established or located in Sweden may also be able to buy compute time on LUMI by contacting SNIC directly. The SNIC allocation

rounds for academic research on LUMI are open for proposals for a limited time twice each year. Applications can be put in through the SNIC User and Project Repository for (SUPR) at <https://supr.snic.se>. The next SNIC round for LUMI will open in the autumn and close in December.

Swedish researchers have a second avenue for applying to use LUMI as researchers from academic institutes and companies established or located in an EU member state or in a country associated with Horizon 2020 can apply to use LUMI through the EuroHPC JU Call for Proposals for Regular Access Mode (see https://eurohpc-ju.europa.eu/eurohpc-ju-call-proposals-regular-access-mode_en). The next EuroHPC JU call closes on the 1st of July 2022, with the following call closing on the 7th of October.



CodeRefinery Update

Johan Hellsvik, PDC

This winter, the CodeRefinery project started up its third phase. The activity was kick-started with a large online workshop on 22-24 and 29-31 March. CodeRefinery implemented a setup whereby the teaching was provided in a format combining live streaming over Twitch (a video streaming service), plus interaction of learners and instructors via HackMD (an editor that allows people to share and collaboratively edit documents), as well as tutoring in Zoom breakout rooms. The combination of these channels enabled close to three hundred people to participate in the workshop and provided a great learning experience – both for the learners attending the workshop on their own and for the learners who were tutored in exercise groups.

Dardel Inauguration

Dirk Pleiter, PDC

On the 25th of January 2022, the first phase of the new Swedish National Infrastructure for Computing (SNIC) supercomputer system at PDC, Dardel, was inaugurated during a virtual event. More than a hundred people joined in from all over Sweden and also from Europe and North America, which demonstrated a strong interest in this new HPE Cray supercomputer.

In a range of presentations, the director of PDC, Dirk Pleiter, the director of SNIC, Lars Nordström, and HPE's vice-president for high-performance computing and artificial intelligence (HPC/AI) in Europe and the Middle East, Mark Armstrong, introduced Dardel as a new facility for computational research in Sweden. Despite all the difficulties that arose due to the pandemic while installing the system, we have succeeded in getting the new supercomputer, which is based on brand new technologies, into stable operation with all performance targets being reached. This major team effort involved the whole staff at PDC, as well as various people from the suppliers.

In October last year, selected researchers were invited to start using Dardel for their research to test the robustness of the new system. The inauguration was an opportunity to highlight some of that research and demonstrate that Dardel is already fulfilling its goal of enabling new research. Philipp Schlatter from the KTH Royal Institute of Technology (KTH) in Stockholm reported on new simulations of the airflow around rotating propellers and showed how well his group's applications scale on Dardel. Olle Eriksson, from



Above: Dirk Pleiter, the PDC director, welcomes everyone to the Dardel Inauguration, KTH, Stockholm, 25 January 2022

Below: Lars Nordström, the director of SNIC, at the Dardel Inauguration, KTH, Stockholm, 25 January 2022



the universities in Uppsala and Örebro, highlighted the use of Dardel for investigating the magnetic properties of materials at different scales. Finally, Patrick Norman from KTH presented initial results using VeloxChem, a new application for theoretical chemistry research. All three of these stimulating research talks nicely showcased the many opportunities that Dardel opens up, which should convince even more researchers to apply to use Dardel's excellent computing resources.

PDC would like to thank all the speakers and also the VIC Studio at KTH, which hosted the inauguration event. We now look forward to more exciting developments when the second phase of Dardel is inaugurated later this year.

For those who were unable to attend the event live, or who would like to refresh their memory of the details, copies of the presentations are available through the links on this page: <https://www.pdc.kth.se/about/pdc-news/dardel-inauguration-new-snic-supercomputer-attracts-strong-interest-1.1140806>.



Above: Mark Armstrong, HPE's vice-president for HPC/AI in Europe and the Middle East speaks over Zoom at the Dardel Inauguration, KTH, Stockholm, 25 January 2022

Lars Nordström, SNIC

The Swedish National Infrastructure for Computing (SNIC) has been responsible for providing large-scale compute and data storage resources, along with training and practical expertise, to support high-quality Swedish academic research since 2003. The situation is about to change as that responsibility will be transitioned to a new national infrastructure from the start of next year. It is likely that the new infrastructure will be known as the National Academic Infrastructure for Supercomputing (NAIS); that is the new organisation that has been proposed jointly by the major Swedish universities. The application is in the process of being evaluated by the Swedish Council of Research Infrastructures with a decision expected in late June.

Academic researchers who are using the SNIC infrastructure are naturally wondering if this upcoming change will impact their research. Putting it briefly, the intention is that, if the new organisation is approved, then SNIC and NAIS would work together to make the transition as seamless as possible from the perspective of the researchers using the resources, with NAIS taking over the existing SNIC resources and continuing to operate them throughout the remainder of their lifetimes. As the high-performance and high-throughput computing landscape is changing so rapidly, there are naturally many factors to take into consideration so that world-class resources will continue to be available for Swedish research.

To keep everyone informed about the transition process, SNIC will be publishing several newsletters over the course of this year. The first newsletter is available from the SNIC website (see https://snic.se/digitalAssets/605/c_605288-l_1-k_snic-newsletter-2022-1.pdf) and provides more details about the transition and the steps SNIC is taking to ensure there will be sufficient capability and capacity in the coming years.

VeloxChem Workshop

Roberto Di Remigio Eikås, ENCCS

VeloxChem is a clean-slate implementation of Hartree-Fock and Kohn-Sham density-functional theory, with a focus on the efficient calculation of molecular response properties. The code is designed to run on modern high-performance (HPC) architectures and, at the same time, to bridge the gap between ease of use and computational efficiency. More information about VeloxChem can be found in the cover article of the PDC Newsletter No. 2, 2019: <https://www.pdc.kth.se/publications/pdc-newsletter-articles/2019-no-2/veloxchem-1.945304>.

The development of VeloxChem is a success story arising from an ongoing collaboration between researchers at the KTH Royal Institute of Technology (KTH) and research software engineers at the EuroCC National Competence Centre Sweden (ENCCS) and PDC. (ENCCS is the Swedish node of the EuroCC project: the European network of 33 competence hubs in HPC for academia, industry, and the public sector.) Together, we have reduced the complexity of obtaining a working copy of the software, through the use of Conda packages, and expanded its use as a foundation for modern courses in computational chemistry.


In the past year and a half, ENCCS and PDC have delivered two training workshops about VeloxChem: “VeloxChem: Quantum chemistry from laptop to HPC” and then “VeloxChem: quantum chemistry towards pre-exascale and beyond”. The training material from both of these workshops is freely available online at <https://enccs.github.io/veloxchem-workshop> and <https://enccs.github.io/veloxchem-hpc> respectively.

The previous workshop (which was held on 6-7 May 2021) focused on the use of VeloxChem as a platform for interactive quantum chemistry computing, through the use of Jupyter notebooks. An interactive platform is appealing both for beginner and experienced quantum chemists. The former can freely explore the relation between the pen-and-paper derivation of established methods and their practical implementation. The latter can

be immediately productive when implementing new methods. In both cases, there is no need to pay a performance penalty for the high-level programming framework as all the numerical heavy lifting is done in C++ and parallelised with MPI and OpenMP. More information about the workshop can be found in the PDC Newsletter No. 1, 2021: <https://www.pdc.kth.se/publications/pdc-newsletter-articles/2021-no-1/veloxchem-quantum-chemistry-from-laptop-to-hpc-1.1094940>.

The latest workshop (which was held on 3-4 March 2022) was geared towards high-performance quantum chemistry and the efficient usage of VeloxChem on the new PDC flagship system, Dardel. We decided to set the stage for our hands-on sessions by providing an overview of some foundational principles underlying HPC, quantum chemistry and performance evaluation. The training material discusses at length the characteristics of the modern CPUs found in newer computing clusters, how these may affect performance, and their influence on programmers’ choices and algorithmic decisions. We included a discussion of performance theory, thus providing an empirical framework for evaluating the performance of VeloxChem on Dardel using real-world workloads.

Overall, both of these events were well received, and thanks to the European network of ENCCS, the events have seen attendance from an international audience. Participants appreciated the ease of use of VeloxChem: the ease of installation, its seamless integration with Jupyter notebooks, and, last but not least, its excellent performance on large HPC clusters.



The screenshot shows a webpage for the VeloxChem workshop. It features a title "VeloxChem: quantum chemistry towards pre-exascale and beyond" and a sub-header "THE LESSON". The main text describes the workshop's focus on quantum molecular modeling and its integration with advanced material design. It mentions that the workshop will highlight the efficient use of the VeloxChem program package on modern HPC architectures, such as the Dardel system at PDC and the pre-exascale supercomputer LUMI. A bullet point indicates that participants will learn how to perform quantum chemical simulations of ground- and excited-

Above: Anyone who would like to learn more about VeloxChem is welcome to work through the online training material from the workshops.



Neko: A Modern, Portable and Scalable Framework for High-Fidelity CFD

Niclas Jansson, PDC; Martin Karp, Jacob Wahlgren & Stefano Markidis, Division of Computational Science and Technology, KTH; Daniele Massaro & Philipp Schlatter, Department of Engineering Mechanics, KTH

With exascale computing capabilities on the horizon, the arena of computational research is transitioning to using systems with heterogeneous architectures. Traditional homogeneous scalar processing machines are being replaced with heterogeneous machines that combine scalar processors with various accelerators, such as graphics processing units (GPUs). Although these new systems offer high theoretical peak performance and high memory bandwidth, complex programming models and significant programming investments are necessary to make efficient use of the enhanced capability of such systems. Furthermore, most of the pre-exascale and exascale systems that are currently planned or that have already been installed (such as Frontier, the HPE Cray EX system at the US Department of Energy, and the LUMI system in Finland) contain a substantial number of accelerators. Thus, the challenges involved in porting scientific codes to these new platforms and tuning codes for these systems can no longer be ignored.


Computational Fluid Dynamics (CFD) is a natural driver for exascale computing as there is a virtually unbounded need for computational resources for accurate simulation of turbulent fluid flow, both for academic and engineering usage. However, established CFD codes build on years of verification and validation of their underlying numerical methods, potentially preventing a complete rewrite of a code base and rendering disruptive code changes a delicate task. Therefore,

porting established codes to accelerators poses several interdisciplinary challenges, from formulating suitable numerical methods and performing hardware-specific tuning to applying sound software engineering practices to cope with disruptive code changes.

To address these challenges and to make it possible to perform high-fidelity fluid simulations on accelerated systems like LUMI and the second phase of Dardel, we have developed Neko, a portable framework for high-order spectral element-based simulations, mainly focusing on incompressible flow simulations. The framework is implemented in the contemporary 2008 version of Fortran and adopts a modern object-oriented approach allowing for multi-tier abstractions of the solver stack and facilitating various hardware backends. Using Fortran as the language of choice instead of languages that have been more popular recently, such as C++ or Python, might at first seem like an odd choice, particularly for developing a new code. However, Neko has its roots in the spectral element code Nek5000 (from the Argonne National Laboratory at the University of Chicago), which was introduced in the mid-nineties and traces its origins to the Massachusetts Institute of Technology's even older NEKTON 2.0. Furthermore, research groups at the KTH Royal Institute of Technology (KTH) have extensively used the scalable Nek5000 code and also further developed its Fortran 77 codebase, thus creating a significant legacy of more than thirty years of verified and validated Fortran code, which, if rewritten in, for example, C++ would have to go through a very expensive and time-consuming revalidation and reverification process before it could be used in production. Therefore, by using modern Fortran, the already validated Fortran 77 kernels can be directly integrated into Neko, with only a minimal revalidation process.

Neko solves the incompressible Navier-Stokes equations, ensuring efficiency when running on systems with single-cores or accelerators via fast tensor product operator evaluations. For high-order methods, assembling either the local element matrix or the full stiffness matrix is prohibitively

Below: The graph below compares the strong scaling of Neko for several different computer backends: Piz Daint (a Cray XC50 equipped with NVIDIA P100 GPUs at CSCS), Beskow (a Crax XC40 with Haswell CPUs at PDC), Vulcan (a cluster using the SX-Aurora TSUBASA at HLRS), Dardel (an HPE Cray EX system with AMD EPYC CPUs at PDC), Alvis (a system at C3SE with A100 GPUs) and DEEP (an internal test system with V100 GPUs for the European project DEEP-SEA). The x-axis shows the number of processing elements (PEs) corresponding to one core or one GPU.



expensive. Therefore, a key to achieving good performance in spectral element methods is to consider a matrix-free formulation, where one always works with the unassembled matrix on a per-element basis. Gather–scatter operations ensure the continuity of functions on the element level, operating on both intra-node and inter-node element data. Currently, Neko uses MPI for inter-node parallelism and parallel I/O for production runs, but one-sided communication options (such as gather-scatter kernels based on Coarray Fortran) are under development.

When designing a flexible and maintainable framework for computational science, a major issue is finding the right level of abstraction. Too many levels might degrade performance, while too few results in a code base with many specialised kernels at a high maintenance cost. The weak form of the equation used in the Spectral Element Method allows Neko to recast equations in the form of the abstract problem to keep the abstractions at the top level and reduce the amount of platform-dependent kernels to a minimum. In Neko, this

is realised using abstract Fortran types, with deferred implementations of required procedures. For example, to allow for different formulations of a simulation's governing equations, Neko provides an abstract type, defining the abstract problem's matrix-vector product. The type comes with a deferred procedure "compute" that would return the action of multiplying the stiffness matrix of a given equation with a vector. In a typical object-oriented fashion, whenever a routine needs a matrix-vector product, it is always expressed as a call to "compute" on the abstract base type and never on the actual concrete implementation. Abstract types are all defined at the top level in the solver stack during initialisation and represent large, compute-intensive kernels, thus reducing overhead costs associated with the abstraction layer. Furthermore, this abstraction also accommodates the possibility of providing tuned matrix-vector products for specific hardware, by only providing a particular implementation of "compute" without having to modify the entire solver stack.


Despite the abstraction, modern Fortran does not provide any easy answers on how to interface with accelerators. Although vendor-specific solutions (such as CUDA Fortran) or directives-based approaches are a popular choice when porting Fortran codes to accelerators, a decision was taken not to rely on these due to portability issues and reduced performance. Instead, Neko uses a device abstraction layer to manage device memory, data transfer and kernel launches from Fortran. Behind this interface, Neko calls the native accelerator implementation written in, for example, CUDA, HIP or OpenCL. Furthermore, if

device-aware MPI is present, it is also exploited to minimise the necessary data movement between the host and device in communication kernels.

The ease of supporting various hardware backends through the solver stack and device layer abstraction is the key feature behind the performance portability of Neko, as illustrated in the figure on the previous page, where the time per time-step for solving the Taylor-Green vortex problem in various architectures is compared. Regardless of the architecture, Neko performs well on traditional general-purpose processors, SX-Aurora vector processors and various generations of accelerators.

Neko is currently used for large-scale direct numerical simulations (DNS) with applications in sustainable transport, studying the flow of air around Flettner rotors and the interaction of the flow with turbulent boundary layers. (A Flettner rotor is a rotating cylinder that will spin around its long axis as air passes across it. They were invented as rotor sails for ships about a hundred years ago and are now being studied as a clean and efficient method of propulsion to save fossil fuels.) A snapshot of these simulations is shown below. Thanks to PDC, Neko has already been ported and tuned for the new AMD MI250X accelerators in Dardel and LUMI; thus, the Neko developers are looking forward to exploiting these powerful systems once they become available to perform direct numerical simulations using previously untenable Reynolds numbers.

Neko is available as open source on GitHub (<https://github.com/ExtremeFLOW/neko>) and is distributed as part of the package manager Spack as "neko".



Above: The image above is a snapshot of our simulation of looking down at the top of a Flettner rotor in a turbulent boundary layer. The colours indicate the magnitude of the velocity of the air flow with violet indicating lower speeds, white showing medium speeds, and orange for high speeds.

Update on BioExcel and PerMed CoEs

Alessandra Villa and Rossen Apostolov, PDC

BioExcel Centre of Excellence

BioExcel@PDC is continuing its broad range of initiatives that support the computational biomolecular research community.

BioExcel Webinar Series

The BioExcel webinar series (<https://bioexcel.eu/category/webinar>) continues to feature notable developments in the field of computational biomolecular research. For example, in February, Paul Bauer (from the Department of Applied Physics at the KTH Royal Institute of Technology) introduced the features that are new in GROMACS for 2022 (<https://youtu.be/yIblJjez6Po>), and in April, Szilárd Páll (from PDC) spoke about GROMACS' performance, flexibility and portability and how advances in these areas can be applied to real-world simulations (<https://youtu.be/rTUz28f8p6g>). If you missed any of the webinars, you can catch up on the videos on BioExcel's YouTube channel: https://www.youtube.com/playlist?list=PLzLqYW5ci-2fanl2RtYHyFvjftu_M4Jwy. Remember to also keep an eye on <http://www.bioexcel.eu/webinars> for upcoming events!

Online Biomolecular Simulation Workshops

Online biomolecular simulation workshops were held in spring, both at the basic and advanced levels (see <https://bioexcel.eu/news-and-events/past-events>). The “physical” (rather than online) BioExcel Summer School will be run again in June this year after a break of two years. For details, see: <https://bioexcel.eu/events/bioexcel-summer-school-on-biomolecular-simulations-2022>.

Biomolecular Simulation Software Use Cases

BioExcel has recently published several use cases involving different types of biomolecular simulation software. These include “Electronic Interaction Phenomena: Proton Dynamics And Fluorescent Proteins”, which looks at using the GROMACS, CP2K, CPMD, MiMiC and pmx software and can be found at [https://www.hpccoe.eu/2022/03/21/electronic-interaction-](https://www.hpccoe.eu/2022/03/21/electronic-interaction-phenomena-proton-dynamics-and-fluorescent-proteins)

[phenomena-proton-dynamics-and-fluorescent-proteins](https://www.hpccoe.eu/2022/03/21/electronic-interaction-phenomena-proton-dynamics-and-fluorescent-proteins), and “Molecular dynamics meet docking: A use case on antibody design”, which features HADDOCK and GROMACS (see https://www.youtube.com/watch?v=_TDKfKX4kwM).

For everyone who is interested in learning more about how to use GROMACS, BioExcel is continuously updating the tutorials at <https://tutorials.gromacs.org>. The latest tutorial is on alchemical transformation using the AWH method (<https://tutorials.gromacs.org/awh-free-energy-of-solvation.html>). If you would like some tips on using social media for communicating about science, check out Michelle Mendoca's webinar: <https://bioexcel.eu/webinar-social-media-in-science-communication-2022-02-17>. Anyone who is active in the area of biomolecular simulations is welcome to join the BioExcel community and participate in the active BioExcel forums (see <https://ask.bioexcel.eu> and also <https://gromacs.bioexcel.eu>).


PerMed Centre of Excellence

PerMed@PDC has been contributing to developing a competency hub in computational personalised medicine (<https://competency.ebi.ac.uk/framework/permedcoe/2.0>). In general, this Competency Hub (see <https://competency.ebi.ac.uk>) is a website where you can browse competencies, career profiles and training resources in the bio-related computational areas to help advance your career. For example, there is a section for professionals in computational personalised medicine. PerMedCoE also recently developed two learning pathways (one based on the Linux command line and the other using Jupyter Notebook) that can help the biomedical community to use high-performance computing (HPC) infrastructures more easily.

Advanced GROMACS Workshop

BioExcel, in collaboration with the CSC – IT Center for Science, Finland, and the EuroCC National Competence Centre Sweden (ENCCS), held an online GROMACS workshop on 7-9 February this year. The workshop was supported by EuroCC and covered various topics that

improve efficiency when using GROMACS: enhanced sampling techniques by Berk Hess and Alessandra Villa (from KTH), QM/MM with CP2K by Dmitry Morozov (from the University of Jyväskylä, Finland), performance optimisation for CPU and GPU by Artem Zhmurov (from ENCCS), automation via BioBB-workflows by Adam Hospital (from the Institute for Research in BioMedicine, Barcelona) and GROMACS Python-API by Eric Irrgang (from the University of Virginia, USA). The workshop targeted advanced GROMACS users. To avoid problems with inexperienced attendees, we provided introduction material on the workshop topics (which is available online at <https://ssl.eventilla.com/advanced-gromacs-2022>). The workshop was run using an open on-demand interface that provided the attendees and trainers with easy access to the Jupyter Notebook environments and to the HPC resources. It was the first time that CSC used this setting and, apart from some small technical problems, overall the attendees found the hands-on approach easy to follow and enjoyed being active in the live documents (using the HackMD platform).



Above: BioExcel core software used in the first use case about antibodies

...continued from page 4

dedicated to making tech spaces more inclusive to women and minorities. She also likes to lift weights and experiment with new cooking recipes.

If you would like to connect with Ruth, her LinkedIn page is: <https://www.linkedin.com/in/ruthjhammond>.

Staff Focus



Peter Larsson

Peter Larsson joined PDC last year as an application expert with the task of being Sweden's contribution to the LUMI User Support Team, a fully remote team distributed across the LUMI consortium countries. In the team, he leads the high-performance computing (HPC) working group and works closely with the HPE/AMD Center of Excellence for LUMI.

Peter received his Ph.D. in physics from Uppsala University in 2009 and went on to work at the National Supercomputer Centre in Linköping as an application expert in computational materials science and also later as a business developer. This was followed by a stint in the HPC industry at Atos in Germany working with AI applications for edge computing. Apart from work, Peter enjoys running, hiking, and studying history in order to understand the future.

PRACE Change of Pace

Michaela Barth, PDC

PRACE is a non-profit organisation which has 25 member countries whose representative organisations have created, and now maintain, a pan-European supercomputing infrastructure that provides access to large-scale computing and data management resources and services for research. The PRACE consortium currently consists of 30 partners from 26 countries. Its budget is 30.1 million EUR with more than 85% of the funding being contributed by the European Commission (EC). To date, the member countries have made two major agreements to provide nationally-funded high-performance computing (HPC) resources to European researchers. The first agreement, PRACE 1.0, ran from 2010-15, and PRACE 2.0 went from 2015-22. In parallel with these agreements, there have been preparatory and implementation phase projects (funded by the EC) which complement the national investments in HPC infrastructure. The sixth of those implementation phases will end in December this year, so let us reflect proudly on the notable achievements of PRACE so far.

The non-profit PRACE organisation can claim its roots both in the HPC in Europe Task Force (HET) and the Distributed European Infrastructure for Supercomputing (DEISA). HET ran from 2004 to 2007 and was all about the big players in European HPC at the time (Finland, France, Germany, Italy, the Netherlands, Spain and the United Kingdom) trying to make Europe competitive within the HPC research arena, establish peer-review processes and obtain EC funding for top-level “Tier-0” national HPC systems. DEISA ran from 2001 to 2011 and, in contrast to HET, consisted of a series of projects which fostered collaboration between

national HPC centres running the next level of systems (“Tier-1”). The centres involved in the DEISA projects worked on a joint benchmark suite, had a dedicated network for a distributed shared file system and established a common operational environment which eventually resulted in the DEISA Extreme Computing Initiative (DECI), which was later incorporated into PRACE.

The best ideas from HET and DEISA were used as a basis to form the non-profit PRACE infrastructure organisation in 2010. There was a short preparatory phase (PP) project. Some of you may remember PDC’s system Povel, which was the PRACE-PP prototype for energy-efficient computing. Since then, there have been six project implementation phases: 1IP-6IP, during which PRACE has matured and become more structured, gradually establishing itself as the leading name for providing HPC research services in Europe. Its processes and services have evolved into a well-oiled machine, and PRACE has become the organisation uniting European HPC centres providing services for research.

Some of PRACE’s key achievements to date include the establishment of a peer-review process based on scientific excellence, and the increase in the competence of the researchers who participated in the extensive training programmes. Other achievements that may be less obvious at first glance, but that are definitely not less important, include the creation of the unified European application benchmark suite (UEABS), and all the applications that have been enabled (95 research projects were supported for more than three months), as well as the many best-practice guides that have been developed. Highlights were documented in hundreds of white papers. Overall, the PRACE operational services were able to support regular calls twice a year (typically for two billion core hours from the 14th call onwards), adding up to a total of 24 regular project access calls. Those calls were supplemented by extra calls on the DECI track and the calls supporting more than 70 small to medium-sized enterprises (SMEs) through the SME HPC Adoption Programme in Europe (SHAPE). PRACE can also be proud of

its activities related to providing information and guidance for decision makers at different levels, including technology monitoring, and giving input on procurements, and prototyping. Those efforts have made important contributions to improving the design and operation of large energy-efficient HPC centres in Europe. These contributions by PRACE have not gone unnoticed: for example, PRACE received the HPCwire Readers’ Choice Award “Best Use of HPC in Energy” for its strong interest in improving the energy efficiency of computing systems and reducing their environmental impact. Recently, PRACE’s incredibly quick reaction in deploying a COVID-19 fast track to provide access to resources for urgent computing research in Europe stood out positively. I have also witnessed PRACE bringing awareness of diversity, equity & inclusion (DEI) into public discussions about HPC human infrastructure at a time when other infrastructure players have not even considered the need to do so.

Looking more closely at the most recent project implementation phase, PRACE-6IP started in May 2019, and the initial plan was for it to run until December 2021; it was extended to June 2022, with a second extension for only a minority of the activities going to December 2022. During the pandemic, PRACE-6IP developed two new Massive Open Online Courses (MOOCs), the training portal was given a major revamp, six technical reports were produced in the area of market and technology watch (<https://prace-ri.eu/infrastructure-support/market-and-technology-watch>), operations groups worked towards a federated approach for the PRACE authentication and authorisation infrastructure (AAI), enabling of applications continued as usual along with benchmarking for 13 application codes (including GROMACS, NAMD, CP2K and Quantum Espresso) with an updated presentation format and on nine different systems. A new version (2.2) of the Unified European Application Benchmark Suite (UEABS) should be available in the PRACE repository very soon (<https://repository.prace-ri.eu/git/UEABS/ueabs>). For five systems, the energy to solution (that is, the total amount of

energy consumed in solving a computational problem) was measured and, for the systems that were in the Green500 list, PRACE’s findings were mostly in line with the Green500 results for those systems. All the benchmarking has led us to the conclusion that, so far, there is no single system that gives the best results on all the benchmarks. So, when applying for allocations, researchers are advised to choose the system that will perform best for their particular applications.

During the final extension of PRACE-6IP, which runs to the end of this year, the best-practice guide for the Nordic pre-exascale system, LUMI, will be completed. There have been delays getting access to LUMI (mostly due to pandemic-related global supply shortages), and PRACE wants to be sure to stick close to its initial paradigm for good best-practice guides, namely “We’ve tested it and it works”, so everything in the guide is based on real hands-on experience. The LUMI guide will be the 29th best-practice guide published by PRACE, and since parts of LUMI are GPU-accelerated, it will focus especially on how to port scientific codes to use GPU accelerators. There are also plans to cover how to write code that is adapted, from the start, for GPU systems. It is vital that researchers who intend to use exascale systems in the future start adapting their code to use GPUs now! While LUMI is a stepping stone towards a true exascale system, further effort is needed to adapt software so it is ready for exascale systems. For that reason, PRACE has been investigating forward-looking software solutions. The high-quality software that has been developed is intended to lay the groundwork beyond PRACE-6IP.

The final PRACE-6IP deliverables are due in December 2022, and the final review for PRACE-6IP is set for February 2023. Those deliverables, including the best-practice guide, are expected to be available on Zenodo soon at <https://zenodo.org/communities/prace>. Also, the UEABS data sets will be moved to Zenodo. The PRACE CodeVault, an open repository for various HPC code samples, will be migrated to GitHub and the PRACE organisation will keep hosting PRACE documents for a certain period after 6IP ends.

What do Tier-0 and Tier-1 really mean?

There is no exact technological definition of what constitutes a Tier-0 or Tier-1 resource; it is a relative scale. Tier-0 refers to the best HPC systems currently available within Europe. Tier-1 refers to resources that are a level below the Tier-0 systems. As HPC centres upgrade their resources, the level of technology associated with the tiers changes over time. Tier-0 and Tier-1 systems are sometimes referred to as “European” and “national” systems, respectively.

The PRACE Distributed European Computing Initiative (DECI) programme provides access to Tier-1 level resources across Europe via a series of competitive calls. This popular programme and the associated tools (such as the portals hosting the proposals, reviews or the accounting data) will come to a halt by the end of June this year, along with the DECI project management data base (DPMDB) and the directory information service containing user and machine account details (LDAP). After supporting 728 projects in 28 countries with a total of 1.6 billion core hours through 17 calls, DECI is coming to the end of its life in this current form; there are no plans for a direct follow-on. However, reviewers have been presented with a document explaining the benefits of DECI, with a view to carrying the most valuable aspects of DECI forward into future calls supported by EuroHPC. (The European High-Performance Computing Joint Undertaking or EuroHPC JU, which started in 2018, is a joint initiative by the EU and European countries to develop a world-class supercomputing ecosystem in Europe.)

In contrast to other European HPC calls, DECI was set up as a resource exchange programme using pre-existing national panels as reviewers. Spanning a wide range of subject areas, DECI has become well known to European computational scientists while naturally supporting collaboration and pooling of competence across Europe and beyond. Between 15% to 30% of the resources were typically made available to “external” projects which came from countries not contributing resources. For the last three calls, 25% of the projects fell into this category, allowing projects of outstanding scientific excellence to be awarded resources irrespective of differences in levels of national funding, and thus supporting excellence in European research. DECI has also played a role in making access to multiple sites easier for users, as well as building up trust between European HPC centres. The resources offered were of the national (Tier-1) type and with a variety of different architectures, leading to a diverse array of available systems, which made it possible to match researchers with the architectures most

appropriate for their simulations, while providing a large range of HPC expertise or filling the gap in between one national system and the next. It would be beneficial if something similar is offered under EuroHPC in the future.

As PRACE 6IP and DECI come to a close, PRACE has been looking into how it can best continue contributing to European HPC use for research. In the EuroHPC work plan, the access and allocation of computing resources, HPC training activities, and the envisaged user forum all play right into PRACE’s strengths; PRACE aims to be using its well-renowned peer-review process to allocate computing time to researchers on the forthcoming pre-exascale and petascale systems of EuroHPC and is considering how to bid in various upcoming calls for EuroHPC activities. PRACE has collected a list of interested partners for future collaborations to participate in those calls and also plans to provide high-level strategic support teams (HLSSTs) for EuroHPC systems; their tasks would include benchmarking, code optimisation and scaling-out applications.

In the future, PRACE aims to deploy and operate a platform for federating supercomputing (including HPC and quantum computing, as well as secure cloud-based services) and data resources for public and private users across the European Union. This platform will provide the framework to deploy basic federated infrastructure services on top of different technology implementations of resources, and enable the evolution, adaptation and combination of such basic services in order to facilitate the development of specific services according to users’ needs, including the research and scientific community, industry (including SMEs) and the public sector.

The draft work plan envisions that funding for facilitating access to the HPC ecosystem through national HPC competence centres will be continued and also proposes that centres of excellence and national HPC competence centres network with each other in the future under the umbrella of the CASTIEL 2 project. One of the first EuroHPC calls in this area involves new algorithms for applications on European exascale

supercomputers (see <https://eurohpc-ju.europa.eu>). In order to achieve a significant increase in the level of technology readiness for exascale computing, PRACE plans to answer this call with a project-based proposal where competence centres would propose short projects to develop novel exascale-ready algorithms and apply for funding for two full-time employees for a year to do the work. Then a PRACE consortium, consisting of the competence centres, would decide which of the projects would be funded.

PRACE envisions a future 3.0 phase where it would develop into a European association for computational science and – with a user-centric approach – change role from computing cycle facilitator to a user-focused service provider. By hosting a European user platform which joins HPC users and communities in Europe into an independent voice, PRACE wants to further develop its strengths in supporting HPC user communities to maximise the impact of their use of the HPC infrastructure. PRACE 3.0 plans to be a bottom-up organisation centred around the HPC user: “Give the HPC users a voice”. Provision of training would also be a high priority. The idea would be for user communities to be more actively involved, for example, by being given the power to openly participate in making decisions about PRACE’s activities and about directions for the future of HPC in Europe. The building blocks of such an undertaking would include inviting all the HPC users in Europe to be involved, and establishing a highly recognised cross-disciplinary scientific journal with a major focus on computational science, supported by an underlying high-quality peer-review process to ensure scientific excellence. Existing established rewards, like the PRACE Ada Lovelace reward and the new PRACE HPC Excellence Award, which recognise groundbreaking research relying on HPC, would continue as attractive and living elements of PRACE 3.0.

The organisational details for the next phase of PRACE are not yet set in stone, but the idea is to have different “chapters” of the infrastructure for major areas of computational research (such as biochemistry, bioinformatics, life sciences,

physiology and medicine, chemical sciences and material sciences) and for membership to be extended so individual people, and not just organisations, can be members of PRACE.

Even though not all the details of the EuroHPC calls are known at this point, PRACE sees itself in a future where it will definitely be responsible for different activities as part of the European HPC ecosystem, especially in user support. PRACE aims to be in a coordination role, acting as a one-stop shop for the European HPC users. In addition, keeping the peer-review processes in place and enabling access to world-class computing and data management resources in Europe are significant parts of the PRACE 3.0 pillars and principles. Fostering international collaboration, continuing to monitor technology, and a thriving HPC ecosystem are also on the map. PRACE 3.0 would also aim to provide a platform for national computing centres to exchange best practices and to discuss ways to improve HPC services. A continued series of infrastructure workshops would make sense too. It is expected that successful and popular exchange and training programmes, like the “Summer of HPC” (SoHPC) would also be revitalised in the future.

At the latest PRACE-6IP all-hands meeting, discussion focused on how to avoid competence loss and keep the whole European HPC infrastructure structured while including both existing and new institutions in PRACE, and doing so in such a way that it will be easy to prepare applications for the upcoming EuroHPC calls expected to be opened later this year. Many voices were in favour of having a reviewed PRACE journal where technical white papers and reports could be published with a DOI, as envisioned in the PRACE 3.0 concepts. In the meantime, you can find details of open calls, including the EuroHPC access calls, at <https://pracecalls.eu>. It is also a good idea to keep an eye on the HPC in Europe portal (<https://hpc-portal.eu>), which collects all the HPC services offered in Europe into a single database. In addition, you can check out the training being offered by the different national competence centres: <https://www.eurocc-access.eu/services/training>.

Boosting AI/ML Research on Dardel

Xavier Aguilar, PDC

The hardware for the second phase of Dardel will arrive later this summer. There will be a new partition comprising 56 graphics processing unit (GPU) nodes, each of which will be configured with one AMD EPYC processor with 64 cores and four AMD Instinct MI250X GPUs. These GPUs can perform both vector and matrix operations. While typical HPC applications use vector operations (and hence express performance in those terms), machine learning (ML) code benefits from matrix operations. The performance of these GPUs equates to 383 TFLOPS for half-precision floating-point format (FP16) and 95.7 TFLOPS for single-precision floating-point format (FP32) per card when using matrix operations. So each node will pack a good amount of computational power, which means the new GPU partition on Dardel will be a platform that is highly suitable for ML workflows.

PDC was provided with a few experimental nodes containing AMD Instinct MI100 cards, a predecessor of the MI250X. While the MI100s cannot be directly compared to the MI250Xs, they serve as a good testbed to try out the AMD software stack for their GPUs while waiting for the final ones to arrive. In this case, we are using the nodes as a testbed for various ML frameworks and workloads. We have installed Tensorflow and Pytorch, the two most used frameworks for ML and Deep Learning (DL), and are testing their functionality as well as their performance, even though the performance observed on the MI100s will not directly relate to the performance that will be provided by the MI250Xs. We are currently testing native installations of the software, however, Singularity has been already deployed on Dardel, and thus, using the GPU nodes for ML/DL will be even easier with the containerised solutions provided directly by AMD. Furthermore, libraries such as RCCL and frameworks such as Horovod will make it possible to use multiple GPUs and multiple nodes at the same time, thereby opening the door to developing and training larger AI models on Dardel.



Preparing GROMACS for Heterogeneous Exascale HPC Systems

Szilárd Páll, PDC, & Andrey Alekseenko, Scilifelab & Department of Applied Physics, KTH

GROMACS is an open-source molecular dynamics (MD) simulation package that is one of the most widely used high-performance computing (HPC) codes in Sweden and worldwide, both in academia and industry. The development is led by a Stockholm-based team from the KTH Royal Institute of Technology (KTH) and Stockholm University (SU), with major contributions recently being developed at PDC. These have included advances in application programming interfaces (APIs), algorithms and parallelisation.


GROMACS has a strong focus on performance thanks to state-of-the-art parallel algorithms and the use of bottom-up performance optimisation. Portability in practice is also a major objective which is achieved by using hardware abstraction layers for single instruction, multiple data (SIMD) units and for graphics processing units (GPUs), and the use of standards-based programming languages and APIs (whenever possible), as well as extensive testing on a wide range of architectures and platforms using automated continuous integration (CI), regular HPC deployments, and integration into multiple Linux distributions.

With PDC's new Dardel system being based on AMD's heterogeneous exascale architecture, which is also used in the LUMI system in Finland and in the Frontier system at the Oak Ridge National Laboratory (ORNL) in the USA (respectively the first and third systems on the recently announced 59th TOP500 list [1]), supporting these architectures in GROMACS has required long-term investments in both algorithms and parallelisation, as well as recent work on adopting new GPU frameworks. GROMACS has embraced GPUs for nearly a decade: the first official release

with native GPU support came out nearly a decade ago in early 2013. At its core, it was a bottom-up redesign of the simulation engine, starting with reformulating key MD algorithms for modern processor microarchitectures combined with a heterogeneous multi-level parallelisation scheme [2].

From the start, the GROMACS MD engine has been heterogeneous, making use of both CPUs and GPUs for performance and flexibility, plus it is multi-level as it directly targets each level of hardware parallelism from SIMD vector units through multicore non-uniform memory access (NUMA) accelerators to multi-node cluster topologies. Together these make it one of the fastest MD engines. Embracing heterogeneity in the early phases of parallelisation design enabled the GROMACS engine to evolve features without needing to port everything to GPUs, as well as making it possible to easily adapt to hardware changes (like faster GPUs or denser heterogeneous nodes), while also maintaining support for a broad range of features and at the same time demonstrating excellent performance on multiple generations of hardware [3,4,5]. This design remains at the foundation of the GROMACS parallelisation that targets various system architectures. However, in anticipation of a changing exascale HPC hardware landscape, we initiated research on algorithms and parallelisation techniques to efficiently target future architectures within the frame of the Swedish Exascale Computing Initiative (SeCI) Swedish Foundation for Strategic Research (SSF) project, which was led by Erwin Laure in collaboration with the GROMACS team.

As a result, further improvements were made to key algorithms; we designed a dual-pair list algorithm for pair interactions which extends the accuracy-based formulation of our original work to improve regularised data reuse. In addition, we started a redesign of the original force-offload heterogeneous setup and GPU-resident parallelisation (see below) was designed and developed in a close codesign collaboration with NVIDIA.



Above: GROMACS Heterogeneous GPU-resident with most computation offloaded. This setup is particularly well-suited for accelerator-dense systems, like Dardel's GPU partition, and is supported in the CUDA and SYCL backends; the latter only lacks the GPU kernels for the bonded force computation (purple) which is computed on the CPU where the other forces (magenta) are done.

Staff Focus



Mustafa Arif

Mustafa Arif recently joined PDC as a systems manager and will be facilitating the operation of the high-performance computing (HPC) systems. Before joining PDC, Mustafa worked at Texas A&M University, Qatar, where he was responsible for managing the HPC and AI infrastructure. His core expertise includes designing architectures for HPC and AI systems and managing such systems, as well as private/public cloud administration and scientific computing consultation. Mustafa holds a master's degree in computer engineering and is a Red Hat Certified Engineer. He has previously managed Cray, HPE and DELL HPC systems and been involved with data centre operations and monitoring, as well as offering training courses on various aspects of HPC, including Singularity containers. At PDC, Mustafa's focus is on assisting the team responsible for operating the HPC system and contributing to new projects; one of them involves setting up a private cloud for HPC users based on OpenStack.

Mustafa enjoys travelling and spends his free time with friends and family. He loves to innovate and has created various open-source embedded systems and 3-D printing projects.

The force-offload scheme relies on shifting the most compute-intensive computations to the GPU, combining it with a multi-level load balancer to optimise CPU–GPU execution overlap to maximise GPU utilisation. It relies on the CPU for integration, which requires frequent CPU–GPU data movement every iteration. With an increasing fraction of computational power being delivered by GPUs, and because CPU–GPU interconnect performance has been lagging behind the actual computational performance, the new GPU-resident scheme was designed with the aim of minimising reliance on CPU resources and stopping the intra-node interconnect from becoming a bottleneck. This algorithm prioritises keeping the GPU busy and avoiding CPU–GPU data movement. To do so, the integration was offloaded, allowing the molecular simulation state to remain resident on the GPU for tens to hundreds of iterations. Importantly, the heterogeneous design is maintained, and the CPU can still be used, but in a reduced role, primarily to execute tasks associated with auxiliary features not ported to the GPU. As an additional benefit, direct GPU communication (also called GPU-aware) could be implemented efficiently (see figure below). This allows data to be transferred directly between GPUs instead of staging through the CPU, which is key to making use of fast GPU interconnects

(like NVIDIA NVLink or AMD Infinity Fabrik links). Direct GPU communication also improves inter-node communication latency, especially on systems like the Dardel GPU partition where the network is directly connected to the accelerators and can free up CPU cycles otherwise used for coordinating data staging.

These long-term algorithmic efforts focusing on heterogeneous architectures have made GROMACS well-prepared for the exascale HPC landscape, but adapting a large codebase to new architectures is not without challenges. The new GPU platforms and the programming models for these proved to require significant effort. Depending on the programming language of the codebase and the desired level of abstraction, a range of choices is available to exploit GPU accelerators: accelerated libraries, preprocessor directives, standard language support, and direct programming languages/APIs. Only the latter is able to provide the level of performance and scheduling control required in GROMACS (and similar codes). Its C++17 codebase puts GROMACS in a favourable position because most direct accelerator programming languages and APIs are built on, or compatible with, C++.

The initial GPU support in GROMACS was built on NVIDIA CUDA, the most robust and best-performing choice at the time. NVIDIA GPUs

have been the dominant accelerator architecture for over a decade, and CUDA is established as a mature and widely used GPU programming framework. However, to ensure portability of the GROMACS codebase, an open standard-based OpenCL backend was added early on in 2015 with AMD and NVIDIA GPU support first, which was extended to Intel later. Although NVIDIA GPUs have continued to dominate, GROMACS has maintained its commitment to portability through its OpenCL backend. A GPU portability abstraction layer has also been gradually developed to reduce code duplication, to abstract details of GPU APIs, and to prepare the code for new backends in anticipation of a diverse accelerator future.

In recent years, both AMD and Intel joined the GPU accelerator market, and both vendors introduced new heterogeneous programming frameworks with OneAPI and ROCm, respectively. For direct programming of GPU accelerators, AMD developed HIP, a CUDA-like language, and ROCm also ships OpenCL support. The Intel OneAPI toolkit offers OpenCL and SYCL, which is based on DPC++. SYCL is a framework based on open industry standards (like OpenCL), although Intel is one of the major vendors behind it, and while SYCL is still evolving fast, it has several implementations (three major ones and another eight in development at the time of writing) and therefore SYCL code can already target a wide range of hardware including AMD, Intel and NVIDIA GPUs [6]. Although ROCm (and hence HIP) has an open-source code-base, it is only supported by a single vendor and software stack (although, with limitations, it can be used as an abstraction on top of CUDA for NVIDIA GPUs). Both HIP and SYCL are based on C++ and use a single-source model.


With heterogeneous HPC systems based on these new architectures on the horizon, GROMACS needed to extend its GPU backend to support the new Intel and AMD GPU architectures. While the existing OpenCL backend was a reasonable option on the new GPU platforms, the C-based OpenCL kernel language has long been cumbersome in the native C++ GROMACS codebase. In addition, first-class OpenCL support in AMD's ROCm

stack was unclear for some time (although more recently support has been confirmed). To maintain portability, achieve good performance, and have good integration in the C++ codebase, an obvious choice would be to adopt both SYCL and HIP.

With the Swedish e-Science Research Centre (SeRC) and Intel forming a oneAPI academic centre of excellence (CoE) in 2020, SYCL adoption started by extending the GROMACS GPU backend with DPC++ support for Intel GPUs, with an initial release in 2021. Based on the experience gained during that endeavour, we decided to aim at also adopting SYCL as the multi-vendor GPU portability backend to replace OpenCL. This was possible thanks to a fast-developing SYCL standard and promising multi-vendor support. We recognised early on that SYCL code developed with the DPC++ toolchain for Intel needs additional effort to ensure its portability. Early work used hipSYCL [7] as a second SYCL implementation to validate and strengthen the portability of the new backend.

The most appropriate way to support AMD GPUs remained unclear and, to decide how to achieve that, in early 2021 we carried out an in-depth assessment of the available solutions, taking into account short- and long-term factors including portability, performance, proprietary versus standards-based solutions, and available development resources. The team was hesitant to take on the major effort of adding an additional GPU backend based on HIP just for AMD GPUs, as that would involve a commitment to maintaining a total of four GPU backends. Although it is deprecated, OpenCL still requires maintenance, and CUDA is expected (at least in the near term) to be essential for compromise-free performance and robustness on the NVIDIA GPUs used by a large part of the GROMACS user-base.

Following a careful technical evaluation and risk analysis, in mid-2021 the GROMACS team decided to use SYCL in production on AMD GPUs through hipSYCL for the upcoming releases, with the primary focus being the Dardel GPU and LUMI HPC architectures [8]. The decision is both a vote for portability and open standards, as well as an effort to maintain a sustainable expansion of the



Above: GROMACS multi-node heterogeneous GPU-resident design with direct-GPU communication


codebase. Instead of pouring the team’s limited development resources into adding multiple new backends, as well as optimising these for the target platforms, we opted to devote more resources to code modernisation and improving the GPU abstraction layer, as well as the robustness of the asynchronous offload parallelisation.

An important feature of hipSYCL is its strong reliance on the platform-native toolchains, making it possible to mix native (HIP or CUDA) code with SYCL and importantly enabling the use of native development tools (like CUDA and HIP profilers and debuggers). Thanks to this, our early performance evaluation prototype, which was developed to inform the decision to adopt hipSYCL, showed good performance and relatively smooth development. The early hipSYCL prototype, which was released in mid-2021, was a joint effort between GROMACS developers at the EuroCC National Competence Centre Sweden (ENCCS), PDC, and the Intel oneAPI CoE. Since then, thanks to being able to devote a significant amount of resources to the SYCL backend, development has focused on the GPU portability layer, feature completeness of the SYCL backend, build system integration, and automated testing.

GROMACS 2022 was released with GPU-resident SYCL support that was close to feature-complete. As the primary hardware platforms were still not available and software stacks were still lacking maturity, performance was tracked, but it was only of secondary focus for this release. Tuning for the target platforms and optimising multi-node parallelisation was left to a later stage. These aspects were focused on during the Dardel HPE hackathon in February 2022, where we greatly improved the performance of


the SYCL compute kernels on AMD GPUs. It is now overall on par with the native HIP kernel performance. (AMD kindly provided access to their internal HIP GROMACS port as a baseline for the performance comparisons.) However, as shown in the figure below, application performance is still lagging behind. Performance analysis and profiling revealed that the source of the overheads (see graphs on the next page) is in the hipSYCL runtime’s scheduling heuristics. The plan is to improve those with the help of the hipSYCL team. Furthermore, a prototype SYCL implementation of the direct GPU communication layer was successfully evaluated using the Cray GPU-aware MPI during the hackathon on PDC’s HPE test platform. This work identified a set of improvements needed in the vendor software stack, as well as in the GROMACS SYCL backend.

The GROMACS project chose a somewhat unconventional path to supporting AMD GPUs, but this made it possible to focus on long-term goals. As our long-term experience with multi-platform OpenCL support shows, sharing as much code as possible across platforms can greatly ease porting and has long-term maintenance advantages. In our experience, it is the GPU API-specific code (like device detection and task scheduling) that has taken the most effort to develop and maintain, but this also lends itself well to being shared across platforms. In comparison, low-level code (like GPU kernels) generally needs vendor or even device-specific optimisation to achieve good performance, whether it is written in HIP, OpenCL, or SYCL. Moreover, relying on open standards-based languages and APIs whenever possible (not just for GPUs), like MPI, OpenMP, C++, OpenCL and SYCL has strong additional benefits. We recognise



Above: Performance analysis revealing hipSYCL runtime scheduling inefficiencies during the 2022 PDC HPE hackathon

Below: GROMACS SYCL kernel (left) and application (right) performance relative to native HIP on AMD MI100. The SYCL kernel performance is on par (or in some cases faster) than native HIP. Application performance (solid line in the graph on the right), especially for small inputs, is still slightly behind.



that hipSYCL adds an extra software layer on top of the native runtime, which can lead to performance trade-offs in the short-term. We believe that the long-term benefits of our approach outweigh such short-term drawbacks, both for developers and users of GROMACS. Furthermore, the option of a different SYCL implementation with AMD support [9] is available, and adding a native HIP backend also remains a long-term option if warranted.

GROMACS is well on track to being ready for the AMD heterogeneous architecture at the heart of the Dardel system at PDC, which will support the molecular simulation research community in being early users of Dardel (and other new systems with similar architectures). This is a brand-new platform with a new software environment, so testing and stabilisation work is still ahead of us.

The GROMACS team maintains a stable hardware-enablement branch based on the 2022 release [10], which includes the majority of the recent SYCL improvements, most of which are too disruptive for a maintenance patch release. This code branch has stronger stability guarantees than the main branch (where active development for the upcoming 2023 release is carried out), and our aim with this branch is to make it possible for researchers to make early use of the branch on AMD and Intel GPUs, and thereby enable early use on the PDC and LUMI systems even before the next major release which is planned for early 2023.

References

- <https://top500.org/lists/top500/2022/06>
- S. Páll & B. Hess. "A Flexible Algorithm for Calculating Pair Interactions on SIMD Architectures." *Computer Physics Communications* 184, no. 12 (December 2013): 2641–50. <https://doi.org/10.1016/j.cpc.2013.06.003>
- S. Páll, A. Zhmurov, P. Bauer, M. Abraham, M. Lundborg, A. Gray, B. Hess & E. Lindahl. "Heterogeneous Parallelization and Acceleration of Molecular Dynamics Simulations in GROMACS." *The Journal of Chemical Physics* 153, no. 13 (7 October 2020): 134110. <https://doi.org/10.1063/5.0018516>
- C. Kutzner, S. Páll, M. Fechner, A. Esztermann, B. L. de Groot & H. Grubmüller. "Best Bang for Your Buck: GPU Nodes for GROMACS Biomolecular Simulations." *Journal of Computational Chemistry* 36, no. 26 (5 October 2015): 1990–2008. <https://doi.org/10.1002/jcc.24030>
- C. Kutzner, S. Páll, M. Fechner, A. Esztermann, B. L. Groot & H. Grubmüller. "More Bang for Your Buck: Improved Use of GPU Nodes for GROMACS 2018." *Journal of Computational Chemistry* 40, no. 27 (15 October 2019): 2418–31. <https://doi.org/10.1002/jcc.26011>
- Khronos SYCL: <https://www.khronos.org/sycl>
- hipSYCL: <https://github.com/illuhad/hipSYCL>
- <https://enccs.se/news/2021/09/gromacs-adopts-hipsy-cl-for-amd-gpu-support>

9. <https://www.olcf.ornl.gov/2021/06/17/argonne-and-oak-ridge-national-laboratories-award-codeplay-software-to-further-strengthen-syclx2122-support-extending-the-open-standard-software-for-amd-gpus>
10. GROMACS 2022-HWE branch: <https://gitlab.com/gromacs/gromacs/-/tree/hwe-release-2022>

This work is licensed under a [CC-BY 4.0 License](https://creativecommons.org/licenses/by/4.0/).

Introduction to PDC Systems Workshop

Xin Li, PDC

The course “Introduction to PDC Systems” is held biannually and serves to introduce researchers to practical information about the high-performance computing (HPC) infrastructure at PDC. The topics that are covered in the course include an overview of the HPC infrastructure, getting an account for the PDC systems, logging in, running jobs, storing data, and compiling code, as well as tips and best practices for using the Bash shell, SLURM script, the Python virtual environment, MATLAB, and EasyBuild.

The first introductory course for 2022 was held online on the 23rd and 24th of March and was attended by 30 participants. The slides from the course are available at <https://pdc-support.github.io/pdc-intro>. Apart from going through the prepared material, course participants were also able to get assistance to log in to the new Dardel HPC system at PDC, practise common procedures and useful commands on the system, and work on several exercises that helped to deepen their understanding of how they can run their calculations on the system.

The next introductory course will be held in November 2022. If you are a new user of the HPC infrastructure at PDC, it is highly recommended that you attend the course in the autumn, and we will do our best to help you speed up the process of becoming comfortable with working in the HPC environment at PDC. For information about this course and other training events run by (or in collaboration with) PDC, you can sign up for the PDC general announcements mailing list <https://www.pdc.kth.se/contact/joining-pdc-mailing-lists-1.736925>, check on the PDC Events webpage <https://www.pdc.kth.se/about/events> or follow PDC on Facebook at <https://www.facebook.com/kth.pdc>.

PDC-Related Events

PDC Summer School 2022

15-26 August 2022

For details, see <https://www.pdc.kth.se/summer-school/2022>.

HPC Sources

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

HPC in Sweden and Scandinavia

- SNIC
<https://snic.se>
- SeRC
<https://e-science.se>
- SeSE
<http://sese.nu>
- NeIC
<http://neic.no>
- ENCCS
<http://enccs.se>

European HPC ecosystem

- HPC in Europe
<https://hpc-portal.eu>
- EuroHPC
<https://eurohpc-ju.europa.eu>
- PRACE
<https://www.prace-ri.eu>
- LUMI
<https://www.lumi-supercomputer.eu>
- ETP4HPC
<https://www.etp4hpc.eu>
- EOSC
<https://eosc-portal.eu>

A selection of projects that PDC is involved with

- BioExcel CoE
<https://bioexcel.eu>
- EBRAINS
<https://ebrains.eu>
- PerMedCoE
<https://permedcoe.eu>
- EXCELLERAT
<https://www.excellerat.eu>
- EOSC-Nordic
<https://eosc-nordic.eu>
- DICE
<https://www.dice-eosc.eu>
- HPC-Europa3
<http://www.hpc-europa.eu>

HPC news sources

- HPCwire
<http://www.hpcwire.com>
- insideHPC
<https://insidehpc.com>