no.2 2022

PDC Center for High Performance Computing

PDC Newsletter

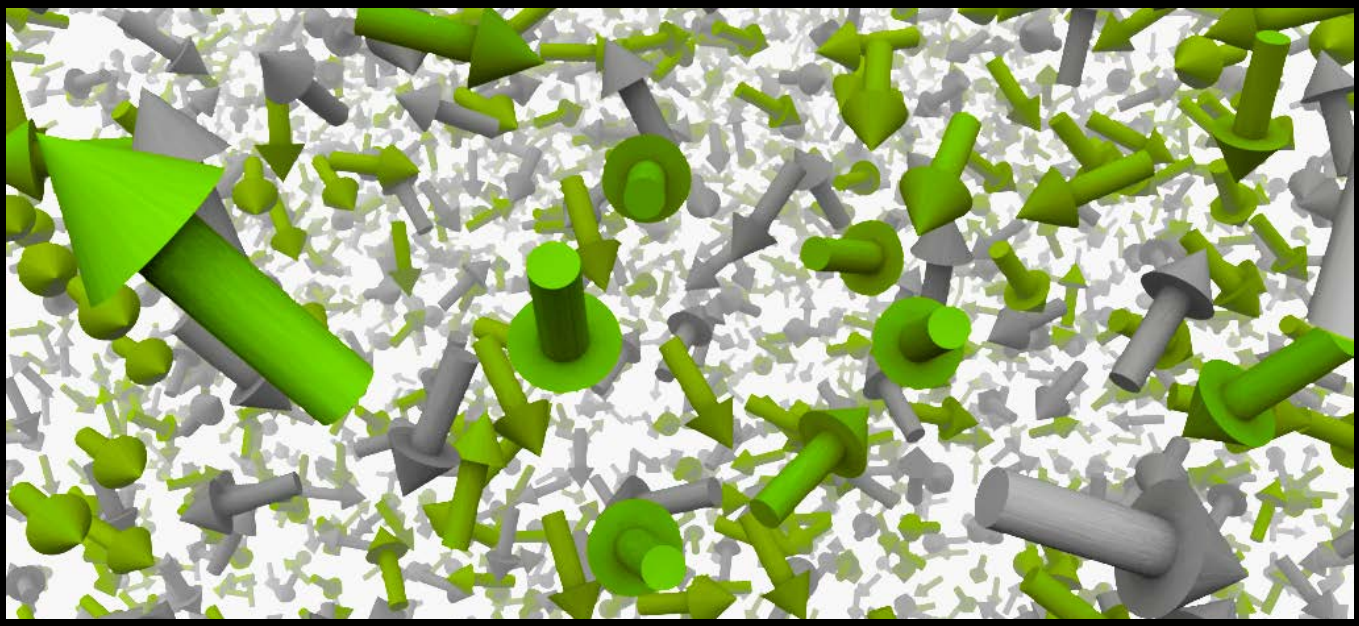

The PDC Newsletter is published by PDC at EECS, KTH.

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

Editor: Dirk Pleiter
Editorial staff: Genet Edmondson
Graphic Design: Maria Malmqvist
E-mail: support@pdc.kth.se

ISSN 1401-9671

Cover

The image on the cover was created from data produced by a simulation using the UppASD code. Images such as this are known as visualisations; this one enables us to visualise what the dynamics of interacting magnetic moments in a magnetic solid would look like.

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

With this second newsletter for 2022, it is time to look back at what has been achieved during the year. While working conditions have improved significantly, thanks to the changing pandemic situation, it has been a challenging year for PDC on various levels.

The ongoing installation of new hardware resources has been a major focus this year. In fact, the last hardware components for the Dardel high-performance computing (HPC) system are arriving at the time of writing. In recent months, there have been many changes in the system during which the team at PDC worked (often at very late hours) to keep the already deployed CPU partition of Dardel running and as stable as possible, despite various issues arising that were outside PDC's control. Although it has not been possible to continuously maintain the target level of availability as a result of these disruptions, we are confident of resuming this level of service once the full upgraded system has passed all the acceptance tests during the first quarter of 2023. A very notable addition to the system was the graphics processing unit (GPU) partition, which pushed Dardel to 68th position on the Top500 list, a list that ranks the systems according to the throughput of floating-point operations while running the high-performance Linpack benchmark (see *Dardel Fastest in Sweden*). Even more notable is that Dardel, with the new GPU partition, is now at position 5 in the Green500 list, which measures power efficiency by dividing the throughput of floating-point operations measured for the Top500 list by the amount of power that is consumed. The significance and importance of operating Dardel as one of the most power-efficient systems in the world are obvious, particularly given the significant increase in electricity prices. In addition, progress is being made on establishing Dardel's little sister, namely the Dardel Cloud, which will be available for users soon (see *Building the Dardel Cloud*). This resource will open up new usage models with researchers being able to deploy virtual servers with a lot of freedom regarding how to configure them. Also, data sharing according to FAIR principles (that is, with data being findable, accessible, interoperable and reusable) will become easier with a new object store.

With all of this new infrastructure which is based on the most recent technologies, support for researchers and engineers who use these infrastructures has to grow to keep pace with the technological hardware developments. The upcoming changes in the Swedish HPC landscape – due to the transition from the Swedish National Infrastructure for Computing (SNIC) to the National Academic Infrastructure for Supercomputing in Sweden (NAISS) at the start

of next year – provide a good opportunity to think about leaping forward in respect to how support is provided for developing and using research software for HPC (see *Supporting Research Software Development*). Such support will be provided to specific research communities through several new EuroHPC Centres of Excellence (see *Successful Applications for EuroHPC CoEs*), and support is also available through ENCCS and SNIC/NAISS. Throughout this year, PDC continued both developing software for researchers, for example, in the area of brain research (see *Describing Connectivity in Neuronal Network Models*), as well as supporting research communities through community training workshops and summer schools (see *UppASD for Simulating Atomistic Spin-Dynamics* and *BioExcel Summer School Live Again!*). A particular priority for next year will be to help researchers leverage the power of the new Dardel GPU partition.

PDC will also continue to strongly engage in broader training and education efforts. For instance, the CodeRefinery training events (see *CodeRefinery Continues*), the series of PDC Summer Schools (see *PDC Summer School 2022*), and the series of introductory workshops (see *Introduction to PDC Systems Workshop*) are being continued. We are happy to be maintaining close cooperation with the EuroHPC-funded Competence Centre ENCCS (see *ENCCS Enters Its Second Phase*), for instance, with a joint GPU hackathon in March 2023. With ENCCS, PDC also shares the intention and ambition to encourage and facilitate wider business use of HPC (see *Presenting Dardel for Businesses*).

This editorial began by describing the significant world-class upgrades to the infrastructure at PDC. These were made possible by venturing to take advantage of the most modern cutting-edge technology. To stay at the forefront of technological developments, it is crucial for KTH to be involved in HPC technology projects funded by EuroHPC (see *Programming Models and I/O with DEEP-SEA and IO-SEA*), as well as being involved in the realisation of new HPC usage models like digital twins (see *Paving the Path for Digital Twins in HPC*).

Let me finish with big thanks to the PDC team, which can be proud of all its achievements. PDC is currently undergoing a major review of its internal structures, and the interviews with all staff members confirmed that PDC is home to a diverse set of highly talented people. This is a strength of the organisation, on which we can build with confidence when we reconvene in 2023.

For the break at the turn of the year, I wish you all a well-deserved peaceful vacation, and I hope all of you will enjoy a good start to the new year and success in 2023.

Dirk Pleiter, Director PDC

In This Issue

Staff Focus

Parasuram Indraganti

Parasuram Indraganti has recently joined PDC's support team as a first-line supporter. He is studying for an M.Sc. in the fluid mechanics track of the Engineering Mechanics program at KTH. He developed an interest in high-performance computing (HPC) while leveraging the power of HPC for computational fluid dynamics (CFD) problems. Parasuram views his work in first-line support as a first step in building a strong foundation for his HPC skills and looks forward to being involved at a more advanced level in the usage and development of HPC tools with an orientation towards CFD applications. Through his current role, Parasuram is gaining a more holistic insight into HPC, not only from the viewpoint of using HPC systems to run simulations. He is also developing an understanding of what goes on in the background to maintain a cluster (from both the hardware and software systems perspectives) and how all these aspects need to be integrated to result in the best performance. On a parallel note, he also feels that this has been a great way to be introduced to the Swedish work system and enjoys working in a multicultural environment, where he learns not only about HPC, but also about different cultures.

...continued on page 18

UppASD for Simulating Atomistic Spin Dynamics

Johan Hellsvik, PDC

The Uppsala Atomistic Spin Dynamics (UppASD) project is one of the ongoing code development projects being run by the Division of Materials Theory at Uppsala University and the Department of Applied Physics at the KTH Royal Institute of Technology. The researchers have developed a program, known as UppASD (https://github.com/UppASD/UppASD), for simulating atomistic spin dynamics at finite temperatures, which makes it possible to describe magnetization dynamics on an atomic level. For example, the cover image (by Anders Bergman from Uppsala University) shows the dynamics of interacting spins in a paramagnetic phase where the system is not long range ordered.

Using UppASD for Simulations

Simulations with UppASD are often performed as part of multiscale modelling of magnetic materials. Modern implementations of first principles density functional theory (DFT) can be used to calculate ground state properties of a crystalline or nanoscale magnetic material, but do not by themselves give access to the dynamics of the spin polarization over extended time and length scales. Atomic magnetic moments can be defined for each lattice site as the integral of the magnetization density over atomic volumes, and thereby constitute a coarse grained representation of the magnetization. Over nanometer length scales, the interactions between magnetic moments are mainly of quantum mechanical origin. The strength of these interactions can be calculated from the DFT ground state solution of the material. The magnetic moments and their interactions are the required ingredients for constructing materials-specific effective magnetic Hamiltonians.

The magnetic phase diagram and thermodynamical properties of a magnetic Hamiltonian can be investigated with techniques for Monte Carlo simulations. Alternatively, the equations of motion of the Hamiltonian can be investigated. In the atomistic spin-dynamics method, the dynamics of the Hamiltonian are treated in the semiclassical limit. This enables simulations that readily can consist of hundreds of thousands of spins evolved over hundreds of picoseconds. An important capability of atomistic approaches is that short-ranged local interactions and correlations, and the spin texture over longer distances, can be studied simultaneously. This is of importance, for instance, in simulations of domain wall motions and for topological excitations [1].

In condensed matter physics, a key knob for tuning the properties of materials is to change their chemical composition. For magnetic materials, doping can be used to alter the interactions

between magnetic moments, as well as the magnetocrystalline anisotropy. Modelling of chemically disordered materials is, in general, more demanding than modelling pristine material with perfect crystalline structure. The UppASD code supports the modelling of disordered material, as was done for instance in modelling of ultrafast magnetization dynamics in amorphous Gd-Fe alloys [2], and the multiferroic phase of the insulating magnet CuO [3]. When chemical disorder is treated on atomic length scales, the computational effort can drastically increase given that averaging needs to be done over simulations of different realizations of the disorder. In methods for simulating coupled spin-lattice dynamics the atomistic spin-dynamics method is augmented to include also the ionic motion degrees of freedom. Support for spin-lattice dynamics simulations has been implemented in UppASD, with modelling having been done for few body systems, as well as for ferromagnetic iron [4].

UppASD Autumn School 2022

A three-day course on modelling atomistic spin dynamics with UppASD was held at the KTH Royal Institute of Technology in Stockholm in October. The school included lectures on the background physics concepts relating to atomistic spin-dynamics, plus tutorials where attendees practised using the UppASD program with assistance from mentors. Amongst the thirty-eight participants at the school, there were early pioneers of the field of atomistic spin dynamics and the lead developers of UppASD, as well as new and more experienced users of the program. The hands-on exercises were run on the CPU partition of the new Dardel supercomputer system at PDC.

UppASD can be built and run on desktop and laptop computers, as well as on supercomputers. The dynamics and characteristics of toy models can sometimes be explored in short and small simulations on, for example, a laptop, whereas more careful investigations of more refined models require extensive computing resources. Consequently it was beneficial for the participants to gain practical experience running UppASD on Dardel during the school's hands-on sessions.

The school was partially sponsored by the Swedish e-Science Research Centre (SeRC) and was organised by PDC (Johan Hellsvik) together with researchers from Uppsala University (Anders Bergman, Manuel Pereiro and Olle Eriksson), KTH (Banasree Sadhukhan, Zhiwei Lu and Anna Delin) and Örebro University (Danny Thonig). Some important parts of the preparations for the school included upstreaming recently developed functionality to the main branch of the UppASD code repository, revising the code documentation and example directories, and writing UppASD tutorial material (https://uppasd.github.io/UppASD-tutorial). This general overhaul of the program resulted in a new release of the code, incrementing the lead version number to six and constituting a good stepping stone for future work on porting selected kernels to GPUs.

References

1. M. Pereiro, D. Yudin, J. Chico, C. Etz, O. Eriksson, and A. Bergman, "Topological excitations in a kagome magnet", Nature Comm. 5, 4815 (2014). https://doi.org/10.1038/ncomms5815
2. R. Chimata, L. Isaeva, K. Kádas, A. Bergman, B. Sanyal, J. H. Mentink, M. I. Katsnelson, Th. Rasing, A. Kirilyuk, A. Kimel, O. Eriksson, and M. Pereiro, "All-thermal switching of amorphous Gd-Fe alloys: Analysis of structural properties and magnetization dynamics", Phys. Rev. B 92, 094411 (2015). https://doi.org/10.1103/PhysRevB.92.094411
3. J. Hellsvik, M. Balestieri, T. Usui, A. Stroppa, A. Bergman, L. Bergqvist, D. Prabhakaran, O. Eriksson, S. Picozzi, T. Kimura, and J. Lorenzana, "Tuning order-by-disorder multiferroicity in CuO by doping", Phys. Rev. B 90, 014437 (2014). https://doi.org/10.1103/PhysRevB.90.014437
4. J. Hellsvik, D. Thonig, K. Modin, D. Iuşan, A. Bergman, O. Eriksson, L. Bergqvist, and A. Delin. "General method for atomistic spin-lattice dynamics with first-principles accuracy", Phys. Rev. B 99, 104302 (2019). https://doi.org/10.1103/PhysRevB.99.104302

Dardel Fastest in Sweden

Gert Svensson, PDC

Dardel is now the fastest high-performance computing (HPC) system in Sweden and fifth in the worldwide Green500 list (https://www.top500.org/lists/green500/2022/11)! In recent months, Dardel has undergone a number of expansions and upgrades. The most significant is that 56 graphics processing unit (GPU) nodes have been installed. They are now in operation and in the process of undergoing acceptance tests.

As previously announced, the plan was for Dardel to have 56 GPU nodes, each with four AMD Instinct™ MI250X GPUs; all of those nodes have now been installed in the system. That increased the speed significantly and resulted in a performance of 8.2 petaflops for the GPUs on the High-Performance Linpack (HPL) benchmark. This means that Dardel's HPL performance went up more than three times.

Dardel is now in 68th place on the latest TOP500 list, which was released in November (see https://www.top500.org/lists/top500/2022/11). Note that the GPU partition of Dardel is referred to as "Dardel GPU" in both the TOP500 and Green500 lists in order to distinguish it from the earlier CPU-only phase of the system, which was previously listed just as "Dardel". The CPU partition of the system is now referred to as "Dardel CPU" and is in 345th place on the TOP500 list. As mentioned earlier, "Dardel GPU" is in fifth place on the Green500 list, and it is interesting to observe that the systems in positions two to seven on the Green500 list are all occupied by HPE Cray EX systems with AMD Instinct™ MI250X GPUs. This clearly demonstrates that this type of architecture is highly energy efficient, measured in floating-point operations per watt.

Below: The new CPU rack has just been moved into place.

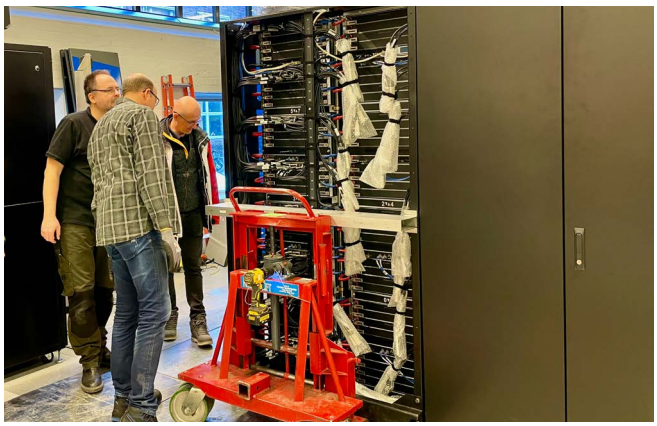

When the previous issue of the PDC newsletter was published, most researchers had been migrated from PDC's previous systems to Dardel, although the Scania partition of Dardel was not fully operational at the time. Thanks to the upgrades and installations, 240 of the Dardel CPU nodes are now serving Scania's research and development. Recently, 468 nodes were added to the system for SNIC academic researchers, including eight two-terabyte "Giant" nodes. In addition, twelve nodes (which are dedicated for research by the Department of Astronomy at Stockholm University) have been installed.

The entire software stack of Dardel has been upgraded to a more current release called Strawberry. At the time of the writing, this release was being tested on an "island" of the Dardel system. Also, the access mechanism for the Lustre file storage system has been changed so it will be compatible with future software versions – more specifically it was changed from Remote Direct Memory Access (RDMA) technology to the Transmission Control Protocol/Internet Protocol (TCP/IP). This may temporarily decrease the speed of the disks until the disk software is also updated, after which the speed should be similar or better than before.

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 now in the Dardel system.

Below: One of Dardel's GPU boards

PDC has been working with Hewlett Packard Enterprise (HPE) to make all these significant changes to Dardel while minimising effects on the system operation. Most of the work has been possible without affecting the researchers using the system.

There are still some major hardware and software upgrades ahead of us. The interconnect in Dardel will soon be upgraded to the next generation of Slingshot network with a speed of 200 gigabit/second. This will require that all the Slingshot cards and some cables be replaced, which will require system downtime of a week.

Other upcoming upgrades include updating the Lustre disk system with 50% more capacity in terms of both size and speed and upgrading the software in the disk system to a new version. As mentioned earlier, that software upgrade should optimize TCP/IP access so the access speed should be similar to that with the previous RDMA access. The Dardel system software will also be upgraded to several new releases that will provide a lot of new functionality, especially for supporting the GPUs.

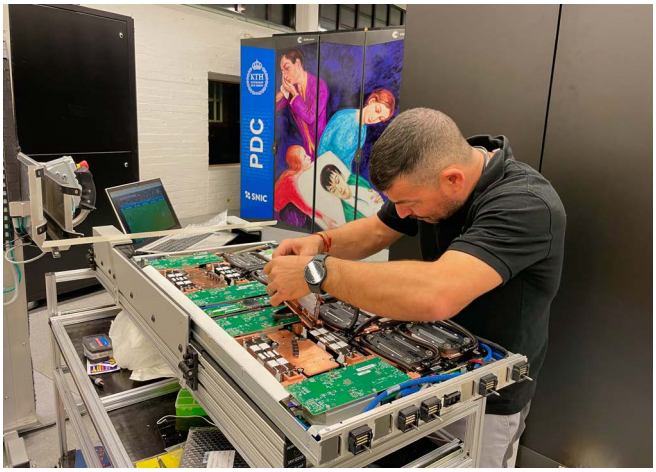

Above: An engineer from HPE is repairing a node.

Building the Dardel Cloud

Ilari Korhonen & Mustafa Arif, PDC

PDC is currently in the process of building a new cloud platform, known as the Dardel Cloud, which is targeted at researchers who wish to build and run their own application stacks in dedicated virtual environments. Users will be able to deploy domain-specific portal services, run less compute-intensive workflows, utilise virtual machines for hosting databases, and implement data curation pipelines and data sharing services as well as interactive and collaborative work using interactive frameworks like Jupyter notebooks. This platform will enable researchers to utilise cloud computing and storage resources through PDC and thus be able to benefit from both the close proximity of high-performance computing (HPC) resources, as well as the expertise in scientific computing at PDC.

In the first stage of the cloud, PDC will offer IaaS (Infrastructure as a Service) in the form of Virtual Machines (VM) of different sizes and operating environments. In the future, we are planning to build upon this by offering more types of services, such as container runtime environments. The Dardel Cloud is expected to be generally available early next year, and the plan is for the first pilot users to be onboarded during the last quarter of this year.

The cloud platform will come with a large-scale object store based on Ceph (a highly scalable open-source software-defined storage solution designed to address today's high-growth storage requirements). The initial gross capacity of this storage will be slightly above 1 PB. This object store will also be accessible through an S3 interface – this is a type of interface that is widely used in public clouds. This will make the Dardel Cloud storage system much more suitable for data sharing in accordance with the FAIR principles, which aim to ensure that data is "Findable, Accessible, Interoperable and Reusable".

Together with our service providers IP Solutions and Engin IT, PDC is working on making the Dardel Cloud a secure compute and storage resource that can be used for research on advanced science challenges together with the Dardel HPC system.

For news about the transition from the Swedish National Infrastructure for Computing (SNIC) to the National Academic Infrastructure for Supercomputing in Sweden (NAISS), see the SNIC Newsletter Winter 2022 on the SNIC home page at https://www.snic.se.

Supporting Research Software Development

Dirk Pleiter, PDC

As the digitalisation of science progresses, an increasing number of research and engineering teams are relying more and more on software [4]. Scientific progress, therefore, depends on this software being continuously enhanced according to the needs of those teams while also being maintained sustainably. As the lifetimes of the computer architectures on which this software is being used are typically much shorter than the lifetimes of scientific software, this software must not only be portable but also performance portable. This means that it must be possible to run the software on different computer architectures and perform at a similar level of efficiency. All these observations are, in particular, true for research software being used on high-performance computing (HPC) systems. While all modern HPC architectures are extremely parallel, the parallelism may be at very different levels. For instance, some systems may come with more but less powerful nodes based on CPUs only, while other systems may feature fat nodes with several graphics processing units (GPUs) that are used as compute accelerators. Therefore, software applications need to be able to adapt to the scale of the different architectures where they may be used and run efficiently irrespective of the inherent parallelism of the system architectures.

Most people in academia would probably agree with the statement that research software is becoming increasingly relevant to research and development. However, there is frequently a dichotomy between being an expert in a given research field and being an expert in developing and/or using HPC software for that particular domain. This distinction gives rise to a need for support to be provided to assist research domain experts in the use and/or development of software for their domain. This article reflects on this evolving need and, in particular, on whether there is currently sufficient support within academia for developing and maintaining such software. For Sweden, now is an ideal time to consider this question. The pending changes to the Swedish HPC research infrastructure offer an opportunity to trigger improvements in how Sweden provides the necessary support for and makes progress in this area. (This task has been referred to as the provision of application expert support services.) Supporting the development of performance portable research software is not a standalone process - it should be regarded as an effort that involves linking research infrastructures with researchers in an optimal way that leads to the production of high-quality research. In the case of an HPC research infrastructure, that link is primarily via the people who develop or maintain suitable research software or who assist the researchers to make the best use of that software on the available HPC systems.

At this stage, there is no simple answer to the question of whether there is adequate support, either globally or in Sweden, for developing research software. This article will focus only on one aspect of this question, namely the role of the people who develop such software. Too often, their role and the significance of their contributions are not fully recognized. Depending on the particular research and engineering domain, opportunities for scientific recognition through publications are often lacking for these people, and career paths are mostly absent. Part of the problem has been that there was no universal concept or term for this type of work. (In the Swedish context, the term application expert has been used for people undertaking, among others, these kinds of tasks.) One big step towards overcoming this situation consisted of establishing the concept of

a research software engineer (RSE). This move towards a global concept started more than ten years ago in the UK, and the momentum of the RSE movement has been growing since then with international attention to the importance of RSEs increasing significantly in recent times. A Nordic RSE initiative was started in 2018 (see https://www.pdc.kth.se/publications/pdc-newsletter-articles/2018-no-1/the-nordic-research-software-engineer-initiative-1.823934), and for more information about the general development of the RSE movement, see the overview in [2]. For details about a very recent activity addressing this area, see the RSE-HPC-2022 workshop (https://us-rse.org/rse-hpc-2022) that was held during the recent SC22 supercomputing conference. Meanwhile, RSEs continue to be a vital part of the research communities utilising HPC, and organisations have been established that have hired dozens of RSEs (explicitly as RSEs or as applications experts or with other job titles).

While promoting the concept of RSEs is an essential step towards creating awareness of the important work done by the many such people in academia, it does not define a specific job profile for an RSE. Although a focus on developing research software is a common denominator of the work done by RSEs, there are many options regarding the roles, skills and primary tasks of any given RSE. In addition, there are many options for where to place RSEs within academic organisations and how to organise the engagement of RSEs with the researchers they support. To complicate matters, what constitutes a good choice for many of these options is highly context-dependent and can, for example, vary significantly with the type of research being done and the area it is in and the computer architectures that are being used. Some of the possible choices are discussed in the following.

Let us first start with the question of how to engage RSEs with research teams. From the perspective of the research team, short-term assignments with a duration of a few months would enable RSEs to join the team temporarily and address tasks with a narrow scope. For example, an RSE could provide advice on porting code from one HPC system to another or on refactoring code (that is, rewriting code to make it portable or making significant changes to the code structure to make it easier to maintain). As part of such assignments, RSEs could assist in the implementation of proofs-of-concept, for instance, to test whether a certain way of rewriting code would have the desired result. While such short-term engagements allow, in principle, for the dynamic assignment of RSEs, long-term assignments will be much more important, particularly those related to the implementation of new codes or the refactorisation of existing codes to support new architectures. For example, assignments to enable codes for GPU-accelerated HPC systems like Dardel or future EuroHPC systems, including the upcoming EuroHPC exascale system JUPITER, will be significant, and such tasks may take several years.

The question of the requisite skill profiles for RSEs is heavily impacted by the underlying question of the expected outcomes of each particular task or assignment. Developing research software for HPC is particularly demanding because of the diversity of skills that is required. Knowledge and experience in best practices of modern software engineering, which includes, for example, the use of modern C++, is only sufficient to implement software components that have already been well specified by others. An understanding of modern computer architectures is needed for performance engineering tasks. However, only the combination of a domain-specific background plus knowledge of relevant numerical algorithms will give an RSE the abilities necessary for playing a leading role in the development of a complex research code. Such a background will also simplify the interactions with researchers in that domain based on a common understanding of the problems that are being addressed with a given code or workflow.

Thirdly, let us consider the question of where to place RSEs within academic organisations. To answer this question, trade-off decisions need to be made. On the one hand, an RSE should work closely with researchers in a specific domain to

develop a good understanding of their needs and requirements. Allowing an RSE to act as a member of a research team or community can be an important factor in attracting people who are highly qualified for such tasks. On the other hand, RSEs also need a range of skills that are relevant across multiple research areas as they relate more to computer systems than specific research topics. Some examples are advanced software engineering skills, or the ability to transform and port code for using compute accelerators like GPUs, adopt software quality standards or implement standards for findable, accessible, interoperable and reusable (FAIR) software. In many cases, these types of skills need to be continuously improved as HPC systems evolve. Hence RSEs would benefit from working together with other RSEs and participating in joint upskilling efforts. In fact, various universities and other organisations have established dedicated organisational units for their RSEs, which have been shown to work successfully. (See [3] for some case studies.) Note that applying the latter approach would still allow for RSEs to be placed in research groups for part of their time. This has worked well for various PDC staff members involved in software development. In this context, engagement in research software development community organisations has also proven to be useful in countries like the UK (see https://society-rse.org) and the US (see https://us-rse.org).

The realisation of the concept of research software engineering has come a long way in the past ten years, and the benefits of research software engineering can now be considered as being well established. However, as a discipline, research software engineering is still in a state of relative infancy with many choices for specific realisations yet to be explored. The academic HPC research infrastructure in Sweden is currently undergoing an important transition towards consolidation of hardware resources. This makes it possible to reduce the costs for the provisioning of hardware. It also allows for a transition towards a more service-oriented provisioning of HPC and other resources based on a separation between the lower-level generic layer of e-infrastructure services (provided by the underlying physical HPC and data storage resources) and the higher-level services, which can be domain-specific services as described in [1]. Such higher-level services can include, for instance, web-based services that are specific to particular research domains and also domain-specific support services (such as those provided by RSEs with expertise in a particular research area). The important task of improving, broadening and further organising the support available to Swedish academic HPC users (by improving the support provided for research software development) needs to be addressed in parallel with the transitions that are happening in the Swedish HPC infrastructure for academic research. Sweden currently has a great opportunity to target and take advantage of synergies between the ongoing evolution of its HPC research infrastructure and this window of opportunity to strengthen the support that is provided to assist researchers to make the best possible use of that infrastructure. Exploiting this opportunity would benefit Swedish research as a whole.

References

1. S. Alam, J. Bartolome, M. Carpene, K. Happonen, J.-C. Lafourcriere, D. Pleiter, Fenix: a Pan-European federation of supercomputing and cloud e-infrastructure services. Communications of the ACM. 2022. https://doi.org/10.1145/3511802
2. J. Cohen et al., The Four Pillars of Research Software Engineering. IEEE Software. 2021. https://doi.org/10.1109/MS.2020.2973362
3. D. S. Katz et al., Research Software Development & Management in Universities: Case Studies from Manchester's RSDS Group, Illinois' NCSA, and Notre Dame's CRC. IEEE/ACM 14th International Workshop on Software Engineering for Science. 2019. https://doi.org/10.1109/SE4Science.2019.00009
4. J. Switters, D. Osimo, Recognising the Importance of Software in Research - Research Software Engineers (RSEs), a UK Example. Publications Office of the European Union, 2019. https://doi.org/10.2777/787013

ENCCS
EuroCC National Competence Centre Sweden

ENCCS Enters Its Second Phase

Lilit Axner, ENCCS

The EuroHPC Joint Undertaking (EuroHPC JU) allows the European Union and the countries participating in the EuroHPC JU to coordinate their efforts and pool their resources to make Europe a world leader in supercomputing (https://eurohpc-ju.europa.eu). One of the aims of the EuroHPC JU is to extend the use of the European supercomputing infrastructure to new public and private user groups and facilitate the uptake of high-performance computing (HPC) skills in European universities and industries through the formation of national competence centres (NCCs) in each European country.

This EuroHPC JU effort has established a large network of competence centres across Europe which is of great value for participating countries. In particular, the NCC network is contributing to

- accelerating the adoption of HPC in the region through a rapidly growing European HPC ecosystem,
- helping countries to strengthen their role as innovation drivers in the global market, and
- supporting and expanding competencies in research and development in academia, industry and the public sector.

Many HPC users are already familiar with the Swedish NCC, which is called ENCCS (https://enccs.se). It was established on the 1st of September 2020, and since then, ENCCS has had a significant impact on the Swedish HPC community and also helped many new researchers and engineers to start using HPC.

ENCCS has built a solid user base and already has a large number of followers from academia, industry, and the public sector. In January 2023, ENCCS will enter its second phase of existence, which consists of three more years and even more ambitious plans compared to the initial two-year start-up phase. ENCCS will continue providing support to researchers from academia, industry, and the public sector (to assist them to utilise the EuroHPC JU supercomputers) and the centre is also planning to organise over 25 training events per year during the second phase. In addition, ENCCS will:

1. reach out to new Swedish users both from new disciplines and from universities and organisations that are currently not the traditional HPC users,
2. analyse and tackle challenges related to software licences,
3. expand the training curriculum, investigate possibilities for HPC certification, and leverage its large international and national network to create an extensive training material database, and
4. contribute its offerings to a large European marketplace for HPC competencies which is being developed in a collaborative effort by all European NCCs.

If you, as an HPC user, are wondering whether EuroHPC JU supercomputers can help you in your research, do not hesitate to get in touch with ENCSS at info@enccs.se. The team will be happy to discuss your research and do its best to help you! You are also welcome to subscribe to the ENCCS quarterly newsletter: https://enccs.se/newsletter.

Want to access HPC computational and storage resources for your project?

Get free help, support and advice!

ENCCS's experts provide hands-on advice to researchers writing EuroHPC applications and can share experience from previous successful projects, for example, about technical proposal requirements. ENCCS can also assist any HPC users who require new technical access mechanisms to EuroHPC resources. Taking advantage of practical advice from ENCCS allows you to focus on your project instead of getting bogged down in application bottlenecks. For more information, see:

https://enccs.se/proposal-support.

Paving the Path for Digital Twins in HPC

Dirk Pleiter, PDC

The concept of digital twins (DT) was established more than twenty years ago. But, over the last five years, the topic has been receiving increasing amounts of attention, both in academia and industry. This can, for example, be seen in a tremendous increase in the annual number of publications [2]. For a digital twin, one creates and operates one or more virtual representations of a physical entity. The virtual representations contain relevant information about the physical entity, which is ideally much cheaper to obtain compared to using the latter. When operating the digital twin, the physical entity and its virtual representations are connected, which means that state information is exchanged in both directions.

The initial focus of research on digital twins was product life cycle management, which covers the creation, production and operation of products [1]. In industry, this is still of significant interest, for instance, in the development of autonomous driving solutions or for the realisation of preventive maintenance concepts. The use of digital twins for large ecosystems is a more recent development. The most outstanding initiative in this direction is Destination Earth (see https://digital-strategy.ec.europa.eu/en/policies/destination-earth), which aims to develop a highly accurate digital model of the earth on a global scale. One goal is to facilitate informed policy decision-making by means of digital twins, as such a twin of the earth would allow us to explore the impact of human activities on natural phenomena and vice versa. One important feature of Destination Earth is the need for high-performance computing (HPC) to model the weather and climate, and the behaviour of the oceans, as well as surface and subsurface water flow, amongst other things.

One effort that is working towards realising the vision of Destination Earth is the BioDT project (https://biodt.eu), where PDC is one of more than twenty partners throughout Europe. The goal of this project is to design and demonstrate a digital twin platform that will improve our understanding of biodiversity dynamics by developing means that provide advanced modelling, simulation and prediction capabilities. The work will be guided by various real-life use cases. One class of use cases is concerned with the question of how the evolution of species is impacted by environmental changes. A more specific question is, for instance, how bird populations will change in the northern part of Europe depending on forest management strategies. Addressing this question requires the integration of forest and biodiversity models, as well as all the necessary environmental data. Another important use case relates to food security. It is being observed that environmental changes are accelerating the decline of pollinators like honey bees. Multiple factors are assumed to impact honey bee dynamics. BioDT plans to port agent-based models to HPC systems and to add support for dynamic model updating based on updated environmental data.

The integration and upscaling of various models, as well as the integration of various data sources from already existing research infrastructures in this field, will be one of the key challenges for BioDT. Once this is achieved, researchers will be able to exploit Europe's most powerful supercomputer system, namely LUMI. At the same time, the project will have paved the path for getting future supercomputing systems ready to deploy and operate digital twins.

References

1. M. Grieves, J. Vickers, Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behaviour in Complex Systems. Springer International Publishing, 2017.
2. M. Sjavor et al., The Digital Twin Concept in Industry – A Review and Systematization, proceedings of the ETFA 2020 conference.

BioExcel Summer School Live Again!

Alessandra Villa and Rossen Apostolov, PDC

Over the last three years, the BioExcel Centre of Excellence for Computational Biomolecular Research (https://bioexcel.eu) has been actively involved in supporting the biomolecular simulation research community. In particular, the BioExcel Summer School is our flagship annual event and is one of the most popular training events organised by the centre.

After having to hold four of the recent BioExcel schools, including the 2020 and 2021 summer schools, remotely due to the COVID pandemic, we were really glad to be able to once more run the summer school with onsite meetings in June this year. The organisers, trainers and attendees were very excited to meet each other face-to-face. It was great to be able to see the audience in person while presenting and to have a coffee break while discussing questions. The school covered a broad range of topics, such as molecular dynamics simulations, biomolecular docking, free energy calculations, advanced sampling methods, workflow solutions with BioExcel Building Blocks (BioBB), and quantum mechanics/molecular mechanics (https://bioexcel.eu/events/bioexcel-summer-school-on-biomolecular-simulations-2022). The summer schools are targeted particularly at early-stage researchers. As with other BioExcel training events, several travel bursaries were offered to assist young researchers to attend this school. You can read about how inspiring and valuable this year's travel grant winners found their experiences at the course (https://bioexcel.eu/hear-from-our-2022-bioexcel-summer-school-travel-grant-beneficiaries). There were also prizes for posters as part of the school and, as is traditional, the winners were given the opportunity to present their work in the BioExcel Student webinars: https://bioexcel.eu/student-webinar-summer-school-2022-edition-2022-09-15.

In addition to the annual schools, BioExcel runs a lot of other initiatives to support the biomolecular simulation community, and PDC has been a leader in many of them.

- The BioExcel webinars feature notable developments in the field of computational biomolecular research, and interested researchers are invited to find out about the latest tools and techniques: https://bioexcel.eu/category/webinar.
- The BioExcel forums, https://ask.bioexcel.eu and https://gromacs.bioexcel.eu, provide support for the community, so check out the existing questions and answers or ask a new question if you need some assistance with your research.
- BioExcel also provides online training materials: you can browse the lectures that are available on the BioExcel YouTube channel https://www.youtube.com/c/BioExcelCoE or view a range of GROMACS tutorials at https://tutorials.gromacs.org.
- In addition, there is a biomolecular competency hub at https://competency.ebi.ac.uk/framework/bioexcel/3.0 where you can browse competencies, career profiles and training resources to advance your career.

Above: Presenters and attendees at the BioExcel Summer School on Biomolecular Simulations 2022, Science and Technology Park of Sardinia, Pula, Italy, 12-17 June 2022

CodeRefinery Continues

Johan Hellsvik, PDC

The CodeRefinery project (see https://coderefinery.org) is a collaboration between Nordic and international partners to provide training to researchers, software engineers and students about best practices relating to research software that is being used within academia.

The project is now well into the first year of the third phase of the collaboration. Over the last two years, CodeRefinery has been pioneering methods for online teaching that make it possible to provide efficient, dynamic, and interactive teaching of computing skills to groups with as many as two hundred or more participants within a single training event. Some of the project's important activities this autumn have included the two-day CodeRefinery work meeting, which was held onsite in Helsinki in September, and the large online CodeRefinery course (with 262 participants!) that was also held in September. If you missed that, keep a watch on the CodeRefinery events page for the next course in March 2023: https://coderefinery.org/workshops/upcoming.

The Swedish National Infrastructure for Computing (SNIC) contributes funds to CodeRefinery to enable staff working at PDC and UPPMAX to participate in the project. At the start of 2023, SNIC will transition into the National Academic Infrastructure for Supercomputing in Sweden (NAISS), and some recent good news is that NAISS has committed to a corresponding contribution over the rest of the project's third phase, which goes into 2025. This will help CodeRefinery to continue being able to offer courses from 2023 on!

Describing Connectivity in Neuronal Network Models

Mikael Djurfeldt, PDC

PDC has recently participated in a project that aimed to improve descriptions of connectivity in neuronal network models. The results have been published in a paper in the journal Public Library of Science (PLOS) Computational Biology [3]. This work was motivated by the ongoing reproducibility crisis in science [4], which also applies to the specific field of computational neuroscience (see, for example, [5]). One of the reasons why results in this field often cannot be reproduced is that the descriptions of the sometimes rather complex simulated models are insufficient and/or ambiguous to such an extent that they cannot be re-implemented.

This is illustrated by the figure on the next page, which shows raster plots of neuronal signalling events ("spikes") of two neuronal networks with "random" connectivity. In part A of the figure, each neuron receives a fixed number of incoming connections from randomly selected pre-synaptic neurons, while in part B each neuron provides connections to a fixed number of randomly selected post-synaptic neurons. Although both networks have "random" connectivity and everything else is equal, the two networks display both different network structure and different dynamics. Clearly, the statement that a neuronal population is connected randomly to another is ambiguous.

The project began with a review of published network models where the implementation of the model could be found in one of the online repositories ModelDB [6] and Open Source Brain [7]. Ambiguous network model descriptions were found in 15 out of the 42 publications that were examined. We also reviewed a set of model description languages, which included those specific to a simulator and those available on multiple platforms. Based on the review, we identified a set of recurring concepts used in network connectivity descriptions and suggested guidelines for how a network could be described unambiguously using mathematics or text, and

Below: Spiking neuron network simulations of a balanced random network with (A) fixed in-degree and (B) fixed out-degree
(Top left) Raster plots show spike times of 50 out of 10,000 excitatory (E) and 50 out of 2,500 inhibitory (I) neurons
(Bottom left) Time-resolved spike rate from spike-count histogram across time with temporal bin width of 5 ms
(Top right) Per-neuron spike rate from spike-count histogram for individual neurons
(Bottom right) Normalised distribution of per-neuron spike rates with bin width of 2/s

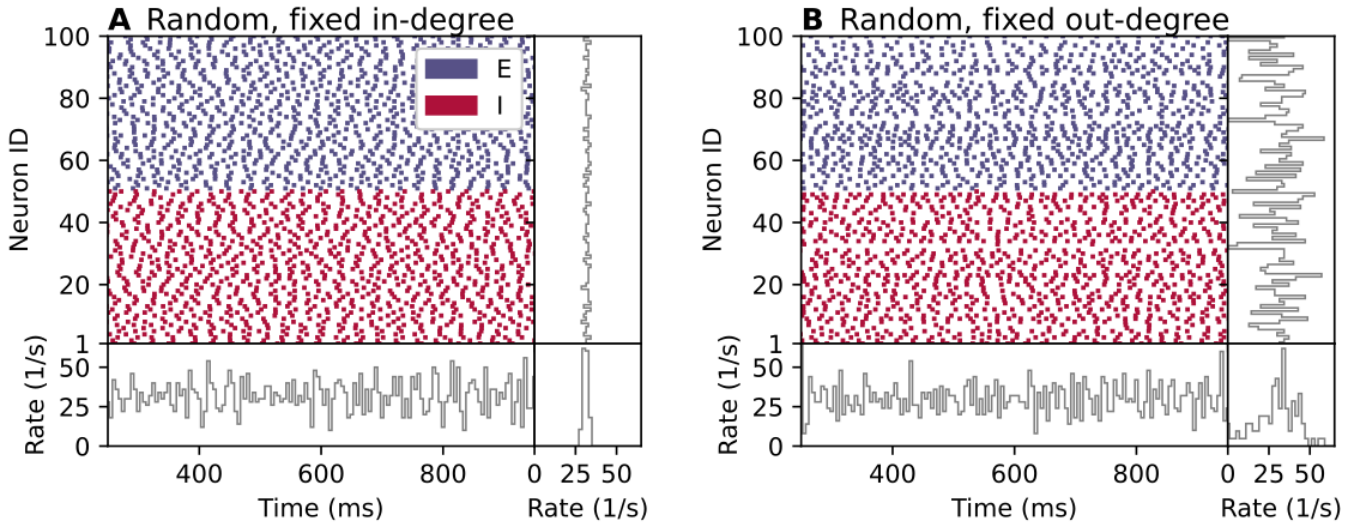

also provided a graphical notation which can be used to unambiguously describe networks in illustrations.

Some examples of concepts that were identified are one-to-one connectivity, where each source neuron is connected to one target neuron (which only receives a connection from that source) and all-to-all connectivity, where each source neuron is connected to every target neuron (see the figure below).

For the mathematical description of connectivity, the Connection-Set Algebra (CSA) [1] was used. It is a mathematical formalism for describing network connectivity using elementary sets of connections and set operators. It has parallelised implementations which help a neuronal simulator instantiate a CSA network

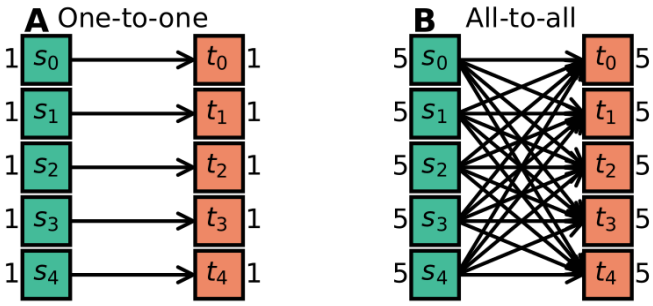

Above: Connectivity patterns reflecting the most common rules The ordered set of sources, S, is depicted by the green squares on the left. They are connected to the ordered set of targets, T, depicted by the orange squares on the right. The respective in- and out-degrees are given next to the nodes.

description efficiently [2]. This makes CSA suitable for highly scalable brain simulations.

In CSA, a connection is represented by a pair of indices *(i, j)* which refer to the entities being connected, usually neurons: *i* is the source neuron and *j* the target neuron. A connection pattern can then be described as a set of pairs of indices. In CSA it turns out to be fruitful to work with infinite sets of indices. For example, the elementary connection set $\delta = \{(0, 0), (1, 1), \ldots\}$ can be used to describe one-to-one connectivity in general, regardless of source and target population size. We can then work with CSA operators on infinite connection sets to extract the actual, finite, connection pattern at the end. Given the (finite) sets of indices $\mathbb{N}_S$ and $\mathbb{N}_T$ of the source and target neuron populations S and T, we can extract the finite one-to-one connection pattern between them through the expression $\delta \cap (\mathbb{N}_S \times \mathbb{N}_T)$ where $\cap$ is the set intersection operator and $\times$ is the Cartesian product.

Another example of an elementary connection set is the set of all connections $\Omega = \{(0, 0), (0, 1), \ldots, (1, 0), (1, 1), \ldots\}$ (where indices in both source and target positions run over all natural numbers). Note that if $S = T$ (connections go from the population onto itself), it is immediately possible to express all-to-all connectivity without self-connections as $\Omega - \delta$ where $-$ is the set difference operator. With its limited set of elementary sets

Below: Different means to describe connectivity of a balanced random network
(A) Network diagram according to our proposed graphical notation
The symbols in the annotations refer to the concepts and not the explicit parameters.
(B) Textual description of the model layout
The subscript "τ E" labels connections from source population E *to target population* $\tau \in$ {E, I}*; the same applies to "τ I" with source population* I*. $K_{\text{in},\tau\,\mathsf{E}}$ and $K_{\text{in},\tau\,\mathsf{I}}$ represent the explicit values used for the in-degrees. Note that an autapse is a connection from a neuron onto itself, while a multapse is two or more connections between the same pair of neurons.*
(C) Table according to the guidelines by Nordlie et al. [8]
(D) Equations according to the Connection Set Algebra (CSA) [1] using the index sets E *and* I

A Illustration

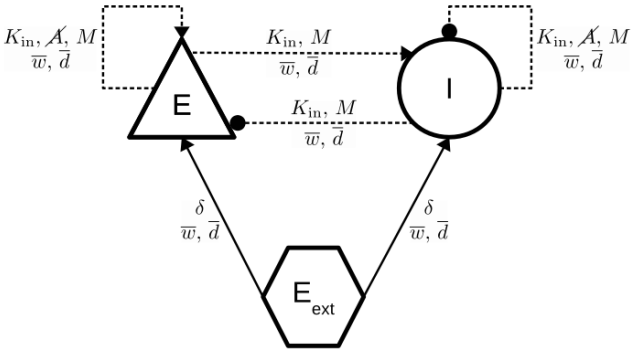

B Text

Each neuron receives $K_{\text{in},\mathcal{T}\,\mathsf{E}}$ excitatory connections randomly drawn from population E and $K_{\text{in},\mathcal{T}\,\mathsf{I}}$ inhibitory connections from population I. Autapses are prohibited and multapses are allowed. Each neuron receives additional input from an external stimulating device.

C Table

Connectivity			
Name	**Source**	**Target**	**Pattern**
EE	E	E	Random, fixed in-degree $K_{\text{in},\mathcal{T}\,\mathsf{E}}$ with multapses (autapses prohibited)
IE	E	I	Random, fixed in-degree $K_{\text{in},\mathcal{T}\,\mathsf{E}}$ with multapses
EI	I	E	Random, fixed in-degree $K_{\text{in},\mathcal{T}\,\mathsf{I}}$ with multapses
II	I	I	Random, fixed in-degree $K_{\text{in},\mathcal{T}\,\mathsf{I}}$ with multapses (autapses prohibited)
Ext	E_{ext}	E ∪ I	One-to-one

D Equation

$$C_{\mathcal{T}\,\mathsf{E}} = \langle \rho_1(K_{\text{in},\mathcal{T}\,\mathsf{E}})\mathbf{M}(\mathsf{E} \times (\mathsf{E} \cup \mathsf{I}))\rangle$$
$$C_{\mathcal{T}\mathsf{I}} = \langle \rho_1(K_{\text{in},\mathcal{T}\,\mathsf{I}})\mathbf{M}(\mathsf{I} \times (\mathsf{E} \cup \mathsf{I}))\rangle$$
$$C_{\mathsf{Ext}} = \langle \mathsf{E}_{\text{ext}} \times (\mathsf{E} \cup \mathsf{I})\rangle$$

and operators, CSA is surprisingly expressive. There are also implementation strategies which handle the infinite sets well.

The figure above shows an example network where the network connectivity is described in four ways, as a graph using our proposed graphical language, as text, as a table and as CSA equations. For details, see [3].

We regard this work as a starting point for discussion. The language and notation that we propose are mostly appropriate for a common class of deterministically and probabilistically connected networks that we find in the literature, and we also, to some extent, address networks embedded in metric space. The notations that are proposed will need to be extended to cover other domains and levels. Nevertheless, we hope that our methods will inspire further work and discussion. An advantage of the CSA notation is that it is readily transferred to software that can instantiate the network with good scaling on a supercomputer.

References

1. M. Djurfeldt, The Connection-set Algebra—A Novel Formalism for the Representation of Connectivity Structure in Neuronal Network Models. Neuroinform. 2012, 10, 287–304. https://doi.org/10.1007/s12021-012-9146-1
2. M. Djurfeldt, A.P. Davison, and J.M. Eppler, Efficient generation of connectivity in neuronal networks from simulator-

independent descriptions. Frontiers in Neuroinformatics. 2014, 8:43. https://doi.org/10.3389/fninf.2014.00043

3. J. Senk, B. Kriener, M. Djurfeldt, N. Voges, H.-J. Jiang, L. Schüttler, et al., Connectivity concepts in neuronal network modeling. PLOS Computational Biology. 2022, 18(9), e1010086. https://doi.org/10.1371/journal.pcbi.1010086
4. J.P. Ioannidis, Why most published research findings are false. PLOS Medicine. 2005 August, 2(8), e124. https://doi.org/10.1371/journal.pmed.0020124. Epub 2005 August 30. Erratum in: PLOS Medicine 2022 August 25, 19(8), e1004085. PMID: 16060722; PMCID: PMC1182327.
5. S.M. Crook, J.A. Bednar, S. Berger, R. Cannon, A.P. Davison, M. Djurfeldt et al., Creating, documenting and sharing network models. Network: Computation in Neural Systems. 2012, 23(4), 131–149. https://doi.org/10.3109/0954898X.2012.722743 PMID: 22994683
6. B.E. Peterson, M.D. Healy, P.M. Nadkarni, P.L. Miller, G.M. Shepherd, ModelDB: An Environment for Running and Storing Computational Models and Their Results Applied to Neuroscience. Journal of the American Medical Informatics Association. 1996, 3(6), 389–398. https://doi.org/10.1136/jamia.1996.97084512 PMID: 8930855
7. P. Gleeson, M. Cantarelli, B. Marin, A. Quintana, M. Earnshaw, S. Sadeh et al., Open Source Brain: A Collaborative Resource for Visualizing, Analyzing, Simulating, and Developing Standardized Models of Neurons and Circuits. Neuron. 2019, 103(3), 395–411. e5. https://doi.org/10.1016/j.neuron.2019.05.019 PMID: 31201122
8. E. Nordlie, M.O. Gewaltig, H.E. Plesser, Towards Reproducible Descriptions of Neuronal Network Models. PLOS Computational Biology. 2009, 5(8), e1000456. https://doi.org/10.1371/journal.pcbi. 1000456 PMID: 19662159

Successful Applications for EuroHPC CoEs

Dirk Pleiter, PDC

The EuroHPC Joint Undertaking called for proposals for new centres of excellence (CoEs) earlier this year. After the procurement of three pre-exascale systems, as well as five petascale systems, EuroHPC is planning the acquisition of the first European exascale system as a next step. Significant efforts are needed to prepare applications that will be able to exploit these supercomputers. The new CoEs will be in charge of developing the necessary applications and workflows and organising the respective science and engineering research communities. Their work will start in January 2023 and continue for four years.

One of the successful applicants was the BioExcel CoE for Computational Biomolecular Research, which will enter a new phase, BioExcel-3, that will focus on the area of life sciences. It will support and maintain key European applications for the simulation of classical atomic models on a molecular level up to single cells. This third phase of BioExcel will be coordinated by PDC and involve eight partners from five European countries.

Two other CoEs will focus on the area of computational fluid dynamics (CFD). A new phase of Excellerat (known as Excellerat P2) will work on the enhancement and optimisation of various engineering applications that can be used for addressing different real-life use cases from the mobility and energy sectors. This CoE is coordinated by HLRS, a supercomputing centre in Germany. Altogether fifteen partners from seven countries will work together.

Finally, the completely new CoE for Exascale CFD (CEEC) will be the second CoE coordinated by PDC. CFD is one of the few application areas with use cases that can only be addressed with exascale supercomputers. Applications that have the potential of reaching exascale performance are part of this CoE. Eight universities and research organisations from five countries will work together on making this happen, guided by use cases with high industrial impact.

...continued from page 4

During his free time, Parasuram likes to explore the beautiful Swedish nature on his bicycle or go on hikes; he also tries to play instrumental music and occasionally experiments with cooking recipes. If you would like to connect with Parasuram, his LinkedIn page can be found at https://www.linkedin.com/in/parasuram-i-v-l-n.

Staff Focus

Johan Larsson

Johan Larsson started working at PDC over the summer and currently works part-time at PDC in first-line support while studying at KTH. Johan mainly assists with account creation and support requests, as well as reviewing the online documentation that is available to help people using the PDC systems. He has a bachelor's degree in engineering physics from KTH and is now pursuing a master's degree, also at KTH, in applied and computational mathematics, where he is specialising in computational mathematics. Johan first came in contact with the supercomputers at PDC during his earlier studies when he attended courses in the field of scientific computing. In his free time, Johan loves to go for a ride on his bike, go running or skiing, watch a tv series or eat something with his friends.

Programming Models and I/O with DEEP-SEA and IO-SEA

Stefano Markidis, CST, KTH

The High-Performance Computing (HPC) group at the KTH Royal Institute of Technology (KTH) does ground-breaking research on programming models and computer architectures to help build the road to exascale supercomputing. Late last year, the group (which is led by Stefano Markidis) successfully completed three high-profile exascale-related projects: EPiGRAM-HS, SAGE2, and VESTEC. All three of those projects were funded by the European Commission. This year the group started two new projects, called DEEP-SEA (Dynamical Exascale Entry Platform) and IO-SEA (Storage I/O and Data Management for Exascale Architectures), which are funded by EuroHPC and the Swedish Research Council (VR).

The three earlier projects paved the way for the work currently being carried out in the new projects. The EPiGRAM-HS project focused on extending and developing the Message Passing Interface (MPI) and GPI-2 programming systems on heterogeneous systems with graphics processing units (GPUs) and Field Programmable Gate Arrays (FPGAs). [GPI-2 is the second generation of the Global address space Programming Interface (GPI) model that was developed to help overcome limitations with using MPI for increasingly large systems.] SAGE2 investigated delivering HPC object storage for exascale systems that would provide efficient I/O on large-scale systems. VESTEC focused on interactive supercomputing enabling in-situ visualisation and data analytics for decision makers.

The two new projects continue European efforts towards developing and implementing a European software stack, including programming models and efficient I/O. DEEP-SEA, which is led by the Jülich Supercomputing Centre in Germany, will provide computation and communication libraries, resource management, and programming abstractions with run-time systems and tools. DEEP-SEA targets the Modular Supercomputing Architecture (MSA) where different components, like standard CPUs, and accelerators, such as GPUs, constitute a complex grid of heterogeneous technologies. IO-SEA, which is led by the French Alternative Energies and Atomic Energy Commission (CEA), will deliver an innovative data management and storage framework based on hierarchical storage management and on-demand provisioning of storage services. The KTH HPC group is contributing to these two projects with the development of programming systems (such as MPI and HDF5, which is the fifth version of the Hierarchical Data Format), DaCe (a parallel programming framework that maps code to high-performance CPU, GPU, and FPGA programs) and the demonstration of applications and use cases.

PDC Summer School 2022

Dirk Pleiter, PDC

For the first time since 2019, this year's PDC Summer School was an in-person event again. With almost fifty participants from fifteen different organisations and seven different countries, attendance has been both high and diverse.

After lectures on parallel algorithms and parallel computer architectures laid a foundation, the main focus of the school was to learn about parallel programming techniques like OpenMP, MPI, GPU programming using CUDA, and Python for high-performance computing (HPC). Furthermore, lectures and practical exercises on software engineering, performance analysis and optimisation were part of the curriculum. For the first time, attendees could use PDC's new flagship system Dardel for the lab sessions.

Niclas Jansson, who once upon a time attended an earlier PDC Summer School, gave a stimulating talk about the new computational fluid dynamics code, Neko, that he has been instrumental in developing, and also spoke about how the knowledge and skills obtained during the school can be used to design and implement highly scalable applications that also provide high performance.

PDC would like to thank Christof Kessler, Ivy Peng, Joachim Hein and Paolo Bientinesi for their lectures, as well as all the PDC staff who helped to realise the summer school by giving lectures, acting as mentors or assisting in other ways.

Above: Some of the PDC Summer School 2022 participants and lecturers at KTH, Stockholm, August 2022

Introduction to PDC Systems Workshop

Xin Li, PDC

The biannual course "Introduction to PDC Systems" serves to introduce PDC's high-performance computing (HPC) infrastructure to new users. In this course we start with an overview of the HPC infrastructure at PDC, and then cover the basic topics, such as getting an account, logging in, running jobs, storing data, and compiling code. In addition, we also cover practical topics including the Bash shell, EasyBuild, Singularity, SLURM script, MATLAB and the Python virtual environment.

The second introductory course for 2022 was held online on the 27th and 28th of October and was attended by eighteen participants. The course was split into lecture sessions and lab sessions; during the lecture session we walked through the slides (available at https://pdc-support.github.io/pdc-intro), while in the lab sessions the participants practiced common procedures and useful commands on the Dardel system. The participants also worked on several exercises to deepen their understanding of how to use the system efficiently.

We aim to continue the introductory courses in 2023, and will do our best to help new users become comfortable and productive working in the HPC environment at PDC. For information about the next introductory courses 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 web page https://www.pdc.kth.se/about/events or follow PDC on Facebook at https://www.facebook.com/kth.pdc.

SLURM settings for hybrid MPI/OpenMP code

- `--nodes` number of nodes
- `--ntasks-per-node` number of MPI processes
- `--cpus-per-task` 2 x number of OpenMP threads (because of SMT)
- `OMP_NUM_THREADS` number of OpenMP threads

Above: The slides for the course include lots of useful information about using the HPC resources at PDC.

Presenting Dardel for Businesses

Gert Svensson and Arash Alizad Banaei, PDC

Dardel is not just used for academic research. Various companies, such as Scania, Sweco, Faurecia Creo, Ishaq Fluid Mechanics, Tyréns, Polarium and Phoenix BioPower, collaborate with KTH and use PDC's resources for aspects of their research programmes. Earlier this year, Arash Alizad Banaei (who is an application expert at PDC) and Gert Svensson (Deputy Director of PDC) provided business users with an overview of the new Dardel system. Personnel from about ten companies attended the presentation, which was based on a similar event for Scania staff that was held late last year.

It was explained that Dardel will have two partitions: the first is primarily based on CPUs and the second, which has been installed late this year, features GPUs. The total estimated HPL performance of the new GPU partition with 56 AMD Instinct nodes, each node with 4 AMD MI250X GPUs, was mentioned as an impressive 8.2 PLFOPS. In addition, on the CPU nodes, the number of cores has been increased to 128. Although the total memory bandwidth per node has almost doubled from that of PDC's previous system, Beskow, the memory bandwidth per core is considerably less. It was highlighted that this can necessitate changes to codes with high memory bandwidth in order to utilise the new system efficiently.

The companies were also informed about other changes in the system environment, including

- the option to login using key pairs instead of Kerberos,
- what shared nodes are and how they work,
- how the new storage system is organised,
- that nodes are not using the previous AFS storage system anymore, and
- details about the new module system,

and participants were shown how to run applications on Dardel.

The presentation concluded with an overview of the computational fluid dynamics (CFD) and mechanical applications that are available on Dardel, since most of the businesses using Dardel are actively doing research in these areas. If your company is interested in using the PDC's high-performance facilities, including the second (accelerated) partition of Dardel, please contact business-unit@pdc.kth.se.

PDC-Related Events

PDC Summer School 2023

last two weeks in August, KTH main campus, Stockholm

Details will be available in the spring at https://www.pdc.kth.se/about/events.

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

- AQTIVATE
  https://aqtivate.ucy.ac.cy
- BioDT
  https://biodt.eu
- BioExcel CoE
  https://bioexcel.eu
- DICE
  https://www.dice-eosc.eu
- EBRAINS
  https://ebrains.eu
- EOSC-Nordic
  https://eosc-nordic.eu
- EUMaster4HPC
  https://eumaster4hpc.uni.lu
- EXCELLERAT
  https://www.excellerat.eu
- PerMedCoE
  https://permedcoe.eu

HPC news sources

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