PDC Center for High Performance Computing

PDC Newsletter 2025

Introducing the New KTH Center for Scientific Computing (KCSC)
Mimer: The Swedish AI Innovation Factory
Computational Methods for Magnetism
Introduction to PDC Systems
VeloxChem at WATOC 2025
BioExcel-3 Activities
ENCCS News
GROMACS Developments
CodeRefinery Enters Phase 4
CEEC Progress
NAISS System News
LUMI Intensive Course
APerOL Minisymposium
Digital Autonomy in Europe
KTH & AstraZeneca Partner to Revolutionise Drug Development with HPC & AI
PDC Summer School 2025
GANANA: A New Europe-India HPC Collaboration Has Been Launched!

Welcome to KCSC

Developing HPC solutions for modelling and simulations in your research

Patrick Norman, Director, KCSC

The PDC Newsletter is published by the PDC Center for High Performance Computing at the School of Electrical Engineering and Computer Science (EECS), KTH Royal Institute of Technology (KTH), Stockholm, Sweden.

PDC operates leading-edge, high-performance computers as easily-accessible national resources. These resources are primarily available for Swedish academic research and education. PDC is hosted by EECS, KTH, and is a member of the National Academic Infrastructure for Supercomputing in Sweden (NAISS).

Editor: Patrick Norman
E-mail: support@pdc.kth.se
ISSN: 1401-9671

Cover

KCSC will use the Teknikringen 14 building at KTH's main campus in Stockholm. The building currently hosts PDC and includes a variety of computer rooms. After renovations, the building will also house KCSC offices, common areas to stimulate spontaneous meetings, and rooms for student education.

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

PDC is delighted to share a particularly exciting announcement! On the 25th of November 2025, Anders Söderholm, the president of the KTH Royal Institute of Technology (KTH), established the KTH Center for Scientific Computing (KCSC). This new centre is dedicated to enabling cutting-edge research at KTH through high-performance computing (HPC). The centre has been under discussion for some time, and now, in what will be the final edition of the PDC Newsletter, it is my pleasure to officially present KCSC (see *Introducing the New KTH Center for Scientific Computing*). The new KCSC centre builds on PDC's long-standing strengths in HPC while placing a more intensive emphasis on software development. Importantly, PDC is incorporated as a core unit within KCSC, thus ensuring the continuity of PDC's long-standing HPC services while also opening up new opportunities for research collaborations.

The vision for KCSC is closely tied to the experiences of my quantum chemistry research group when I first became the director of PDC. At that time, my group relocated to the building where PDC is housed at Teknikringen 14. Initially, it was a purely practical decision to avoid splitting my time between two offices. In practice, however, the real benefit was the unique collaboration it fostered between my group and PDC's HPC experts. That proximity and shared expertise gave a tremendous boost to the development of the open-source quantum chemistry package VeloxChem, which was created by my group. (See *VeloxChem at WATOC 2025* and *KTH & AstraZeneca Partner to Revolutionise Drug Development with HPC & AI* for information about industry collaborations using VeloxChem.)

This model of cooperation has already proven its worth at KTH in other research areas. The development of the world-leading molecular dynamics package GROMACS and the computational fluid dynamics package Neko are both the result of close partnerships between researchers from those fields and HPC experts. GROMACS has become the most widely used molecular dynamics software globally, while Neko was a finalist for the prestigious Gordon Bell Prize in 2023. These examples demonstrate the extraordinary potential of bringing together complementary expertise. With the establishment of KCSC, I want to extend the same kind of opportunities to other KTH research groups. If your research involves advanced HPC or has demanding software needs, I warmly invite you to contact us at Teknikringen 14. We are making office space available within KCSC to foster exactly this type of collaboration between faculty researchers and HPC specialists – renovating the building to create more offices at KCSC for research groups is likely to be completed in 2027.

The launch of KCSC is happening at a pivotal time for Swedish HPC. In January 2023, the National Academic Infrastructure for Supercomputing in Sweden (NAISS) replaced the Swedish National Infrastructure for Supercomputing (SNIC), consolidating HPC hardware and funding into a single national framework supported by the Swedish Research Council. The transition has already been a great success: for the first time, Sweden has secured significant EuroHPC funding, which will support a medium-scale EuroHPC system named Arrhenius (see *NAISS System News*), and, in addition, Sweden has been awarded EuroHPC funding for a new AI innovation factory, Mimer (see *Mimer: The Swedish AI Innovation Factory*). Previously, Swedish universities contributed collectively to the financing of SNIC. Now KTH's contributions will instead be directed to KCSC, thus strengthening software development and adaptation at KTH for research that relies on HPC.

At PDC itself, a number of important changes are also taking place. On the hardware side, NAISS funding for Dardel will conclude at the end of 2026. While this means that PDC will no longer host NAISS flagship computers, we will continue to operate large-scale compute and storage systems for KTH and partners, as well as smaller-scale systems to promote software readiness for flagship software programs developed at KCSC. There have also been transitions among our staff. Gert Svensson, PDC's first employee and long-serving deputy director, has stepped down from that role and is now sharing his time between PDC and NAISS, where he is leading the procurement and deployment of the Arrhenius system. Rossen Apostolov, who contributed greatly to PDC's European projects over many years, has joined NAISS as the director of the Swedish AI Innovation Factory, Mimer.

We look forward to hearing from KTH colleagues interested in collaborating within KCSC and wish you all a successful semester!

Patrick Norman, Director KCSC

Above: The fun and inspirational VeloxChem booth at the 13th Triennial Congress of the World Association of Theoretical and Computational Chemists (WATOC 2025), Oslo, Norway, 22-27 June 2025

In This Issue

Editorial
Patrick Norman 2

Introducing the New KCSC
Patrick Norman 4

Staff Focus
Cathrine Bergh 5
Stefan Fleischmann 8
Andrey Alekseenko 11

Mimer: The Swedish AI Factory
Rossen Apostolov 6

Computational Methods for Magnetism
Anders Bergman et al. 8

Introduction to PDC Systems
Parasuram Indraganti 13

VeloxChem at WATOC 2025
VeloxChem Team 14

BioExcel-3 Activities
Alessandra Villa 15

ENCCS News
Apostolos Vasileiadis 16

GROMACS Developments
Alessandra Villa 17

CodeRefinery Enters Phase 4
Johan Hellsvik 18

CEEC Progress
Niclas Jansson 18

NAISS System News
Gert Svensson 19

LUMI Intensive Course
Jonathan Vincent 20

APerOL Minisymposium
Andrey Alekseenko 20

Digital Autonomy in Europe
Andrey Alekseenko 21

KTH & AstraZeneca Revolutionise Drug Development with HPC & AI
Mårten Ahlquist 22

PDC Summer School 2025
Alessandra Villa 22

GANANA
Cathrine Bergh 23

PDC-Related Events 24

HPC Sources 24

Introducing the New KTH Center for Scientific Computing (KCSC)

Patrick Norman, Director, KCSC

The KTH Royal Institute of Technology (KTH) is establishing a world-leading centre, KCSC, for scientific software development in high-performance computing (HPC) environments. The centre will meet the increasing needs for modelling and simulations in prominent areas of research at KTH (including the life and materials sciences and engineering) and also strive to make HPC resources easy to use and accessible to all researchers at KTH.

What is KCSC?

KTH is taking a forward-looking step to support researchers at KTH by establishing the "KTH Center for Scientific Computing", which will be known as KCSC. The centre will focus on the scientific software component of high-performance computing in Sweden. This will complement other HPC support for research provided through the National Infrastructure for Supercomputing in Sweden (NAISS), such as physical computing and data storage resources along with support from application and systems experts.

KCSC will be a physical centre on the main KTH campus in Stockholm at Teknikringen 14 where domain scientists will meet system experts and research software engineers who stay attuned to hardware developments and modern (parallel) programming practices. The building already features state-of-the-art computer room facilities, which serve the PDC Center for High Performance Computing (PDC), and will undergo a renovation to incorporate PDC's operations with the holistic needs of KCSC. The result shall be an open and welcoming environment for students, teachers, researchers, and technical staff members centred on the topic of developing novel and innovative mathematical and computer science methods enabling application impact in natural sciences and engineering. PDC will continue to operate as part of KCSC.

Why do KTH Researchers Need KCSC?

KTH is one of the foremost technical research universities in the world, coming in equal thirty-seventh place for Engineering in the Times Higher Education World Ranking of Universities 2025. KTH's excellence in science and engineering research has been facilitated for many decades by continually advancing first-class computational and data-processing resources – these have been provided by KTH in partnership with various national infrastructures and organisations that support Swedish research reliant on HPC.

Several research groups at KTH are involved in world-leading development of outstanding globally-used software for scientific research in domains such as computational fluid dynamics, quantum chemistry and molecular dynamics. And many more researchers at KTH perform first-class research using these software applications (and others) to stretch the boundaries of what can be studied in their research domains.

As the speed and design of computer systems continue to increase, the complexity of models that can be simulated, and the quantities of data that can be analysed and utilised for modelling, continue to grow. To remain competitive and to utilise the full power of available supercomputing systems, KTH researchers must be able to run scientific software that is flexible and easy to use.

The software also needs to be easily adapted to run efficiently on the even more powerful systems that will be used in the future. Obviously, from a cost perspective, it makes sense to run performance-optimised software to make the most of every computing clock cycle. However, experience shows that there is generally room for improving the degree to which research software is optimised for the systems on which it is run. To extend the realms of what can be modelled so that Swedish research remains competitive, researchers need access to optimised software in their domains. In actual practice though, there are not many general-purpose domain software programs in natural sciences and engineering fields that achieve international recognition and usage. KTH is, in fact, known for the development of no less than three such programs, which is an outstanding achievement based on very long-term efforts and dedication. These three applications will form the software foundation of KCSC. They cover the domains of molecular dynamics (GROMACS), quantum chemistry (VeloxChem), and computational fluid dynamics (Neko), all of which are pertinent to prominent research activities at KTH.

KCSC Facilitating Excellent Research into the Future

The KCSC initiative at KTH is indeed an exciting development that will drive excellence in research because the work at KCSC will be defined by scientifically driven target challenges, potentially multidisciplinary or transcending the schools of KTH. KCSC will be in complete control of the scientific software development and distribution cycle. So, rather than having to wait and see what types of modelling new software releases have to offer, we can steer the software development according to our scientific ideas, which more often than not come from the experimental side. This means the software can be developed to meet what is needed for the research, rather than research being adapted to what the software allows.

Furthermore, since KCSC staff will have detailed knowledge of both the system and software components of modelling, the centre is uniquely equipped to dramatically lower the access barrier to the Swedish and European HPC infrastructures and empower researchers without HPC expertise, such as experimentalists, with the means to interactively explore ideas. Two external project grants from the Swedish Energy Agency and Vinnova recently awarded to KTH researchers are testaments to the novel opportunities that KCSC will provide towards competitiveness for research funding.

If your research at KTH could benefit from a component of modelling and simulation, you are encouraged to contact KCSC! We are in the phase of organisational transformation and building renovation, but please contact any member of staff at PDC (see https://www.pdc.kth.se/about/contact/contacting-pdc-staff-members-1.760894), and we will organise a meeting to give you more information and discuss opportunities.

Staff Focus

Cathrine Bergh

Cathrine Bergh recently joined PDC as a postdoctoral researcher. She earned her Ph.D. in theoretical biophysics in 2023 from the KTH Royal Institute of Technology, after which she did a year of postdoctoral research, with Prof. Lindahl at KTH. Her doctoral research focused on mapping the free energy landscapes of ion channels using large-scale molecular dynamics simulations and Markov state modelling, finally utilising these models to make new predictions of ion channel function. She also contributed to GROMACS software development.

At PDC, Cathrine's responsibilities include project management and workflow development for AI-driven immunogenic peptide prediction in the GANANA project, teaching and tutorial development for BioExcel, and perhaps software development in GROMACS as well.

In her free time, Cathrine plays violin in a symphony orchestra and also enjoys playing classical piano, drawing, writing, travelling, and occasionally horseback riding when the opportunity arises.

Mimer: The Swedish AI Innovation Factory

Rossen Apostolov, Director, Mimer

In Scandinavian mythology, the Well of Mimer was a source of wisdom and understanding, where Odin sacrificed an eye in exchange for knowledge. Inspired by this tale, Mimer – The Swedish AI Innovation Factory (https://mimer-ai.eu) positions itself as a modern-day wellspring of insight, designed to transform Sweden and Europe into leaders in artificial intelligence (AI) innovation. Anchored in the framework of the European High Performance Computing Joint Undertaking (EuroHPC JU) and hosted by the National Academic Infrastructure for Supercomputing in Sweden (NAISS), Mimer is bringing universities, research institutes, and industry partners together to establish a powerful hub for AI-driven science, business, and societal development.

A Hub for Knowledge and Innovation

Mimer is built on a simple but bold vision: to make AI fundamental to European industry and society. The factory focuses on immediate opportunities where AI can act as a "cognitive companion" – complementing human expertise and enhancing existing processes. Its mission is to catalyse the adoption of AI across various sectors, reduce reliance on foreign cloud providers, and establish European sovereignty in critical digital infrastructure.

At its core, the Mimer AI Factory rests on three roots: academic research, industrial innovation, and enabling infrastructure. This mirrors the Norse world tree Yggdrasil, whose three deep roots sustain the growth of an entire ecosystem. Through this approach, Mimer will ensure that scientific excellence, industrial competitiveness, and societal needs advance together.

Addressing an Urgent Demand

Now is a crucial time to take action. The surge in demand for AI capabilities has outpaced the requirements for traditional supercomputing, with industry now outstripping academia in its need for large-scale data handling and model training. In particular, Swedish small to medium-sized enterprises (SMEs) struggle due to having limited access to AI infrastructure and expertise, which leaves them dependent on costly and non-European cloud services that, in turn, raise issues of compliance, sovereignty, and data security.

Mimer is addressing these challenges by establishing a dedicated AI supercomputer for Sweden. Unlike traditional computing systems, this resource will be optimised specifically for running AI workloads. It will be equipped with powerful GPUs, fast interconnects, large-scale storage, and seamless cloud-style access. It will also integrate with global model repositories allowing researchers and companies to tap into AI "on demand" while benefiting from sovereign European infrastructure.

Anchored in Sweden's Research and Innovation Landscape

Sweden is uniquely positioned to lead this effort. Over the past decade, the Knut and Alice Wallenberg Foundation has invested heavily in programs such as WASP (about AI, autonomous systems, and software), DDLS (on data-driven life sciences), and WISE (on materials science for sustainability). These initiatives, each on the scale of hundreds of millions of euros, have cultivated world-class research and industry collaborations. The government has also supported large-scale life sciences infrastructures like SciLifeLab, while Sweden's thriving startup ecosystem (which is home to companies like Spotify, Klarna, and Mojang) demonstrates a proven ability to scale global innovations. Lovable, a Swedish online AI-based website creator that allows users to build websites from scratch using just prompts and no code, broke the world record for the fastest-growing startup!

Mimer builds on this foundation by combining the national NAISS supercomputing infrastructure, hosted at Linköping University, with the industry-oriented expertise of the Research Institutes of Sweden (RISE). This partnership ensures that both academic excellence and SME innovation are served, with co-working hubs being established in Linköping and Stockholm and support extended across the country.

Strategic Objectives

Mimer's objectives are clear and ambitious.

- ***Catalyse growth and competitiveness:*** By lowering barriers to AI adoption, Mimer will empower SMEs, public agencies, and large enterprises alike to leverage AI for productivity and innovation.
- ***Transform scientific computing:*** Through its "AI4Science" agenda, Mimer will integrate AI with physics-based simulations, ensuring Europe remains at the forefront of global discovery.
- ***Train the AI workforce:*** Comprehensive education and hands-on programs will create a pipeline of skilled researchers, professionals, and entrepreneurs equipped for the AI age.
- ***Provide AI infrastructure that is ethical and sovereign:*** Making secure, EU-compliant data management accessible will ensure that sensitive information in sectors like healthcare or defence can be handled responsibly.

Target Sectors

Mimer prioritises four key domains where Sweden already holds competitive advantages.

1. ***Life Sciences*** – advancing personalised medicine, drug discovery, and health data analysis while safeguarding patient privacy
2. ***Material Sciences*** – accelerating discovery of sustainable materials by combining quantum chemistry, AI models, and industrial process design
3. ***Autonomous Systems*** – supporting robotics, drones, and self-driving vehicles, in collaboration with world-leading Swedish research and companies
4. ***Gaming Industry*** – harnessing AI to drive content creation, immersive experiences, and real-time simulations in one of Sweden's most dynamic export sectors

These focus areas provide a strong base for broader cross-sector applications, from climate modelling to public administration.

Phased Roadmap

The implementation of Mimer will follow a structured three-phase roadmap.

1. ***Establishment (Years 1–2):*** Build the AI supercomputer, recruit key staff, onboard early users, and launch training programs. Partnerships with community platforms and model repositories are solidified.
2. ***Expansion (Years 2–3):*** Co-working hubs open in Linköping and Stockholm, roadshows attract industrial users, and broad training delivery reaches academia, SMEs, and government agencies. Pilot programs with large enterprises test cost-recovery and inference services.
3. ***Full Operation (Years 3–5):*** Stable large-scale use across sectors, integration into public services, commercialisation of models, and long-term sustainability through blended funding. By this stage, Mimer is a central node in the European network of AI factories, contributing to pan-European standards and collaborations.

A European Node of Excellence

Beyond its national mission, Mimer will play a vital role in the broader EuroHPC ecosystem. It will be deeply networked with other European AI factories, enabling the sharing of knowledge and resources across borders and also contributing to the establishment of a "European AI grid". This approach will strengthen not only Sweden's competitiveness but Europe's digital sovereignty and leadership.

Conclusion

Mimer is more than just a supercomputer or a training programme. It is a national and European strategy in action — one that seeks to empower society with trusted, accessible, and cutting-edge AI. By combining the strengths of academia, industry, and government, Mimer will ensure that Sweden remains a world leader in AI-driven innovation while safeguarding European independence in a field that is rapidly becoming foundational for the future.

Like its mythological namesake, Mimer aims to be a wellspring of wisdom for generations to come – a source where knowledge, technology, and collaboration converge to shape a resilient and competitive Europe.

Staff Focus

Stefan Fleischmann

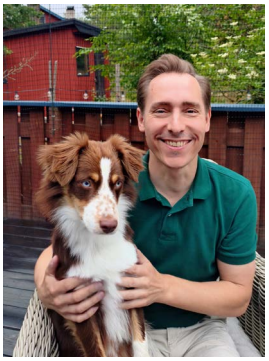

Stefan Fleischmann joined PDC in 2024 as an HPC systems manager. He earned his B.Sc. in physics from Paderborn University in Germany. He then continued his studies in theoretical physics at KTH, where he started working on improving the GROMACS continuous integration (CI) infrastructure. In 2014 he started working at SciLifeLab in Solna, where his responsibilities included, amongst other things, setting up and maintaining compute and storage systems for biophysics research groups, the National Bioinformatics Infrastructure Sweden (NBIS), and the Cryo-EM infrastructure unit that was established in 2016. Stefan also continued work on maintaining the GROMACS CI infrastructure.

Stefan is now working on a storage pilot project at PDC and collaborating with some of PDC's application experts to establish automated continuous benchmarks for PDC's flagship research software packages. He has also joined the PDC staff who contribute to NAISS and started working on systems for the upcoming Arrhenius supercomputer.

Computational Methods for Magnetism

Anders Bergman, Olle Eriksson, Maryna Pankratova, Manuel Pereiro & Nastaran Salehi, Uppsala University; Anna Delin, Johan Hellsvik, Mariia Mohylna & Qichen Xu, KTH; Danny Thonig, Örebro University

Introduction

Magnetic materials are everywhere. For example, there are several dozen permanent magnets in a modern car and around ten in a smartphone. Decades of research and development of magnetic materials have given us today's sophisticated electric motors, generators, speedometers, media for data storage, door locks, speakers, microphones, sensors, camera stabilisation, and wireless charging. How many magnets did you use today?

However, there is more. The electron spin, which is the quantum mechanical origin of magnetism, can also be used to transport and process information in the form of spin currents or spin waves. This quickly developing technological field is called spintronics and emerged from discoveries in the 1980s concerning spin-dependent electron transport phenomena in solid-state devices. The Nobel Prize-winning discovery of giant magnetoresistance (GMR), and the related tunnelling magnetoresistance (TMR), have revolutionised hard disk drive and sensor technology due to the ability to detect minute magnetic field variations.

Discovering new magnetic materials and spintronics phenomena is a very active field of research nowadays, and high-performance computing has become an indispensable tool to accelerate this research. In this article, we highlight recent and ongoing contributions from our research groups in developing methods and tools for data-driven e-infrastructure to predict and compare commonly available experimental data for magnetic materials. This involves novel methods inspired by genetic algorithms (that are typically found in biological systems) being used to identify magnetic configurations, a method to simulate magnetisation dynamics where spin-lattice effects are included, multiscale modules in simulations, and the connection to ab-initio electronic structure theory based on density functional theory (DFT), as well as exploration of new computational architectures. (Walter Kohn shared the Nobel prize for chemistry in 1998 for his work on DFT.)

Ab-initio Theory of Exchange and Atomistic Spin-Dynamics

Simulations of magnetism rely on having a good estimate of the exchange- and Dzyaloshinskii-Moriya (DM) interactions of the material in question, in addition to dipolar energy and the magneto crystalline anisotropy. The dipolar interaction is a classical one, for which analytical expressions are known, while all other effects have their origin in quantum mechanics. For simulations on an atomistic level, where the magnetism is represented on the atomic scale,

this implies that access to a microscopic spin-Hamiltonian (H) on atomic scale is needed. This normally involves the Heisenberg Hamiltonian for exchange, where the energy cost of rotating two magnetic moments is governed by a scalar number: the bilinear Heisenberg exchange parameter. In a more general form, this interaction is represented in matrix form, where the DM interaction is also present as an anisotropic, antisymmetric component. Due to seminal work in the middle of the 1980s, it is possible to directly calculate these interaction terms, via an expression that relies on the master equation of DFT, the Kohn-Sham equation. From the derivative of H with respect to the rotations of a magnetic moment, one can then calculate a local exchange-dominated field $\boldsymbol{B}_i$ as $\boldsymbol{B}_i = -\mathrm{d}H/\mathrm{d}\mathbf{m}_i$ that would provide a torque on the moment $\mathbf{m}_i$ if it were to deviate from the orientation of $\boldsymbol{B}_i$. Note that this is analogous to calculations of, for example, the force of an atom, if it were to deviate from its equilibrium position, via the expression $\boldsymbol{F}_i = -\mathrm{dE}/\mathrm{d}\mathbf{u}_i$, where E is the energy landscape generated by the positions, $\mathbf{u}_i$, of all atoms of the solid. This forms the basis for molecular dynamics simulations where, once the forces have been obtained, Newtonian mechanics is invoked to study the dynamics. In the field of magnetism, the situation is a bit more complex since $\mathbf{m}_i$ is an angular momentum, which causes the dynamics of any deviation between $\boldsymbol{B}_i$ and $\mathbf{m}_i$ to be expressed in terms of a torque. The resulting master equation of the dynamics is normally referred to as the Landau-Lifshitz or Landau-Lifshitz-Gilbert (LLG) equation and can be written as

$$\mathrm{d}\mathbf{m}_i/\mathrm{d}t = -\gamma\, \mathbf{m}_i \times \boldsymbol{B}_i - \alpha\, (\gamma/m_i)\, \mathbf{m}_i \times [\mathbf{m}_i \times \boldsymbol{B}_i]$$

where γ is the gyromagnetic ratio and α is the Gilbert damping parameter. The first term of the LLG equation describes a conservative precession around the local magnetic field $\boldsymbol{B}_i$, while the second term describes the dissipative motion of aligning the moment $\mathbf{m}_i$ with $\boldsymbol{B}_i$. In order to incorporate thermal effects into the equations of motion, a Langevin-based approach can be used where a temperature-dependent stochastic contribution $\boldsymbol{b}_i$ is added to the local field, keeping the remaining parts of the LLG equation intact. The equations of motion can then be integrated with standard numerical integrators, for example, from the Runge-Kutta family.

Multiscale Simulations of Magnetism

Multiscale modelling is a powerful approach in computational physics and materials science that enables the seamless integration of simulations across different length and time scales, and it plays a particularly vital role in the study of spin dynamics. The core idea of multiscale methods is to bridge atomistic and continuum descriptions, each tailored to a specific scale, within a unified computational framework. This is crucial in magnetic systems, where microscopic interactions at the atomic level, such as exchange coupling, anisotropy, and Dzyaloshinskii–Moriya interaction, determine macroscopic behaviours like domain wall motion or skyrmion dynamics. Thus, the atomistic spin dynamics region models individual atomic moments and their interactions using the LLG equation, allowing for the resolution of local fluctuations and thermal effects. Conversely, the micromagnetic model describes magnetisation as a continuous field, enabling the study of larger domains with a reduced computational cost. The interface between these regions is managed through a combination of padding atoms, finite-difference grid interpolation, and damping bands to prevent unphysical reflections of high-frequency spin waves. Coarse graining further smooths the transition by gradually adjusting the interaction range near the interface. The framework developed in this work supports simulations involving complex magnetic properties [1,2]. It reproduces key magnetic phenomena such as spin wave interference, domain wall pinning, and skyrmion robustness in the presence of defects. For instance, the double-slit experiment and the tetrahedral anisotropy cluster embedded in a hybrid domain revealed the interplay between wave-like and topological excitations. Analytical models validate the observed behaviours, underscoring the physical accuracy of the multiscale approach. The implementation is extensible, designed to operate efficiently on large systems, and is available

within the UppASD [7] simulation suite. Overall, this multiscale framework not only deepens our understanding of magnetism across scales but also lays the groundwork for future developments, including adaptive interfaces and finite element discretisation, which promise even greater flexibility and fidelity in simulating real-world magnetic materials.

Ultrafast magnetisation dynamics — the ability to change the magnetisation in a material following excitation by a femtosecond laser pulse — has opened up exciting technological possibilities like all-optical magnetic switching, where data bits can be written using light instead of magnetic fields and terahertz-speed memory devices. However, describing what actually occurs in the material in light-driven transient magnetisation dynamics has turned out to be very complex. The electrons, spins and lattice motions interact in highly nontrivial ways. Thermodynamically, one can describe the interactions by introducing a temperature for each degree of freedom or reservoir (electrons, spins, and lattice) together with an interaction pathway between each of the reservoirs (for example, the electron-phonon interaction), leading to the so-called three-temperature model. By realising that the heat distributions of each reservoir should be calculated self-consistently, we have developed a method that describes the process much better than the conventional three-temperature model. We call our new method the heat-conserving three-temperature model (HC3TM) [3].

From Biology to Magnetism

We have developed a novel optimisation technique inspired by genetic algorithms, called the Genetic Tunnelling Optimiser (GTO) [4]. This method efficiently locates global minimum energy states in complex magnetic systems by effectively navigating potential energy surfaces characterised by numerous local minima. Our approach achieves significant parallelism, allowing hundreds of independent Markov Chain Monte Carlo (MCMC) processes, each targeting local minima, to run concurrently. This parallel execution accelerates the main optimisation process. The figure below illustrates how the GTO method can effectively utilise the CPU nodes of the NAISS Dardel system at PDC for large-scale concurrent computations for two spin system sizes.

In order to discover long-lived, metastable spin textures, we utilise a neural-network-based approach. The method has successfully identified new skyrmionic metastable state spin configurations with high topological charges in Pd/Fe/Ir(111) [5]. By leveraging Dardel's graphical processing unit (GPU) nodes with AMD MI-250X GPUs, researchers can now efficiently explore increasingly complex potential energy landscapes and accelerate discoveries for non-trivial spin textures. The figure on the next page shows a subset of identified topological metastable spin textures in the Pd/Fe/Ir(111) system in the presence of an external magnetic field with a strength of 3.5 Tesla.

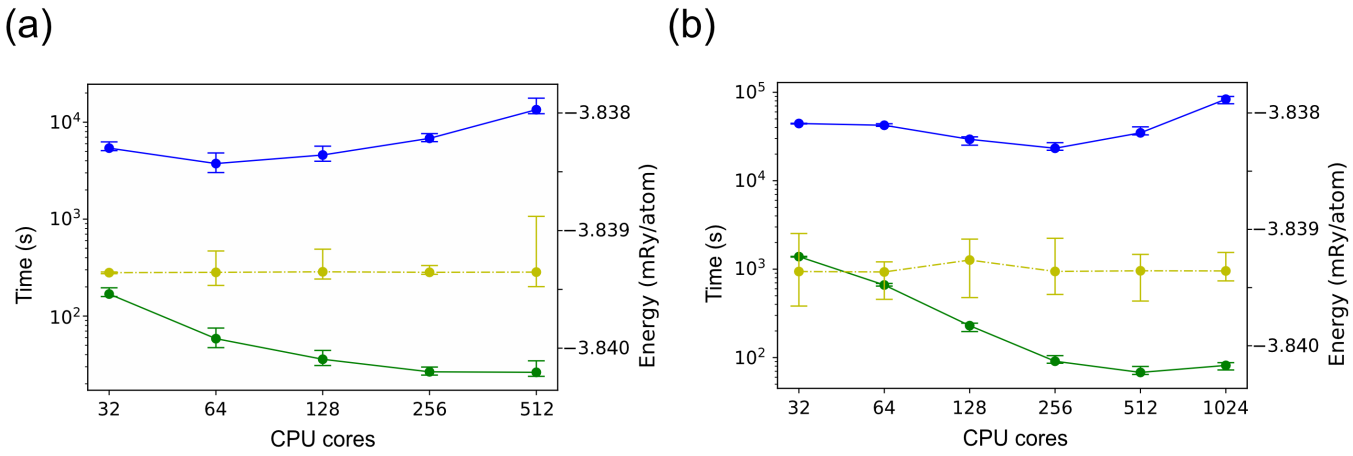

Above: Graph (a) shows the user execution time (green symbols), core execution time (blue symbols), and predicted ground-state energy (yellow symbols) for the first-generation GTO optimisation as a function of CPU core count for a 100×100 spin system. Graph (b) presents similar data for a 200×200 spin system.

Below: The images below show (a) a skyrmion, that is, a noncollinear spin texture with topological charge, (b) a skyrmionium and (c) an antiskyrmon, with the remaining panels (d)-(o) displaying magnetic textures with high-order topological charges varying between Q = −2 and Q = −13.

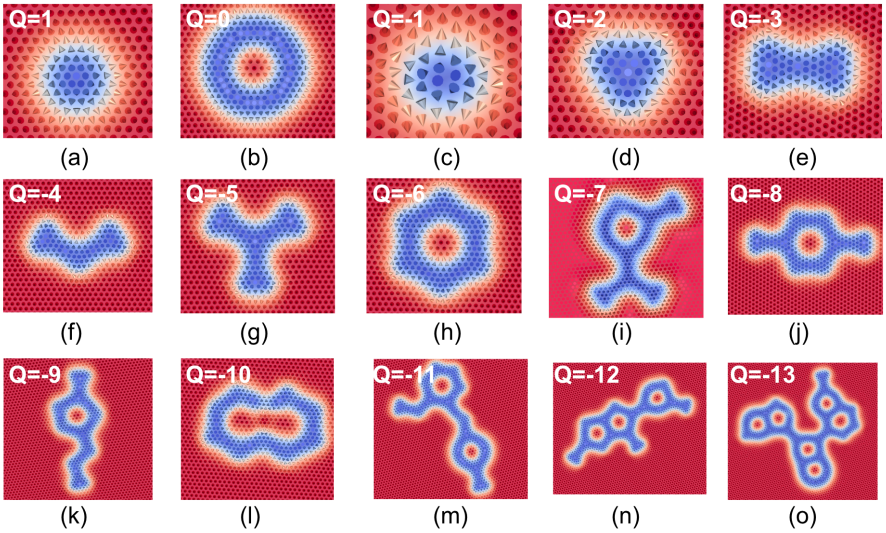

To address the significant volumes of data produced by multiscale magnetic simulations, we have developed SpinView, an interactive visualisation tool designed for filtering, analysing, and presenting vector field data. SpinView enables rapid generation of high-quality publication-ready images, videos, or interactive web pages [6]. It seamlessly integrates with various spin dynamics simulation tools and efficiently manages the visualisation of large-scale long-time simulation trajectories directly within cloud-based workflows.

Turning to the GPU Architecture

The previously mentioned long-lived spin textures can range in size from nanometres to even micrometres, involving thousands to millions of atoms. Understanding how these structures behave over time – their dynamics – is also crucial for potential technological applications. The dynamics is often studied using atomistic spin dynamics simulations. To model systems at these large scales more efficiently, we have adapted our computational methods to run on GPUs. We are working on the development of GPU code backends for modelling magnetisation dynamics on microscopic length scales within the formalism of atomistic spin-dynamics (ASD) [7,8] and spin-lattice dynamics (SLD) [9]. The ASD formalism is a well-established method for ab initio modelling of magnetic material physics, enabling numerical simulations of the properties of magnetic materials. It provides theoreticians and experimental physicists with access to a wide range of physical observables, both for systems in thermal equilibrium and for systems driven out of equilibrium by thermal or electromagnetic pulses or other forms of excitation. The SLD formalism extends ASD by incorporating ionic displacement into the set of state variables, allowing for simultaneous modelling of lattice and spin dynamics.

Staff Focus

Andrey Alekseenko

Andrey Alekseenko joined PDC at the beginning of 2025, although he has been a familiar face for several years due to his collaborations with PDC researchers while at SciLifeLab. Andrey earned his Ph.D. in mathematical modelling from the Moscow Institute of Physics and Technology in 2017. He completed his first postdoctoral position at the Applied BioComputation Group at Stony Brook University before joining KTH and SciLifeLab in 2020. At KTH, Andrey initially focused on adding a SYCL backend to GROMACS, a project that began with a focus on upcoming Intel GPUs. The scope of his work soon expanded to include using SYCL to run on AMD GPUs, which are utilised in both the Dardel system at PDC and the LUMI pre-exascale system at the CSC – IT Center for Science in Finland. His current work at PDC involves enhancing GPU support in GROMACS. Andrey is also a member of the SYCL and OpenCL Advisory Panels and is actively involved in advancing these open standards. In his leisure time, Andrey enjoys skiing and biking.

Below: The graph shows the dependence of the total simulation time on the number of simulated atoms (N) for Intel Core i9-9800X@3.6GHz and Intel Core i9-11900@2.5GHz for CPUs (yellow and blue squares, respectively) and NVIDIA GeForce RTX 2080 Ti, NVIDIA GeForce RTX 3080, and NVIDIA A4000 for GPUs (purple, red, and green circles).

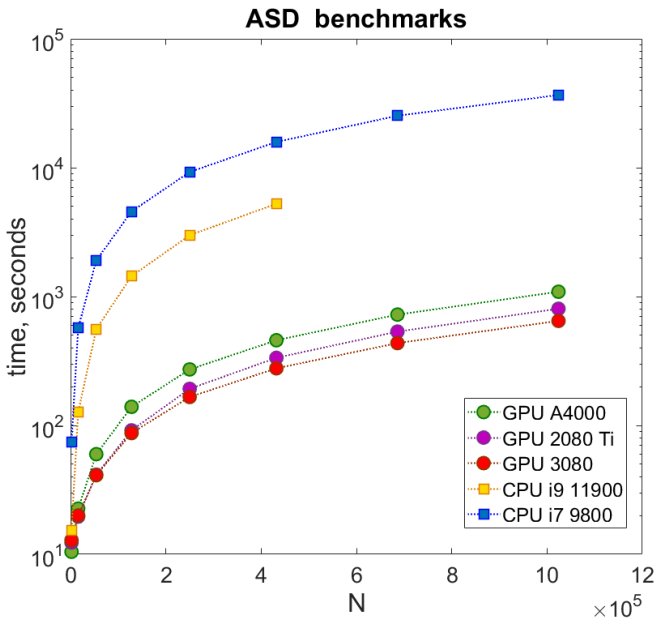

In the preliminary implementation, kernels for magnetic field calculation for bilinear spin-spin pair interaction and one numerical integrator have been implemented with the CUDA programming model for use on single NVIDIA GPUs. Benchmarks for the simulation of bulk body-centred cubic (bcc) Fe firmly establish the proposition of high speed-up for magnetisation dynamics simulation on GPU hardware. The results of these simulations are presented in the figure above for several widely available hardware devices.

While the performance of the code at a relatively small number of spins, N, is comparable across all the hardware, a steep increase in execution time is observed for both CPUs, with the i9-11900 significantly outperforming the i9-9800X. Despite both CPUs having relatively high clock speeds, they fall behind all the GPUs, even at moderate N values, due to the massive parallelism of the ASD algorithm. The difference in performance reaches a 60-fold speed-up between the slowest and fastest devices for N of the order of a million. In contrast, the difference in execution time between the GPUs is relatively less significant, with the RTX 3080 outperforming its competitors, likely due to its greater number of CUDA cores. The superior performance of the GPUs highlights the compatibility of spin-related problems with GPU architectures and points to the future direction for the increasingly computationally demanding field of magnetic simulations.

In order to be able to take on modelling on the same footing as magnetic and phononic degrees of freedom in solids, we have developed a method for coupled spin-lattice dynamics simulation [9]. The method is based on a coupling of atomistic spin dynamics and molecular dynamics simulations, expressed through a spin-lattice Hamiltonian, where the bilinear magnetic term is expanded up to second order in displacement. The motion of the atomic spins is described by the LLG equation and the ionic dynamics follow Newtonian dynamics.

A particular challenge for spin-lattice dynamics simulations is that the range of intrinsic and extrinsic system frequencies spans many orders of magnitude, with the consequence that the computational effort can become an order of magnitude larger than for a pure lattice or pure spin dynamics simulation of a system. Moreover, given that regular periodic boundary conditions violate the conservation of mechanical angular momentum, capabilities for the simulation of larger systems with open boundaries become of importance. Hence, we have identified the need for capability enhancement by means of implementing a GPU backend for SLD.

Conclusion and Outlook

Magnetic materials are essential for future sustainable technological development and are widely used in every aspect of our everyday lives, ranging from magnetic memories to wind power plants. To accurately describe experimental observations and predict magnetic properties for various applications, one needs to perform extensive simulations across different time and length scales. Many magnetic phenomena (such as ultrafast demagnetisation) need to take lattice dynamics into account in addition to the study of magnetic systems, or it may be necessary to consider the dynamics of millions of atoms to predict the behaviour of non-trivial magnetic

textures, which are an integral part of prospective spintronic devices.

Multiple powerful approaches have been recently proposed in our groups to describe various magnetic phenomena, such as coupled atomistic spin-dynamics simulations, neural network-based algorithms for searching for metastable spin textures, and a framework for multiscale modelling to bridge atomistic and continuum descriptions. These developments facilitate a deeper understanding of various magnetic effects and result in good agreement with experimental observations, such as the robustness of skyrmions in the presence of defects, timescales of ultrafast magnetisation dynamics of ferro- and antiferromagnets, and many others. To improve computational efficiency and data treatment for those resource-hungry simulations, we have recently proposed optimisation techniques enabling faster and more efficient simulations leveraging existing GPU nodes. The development of the SpinView software expedites the fast treatment of a large amount of data from magnetic simulations. Proposed methods and optimisation techniques provide a strong basis for further progress in magnetism and materials science research.

References

1. N. Salehi et al., "Micromagnetic-atomistic hybrid modeling of defect-induced magnetization dynamics", preprint arXiv:2504.14019 (2025).
2. E. Méndez et al., "A multiscale approach for magnetization dynamics: Unraveling exotic magnetic states of matter, Phys. Rev. Research 2, 013092 (2020).
3. M. Pankratova et al. "Heat-conserving three-temperature model for ultrafast demagnetization in nickel", Phys. Rev. B 106, 174407 (2022).
4. Q. Xu et al., "Genetic-tunneling driven energy optimizer for spin systems", Comm. Phys. 6, 239 (2023).
5. Q. Xu et al., "Metaheuristic conditional neural network for harvesting skyrmionic metastable states", Phys. Rev. Research 5, 043199 (2023).
6. Q. Xu et al., "SpinView: General Interactive Visual Analysis Tool for Multiscale Computational Magnetism", preprint arXiv:2309.17367 (2023).
7. UppASD code https://github.com/UppASD/UppASD
8. O. Eriksson et al., "Atomistic spin-dynamics: foundations and applications,", Oxford University Press, Oxford, UK, (2017).
9. J. Hellsvik et al., "General method for atomistic spin-lattice dynamics with first-principles accuracy", Phys. Rev. B 99, 104302 (2019).

Introduction to PDC Systems

Parasuram Indraganti, PDC

The "Introduction to PDC Systems" course is typically held twice a year to help researchers who have some experience using high-performance computing (HPC) systems, as well as researchers who are new to using HPC, to start utilising the Dardel system at PDC to boost their research.

In 2025, the course was run in March and early October. The courses were attended by researchers ranging from master's and doctoral students to senior researchers from institutions such as the Karolinska Institute, SciLifeLab and KTH. The courses introduced participants to a host of topics covering the architecture of HPC systems (particularly Dardel), logging in to Dardel and the structure of the file systems. The participants also learned to make good use of the Linux shell and software environment on Dardel, plus how to schedule jobs, including best practices for submitting jobs and compiling code on Dardel. Overall, the courses were great experiences that smoothed the path for the researchers in the process of beginning to use Dardel. The introductory course will be offered again in spring 2026, possibly in another format aligned with changes being made to HPC training approaches by the National Academic Infrastructure for Supercomputing in Sweden (NAISS). So, if you would like to accelerate your research using Dardel, keep an eye on https://www.pdc.kth.se/about/events for details of upcoming courses. You can also sign up for the PDC Announce mailing list to be sent details (see https://www.pdc.kth.se/about/contact/joining-pdc-mailing-lists-1.736925).

VeloxChem at WATOC 2025

VeloxChem Team, PDC & KTH

At this year's World Congress of Theoretical Chemists (WATOC) 2025, VeloxChem was proud to sponsor the event and showcase the latest developments. The VeloxChem booth featured live demonstrations where visitors could try out VeloxChem using Jupyter Notebooks, focusing on how the Python interface streamlines the preparation and execution of computational chemistry workflows.

During these sessions, we highlighted recent advances in VeloxChem, particularly its new GPU support, which enables users to handle much larger molecular systems than previously. We discussed real-world examples, including recent research on G-quadruplex DNA, which demonstrates how VeloxChem performs on modern high-performance computing (HPC) and GPU infrastructures [1].

We also introduced participants to both the comprehensive documentation on the VeloxChem website [2] and a collection of Jupyter Notebooks that are available in a GitHub repository[3]. These resources are designed to help users get started quickly and make the most of the software's capabilities. As an extra incentive, we invited attendees to try installing VeloxChem themselves – those who did so at our booth received a VeloxChem t-shirt as a thank you for engaging with our open-source community.

In addition, visitors also tried out recent features in VIAMD [4], such as the ability to read VeloxChem output files, plot molecular orbitals, and interactively analyse spectra. These improvements make it easier for users to visualize and interpret their results directly within the software, which supports both research and teaching applications.

Many conference participants were impressed by the performance of our tools on large systems and were pleasantly surprised to learn that both VeloxChem and VIAMD are open source and free. Several researchers mentioned that having access to such powerful, freely available software significantly lowers barriers for newcomers entering the field. Others appreciated the intuitive workflows and the way the tools integrate visualization and analysis, making complex data more accessible.

Below: VeloxChem booth at the 13th Triennial Congress of the World Association of Theoretical and Computational Chemists (WATOC 2025), Oslo, Norway, 22-27 June 2025

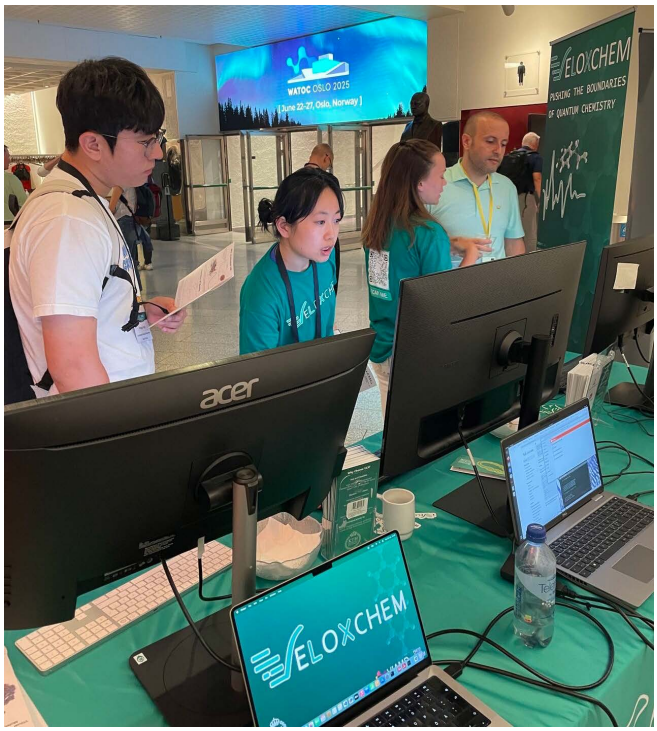

Overall, the event was a valuable opportunity to connect with the broader computational chemistry community, receive constructive feedback, and discuss future directions for both VeloxChem and VIAMD. The positive reception and interest from attendees have motivated us to continue developing these tools in collaboration with users worldwide.

References

1. X. Li, M. Linares, P. Norman, "VeloxChem: GPU-accelerated Fock matrix construction enabling complex polarization propagator simulations of circular dichroism spectra of G-quadruplexes". J. Phys. Chem. A (2025) 129(2), 633–642. https://doi.org/10.1021/acs.jpca.4c07510
2. https://veloxchem.org
3. https://github.com/VeloxChem/vlx-notebook
4. https://github.com/scanberg/viamd

BioExcel-3 Activities

Alessandra Villa, PDC

BioExcel-3 is running a lot of initiatives to support the biomolecular simulation community.

- The **webinars** (https://bioexcel.eu/category/webinar) feature notable developments in the field of computational biomolecular research. Have a look at the playlist here: https://www.youtube.com/playlist?list=PLzLqYW5ci-2fanl2RtYHyFvjftu_M4Jwy.
- The **forums** support the user communities of the four core applications: GROMACS, HADDOCK, BioBB, and PMX (see https://ask.bioexcel.eu and https://gromacs.bioexcel.eu).
- The **competency hub** contains competency information, career profiles and training resources that you can browse to get help to advance your career (see https://competency.ebi.ac.uk/framework/bioexcel/3.0).
- The **ambassador program** (https://bioexcel.eu/about/governance/ambassador-program) provides specialist training tailored to the needs of researchers in specific areas.

BioExcel Ambassador Program

PDC leads the BioExcel Ambassador Program. One of the core activities within the program is organising tailored, user-defined workshops. In recent months, BioExcel has been streamlining the process of identifying researchers' needs and organising training that covers those requirements. Several key steps have been established.

1. Ambassadors from geographically close and/or neighbouring countries work together to choose a suitable host country for training in that region.
2. The ambassadors undertake outreach activities within their own research communities to understand the training needs of those local researchers. To facilitate this, BioExcel has developed a training portfolio and survey that the ambassadors can distribute to their local research communities. Then the workshop topic will be selected to target the most pressing needs of the researchers in that region.
3. EuroHPC National Competence Centres in the relevant countries are involved in supporting the organisation of the workshop and providing access to suitable computing infrastructure.
4. Overall, the workshop organisation is driven by the local research institutions and EuroHPC National Competence Centres, while BioExcel mainly contributes by providing trainers and training materials, as well as sharing know-how concerning the organisational aspects of running workshops.

The first workshop in the BioExcel Ambassador Program was known as the "Carpathian edition" and was held in Bratislava, Slovakia, in October 2023. It was a great success. The second workshop (or "Atlantic edition") was held in Coimbra, Portugal, in November 2024 and was also very popular. In May this year, the third and very valuable "Balkan edition" workshop was held in Sofia, Bulgaria. The next one (the "Adriatic edition") will be held in Ljubljana, Slovenia, in February 2026. For further details and to register to attend the workshop, see: https://bioexcel.eu/events/bioexcel-eurocc-workshop-in-slovenia.

ENCCS
EuroCC National Competence Centre Sweden

ENCCS News

Apostolos Vasileiadis, ENCCS

The spring was a busy time for the EuroCC National Competence Centre Sweden (ENCCS). We continue to provide a platform for industry, academia and the public sector to access high-performance computing (HPC) infrastructure to satisfy their research and development needs. At the same time, training is still a staple in our activities, and ENCCS organised several teaching events during the first half of the year. In addition, two HPC experts joined the ENCCS team in the spring, and another will join us after the summer. Meet them (and the rest of ENCCS) on our team page (https://enccs.se/people)!

Industry Support

ENCCS continues to support companies in satisfying their compute needs using EuroHPC systems. We have recently helped a KTH-incubated startup Mappi that develops personalised learning paths in mathematics (https://enccs.se/success-story/2025/05/mappi-ai-math-education); they fine-tuned some open source AI models to create graphics explaining complex mathematical concepts. Another relevant use case was tackled by the AI consultancy company 42 labs (https://enccs.se/success-story/2025/01/42labs-accesses-marenostrum-5-to-develop-a-swedish-llm). They used the MareNostrum 5 system at the Barcelona Supercomputer Centre to fine-tune a Swedish-speaking large language model (LLM).

Access to EuroHPC Systems

Researchers in the public sector, industry and academia can continue to access EuroHPC systems through various calls (https://eurohpc-ju.europa.eu/supercomputers/supercomputers-access-calls_en) with different allocation sizes. Participating in the smaller Development Access and Benchmark Access calls can pave the way to utilising the Regular Access and Extreme Scale Access modes, with substantial compute resources being provided. The new AI for Science and Collaborative EU Projects Access mode (https://eurohpc-ju.europa.eu/eurohpc-ju-call-proposals-ai-science-and-collaborative-eu-projects_en) in particular provides large amounts of GPU compute resources for cutting-edge research in foundational models and the ethical use of AI.

Training Events

The autumn featured an inspiring line-up of training events organised by ENCCS and its European partners, including the yearly iteration of the Quantum Autumn School that was held over five days in November. If you are interested in attending future training events, keep an eye on the ENCCS events page (https://enccs.se/events).

Note that a new HPC in Europe portal (https://hpc-portal.eu) serves as a gateway to help you keep up with the European network of competence centres and centres of excellence.

Mimer, the Swedish AI Factory

In parallel with national competence centres, the EuroHPC Joint Undertaking has awarded funding for seven new AI-oriented systems in Europe. Sweden was awarded one of these. The system will be known as Mimer (https://mimer-ai.eu) and will be co-hosted by Linköping University and the Research Institutes of Sweden (RISE). This project is happening in parallel with ENCCS and, in addition to providing the hardware, expertise and training will be offered in connection to it. (For more information about what the project will encompass, see *"Mimer: The Swedish AI Innovation Factory"*.)

Stay tuned with these exciting developments through the ENCCS website (https://enccs.se) and LinkedIn page (https://www.linkedin.com/company/enccs)!

FAST. FLEXIBLE. FREE.
GROMACS

GROMACS Developments

Alessandra Villa, Andrey Alekseenko & Szilárd Páll, PDC

The GROMACS 2025 series has been released this year: version 2025.0 was released on the 11th of February 2025, 2025.1 in March, and 2025.2 in May (see https://manual.gromacs.org/2025.2/release-notes/2025/major/highlights.html).

GROMACS 2025 is mainly characterised by the inclusion of a (limited) interface to the PLUMED package (https://www.plumed.org), more support for hydrodynamics simulations and for force-field forms including neural-network potentials, plus improvements in performance.

A PLUMED interface using the MDModules interface is being shipped with GROMACS. This functions in a similar way to the Colvars interface that was implemented for GROMACS in 2024. Now it is only necessary to link GROMACS with PLUMED; special code hacking is no longer required! This guarantees that PLUMED sees an up-to-date state with integration on the GPU. Note that the PLUMED interface is currently limited to certain features at present, whereas the entire Colvars package is included in GROMACS. Also neural network potentials are implemented via an interface, in this case with PyTorch. More specifically, PyTorch is used to train the model, then the trained model is exported and the coordinates are fed into GROMACS via TorchScript and the PyTorch C++ API. Currently, neural network potential/molecular mechanics (NNP/MM) simulations using mechanical embedding schemes are also possible.

GROMACS' recent performance improvements include automatically turning on more GPU features, an initial version of a HIP backend that targets AMD GPUs and more OpenMP parallelisation, plus GPU-direct communication is now used by default when the MPI library supports it. In addition, enhanced particle-particle (PP) halo exchange using GPU kernel-initiated communication has been implemented using NVSHMEM, thereby improving performance when scaling to multiple NVIDIA GPUs.

Moreover, GROMACS 2025 has the option to have different CMAP types for each amino acid, which will make it possible to use the most current Amber protein force field (that is, ff19SB) soon. Also, there is support for volume scaling of centres of mass of more than one position restraint group, which has enabled shear with walls and constant pressure in 2025.

Now the developers are working to bring new features into the 2026 release. You can follow the development of GROMACS by looking at the GROMACS GitLab repository (https://gitlab.com/gromacs) or at the developer discussions section of the GROMACS discussion forum (http://forums.gromacs.org). Finally, GROMACS developer calls are held every second Wednesday (alternating between 11:00 and 17:00 CET/CEST) and everybody is welcome to join these calls. Information about how to join the calls is posted on the forums.

As always, if you need some help with your simulations, you are welcome to sign up in the user discussion section of the GROMACS user forum (which can be found at http://forums.gromacs.org) and ask your questions there. You can also keep yourself updated by following the latest GROMACS news on https://www.gromacs.org.

To find out more about GROMACS 2025, you can watch the "What's new in GROMACS 2025" webinar (see https://bioexcel.eu/webinar-whats-new-in-gromacs-2025-2025-2-18).

If you are interested in attending the next "Learn to Code in GROMACS" workshop, keep an eye on https://bioexcel.eu/events for announcements.

CodeRefinery Enters Phase 4

Johan Hellsvik, PDC, Diana Iușan, Uppsala University and Samantha Wittke, CSC – IT Center for Science

CodeRefinery is a project and collaboration network to help people improve their scientific software development skills. The training aspect of CodeRefinery focuses on teaching how to use state-of-the-art tools and practices for modern collaborative scientific software development. CodeRefinery's training activities are provided in the form of large online workshops as well as in-person events.

Initially launched in 2016, the CodeRefinery project has now entered phase 4. PDC has been a partner in the project since its inception. At the end of May there was a two-day kick-off for phase 4 in Gothenburg. CodeRefinery is now a consortium of nine partner organisations across the Nordic countries with the project owner being the Nordic e-Infrastructure Collaboration (NeIC). The partners by country are as follows.

- Finland: CSC - IT Center for Science and Aalto University
- Norway: Norwegian University of Science and Technology (NTNU), Sigma2/Norwegian Research Infrastructure Services (NRIS) and the University of Oslo
- Sweden: EuroCC National Competence Center Sweden (ENCCS) and the National Academic Infrastructure for Supercomputing in Sweden (NAISS)
- Denmark: Danish e-Infrastructure Consortium (DeiC)/Center for Humanities Computing (CHC) at Århus University
- Iceland: University of Iceland

Anyone who is interested in attending any of the CodeRefinery events can find the latest news on the project's website (https://coderefinery.org) or sign up for the newsletter (https://coderefinery.org/about/newsletter). CodeRefinery also invites everyone who is interested in joining the community to participate in our Zulip chat (https://coderefinery.zulipchat.com). This chat is shared with the Nordic Research Software Engineering community (https://nordic-rse.org).

CEEC Progress

Niclas Jansson, PDC

For centuries, scientific discovery has relied on experiments to understand new theoretical models. However, the advent of powerful supercomputers and efficient numerical methods has made it possible to perform simulations instead. Computational Fluid Dynamics (CFD) provides a powerful tool for qualitative and quantitative analysis of fluid flow by solving a set of equations on a computer. This enables scientists and engineers to perform experiments in a virtual laboratory and study a variety of fluid flows encountered in everyday life. That is why CFD is used today in various disciplines, from basic scientific research to engineering applications. However, unlike many other scientific fields, there is virtually no upper limit to the problem size in fluid dynamics, thus making detailed simulations exceedingly expensive to perform. The main goal of the Centre of Excellence in Exascale CFD (CEEC) is to address extreme-scale computing challenges to make it possible to use large-scale, accurate and cost-efficient high-fidelity CFD simulations for both academic and industrial applications.

At the end of 2024, CEEC reached the halfway mark of the project. During the first two years of CEEC, we completed the initial requirements phase, defining lighthouse use cases, which are a central part of CEEC around which the rest of the centre of excellence is organised. They provide the necessary requirements for our algorithmic work packages. We have also been starting to see what the end of the project's development phase might look like, where CEEC has made significant progress in ensuring exascale readiness of the codes supporting the lighthouse cases. All of the core codes have now been ported and tuned for accelerators and are running on all of the currently installed EuroHPC Joint Undertaking (EuroHPC JU) pre-exascale machines; some are even participating in the exascale pilot for the first European exascale machine, JUPITER.

A large part of the current development phase of CEEC is dedicated to improving algorithms for efficient exploitation of exascale architectures and also to enhancing the techniques and technologies required for our lighthouse cases. At the halfway mark, these algorithmic developments are being integrated into our core codes to support the execution of the lighthouse codes. For techniques and technologies, we see a similar trend: lighthouse cases are being coupled with machine learning interfaces, workflow environments, and enhanced visualisation and data management tools.

CEEC is now working on the second half of the project goals. We are entering the demonstration phase and have been working on finalising the integration and implementation of our work with the lighthouse cases. We will be moving forward to demonstrating the capabilities on available pre-exascale and exascale systems.

The work done by CEEC so far has been highly successful. In particular, some of the early developments in CEEC formed a key part of work that reached the final of the Gordon Bell Prize in 2023. This prize is awarded each year to recognise outstanding achievement in high-performance computing (HPC), such as innovation in applying HPC to applications in important engineering science, and large-scale data analytics problems.

NAISS System News

Gert Svensson, NAISS & PDC

Arrhenius

The procurement process for Arrhenius, the new EuroHPC system for Sweden, was completed early this summer. It will be an HPE system quite similar to the Dardel system at PDC and will have the same NVIDIA Grace Hopper GPU nodes as in the Dardel expansion (which is described below) plus AMD CPUs that are more recent and powerful than those in Dardel.

Arrhenius will have a CPU partition with 424 AMD Turin 128-core CPUs and a GPU partition with 382 Grace Hopper nodes. The GPU nodes will each have four NVIDIA GH200 Grace Hopper Superchips, and each of those superchips will have a Hopper GPU that is tightly connected with a Grace ARM CPU. The GPU partition is expected to have a High-Performance Linpack (HPL) benchmark performance of more than 60 PFLOPS. Storage will be provided by a 29PB fast parallel Lustre file system. Arrhenius will also have a partition for cloud computing and a partition for sensitive data, each with its own storage.

The system is currently being installed in Linköping and will be completed in the spring of 2026. Arrhenius is planned to replace existing NAISS systems, such as Tetralith (in mid-2026) and Dardel (at the end of 2026).

Dardel

In the previous edition of the PDC Newsletter, we announced that the Dardel system would get more disk capacity plus some more nodes with NVIDIA GPUs. Many researchers requested NVIDIA GPU nodes so they could use applications that are only available for the NVIDIA software stack with CUDA. These upgrades have now been installed, and the new storage and nodes are operational. The Lustre file system has been expanded with an additional 4.7 PB of hard disk storage and 260 TB flash storage. This will give a total usable Lustre capacity of around 22 PB.

Eight Grace Hopper nodes have been added to the system, resulting in a total of 32 superchips. One of the nodes is reserved as a login node.

LUMI Intensive Course

Jonathan Vincent, PDC

Between the 3rd and 7th of March this year, PDC hosted an intensive course on using the CPU and GPU partitions of the LUMI system (LUMI-C and LUMI-G). LUMI is a EuroHPC Joint Undertaking pre-exascale supercomputer system that is based in Kajaani, Finland. Swedish researchers can access the LUMI system using time allocations that are decided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS). (For information about applying for time via NAISS, see https://www.naiss.se/resource/lumi.)

The course was run by the LUMI HPE Centre of Excellence (HPE CoE), AMD and the LUMI User Support Team (LUST). It was designed for researchers who were either working on ongoing projects on LUMI or who had submitted a project proposal for access to LUMI. The course was split into two parts, with the first two days covering an initial introduction to LUMI, and the final three days going into more detailed and advanced usage of the system. Researchers could attend either or both parts of the course, which was offered in person at the KTH main campus in Stockholm and also online, with a peak of seven onsite users and forty-five online users attending.

If you are interested in attending training events on using LUMI in the future, keep an eye on https://www.lumi-supercomputer.eu/events.

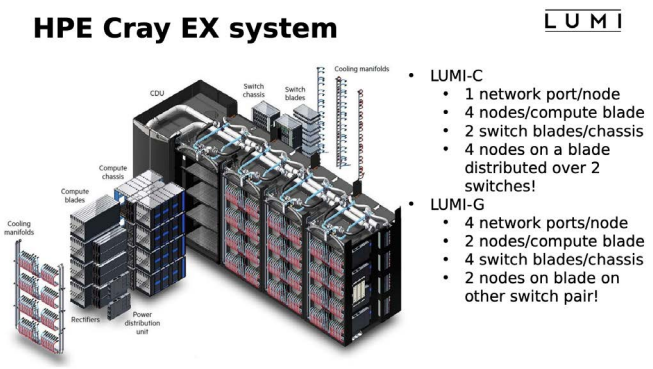

Above: Slide from LUMI Intensive Course

APerOL Minisymposium

Andrey Alekseenko, PDC

The third "Application Perspective on SYCL" (APerOL) minisymposium took place on the 18th of June at the Platform for Advanced Scientific Computing 2025 (PASC25) conference in Brugg, Switzerland. The minisymposium was organised by Andrey Alekseenko and Szilárd Páll from PDC with speakers representing a diverse range of institutions and disciplines: Peter Thoman (University of Innsbruck), Markus Büttner (University of Bayreuth), Kevin Harms (Argonne National Laboratory), and Timothée David-Cléris (Grenoble Alpes University).

Aligning with the conference's theme of how supercomputing can be utilised to help address the seventeen UN Sustainable Development Goals (https://sdgs.un.org/goals), the minisymposium demonstrated SYCL's role in addressing three critical sustainability challenges: accessibility through broader hardware compatibility, energy efficiency by enabling optimal hardware selection, and long-term sustainability through reduced code maintenance and enhanced reusability across increasingly diverse hardware landscapes.

Application developers shared compelling success stories in areas including computational

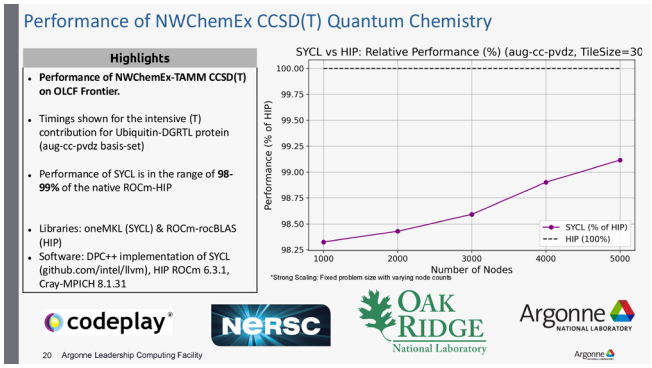

Above: Slide from the talk "Experiences with using SYCL for NWChemEx on Exascale Supercomputers" by Abhishek Bagusetty and Kevin Harms (https://doi.org/doi:10.5281/zenodo.15748782) at APerOL, 18 June 2025

chemistry, astrophysics, and coastal modelling. Highlights included near-exascale performance of the Shamrock Smoothed Particle Hydrodynamics (SPH) code on NVIDIA, AMD, and Intel GPUs; coastal simulation framework portability from CPUs to Field Programmable Gate Array (FPGAs); challenges and successes maintaining the NWChemEx computational chemistry code on top US Department of Energy (DOE) supercomputers; and productivity improvements offered by the Celerity distributed programming framework.

The event fostered a collaborative atmosphere with lively discussions between researchers and the audience. This spirit of collaboration and idea exchange demonstrated growing momentum behind open standards in scientific computing, positioning portable programming models as essential tools for addressing global challenges through computational science.

The presentations from the event are available, along with the presentations from the previous two APerOL minisymposia, at https://kth-pdc.github.io/aperol.

Digital Autonomy in Europe

Andrey Alekseenko, PDC

The newly launched Digital Autonomy with RISC-V in Europe Specific Grant Agreement 1 project (DARE SGA 1) aims to increase European autonomy in high-performance computing (HPC) and AI by advancing open source technologies that will power future European supercomputers without having to rely on external input. In these times, it is crucial for Europe to build and maintain its own resources in this field. Developing homegrown computing solutions requires tight collaboration and co-design along the full stack from hardware to system software through HPC libraries and compilers to applications. The work of DARE is based on the open source RISC-V instruction set architecture for computer chips. The idea is for Europe to develop and produce its own RISC-V processors and accelerators for building future supercomputers, along with developing a complete, optimised, competitive and sustainable HPC/AI software stack for the DARE hardware.

Below: DARE kickoff, Barcelona, Spain, 6 March 2025

PDC is participating in the DARE project (https://www.pdc.kth.se/about/pdc-news/kth-pdc-partner-in-newly-launched-digital-autonomy-project-1.1391064) as one of the main developers of GROMACS, a widely used code for high-performance molecular dynamic simulations. The project started several months ago, but the collaborative work is already in full swing. Along with partners from the Barcelona Supercomputing Center, Tampere University, Chalmers University of Technology and the CSC – IT Center for Science, PDC is working on the co-development of SYCL runtimes and compilers for targeting RISC-V vector extensions. (The vector extensions in the RISC-V architecture provide a powerful way to perform parallel computations which enhance performance for tasks like scientific simulations and machine learning.) In addition, PDC is working with the EU-funded OpenCHIP company (https://openchip.com) on refining the desired hardware specifications to ensure that GROMACS will run efficiently on the new vector accelerator (known as VEC) that is being developed by DARE for HPC and AI.

For more information about the DARE SGA 1 project, see the project's home page https://dare-riscv.eu and the European Commission (EC) page about the project funding https://cordis.europa.eu/project/id/101202459.

KTH & AstraZeneca Partner to Revolutionise Drug Development with HPC & AI

Mårten Ahlquist, KTH, & Christof Jäger, AstraZeneca

The KTH Royal Institute of Technology and the pharmaceutical giant AstraZeneca have launched a cutting-edge collaboration to accelerate drug development using high-performance computing (HPC), quantum chemistry, and machine learning (ML). At the heart of this initiative is VeloxChem, a next-generation quantum chemistry software developed at the PDC Center for High Performance Computing at KTH.

As drug development increasingly relies on computational methods, VeloxChem represents a major leap forward. Written from scratch to exploit modern supercomputers, including GPU-enabled systems like the National Academic Infrastructure for Supercomputing in Sweden (NAISS) Dardel supercomputer system hosted at KTH, VeloxChem can handle molecular simulations on a vastly larger scale than traditional tools. Its hybrid design – using Python for user interaction and high-performance languages like C++, HIP and CUDA for intensive computations – makes it both powerful and user-friendly.

The collaboration aims to build a cloud-based environment where AstraZeneca researchers can access VeloxChem remotely to perform complex simulations and train Artificial Intelligence (AI) models. This setup will support tasks such as predicting drug reactivity, optimising synthesis pathways, and modelling molecular interactions – all critical stages in the pharmaceutical pipeline. AstraZeneca, with a strong presence in Sweden and a growing focus on AI-driven research, plans to use VeloxChem to generate large, high-quality datasets which are needed to train graph-based neural networks. These models could significantly improve predictions of molecular properties like stability, selectivity, and toxicity, cutting both time and cost in drug discovery.

With support from Vinnova's "Advanced Digitalisation" programme for a pilot study, the project's first milestone is a cloud-accessible prototype of VeloxChem running on the Dardel supercomputer, complete with a user-friendly interface and secure access. A workflow for calculating reactivity parameters in organic molecules will also be developed, laying the groundwork for broader applications such as electrosynthesis and crystallisation prediction.

Long-term goals include expanding the solution to other industrial sectors and enabling more secure handling of confidential data. KTH researchers believe this collaboration will also push academic software development to better align with real-world needs. This strategic partnership showcases how academia and industry can unite to harness Sweden's national HPC infrastructure, driving innovation in drug development and beyond.

PDC Summer School 2025

Alessandra Villa, PDC

Each year in August, the PDC Center for High Performance Computing (PDC) and the KTH School of Electrical Engineering and Computer Science (EECS) hold the PDC summer school "Introduction to High Performance Computing". The school is part of the Swedish e-Science Education (SeSE) programme and is supported by the Swedish e-Science Research Centre (SeRC). Participants learn the skills needed to utilise high-performance computing (HPC) resources for research. The school consists of lectures and guided hands-on lab exercises using the Dardel HPC system, which is hosted at PDC.

There are lectures on research software engineering, modern HPC architectures, and both parallel and GPU programming, along with hands-on practice using OpenMP-CPU, OpenMP-GPU,

Above: PDC Summer School, 22 August 2025, Stockholm

and MPI, plus programming AMD GPUs with HIP and SYCL, given by national and international speakers. This year the speakers were Ana Lucia Varbenescu (University of Twente), Erwin Laure (Max Planck Computing and Data Facility), Tim Dykes and Harvey Richardson (HPE HPC/AI EMEA Research Lab), Samantha Wittke (CSC – IT Center for Science), and Andrey Alekseenko and Niclas Jansson (PDC).

The school started with an online introduction for two half-days, ensuring that the attendees had the basic knowledge required to work in an HPC environment. PDC staff Parasuram Indraganti, Arash Banaei, Karl Johan Westrin, Ilari Korhonen, and Johan Hellsvik, lead by Henric Zazzi, helped the attendees to become familiar with the HPC environment at PDC and to be able to log in and use Dardel. On Monday the following week, the PDC Director, Patrick Norman, opened the on-site part of the school at the KTH main campus. As well as lectures and hands-on sessions, the five days at KTH included short talks on important HPC topics, round-table discussions, flash talks by the participants and a tour of the PDC computer hall and Dardel guided by Luca Manzari from PDC.

Two years ago, short talks on sustainability aspects in supercomputing – plus case studies in scientific disciplines like quantum chemistry, fluid dynamics and molecular dynamics – were added to the programme of the school. The participants perceived these initiatives very positively. Thus, this year the program included a sustaINABILITY lecture by Luca Manzari and lectures on the scientific software VeloxChem, GROMACS and Neko from an HPC perspective given by Xin Li, Szilárd Páll and Niclas Jansson (all from PDC).

To promote interactions and networking between participants and trainers, this year we introduced round-table discussions on parallel and GPU programming, plus career paths in HPC. Attendees were invited to give a short presentation on themselves, their project and their interest in HPC. This year, there were around 30 attendees from different parts of Europe, at various stages of their careers with a high representation of Ph.D. students from Swedish academia.

The summer school has been held since 1996 and has always been a collective effort. A big thanks to all the PDC staff over the years who have made the summer school possible and successful!

GANANA: A New Europe-India HPC Collaboration Has Been Launched!

Cathrine Bergh, PDC

India has made significant strides in high-performance computing (HPC) over the last decade, including the deployment of seven large-scale HPC systems (with performance above 1 PFLOPS) and ambitious future plans for developing a completely indigenous processor. The newly launched GANANA project aims to unite Europe's mature HPC competency with India's rapidly growing capabilities in order to tackle shared scientific problems of high societal impact.

GANANA was officially launched on the 1st of February this year and is centred on three different scientific domains: the life sciences, geographical hazards, and weather and climate modelling. These fields all share a great need for efficient computation to generate predictions of high societal impact, especially in the face of various disasters, such as diseases, earthquakes or harsh weather conditions. However, fast and accurate

Above (from left to right): Rossen Apostolov (GANANA Project Director), Sunita Verma (India Ministry of Electronics and IT), Anders Söderholm (KTH President) and Juan Pelegrin (European Commission), 6-7 May 2025, Stockholm

Below: GANANA project launch, 6-7 May 2025, Stockholm

calculations require optimised and bug-free computer software. GANANA therefore aims to create new avenues for collaboration between researchers already working on HPC codes within these three scientific fields, for example, by gathering domain expertise from different institutions and transferring knowledge between domains that share many computational approaches but which, by tradition, do not interact much. To facilitate collaboration, GANANA also includes a program to exchange both personnel and HPC resources between Europe and India. This way, we can learn more about each other's expertise, approaches, and future needs.

KTH and PDC, as the project coordinator, spent much of the spring this year setting up channels for collaboration and getting the GANANA project going. This work finally culminated in the GANANA Launch Symposium held in May 2025, where we had the pleasure of inviting representatives from all 10 European partners, three Indian partners, and other stakeholders to Stockholm to discuss the future of the project. We were pleased to witness many constructive discussions and concrete plans being made. We now look forward to rolling up our sleeves for the next phase of the project and beginning the journey towards creating a long-lasting collaboration between Europe and India for future impactful scientific high-performance computing. KTH and PDC are coordinating the project and will primarily contribute to accelerating and enhancing the GROMACS molecular dynamics software (through implementing fixed precision coordinates, modernising communication algorithms, and improving software/hardware co-design) and to the development of a framework for AI-driven immunogenic peptide prediction.

For more information about the GANANA project, see the project's website at https://ganana.eu.

PDC-Related Events

BioExcel Summer School on Biomolecular Simulations 2026
7-12 June 2026, Pula, Sardinia, Italy
https://bioexcel.eu/events/summer-school-on-biomolecular-simulations-2026

HPC Sources

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

HPC in Sweden & Scandinavia

- ENCCS
  http://enccs.se
- Mimer AI Factory
  https://mimer-ai.eu
- NAISS
  https://naiss.se
- NeIC
  http://neic.no
- SeRC
  https://e-science.se
- SeSE
  http://sese.nu

European HPC ecosystem

- EOSC EU Node
  https://open-science-cloud.ec.europa.eu
- ETP4HPC
  https://www.etp4hpc.eu
- EuroHPC
  https://eurohpc-ju.europa.eu
- HPC in Europe
  https://hpc-portal.eu
- LUMI
  https://www.lumi-supercomputer.eu
- PRACE
  https://www.prace-ri.eu

A selection of projects that PDC is involved with

- BioExcel CoE
  https://bioexcel.eu
- CEEC
  https://www.ceec-coe.eu
- EBRAINS
  https://ebrains.eu
- EPICURE
  https://www.epicure-hpc.eu
- EUMaster4HPC
  https://eumaster4hpc.uni.lu
- EXCELLERAT
  https://www.excellerat.eu
- GANANA
  https://ganana.eu
- Plasma-PEPSC CoE
  https://plasma-pepsc.eu

HPC news sources

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