PDC Center for High Performance Computing

PDC Newsletter 2/24

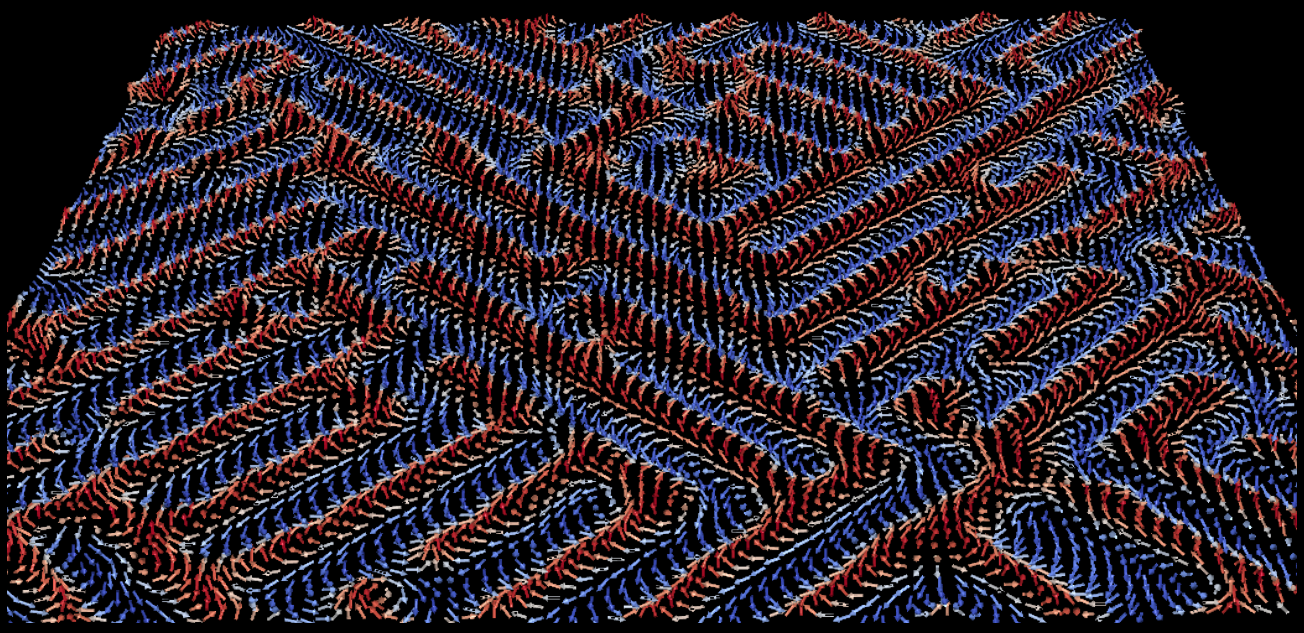

Patrick Norman, Director, PDC

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

The cover shows the spin-spiral magnetic state, with 90-degree domains forming a “labyrinth”, of a Mn monolayer on a W(110) surface modelled using atomistic spin dynamics in the UppASD software and based on magnetic interactions calculated using the RSPt software.

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

As called for at this time of year, here follows a Christmas stocking filled with presents for your enjoyment in the form of articles describing recent events and scientific advances centred around PDC. After having read the draft manuscript of this newsletter (with an ever so biased perspective), I was filled with pride by the diversity of activities shown, and, in composing the editorial, I had the opportunity to reflect on the important role PDC plays for research and development in Sweden.

The *cover article* demonstrates in silico material design by means of scale-bridging modelling based on first principles in the first layer and carried out on contemporary supercomputer hardware. The advancements in hardware continue to enable ever more realistic simulations of the real world, but it is well worth reminding ourselves that the development of methods and software may well be equally, if not even more, important. One aspect of software development is to make more efficient use of complex hardware, and this is demonstrated well by the work on the *GROMACS program for execution on AMD GPUs*. The importance of such optimisation work cannot be overestimated in view of the high costs involved with a high-performance computing (HPC) infrastructure. Another aspect of software (in combination with method) development is to enable new scientific discoveries. We see this demonstrated in the developments reported for the *VeloxChem program for embedded quantum mechanical modelling of truly large-scale and complex molecular systems*.

On the education and training side, the 2024 *PDC Summer School* continued an annual tradition of the highest value that has educated multiple generations of scientists in parallel programming – I attended this school as a Ph.D. student back in 1997! Likewise, this year’s *BioExcel summer school and conference* continued BioExcel’s excellent educational offerings for biomolecular researchers.

As the leading technical university in Sweden, KTH is important to Swedish industry, and we are proud of our industrial collaborations.

Above: BioExcel Summer School on Biomolecular Simulations 2024

Below: 2nd BioExcel Conference on Advances in Biomolecular Simulations, 2024

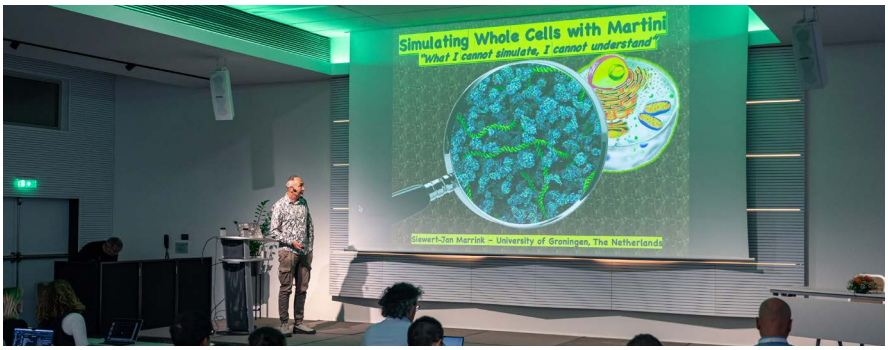

In October, PDC was welcomed to Scania for a day filled with presentations giving insights into the use of computational fluid dynamics (CFD) simulations for the construction of trucks. Apart from being highly inspirational, I hope such interactions can lead to improved services and collaborations in the future. In this newsletter, we find reports on expanded industrial collaborations with specific mentions of *Sandvik Coromant* and *Elfly*.

"Nothing lasts forever but the certainty of change." Bruce Dickinson

The time has come for KTH and PDC to thank Prof. Lennart Johnsson for his services as chairman of the board. Lennart received a Ph.D. degree in 1970 at Chalmers University of Technology and has since had an outstanding international career with positions at UCLA, Caltech, Yale, and Harvard. He is currently based at the University of Houston. I have personally only known Lennart during the past few years. With his wisdom and knowledge, he has guided my work at PDC. He has been nothing but kind and generous and very caring, not only for PDC as an organisation, but also for its staff members.

At the same time, we welcome our new board chairman, Dr. Per Öster. Per is the director of the Advanced Computing Facility at the CSC – IT Center for Science in Finland, and he has served as deputy director of PDC in the past. Apart from his directorship at CSC, Per holds several positions of trust in the landscape of European HPC, and KTH believes that this will fit well with the international profile of PDC. I am looking forward to working with Per and benefitting from his expertise and network.

Patrick Norman, Director PDC

In This Issue

Scale-Bridging Modelling and Designing New Materials in a "Virtual Lab"

Vladislav Borisov, Department of Physics and Astronomy, Uppsala University

Motivation

One challenge of modern materials science is to find new functional materials (such as magnets, superconductors and mechanically robust alloys) with better performance. Theoretical simulations have proven to be very helpful in this respect by providing a sort of "virtual lab" environment (as illustrated in the diagram below) where one can study different materials, make predictions about their properties based on computer simulations and find ways of improving their properties much faster than in a physical laboratory setting. The underlying equations describing any material are those of quantum mechanics where the behaviour of all interacting electrons and atomic nuclei in the material can be, in principle, predicted almost exactly. However, for realistic models of materials, the quantum-mechanical calculations are too expensive, even for modern supercomputers, and hence simpler so-called effective models become necessary to study the systems that comprise those materials on different length scales, bridging the microscopic (atomic scale) and macroscopic (device scale) descriptions. This article gives a brief overview of different steps of scale-bridging modelling in the context of magnetic materials. (More in-depth information is available in the references [1], [2] and [3]).

Scale-bridging Methodology and Examples

1. Electronic Properties

When it comes to studying a particular material, the minimal information that should be provided as input for computer simulations is its crystal structure and chemical composition. From

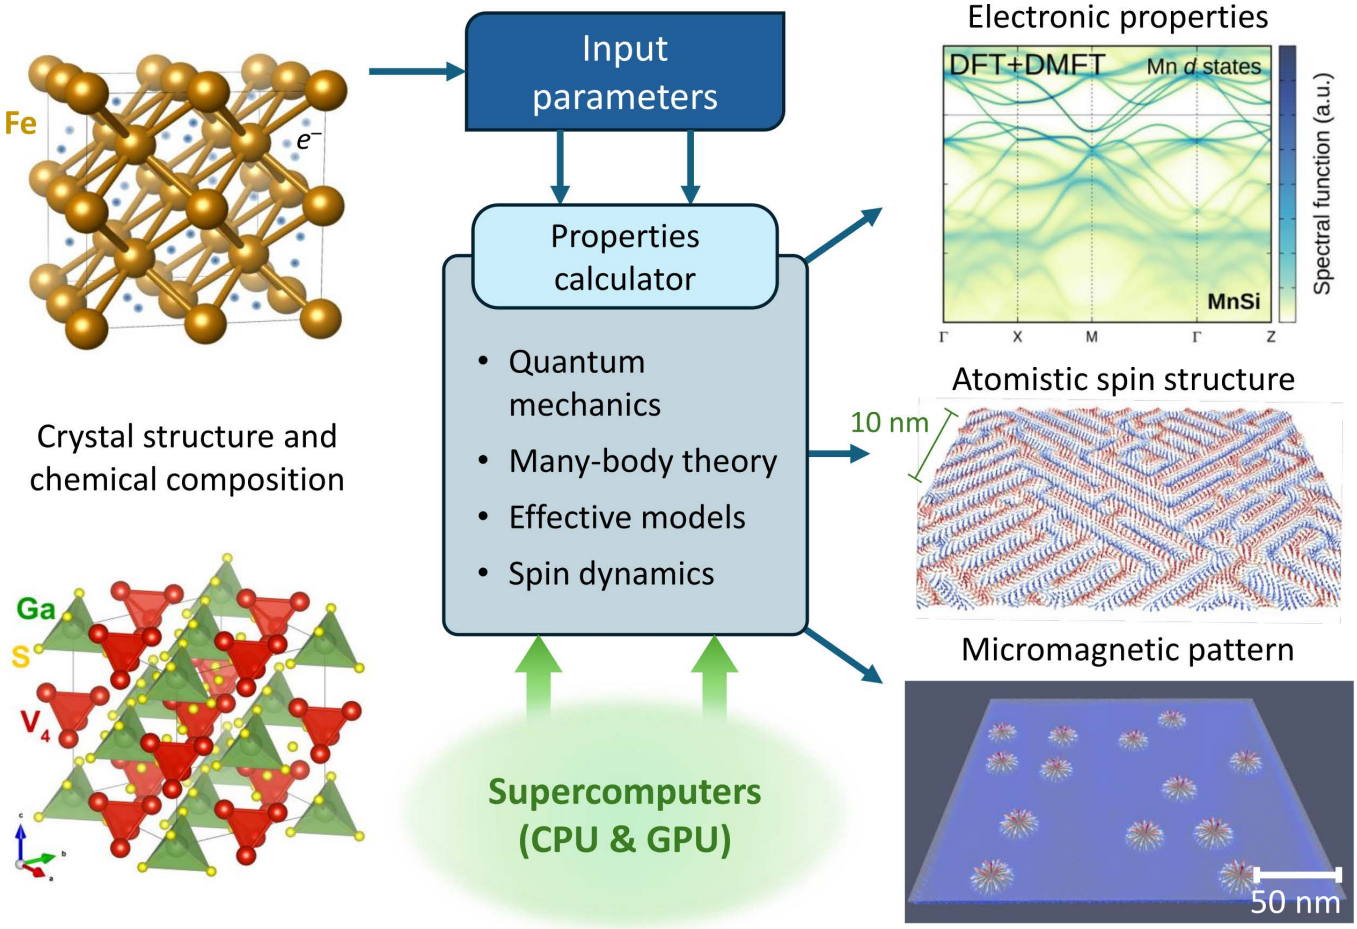

Above: The diagram illustrates the schematics of a scale-bridging modelling workflow with different calculation steps corresponding to length scales typically between a few Angstrom and several micrometres. Parts of this figure are reproduced from references [5] and [11].

this, it is possible to write down the full quantum Hamiltonian of the system that can be solved using, for example, density functional theory (DFT). The latter makes it possible to calculate, so far approximately but quite accurately, the key electronic properties of the system, such as whether it is metallic or insulating, what kind of chemical bonding it has and whether there are finite magnetic moments. The sole reason why DFT only gives approximate results for the calculation of ground-state properties at present is that the exact expression for the exchange-correlation functional is not known yet. However, there are a number of well-established approximations to this functional which have been thoroughly tested over recent decades, for example, local-spin density and generalised-gradient approximations, with which quantities like equilibrium crystal volume and magnetisation can be predicted within a 5% deviation from the measured values. Also, in the case of transition metal systems where correlations can be moderate or strong, the description of electronic properties can be much improved by using dynamical mean-field theory (DMFT) on top of DFT, which is the idea of the DFT+DMFT method [4]. The resulting electronic properties can change in terms of bandwidth and lifetime of electronic states, which has impact on all the other properties, especially the magnetic ones, as elaborated upon, for example, in recent work [5].

In the next step, since we are interested in magnetic properties on the larger scale, interactions between the atomic magnetic moments are calculated using the Lichtenstein-Katsnelson-Antropov-Gubanov (LKAG) approach [6] for different pairs of atomic moments at varying distances. The computational efficiency of this approach is characterised by the fact that one can use the chemical unit cell of the chosen system to address the magnetic interactions at much larger distances. For example, for ferromagnetic iron where the interatomic distance is around 2.8 Å, interactions for distances up to 15 Å can be calculated routinely and reliably using the RSPt software [7,8] where the LKAG approach is available. Such calculations would otherwise require large

Below: The graphs below show the computational time required for calculating nearest-neighbour Heisenberg interactions in the body-centred cubic (bcc) structure of iron versus the number of k-points (number of cores is fixed at 64) and the number of cores for fixed k-points. The benchmarks were performed with the RSPt code on the Tetralith system at the National Supercomputing Centre (NSC) in Linköping.

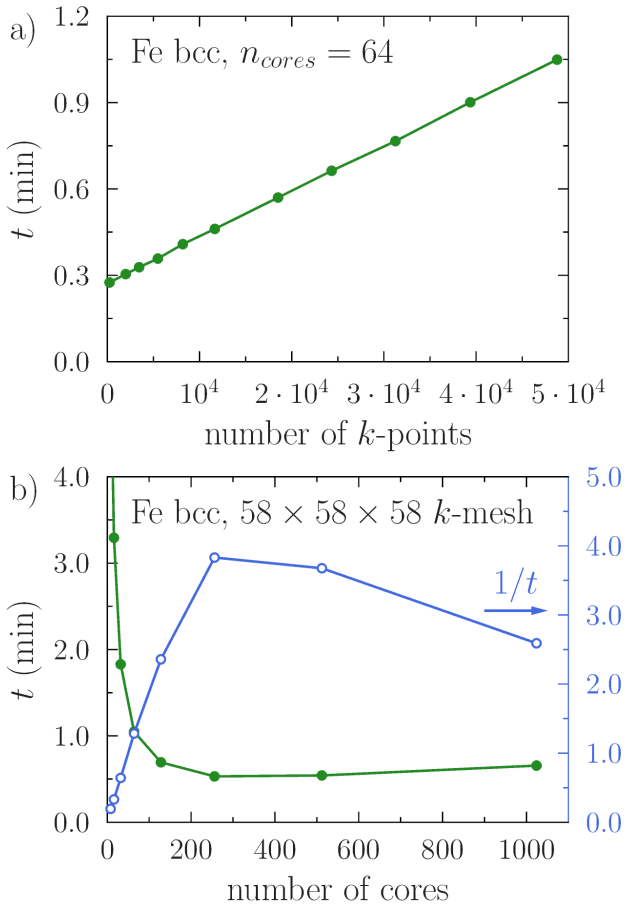

supercells, leading to high computational costs. The limiting value of real-space distance between interacting spins which can be treated in the LKAG method is largely determined by the number of *k*-points in the reciprocal space that are used in the DFT calculation: the magnetic interactions for larger distances converge to stable values for a sufficiently fine *k*-mesh. From the top graph above, which is plotted for the simple example of iron, one can see close-to-linear scaling of the calculation time with the total number of *k*-points. This is to be expected since the magnetic interaction between two spins is evaluated from a sum over different *k*-points, which should also theoretically be parallelisable almost linearly with respect to the number of cores. In the bottom graph, however, one can observe linear scaling for up to around 256

cores, which was tested with the RSPt code. For a larger number of cores, there is a considerable decrease in parallelisation efficiency, most likely due to additional overhead time, for example, internode connections. For systems that are more complex than iron, one can expect better scaling, that is, up to a higher number of cores. For example, we looked at the Co/Pt(111) system from [5] where the Co monolayer is magnetic and is located on the Pt(111) surface. In the graph below, the inverse computational time is plotted versus the number of cores for the calculation of exchange parameters for this system on the Tetralith supercomputer at the National Supercomputing Centre (NSC) in Linköping (1,800 *k*-points), as well as on the old and new nodes of the Dardel system at PDC for 1,800 and 6,075 *k*-points. (The old nodes refer to Dardel nodes before the interconnect upgrades in early 2023 and the new nodes refer to the nodes after the upgrade to Slingshot 11.) Here one can observe better scaling with respect to the number of cores and *k*-points than for iron, leading to parallelisation benefits even for 1,024 cores.

Since electronic correlations in certain systems can be more complicated than in iron, advanced methods like DFT+DMFT might be necessary to calculate the magnetic interactions more accurately. In recent work [5], it was demonstrated

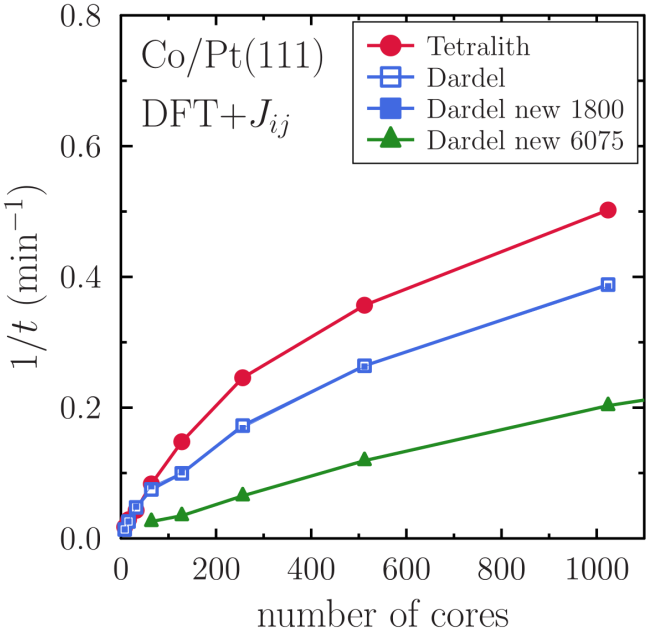

Above: Scaling benchmark for the calculation of magnetic interactions in Co monolayer on Pt(111) surface using the RSPt code

how important such correlation effects can be for different types of systems (including Co/Pt(111) mentioned above), and their possible origins in terms of electronic properties were discussed.

2. Atomistic spin dynamics

Once reliable values of magnetic interaction parameters are obtained, one can construct an effective spin model, such as the Heisenberg or extended Heisenberg model, that abstracts away from the electronic description of the system and pictures every atom as a well-defined magnetic moment that interacts with its neighbours. The model then describes the magnetic energy of the system due to these interactions and makes it possible to calculate the torques acting on each spin if the system is out of equilibrium. Based on that and the Landau-Lifshitz-Gilbert (LLG) equation, one can simulate the time evolution of all the magnetic moments as the system moves towards equilibrium, which is the main subject of *atomistic spin dynamics* (see book [1]). This can be useful when searching for an equilibrium magnetic configuration. An example is shown on the cover which depicts the spin-spiral magnetic state of a Mn monolayer on a W(110) surface obtained from *simulated annealing* using the UppASD software [1]. In this method, one starts from a random magnetic configuration in a large supercell with thousands or even millions of spins (see the next section for further details on computational scaling) at high temperature and then calculates the spin dynamics while gradually lowering the temperature of the system until, for example, zero temperature is reached. The system usually ends up in one of the many local energy minima which can represent some of the features of the actual lowest energy state; in this case, it is the spin-spiral stripe-shaped domains that are at 90 degrees to each other and can coexist because they have the same energy. These domains have also been measured in experiments [9] on a similar length scale of a few nanometres. As discussed, for example, in review [3], longer annealing time can help to get closer to the lowest energy state in annealing simulations. In the case of Mn/W(110), this increases the fraction of spin-spiral domains of a certain orientation, that

Below: These images show emergent topological magnetism in B20 multilayers. In the lower plots, the atomic spins are indicated by arrows. This figure is reproduced from [10].

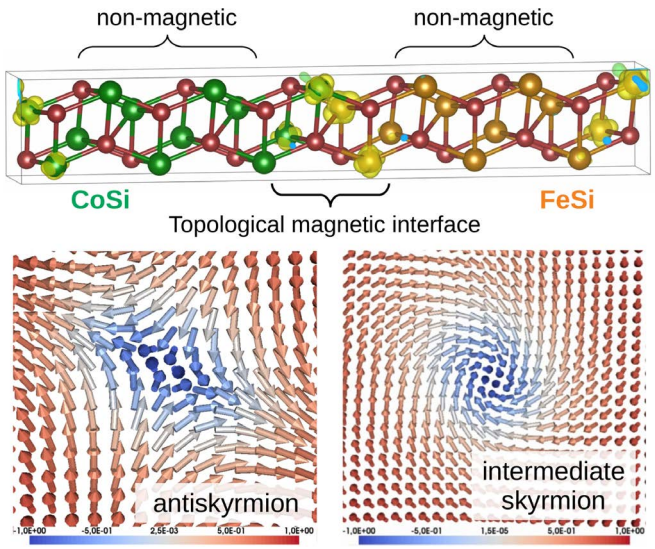

is, it makes the system closer to a single-domain state. In general, for more complicated systems (especially with frustrated magnetic interactions), the problem of finding the global energy minimum is highly non-trivial, and simulated annealing can only provide approximate answers and insights, which are nevertheless still useful.

Another example is from a recent study [10] which makes predictions for multilayers formed from the B20 compounds FeSi and CoSi (see figure above) that have not yet been confirmed experimentally. Theory suggests that, although neither compound shows any long-range magnetic order in bulk, the presence of nanoscale interfaces in FeSi/CoSi structures leads to the emergence of magnetic order near the interfaces. Furthermore, even though bulk FeSi and CoSi are already chiral in the bulk, their chirality properties are modified considerably by the interface effects, leading to stabilisation of topological magnetic textures of different sorts: antiskyrmions, intermediate skyrmions and antiferromagnetic skyrmions (some of which are shown in the figure above). This variety can be traced back to the structure of chiral magnetic interactions calculated using the LKAG approach, which turns out to be very different from what is typically observed in B20 magnets, such as FeGe. Different possibilities for topological magnetism in these systems were explored in [10] via extensive atomistic spin dynamics for different B20 multilayers using the UppASD software and large simulation cells with 875,000 to 2,000,000 spins, depending on the interface configuration being studied. This motivates future experiments and studies of similar nanostructures with emergent interface phenomena.

It is noteworthy that from spin dynamics simulations one can not only get the picture of atomic spins in real space forming different magnetic patterns (ferromagnetic, antiferromagnetic or non-collinear ones, as shown in the figures on page 4 and to the left) but also obtain information on magnetic correlations and excitations in the system that can be compared with experiment. This comparison is important for validating the accuracy of the constructed effective spin model and theory approximations in general.

3. Micromagnetics

Despite the power of the spin dynamics method, it has limitations in terms of the system size that can be simulated within a reasonable computational time. On modern supercomputer architectures, it is feasible to simulate spin systems with up to a few million spins, where simulated annealing may take up to a few weeks of computational time, depending on the complexity of magnetic interactions. To overcome these limitations and increase the length scale of simulations, one can switch to the micromagnetic description where the magnetisation in the material is assumed to vary slowly between neighbouring atomic spins and is effectively treated as a continuous variable. At the same time, the LLG equation is used to describe the magnetisation dynamics at finite temperature, just with different system parameters (like spin stiffness, DM spiralisation, and magnetisation density) that can be calculated from the atomistic interactions and properties. The size of the system that can be simulated micromagnetically is larger (above several micrometres) than the upper limit for the atomistic spin dynamics. However, the atomistic approach is more accurate and general, while the micromagnetic approach is limited to systems with slowly varying magnetisation.

One example is discussed in the theoretical study [11] of multiferroic skyrmionic spinel

Below: These images show the crystal structure (a) and the two magnetic phases: spin spiral (b) and skyrmions (c) in lacunar spinel GaV_4S_8. The images are reproduced from [11].

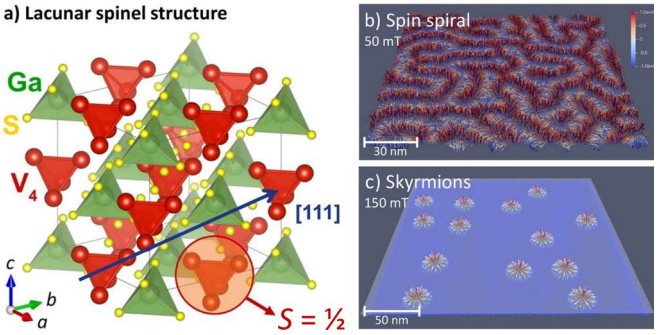

GaV_4S_8 (see figure above) which is known to host Néel skyrmions of around 20 nm size in bulk, a rather unique property of lacunar spinels [12]. The goal of that work was to understand the magnetic state of the V_4 clusters which play the role here of effective spins interacting through the Heisenberg and Dzyaloshinskii-Moriya interactions that are usually also important for other topological magnets [13]. By calculating the micromagnetic parameters from first-principles atomistic data for two cluster configurations, which is widely discussed in the literature, it was found that the DM interaction is more pronounced in the so-called distributed-moment state where all V sites have a sizeable magnetic moment. The micromagnetic simulations for the two cluster states showed that this distributed-moment state much better supports the spin spiral and Néel skyrmions observed in experiments. The two visualisations on the right of the figure above indicate that simulated annealing in micromagnetic simulations does not lead to the lowest-energy state (ground state) but to a local energy minimum which can be fairly close to the ground state. For example, instead of an ordered skyrmion lattice, individual skyrmions are obtained at different positions, dependent on the initial random magnetic configuration. However, further analysis demonstrated that setting the local magnetic anisotropy to zero and annealing the magnetic state can actually stabilise a relatively ordered periodic skyrmion lattice, which preserves its structure even after the local anisotropy is restored in the simulation. This means that approaching the lowest energy state for a given magnetic system is highly non-trivial and depends on the value of magnetic interactions that can create energy barriers between different local energy minima. Despite this issue, analysing local energy minima obtained routinely in atomistic spin-dynamics and micromagnetic simulations can provide useful information on the magnetic properties that are realised in the studied material, as discussed in this article for the examples of Mn/W(110) bilayer [3], B20 multilayers [10] and GaV_4S_8 spinel [11].

Importantly, multiscale methods [14] can be applied where different parts of a given magnetic system are described either atomistically (for example, regions around point defects) or micromagnetically (the rest of the system) and the two descriptions are smoothly connected. This powerful combination of methods increases the accuracy and computational efficiency of modelling complex magnetic systems and makes it possible to discover qualitatively new phenomena (see [14]).

To summarise, the scale-bridging approach developed by a vast research community enables the modelling of magnetic materials on different length scales where model parameters are material-specific and effectively originate from accurate quantum-mechanical descriptions of the underlying crystal lattice. This approach takes full advantage of modern CPU- and GPU-based supercomputing architectures, the performance of which makes it possible to simulate realistic material models and push the boundaries of materials science.

Acknowledgements

Fruitful discussions with Dr. Olle Eriksson and Dr. Anna Delin are gratefully acknowledged.

References

1. O. Eriksson, A. Bergman, L. Bergqvist, J. Hellsvik, “Atomistic Spin Dynamics: Foundations and Applications, Atomistic Spin Dynamics: Foundations and Applications”, Oxford University Press, Oxford, UK, (2017).
2. A. Szilva, Y. Kvashnin, E. A. Stepanov, L. Nordström, O. Eriksson, A. I. Lichtenstein, M. I. Katsnelson, “Quantitative theory of magnetic interactions in solids”, Rev. Mod. Phys. 95, 035004 (2023).

3. V. Borisov, “From electronic structure to magnetism and skyrmions” (Invited Topical review), Electron. Struct. 6, 023002 (2024).
4. G. Kotliar, S. Y. Savrasov, K. Haule, V. S. Oudovenko, O. Parcollet, C. A. Marianetti, “Electronic structure calculations with dynamical mean-field theory”, Rev. Mod. Phys. 78, 865 (2006).
5. V. Borisov, Y. O. Kvashnin, N. Ntallis, D. Thonig, P. Thunström, M. Pereiro, A. Bergman, E. Sjöqvist, A. Delin, L. Nordström, O. Eriksson, “Heisenberg and anisotropic exchange interactions in magnetic materials with correlated electronic structure and significant spin-orbit coupling”, Phys. Rev. B 103, 174422 (2021).
6. A. I. Liechtenstein, M. I. Katsnelson, V. P. Antropov, V. A. Gubanov, “Local spin density functional approach to the theory of exchange interactions in ferromagnetic metals and alloys”, J. Magn. Magn. Mater. 67, 65 (1987).
7. J. M. Wills, B. R. Cooper, “Synthesis of band and model Hamiltonian theory for hybridizing cerium systems”, Phys. Rev. B 36, 3809 (1987).
8. J. M. Wills, O. Eriksson, P. Andersson, A. Delin, O. Grechnyev, M. Alouani, “Full-potential electronic structure method, in Energy and Force Calculations with Density Functional and Dynamical Mean Field Theory”, Springer Series in Solid-State Science Vol. 167, Springer, Berlin, Heidelberg, 2010.
9. P. Ferriani, K. von Bergmann, E. Y. Vedmedenko, S. Heinze, M. Bode, M. Heide, G. Bihlmayer, S. Blügel, R. Wiesendanger, “Atomic-scale spin spiral with a unique rotational sense: Mn monolayer on W(001)”, Phys. Rev. Lett. 101, 027201 (2008).
10. V. Borisov, A. Delin, O. Eriksson, “Tunable topological magnetism in superlattices of nonmagnetic B20 systems”, Phys. Rev. B 110, L060407 (2024), Editors’ Suggestion Letter.
11. V. Borisov, N. Salehi, M. Pereiro, A. Delin, O. Eriksson, “Dzyaloshinskii-Moriya interactions, Néel skyrmions and V_4 magnetic clusters in multiferroic lacunar spinel GaV_4S_8”, npj Computational Materials 10, 53 (2024).
12. I. Kézsmárki, S. Bordács, P. Milde, E. Neuber, L. M. Eng, J. S. White, H. M. Rønnow, C. D. Dewhurst, M. Mochizuki, K. Yanai, H. Nakamura, D. Ehlers, V. Tsurkan, A. Loidl, “Néel-type skyrmion lattice with confined orientation in the polar magnetic semiconductor GaV_4S_8”, Nat. Mater. 14, 1116-1122 (2015).
13. N. Kanazawa, S. Seki, Y. Tokura, “Noncentrosymmetric magnets hosting magnetic skyrmions”, Adv. Mater. 29, 1603227 (2017).
14. É. Méndez, M. Poluektov, G. Kreiss, O. Eriksson, M. Pereiro, “Multiscale approach for magnetization dynamics: unraveling exotic magnetic states of matter”, Phys. Rev. Research 2, 013092 (2020).

Course Recap: Introduction to PDC Systems

Juan De Gracia Triviño, PDC

On the 27th of September 2024, PDC hosted an “Introduction to PDC Systems” course. This was run as an online session and was designed for beginners in high-performance computing (HPC) and for experienced users interested in exploring the Dardel system, including its new GPU partition. Led by PDC experts, this course offered a comprehensive look at Dardel’s infrastructure, covering essential topics such as setting up an account, scheduling jobs with Slurm, data storage, and software environments.

The course attracted around 25 attendees who actively participated, posing questions on a variety of topics. Highlights included sessions on efficient login practices, the use of ThinLinc, and managing Python environments with Jupyter notebooks and Conda. Additionally, the final segment introduced Singularity, focusing on containerised workflows, which are an increasingly relevant topic for researchers working in diverse computational environments.

Participants left with enhanced confidence in using Dardel, better prepared to leverage HPC for their research.

Elfly Collaborates with PDC on Electric Aircraft Design

John Larson, Elfly

The Elfly Group is a Norwegian electric aircraft company that is developing the world's first fully electric seaplane, the Noemi, with the dual aims of improving transport (for people and cargo) in coastal areas and saving three megatons of CO_2 emissions by 2050.

The Noemi is to be a twin-motor aircraft with a range of 200 km that will be powered by lithium batteries and an electric propulsion system. The plane will be able to carry up to nine passengers or could be modified for various purposes, such as carrying cargo.

As seaplanes can take off from and land on water, they can be far more convenient and feasible for coastal communities, where there may not be any land nearby that is suitable for runways for standard aircraft. Seaplanes can also travel much faster than boats or ferries. Overall, the use of non-polluting short-haul electric seaplanes could be a game changer for coastal communities (on mainlands, islands and inland around fjords and lakes). Since aviation is responsible for 5% of global greenhouse gas emissions and transportation overall for 20%, Elfly hopes to create the greatest possible gains in our shared battle against climate change by targeting emissions from fossil-fuel-based short-haul commuter and cargo flights in coastal areas.

Elfly is using the Dardel system at PDC to run detailed computational fluid dynamics (CFD) simulations which are used to evaluate and optimise the design of the aircraft. This ultimately improves the performance of the aircraft and is more efficient than building multiple physical models and evaluating them, for example, in wind tunnel experiments.

For details about the exciting developments at Elfly, see https://el-fly.no, and if your company would like to use high-performance computing resources at PDC, you are welcome to contact the PDC Business Unit (business-unit@pdc.kth.se).

Dardel Expanded with More Disk Space and NVIDIA GPUs

Gert Svensson, PDC

The Dardel system is now finally fully accepted! Due to various delays with deliveries during the installation period of the whole system (which spanned several years), PDC has received substantial compensation over that time. The latest compensation for delays, along with some unused funding from the initial investment, will now be used to expand Dardel with additional fast Lustre storage and some nodes with NVIDIA Grace Hopper GPUs.

The Lustre file system will be expanded with an additional 4.7 PB of hard disk storage and 260 TB flash storage. This will give a total Lustre capacity of 23.6 PB.

Many researchers who use Dardel have asked for NVIDIA GPU nodes to be made available at PDC. This is because many research applications are only available for the NVIDIA software stack with CUDA. In response to these requests, eight nodes will be added to Dardel, each with four Hopper (H100) GPUs and four NVIDIA Grace ARM CPUs. This will add a total of 32 NVIDIA Grace Hopper Superchips. Each of these nodes will have four Slingshot ports and the total expansion is expected to increase the performance of the system by around 1.3 PFLOPS, as measured with the high-performance Linpack (HPL) benchmark. For more information about the Grace Hopper Superchips, see https://www.nvidia.com/en-gb/data-center/grace-hopper-superchip.

Above: The photo shows an HPE Grace Hopper Superchip board with two nodes. Each node contains four NVIDIA Superchips, each with a CPU and a GPU. The board is made by HPE using the NVIDIA Grace Hopper Superchips.

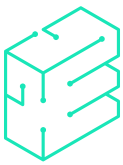

EPICURE

PDC & Sandvik Coromant Collaborate via EPICURE

Jonathan Vincent & Peter Larsson, PDC, with Martina Lattemann, Sandvik Coromant

The EPICURE project began in February 2024. The project was initiated by the European High Performance Computing Joint Undertaking (EuroHPC JU) to improve the support services available to projects that are awarded access to EuroHPC systems. At present, the entities that host the systems are responsible for providing basic-level help desk support for daily operational issues and problems. The idea is for EPICURE to provide higher levels of support. One way EPICURE is providing top-level support is through application support teams (ASTs), which are based at organisations that host EuroHPC systems. These teams work on application porting, optimisation and scalability improvements. Researchers whose projects are awarded an allocation on one of the EuroHPC systems are offered access to EPICURE's support services.

KTH is participating in EPICURE with PDC coordinating the activities. PDC's first EPICURE collaboration was with Sandvik Coromant, a Swedish company that supplies high-precision cutting tools and services to the metal cutting industry. Sandvik Coromant's tools are not just for cutting straight lines in metal: they cover a wide range of tasks involving the manufacture of metal items, including drilling and turning, as well as adding threads or grooves to different types of metal. There are continuous developments across the manufacturing industry which result in a need for challenging machining operations. Sandvik Coromant needs to keep adapting to those evolving needs. In addition, an increased focus on sustainability and recyclability is in the forefront when it comes to developing new products and production techniques. Applying computational tools for product development will lead the way for Sandvik Coromant in facing those challenges, as well as helping to shorten the time to market.

Under the umbrella of the EPICURE project, PDC staff members Jonathan Vincent and Peter Larsson worked with researchers at Sandvik Coromant to help them use the EuroHPC pre-exascale supercomputer system, LUMI, which is hosted in Finland. Jonathan and Peter ensured that the intended calculations on the LUMI supercomputer could be run efficiently by configuring and installing the necessary software packages. The difficult part was to establish an environment to run the workflow scripts passing information between the software packages. The collaboration has now been completed. It involved using the VASP code to generate large training sets based on density functional theory (DFT). These training sets were used together with the PACEmaker code, which uses machine learning (ML) to combine the many data points generated by VASP into a complete atomic cluster expansion (ACE) potential surface.

The resulting potentials were then used to perform large-scale molecular dynamics (MD) simulations with the LAMMPS code. These were used to understand the behaviour of the underlying materials under realistic conditions, such as those that occur during cutting operations. In this way, potential new designs or materials for tools can be investigated without having to build multiple prototypes. The active learning capabilities of PACEMaker identify unknown atomic configurations during MD simulations and retrain the ACE potential "on the fly".

For further information about the EPICURE project, see the website: https://epicure-hpc.eu, and for details about how PDC can help researchers (who have been allocated projects on LUMI and other EuroHPC systems) through EPICURE, please contact Rossen Apostolov (rossen@kth.se). Businesses and companies that do not currently have allocations on EuroHPC systems and that would like assistance directly from PDC (with high-performance computing for R&D activities) are warmly invited to contact the PDC Business Unit (business-unit@pdc.kth.se).

VeloxChem: Complex Polarisation Propagator Simulations of Circular Dichroism Spectra on GPUs

Xin Li, PDC

Theoretical studies of microscopic structures and chemical and physical processes in molecular systems provide valuable insight into the understanding of macroscopic functions by relating the molecular and electronic structures to properties and observables. Among the theoretical methods for performing such simulations, time-dependent density functional theory (TDDFT) is perhaps the most impactful due to its ability to study large, complex systems with manageable computational cost. However, in systems with multiple interacting chromophores (parts of molecules within a compound that absorb some frequencies of light and hence make the compound appear coloured), TDDFT calculations remain challenging due to the requirement of simulating a large quantum mechanical (QM) region and the difficulty of tackling the many systems within that region that are close to each other and also in excited states.

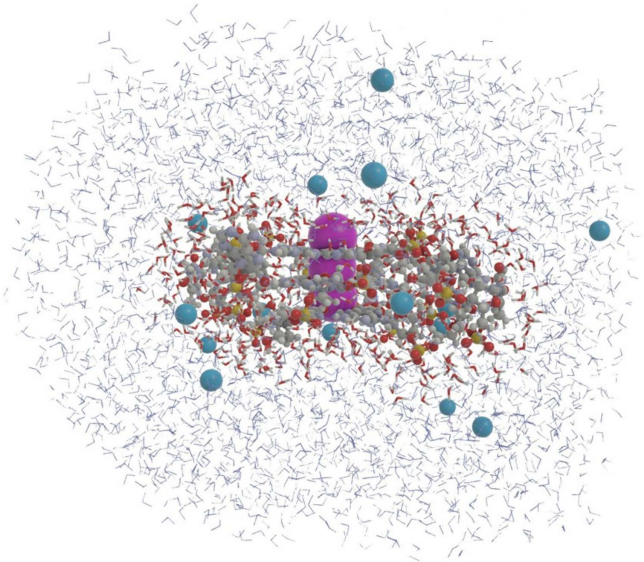

Above: Side view of the large model that includes the whole G-quadruplex and surrounding water molecules and ions

Below: The graph illustrates circular dichroism spectra for three models. Solid lines refer to broadened spectra from TDDFT, while circles refer to spectra from CPP. The inset shows the experimental spectrum [2].

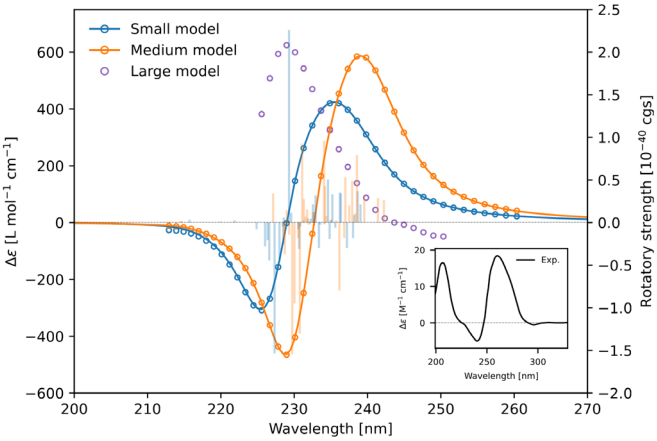

The complex polarisation propagator (CPP) formulation of response theory provides an alternative approach for studying a discretised frequency window of interest within the spectrum of light that is being emitted by the systems that are under investigation. The convergence of CPP can also be made robust and efficient by solving many response equations simultaneously. Based on recent developments in the VeloxChem quantum chemistry program (for calculating molecular properties and simulating a variety of spectroscopies), we have implemented real and complex linear responses where GPU-accelerated CPP calculations are realised by efficient evaluation of symmetric and antisymmetric auxiliary Fock matrices. The details of the implementation are summarised and documented in our preprint on ChemRxiv [1].

As an example, we carried out CPP calculations for the electronic circular dichroism spectrum of a G-quadruplex. We used three models of different sizes, and the largest model includes the whole G-quadruplex and surrounding water molecules as the QM region, plus an additional layer of water molecules as the molecular mechanical (MM) region, as shown in the figure to the left. All three models present clear bisignated signal characteristics of the G-quadruplex (see graph above); however, only the largest model reproduced the experimentally observed small negative band in the long-wavelength region. Our analysis of individual orbital excitations suggests

that the negative band can be ascribed, at least partially, to transitions centred on the thymine nucleotides as well as the presence of potassium ions [1].

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". ChemRxiv. (2024) https://doi.org/10.26434/chemrxiv-2024-rk6w2
2. J. Kypr, I. Kejnovska, D. Renciuk, M. Vorlickova, "Circular dichroism and conformational polymorphism of DNA". Nucleic Acids Res. (2009) 37, 1713–1725.

PDC Summer School 2024

Alessandra Villa, Niclas Jansson & Stefano Markidis, PDC

Every August the PDC Center for High Performance Computing (PDC) and the School of Electrical Engineering and Computer Science (EECS) at the KTH Royal Institute of Technology run 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). The school is open to academic and industry researchers and focuses on the skills that researchers need to effectively utilise high-performance computing (HPC) resources for research and/or development. It consists of lectures and guided hands-on lab exercises using the Dardel HPC system hosted at PDC.

This year's school ran for a week from 19-23 August on site at the KTH campus in Stockholm. The school started with a lecture by Ana Lucia Varbenescu, University of Amsterdam and University of Twente, Netherlands, on modern HPC architectures, followed by a lecture on research software engineering for HPC by Radovan Bast from the University of Tromsø, Norway. The first day closed with a lecture on sustainability aspects in supercomputing by Stefano Markidis, PDC. On Tuesday, Niclas Jansson, also from PDC, told the attendees about OpenMP-CPU and OpenMP-GPU. Tim Dykes and Harvey Richardson, from the HPE HPC/AI EMEA Research Lab, gave an overview on GPU programming and Joan Vinyals Ylla Català, Barcelona Supercomputing Center, Spain, spoke about performance analysis. At the end of the day, Luca Manzari from PDC led a tour of Dardel and the PDC supercomputer hall. Wednesday was fully dedicated to advanced GPU programming, debugging, and profiling with Tim and Harvey. During the last two days, the attendees learnt about MPI from Erwin Laure, Director of the Max Planck Computing and Data Facility, Germany. Case studies in various scientific disciplines helped to illustrate practical research applications of the topics: in particular, the scientific software GROMACS by Szilárd Páll (PDC), VeloxChem by Xin Li (PDC) and Neko by Niclas Jansson (PDC). Finally, PDC's director, Patrick Norman, presented PDC and the Swedish HPC Landscape. The complete schedule of the school and links to materials can be found here: https://www.pdc.kth.se/summer-school/2024/timetable-1.1319716.

The school brought together around 40 attendees at different stages of their careers and coming from different parts of Europe and the USA with a high representation of Ph.D. students from Swedish academia. Overall, the attendees ranked the school between very good and excellent. Some found it a bit too fast or too short. Indeed the PDC summer schools usually run for two weeks, but this year, due to dates overlapping with other events, the school was concentrated into one week. A big thanks to all the PDC staff who helped with mentoring during the summer school and who supported the school attendees on the Dardel supercomputer.

Above: PDC Summer School "Introduction to High Performance Computing", 19-23 August 2024

ENCCS
EuroCC National Competence Centre Sweden

ENCCS News

Apostolos Vasileiadis, ENCCS

Autumn at ENCCS has been brimming with events and news. ENCCS continues to support R&D in industry, the public sector and academia by providing high-performance (HPC) training and support. The team has been expanded with two HPC experts and a business developer. (For more about the team, see https://enccs.se/people.)

New Podcast: "Supercomputing in Europe"

ENCCS recently initiated a new podcast, "Supercomputing in Europe", in which people from the EuroCC project (https://www.eurocc-access.eu) and its centres of excellence (https://www.hpccoe.eu) talk with guests from around Europe about their work and give us interesting insights into the role of supercomputing in multiple fields, such as the life sciences, machine learning, engineering, data science, quantum computing and more. You can listen to the podcast on:

- Spotify: https://open.spotify.com/show/4pZ7nfUZTs3tDyfWVFAR8q?si=bd6153dc25824263,
- RSS feed to add to your favourite app: https://anchor.fm/s/f01f82a4/podcast/rss, and
- Apple podcasts: https://podcasts.apple.com/se/podcast/supercomputing-in-europe/id1768782069.

Focus on Artificial Intelligence (AI)

ENCCS, together with PDC, Stockholm AI, and Hyperight, recently organised the Supercomputing for AI seminar (for details, see https://enccs.se/events/supercomputing-for-ai). It included presentations and a hands-on session on how to port AI workflows into HPC systems using the Dardel system at PDC. There will be more occasions similar to this one in the near future.

ENCCS also co-organised and participated in the Nordic Industry Days conference in Copenhagen, with talks and discussions with experts from multiple industrial sectors, providing important insights and feedback on industrial AI and HPC needs and challenges. For a great overview of the conference, see: https://enccs.se/news/2024/09/nordic-industry-days-2024-in-retrospect.

Collaboration with PDC

PDC and ENCCS are now actively collaborating to enhance the support they provide to companies using high-performance computing. PDC is expanding its service portfolio to meet the unique needs of industry, offering resources and expertise to help businesses push the boundaries of R&D. Together, PDC and ENCCS are seeking a VIP industry partner to participate in an exclusive pilot project, where the selected company will receive free access to compute resources along with dedicated support. In return, the partner will provide valuable insights to help tailor the PDC and ENCCS services into a truly customer-centric offering that aligns even better with industry needs. If your company is interested in this opportunity, you are welcome to contact us at info@enccs.se. More information about the collaboration can be found in this article: *"PDC-ENCCS Collaboration Seeking Industry Partner"*.

Upcoming Events

Stay tuned to upcoming ENCCS events and the latest news by following ENCCS on LinkedIn (https://www.linkedin.com/company/enccs).

FAST. FLEXIBLE. FREE.
GROMACS

Learn to Code in GROMACS Online Workshop

Alessandra Villa, PDC

The 2024 edition of the workshop "Learn to code in GROMACS" (https://www.gromacs.org/workshop.html) was held over three days from the 10th to the 12th of September. The workshop brought together forty-four attendees and ten mentors from different countries, both from academia and industry, for this online event. Forty-four people attended the online lectures, while only a restricted number of attendees (eleven) were selected for the hands-on session to guarantee an optimal ratio between attendees and mentors (one-to-one).

On the first day, Sebastian Wingbermuehle (https://www.biophysics.se/index.php/members/sebastian-wingbermuhle) introduced "GROMACS GitLab and version control", and, in the afternoon, Mark Abraham (https://www.linkedin.com/in/mark-abraham-5a329599) spoke about "Software structure and interface". The following day, Andrey Alekseenko (https://www.biophysics.se/index.php/members/andrey-alekseenko) gave a presentation on "Testing and testing infrastructure", and, on the last days, Berk Hess (https://www.biophysics.se/index.php/members/berk-hess) spoke about "Everything around coding", giving a historical overview of GROMACS developments. The event closed with a roundtable discussion featuring a panel with Andrey Alekseenko, Berk Hess, Mark Abraham and Magnus Lundborg (https://www.biophysics.se/index.php/members/magnus-lundborg).

Mentoring took place online via breakout rooms. There were a total of eleven mentors who were from Sweden, Germany, France and the USA. During the hands-on session, the attendees each worked closely together with a mentor to address a pre-selected GROMACS issue. The issues were labelled as "2024 workshop" in the GROMACS GitLab repository (https://gitlab.com/gromacs/gromacs). The scope was such that the attendees could participate in, and directly experience, the whole review process. At the end of the workshop, five out of over fourteen selected issues had been closed. To avoid attendees getting stuck on a problem, they were asked to report the status of the issue they worked on at the beginning of each session. According to the type of problem, one or more mentors met with each attendee in a breakout room. We observed that assigning a reference mentor to each issue encouraged the attendees to ask questions. The lectures were recorded and posted on the BioExcel YouTube channel: Learn to Code in GROMACS (https://www.youtube.com/playlist?list=PLzLqYW5ci-2d82Pbcv9r712NQlCgjF-2E) and were uploaded on Zenodo Workshop: Learn to code in GROMACS (2024) (https://zenodo.org/records/13739992). We hope that the collected material will be re-used in the future to build self-learning tutorials.

The organisers would like to extend a big "thank you" to Andrey, Berk, Magnus, Mark and all the other mentors: Amr Alhossary (https://gitlab.com/aalhossary), Eliane Briand (https://gitlab.com/ElianeBriand), Giacomo Fiorin (https://research.ninds.nih.gov/staff-directory/giacomo-fiorin-phd), Hubert Santuz (https://gitlab.com/HubLot), Lukas Müllender (https://gitlab.com/lmuellender), Michele Pellegrino (https://gitlab.com/michele.pellegrino), and Vedran Miletic (https://vedran.miletic.net).

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

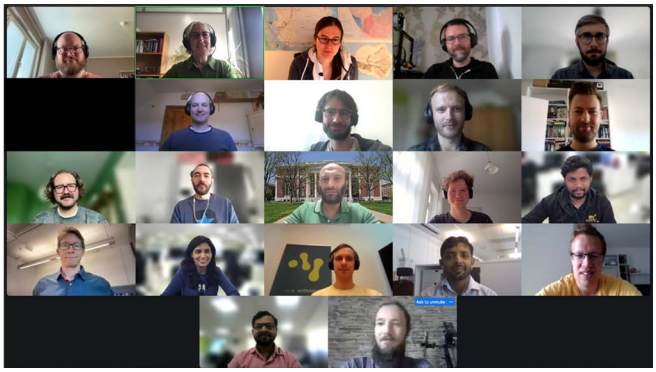

Above: Learn to code in GROMACS online workshop, 10-12 September 2024

New Director, Partners and Branches for NAISS

Björn Alling, NAISS

New NAISS Director: Professor Erik Lindahl

Linköping University's Vice-Chancellor, Jan-Ingvar Jönsson, appointed Professor Erik Lindahl as the new director of the National Academic Infrastructure for Supercomputing in Sweden (NAISS) on the 11th of November 2024. Eric will begin his role as the director on the 1st of January 2025. Erik has been involved in NAISS from its start as deputy chair of the steering committee. He is a leading scientist in the field of biophysics and has extensive experience within Swedish and international high-performance computing (HPC) infrastructures. Björn Alling will continue as acting director until the end of this year and will then continue to work for NAISS in a role supporting Erik from the start of 2025.

Upcoming NAISS & EuroHPC System Arrhenius

NAISS and the European High Performance Computing Joint Undertaking (EuroHPC JU) have agreed on the technical specifications for the procurement of the Arrhenius system and the call for proposals has gone out to the vendors. The deadline for submissions is set to the 31st of January 2025. The plan is for the CPU and storage parts of Arrhenius, including sensitive data, to be brought online during the autumn of 2025. Due to a very high demand in the market for relevant components, the GPU part of the system is expected to be delivered a few months later.

NAISS Partners and Branches

Luleå Technical University signed up as the eleventh NAISS partner on the 1st of January 2024, and Linnaeus University in Småland became the twelfth partner university on the 1st of July, while the original ten partner universities have been on board since the start of NAISS in 2023.

Branch agreements for NAISS branches at the KTH Royal Institute of Technology, Linköping University, Lund University and Umeå University have now been signed. Chalmers and NAISS are practically in agreement and some more negotiations will be held with Uppsala University. There have also been signals from one of the entry-level NAISS partners that they would like to level up and start a NAISS branch in the near future.

Build Systems Course and Hackathon

Johan Hellsvik, PDC

In October this year, PDC hosted a combined build systems course and hackathon. The event was organised in collaboration with the EuroCC National Competence Centre Sweden (ENCCS), the CodeRefinery project, the CSC – IT Center for Science, Finland, and the Aalto Scientific Computing research infrastructure, Finland.

Modern computer hardware has developed towards heterogeneous architectures that combine the compute capabilities of multicore central processing units (CPUs) and graphics processing units (GPUs), non-uniform memory access (NUMA), and high speed network interconnects. On the one hand, this offers the prospect of extreme computing power, and, on the other hand, it definitely presents a challenge for developers and users of scientific software. To overcome the challenges of adapting existing code (or writing new code) to run efficiently on heterogeneous systems, it is important to be able to build and install programs and libraries of code on these systems. There are various tools and frameworks available to help with these tasks. The scope of the course was to provide an overview and hands-on experience of commonly used tools and frameworks for building scientific software on Linux-based computers, with an emphasis on code written in compiled languages.

The course part of the event ran over four half days with a curriculum featuring an introduction to compilers, linkers, and libraries, the GNU make and autotools, CMake, Spack, EasyBuild, and Singularity/Apptainer. In the exercise sessions, the participants got the chance to try out these

tools for a variety of open source code programs. The hackathon was at first intended as an onsite event over three days, but was later rescheduled to become an online activity with one participating team, with which we are working to get a multiple-program workflow to be able to run in sync as multinode jobs on the Dardel supercomputer.

BioExcel Activities

Alessandra Villa & Rossen Apostolov, PDC

Summer School

The BioExcel Centre of Excellence for Computational Biomolecular Research (BioExcel CoE) has been actively involved in supporting the biomolecular simulation community for many years now. The BioExcel Summer School, which targets early-stage researchers, is one of the centre's most successful initiatives. To get a feel for the school, see the video: https://www.youtube.com/watch?v=BTB1J4S_cAU.

The school has been run very successfully multiple times between 2018 and 2024, and, as per tradition, it took place in Sardinia, Italy, as an on-site event at the research hub Sardegna Ricerche (also known as the Sardinian Agency for Research and Technology Development). This year's summer school (https://bioexcel.eu/events/bioexcel-summer-school-on-biomolecular-simulations-2024) covered topics like molecular dynamics simulations, biomolecular docking, free energy calculations, advanced sampling methods, and BioExcel Building Blocks (BioBB). This year we also had guest lectures on quantum mechanics/molecular mechanics and AlphaFold. In addition, there were poster and career sessions.

This year there were thirty participants at the school who represented fifteen European countries and seventeen nationalities. The lecturers and participants were asked to introduce themselves on the BioExcel forum well in advance to break the ice and promote future interactions. To aid the poster-related discussion and future collaboration, we asked for the posters to be uploaded with their abstracts on the forum and we had flash talks, where the participants were asked to introduce themselves and pitch their posters. The school received an excellent average overall rating. The next BioExcel Summer School will run from 8-13 June 2025 (https://bioexcel.eu/events/summer-school-on-biomolecular-simulations-2025). See you in Sardinia next June!

2nd BioExcel Conference

The 2nd BioExcel Conference on Advances in Biomolecular Simulations was held in the beautiful city of Brno, Czech Republic, from 20–23 October. The biomolecular research community gathered to explore the latest trends, updates and challenges in the fields of integrative modelling, free energy and drug design, workflows, automation and data integration. The line-up of invited speakers, all top scientists in their respective fields, provided a comprehensive programme. Students had an opportunity to showcase their work during three poster sessions. The local organising partner, Scientifika, provided excellent facilities and gourmet catering and ensured the event ran smoothly. The conference culminated with a mesmerising presentation by Rommie Amaro in the local observatory, projecting striking biomolecular machinery images on the large ceiling. All the participants, students, speakers and guests told us over and over that: "This was the best conference I've ever been to!".

BioExcel Initiatives

BioExcel has run a lot of initiatives to support the biomolecular simulation community and PDC has been a leader in many of them, including the Summer School and Ambassador Program.

Ambassador Program

The BioExcel Ambassador Program (https://bioexcel.eu/about/governance/ambassador-program) aims to establish strong connections with the local communities in EuroHPC and increase the uptake of high-performance computing (HPC) in general and BioExcel tools and expertise. One of the activities within the ambassador program is co-organising BioExcel workshops and training events tuned to the needs of the local communities. For this we have built a training portfolio (https://bioexcel.

eu/services/training) and developed protocols to organise workshops in collaboration with national competence centres (NCCs) and local ambassadors. The next ambassador-driven workshop will take place in Coimbra in November (https://bioexcel.eu/events/bioexcel-eurocc-workshop-in-portugal).

BioExcel Webinars

The BioExcel webinar series (https://bioexcel.eu/category/webinar) features notable developments in the field of computational biomolecular research. Have a look at the current playlist here: https://www.youtube.com/playlist?list=PLzLqYW5ci-2fanl2RtYHyFvjftu_M4Jwy.

BioExcel Forums

The BioExcel forums (https://ask.bioexcel.eu and https://gromacs.bioexcel.eu) are provided to support the research community.

Competency Hub

The BioExcel Competency Hub (https://competency.ebi.ac.uk/framework/bioexcel/3.0) is available for browsing competencies, career profiles and training resources to advance your career.

GROMACS Performance Optimisation on AMD GPUs

Szilárd Páll, PDC, & Andrey Alekseenko, KTH/SciLifeLab

GROMACS, which is a widely used molecular dynamics (MD) simulation engine, has seen significant performance gains on AMD GPU-based heterogeneous systems (like the Dardel system at PDC and LUMI in Finland) thanks to optimisations made to the SYCL backend. This builds upon our previous work enabling GROMACS on AMD GPUs.

The packed math instruction optimisations have been undertaken to better leverage the single-precision floating-point hardware capabilities of the AMD MI250X GPUs. While these should ideally be transparently generated by the compiler, manual optimisations have been added to GROMACS 2024 to make sure that performance is not dependent on optimal code generation.

A key focus area has been SYCL runtime performance. We identified that the way the SYCL runtimes handle task launches was hindering GROMACS's highly latency-sensitive scheduling. To address this, collaborations with developers of AdaptiveCpp and oneAPI DPC++ runtimes aimed to optimise performance. Additionally, we actively participated in the design of the next SYCL standard to incorporate these findings.

GROMACS primarily relies on AdaptiveCpp for AMD GPUs. Improving its performance was crucial. By default, AdaptiveCpp uses a deferred task launch strategy, caching SYCL API calls for later submission, allowing for runtime optimisation but introducing significant latency. This is detrimental to GROMACS's latency-sensitive MD engine. Delays in submitting tasks lead to GPU starvation, particularly when CPU operations (like MPI calls) require CPU-GPU synchronisation.

To address low-latency submissions, the AdaptiveCpp 23.10 release introduced an instant submission mode, improving performance by up to 22% compared to the optimised cached mode. To benefit from this, users only need to recompile GROMACS with a recent AdaptiveCpp version enabling instant submission; there is no need to change the GROMACS code. This allows users of earlier GROMACS releases to gain performance improvements as well.

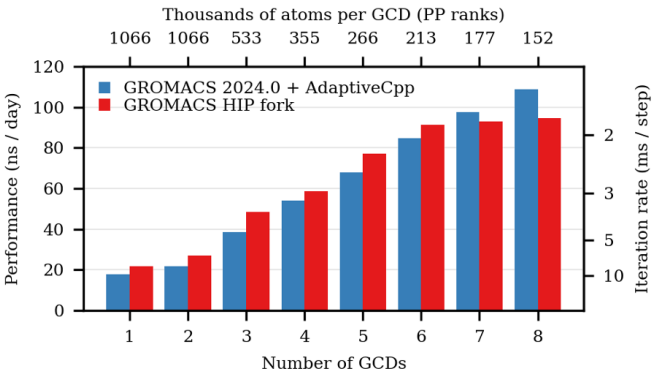

Above: The graph shows the application performance (ns/day) and the corresponding iteration rate (ms/step) for the STMV benchmark running on different numbers of GCDs within a single node. (An MI250X GPU consists of two GCDs.) The performance of GROMACS 2024.0 using AdaptiveCpp 23.10.0 with instant submission mode is compared to that of AMD's GROMACS HIP fork.

Our long-term work on readying GROMACS for AMD GPUs using SYCL was presented at the 2024 Cray User Group Conference (see [1]). The paper provides a detailed analysis of node-level kernel and runtime performance, sharing best practices for using SYCL as a performance-portable GPU framework within the high-performance computing (HPC) community. Performance demonstrations are provided for Cray EX235a machines with MI250X accelerators, illustrating that portability can be achieved without sacrificing significant performance.

References

1. A. Alekseenko, S. Páll, and E. Lindahl, "GROMACS on AMD GPU-Based HPC Platforms: Using SYCL for Performance and Portability." arXiv, May 2, 2024. https://doi.org/10.48550/arXiv.2405.01420.

PDC-ENCCS Collaboration Seeking Industry Partner

Michaela Barth, PDC

As part of the EuroCC project, which leads the initiative to form and coordinate competence centres in high-performance computing (HPC) across Europe, the EuroCC National Competence Centre Sweden (ENCCS) provides free high-performance computing training and support for industry, academia and public administration. ENCSS is hosted by the RISE Research Institutes of Sweden and Linköping University with funding coming from the European High-Performance Computing Joint Undertaking (EuroHPC JU), the Swedish Research Council and the Swedish Innovation Agency (Vinnova).

Since May 2023, PDC and ENCCS have been investigating whether they could work together to improve the advanced HPC training they each offer, especially in terms of supporting industry in its use of HPC resources. In June this year, ENCCS and PDC proposed a strategic collaboration focusing on supporting business HPC users and offering more advanced types of training tailored to industry research needs. The decision to proceed with the collaboration was made in August.

PDC offers HPC access to business customers according to a pay-per-use model (see https://www.pdc.kth.se/hpc-services/prices-for-using-pdc-resources-1.932289). Compared to standard cloud computing offerings for industry, using the HPC resources at PDC that are dedicated for industry is generally cheaper, since PDC's pricing reflects the actual mean cost and energy consumption of using the resources. (As part of a government agency, PDC does not make any profit, nor is it allowed to subsidise costs.) The goal of the PDC pricing model has always been to be easy to understand, predictable and without any hidden costs.

PDC aims to foster the best possible climate to enable top-of-the-line industrial research and development (R&D) while, at the same time, working together with our industrial partners to increase their and our competence in leading industrial software and to keep each other updated on trends, as well as developing future use cases and associated requirements. Current business collaborators appreciate the high level of professional competence at PDC, both when it comes to application software and making the best possible use of the underlying hardware. Some of our industrial collaborators have worked for years on different generations of PDC's resources building up mutual trust. Often, the allocated resources at PDC are like a flexible extension of the company's resources as they are used for burst-type computing when in-house resources reach their limits ahead of critical deadlines, but they are also used when developing and testing completely new R&D scenarios and simulations further away from regular day-to-day production and standard R&D. Here the usage of the Dardel resources provides a new sense of freedom and the possibility to explore new, future-relevant research areas without disturbing the core business computational tasks.

Another selling factor for choosing PDC, which has gained more weight recently, is the drive towards sustainability in business. PDC has been a trail-blazer in re-using heat from supercomputers since 2009. Instead of using about 30% more energy (and money) to cool the Dardel systems,

both the air and liquids in the system are cooled down via heat exchangers and connected to the KTH heat pump (which was installed in 2015) in order to heat the surrounding KTH campus buildings. Together with KTH embracing the UN Sustainable Development Goals, PDC's users can sleep soundly knowing that 100% of the electricity that PDC uses is from renewable sources, and thus all businesses buying computing time from PDC already fulfil future industry requirements on sustainability. The Dardel GPU partition is seventeenth in the world on the November 2024 Green500 list of the most energy-efficient systems and is the highest-ranked system in Sweden on the list.

ENCCS does a marvellous job of supporting Swedish industry, especially in onboarding newcomers and assisting them in making efficient use of the available EuroHPC top-tier computational resources. ENCCS offers free HPC consultancy and training, assistance in writing proposals to apply for access to EuroHPC systems and in-depth support for proof-of-concept projects. As a one-stop shop for HPC in Sweden, ENCCS also connects companies to HPC and artificial intelligence (AI) service providers and funding agencies in the EuroHPC ecosystem. Industry researchers can use EuroHPC resources free of charge for open science after application. However, for computational work which cannot be classified as open science or where it is crucial to have a guarantee of being able to continue using the computational resource for the full lifetime of a particular research project, PDC may be the better alternative.

As mentioned in the article about *"ENCCS news"*, PDC and ENCSS are looking for an industrial partner for an in-depth pilot project. The idea of the pilot is for the company to undertake a research project where PDC and ENCCS provide free HPC resources and also support (for example, from PDC/ENCCS research software engineers and application experts) to use those HPC resources. At the same time, the company will assist PDC and ENCCS to determine how the HPC services that they currently offer to support industry research could best be developed to meet the current and ever-evolving needs of industry. If your company is interested in this exciting opportunity, please contact PDC and ENCCS at info@enccs.se and mention that you are interested in the pilot project.

PDC-Related Events

BioExcel Workshop: Balkan Edition
21-22 May 2025, Sofia, Bulgaria
Details will be available soon at https://bioexcel.eu/events.

BioExcel Summer School
8-13 June 2025, Sardinia, Italy
https://bioexcel.eu/events/summer-school-on-biomolecular-simulations-2025

HPC Sources

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

HPC in Sweden & Scandinavia

- ENCCS
  http://enccs.se
- 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
- PerMedCoE
  https://permedcoe.eu
- Plasma-PEPSC CoE
  https://plasma-pepsc.eu

HPC news sources

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