no.2 2020

PDC Center for High Performance Computing

PDC Newsletter

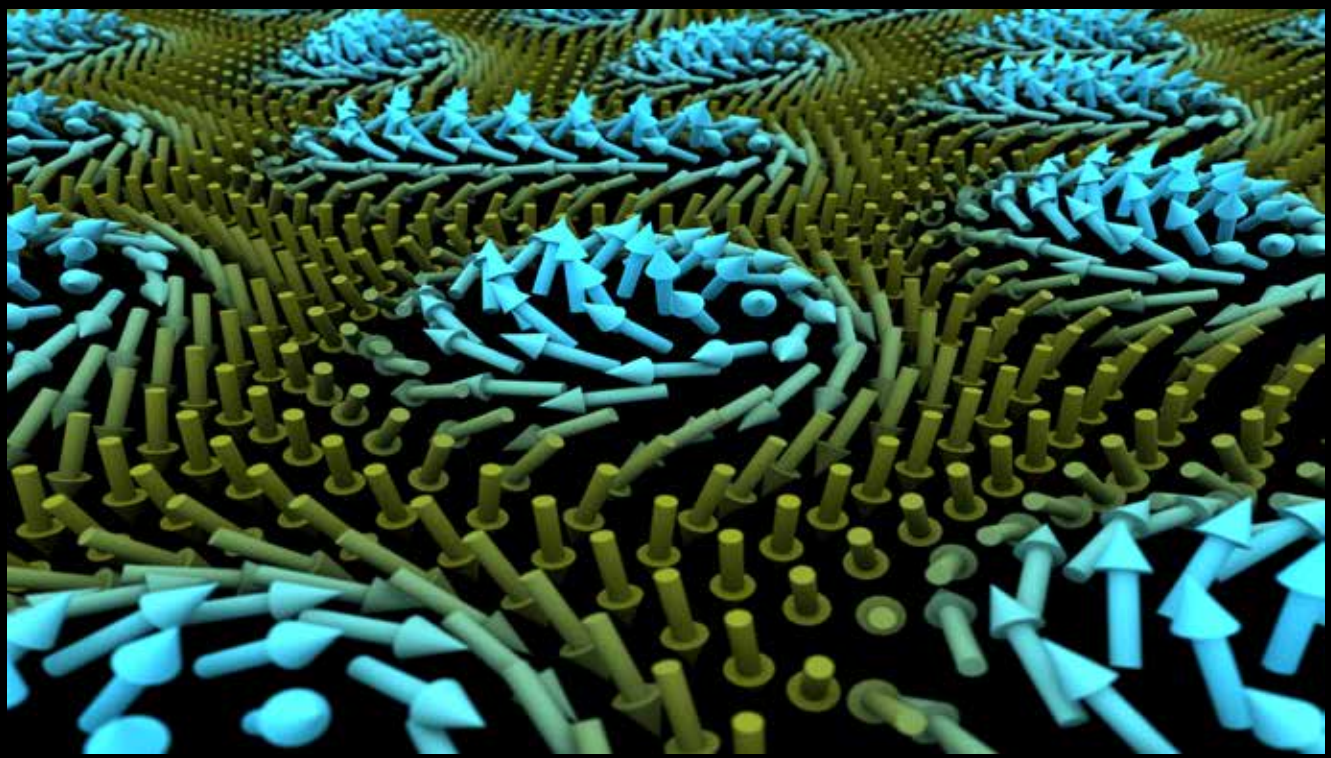

The PDC Newsletter is published by PDC at EECS, KTH.

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

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

ISSN 1401-9671

Cover

The cover figure shows a lattice of magnetic skyrmions as obtained from atomistic spin dynamics simulations. Magnetic skyrmions are particle-like topologically-protected solitons and have been touted as a possible basis for future low-energy magnetic memory applications.

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

This will be my second and also last editorial for the PDC Newsletter as KTH has recruited a professor and new director for PDC who has a European high-performance computing (HPC) research profile. Given Europe's recent and ongoing advancements in the field of HPC, it is of strategic importance for PDC and Sweden alike to be represented on the international arena and, not least for this reason, it will be with pleasure that I hand over the baton to the next director in line by early next year.

But let us not get ahead of ourselves as there are currently important matters to conclude before we reach the holidays and the end of the year. First and foremost, PDC has reached the final stage of the complex procurement process to find a replacement for the Beskow cluster. By the time you read this newsletter, we will have awarded the contract to the winning vendor and started the planning for the system installation. I can already say, however, that the procurement process has been professionally driven forward by project leader Gert Svensson to a point where it is clear that, no matter which vendor's bid wins the competition, we in the applied HPC research community can look forward to a system that will enable our research to achieve even higher levels of excellence. This newsletter contains a report by Gert on the current state of affairs in the procurement process.

It comes as no surprise that the type of hardware-driven scientific advancements that we are presently seeing around the world do not come without effort. We are witnessing a paradigm shift in HPC hardware solutions that fundamentally impacts our algorithmic design strategies and the accompanying software implementations. If we take my own specific research field of chemical sciences as an example, it is noted that the U.S. Department of Energy (DOE) announced plans in early November to provide up to 32 million USD for the development of what they refer to as "sophisticated software", which is intended to take advantage of the rapidly advancing supercomputing capabilities of the DOE National Laboratories. Another related example is the formation of the Center for Accelerated Application Readiness (CAAR) at the Oak Ridge Leadership Computing Facility with very similar goals but for a wider target.

For Swedish researchers to remain at the forefronts of our respective research fields, Sweden must take similar action and facilitate the required software transformation in e-sciences. Here I believe that the application experts at the Swedish supercomputer centres can and should be allowed to play a central role, which requires long-term commitments and close collaboration with the teams of software developers that exist at our universities. PDC has been acting accordingly and has formed these types of strong bonds with the teams developing GROMACS (molecular dynamics), VeloxChem (quantum chemistry), and Nek5000 (fluid dynamics). Another welcome instrument has emerged from Sweden's participation in the EuroHPC Joint Undertaking, namely the EuroCC National Competence Centre Sweden (or ENCCS). An overview of the ENCCS and the upcoming events that the centre is organising is presented in this newsletter.

PDC is also involved in the European Centre of Excellence for Engineering Applications (EXCELLERAT) project and, in one of the two main articles in this newsletter, we are given an insight into the work aimed at the preparation of an in situ instrumentation for the Nek5000 program to perform the time evolution analysis of coherent structures. The cover article is concerned with the microscopic understanding of magnetic materials by combining first-principles density functional theory (DFT) and the Heisenberg Hamiltonian. The resulting multiscale approach is presented and an outlook for the importance of spin dynamical phenomena in future technologies is provided. Enjoy reading about these developments!

With that, it is time for me to close this editorial and thank you for all the exciting research that you have produced during 2020 with the resources at PDC. It is your work that makes it worthwhile for us to try to provide the best possible services with the means available. Conversely and, in the spirit of the upcoming holidays, I encourage you to give a thought to the PDC staff members who work hard, and often outside office hours, to enable your research. Myself, with one foot in each camp, I address all of you in a collective: Merry Christmas and a Happy New Year!!!

Patrick Norman, Interim Director PDC

Save the Date

In This Issue

Staff Focus

Arash Alizad Banaei

Arash Alizad Banaei has recently joined PDC as an application expert in computational fluid dynamics (CFD). Arash finished his PhD in fluid mechanics at KTH last year, and before that he had completed both a master's and a bachelor's degree in fluid mechanics. Arash has been developing several high-performance CFD tools for simulating multiphase flows and in particular for a dispersed elastic solid phase in fluid flows. He has also worked with a number of commercial CFD tools that are used extensively in the industry. He believes that to have an efficient CFD code, both a proper numerical method and a good parallel computing algorithm is needed and that combining these two is the real art of being a CFD application expert.

Arash will be assisting PDC users with the CFD tools and software available on the PDC systems. He will also be doing research on improving the computational efficiency of CFD codes, as well as sharing his knowledge with other researchers who attend PDC courses and social meetings. In his spare time, Arash enjoys cooking and working out. He also plays floorball at KTH once a week.

Dynamics of Magnetic Materials

Anders Bergman, Anna Delin[1,2], Manuel Pereiro, Erik Sjöqvist, Danny Thonig[3], Olle Eriksson[3], Department of Physics and Astronomy, Uppsala University

Introduction

Magnetism and magnetic excitations of solids have traditionally been analyzed by two seemingly different approaches: electronic structure theory based on density functional theory [1,2] and the Heisenberg Hamiltonian [3]. The former has been used to calculate (in ab initio fashion) the magnetic moments, the easy axis magnetization and the magnetic ordering [4] of solids, while the latter is typically used to analyze magnon (or spin wave) dispersions and to estimate ordering temperatures [5]. The Heisenberg Hamiltonian, and modifications of it, have also been adopted to provide microscopic understanding of exotic magnetic phenomena, such as spin-glass formation [6], as well as the Kosterlitz-Thouless transition [7]. Analyzing magnetic phenomena by means of the Heisenberg Hamiltonian has proven to be extremely successful, with the gnawing shortcoming of being a parametrized theory: the interactions used in this Hamiltonian have hitherto been fitted to experimental data or guessed.

These two seemingly disparate approaches have recently been merged into what may best be described as a multiscale approach for evaluating basically *all* magnetic phenomena from one common platform. This article describes the essentials of this approach, and provides numerical examples of calculating magnetic ground state properties (as well as results for the excited state) for a range of materials. The theoretical toolkits that have emerged from these considerations are now available in software [8,9] that is installed at all Swedish National Infrastructure for Computing (SNIC) supercomputer centres. The theory is therefore ready to tackle any challenge associated with new and exotic magnetic phenomena, such as those that will emerge from large scale facilities like the European Spallation Source (which is a multi-disciplinary research facility currently being constructed to provide unique information about the structure and properties of biology, chemistry, physics, and engineering materials by using what will be the world's most powerful pulsed neutron source).

Atomistic Spin Dynamics

The accuracy of Density Functional Theory (DFT) calculations and the efficiency of analyzing magnetism using the Heisenberg Hamiltonians can be combined with the atomistic spin dynamics method (ASD) [10]. As the name indicates, the quantum-mechanical spin density that originates from the distribution of the electrons in the material is treated on an atomic level. This is done by assuming a model where each atom i is associated with a local magnetic moment, described as a vector, $\boldsymbol{m}_i$. The interplay between the local moments is then described with the Heisenberg model. This model is based on pair-interactions, J_{ij}, between moments $\boldsymbol{m}_i$ and $\boldsymbol{m}_j$ on sites i and j, so that the energy can be written as

$$H = -\frac{1}{2}\sum_{i \neq j} J_{ij}\boldsymbol{m}_i\boldsymbol{m}_j$$

In this context, the pair-interactions are called exchange interactions since they stem from the energy associated with exchanging spins in a system, according to the Pauli exclusion principle. The Heisenberg model can be considered in a fully quantum mechanical fashion, where $\boldsymbol{m}_i$ represents a spin operator on site i, multiplied by the g-factor. Here we will adopt a classical (or semi-classical) approach, where $\boldsymbol{m}_i$ represents the magnetic moment at site i, and the moment can point in any direction. Assuming that the semi-classical Heisenberg model can provide an adequate approximation of a system, the critical part for obtaining a realistic and material-specific description is to determine the exchange interactions J_{ij} accurately. This is where the coupling to DFT calculations enters since there now exist several methodologies to calculate the needed set of exchange interactions for almost any material.

With a realistic Heisenberg model, the energy for any given magnetic configuration (that is, any ensemble of directions of the moments in the system) can be accurately calculated for large systems containing billions of atomic moments. This can be used in combination with finite-temperature Monte Carlo simulations to determine the magnetic ground state of the system, as well as phase transition temperatures (in other words, temperatures at which the system changes from an ordered to a disordered state).

In order to properly describe excited states of magnetic materials (that is, what happens to a magnet when an external field, a laser pulse, or an electric current is applied), one needs to describe the magnetization dynamics (which means the time dependence of the magnetic moments). As has been shown – both from equivalents to phenomenological theories and rigorous derivations from first principles – the local magnetic moments follow an equation of motion that depends on the effective magnetic field that acts on each moment. The equation of motion for the moments can be written as

$$\frac{d\boldsymbol{m}_i}{dt} = -\gamma \boldsymbol{m}_i \times \boldsymbol{B}_i - \alpha \frac{\gamma\, \boldsymbol{m}_i \times (\boldsymbol{m}_i \times \boldsymbol{B}_i)}{|\boldsymbol{m}_i|}$$

which is an atomistic version of the Landau-Lifshitz-Gilbert (LLG) equation. Normally it is formulated on a macro-spin level, and is then used to describe the magnetization on a more coarse-grained level. In the atomistic LLG equation, the gyromagnetic ratio γ and the Gilbert damping α also come into play, where the latter determines the rate of dissipation of energy and angular momentum. The remaining quantity to be determined is the effective magnetic field $\boldsymbol{B}_i$ which can be obtained from the Hamiltonian as $\boldsymbol{B}_i = -\partial H/(\partial \boldsymbol{m}_i)$. Temperature effects can be included in the magnetization dynamics by means of Langevin dynamics which are obtained by adding temperature-dependent stochastic contributions to the effective field.

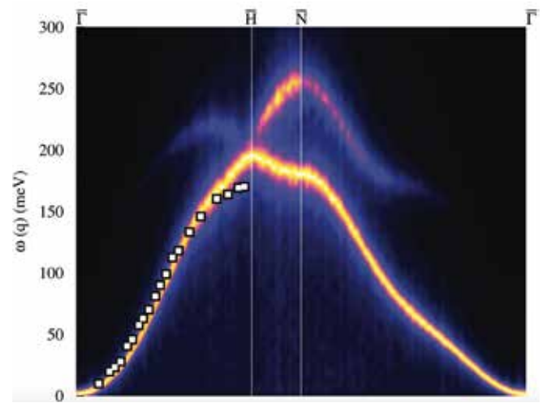

Above: Spin wave dispersion spectra obtained from ASD simulations of 2 ML Fe/W(110) at T = 300K and α = 0.01. Experimental values are marked by white squares (see [16]).

[1] Also at Department of Applied Physics, School of Engineering Sciences, KTH Royal Institute of Technology

[2] Also at Swedish e-Science Research Centre (SeRC), KTH Royal Institute of Technology

[3] Also at School of Science and Technology, Örebro University

These atomistic spin dynamics have been implemented, together with a collection of Monte Carlo routines and related functionalities, in the code UppASD, available at [9], which runs readily on high-performance computing (HPC) resources available at PDC and other SNIC centres. The algorithms have been parallelized using both OpenMP and MPI, and can also run efficiently on GPU resources. However, due to the efficient formulation and implementation of the method, it is possible to run the majority of simulations on a single node/workstation.

The ASD framework can be used to extract a multitude of relevant physical properties for magnetic systems. As mentioned earlier, the critical temperature where a phase transition will occur is one of these central properties. Studies of real-time dynamics can also be done with ASD. Notable examples of this include simulations of complex ultrafast demagnetization processes that occur when materials are exposed to high-fluency laser pulses and studies of the dynamics of topological magnetic solitons in the presence of electric currents and magnetic fields. In this article, we highlight a few examples of previous and ongoing studies of relevant magnetization dynamics where UppASD has been employed: magnon and skyrmion dynamics.

Magnetic Excitations

In recent years, the magnetism of ultrathin films on certain substrates has become experimentally accessible, in particular, by means of spin-polarized electron energy loss spectroscopy (SPEELS) or scanning-tunnelling microscopes (STM). This has been an important development, since it provides the means to extract magnetic excitation energies from nano-sized samples, which more traditional methods (like neutron scattering) cannot do. Examples of such ultrathin films are a single atomic layer (monolayer or ML) of Co on Cu(001) [11], Fe on Cu(001) [12] or multiple layers (1-5 ML) of Fe on W(110)[13]. Experimental accessibility to these systems has been opening up new views on whether the semi-classical atomistic model discussed earlier is accurate enough to describe

Below: The plot below shows calculated spin wave asymmetry for the magnon spectrum of 2 ML Fe/W(110), using exchange interactions calculated from DFT. The experimental values were obtained by Zakeri et al. (see [17]).

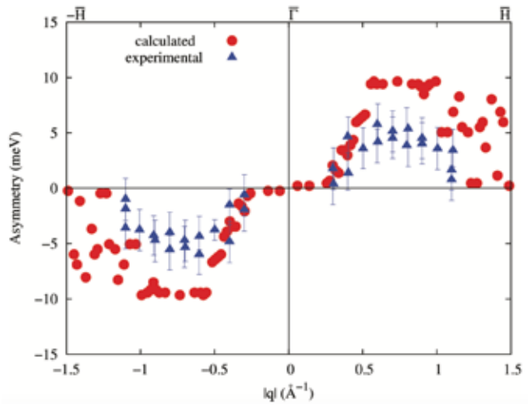

the experimental trends or to even predict magnetism in different materials.

Magnetism in ultrathin magnetic films can show exotic ground states arising from finite temperature and relativistic spin-orbit coupling effects. These effects also cause unexpected dynamics of the atomistic magnetic moments and an anisotropic magnon propagation. Magnons are collective, quasi-particle wave excitations of the atomistic magnetic moments [10]. They carry a certain energy ω related to the precession frequency of the individual spins in a spin wave, as well as momentum q related to the wave length and the propagation direction of the magnon. The relation between the energy and momentum of a magnon, known as magnon dispersion $\omega(q)$, can be measured in the SPEELS experiments mentioned previously and is a measure of the exchange interactions between the atomic magnetic moments. The magnon dispersion is, coincidentally, one measurable for which ASD is perhaps the best-suited theoretical tool. As a showcase, we have studied magnons in 2 ML Fe on W(110) [14] (see figure on the previous page) and in the lanthanide Gd [15] (see figure above) and achieved not only excellent qualitative, but also quantitative, agreement with experimental results.

For both of our examples, the interactions between the atomistic spins were calculated from first-principle DFT methods, where the only

Below: Spin wave dispersion spectra (black curve) obtained from ASD simulations of Gd at T = 300K and α = 0.01. Experimental values are marked by red dots (see [18]).

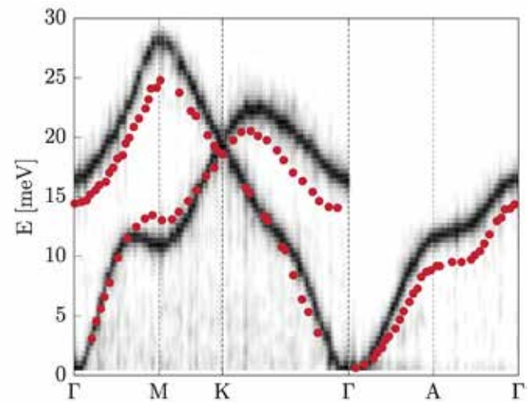

parameter to the simulation is the lattice constant of the crystal structure. When these interactions were included in our spin dynamics simulations, we observed a softening of the magnons for 2 ML Fe/W(110) with respect to the Fe-bulk spin wave dispersion. This is the result of weaker interactions between the atomic magnetic moments due to a lower level of coordination at the surface. Furthermore, our calculations revealed that spatial-anisotropic interactions (the so-called Dzyaloshinskii-Moriya interactions) are present in 2 ML Fe/W(110) and that they cause an asymmetry, in other words, ω(q) differs from ω(-q) around the Γ-point in the magnon dispersion, which represents the case where all spins have coherent precession. This asymmetry depends strongly on the magnon momentum (shown in the figure on the previous page), which was first measured in SPEELS experiments [17] (blue dots) and later confirmed by our model (red dots).

As a second example, we focused on the magnetism of lanthanides, with gadolinium (Gd) as a representative of this class. Lanthanides are particular due to the localization and correlation of the electrons that are responsible for the magnetism (the 4f shell). As a consequence, the interaction between the magnetic moments is weak since it is related to the small wave function overlap of the 4f shell of two separate lanthanide atoms. This makes it demanding to calculate with first principle methods, due to the high computational precision needed to calculate the interaction and hence the respective magnon dispersion (see the figure to the left). The dispersion is seen to follow a parabolic dependence. This is clear evidence of the colinear ferromagnetic ordering, which is intrinsic to Gd at low temperatures. As in our previous example, we obtained excellent agreement between experiment (red dots) and theory (gray background), which highlights the high quality of the interactions obtained between the magnetic moments in Gd, as well as the accuracy in extracting the magnon dispersions from ASD simulations.

Topological Magnets

The presence of defects or other symmetry-breaking inclusions can strongly influence dynamical properties at both microscopic and macroscopic length scales. The ASD framework is ideally suited for resolving the magnetization at an atomic level but in order to access more mesoscopic length scales (which here means little more than a micrometer scale), one needs to include another level of multiscale simulations. We have recently implemented a multiscale method [19] capable of bridging atomic and micrometer length scales. This technique is capable of solving the Landau-Lifshitz-Gilbert equations efficiently in two regions of a magnetic material – the mesoscopic and the atomistic regions, which are coupled in a seamless way. The method has initially been applied to studies of topological magnets where magnetic whirls, called skyrmions, can occur. Skyrmions have many intriguing properties, including ideally topological protection and a particle-like behaviour, and have been touted as a viable basis for future data processing technologies. Our multiscale study on the transport behaviour of skyrmions in the presence of defects is best presented in the video at https://play.kth.se/media/PDC+Newsletter+2020+No.+2A+Dynamics+of+Magnetic+Materials+-+supplementary+video/0_wxolmrm2.

Quantum Information

A central concept in quantum information science and technology is the possibility to create and manipulate quantum entanglement. This purely quantum-mechanical feature implies

stronger-than-classical correlations that can persist over long distances and even in noisy environments, such as in the case of excitation energy transfer in biological systems. It is used as a resource in quantum communication and is a key ingredient for the quantum parallelism that is used in quantum computers.

Driven by potential applications for information technology, there is a growing recent interest in finding methods to prepare, control, and verify entanglement amongst collective modes in the solid state. In particular, magnons can show entanglement and can also entangle with other degrees of freedom, such as phonons and photons.

A long-term goal of our research is to explore how magnonic degrees of freedom can be used in hybrid settings, involving other quantum-mechanical degrees of freedom, for building quantum networks for quantum information processors and quantum communication devices. In particular, we have in mind a new form of hardware for quantum computation in materials – a quantum version of "magnonics" – where entangling logical gates between localized qubits can be mediated by means of entangled magnon-pairs.

Antiferromagnets are candidate systems for magnon entanglement, as their magnetic spin lattice naturally divides into a pair of sublattices. In a paper by Azimi Mousolou et al. [20], it has recently been demonstrated that these sublattices are indeed entangled, even in the vacuum state where no magnons are present. This mode entanglement is a quantifiable resource (see the figure to the right) that can be very large, in particular for long-wavelength magnon modes. To show this, we have considered a generic lattice of spins with Heisenberg and Dzyaloshinskii-Moriya interaction, as well as anisotropy terms, for which an intricate layer of entanglement between the two sublattices can be seen. This intrinsic entanglement of the antiferromagnet can potentially be tested by letting the collective spin modes interact with an external source of, for example, photons or photon polarization, which can be measured by techniques known in quantum optics. In this way, the magnon

Below: The entropy of entanglement E_o^{ab} between the two sublattices of an antiferromagnet with nearest-neighbour Heisenberg interaction. For this system, E_o^{ab} depends only on a parameter $|\gamma(k)|$ that is determined by the geometry of the spin lattice, but is independent of the spin value at each lattice site. The different curves in the main plot correspond to zero (black curve), one (brown), and two (orange and blue) magnons. The inset depicts the magnon dispersion of $SrMnO_3$ for a selected path of q along high-symmetry directions of the Brillouin zone. The width of the curve depicts the entropy of entanglement. (The picture is taken from [20].)

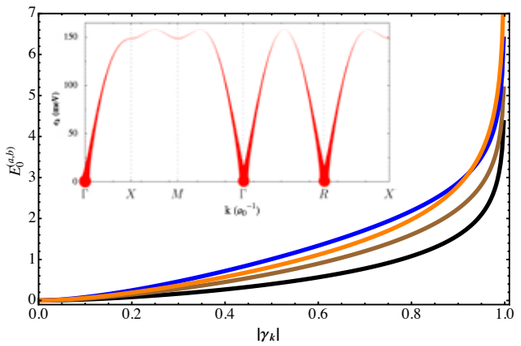

mode entanglement can not only be tested experimentally, but can also act as a resource for photon entanglement. This particular study was performed using analytical methods based on the Heisenberg model but ASD studies on these materials are forthcoming.

Outlook

Looking forward, we believe that spin dynamical phenomena will come to play an important role in the development of several key future hardware technologies for computing. The time-integrated amount of stored information is doubling roughly every eighteen months and, since the majority of the world's information is stored in magnetic media, the possibility to write information to (and retrieve information from) a magnetic material at ever greater speed, and with lower energy consumption, has obvious benefits for our society. Hence the seemingly simple switching of a magnetic unit – a magnetic bit – is a crucial process which defines how efficiently information can be stored and retrieved from a magnetic memory. From an application point of view, it is apparent that it is advantageous to be able to switch the magnetization of a bit as fast as possible while minimizing energy losses. All-optical control of magnetism has for this reason become an intense research field, with reports on remarkably fast switching times. Surprisingly, the mechanism behind the ultrafast switching is still unknown. Our recent work that connects the spin motion and lattice vibrations [21] aims to answer this question.

Topological magnets such as skyrmions have also been mentioned as promising for future magnetic memories and in addition they have other properties which make them suitable for related processes. Neuromorphic computing, where features of the human brain are used as inspiration for efficient data processing, is currently becoming an increasingly active field of research. Magnetization dynamics play a central role in spintronic-based neuromorphic computing [22]. Here, magnetic topological textures such as skyrmions or domain walls take the role as neurons and magnetic tunnel junctions can function as both synapses and neurons.

As indicated in the Engtanglement section earlier, quantum entanglement in magnetic materials can be used to link several small quantum processors for large-scale quantum computing and also for interfacing a quantum computer, for example, arrays of spins or ions, with optics for long-distance quantum communication [23].

References

1. P. Hohenberg, W. Kohn, Phys. Rev. 136, B864-871 (1964)
2. W. Kohn, L. Sham, Phys. Rev. 140 A1133-1138 (1965)
3. See e.g. N. Ashcroft and D. Mermin, Solid State Physics, Saunders College Publishing (1976)
4. J. Kubler, Theory of itinerant electron magnetism, Oxford Science Publications (2000)
5. D. Landau, K. Binder, A guide to Monte Carlo simulations in Statistical Physics, Cambridge University Press (2005)
6. G. Parisi, P. Urbani, F. Zamponi Theory of Simple Glasses. Cambridge University Press (2020)
7. J.M. Kosterlitz, D.J. Thouless, Ordering, metastability and phase transitions in two-dimensional systems. Journal of Physics C: Solid State Physics. 6 (7): 1181–1203 (1972)
8. https://www.physics.uu.se/research/materials-theory/ongoing-research/code-development/rspt-main
9. https://www.physics.uu.se/uppasd, https://github.com/UppASD/UppASD
10. O. Eriksson, A. Bergman, L. Bergqvist and J. Hellsvik, Atomistic Spin-Dynamics, fundamentals and applications (Oxford university press, 2017)
11. R. Vollmer et al., Phys. Rev. Lett. 91, 147201 (2003)
12. S. Achilli et al., J. Phys. Cond. Matt. 19 305021 (2007)
13. J. Prokop et al., Phys. Rev. Lett. 102, 177206 (2009)
14. L. Bergqvist, A. Taroni, A. Bergman, C. Etz, and O. Eriksson, Phys. Rev. B 87 144401 (2013)
15. I. L. M. Locht, Y. O. Kvashnin, D. C. M. Rodrigues, M. Pereiro, A. Bergman, L. Bergqvist, A. I. Lichtenstein, M. I. Katsnelson, A. Delin, A. B. Klautau, B. Johansson, I. Di Marco, and O. Eriksson, Phys. Rev. B 94 085137 (2016)
16. W. X. Tang et al., Phys. Rev. Lett. 99, 087202 (2007)
17. Kh. Zakeri et al., Phys. Rev. Lett. 104, 137203 (2010)
18. J. Jensen and A. R. Mackintosh, Rare Earth Magnetism (Clarendon Press, Oxford, U.K., (1991)
19. E. Mendez, M. Poluektov, G. Kreiss, O. Eriksson, and M. Pereiro. Phys. Rev. Research 2, 013092 (2020)
20. V. Azimi Mousolou, A. Bagrov, A. Bergman, A. Delin, O. Eriksson, Y. Liu, M. Pereiro, D. Thonig, and E. Sjöqvist, https://arxiv.org/abs/2006.03479 (2020)
21. J. Hellsvik, D. Thonig, K. Modin, D. Iuşan, A. Bergman, O. Eriksson, L. Bergqvist, and A. Delin, Phys. Rev. B 99, 104302 (2019)
22. S. Bose, PRL 91, 207901 (2003)
23. J. Grollier et al., arXiv:2007.06092 (2020)

Staff Focus

Johan Hellsvik

Johan Hellsvik has a background in theoretical condensed matter physics working on topics such as spin dynamics, multiferroic materials, ultrafast magnetization dynamics, and frustrated magnetism. For the last three years he has been working at Nordita and KTH within The Swedish QuEST for BIFROST and Novel Quantum Materials collaboration, performing modelling of magnetic excitations

…continued on page 11

Alvis System Now Available for AI/ML Research

Thomas Svedberg, C3SE

Alvis is a new system for Artificial Intelligence and Machine Learning (AI/ML) research which is located at the Chalmers Centre for Computational Science and Engineering (C3SE). C3SE is hosted at Chalmers University of Technology in Gothenburg and is one of the six Swedish National Infrastructure for Computing (SNIC) centres for scientific and technical computing. The new Alvis system is a part of SNIC and is financed by SNIC together with the Wallenberg Artificial Intelligence, Autonomous Systems and Software Program (WASP).

Researchers from Swedish universities and research institutes may use Alvis. If you are interested in using Alvis in your research, see the instructions on the C3SE website about how to apply for access: https://www.c3se.chalmers.se/documentation/getting_access, or see the SNIC AI/ML rounds page on the SNIC User and Project Repository (SUPR) website: https://supr.snic.se/round. Note that, as with other SNIC resources, researchers need to be a member of a corresponding SNIC project in order to use the system.

Alvis is based on Graphical Processing Unit (GPU) accelerator cards, and the system contains several types of compute nodes, all of which have multiple NVIDIA GPUs. The system is being installed in two phases. Phase I of the project is an initial system that was installed at C3SE during the summer of 2020. Phase I of Alvis became fully operational at the end of August, after a preliminary test period. Applications for research projects to use Phase 1 of Alvis opened in mid-August 2020. SNIC has already been receiving applications for projects to use Alvis. Since Alvis is the only system within SNIC that is dedicated solely for AI/ML research, C3SE expects to see a lot more applications from AI/ML researchers wanting to use Alvis!

Below: Alvis system

Phase II will actually be the major part of Alvis – this part of the system will be designed and built based on the experiences (and resultant recommendations) from operating and using Phase I of the system. The plan is that Alvis will be extended early in 2021, so Phase II is expected to be in operation by the end of summer next year.

Phase I of the Alvis system consists of a login node, a group of high-performance GPU compute nodes (Phase Ia) and a group of capacity GPU compute nodes (Phase Ib). Phase II is under discussion at the moment. It will most likely be in line with Phase I and consist of a group of capacity nodes and one (or several) groups of performance nodes, as it has been observed that different fields of AI/ML research and research using AI/ML methods have varying needs in terms of the types of nodes. The specifications of the Alvis Phase I nodes are as follows.

Login node

- 4 × NVIDIA Tesla T4 GPUs with 16GB RAM
- 2 × 16 core Intel(R) Xeon(R) Gold 6226R CPU @ 2.90GHz (total 32 cores)
- 768GB DDR4 RAM

Phase Ia

12 high-performance GPU compute nodes with

- 2 × NVIDIA Tesla V100 SXM2 GPUs with 32GB RAM, connected by NVLink
- 2 × 8 core Intel(R) Xeon(R) Gold 6244 CPU @ 3.60GHz (total 16 cores)
- 768GB DDR4 RAM

5 high-performance GPU compute nodes with

- 4 × NVIDIA Tesla V100 SXM2 GPUs with 32GB RAM, connected by NVLink
- 2 × 16 core Intel(R) Xeon(R) Gold 6226R CPU @ 2.90GHz (total 32 cores)
- 768GB DDR4 RAM

Phase Ib

20 capacity GPU compute nodes

- 8 × NVIDIA Tesla T4 GPUs with 16GB RAM
- 2 × 16 core Intel(R) Xeon(R) Gold 6226R CPU @ 2.90GHz (total 32 cores)
- 576GB DDR4 RAM (1 node with 1536GB)

For more information about Alvis and how to use it, see the C3SE website (https://www.c3se.chalmers.se/about/Alvis) and, in particular, the section on “HPC and AI software”.

Staff Focus

…continued from page 10

of organometallic materials with a connection to inelastic neutron spectroscopy, aiming at early use of the BIFROST spectrometer at the European Spallation Source (ESS).

Johan has experience of development of scientific software, with skills in designing and implementing programs for high-performance computers using in particular the language Fortran90 in combination with parallelization using message passing interface (MPI) and shared memory (openMP). He is the lead developer of the magnetism branch of the Organic Material Database (OMDB), a platform for virtual material science experiments.

Now joining PDC as an application expert in materials theory, Johan will provide user support to the research community on computational physics in general, and with a niche towards modelling of condensed matter physics and quantum materials. As an application expert he will facilitate the use of software and services for modelling of condensed matter physics and measurement data from ESS neutron experiments. Modelling of magnetic and quantum materials can be pursued as part of the planning for experiments, and moreover also be a valuable tool directly in connection to the, typically week-long, allocations of measurement time at a neutron instrument.

In his spare time Johan enjoys choir singing, often in a language as alive and kicking as Fortran, namely Latin.

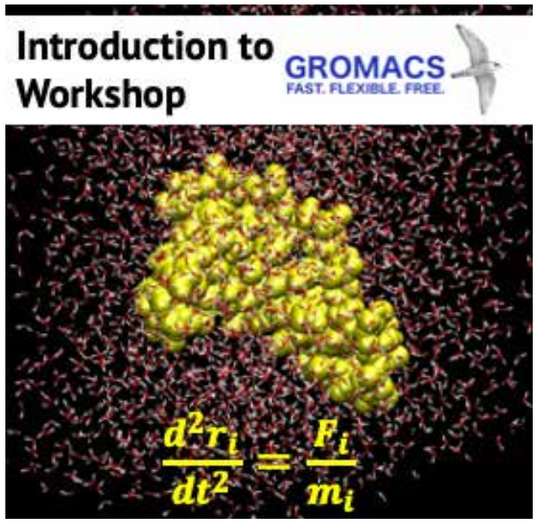

Introduction to GROMACS Workshop

SNIC/PRACE Workshop in Collaboration with BioExcel: 3-4 September 2020

Alessandra Villa, PDC

The first online "Introduction to GROMACS" workshop brought together around 50 participants from all five continents, from both academia and industry, and the Stockholm GROMACS team. Erik Lindahl gave a unique introduction to "Modern Molecular Dynamics Simulation (with GROMACS)", Berk Hess provided a solid introduction to free energy calculations, and Alessandra Villa covered the basics of molecular simulation. Paul Bauer and Szilárd Pall spoke about GROMACS features and development, as well as parallelization and performance.

The workshop was held online due to COVID-19 restrictions. Thanks to PDC Support, we were able to help participants set up their home computers so they could participate in the tutorial. Zoom was used to run the workshop remotely. Even though we could not meet the participants personally, lively discussions took place at all the sessions via Google Docs live documents. The tutorial sessions went smoothly: the trainees got support through Zoom chat and really enjoyed using Jupyter notebooks for the tutorial session. We look forward to meeting you at the next in-person workshop!

Welcome to the First Online Get-Together of the Nordic RSE Initiative!

30 November - 2 December 2020

You are invited to attend the first online get-together of the Nordic Research Software Engineer initiative: https://nordic-rse.org/events/2020-online-get-together.

Are you employed to develop software for research? Or are you simply spending more time developing software than conducting research? Then you have much in common with a growing international network of research software engineers (RSEs). The Nordic-RSE initiative aims to build a community of RSEs in the Nordic countries, plan meetings and workshops where knowledge can be shared, organize a biannual conference, and provide assistance in starting local RSE groups or hiring RSE staff in Nordic universities.

This online event will start with two invited talks on the 30th of November to set the stage, and will be followed by free-form conversations, short talks, workshops, or other types of contributions on the 1st and 2nd of December. Do you have an idea for a technical talk or short workshop that might not fit into a conference in your academic domain but might appeal to other technical people? If so, please consider submitting a lightweight abstract! If you are unsure whether your idea is within the scope, you can get in touch (contact@nordic-rse.org) to discuss it or browse the programme of the last UK-RSE conference in 2019 (https://rseconuk2019.sched.com) to see the list of RSE topics discussed there. You can, of course, also register without presenting! More details can be found on the event page: https://nordic-rse.org/events/2020-online-get-together.

Join us for the first Nordic-RSE get-together to develop your skills, meet like-minded people, learn about new job opportunities or find people to hire, and join an advocacy network for better research practices!

ENCCS
EuroCC National Competence Centre Sweden

Training Activities Offered by the EuroCC National Competence Centre Sweden

Thor Wikfeldt, ENCCS

The EuroCC National Competence Center Sweden (ENCCS) started on 1 September 2020 as one of 33 national nodes of the EuroCC project. The mission of ENCCS is to develop competence, knowledge and support in Sweden to enable academic and industrial researchers and high-performance computing (HPC) users to take advantage of forthcoming (pre-)exascale EuroHPC resources as well as modern artificial intelligence (AI) and high-performance data analytics (HPDA) methodologies.

A key part of this mission will be to deliver training events for HPC/AI/HPDA users and developers who aim to scale up their workloads or adapt their HPC/AI/HPDA software to new or different hardware platforms. ENCCS and SNIC will be collaborating and finding synergies in these activities within Sweden, for example by coordinating events, harmonizing training material and working together on outreach and dissemination. The training portfolio of ENCCS will largely be focused on topics which are rarely covered in existing SNIC training, including intermediate/advanced level training in MPI, OpenMP and CUDA, but will also include introductory/intermediate level training in AI and HPDA methods. If these efforts go according to plan, a coherent HPC training curriculum will emerge with regular courses following a clearly defined level specification which will assist researchers to develop from novice HPC users to parallel programming gurus!

ENCCS is also planning to host three hackathons in the coming two years. Research groups developing HPC code will be invited to submit proposals for hackathon projects and the proposals that are accepted will be assigned one or two mentors who will work with the group for a week. ENCCS aims to host one hackathon in collaboration with NVIDIA that will focus on GPU computing and another, together with Intel, on OpenMP in the first year, and a hackathon dedicated to AI in the autumn of 2022.

A tentative plan for upcoming ENCCS events is as follows.

- October 2020: Fundamentals of CUDA C/C++
- December 2020: Advanced MPI
- January 2021: Advanced GROMACS topics
- February 2021: Practical Deep Learning
- March 2021: Nek5000 training
- March 2021: GPU bootcamp with NVIDIA
- April 2021: OpenMP hackathon with Intel

To stay up to date on ENCCS activities, you can visit our website at https://enccs.se, subscribe to our newsletter (http://eepurl.com/heQCiv) or follow the project on Twitter (https://twitter.com/EuroCC_Sweden) or LinkedIn (https://www.linkedin.com/company/enccs).

SeRC Code Repository

Olivia Eriksson & Apostolos Vasileiadis, KTH

SeRC has recently opened a code repository where researchers can access software developed by SeRC. The repository is available online at https://e-science.se/code-repository. The computer programs that are included cover a range of research areas, such as molecular dynamics, bioinformatics, cancer screening, materials science, climate modelling, quantum chemistry and vizualisation. A few examples include the molecular dynamics simulation program GROMACS (which has more than 40,000 citations), the astrovizualisation tool OpenSpace and climt, a toolkit for building Earth system models in Python.

All code in the repository is open-source and may be used free for non-commercial research, with proper acknowledgment. However special rules may apply in some cases. These can be found in the documentation section for each piece of software.

Tracking Large Data Sets

Miguel Zavala-Aké, Niclas Jansson, Mohamad Rezaei, Marco Atzori & Philipp Schlatter, KTH, & Erwin Laure, Max Planck Computing and Data Facility

Introduction

Handling massive amounts of data is a vital issue in the Exascale age. In some areas, such as N-body cosmology, meteorology, Monte Carlo simulations or ocean eddy analysis, this issue is addressed by processing the data at the same time as it is generated. PDC is involved in the quest to overcome the difficulties associated with dealing with huge sets of data as part of its role in the European Centre of Excellence for Engineering Applications (EXCELLERAT) project.

In a majority of engineering and scientific applications, real-time processing aims to prepare data for visualization. (Visualization means converting the data into graphical representations, such as images, maps or charts, that make it easier to see patterns or trends in the data). However, reducing the volume of the data – by means of compression or feature extraction – is also important. Compressing the data is appropriate when a combination of post hoc and in situ processing is used: post hoc processing means the data is processed after it is produced, while in situ means the data is processed while it is being created. If the data-processing strategy is known in advance, in situ feature extraction is a more appropriate approach.

PDC's efforts have been focused on the challenges related to data reduction via the real-time handling of algorithms already available in the Visualization Toolkit (VTK) software. This data handling is performed at PDC using a high-performance computing (HPC) analysis tool, known as PAAKAT (which means “looking at” in the Mayan language).

An HPC In Situ Visualization Tool

The PAAKAT library has been designed as an HPC tool which encourages scalability and portability of in situ analysis in large-scale simulations. The emphasis is on reducing the output data of such simulations during run-time

Below: Compilation times for ParaView 5.6 without the graphical user interface (top) and modified ParaView (bottom)

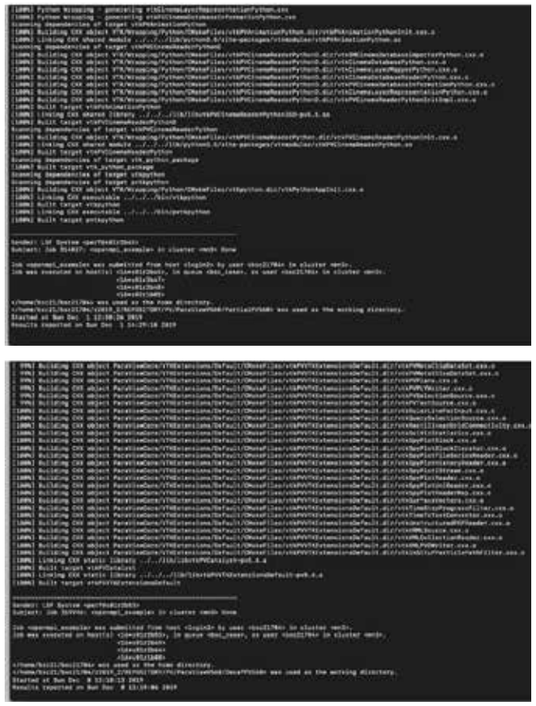

by using algorithms already available in VTK. The main difference regarding the great deal of effort made to develop software specializing in the solution of in situ visualization and analysis is related to the fact that PAAKAT encourages scalability and portability. This has been done by focusing on data arising from VTK filters while it obviates the need for rendering in the Paraview source code (version 5.6). These modifications encourage the use of the C++ VTK API, so that the need for third-party software components is reduced. As a consequence of these modifications, filters must be implemented using C++ instead of the Python scripts created by ParaView.

The screenshots above show the compilation times for two different ParaView setups. The first case corresponds to ParaView 5.6 with Python where renderings are considered but without the graphical user interface. In the second case, the modified ParaView is compiled. While in the first case the compilation time was about 99 minutes, in the second case it only took around 9 minutes. In both cases, 64 cores were used. As future work, more computer systems and compilers must be tried to investigate reductions in compilation time.

A Parallel Time Tracking Algorithm

This section briefly describes an example which uses the modified version of ParaView. This example features a parallel tracking algorithm which has been implemented exclusively using the C++ VTK API. The goal is to analyze the time evolution of coherent structures in a given turbulent flow. The algorithm is divided into three main parts (see below). Firstly, a scalar function $f(r,t) = c_O$ defines a set of points (isosurfaces) which take on a constant value c_O. Then, all these points are grouped in subsets (clusters), each of them characterized by a unique number. Finally, a search for overlappings between subsets belonging to different time steps is performed. Subsets are considered to be connected if overlappings between them exist. The search for overlappings is repeated each time step, so that the connections that are found make it possible to track the time for each cluster. The following subsections give details for each of these.

Isosurfaces and Clustering

Given a scalar field defined by a function, $f(r,t)$, isosurfaces and clusters are obtained using algorithms that can be found in the Visualization Toolkit library. Thresholding is performed by a recursive algorithm which allows users to identify a portion of an isosurface (cluster) through a unique number. Then these sets of numbers can be used to separate, and carefully investigate, each cluster.

Temporal Connectivities

In order to track each cluster over time, connectivities between clusters allocated in different time steps must be established. Here these connectivities are established in three steps. Firstly, each cluster is confined within the smallest box (aligned to the Cartesian coordinate system) which can hold all its points. Then, a search for overlappings is performed between a bounding box belonging to the current time step and bounding boxes held in a previous time step. Finally, if two bounding boxes (allocated in different time steps)

1. For two consecutive time steps t_i and t_{i-1} to calculate
   a. threshold
   b. clusters
2. For time step t_i:
   a. Overlapping
      i. Octree at t_{i-1}
      ii. (Minium) bounding boxes at t_i
      iii. Octree-bounding boxes intersection
      iv. Boolean operations on clusters
   b. Cluster interaction graph

Above: Tracking algorithm

Staff Focus

Niklas Karlsson

Niklas Karlsson has recently joined PDC as a systems manager and will be focusing on infrastructure. He most recently comes from the consulting operation, MSC Solutions, and has been involved in system administration with a Linux focus for ten years. Prior to that, he had a background in computer networking, having spent two years working for Cisco Systems early in his career, and technical support.

In his free time, he enjoys role-playing and board games as well as science fiction, both in written form and on the small or the big screen.

overlap, then a Boolean operation is performed between the clusters contained by each bounding box. A positive result from the Boolean intersection of these clusters means that overlappings between these clusters exist. Such an overlapping can be understood as indicating that a temporal connection exists between those two clusters.

Overlapping

In the simplest cases, a cluster allocated in a given time could be either completely unconnected to any other clusters or connected with only one previous cluster. In the first case, it is possible to assume that a new cluster has emerged, while in the second case, the tracked cluster has only suffered small changes. In more complicated cases, multiple connectivities indicate that the current cluster results from the merging of multiple clusters. Other possible kinds of connections (merging and splitting) are depicted in the figure below.

Merging connectivities take place when, in a given time, two or more previous clusters overlap a single cluster, so that some of the previous clusters seem to have disappeared. Now, these merges could be either total or partial. A total merge occurs when each of the previous clusters is connected to each of the current clusters. A partial merge takes place when at least one previous cluster does not overlap any current cluster. This means that some of the previous clusters are temporally unconnected. In this situation, it could be considered that the unconnected clusters were either embedded (by a contiguous cluster) or vanished (perhaps due to the nature of the physics of the underlying problem). Finally, it is worth noting that, when merges take place, the total number of current clusters is lower than the number in the previous time step.

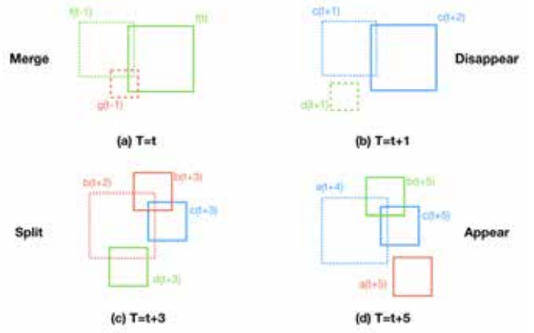

Above: Temporal connectivities which correspond to the stages in the algorithm on the previous page. Continuous lines are assigned to clusters allocated in current time steps, while dashed and dotted lines are used for clusters in previous time steps. In all cases, a letter is used to identify each cluster, and the subscript indicates the time step to which the cluster belongs.
(a) ***Total merging:*** *Previous clusters f_{t-1} and g_{t-1} give rise to current cluster f_t.*
(b) ***Partial merging:*** *Cluster d_{t+1} has disappeared, while clusters c_{t+1} and c_{t+2} are connected.*
(c) ***Total splitting:*** *Cluster b_{t+2} gives rise to clusters b_{t+3}, c_{t+3}, and d_{t+3}.*
(d) ***Partial splitting:*** *Cluster a_{t+4} is connected to clusters b_{t+5} and c_{t+5}, while cluster a_{t+5} has appeared.*

Unlike merging, splitting connectivities takes place when a single previous bounding box overlaps two or more current bounding boxes. In these

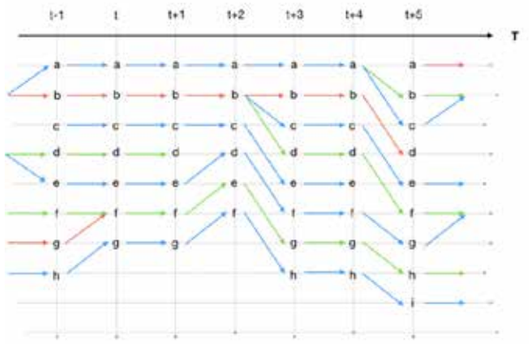

Above: This time tracking graph depicts seven time steps (or stages). Lowercase letters are used as cluster identifiers. At each stage, arrows are used to indicate input and output connections. Input connections show the connectivities between clusters held in the previous time step and the clusters in the current stage. Output connections show the connectivities between the clusters held in the current stage and the clusters in the next time step. For instance, at stage t+2, there are six clusters, six input-connections and eight output-connections. Input-connections show that the last three clusters have been renamed, in other words, the previous clusters identifiers e, f and g correspond to d, e and f in this stage. This is due to that fact that cluster d_{t+1} has died out, that is, it has become unconnected. In the same stage, output-connections show the total splitting of cluster b, see (c) in the figure to the left. This leads to a situation where the last four clusters (c, d, e, and f) must be renamed, since two new identifiers (c and d) have been added. Now time tracking of each cluster could be followed. For instance, at stage t + 2, cluster d results from cluster a splitting which takes place at step t – 2 ($d_{t-2} \rightarrow e_{t-1} \rightarrow e_t \rightarrow e_{t+1} \rightarrow d_{t+2}$). This same cluster ($d_{t+2} \rightarrow f_{t+3} \rightarrow f_{t+4} \rightarrow g_{t+5} \rightarrow f_{t+6}$) is merged with a portion of cluster b ($b_{t+2} \rightarrow d_{t+3} \rightarrow d_{t+4} \rightarrow f_{t+5} \rightarrow f_{t+6}$).

cases, new clusters seem to have emerged, so that the total number of current clusters is larger than in the previous step. Total splitting occurs when each cluster that emerges is connected with some of the previous clusters. In contrast, when partial splitting occurs, any unconnected clusters that arise are supposed to be entirely new.

The set of temporal connectivities established at each time step can be used to compose a graph which shows how the interaction between clusters evolves over time. In the next section, this is discussed in depth.

Time Tracking Graph

Merging and splitting of clusters could happen at any time and anywhere. A time tracking graph (as shown in the figure on the right on page 16) helps us to take an in-detail look at when merging and splitting take place. In these graphs, the set of cluster identifiers are represented as vertices while edges indicate the temporal connectivities resulting from searching for overlappings (see the previous section). Vertices are grouped according to the time step (or stages) to which they belong and, in turn, these time steps are organized in chronological order. In this way, edges can only traverse from one time step to another.

In the figure on the right on page 16, temporal connectivities shown in the figure on the left on page 16 are used to depict a time tracking graph. Cluster identifiers are placed from top to bottom, while time flows from left to right. Lowercase letters are used to identify each

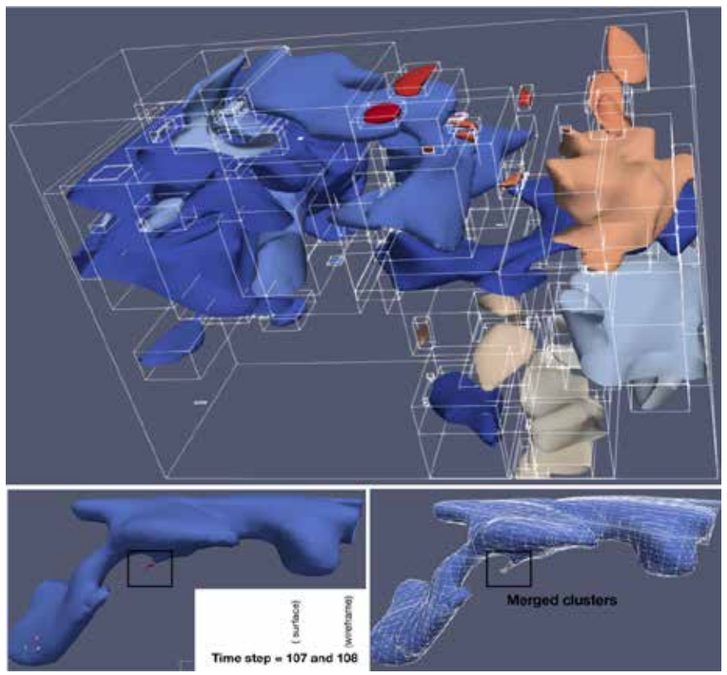

Above: Time evolution of coherent structures – Clusters at c_o = 2.0 at time steps 107 and 108

Celebrating Five Years of BioExcel

Rossen Apostolov, PDC

The BioExcel Centre of Excellence for Computational Biomolecular Research started on the 3rd of November 2015 and has now been running for five years. KTH/PDC is the coordinator of the centre and, in addition to managing the project, provides expertise in molecular dynamics simulations and their efficient use on HPC systems.

To celebrate its 5th anniversary, BioExcel has launched a blog and social media campaign with a video, highlighting its journey and team. Please use the links below to find out more about BioExcel's achievements and meet some of the people behind its successes!

BioExcel Blog
bioexcel.eu/celebrating-5-years-of-bioexcel

BioExcel on Twitter
twitter.com/BioExcelCoE/status/1323248834701676550

BioExcel on LinkedIn
www.linkedin.com/feed/update/urn:li:activity:6729014964100571136

Facebook
fb.watch/1vLAM2W9zL

vertex, while arrows are used for edges. In each stage, arrows indicate input and output connections. An input connection relates a previous time step with the current stage, while an output connection relates the current stage with the next time step. Now, the four stages (t, $t+1$, $t+3$, and $t+5$) shown in the figure on the left on page 16 can be found here. At stage t, seven clusters exist, along with eight input connections and seven output connections. Each cluster is simply connected, with the exception of cluster f, which has two input connections ($f_{t-1} \rightarrow f_t$ and $g_{t-1} \rightarrow f_t$) that arise from the merging of clusters f and g at time step $t-1$. Finally, it is worth noting that, due to the merging, in this stage the number of clusters has decreased with respect to the previous time step. As a consequence, the last cluster at $t-1$ has been renamed. This means that, from the search for overlappings, cluster h_{t-1} and g_t are simply connected, in other words, these two clusters are the same. The previous procedure can be repeated for the rest of the stages. At $t+1$, the number

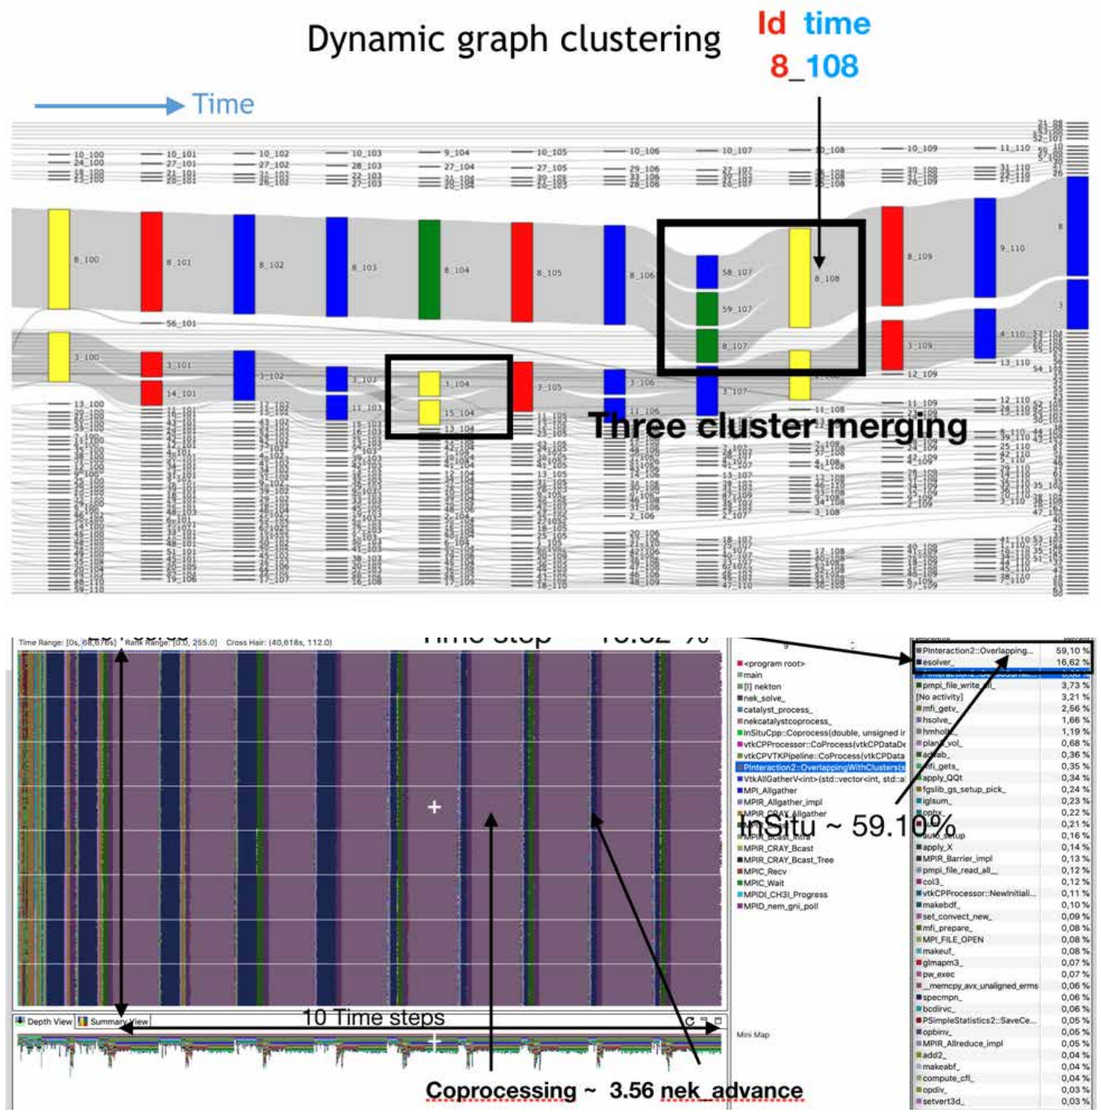

Above: Time evolution of coherent structures – (top) Dynamic graph clustering, and (bottom) Parallel execution unfolds over time

of clusters remains unchanged. For the next stage, the number of clusters has decreased, due to the fact that cluster d_{t+1} has died out.

Results

In the next case study, the goal is to analyze the time evolution of coherent structures in a turbulent flow. The numerical simulation makes use of 2×10^6 elements, 1×10^2 time steps, 256 MPI processes, and the parallel code Nek5000 to solve a direct numerical simulation (DNS) of a turbulent flow at a friction Reynolds number $Re_\tau = 180$.

The figure on page 17 shows a set of clusters (coherent structures) at time step 107, along with three extracted clusters (8, 58 and 59) and their evolution towards a simple cluster (8 at time step 108). Their evolution over ten time steps (dynamic graph clustering) and the corresponding chronology in which the parallel workflow of the tracking algorithm is executed (parallel activity trace) are shown in the upper and lower parts of the figure on page 18 respectively.

The time tracking graph related to this case is plotted by using a Sankey diagram in which temporal connectivities between clusters allocated in different time steps can be represented. In the upper part of the figure on page 18, time flows from left to right, clusters are represented as vertices (coloured rectangles) with temporal connectivities (thick gray lines) as edges between them. Each vertex is labelled as *id_time*, which corresponds to its cluster identifier *id* (as given by vtkConnectivityFilter) and the current time step time. Thus, the merging of clusters shown in the figure on page 17 is represented in the graph as three vertices (8_107, 58_107 and 59_107) connected with a simple one (8_108) through three edges. Taking as a reference time t=107, it is possible to see that these clusters arise from the splitting of cluster 8 at time t=106. In the same way, the evolution of each and every one of the clusters can be followed over the simulation time.

The parallel activity trace shown in the lower part of the figure on page 18 gives a first performance analysis of the tracking algorithm. In the trace that is shown, 256 MPI processes are used to execute ten time steps, which correspond to steps from 100 to 109. For these steps, the statistics view shows that the execution time of the in situ analysis (vtkCPVTKPipeline :: CoProcess ~ 59.10%) is 3.56 times bigger than for the Nek5000 solver (esolver ~ 16.62%).

Conclusions

As part of the EXCELLERAT project, an in situ instrumentation for the code Nek5000 has been prepared and used to perform the time evolution analysis of coherent structures. For this instrumentation, the HPC in situ analysis tool PAAKAT has been used. The performance analysis shows a low increase in the total execution time (which includes simulation and analysis times). In addition, it should be considered that in post hoc processing the simulation program needs to stop and wait for the application of some data reduction method, which increases the total execution time. The ongoing work considers efficiency and scalability tests in different exascale machines.

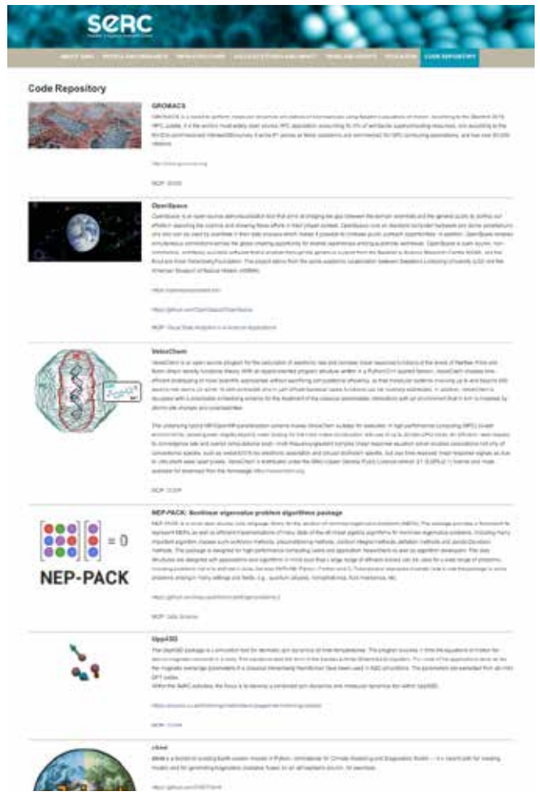

Above: SeRC Code Repository - for details see page 13.

Procurement of New PDC System

Gert Svensson, PDC

As mentioned in the previous PDC Newsletter, PDC received a substantial grant (in total 170 million SEK) from the Swedish National Infrastructure for Computing (SNIC) to install and operate a new general-purpose high-performance computing (HPC) system for academic research. The process of procuring the new system had just started in the spring with an initial invitation being published and submitted to a range of potential vendors in April. The invitation roughly described the desired system, as well as specifying some requirements that any companies submitting tenders would need to satisfy. You may remember that the earlier article explained that the new system will have one partition using only central processing units (CPUs) and another partition which will be equipped with graphics processing unit (GPU) accelerators. In addition, the system will include a fast Lustre storage subsystem.

During May a large number of vendors replied that they were interested in bidding, and also provided information about the economic status and technical capacity of their businesses, so that PDC could judge if their companies were financially and technically suitable to qualify as potential suppliers of the new system.

In June PDC sent a request for proposals (which were due at the end of August) to all the companies that had qualified. The request contained a detailed technical and commercial description of the requirements for the new system, together with the benchmark suite and test cases, as well as detailed information about how the bids would be evaluated. Several vendors asked us to prolong the proposal period due to the COVID-19 situation, so the due date was moved three weeks later. Also many of the vendors that had initially expressed interest dropped out and did not submit proposals.

By late September PDC had received some exciting proposals from various vendors. Since then, things have been hectic: reviewing the bids, rectifying mistakes in the proposals and benchmarks, and holding discussions with the vendors to improve the bids. Now we are in the final stage of the procurement process. At the time of writing, the remaining vendors had just been asked to submit their best and final offers. So, by the time this newsletter is published, the supplier of the new system may have been selected. You can rest assured that PDC will soon have an excellent new system for the coming years!

PDC-Related Events

PDC Summer School 2021

last two weeks of August 2021, KTH, Stockholm

For details of the next summer school, watch https://www.pdc.kth.se/about/events.

HPC Sources

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

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

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

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

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

HPCwire
http://www.hpcwire.com/events

NeIC
http://neic.nordforsk.org

PRACE
http://www.prace-ri.eu/HPC-access
http://www.training.prace-ri.eu
http://www.prace-ri.eu/events
http://www.prace-ri.eu/news

SeSE
http://sese.nu

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

XSEDE
https://www.xsede.org

PDC's current flagship system, Beskow, is due to be replaced during 2021. It has been in operation since 2014.