PDC Center for High Performance Computing

PDC Newsletter 1/24

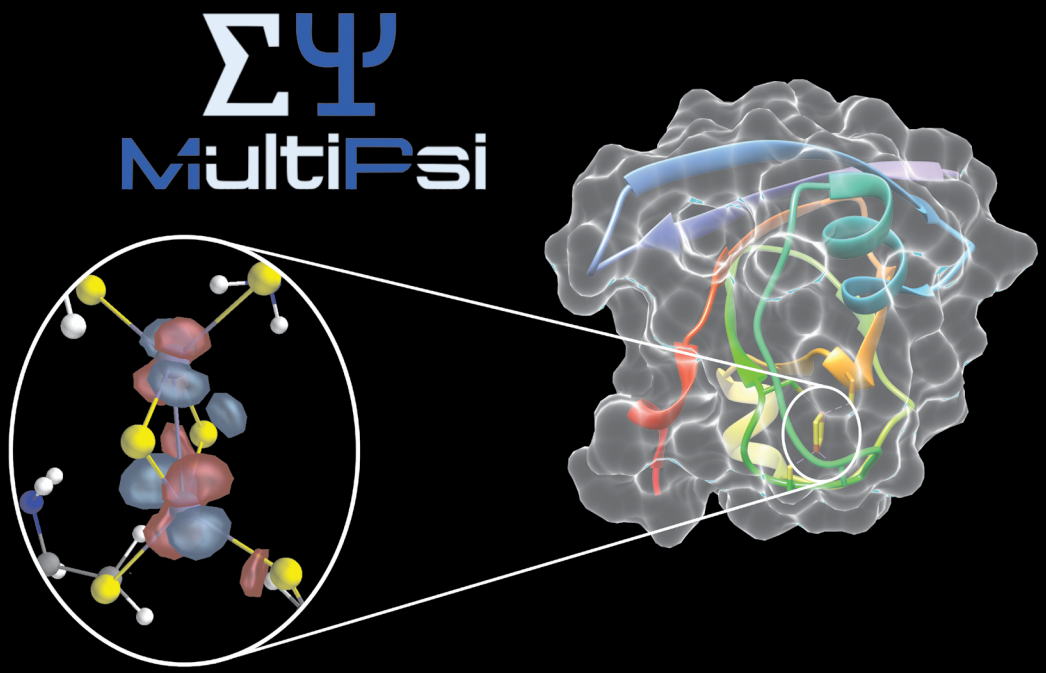

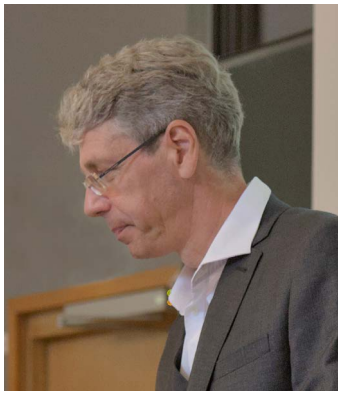

Patrick Norman, Director, PDC

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 image on the cover displays the ferredoxin protein that was computed fully quantum-mechanically using the program MultiPsi. The zoomed-in region shows the iron-sulphur [Fe_2S_2] active site and one of the ten orbitals describing the antiferromagnetically coupled iron spins.

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

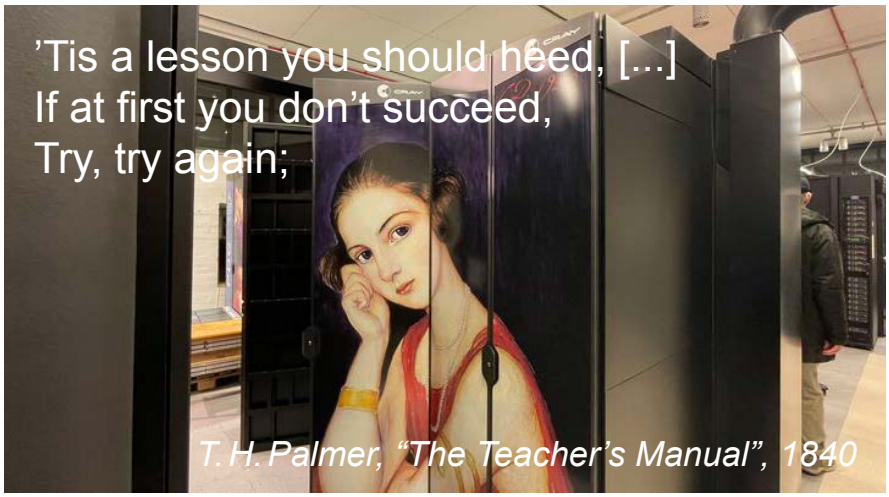

After an all too long wait, the summer season has suddenly taken us by surprise yet again and filled our landscape with bright colours and sweeping warm winds. No matter the hardships one faces at work, merely travelling back home while the evening skies are blue and full of light fills the spirit with a sensation of optimism, and I dare say that has been needed very much in recent months. There is no denying that the Dardel system has been less than cooperative when undergoing changes and upgrades in the past. After what felt like endless discussions with the vendor, the decision was finally made to change the cluster management software for the system from CSM to HPCM (see *"Dardel Status"*). By the time you read this editorial, we will know if our optimistic expectations for a more stable and easily managed system have been realised, as that is certainly what we are aiming for. The number of work hours and amount of stress leading up to this software transformation of Dardel are difficult to put into words, but having watched this from the ringside, I wish to express my deepest gratitude to everyone at PDC who has been involved in implementing this upgrade for their hard work and dedication.

We are looking forward to entering a period of normality in terms of system operation so that overworked staff can recover after we have crossed this bridge to stability. Reaching the other side of this transitional bridge in our operations will allow PDC to concentrate on introducing more science-enabling interfaces, services, and software tools onto Dardel. Of note is the fact that the KTH Royal Institute of Technology (KTH) and PDC are collaborating with the vendor of the overall Dardel system, namely Hewlett Packard Enterprise (HPE), and also with AMD, which produces the processors used in the Dardel system, to enhance the performance of selected scientific applications including GROMACS, VeloxChem

and Neko (see *"KTH Collaborations with AMD and HPE in the Development of Scientific Applications"*). Our flagship software programs are often highlighted (and rightfully so!) and, in this context, I would like to take the opportunity to mention the rapid development of the VIAMD software for visual analysis of classical – and now also quantum mechanical – simulations of molecular systems (see *"VIAMD Update"*).

It is very rewarding to see how the KTH investments in high-performance computing (HPC) software developments are impacting activities elsewhere, and it is with great pleasure I have finished reading the cover article of this newsletter that presents the impressive work of a research group at Lund University. In a short period of time, they have, with the use of the VeloxChem program as a platform, created the MultiPsi program and thereby reached and surpassed the state-of-the-art in electronic structure theory simulations of strongly electron-correlated molecular systems (see *"MultiPsi: Quantum Chemistry on HPC for Photochemistry and Inorganic Chemistry"*). Many of you may have heard me say this before, maybe even more than once, but I find this work to be yet another manifestation of the fact that we are now in a position where we can adapt our software according to the needs of our research, rather than having the bounds of our research being delimited by the software that is available on the broader market.

A concrete piece of evidence demonstrating that the KTH investments in HPC are sound is the remarkable success that has been noted for KTH in attracting external funding via a range of EuroHPC Centres of Excellence (CoEs) and other international projects, several of which are addressed in the present newsletter, such as the new EPICURE project (see *"EPICURE: High-Level Support for Research with Top-Flight EuroHPC Systems"*), as well as familiar faces like BioExcel (see *"BioExcel News"*). The GROMACS and Neko programs are both vehicles for their respective CoEs and PDC is now looking for talented research software engineers to strengthen these two teams. (Information about currently advertised jobs can be found at https://www.kth.se/lediga-jobb?l=en.) On this optimistic, forward-looking note, I close this editorial and wish you all a wonderful summer.

Patrick Norman, Director PDC

Save-the-Date:

NAISS User Forum & All-Hands Meeting

1-3 October 2024, Uppsala

https://www.naiss.se/event/user-forum-2024

https://www.naiss.se/event/all-hands-meeting-2024

In This Issue

MultiPsi: Quantum Chemistry on HPC for Photochemistry and Inorganic Chemistry

Mickael Delcey, Division of Computational Chemistry, Lund University

Modelling Chemistry

Computational chemistry has revolutionised science and industry and helped to improve our understanding of chemical processes. Instead of relying solely on trial and error in the lab, researchers can now predict how molecules will behave under different conditions using computer simulations, thereby accelerating, for example, drug design or the development of materials for applications ranging from electronics to energy storage. These advancements have only been made possible by simultaneous developments of new computational chemistry methods, more powerful computers and more efficient implementations that leverage the power of emerging computer architectures. Thus, while the first quantum descriptions of atoms and simple molecules date back a hundred years, it is only in recent decades that quantum chemistry has become a realistic tool for industry.

One Size Does Not Fit All

Quantum chemistry encompasses a broad collection of methods, each with its own strengths and weaknesses. Some of this diversity simply corresponds to increasingly sophisticated approximations that yield better results at a higher computational cost. However, the choice of an appropriate method often depends on the specific system or property of interest. What works well for one molecule or chemical property may fail dramatically in another case. One important example is the approximation that the electronic structure of a molecule is well described by a single electron configuration. This is the case for most organic molecules: molecules composed mainly of carbon, hydrogen, oxygen and nitrogen, which are prevalent in living beings. In these cases, computationally inexpensive methods, like perturbation theory or density functional theory (DFT), are very effective. However, these approximations sometimes break down, in particular for molecules containing transition metals like iron, manganese, or chromium. Yet these atoms are crucial in solid-state materials and also in biological systems. Indeed, even if they represent a small fraction of the elements in living beings, transition metals are estimated to be present in a third of all proteins and enzymes where they are often key to their roles. Transition metals are, in particular, essential to some of the most important bioreactions, such as photosynthesis or nitrogen fixation (a reaction in some bacteria which contributes to soil fertilisation). A similar breakdown of these approximations can occur in photochemistry, for example, at the critical moment when an excited molecule relaxes back into its ground state without emitting light. In these cases, the correct wave function can be built using a (hopefully small) number of configurations, giving rise to so-called multiconfigurational methods.

Modularity

The abundance of quantum chemistry methods has led to the implementation of numerous algorithms in most quantum chemistry software. Consequently, these programs become large and unwieldy, posing challenges for long-term maintenance and hindering accessibility for new students and researchers. To address this issue, there is a growing push towards developing small modular programs that excel at specific tasks, promoting consistency and interoperability [4]. This shift aims to streamline code development and usage, thus reducing redundancy and facilitating collaboration within the research community.

Some years ago, development of the VeloxChem program [6] started at the KTH Royal Institute of Technology (KTH), with active support from PDC staff and promised a highly efficient code for all types of chemical properties at the Hartree-Fock and DFT level. This initiative paid off, and some of the code's impressive performances have been reported in previous PDC newsletters. This world-leading development inspired a number of groups, including ours, to develop additional libraries covering alternative electronic structure methods.

Thus, while VeloxChem covers Hartree-Fock and DFT (the most inexpensive ab initio quantum chemistry methods), Gator [5], which was developed in partnership with Heidelberg University, expands the functionality with the algebraic diagrammatic construction (ADC) method, which is a more expensive but reliable wave function method for organic molecules. Similarly, the MultiPsi [1] library, which was developed mostly at Lund University, expands the functionality towards multiconfigurational methods to describe, among others, metal complexes.

Goals and Ambitions for MultiPsi

Multiconfigurational methods can often be thought of as generalisations of standard single configurational methods and have, in principle, a computational cost that is only moderately higher. Yet, they have received less attention and, for this reason, lag behind in efficiency. One of the main aims of MultiPsi is to enable truly large-scale multiconfigurational calculations and, in particular, to be fully adapted to modern heterogeneous parallel computer architectures. The program should be, for instance, efficient enough to simulate large metalloproteins or the dynamics of medium-sized molecules after light absorption for nanoseconds or longer.

But beyond this, these methods are also known to be more challenging to use and develop and are rarely presented in great detail in standard textbooks. For this reason, education is a significant aspect of making these methods more widely used and promoting further developments. MultiPsi was thus designed with usability as a fundamental property as well as high interactivity. With this mindset, along with the rest of the VeloxChem community, we have set up a collection of Jupyter notebooks forming an e-book meant for teaching quantum chemistry at all levels, which we call the e-Chem initiative [2].

Finally, MultiPsi is not meant to simply implement existing methods in an efficient way but to be a development platform for modern multiconfigurational quantum chemistry. Its structure should make it easy to develop new code, from early prototypes to final implementations, even for new students who have limited prior programming experience.

Structure of MultiPsi

In practice, MultiPsi, like all of our programs, has a two-layer programming structure. The upper layer, written in Python, is where most of the quantum chemical methods are implemented, and it also manages the use of the available hardware

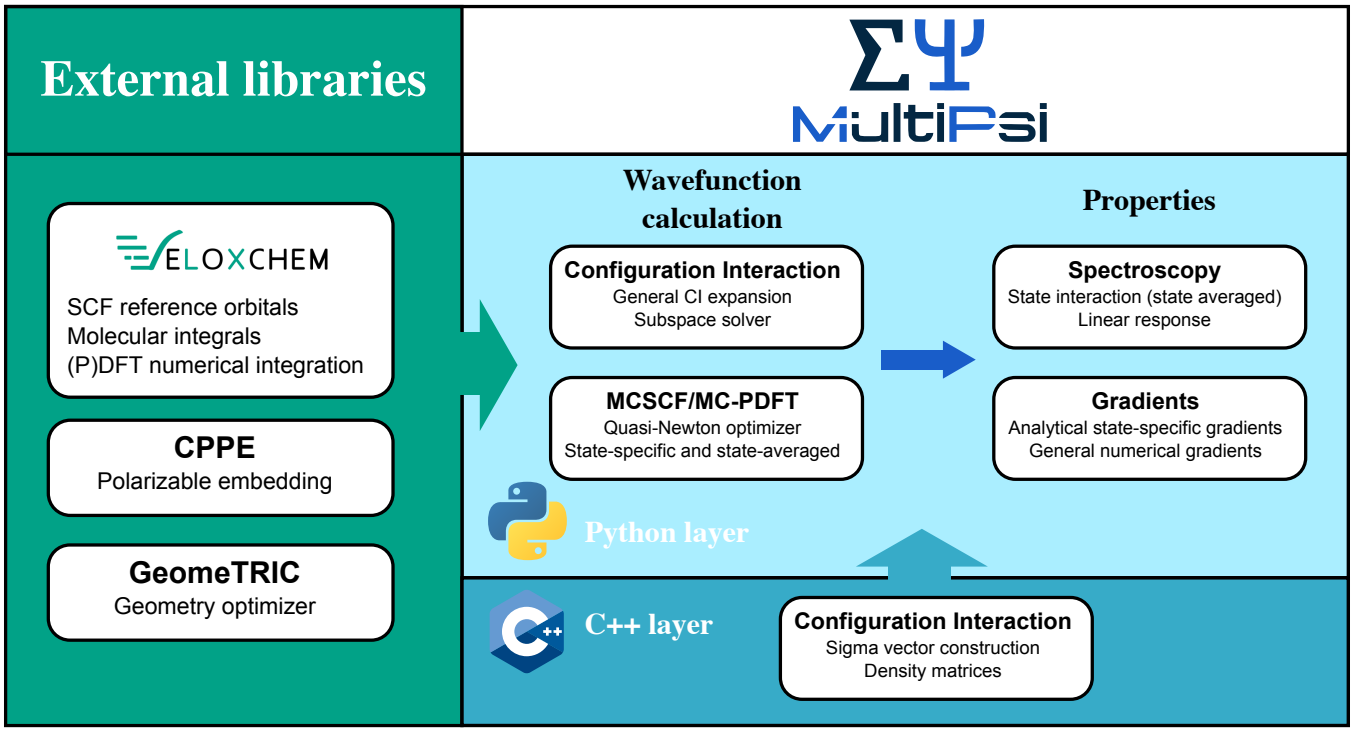

Above: Code structure of MultiPsi

resources using message passing interface (MPI) communicators. The lower layer, written in C++, contains the compute-intensive core functionalities. Finally, the code uses a number of external libraries, including VeloxChem.

We can illustrate this structure for the multiconfigurational self-consistent field method (MCSCF), which is one of the fundamental multiconfigurational methods in much the same way as the Hartree-Fock method is for single configurational methods. The MCSCF algorithm is implemented in Python but requires two computationally intensive functionalities: the evaluation of the electron repulsion integrals (ERIs), which represent the interactions between electrons, and the configuration interaction (CI) to evaluate the relative weights of all configurations. The ERIs are a central quantity in all of quantum chemistry, as they allow us to transform the differential equations of quantum mechanics into a linear algebra computer program. They are expensive and numerous, scaling as N^4 with the system size, though efficient screening methods can reduce this scaling significantly. Their evaluation is performed by C++ routines in VeloxChem, but, through Pybind11, these routines are exposed to the Python layer and can then easily be used by MultiPsi by importing the VeloxChem module. The CI functionalities are, on the other hand, exclusive to MultiPsi and implemented in their own C++ layer. The overall structure can be seen in the figure on the previous page.

The modularity of the design means, however, that every part can easily be substituted. For example, MultiPsi has been used in a prototype of hybrid quantum computing code [3] by substituting its CI module with the module QiSkit from IBM, which is able to perform similar function but on a quantum computer. Writing this prototype took less than a day for a student who was not familiar with MultiPsi's code prior to doing this!

Efficiency

MultiPsi has been written from scratch around heterogeneous parallel computer architectures. Taking MCSCF as an example again, as discussed above, the cost is essentially the sum of an ERI-driven part and a CI part. By using the general ERI evaluator from VeloxChem, MultiPsi inherits its efficiency with only a little work being needed to control the task distribution. This gives the parallel performance shown in the figure above. This figure shows both the parallel efficiency for the first iteration and for the total calculation. The timing for the first iteration is almost exclusively made up of the ERI evaluation, while the (partly single-node) convergence algorithm starts to dominate for a large number of nodes due to Amdahl's law.

Below: Parallel performance to simulate an X-ray absorption spectrum at the MCSCF level using MultiPsi Taken from Ref. [7]

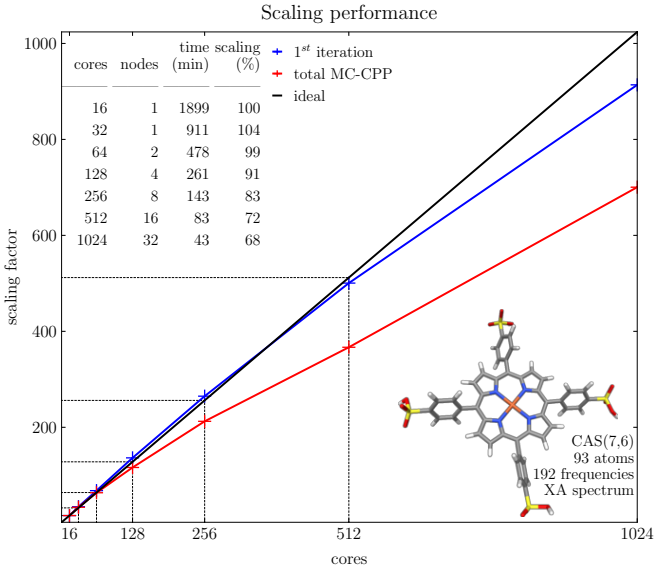

The second part of the cost is the CI, which depends on the number of configurations. Configurations are typically chosen by allowing all possible distributions of a subset of electrons in a subset of orbitals in the system. Due to the combinatorial nature of this technique, the number of configurations grows very fast (factorially) with the number of electrons and orbitals. As a result, while the most expensive operation in a CI calculation is a simple matrix-vector product, the size of the vector (the number of configurations) can quickly reach billions of double-precision numbers. For such sizes, the CI matrix is completely impossible to store, even taking advantage of its significant sparseness. Instead, it is computed on the fly and immediately multiplied by the CI vector. The choice of algorithm then becomes a fine balance between vectorisation and maximal utilisation of the sparseness.

Below: Scaling of the wall time of the CI step in two different calculations, with the number of OpenMP threads or MPI ranks
(a) Scaling of the wall time for a CI calculation of more than 82 million configurations with increasing number of cores
(b) Scaling of the wall time a CI calculation of more than 1.1 billion configurations with increasing number of MPI ranks, with 64 cores each.

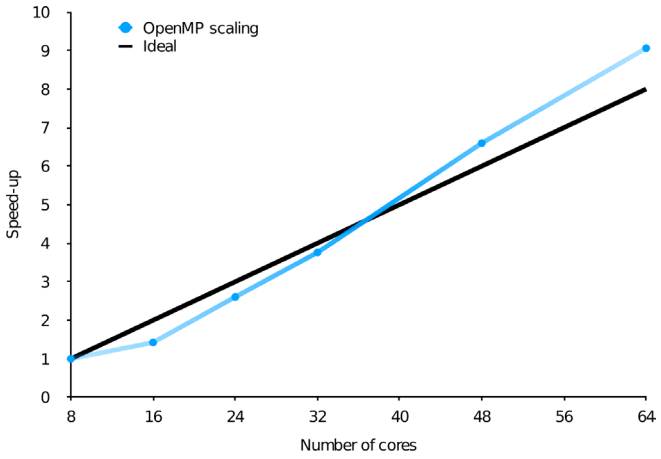

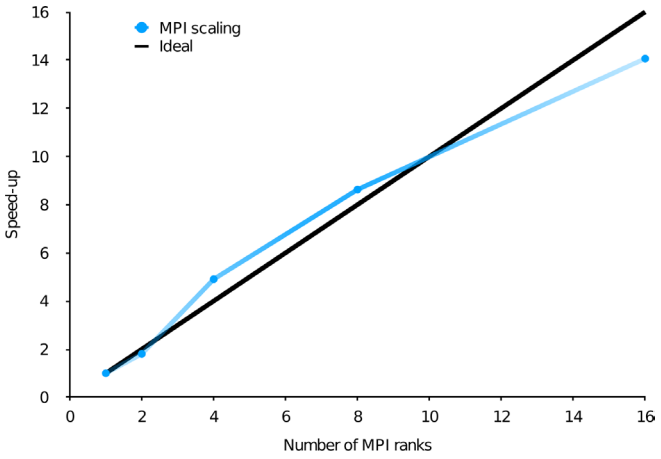

Still, with billions of configurations, even the CI vector can become a memory bottleneck. Thus, when designing parallel algorithms, it is essential for these vectors to be stored in the distributed memory. This, however, necessarily leads to high inter-node communication since in a matrix-vector multiplication, in principle, each resulting vector element has contributions from all the starting vector elements. This can be greatly reduced by making use of the predictable sparsity in the CI matrix, which introduces some level of locality in the data. The resulting OpenMP and MPI parallel efficiency using Dardel are shown in the figures above on the left (a) and right (b), respectively.

In the end, this efficiency has allowed MultiPsi to shatter several records. In our release paper, we broke the record for the largest CI optimisation. Using 128 nodes on the Dardel system at PDC, we optimised the CI weights for more than 400 billion configurations. At this size, a single CI vector is about 3TB of distributed memory, and a couple of them need to be stored at a given time. In addition, we also significantly pushed the limit of the largest molecule ever computed by MCSCF. We have now demonstrated MultiPsi's ability to simulate a biomolecule of nearly 1,500 atoms using only 256 cores [8]. It is generally considered that 1,000 atoms are sufficient for a quantum chemical calculation, as atoms far apart from the region of interest can then be advantageously modelled using significantly cheaper classical molecular mechanics – a form of multiscale modelling. With these performances, MultiPsi is arguably the most efficient platform for multiconfigurational calculations in the world.

Outlook: Towards New Science

MultiPsi is already available through Conda for testing and is going to be released officially under the GNU Lesser General Public License version 2.1 (LGPLv2.1) sometime this year. Currently, it offers mostly energy and properties (in particular spectroscopy simulations) with the MCSCF method. However, as discussed earlier, this is not enough. Multiconfigurational methods have also lagged behind in functionalities and method availability and require new developments.

In particular, while DFT has been the most used method in quantum chemistry for decades, there has not been a satisfying multiconfigurational variant, despite decades of development. We believe we have finally developed such variant, a true generalisation of DFT to many configurations [8]. We showed that the method indeed extends the accuracy of DFT to systems requiring multiconfigurational methods. In addition, as can be seen in the figure on the next page, for large molecules, this method can be as efficient as DFT, despite it providing a better description of the multiconfigurational nature of the metal centre. This record performance is also thanks to help

Below: Calculation time for MCSCF and our new multiconfigurational DFT (MC-PDFT) compared to standard DFT, for increasing model sizes of a Ferredoxin protein The number of configurations is kept constant and centred around the metal active site.

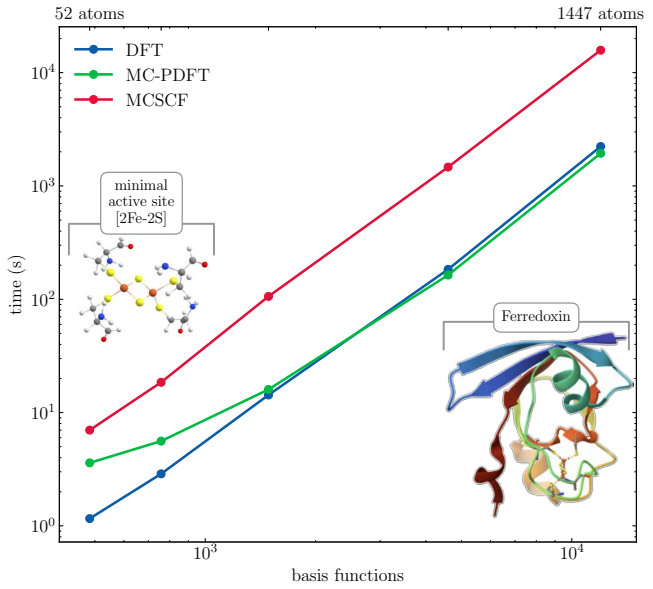

from a PDC software specialist to parallelise the (P)DFT numerical integration.

Developments are ongoing to establish this new method and implement property calculations with it. With this, MultiPsi will finally offer quantum chemists a realistic alternative to DFT for metal complexes and other molecules where the latter may not be accurate enough.

Acknowledgements

The VeloxChem and MultiPsi projects are supported by a collaboration with the PDC Center for High Performance Computing. The results presented in this article were obtained with the use of computational resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS).

References

1. M. G. Delcey, "MultiPsi: A python-driven MCSCF program for photochemistry and spectroscopy simulations on modern HPC environments". WIRES Comput. Mol. Sci. 13.6 (2023), e1675. https://doi.org/10.1002/wcms.1675
2. T. Fransson et al., "eChem: A Notebook Exploration of Quantum Chemistry". In: Journal of Chemical Education" 100.4 (2023), pp. 1664–1671. https://doi.org/10.1021/acs.jchemed.2c01103
3. J. de Gracia Triviño, M.G. Delcey, G. Wendin, "Complete Active Space Methods for NISQ Devices: The Importance of Canonical Orbital Optimization for Accuracy and Noise Resilience". Journal of Chemical Theory and Computation 19.10 (2023). PMID: 37103120, pp. 2863–2872. https://doi.org/10.1021/acs.jctc.3c00123
4. S. Lehtola, "A call to arms: Making the case for more reusable libraries". The Journal of Chemical Physics 159.18 (Nov. 2023), p. 180901. issn: 0021-9606. https://doi.org/10.1063/5.0175165
5. D. R. Rehn et al., "Gator: A Python-driven program for spectroscopy simulations using correlated wave functions", WIREs Comput. Mol. Sci. 11.6 (2021), e1528. https://doi.org/10.1002/wcms.1528
6. Z. Rinkevicius et al., "VeloxChem: A Python-driven density-functional theory program for spectroscopy simulations in high-performance com- puting environments". WIREs Comput. Mol. Sci. 2020, 10:e1457. https://doi.org/10.1002/wcms.1457
7. M. Scott, M. G. Delcey, "Complex Linear Response Functions for a Multiconfigurational Self-Consistent Field Wave Function in a High Performance Computing Environment". Journal of Chemical Theory and Computation 19.17 (2023), PMID: 37596971, pp. 5924–5937. https://doi.org/10.1021/acs.jctc.3c00317
8. M. Scott et al., "Variational Pair-Density Functional Theory: Dealing with Strong Correlation at the Protein Scale". Journal of Chemical Theory and Computation 20.6 (2024), PMID: 38217859, pp. 2423–2432. https://doi.org/10.1021/acs.jctc.3c01240.

Dardel Status

Gert Svensson, PDC

In the last issue of the PDC Newsletter, it was mentioned that PDC had a test rack with a different cluster manager from Hewlett Packard Enterprise (HPE) called HPCM (https://www.

pdc.kth.se/publications/2023-no-2/dardel-status-1.1291982). HPCM is a less complex and more stable cluster management system that has worked well on other cluster installations for many years. PDC considered the results of tests with the rack to be promising and decided to convert Dardel to use HPCM as this change should make it easier to keep the entire Dardel software stack up to date while maintaining stability.

The last step in the conversion was to shut down the old system and install a more modern software stack under HPCM. At the time of writing, this process was in its final steps. Almost all system software has been updated, including the Cray Programming Environment, the Slingshot software, the Linux software on the compute nodes and the AMD ROCm package. This means that all applications executing on multiple nodes must be recompiled. PDC has already installed the applications that are used most.

The package for the remote graphical desktop system ThinLinc is also being updated. PDC is installing a new version of ThinLinc, as well as fixing several of the problems with the previous installation. For users with less experience of Linux, ThinLinc is an easier alternative for logging in with SSH and using Linux shell commands. ThinLinc allows users to interactively perform data visualisation and simulation in a virtual desktop environment, thereby minimising the need to offload data from Dardel for simulations. With this upgrade, PDC will be offering support for additional applications, such as JupyterLab, Jupyter Notebook, MATLAB, Ansys Fluent, and VIAMD amongst others.

DEEP-SEA Wrap-Up: Advancing Software for European Exascale Systems

Stefano Markidis, PDC

The DEEP-SEA project was successfully completed earlier this year, culminating in a final review meeting in April 2024. DEEP-SEA, or "DEEP—Software for Exascale Architectures", has focused on creating an advanced programming environment for the next generation of European exascale systems, including JUPITER, the first European exascale supercomputer. The project has successfully adapted the software stack to support the highly heterogeneous computing and memory configurations expected in future high-performance computing (HPC) environments. One of DEEP-SEA's main accomplishments has been adapting all levels of the software stack. This includes low-level drivers, computation and communication libraries, resource management systems, and programming abstractions with associated runtime systems and tools. These components have been developed to ensure integration and optimisation across current and future European HPC architectures and systems. The KTH Royal Institute of Technology team involved in DEEP-SEA delivered tools and libraries, such as a Fast Fourier Transform (FFT) library based on the Data-Centric (DaCE) parallel programming framework for the GROMACS molecular dynamics code. This library enhances the performance and scalability of GROMACS simulations. Additionally, performance-portable matrix-multiply kernels have been developed for the Neko computational fluid dynamics (CFD) code, enabling efficient utilisation of the heterogeneous hardware configurations provided by future exascale systems. The insights and software technologies developed during DEEP-SEA will continue to influence the evolution of European exascale computing.

Above: JEDI, the first module of the JUPITER exascale system, at the Jülich Supercomputing Centre

Porting a Quantum Computer Simulator to AMD GPUs on Dardel

Stefano Markidis, PDC

Quantum computing is an important and emerging technology that solves problems beyond the reach of classical computers. Its applications are growing from cryptography to quantum chemistry simulations and machine learning. Quantum computer simulators have become essential tools for designing, validating, and optimising quantum hardware, software, and algorithms. Among the various quantum computing simulation techniques, the state vector simulator is particularly significant for quantum algorithm development. It represents the quantum state of a system as a complex vector, enabling continuous inspection and monitoring during circuit execution. As quantum computing simulation techniques evolve, the need for high-performance quantum computer simulators that are capable of leveraging high-performance computing (HPC) systems has intensified.

GPUs are an important HPC technology as they provide significant acceleration in quantum computer simulations and often outperform CPUs by over 40 times. While it is possible to run quantum computer simulations on NVIDIA GPUs, thanks to tools like "cuQuantum" (which is based on NVIDIA's CUDA interface that makes it possible for software to use particular types of GPUs for general-purpose computing), there has been a notable absence of similar support for AMD GPUs. This is a critical gap, especially given that many of the world's leading supercomputers (including Dardel at PDC, Frontier at the Oak Ridge Leadership Computing Facility in the USA, and LUMI at the CSC – IT Center for Science in Finland) are equipped with multiple AMD GPUs per node.

To address this gap, the challenge of porting the Google "qsim" quantum computer simulator to utilise AMD GPUs, particularly for the Dardel supercomputer, has been undertaken. By developing a "qsim" backend that leverages the AMD HIP (Heterogeneous-Compute Interface for Portability) programming interface and tools, the aim is to extend qsim's capabilities to include AMD GPU execution. The work bridges the gap for AMD GPU support in quantum simulations and includes a comprehensive performance evaluation comparing the AMD GPU backend with the existing NVIDIA GPU backend. This development is a significant step towards enhancing the versatility and performance of quantum computer simulators on supercomputers based on AMD GPUs, such as Dardel. More details about this work can be found in *"Enabling Quantum Computer Simulations on AMD GPUs: a HIP Backend for Google's qsim"* at https://doi.org/10.1145/3624062.3624223.

Below: Execution time of the qsim state vector simulator on AMD Trento CPU and AMD MI250X GPU varying the number of fused gates

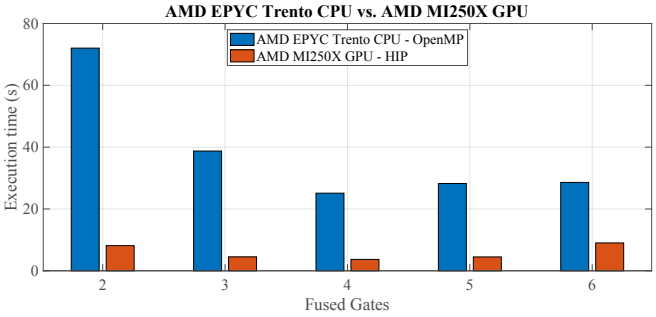

KTH Collaborations with AMD and HPE in the Development of Scientific Applications

Gert Svensson, PDC

The KTH Royal Institute of Technology (KTH) has recently entered into two collaborations, one with AMD and one with HPE, to enhance selected scientific applications where the code development is led by KTH and PDC, together with various other European institutes. The KTH collaboration with AMD focuses on the world-leading molecular dynamics code GROMACS and the partnership with HPE is starting on the VeloxChem quantum chemistry code, but the aim of both collaborations is more general and will include other applications, such as the Neko computational fluid dynamics code. The primary aim of the collaborations is to

optimise the performance of the research software when it is running on systems powered by modern GPU accelerators, such as the AMD Instinct™ MI250X GPU accelerators in the Dardel HPE Cray EX supercomputing system, which is hosted at PDC and funded primarily by the National Academic Infrastructure for Supercomputing in Sweden (NAISS).

The collaboration agreement between KTH, PDC and AMD started in December last year with one AMD software engineer (with a strong background in computational chemistry and molecular modelling software) working closely with research teams at KTH. The development of GROMACS is led by Berk Hess and Erik Lindahl from the Department of Biophysics at KTH and the Science for Life Laboratory (SciLifeLab) in Stockholm and the AMD software engineer is currently working with GROMACS developers at KTH Biophysics. Up to now, AMD has had its own GROMACS fork that relied on a HIP accelerator backend to sustain optimal performance on AMD Instinct accelerators. The AMD GROMACS backend comes in a stand-alone mode, complementary to the existing upstreamed CUDA and SYCL backends. KTH and AMD will now work together to mainstream, optimise, and maintain the HIP backend, thus enabling the GROMACS community to get optimal performance from AMD Instinct-powered devices. The KTH-AMD collaboration will also work on other world-leading software applications that are developed through PDC, such as Neko (a portable and scalable framework for high-fidelity computational fluid dynamics) and VeloxChem (a quantum chemistry application developed for the calculation of molecular properties and simulation of a variety of spectroscopies).

KTH's collaboration with HPE just started in May and will focus on VeloxChem. This work will be led by the HPE Centre of Excellence (CoE) for Europe, the Middle East and Africa (EMEA), which is based in Grenoble, France.

Patrick Norman, Director of PDC and head of the VeloxChem team, said "*PDC welcomes this joint effort with AMD and HPE on software optimisation as it will directly benefit both the Swedish and European research communities. This collaboration also aligns well with PDC's long-term and more general commitment to the development of a range of scientific software applications to make sure they run efficiently on all architectures, including the Dardel system.*" Several aspects are being emphasised in the collaboration, including supporting ongoing design and development activities such as fine-tuning these applications and making sure they work on current and future GPU accelerators, as well as developing new code features to benefit scientific research.

FAST. FLEXIBLE. FREE.
GROMACS

News from GROMACS

Magnus Lundborg, Department of Applied Physics, KTH, & Alessandra Villa, PDC

The end of February saw the minor release of GROMACS 2024.1, which comes with a number of bug fixes. Note that none of the issues that were addressed in that release were expected to have caused any errors in previous simulation results. At the start of May, there was another minor release: GROMACS 2023.5. Now the developers are working hard to bring new features into the 2025 release. To follow the development of GROMACS, you can look at the GROMACS GitLab repository: https://gitlab.com/gromacs. If you are interested in coding in GROMACS there will be a second edition of the "Learn to code in GROMACS" workshop in September. It will be held online, and you can check out the details here: https://www.gromacs.org/workshop.html. Registration will be opened nearer the date.

As always, you are welcome to the GROMACS user forum https://gromacs.bioexcel.eu if you need help with your simulations. If you would like to provide feedback to the developers, we will run a new GROMACS user survey in June. Lastly, you can keep yourself updated by following GROMACS news on https://www.gromacs.org.

EPICURE: High-Level Support for Research with Top-Flight EuroHPC Systems

Rossen Apostolov, PDC

Most scientists and engineers who use high-performance computing (HPC) systems to further their research or R&D activities are well aware that the need for advanced levels of support is growing as the top-flight HPC systems approach exascale performance. Tasks that enable users to exploit the full capacity of these systems (or as close to it as possible) are becoming progressively more vital. Some typical examples include adapting code to run on a different HPC system, refactoring code (rewriting or simplifying parts of it) to make it more efficient, and scaling code to make it run on larger numbers of parallel processing components (such as CPUs or GPUs).

Previous articles in the PDC newsletter, such as https://www.pdc.kth.se/publications/2023-no-2/research-software-engineer-teams-organising-the-most-advanced-level-of-user-support-at-the-serc-universities-1.1291981 and https://www.pdc.kth.se/publications/2022-no-2/supporting-research-software-development-1.1208807, have highlighted the need for top-level user support and establishing a profession for people with the advanced training, experience and subject area expertise to be able to provide these kinds of high-level services. The articles mention various considerations for Sweden in terms of providing advanced user support. The EuroHPC Joint Undertaking (EuroHPC JU, https://eurohpc-ju.europa.eu), which is responsible for providing a super-/quantum-computing and data infrastructure ecosystem for research in the European Union, also recognises the importance of high-level support for users of their top tier pre-exascale and petascale systems and also for coming exascale systems.

To meet this need, the EUROHPC JU has recently initiated a new project, called EPICURE, to provide top-level support for the researchers who use the EuroHPC top-tier systems. The project is being undertaken by a consortium, led by the CSC – IT Center for Science in Finland, which includes 16 partners from 14 different countries: the Barcelona Supercomputing Centre (Spain), the Cineca consortium (Italy), the VSB - Technical University of Ostrava (Czechia), the Institute of Information Sciences (Slovenia), LuxProvide (Luxembourg), the Institute for Systems and Computer Engineering, Technology and Science (INESC TEC, Portugal), Sofia Tech Park (Bulgaria), the University of Antwerp (Belgium), the KTH Royal Institute of Technology (Sweden), the AGH University of Krakow (Poland), the Technical University of Denmark (Denmark), the Jülich Research Centre (Germany), and GENCI (France). There are also two other organisations affiliated with the consortium: the Jožef Stefan Institute (Slovenia) and CINES (France).

The EPICURE project was launched on the 1st of February 2024 and will establish Application Support Teams (ASTs) based at organisations that host EuroHPC system now or in the future. These teams will work on application porting, optimisation and scalability improvements. In addition, EPICURE will also organise specialised events including training events, workshops, webinars and hackathons. Researchers whose projects are awarded an allocation on one of the EuroHPC systems will be offered access to EPICURE's support services.

EPICURE will also develop a European HPC Application Support portal that will serve as a single point of contact for researchers using HPC in both the public and private sectors, including small to medium-sized enterprises (SMEs). This portal will provide comprehensive information on EuroHPC systems, their architectures, access mechanisms, and the range of support services available.

KTH's efforts within EPICURE will be coordinated by PDC. PDC will be involved in all activities of the project with a special focus on code porting, enabling, and scaling, as well as the provision of technical expertise for the optimisation of the cases. For more information about EPICURE, see https://www.epicure-hpc.eu.

Running GROMACS Efficiently on LUMI Workshop 2024

Alessandra Villa & Szilárd Páll, PDC

The "Running GROMACS Efficiently on LUMI Workshop 2024" (https://bioexcel.eu/events/running-gromacs-efficiently-on-lumi-workshop-2024) was held online on the 24th and 25th of January 2024. The workshop came into being after brainstorming between GROMACS developers at PDC and the KTH Royal Institute of Technology (KTH) and the application expert teams at the CSC IT – Center for Science (CSC) in Finland about how to address the challenges of onboarding new users to the LUMI GPU partition (which is usually referred to as LUMI-G).

The workshop was run in collaboration with the BioExcel Centre of Excellence for Computational Biomolecular Research (BioExcel CoE) and gave practical tips on how to run GROMACS simulations efficiently on LUMI-G, that is, on AMD GPUs. The participants also learned how to assess and tune the performance of GROMACS. In addition, the course provided an overview of the LUMI architecture and heterogeneous parallelisation in GROMACS with special attention paid to AMD GPUs. In total, the workshop brought together 40 attendees from different European countries, with participants coming both from academia and the business/industry world.

There were two days full of lectures and hands-on sessions with a total of eight lecturers/trainers from CSC and KTH. The workshop started with an introduction to the LUMI architecture by Rasmus Kronberg from CSC (https://zenodo.org/records/10683366/files/lumi-arch.pdf) followed by a brief introduction to GROMACS by Alessandra Villa from PDC (https://zenodo.org/records/10683366/files/Short_Intro_GROMACS_LUMI.pdf). Then, Szilárd Páll from PDC addressed diverse aspects of GROMACS parallelisation/heterogeneous/GPU algorithms in depth (https://zenodo.org/records/10683366/files/LUMI-WS_Szilard-Pall_GROMACS-parallelization.pdf) and Andrey Alekseenko of the Department of Applied Physics at KTH spoke about AMD GPU support in GROMACS (https://zenodo.org/records/10683366/files/AMD-GPU-support-in-GROMACS.pdf). After the lunch break, Rasmus Kronberg from CSC described how to run programs on LUMI (https://zenodo.org/records/10683366/files/lumi-batch.pdf) and the first hands-on session started.

The following day opened with a lecture by Szilárd Páll on assessing and tuning the performance of GROMACS simulations (https://zenodo.org/records/10683366/files/LUMI-WS_Szilard-Pall_Asessing_and_tuning_GROMACS_performance_on_heterogeneous_systems.pdf), which was then followed by three consecutive hands-on sessions about GPU-accelerated simulations, scaling GROMACS across multiple GPUs, and ensemble parallelisation across multiple GPUs (https://zenodo.org/records/10683366/files/Running_molecular_dynamics_simulations_with_GROMACS_on_LUMI-HackMD.pdf).

Due to the similarity between the Dardel and LUMI architectures, the materials from this workshop will be very useful to Dardel users. If you would like to learn more, you can find the lecture material at https://zenodo.org/records/10683366 and the batch scripts and reference log files for the hands-on exercises can be downloaded from https://github.com/Lumi-supercomputer/gromacs-on-lumi-workshop/releases.

Above: LUMI system hosted at CSC in Finland

BioExcel News

Alessandra Villa & Rossen Apostolov, PDC

Since late 2015 when the BioExcel Centre of Excellence (CoE) for Computational Biomolecular Research (https://www.bioexcel.eu) was first established, the CoE has been supporting biomolecular researchers in their computational research. BioExcel works on improving the usability of technologies for biomolecular researchers and also the efficiency and scalability of important software packages for biomolecular research (like HADDOCK and GROMACS). In addition, BioExcel provides training to help life sciences researchers make good use of the available software and e-infrastructures and to help them to be aware of the "best practices" for the relevant combinations of software and hardware.

Following the success of the first BioExcel conference – which was held online in 2021 as a collaboration with the European Molecular Biology Organization (EMBO) – BioExcel is running its second conference this summer on "Advances in Biomolecular Simulations".

In the autumn, BioExcel will also organise the second event of the BioExcel Ambassador Program, the Atlantic Edition (involving Portugal, Spain, France and Ireland). The workshop will have lectures/tutorials on the four BioExcel core application (GROMACS, PMX, HADDOCK and BioBB library) and on AlphaFold (thanks to the recent agreement between BioExcel and the European Bioinformatics Institute, also known as EBI). The conference will be held at the University of Coimbra in collaboration with Portuguese and French National Competence Centres (NCCs).

BioExcel Webinar Series

The BioExcel webinars (see https://bioexcel.eu/category/webinar) feature notable developments in the field of computational biomolecular research. Amongst other speakers this spring, BioExcel has hosted:

- Mathieu Linares (PDC) who, together with Robin Skånberg from Linköping University, spoke about the VIAMD software for visual analysis of molecular dynamics,
- Magnus Lundborg (KTH) revealed the new features and developments in the latest GROMACS 2024 release, and

Current and past ambassador-driven activities

Carpathian Edition	Austria,Czechia, Hungary, Slovakia	Autumn 2023 (18-19 Oct)
Atlantic Edition	France, Spain, Portugal, Ireland	Autumn 2024
Balkan Edition	Bulgaria, Romania, North Macedonia, Serbia	Spring 2025
Adriatic Edition	Slovenia, Croatia, Montenegro, Italy	Autumn 2025
Aegean Edition	Turkey, Greece, Cyprus	Spring / Autumn 2026
Aurora Edition	Denmark, Norway, Finland, Iceland, Sweden	Spring / Autumn 2026

Above: Now BioExcel is currently organising the Atlantic edition of the Ambassador Program and, before summer, the centre will start gathering countries for the Balkan edition.

- the COLVARS development team who discussed the use of enhanced-sampling in collective variable space with the Colvars library, which is available in standard GROMACS 2024 releases.

In addition, at the time of writing, it is planned that on the 28th of May, Adam Hospital from BioExcel will speak about "Using interactive Jupyter Notebooks and BioConda for FAIR and reproducible biomolecular simulation workflows".

Parallel Programming Using Message Passing with MPI Online Course

Juan De Gracia Triviño, PDC

As part of the National Academic Infrastructure for Supercomputing in Sweden (NAISS) training programme, PDC collaborated with the High Performance Computing Center North (HPC2N) at Umeå University and the Centre for Scientific and Technical Computing (LUNARC) at Lund University to offer an online course on parallel programming using message passing with MPI at the end of last year. The course concluded with remarkable success, drawing keen interest from Swedish academia and industry. Designed for beginners with no previous experience in parallel computing, the course offered an engaging introduction to the world of message passing and distributed memory computing, both of which are cornerstones of high-performance computing (HPC) across a variety of computer architectures.

About forty people participated in the course, and they displayed a strong preference for Python (mpi4py) over traditional languages like C++ and Fortran, with the latter being the least popular amongst the participants. This trend underscores the growing appeal of Python in the parallel computing domain, reflecting its accessibility and the expanding ecosystem of scientific libraries.

The course explored the essentials of message passing, a programming model that has gained widespread adoption in massively parallel HPC. By covering topics from point-to-point communications to non-blocking and collective communications calls, the course equipped learners with the foundational knowledge and skills to navigate the complexities of current computer architectures. From multi-core desktops to the fastest HPC systems in the world, message passing is pivotal in harnessing the potential of several hundred thousand processing elements.

Theoretical lectures were complemented by live demonstrations and practical sessions, enriching the participants' learning experience and deepening their understanding of the material. This hands-on approach ensured that by the end of the course, attendees were not only familiar with the key MPI calls but also capable of writing their own MPI programs at an intermediate level.

The successful conclusion of this course highlights the growing interest and importance of parallel computing skills in today's technology-driven landscape. PDC was encouraged by the enthusiastic participation and the positive feedback received, and we look forward to offering further courses that meet the evolving needs of the global computing community.

ENCCS
EuroCC National Competence Centre Sweden

ENCCS Co-Organises Nordic HPC Conference "Supercomputing, The Gateway to AI"

Thor Wikfeldt, ENCCS

Supercomputing is no longer confined to academic research; an increasing number of private companies and public sector authorities are harnessing high-performance computing (HPC) systems for their research, development, and innovation endeavours. Recent advancements in large-scale artificial intelligence (AI) models are particularly drawing attention to HPC from entirely new potential user groups in both the private and public sectors. These groups realise that relying solely on cloud computing may not be economically feasible for their projects. HPC offers them the ability to think bigger!

The EuroCC National Competence Centre Sweden (ENCCS) serves as the central hub for private companies and public sector authorities in Sweden that are interested in leveraging HPC for their projects. In recent months, ENCCS has assisted numerous companies and public sector organisations to apply for and then utilise free supercomputing access to various EuroHPC systems. These initiatives have accelerated development projects on diverse topics such as advanced stainless steels, advanced X-ray source technologies, urban air pollution modelling, digitalisation of historical weather observations, AI-powered game development, and real-time adaptation of AI models for autonomous driving.

Now, ENCCS is teaming up with EuroCC colleagues from other Nordic countries to organise a conference on 2-3 September 2024 aimed at professionals and leaders in industry, high-tech start-ups, and small and medium-sized enterprises (SMEs) interested in leveraging AI, HPC, and quantum computing technology to drive their businesses forward. Set against the backdrop of Copenhagen's charming Tivoli Gardens, the conference promises two days of insightful discussions, networking opportunities, and knowledge sharing. The conference's theme is how supercomputing serves as an enabling technology for large-scale AI and will delve into how AI and HPC can help companies succeed in the digital age.

If you are part of a company and believe AI and HPC could enhance your product development, sign up for this free conference and join us in Copenhagen for inspirational talks, panel discussions, and networking opportunities! And if you are not working in industry, spread the word to your friends who are! Registration and further details can be found at https://enccs.se/events/nordic-industry-days-2024.

Introduction to PDC Systems Course

Juan De Gracia Triviño, PDC

Welcome to the first "Introduction to PDC Systems" course for 2024! This spring, the event was fully booked, welcoming 40 participants from a variety of backgrounds for a comprehensive two-day workshop.

The first day was dedicated to providing a thorough overview of PDC, with sessions covering the basics of Linux and the Slurm queuing system, alongside an introduction to the graphical interface with ThinLinc. Reflecting the increasing popularity of programming tools, we also introduced Python and Jupyter notebooks, as well as Singularity for managing containers on Dardel.

On the second day, we shifted our focus to technical skills, including the fundamentals of CPU and GPU compilation on Dardel, equipping our more advanced users with the necessary skills to compile their own software. For the first time, we expanded our curriculum to include an introduction to various scientific fields where our application experts can provide targeted support. This new section is in line with the strategic shift of the National Academic Infrastructure for Supercomputing in Sweden (NAISS) towards offering support based on areas of expertise rather than being cluster-specific. The areas of expertise presented were machine learning, molecular modelling, materials theory, computational fluid dynamics, and bioinformatics. We also took the opportunity to showcase flagship software that is developed or co-developed at PDC, such as VeloxChem, GROMACS, and Neko.

The participation level was good, demonstrating strong engagement throughout the event. The feedback we received has been invaluable and will be instrumental in shaping our training programs in the coming years. PDC is committed to continuously improving our offerings, and the organisers thank all participants for their active involvement and insightful contributions to a successful introduction. We look forward to seeing you at further training events later this year! Stay tuned to the PDC Events web page https://www.pdc.kth.se/about/events or follow PDC on Facebook at https://www.facebook.com/kth.pdc.

NAISS

The Work of NAISS

Björn Alling, NAISS

The National Academic Infrastructure for Supercomputing in Sweden (NAISS) has compiled its annual report for 2023, which demonstrates the significance of NAISS resources for researchers in Sweden. Some highlights from the report are that 7,630 people were active in the infrastructure last year, and one in three of the users were female. Also, NAISS allocated resources to 1,924 projects spanning 32 research areas from across all the large research universities in Sweden, and more than 1,300 publications acknowledged NAISS or its predecessor, the Swedish National Infrastructure for Computing (SNIC), in 2023.

Preparations for the new Swedish-EuroHPC system Arrhenius are ongoing. As it is such a large undertaking, there are many juridical and economic details, in addition to technical specifications, that have to be sorted out. The official procurement process will start soon.

The kick-off meeting for NAISS' new way of organising user support was held in April. The new approach is based on a collaboration between NAISS and eleven large universities in Sweden. The goal is to create a unified support organisation where each user can be helped by the expert with the most relevant knowledge, regardless of where the researcher or expert is working or where the hardware might be situated. NAISS is in the process of finalising bilateral branch agreements for organising this distributed expert support.

The NAISS User Forum and All-Hands Meeting will be held in Uppsala on the 1st-3rd of October. You can register now at https://www.naiss.se/event/user-forum-2024 and https://www.naiss.se/event/all-hands-meeting-2024. We hope to see you all there!

Below: Illustration of the new functionalities in VIAMD for calculating and plotting coordinates and densities across a box as exemplified on a sodium channel (Dataset courtesy of Lucie Delemotte's Lab, Science for Life Lab/KTH)

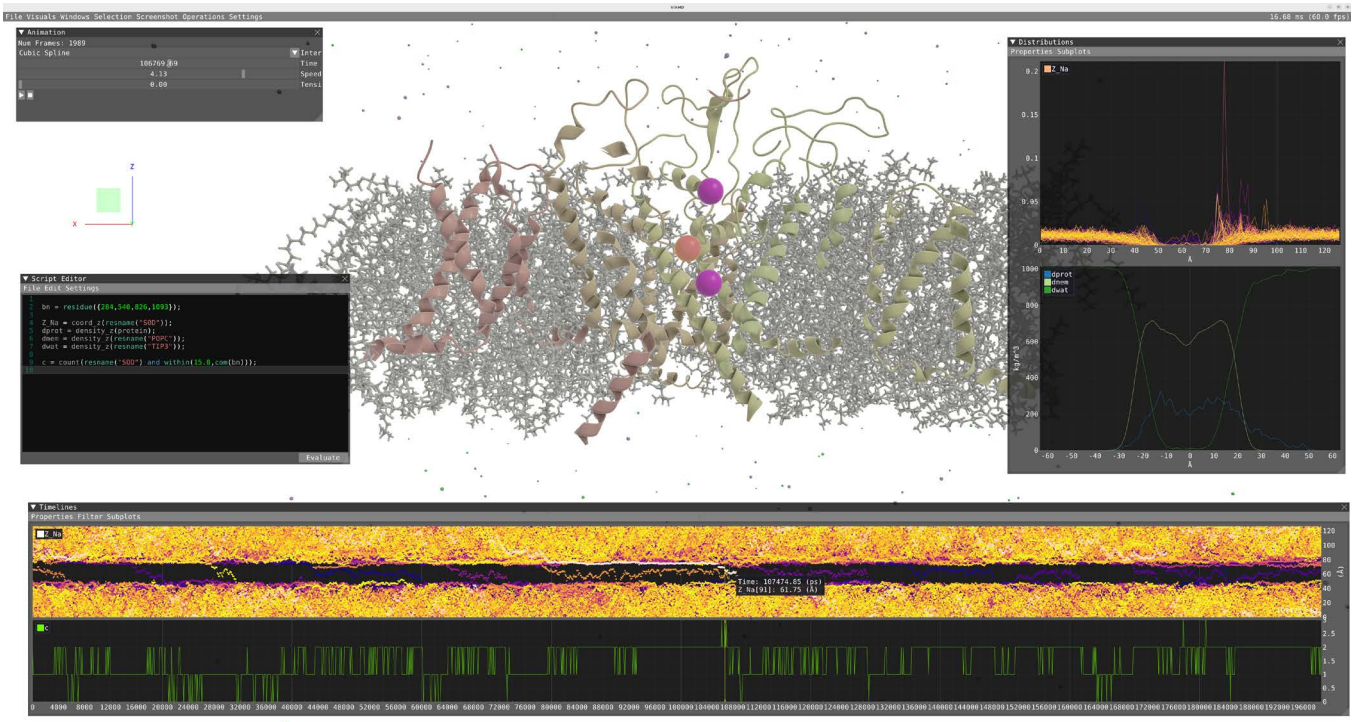

VIAMD Update

Mathieu Linares, PDC

Since the publication of an article in November 2023 which presented the VIAMD code (*"VIAMD: a Software for Visual Interactive Analysis of Molecular Dynamics"* – see https://doi.org/10.1021/acs.jcim.3c01033), the VIAMD development team have implemented several new features including handling of trajectories from the Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS) molecular dynamics program and new functionality to extract, calculate and plot coordinates and densities across a box as exemplified on the ion-conducting membrane shown above.

The team is currently working on adding support for VeloxChem files to VIAMD in order to efficiently evaluate and render molecular orbitals under different modalities, making it possible, for instance, to undertake interactive analysis of absorption spectra of large molecular systems in terms of natural transition orbitals.

When it comes to dissemination, the team has created a range of tutorial videos that are available on a YouTube playlist (https://bit.ly/4aRsPrh) and an X account: @VIAMD_(https://x.com/VIAMD_). If you would like to start using VIAMD, you can find it on Github (https://github.com/scanberg/viamd).

Hybrid Workshop: Space Plasma Simulations with Vlasiator on LUMI

Stefano Markidis, PDC

The Plasma-PEPSC Centre of Excellence, in collaboration with the EuroCC National Competence Centre Sweden (ENCCS), hosted a hybrid training workshop from 13-15 February in Stockholm. This workshop was designed to bring together code developers, researchers, and research software engineers engaged in space plasma simulation technologies. The workshop participants had the unique opportunity to learn how to use and program the Vlasiator package (a state-of-the-art tool for plasma simulation and analysis especially designed for space simulations) on the EuroHPC LUMI supercomputer. Expert Vlasiator developers from the University of Helsinki and the Finnish Meteorological Institute provided comprehensive training throughout the workshop on utilising the Vlasiator code's

Below: Slide from Space Plasma Simulation workshop

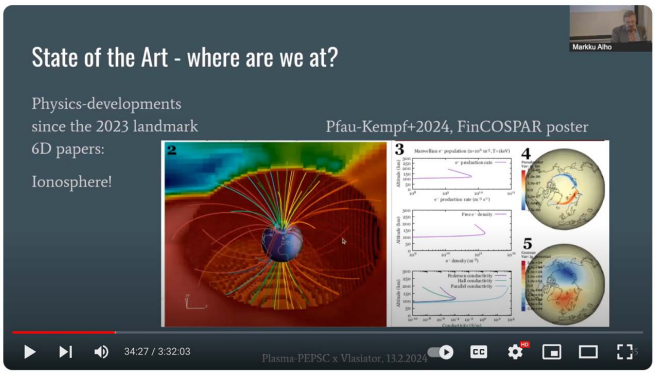

capabilities, including data analysis tools to understand space simulations. The sessions were structured to cater to both beginner users and experienced space researchers. In addition to hands-on training on using Vlasiator on the LUMI supercomputer, the workshop featured a series of presentations and discussions highlighting the latest advancements in plasma simulation techniques for space simulations. Attendees from the space physics and space weather community at the KTH Royal Institute of Technology and Uppsala University could network with peers, share their experiences of space simulations, and explore potential collaborations across different plasma physics and simulation projects. The hybrid format of the event also allowed remote participants from France and Germany to engage in the training activities, broadening the workshop's reach and impact. The lectures about using Vlasiator on the LUMI system are collected on the Plasma-PEPSC YouTube channel (https://www.youtube.com/@PlasmaPEPSC).

VeloxChem: Electron Repulsion Integrals and Fock Matrix Formation on GPUs

Xin Li, PDC

Quantum chemistry offers invaluable insight into and understanding of the complexities of electronic structure and can, therefore, serve as a very powerful tool in scientific research and theoretical design of molecules and materials. The major computational complexity in quantum chemistry lies in the evaluation of the electron repulsion integrals (ERI) and the formation of the so-called Fock matrix, which formally scales as $O(N^4)$, with N being the system size. In practice, this time-consuming task can be done much more efficiently by taking advantage of the screening of the ERIs, and utilisation of GPUs can provide significant speedup.

The VeloxChem program [1] is an open-source quantum chemistry software application developed at the KTH Royal Institute of Technology (KTH), including at PDC. Recently, the VeloxChem team at KTH and PDC implemented ERI evaluation and Fock matrix formation on GPUs to further push the limit of the sizes of the systems that can be routinely studied. In our implementation, the ERIs are evaluated by the Obara-Saika recurrence relation [2], and, on top of that, the formation of the Fock matrix is done by contracting the ERIs with the density matrix. The Fock matrix can be split into two contributions due to different contracting patterns, which are called the Coulomb and the exchange contributions, respectively. To take advantage of the $O(N^2)$ scaling of the Coulomb contribution and the $O(N)$ scaling of the exchange contribution, we implemented the Fock matrix formation based on the direct self-consistent field implementation [3] and the pre-selective screening approach [4].

We tested the performance of the Fock matrix formation and ERI evaluation in VeloxChem by running Hartree-Fock calculations for a series of water clusters on Dardel. The largest water cluster used in the benchmark contains 3,879 atoms, and the number of contracted basis functions exceeds 31,000 with the def2-SVP basis set [5]. By plotting the time spent in Fock matrix formation with respect to the size of the system (indicated by the number of contracted basis functions), we can see that the scaling of the computational cost on GPUs is $O(N^{1.76})$, which falls between the expected scaling of the Coulomb and exchange contributions. Due to such beneficial scaling, the formation of the Fock matrix for the largest water cluster took around 850 seconds on a Dardel GPU node. This opens up the possibility of routine study of large and complex chemical systems.

Below: Benchmark of Fock matrix formation on Dardel CPU and GPU compute nodes using water clusters and the def2-SVP basis set

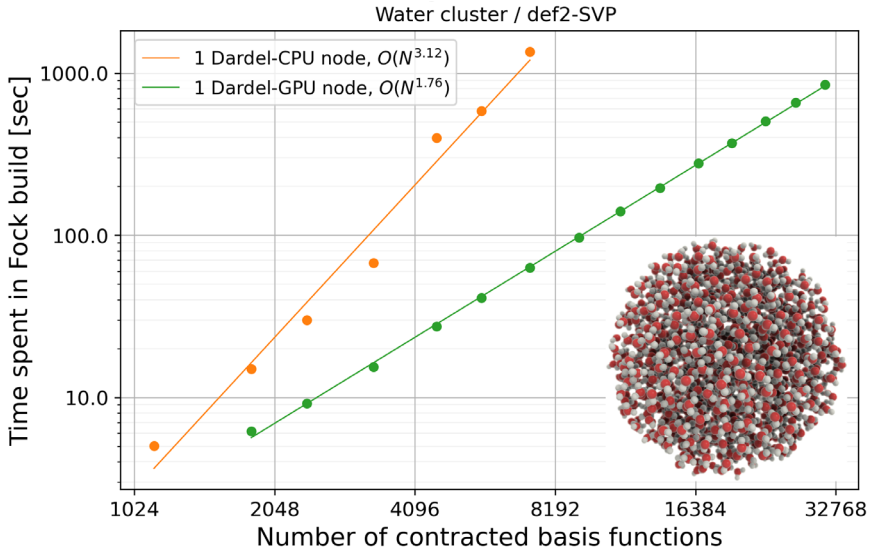

References

1. Z. Rinkevicius, X. Li, O. Vahtras, K. Ahmadzadeh, M. Brand, M. Ringholm, N. H. List, M. Scheurer, M. Scott, A. Dreuw, P. Norman, "VeloxChem: a Python-driven density-functional theory program for spectroscopy simulations in high-performance computing environments". WIREs Comput. Mol. Sci. 2020, 10:e1457. https://doi.org/10.1002/wcms.1457
2. S. Obara, A. Saika, "Efficient recursive computation of molecular integrals over Cartesian Gaussian functions". J. Chem. Phys. 1986, 84, 3963-3974, https://doi.org/10.1063/1.450106
3. I.S. Ufimtsev, T. J. Martinez, "Quantum Chemistry on Graphical Processing Units. 2. Direct Self-Consistent-Field Implementation". J. Chem. Theory Comput. 2009, 5, 1004-1015, https://doi.org/10.1021/ct800526s
4. J. Kussmann, C. Ochsenfeld, "Pre-selective screening for matrix elements in linear-scaling exact exchange calculations". J. Chem. Phys. 2013, 138, 134114, https://doi.org/10.1063/1.4796441
5. F. Weigend, R. Ahlrichs, "Balanced basis sets of split valence, triple zeta valence and quadruple zeta valence quality for H to Rn: Design and assessment of accuracy". Phys. Chem. Chem. Phys., 2005, 7, 3297-3305, https://doi.org/10.1039/b508541a

PDC-Related Events

BioExcel Conference: 21-23 Oct. 2024, Brno, Czech Republic
https://bioexcel.eu/events/2nd-bioexcel-conference-on-advances-in-biomolecular-simulations

BioExcel Workshop: 21-23 October 2024, Madrid, Spain
https://bioexcel.eu/events/workshop-upskilling-and-training-users-in-science-and-research

Learn to code in GROMACS: 10-12 September 2024 (online)
https://www.gromacs.org/workshop

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 https://eosc-portal.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