
ASHPC22
Austrian-Slovenian HPC Meeting
Grundlsee: May 31 – June 2, 2022

Austrian-Slovenian HPC Meeting 2022 – ASHPC22

Grundlsee, May 31 – June 2, 2022

https://ashpc.at

Welcome to ASHPC22

We would like to welcome you all at the second Austrian-Slovenian HPC Meeting – ASHPC22.

ASHPC22 continues the tradition of the annual Austrian HPC Meetings (held 2015 – 2020) and the first Austrian-Slovenian HPC Meeting held 2021: Bringing together scientists and technicians with a background and interest in supercomputing. High-Performance Computing (HPC) operates at the limits of computationally feasible problems and helps to conquer new territory of science. Learning about current limitations and exchange of ideas how to address these issues is key for the further development in scientific and technological competitiveness. Therefore, ASHPC22 is an excellent opportunity to learn about the latest research results and to exchange ideas between the users and providers of HPC resources.

We are happy that after a long time with remote meetings we are again able to hold an in-person meeting in Grundlsee, surrounded by splendid mountains and marvelous lakes.

ASHPC22 will feature a full program with keynote and contributed talks as well as poster presentations and you will be able to engage in lively discussions with the Austrian and Slovenian HPC community and our international guests.

Financial support for ASHPC22 is provided by EuroCC Austria and by the Slovenia National HPC Competence Centre.

ASHPC22 is organized by EuroCC Austria – National Competence Centre for Supercomputing, Big Data and Artificial Intelligence, Austria, and EuroCC Slovenia in cooperation with the Vienna Scientific Cluster (VSC), Austria, the Research Area Scientific Computing in Innsbruck, Austria, the Institute of Information Science in Maribor (IZUM), Slovenia, and the Slovenian consortium for high-performance computing (SLING).

We wish you an inspiring and enjoyable time at this year's ASHPC22 in Grundlsee and online.

The Organizing Committee:

Mariella Wittich (meeting chair): EuroCC Austria – National Competence Centre for Supercomputing, Big Data & Artificial Intelligence, Austria

Eduard Reiter: Research Area Scientific Computing, University of Innsbruck, Austria
Davor Bračko: Institute of Information Science (IZUM), Slovenia

The Programme Committee:

Andrej Filipič (programme chair): Department of Experimental High Energy Physics – F9, Jozef Stefan Institute, Slovenia

Uroš Lotrič: Faculty of Computer and Information Science, University of Ljubljana, Slovenia
Zoran Ren: University of Maribor, Slovenia
Claudia Blaas-Schenner: VSC Research Center, TU Wien, Austria
Christoph Dellago: Faculty of Physics, University of Vienna, Austria

Schedule

Tuesday, May 31, 2022

Start		Title
07:30	BREAKFAST	
09:00	Claudia Blaas-Schenner	WELCOME to the EuroHPC Access Workshop
09:05	Andrej Filipčič	Overview of (already existing) EuroHPC systems
09:20	Alessandro Marani	LEONARDO - current status and what it is to us (the consortium)
09:30	Klara Meštrović & Krishnakshi Bhuyan	Calls for application and review process
10:00	Philipp Gschwandtner	Tipps and tricks from the user perspective with PRACE applications
10:30	COFFEE	
11:00	Tomáš Kozubek & Teo Prica	High-level support on sites, how to choose the appropriate HPC system and call
11:30	EuroHPC-Workshop	Round table with officers and experts, Q&A
12:00	LUNCH	
13:00	Andrej Filipčič	WELCOME to ASHPC22
13:15	Evangelos Floros	EuroHPC Joint Undertaking State-of-Play (KEYNOTE TALK)
14:00	Bernhard Semlitsch	Performance Modelling of Ship Propellers by Computational Flow Simulation
14:20	Felix Reuß	Crunching Petabytes with the VSC: The processing and analysis of global satellite imagery
14:40	Soner Steiner	Collaboration with the SME TAILSIT during the SHAPE project PARTS: Electromagnetic simulations with the finite/boundary element method for large systems using HPC
15:00	COFFEE	
15:30	Harald Höller-Lugmayr	From silicon to silicon
15:50	Peter Kandolf	Austrian DataLAB and Services
16:10	Alois Schlögl	Where is the sweet spot? A procurement story of general purpose compute nodes
16:30	BREAK	
16:45	Dejan Valh	Vega EuroHPC, 1st year in service
17:05	Siegfried Höfinger & Jan Zabloudil	Next generation Vienna Scientific Cluster (VSC-5)
17:25	Markus Hickel	VSC Storage, Present and Future
17:40	Irene Reichl & Simeon Harrison	Current Developments on the Vienna Scientific Cluster (VSC)
18:00	Alexander Ostermann (moderator)	Plenary discussion with users about their needs and wishes for future HPC systems
19:00	DINNER	

Wednesday, June 1, 2022

Start		Title
07:30	BREAKFAST	
09:00	Lubomir Riha	Selected Cases of GPU Accelerated Parallel Applications at IT4Innovations (KEYNOTE TALK)
09:40	Uroš Lotrič	Energy-Efficient Computing with Approximate Tensor Core Units
10:00	Josef Weinbub	Shared-Memory Fast Marching Method for Re-Distancing on Hierarchical Meshes
10:20	Ioannis Vardas	mpisee: MPI Profiling for Communication and Communicator Structure
10:40	Sascha Hunold	MPI Performance Tools under the Microscope: A Thorough Overhead Analysis
11:00	COFFEE	
11:30	Sarah Stryeck	NCC Austria - A National Competence Center for High-performance computing, High-Performance Data Analytics and Artificial Intelligence
11:50	Pavel Tomšič	Slovenian National Competence Centre HPC
12:10	Ferenc Szani	Vision of a HPC NCC today
12:30	LUNCH	
14:00	Albert Frisch	HPC-Integration of an Ion-Trap Quantum Computer (KEYNOTE TALK)
14:45	Markus Wallerberger	Romeo and Julia: HPC, ranking, and the many-electron problem
15:05	Philipp Gschwandtner	The Cluster Coffer: Teaching HPC on the Road
15:25	COFFEE	
15:55	Richard Gerber	NERSC, Perlmutter and HPC in the U.S. Department of Energy Office of Science (KEYNOTE TALK)
16:40	BREAK	
16:45	POSTER LIGHTNING TALKS	
17:10	POSTER SESSION (parallel)	
	László Ligeti	Case study: How to select the right Software mix for SMEs
	Vladislav Kashansky	Heterogeneous Workflows Scheduling in the Computing Continuum Systems (ONLINE)
	Jakob Merljak	arcControlTower for distributed HPC job management
	Renáta Rusková	HPC in design of devices for chiral nanotechnology (ONLINE)
	Dušan Račko	The free volume in PVME-water mixture as obtained from HPC simulations (ONLINE)
	Shokirbek Shermukhamedov	Neural network potentials for fusion material research

	Eduard Vorobyov	Computing the gravitational (electrostatic) potential on nested Cartesian meshes using the convolution method (ONLINE)
	Stefano Elefante	Benchmarking using Relion GPU workloads
	Leon Deutsch	Catch22 = Biomed 2022 (ONLINE)
	Leon Kos	Interactive Hands-on Introduction to Parallel Programming
	Pavel Tomšič	Erasmus+ project: SCtrain
17:10	LIVE DEMO (parallel)	
	Philipp Gschwandtner	The Cluster Coffer: Teaching HPC on the Road
18:00	DISCUSSION	
19:00	DINNER	

Thursday, June 2, 2022

Start		**Title**
07:30	BREAKFAST	
09:00	Matteo Ambrosetti	Multiscale Methodologies for Electrolyte Characterization (KEYNOTE TALK)
09:45	Janez Povh	High-performance electricity consumption prediction
10:05	Aleš Zamuda	Speeding up Vectorized Benchmarking of Optimization Algorithms
10:25	COFFEE	
11:00	Dominik Gehringer	Modelling-aided materials discovery at Montanuniversität Leoben
11:20	Gerhard Kahl	On the emergence of quasi-crystalline structures in a Wigner bilayer system
11:40	Davor Sluga	Development of GPU accelerated molecular software (CmDock) for efficient high-throughput virtual screening
12:00	LUNCH	
13:30	Domen Verber	High-Performance Computing with Relational Database Management Systems
13:50	Giovanna Roda	Distributed computing for everyone
14:10	Jean Pierre Panziera	Atos - Future HPC technologies (KEYNOTE TALK)
14:55	CLOSING (tentative)	
15:00	COFFEE (tentative)	
15:30	END of ASHPC22 (tentative)	

Contents

EuroHPC Access Workshop

The European High Performance Computing Joint Undertaking (EuroHPC JU; `https://eurohpc-ju.europa.eu`) currently builds up a world-class supercomputing ecosystem in Europe, consisting of five petascale and three pre-exascale systems that have recently started or are just about to start regular user operation. The next Call for Hosting Entity for a High-End Supercomputer has closed and we envisage more and more computing capabilities to be available for the benefit of the European Research and Development community in the near future. The objective is to reach exascale capabilities, i.e., 10^{18} Flop/s, by 2022/2024 and build 'hybrid' machines that blend the best of quantum and classical HPC technologies with the first state-of-the-art pilot quantum computers by 2025.

In addition, the EuroHPC JU aims to widen the use of HPC infrastructures to a large number of public and private users wherever they are located in Europe and support the development of key HPC skills, education and training for European science and industry, e.g., by creating a network of national HPC Competence Centres (`https://www.eurocc-access.eu`) to ease access to European HPC opportunities in different industrial sectors and deliver tailored solutions, as well as setting up the first pan-European Master of Science programme for HPC to develop HPC talents in Europe (`https://eumaster4hpc.uni.lu`).

In this workshop you can learn and discuss about benefiting from the new European HPC Systems. The EuroHPC Access Workshop will offer information for organizational decision-making, the application process and end-user applications.

This page has been intentionally left blank.

Keynote Talk:

EuroHPC Joint Undertaking State-of-Play

Evangelos Floros

EuroHPC Joint Undertaking, European Union; Evangelos.FLOROS@eurohpc-ju.europa.eu

The European High Performance Computing Joint Undertaking (EuroHPC JU) is a legal and funding entity, established by the European Commission (EC) in 2018 [1]. The goal of the JU is to coordinate actions, resources and policies between the European Commission and EU member states, in order to pursue in a concerted manner, the development of HPC in Europe, making the EU a global leader of supercomputing. Members of the JU are the European Commission, 32 Participating States and 3 private members.

EuroHPC's activities evolve around 2 main pillars: a) Infrastructure and b) Research & Innovation.

The goal of the Infrastructure pillar is the development of a world-class, state-of-the-art supercomputer ecosystem, able to satisfy the computational requirements of European HPC applications. For this purpose, EuroHPC is procuring supercomputing systems which are installed in Participating States across the EU. Currently, the JU has successfully procured, deployed and has made operational 5 systems with 3 additional systems under installation.

EuroHPC implements an ambitious research and innovation program which covers all aspects of the HPC-domain: development of novel technologies, advancement of algorithms and applications, supporting applications to move to exascale etc. Many actions are also aiming towards improvement of skills and training.

EuroHPC's mandate was renewed in 2021 [2] through an updated and expanded program and a committed budget from the European Union that will exceed 7 billion Euro. The two core pillars of Infrastructure and R&I are now complimented with activities in three additional areas: Infrastructure Federation, Skills & Training and International Collaboration.

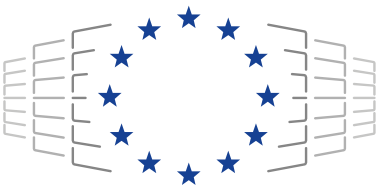

EuroHPC
Joint Undertaking

References

[1] Council Regulation (EU) 2018/1488 of 28 September 2018 establishing the European High Performance Computing Joint Undertaking

[2] Council Regulation (EU) 2021/1173 of 13 July 2021 on establishing the European High Performance Computing Joint Undertaking and repealing Regulation (EU) 2018/1488

Performance Modelling of Ship Propellers by Computational Flow Simulation

Bernhard Semlitsch

Institute for Energy Systems and Thermodynamics, TU Wien, Getreidemarkt 9, 1060 Wien, Austria

Climate change causes more extreme weather conditions, i.e. hotter summers and colder winters, and less distributed rain. That has consequences on the amount of water carried in rivers. Because of the shallow water levels occurring during longer-lasting dry seasons, the draft of ships must be restricted to prevent damage to the riverbed. Under such conditions, cargo ships are forced to limit the load below their capacity. The reduction of inland shipping has wide-ranging consequences because some ships can transport 3,000 tons of cargo, equivalent to 150 trucks. Road transportation generates more CO_2eq emissions than inland shipping [1]. Therefore, inland vessels with a lower draft are in demand, which can continue the undisturbed transport even in low water periods.

One of the determining factors for the draft of the ship is its propeller, which has to be fully submerged and capable of pushing the vessel with a certain thrust. The thrust is proportional to the quadratic forward velocity of ships, the density of the river water, and the propeller area. Hence, the propellers are designed with the largest possible cross-section because the ship shall reach its target as fast as possible. The diameter of the propeller defines, therefore, the draft of inland vessels. This project aims to replace the large propeller with two smaller propellers while the thrust and the total propeller area are preserved. Thereby, the ship's draft can be reduced under the challenge of providing the same thrust generated by a single propeller. The two propellers are arranged close to each other, enabling the utilisation of the swirling energy that is lost with single propellers. The generated swirl component of the individual propellers shall be converted into thrust using the fluid interactions and a surrounding nozzle.

Numerical flow simulations are employed to find the optimal propeller arrangement and nozzle shape, reducing the ship's draft. The simulation methodology of the flow phenomena occurring around the propeller must correctly replicate turbulent flow features and cavitation in a multiphase environment. Further, the rotation of the propellers needs to be modelled with respect to the static geometries. The close arrangement of the propellers represents the challenge for the numerical mesh generation because the conventional sliding-mesh approach cannot handle such tight spacings. Another methodology is the overset grid technique [2], where the flow quantities are computed on overlapping meshes moving independently on top of each other. The connection between the grids is built by interpolating the cell and point values. The steep pressure gradients occurring due to cavitation and the turbulent flow demand a fine mesh resolution. Therefore, the governing equations, i.e. the Navier-Stokes equations, are numerically solved with high accuracy on fine overset grids to meet the associated requirements.

The application of such numerical shape optimisations for the industrial design of ship propellers is presented. The computational requirements of the individual numerical models and assumptions are stated and analysed. The need for high-performance computers is discussed when the simulation complexity of optimisation steps considering the correct physical representation is as high as for ship propeller design.

We acknowledge funding as a PRACE SHAPE project by the EU's Horizon 2020 research and innovation programme (2014-2020) under grant agreement 823767. The work is achieved using the resources of VSC-4 provided by the Vienna Scientific Cluster (VSC).

References

[1] Van Fan, Y., Perry, S., Klemeš, J.J. and Lee, C.T., J. Clean. Prod., **194**, 673 (2018)

[2] Tang, H. S., Jones, S. C., and Sotiropoulos, F., J. Comput. Phys., **191**(2), 567 (2003)

Crunching Petabytes with the VSC: The processing and analysis of global satellite imagery

Felix Reuß[a], **Bernhard Bauer-Marschallinger**[a], **Claudio Navacchi**[a], **Florian Roth**[a], **Mark Tupas**[a], **Christoph Reimer**[b], **Senmao Cao**[b], **and Wolfgang Wagner**[a,b]

[a] *Department of Geodesy and Geoinformation, TU Wien, Austria*
[b] *EODC Earth Observation Data Centre for Water Resources Monitoring, Austria*

Earth observing satellites and their imagery deliver rich and global data about our planet's surface. Microwave remote sensing data, and in particular those acquired by synthetic-aperture radars (SARs), constitute an important input for a wide range of meteorological or geophysical applications. These include environmental monitoring, soil moisture retrieval, flood mapping, yield prediction, and many more. In contrast to optical remote sensing, a SAR sensor can operate under all weather and light conditions, thus allowing gap-free measurements and denser time series. The current ESA Sentinel-1 mission with its SAR satellites offers global data with an unparalleled high temporal and spatial resolution, down to 10 m pixel sampling [1].

The Microwave Remote Sensing research group from the TU Wien GEO Department (TUW) in cooperation with the Earth Observation Data Center (EODC), processed the entire Sentinel-1 data archive to provide analysis ready data (ARD) to users [2]. This effort was carried out on the Vienna Scientific Cluster (VSC) and included the processing of 1.5 PB input satellite data to a 700 TB ARD archive.

TUW and EODC utilize the VSC infrastructure in numerous services and projects to provide services based on these data. For example, for the Sentinel-1 Global Backscatter Model [3], we processed a raw data volume of 360 TB, via 700 TB intermediate data, to a total of 80 TB final global mosaic layers. These describe the entire world's land surface and are visualized in a data viewer to allow easy access to the users. Within the Copernicus Emergency Management Service global flood monitoring project, TUW and EODC have a major contribution in providing input parameters processed on the VSC. Here, we retrieved several global parameters from Sentinel-1 time series serving as an input to flood mapping algorithms.

To allow the (time-)efficient processing of the tremendous amount of data, TUW developed processing workflows optimized for parallelization on high performance clusters (HPC) such as VSC. However, we're still facing challenges regarding efficient reading and writing of data, workload distribution, and output quality assessment. Within our contribution, we aim to share our gained experience and insights on processing large amounts of geospatial data on the VSC cluster, and seek for exchange with the HPC-community to improve our processing chains.

References

[1] Torres, R.; Snoeij, P.; Geudtner, D.; Bibby, D.; Davidson, M.; Attema, E.; Potin, P.; Rommen, B.; Floury, N.; Brown, M.; et al. GMES Sentinel-1 Mission. Remote Sens. Environ. 2012, 120, 9–24, doi:10.1016/j.rse.2011.05.028

[2] Wagner, W.; Bauer-Marschallinger, B.; Navacchi, C.; Reuß, F.; Cao, S.; Reimer, C.; Schramm, M.; Briese, C. A Sentinel-1 Backscatter Datacube for Global Land Monitoring Applications. Remote Sens. 2021, 13, 4622, doi:10.3390/rs13224622

[3] Bauer-Marschallinger, B.; Cao, S.; Navacchi, C.; Freeman, V.; Reuß, F.; Geudtner, D.; Rommen, B.; Vega, F.C.; Snoeij, P.; Attema, E.; et al. The Normalised Sentinel-1 Global Backscatter Model, Mapping Earth's Land Surface with C-Band Microwaves. Sci. Data 2021, 8, 277, doi:10.1038/s41597-021-01059-7

Collaboration with the SME TAILSIT during the SHAPE project PARTS: Electromagnetic simulations with the finite/boundary element method for large systems using HPC

Soner Steiner

VSC Research Center / EuroCC Austria, TU Wien, Austria

TAILSIT is a company based in Styria, Austria, that produces custom-fit simulation software tools for electromagnetic problems and structural analysis. During this SHAPE project TAILSIT's electromagnetic simulation software has been ported to HPC machines with the support from an expert from the VSC Research Center at TU Wien and using the computing resources of the Vienna Scientific Cluster (VSC, https://vsc.ac.at). A careful performance analysis revealed several bottlenecks and limitations that were addressed by implementing optimised communication strategies with the Message Passing Interface (MPI). While the previous version of TAILSIT's software simulation tools relying on shared-memory parallelism was restricted to typically less than one million surface degrees of freedom, the new and optimised version employing distributed-memory parallelisation with MPI allows to treat problems up to $50{\cdot}10^9$ surface degrees of freedom. In addition, access to the HPC resources allowed to test and further optimise the shared-memory version of the fast multipole method FMM library. A detailed description of the project is given in [1].

This SHAPE project allowed for a giant leap from running simulations on desktop workstations to an HPC version of the software being able to not only treating much larger problem sizes than previously possible but also providing much faster time to solution. Since larger models become more and more relevant to TAILSIT's customers and being aware of their customer's interest in 'faster code', the new HPC version helps TAILSIT to increase the competitiveness of their software library and puts TAILSIT in a much better market position.

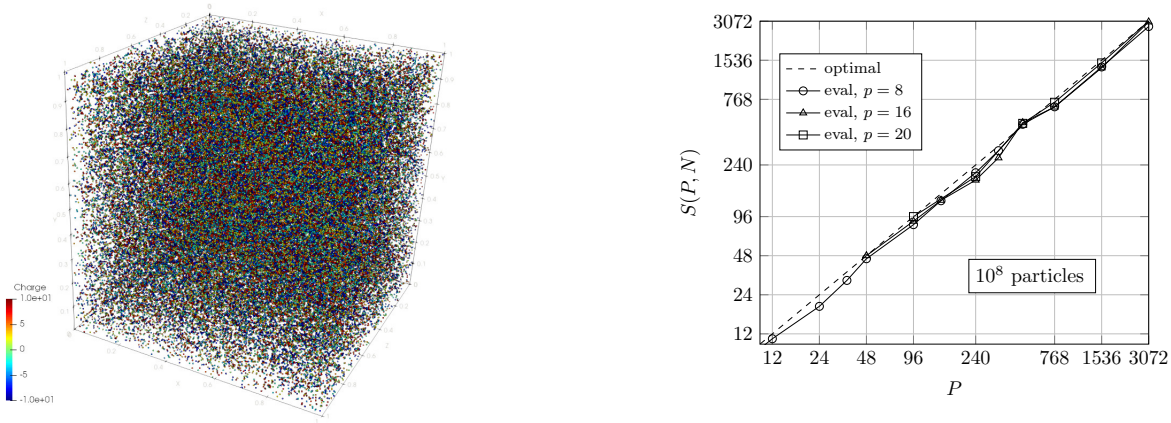

Fig. 1: Randomly charged point cloud in an unit cube (left). Parallel speedup of the evaluation of the potential with $N = 10^8$ particles for different FMM expansion degrees p (right).

The results shown in Fig. 1 are accumulated in a white paper [1] and will be online in the near future.

We acknowledge funding as a PRACE SHAPE project by the EU's Horizon 2020 research and innovation programme (2014-2020) under grant agreement 823767. The work is achieved using the resources of VSC-4 provided by the Vienna Scientific Cluster (VSC).

References

[1] Zechner J., Kielhorn L., Rüberg T. and Steiner S., to be published at
`https://prace-ri.eu/training-support/technical-documentation/white-papers/shape-white-papers/`

From silicon to silicon

Harald Höller-Lugmayr **and Martin Jurkovič**

IMS Nanofabrication GmbH, Brunn am Gebirge, Austria

In this work we present how an Austrian HPC application is essential for manufacturing the core building blocks of most high end commercial HPC hardware worldwide.

The IMS MBMW (Multi Beam Mask Writer) is a piece of equipment that went from proof of concept to one of the most demanded supplies for semiconductor fabs in less than ten years [1]. The already strong interest into this high end tool is further being fueled by an expansion course in capacities driven by the big players in the semiconductor industry: TSMC, Samsung and Intel Corp. [2].

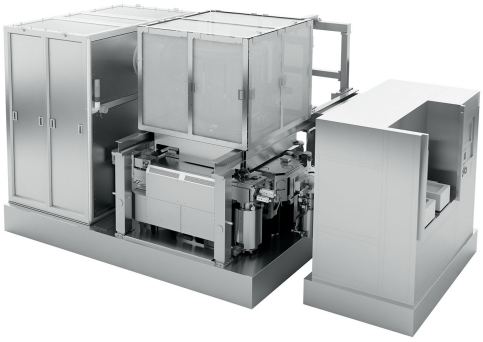

Fig. 1: Rendering of the IMS MBMW. This rendering does not show the HPC cluster.

The core technology of the MBMW, the in-house developed blanking chip [3], translates a constant stream of data provided by the HPC application to a modulated ray of a quarter million electron beams. This beam of electrons is printing a photomask used in a processing step of photolithography to building layers of semiconductor chips.

The data stream for the MBMW has highest demands on throughput, accuracy, reproducibility, resilience and robustness while all of this happens in a strictly IP protected environment. Already small defects on a modern photomask are a major issue for the customers. Pauses during writing – even if parts of the HPC clusters hardware were failing – are to be avoided and counter acted.

The HPC data pipeline consists of a heterogeneous system with 80 nodes of pure CPU computing and another 20 nodes which contain Nvidia GPUs. For a full mask up to 3×10^{15} pixels are generated on demand and printed in less than 10 hours while at no time this data is (or could be) stored on disks. InfiniBand connects the data source, i.e. the HPC cluster and the data sink, a custom FPGA based device.

In this talk, we will give an introduction to the technology while focusing on the peculiarities of the HPC application, that is an in house developed InfiniBand library based on the verbs API and GPU computing based on the CUDA driver API. The software is purely optimized for latency and data rate, while runtime, memory and energy footprint are neglected.

References

[1] Lapedus, Mark, Semiconductor Engineering, November 2018, *Multi-Beam Mask Writing Finally Comes Of Age*, `https://semiengineering.com/multi-beam-mask-writing-finally-comes-of-age/`

[2] Shamoun, Bassam and Chandramouli, Mahesh and Bin, Liu and Juday, Reid K., and Bucay, Igal and Sowers, Andrew T. and Abboud, Frank E., SPIE, **Novel Patterning Technologies 2021**, pages 64 – 82, *Multi-beam mask writer in EUV era: challenges and opportunities*

[3] Witt, Martin, Fraunhofer-Institut für Siliziumtechnologie ISIT, *TROM2 chip: A chip for mask production*, `https://www.isit.fraunhofer.de/en/micro-fabrication-technologies/Lighthouse-project/TROM.html`

Austrian DataLAB and Services

Peter Kandolf

VSC Research Center, TU Wien
Information Technology Services, University of Innsbruck

ADLS – Austrian DataLAB and Services project is an initiative funded by the Austrian government to foster collaborative approaches between our partner universities in the fields of data science and high performance computing. The explicit focus lies on increasing usability, lowering the learning curve and hence saving time for researchers, teachers and students in the course of using computing resources.

The project deals with the extension and integration of digital services (application programmes that enable the processing and analysis of different types of data) and the provision of templates for future digital services and platforms – building on existing infrastructure. Access to these services will be provided through graphical user interfaces and interactive platforms. In the background, these applications access federated resources for processing and storing data.

Inter-university access to research data and networking are not always easy, this project aims to change that. Cooperation between universities should be promoted and thus time can be saved. Different target groups should benefit equally from the project, therefore various areas of application have been identified, which will be considered in the implementation.

The services, including the infrastructure, support digitalization in teaching and research by simplifying the use of computing resources and algorithms. This makes them accessible to a broader group of researchers, teachers and students, especially for the humanities and social sciences.

In this talk we will present the current state of the project, give you a first view on the already existing services and the underlying architecture as well as an outlook on features to come.

Where is the sweet spot? A procurement story of general purpose compute nodes

***Alois Schlögl*, Andrei Hornoiu, Stefano Elefante, and Stephan Stadlbauer**

Institute of Science and Technology Austria (ISTA), Klosterneuburg, Austria

Introduction: Procurement of hardware is a critical part for running a HPC cluster. Questions about the number of nodes, selection of CPU's, how much RAM is needed, budget and cost constraints need to be decided. Here, we discuss a case study that shows a significant shift in the optimal design of a general purpose compute nodes.

Method: Some older computer nodes purchased in 2013 that were part of a cluster of about 200 nodes [1,2], had to be replaced. This concerned 16 nodes, each with 128GB RAM, Dual-socket CPU Intel Xeon E5-2630 (S2:C8:T2, 2.4Ghz), Infiniband (FDR), 1 TB HDD, each. That is a total of 2 TB RAM, and 512 CPU threads, and 8 GB/CPUcore.

In order to find the optimum solution, we did not ask for a specific number of nodes of a certain size, but set a budget limit and asked several vendors to provide a solution with the best performance to cost ratio. The minimum requirements were 8 GB/CPUcore, amd64 with AVX2 instruction set, EDR Infiniband. As a metric for performance we used the product of "number of CPU cores" times "clock rate". Because this was a small tender, infrastructure costs like energy consumption, switches, etc. were not considered in the evaluation.

Results: After receiving offers from four different vendors, the best offer contained five nodes, with 1TB RAM, Dual-Socket CPU AMD EPYC 7662 64-Core (S2:C64:T2, 2GHz), HDR Infiniband, 2TB HDD per node; in total these are 5 TB of RAM, and 1280 CPU threads. The costs were slightly smaller than for the 2013 systems. The big surprise is that the number of nodes reduced from 16 to 5, and the nodes looked more like nodes we previously considered "big-memory machines".

The total amount of memory as well as the number of CPU cores increased by 2.5x, less network ports and hard discs were needed, the power consumption per node was about the same, which means total power consumption was reduced (1/3 x), and less overhead from the operating system is needed, and costed 45 Euro/(Core x GHz).

Conclusion: These nodes are in production since Oct 2021 and users are quite satisfied with that hardware solution. Several MPI workloads have now been moved to these machine and jobs that previously required 8 nodes, now they have the possibility to run on a single node.

This is a significant shift in the optimal architecture for the general purpose compute node. Many workloads that were running in distributed mode, can now run on a single node. The advantages are obvious, (i) addition performance gains because latency caused by network transfer can be avoided, (ii) simpler software development, and (iii) less infrastructure costs (e.g. number of switch ports needed, number of power supply units, number of hard disks). The distinction between standard compute nodes, and big memory machines has also become less significant, or in other words the definition of specialized big-memory machines will shift to several Terabyte of RAM per node.

References

[1] Schlögl, A., and Kiss, J. Scientific Computing at IST Austria. p.28, Booklet AHPC (2017)

[2] Schlögl, A., Kiss, J., and Elefante, S. Is Debian suitable for running an HPC Cluster? p.25, Booklet AHPC (2019)

Vega EuroHPC, 1st year in service

Andrej Filipčič[a] **and *Dejan Valh***[b]

[a] *Jožef Stefan Institute, Jamova 39, 1000 Ljubljana, Slovenia*
[b] *IZUM, Prešernova Ulica 17, 2000 Maribor, Slovenia*

Vega (figure 1), the first of eight new EuroHPC supercomputers, started operation in March, 2021. With 6.9 PFlops of aggregated performance on CPU and GPU partitions it provides a significant computing power to Slovenian scientists and international projects where the researchers from Slovenia participate. One third of the system is dedicated to European users applying for resources through the development and benchmarking, and regular EuroHPC allocation calls. Atos Bullsequana XH200 architecture with 2040 AMD-7H12 CPUs and 240 NVidia-A100 GPUs provides one of the best solutions for extreme performance and power efficiency with PUE 1.05. The stable operation was reached in less than one month after commissioning with availability of 98% within the last year. More than 150 Slovenian scientists have used Vega till now for their research, while the EuroHPC Joint Undertaking share was allocated to almost 50 projects.

Apart from CPU or GPU intensive workloads, Vega was also designed to provide efficient data processing, storage and data transfer capabilities. The fast storage is provided by DDN Lustre array of NVMe disks, while the underlying large capacity storage is managed by Ceph. The latter provides RBD images to virtual machines, the shared CephFS filesystem to projects requiring PBytes of storage, and ObjectStore accessible through S3 interface. The dCache distributed storage (in deployment), will aggregate the storage centers at IZUM(Vega), ARNES and Jožef Stefan Institute to a single point of access and management, suitable for long term storage including tape archive capabilities. Large scientific projects such as ATLAS at CERN, Belle-2 in Japan, Cherenkov Telescope Array, Vera Rubin Observatory, use integrated computing and storage resources of hundreds of computing centers worldwide. As such, Vega provides Computing Element (ARC-CE) and storage endpoints (dCache) for transparent integration of resources in the project workflow and data management systems.

The Proxmox virtualization partition provides resources to deploy virtual machines and containers in high-availability mode with live migration of instances using dedicated Ceph storage. Services such as data transfer agents, remote job submission endpoints, S3 endpoints, monitoring infrastructure, squid proxies, kubernetes control planes are deployed in a horizontally scalable mode and are able to fully exploit the available WAN bandwidth of 200Gb/s. In addition, user-space virtual machines can also be provided on demand.

Fig. 1: Vega EuroHPC.

One of the main goals of Vega is to provide all the infrastructure the user communities might need for their complex workflow organization and data management (HPC Cloud), including the infrastructure for open science in Slovenia.

References

[1] Vega description: `https://www.izum.si/en/hpc-en/`

[2] Vega user documentation: `https://doc.vega.izum.si/`

Next generation Vienna Scientific Cluster (VSC-5)

Siegfried Höfinger[a] **and** ***Jan Zabloudil***[a,b]

[a] *VSC Research Center, TU Wien*
[b] *BOKU Wien*

Following a competitive procurement process for a new cluster system, a contract has been awarded to the German company Megware. The installation of the system was carried out in Q1-Q2/2022 and took longer than anticipated because of the well-known shortage of electronic components worldwide.

The new cluster system, VSC-5, features a hardware partition with more than 700 CPU nodes and one with 60 dual GPU nodes connected via HDR infiniband. As processing units 2x AMD Epyc CPUs of Milan Architecture with 64 cores are deployed on all nodes, while each GPU node is additionally equipped with two NVIDIA A100 (40GB) cards. For energy efficiency the CPU nodes use a direct hot water cooling system developed by Megware which can be operated at temperatures up to 50 degree Celsius. The existing Spectrum Scale installation will continue to serve as the central storage system.

In the presentation we will discuss details of the delivered hardware and present results of benchmarking the system. From a user perspective there will be changes in the deployed OS, some adaptions of SLURM and enhancements in the application software environment to take care of the differences in processor directives between AMD and Intel.

Fig. 1: Snapshots of compute nodes of VSC-5.

VSC Storage – Present and Future

Markus Hickel

VSC Research Center, TU Wien

VSC offers several different storage targets to satisfy the storage needs of VSC users. Due to increasing issues with aging hardware and software compatibility, the BeeGFS parallel file systems will be retired as soon as VSC-3+ reaches end of life. The storage infrastructure acquired with VSC-4 will serve as replacement for the BeeGFS file system and is going to be used for our new flagship VSC-5.

IBM Spectrum Scale was chosen to serve as parallel file system. The hardware is comprised of 6 NSD-Servers, 6 HDD storage arrays and two IBM flash systems 7200. Each NSD-Server is equipped with 256 GB of memory, 1 Intel Omni-Path and 1 Inifiniband HDR host channel adapter. The additional Infiniband HDR HCA is necessary to provide connectivity to the new VSC-5 System. Each HDD storage array is directly attached to two NSD-Servers via 12 Gb SAS and contains 120 disks with 12 TB each providing a total of 6 PB usable storage capacity. Each IBM flash system is connected to 3 NSD-Servers via 6 32GB Fibre Channel links and provides roughly 1 PB of pure flash storage. This is a significant upgrade to the previously used IBM flash system 9100, which provided 150 TB of usable storage capacity.

This talk will focus on our experience gained with Spectrum Scale over the last 2 years of operation and the implementation of the new flash storage. We decided to add the flash storage as a separate storage pool to the present data file system to ensure seamless integration without any user intervention or manual data migration. This is achieved by using Spectrum Scale file placement and file management policies. These policies allow to place newly created files in a specific storage pool and move them to different storage pools throughout their lifecycle.

Fig. 1: Spectrum Scale storage infrastructure

Current Developments on the Vienna Scientific Cluster (VSC)

Josef Beiglböck[a], Sebastian Frank[b], *Simeon Harrison*[b], Markus Hickel[a], Siegfried Höfinger[a], Katrin Muck[a], and *Irene Reichl*[a,b]

[a] *VSC Research Center, TU Wien*
[b] *EuroCC Austria, VSC Research Center, TU Wien*

JupyterHub on VSC

The JupyterHub running on VSC-3 and VSC-4 has gone from the beta testing phase to being fully operational and accessible to the public. Users can now choose between different pre-installed singularity images, their own custom built singularity image or a virtual environment. VSC users are now able to install their packages, deploy their notebooks, train machine learning models, submit jobs and visualise their data without the hassle of the Linux command line. We will demonstrate the ease of the current state by training and testing a deep learning model for road sign recognition.

Interactive Nodes in the VSC Environment

Interactive node access with powerful 3D rendering capacities is crucial to engineering applications in HPC environments. Moving data between cluster and local machine for pre and post processing is not only unnecessarily time-consuming, it may even be unfeasible for large amounts of data. The rack work stations instead set up as interactive login nodes enable that a case can be set up, submitted and post processed directly on the cluster. When the job starts running, the engineer can check convergence and stability of the case and modify parameters if necessary. We will demonstrate how these nodes are integrated in the VSC environment and how to use them.

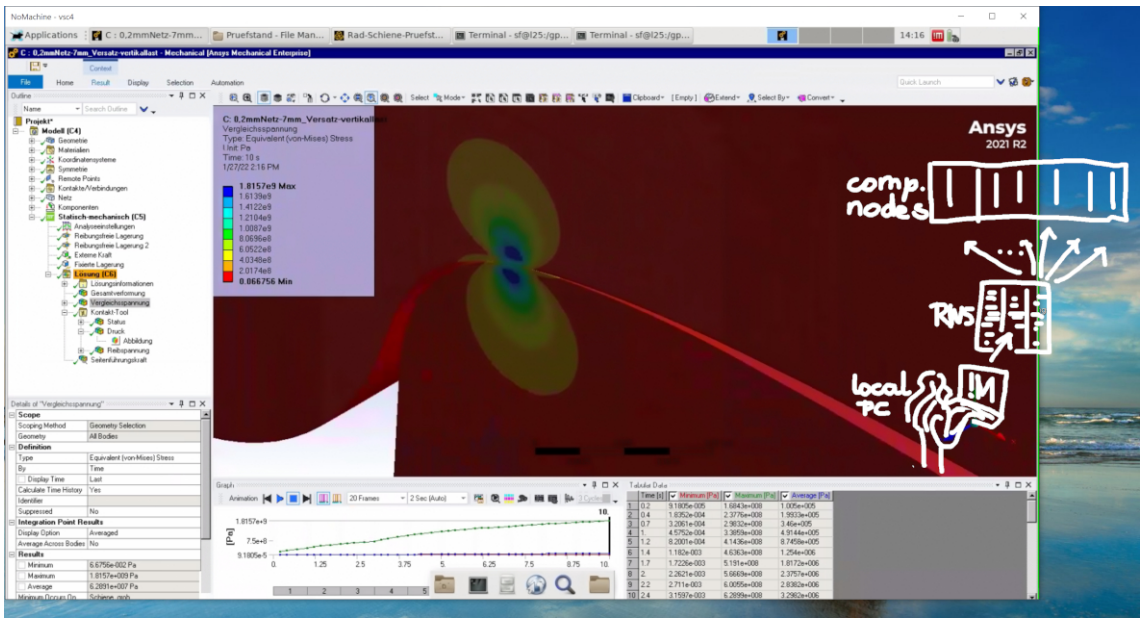

Fig. 1: Running ANSYS on the interactive nodes through noMachine.

Keynote Talk:

Selected Cases of GPU Accelerated Parallel Applications at IT4Innovations

Lubomir Riha

IT4Innovations, VSB - Technical University of Ostrava

In this talk I would like to present selected GPU accelerated applications we have developed or ported at IT4Innovations National Supercomputing Centre in Czech Republic.

One of the applications that we have successfully ported to multi-GPU nodes with fast NVLink interconnect is the GPU accelerated Blender Cycles renderer. Our proposed memory management method enables to distribute massive scenes among memories of multiple GPUs. By duplication of carefully selected very small amount of data, between 1 and 5%, we are able to achieve performance very close to the case where entire scene data are fully duplicated [1].

Another application entirely developed by one of our teams is the ESPRESO FEM [2] which includes the massively parallel FETI based linear solver. We have developed a method that converts sparse data structures of the finite element-based method into dense ones, without significant increase in its size. This enables us to successfully port the FETI solver to GPUs. This work is part of the EuroHPC EUPEX project, which aims to design, build, and validate the first EU platform for HPC.

I will also present a work of our team that works on Energy Efficiency of both CPU only as well as GPU accelerated nodes. Our method is able to save energy by tuning the hardware to the specific needs of the current workload [3].

Finally, IT4I is also part of the H2020 POP2 CoE, and I would like to present our services, which include analysis of parallel as well as GPU accelerated applications, identification of parallel bottlenecks, and providing help with code optimization.

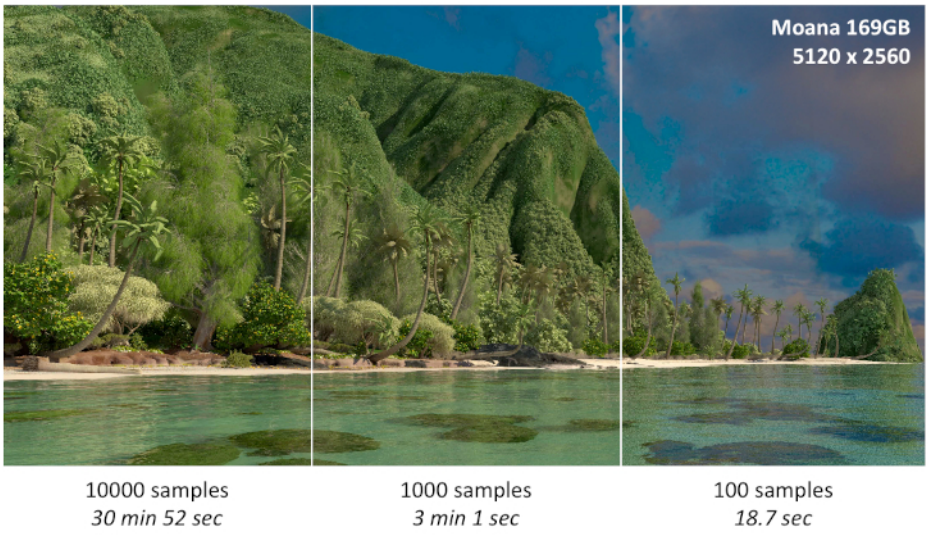

Fig. 1: Examples of rendering of massive Moana Island Scene dataset using modified Blender Cycles path tracer. The rendering time is measured on an NVIDIA DGX-2 machine with 16 GPUs.

References

[1] Jaros, M., Riha, L., Strakos. P., and Spetko, M. GPU Accelerated Path Tracing of Massive Scenes. ACM Trans. Graph. 40, 2, Article 16 (April 2021)

[2] Riha, L., Merta, M., Vavrik, R., et al. A massively parallel and memory-efficient FEM toolbox with a hybrid total FETI solver with accelerator support. The International Journal of High Performance Computing Applications. 2019;33(4):660-677

[3] Spetko, M.; Vysocky, O.; Jansik, B.; Riha, L. DGX-A100 Face to Face DGX-2—Performance, Power and Thermal Behavior Evaluation. Energies 2021, 14, 376

Energy-Efficient Computing with Approximate Tensor Core Units

Ratko Pilipović, Patricio Bulić, and *Uroš Lotrič*

University of Ljubljana, Faculty of Computer and Information Science

High-performance computing systems must simultaneously address performance, energy efficiency, and cost constraints. Approximate computing is a new emerging paradigm that could help tackle the constraints and lead to energy-efficient computing by inducing acceptable computation errors. We can apply approximate computing on many levels of an application: it could be at the software level to reduce the calculation when sufficient precision is reached, or even at the hardware level, where we replace an exact circuitry with the approximate one.

Many error-resilient applications tolerate noise and measurement errors in data and could also handle errors in computation, deep neural networks being among the most well-known. The essential operation in deep neural networks is the general matrix multiplication (GEMM) comprised of matrix multiplication and accumulation. The importance of its fast execution has led graphics accelerator manufacturers to integrate tensor core units into their products.

We introduce an approximate tensor core, which employs approximate multipliers to perform the GEMM operation on four-by-four matrices given in sixteen-bit signed fixed-point format. We implemented the approximate tensor core unit with several state-of-the-art approximate multipliers: logarithmic, radix-based and hybrid (Fig. 1). The synthesis of the approximate tensor core unit to the 45 nm Nangate Open Cell Library revealed that it consumes only up to 36% of the area and 25%of the energy required by the exact tensor core unit and at the same time speeds up the computation by reducing the propagation delay.

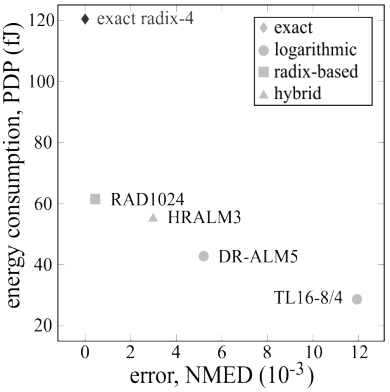

Fig. 1: Characteristics of the fixed-point multipliers used in the approximate tensor cores.

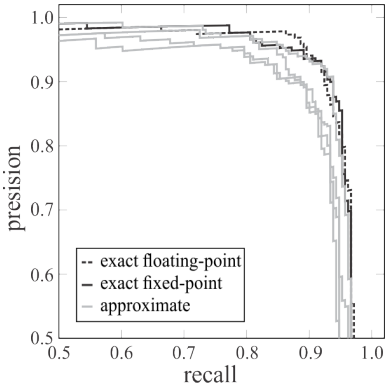

Fig. 2: Precision-recall curves of the approximate tensor cores on the honeybee dataset.

We evaluated the tensor core unit with the YOLOv4-tiny convolutional neural network object detector implemented in Keras/Tensorflow framework, where we simulated the approximate tensor core unit by replacing the exact floating-point multiplier with the approximate fixed-point multipliers. We tested detectors' performance on the honeybee detection dataset, where the detectors employing the approximate fixed-point tensor cores were only slightly less accurate than those using the floating-point tensor cores (Fig. 2).

The case study shows that the approximate tensor core unit has potential in error-resilient applications where it can use significantly less silicon area and energy for the task. Artificial intelligence with deep neural networks has become more important in everyday life and sciences, so high-performance systems empowered with approximate accelerators might appear soon, maybe through the European processor initiative.

References

[1] Pilipović, R., Risojević, V., Božič, J., Bulić, P., and Lotrič, U., An approximate GEMM unit for energy-efficient object detection, Sensors, **21**(12), 1-19, `https://doi.org/10.3390/s21124195` (2021)

Shared-Memory Fast Marching Method for Re-Distancing on Hierarchical Meshes

Michael Quell[a], **Andreas Hössinger**[b], **and *Josef Weinbub***[a]

[a] *Christian Doppler Laboratory for High Performance TCAD, Institute for Microelectronics, TU Wien, Austria*
[b] *Silvaco Europe Ltd., UK*
josef.weinbub@tuwien.ac.at

The fast marching method (FMM) is used in level-set based topography simulations to re-establish the signed-distance property of the level-set function, i.e., calculating the signed-distance to the zero-level-set. The signed-distance property is essential for numerically robust topography simulations. The typical geometries encountered when simulating feature-scale manufacturing processes of microelectronic devices consist of large flat regions and sharp corners which require a high spatial resolution. Considering that level-set solvers favor structured meshes because of efficient finite difference schemes, hierarchical meshes are inevitable to efficiently resolve small features in large simulation domains.

Hierarchical meshes naturally enable parallelism by allowing for independently executing the FMM on each mesh, followed by a synchronized exchange which then requires a partial restart of the FMM. However, the parallel speedup of this approach is often limited by the number and size of the given hierarchical meshes, especially for a high number of threads because of load imbalances. This bottleneck is overcome by a recent developed block decomposition approach, which allows for a higher parallel speedup [1]. The decomposition affects only the re-distancing, because the hierarchical meshes are already tailored to other computational tasks in the simulation flow, e.g., advection. To gain the best performance the block size and the stride width, i.e., the frequency of synchronization steps between meshes have to be tuned. Currently, the optimal decomposition size has to be determined on a case by case basis. To that end, we present an extensive parameter study, providing insights into the parameter selection.

Figure 1 shows the speedup of the re-distancing for the *Mandrel* geometry, measured on a single node of VSC-4. The *Mandrel* geometry consists of two trenches in an otherwise flat rectangular domain. By tweaking the stride width a serial speedup of up to 1.2 is possible. In case no decomposition is employed the peak speedup is 7.6 for 16 threads, however, using decomposition with a block size of 50 and stride width of 5.0 a speedup of 15.4 is achieved for 16 threads. In the decomposition case the speedup is further increased to 17.5 for 24 threads.

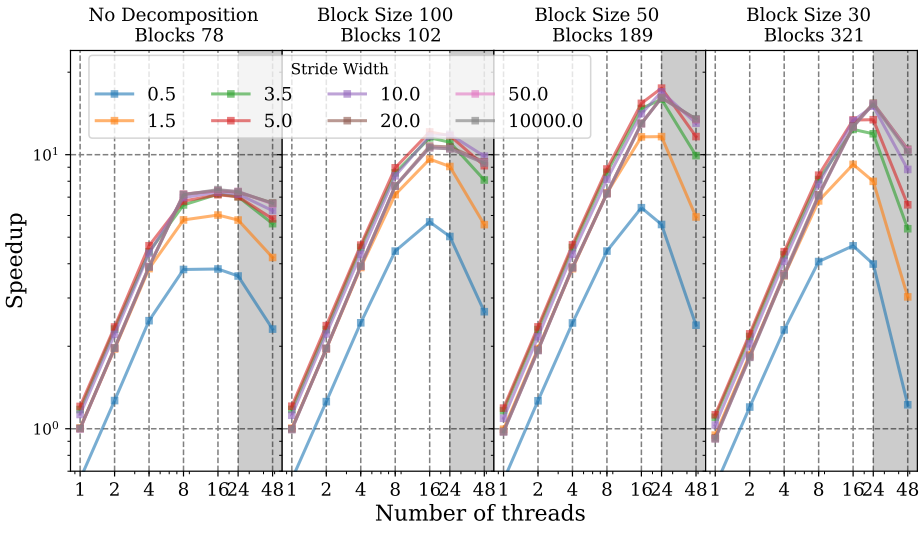

Fig. 1: Speedup of parallel re-distancing in dependence of block size and stride width. Reprinted with permission from [1]. Copyright 2021 Authors, licensed under the CCA 4.0 International License.

Acknowledgment. The financial support by the Austrian Federal Ministry for Digital and Economic Affairs, the National Foundation for Research, Technology and Development and the Christian Doppler Research Association is gratefully acknowledged. The computational results presented have been achieved using the Vienna Scientific Cluster (VSC).

References

[1] Quell, M., Diamantopoulos, G., Hössinger, A., and Weinbub, J., J. Comput. Appl. Math. **392**, 113488 (2021)

`mpisee`: MPI Profiling for Communication and Communicator Structure

***Ioannis Vardas*, Sascha Hunold, Jordy I. Ajanohoun, and Jesper Larsson Träff**

TU Wien, Faculty of Informatics, Austria

MPI Application Profiling: Profiling MPI applications is a fast and lightweight method for obtaining information about their characteristics, such as the type of communication and time. Another method for obtaining more detailed information such as timestamps of events is tracing. However, tracing introduces more overhead both in time, memory and storage compared to profiling. Common profiling tools profile on a per-process basis disregarding the different communicators of the MPI applications. The communicator is an essential abstraction of MPI which defines a communication domain were all types of communication take place, both point to point and collective. Furthermore, a communicator, can carry information about the virtual process topology which represents the communication pattern of the application. We present a novel MPI profiler (`mpisee`) that separates and records all type of communication per communicator in the application. We refer to this method of profiling as communicator-centric. Our aim is to gain insight regarding the communication within different communicators that can carry various virtual process topologies, to enable improved solutions to the communicator reordering problem of MPI.

MPI Communicator Concept: The MPI standard enables the creation different communication domains by partitioning the MPI processes into subsets (Fig. 1). Each of these subsets is called communicator, where both collective and point-to-point communication take place. The communicator that contains all the processes is called `World` and every other communicator is a sub-communicator of `World`.

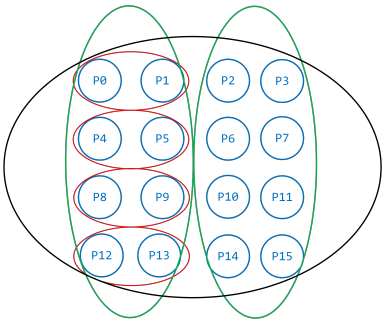

Fig. 1: Partioning of 16 MPI processes (blue circles) into six different communicators.

Communicator-centric Profiling: We developed a novel lightweight MPI profiling tool, called `mpisee`, which is able to obtain information on the data volume, the time spent, and the frequency of MPI calls per communicator. The tool relies on the PMPI, which is the MPI standard profiling interface. With our unique communicator naming scheme the tool is able to distinguish between the different communicators and supports the 35 most commonly used communication calls [1]

Use Case: `mpisee` reveals the time spent in every MPI primitive for each communicator. Fig. 2 shows a sample profile for the SPLATT application, where the majority of time was spent in the `World` communicator by Allreduce, Bcast, Barrier and Recv operations. The second most used communicator is the Cartesian communicator `W_a1`(256), the number in the parenthesis is the number of processes in the communicator. Finally, only Alltoallv is used in the rest of the communicators. Such information cannot be obtained from other MPI profilers such as Score-P or mpiP [2] and tracing tools.

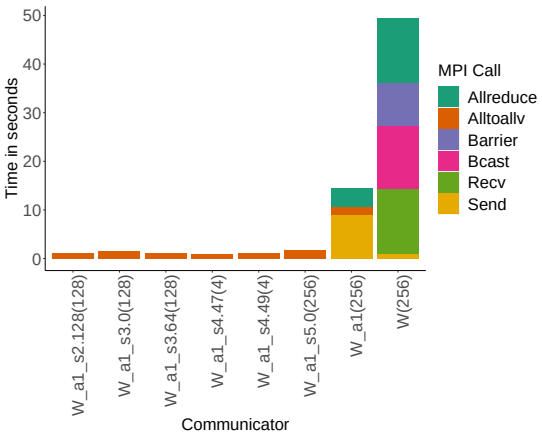

Fig. 2: Time spent (y-axis) in each communicator (x-axis) of MPI calls for SPLATT.

References

[1] Chunduri, S., Parker, S., Balaji, P., Harms, K. and Kumaran, K., SC'18: Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis, Article No.: **30**, (2018)

[2] Padua, D., Encyclopedia of Parallel Computing, 1637 (2015)

MPI Performance Tools under the Microscope: A Thorough Overhead Analysis

Jordy I. Ajanohoun, Ioannis Vardas, Jesper-Larsson Träff, and *Sascha Hunold*

TU Wien, Faculty of Informatics, Austria

MPI Performance Analysis: Profiling and tracing tools are important for pinpointing performance issues of MPI applications at large scale. Profilers produce an aggregated view of the time spent in different parts of an MPI application, whereas tracing tools provide an in-depth view of individual events. MPI profiles are relatively small in size compared to traces, and obtaining profiles typically entails a smaller overhead than recording a full application trace.

Goal: In our work, we set out to precisely **examine the overheads** associated with different **performance analysis tools**. We classify the tools into (1) hardware performance counters (HPCs), (2) profiling, and (3) tracing tools.

The Overheads of MPI Performance Tools: We identify three types of overheads of typical MPI performance tools, which are depicted in Fig. 1. Overhead 1 may occur when performance analysis tools intercept the `MPI_Init()` and the `MPI_Finalize()` functions to perform pre- and post-processing. Overhead 2 refers to the overhead when performance tools slow down the MPI application itself. This may happen when tools intercept individual functions and record timing statistics before and after these functions are executed. We denote the wrapping overhead as the third type of overhead. This overhead only occurs for performance tools that wrap the MPI application to be analyzed. These tools are typically started as the main program and they will start the MPI application themselves (e.g., hardware performance counter tools, such as Likwid).

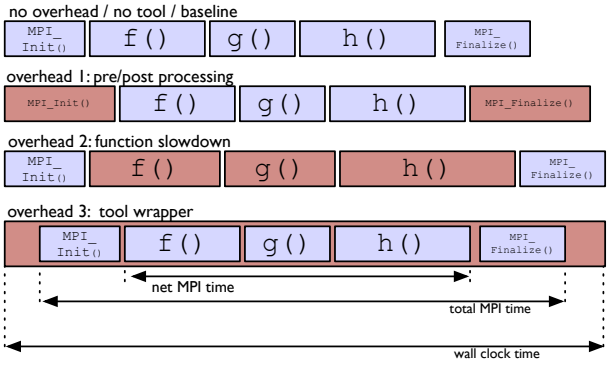

Fig. 1: The three types of overheads investigated.

Comparison of MPI Performance Tools: We have performed a thorough experimental examination of the most commonly used tools for MPI performance analysis. Fig. 2 shows the overheads introduced by different performance analysis tools when analyzing the AMG proxy application [1] with `8x32` processes and MPICH 3.4.2. The *baseline* denotes the runtime of AMG without the utilization of any performance tool. We can observe that several tracing tools cause a significant overhead, which users need to be aware of.

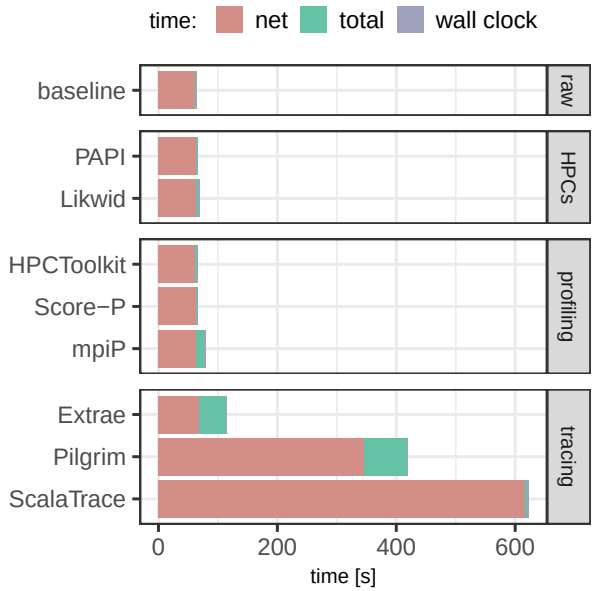

Fig. 2: Runtime decomposition for AMG.

References

[1] Gahvari H., Baker A. H., Schulz M., Meier Yang U., Jordan K. E., and Gropp W., ICS 11 Proceedings of the 2011 ACM International Conference on Supercomputing, 172, (2011)

EuroCC Austria – A National Competence Centre for High-Performance Computing, High-Performance Data Analytics and Artificial Intelligence

Sarah Stryeck[a,d] **and Markus Stöhr**[b,c,d]

[a] *Graz University of Technology, Austria*
[b] *University of Natural Resources and Life Sciences, Vienna, Austria*
[c] *VSC Research Center, TU Wien, Austria*
[d] *EuroCC Austria*

EuroCC (European Competence Centres) [1] is a European project funded by the European High Performance Computing Joint Undertaking (EuroHPC JU) to support research and innovation in High-Performance Computing (HPC), High-Performance Data Analytics (HPDA) and Artificial Intelligence (AI) in 33 European countries. Within the frame of the project, the National Competence Centre in Austria – EuroCC Austria [2] – was founded in September 2020 with national co-funding from the Austrian Federal Ministry of Education, Science and Research (BMBWF).

The mission of EuroCC is to foster HPC/HPDA/AI in general but also to strengthen existing synergies by collaborating with academia, businesses, industry and the public sector. The focus of EuroCC Austria is to improve the services for our already existing user base from academia. An additional goal for EuroCC Austria is to reach out to start-ups, SMEs and big industry as a new user base. The goal is to help them to identify their yet unknown needs and to develop dedicated services for industry. One major activity in business development is to raise awareness about possibilities offered by HPC and, in a next step, to offer support to initiate HPC related projects. That can range from giving advice on how to apply for government subsidies to actually running simulations on a supercomputer. A central goal is to map these industrial needs to the already available competences in academia. The role of EuroCC Austria is to do a matchmaking - this means that, if a certain need in industry is identified, we aim to connect this stakeholder to experts from academia or at a later stage from specialized SMEs/companies with the respective competence. In addition, we also identified certain gaps where we lack competences in Austria. In these cases, we use our international network through the EuroCC project to bring those competences or services to Austria, or we are working on new training programs to improve and strengthen the required skills in Austria and beyond. Within the project, we were already able to deliver various trainings for HPC, to organize and contribute to events in these areas and to get an overview about the Austrian landscape in HPC, HPDA and AI. With this presentation we would like to present to you EuroCC Austria, show what we have achieved so far, our currently available services as well as the future outlook for the competence centre.

References

[1] `https://www.eurocc-access.eu` (31.01.2022)

[2] `https://www.eurocc-austria.at` (31.01.2022)

Slovenian National Competence Centre for HPC

Pavel Tomšič

University of Ljubljana, Faculty for Mechanical Engineering, Aškerčeva 6, 1000 Ljubljana, Slovenia

Introduction

National Competence Centre HPC (NCC) is an H2020 project (EuroCC) and an activity of the Slovenian national supercomputing network SLING [1]. The mission of the NCC is to promote the use of supercomputer capabilities for research in science, industry, and academia, especially through raising the level of user knowledge and general awareness of the benefits of the use of such technology.

NCC provides users with support in the usage of HPC and HDPA technology for different levels of knowledge and areas of application, the possibility of acquiring knowledge at various training and seminars, and assistance in the use of technology for medium and small enterprises.

Providing customer support

For existing and future users of supercomputer capabilities, a multi-level system of user support has been established. One of our goals is to include as much expertise from Slovenian research communities as possible in the operation of NCC and thus ensure a high level of support in various research areas and a unified and comparable user experience on all freely available national supercomputer infrastructure and national EuroHPC systems.

In 2021 we have successfully attracted more than 900 new users to our HPC infrastructure. In the framework of High-level support, we have established an application portal that includes detailed documentation for configuring and running applications on HPC.

Organization of training, seminars and conferences

We have successfully established a training and skill development program, from short (half-day) to multi-day training courses to organising short seminars covering specific topics presented by academic researchers. Our training activities are divided into three levels, according to the knowledge transfer to be obtained: (i) the basic level focuses on HPC literacy and general HPC-oriented themes; (ii) the advanced level covers different research fields in science, technology, engineering, and mathematics, while (iii) the higher level tries to raise competences among regular HPC users.

During the year 2021, we have successfully performed a total of 34 activities, 19 training courses and 15 seminars. The distribution of the courses was 10 trainings on the basic level, 5 trainings on the advanced level and 4 trainings on the higher level. Altogether we have educated 444 individuals, and the average attendance was just shy of 27. 287 participants (64,6%) have joined only one event, and 157 (35,4%) have joined more than one event. 15% of participants have joined more than two events, with several attending more than five events. This indicates the appropriate orientation of the program, as it clearly shows that we can provide motivation for upgrading knowledge through various events.

We have also organized *Exascale Computing Day* or *Open Day of Slovenian Supercomputing Centres*, where we invited members of the general public to visit our biggest HPC centres. The event was quite successful, having around 600 visitors, mostly from Slovenian high schools. Participants were able to visit supercomputer centres, and the event was also streamed on-line.

References

[1] SLING: the Slovenian NCC website, `https://www.sling.si/sling/en/`, accessed in April 2022

EuroCC project has received funding from the European High-Performance Computing Joint Undertaking (JU). The JU receives support from the European Union's Horizon 2020 research and innovation program and 33 European countries.

Vision of a HPC NCC today

Ferenc Szani

Governmental Agency for IT Development (KIFÜ) - Hungarian NCC

Rethinking the development of competences related to use of HPC in educational management, planning, control and learning processes is inevitable due to the evolution of educational technology, both at the Hungarian and the international level.

The framework: OER

Open Educational Resources (OER) is based on the 2012 UNESCO Paris Declaration. The amount of OER elements in education systems should increase step by step. The technologies and capabilities called HPC today, represent the standard IT level of the future, so the competences needed to use them must be provided openly.

The vision

National Competence Centers (NCC) and related university relationships are not options in which a university or an NCC may or may not participate, but it is self-evident. Geographical borders do not impose a barrier for education, therefore costs can be reduced and cooperation can be expected to take place without central funding. This is essential for NCCs, to become self-sustaining HPC knowledge centers. To do this, of course, one also needs to develop an appropriate service portfolio. One of the priorities in the service portfolio is to relocate HPC education and competency development systems beyond the Gutenberg galaxy.

The following steps were and will be needed next year:

- HPC NCC should operate as a separate economic, and professional unit.
- Develop copyright, legal and financial environment. As much content as possible, including existing content, should be in a creative common license that can be used for further content development, even internationally.
- Provide ongoing “Train the Trainer” programs for instructors.
- We will convert our most successful courses into English (M)OOC courses with good lecturers and upload them to an international course provider.
- We will create collaboration hub for industry players such as SMEs and develop content and other educational resources that can be used immediately by them.
- Fresh, authentic content sources will be developed and their virtual environment under the brand name of our institution.
- A single certified/accredited HPC examination center will be established, which will be entitled to check and map the competences of HPC CC staff and users. Of course, this will require an internationally accepted HPC competence framework.

In this way, the Hungarian NCC will become a modern scientific hub, adapting to needs of the modern age and evolving from a reactive, adaptive environment to a managing environment.

Keynote Talk:

HPC-Integration of an Ion-Trap Quantum Computer

Albert Frisch[a] **and Thomas Monz**[a,b]

[a] *Alpine Quantum Technologies, 6020 Innsbruck, Austria*
[b] *Institut für Experimentalphysik, Universität Innsbruck, 6020 Innsbruck, Austria*

In recent years, we have seen first developments of quantum computing (QC) systems towards powerful solvers for a variety of real-world applications. Traditionally, many of these same calculations are solved using conventional high-performance computing (HPC) systems, which have progressed steadily through decades of hardware and software improvements. They define the pinnacle of modern computing by drawing on massively parallel processing and the usage of specialized accelerators. For near-term integration of QC devices with HPC nodes, which will have quantum resources available within a novel type of accelerator, several considerations and requirements have been outlined for use cases, macro-architecture, micro-architecture and programming models [1].

At AQT we are working on bringing our QC systems based on trapped ions closer to existing HPC infrastructures by integrating required hardware devices, like ion trap devices, laser systems, optical references and control electronics, into standard data center racks.

A strong focus lies on the generation of analog carrier signals for qubit manipulation with highest fidelities, but also on the orchestration of individual devices to ensure ideal qubit environmental conditions. Mastering both in-band as well as out-of-band system control allows us to achieve highest system performance and optimal usage of quantum resources available. The development of appropriate connectivities, interfaces and programming models integrates the QC node into existing HPC environments and makes a quantum processing unit (QPU) available to the user as an accelerator for solving specialized computational tasks.

In this talk, I will first give a general introduction into quantum computing basics and give an overview of current ideas and our long-term vision of applications for hybrid HPC-QC systems. Then I will present our current steps towards HPC-QC integration. Finally I will briefly highlight some technical challenges and requirements and show the status and performance of our QC systems at AQT.

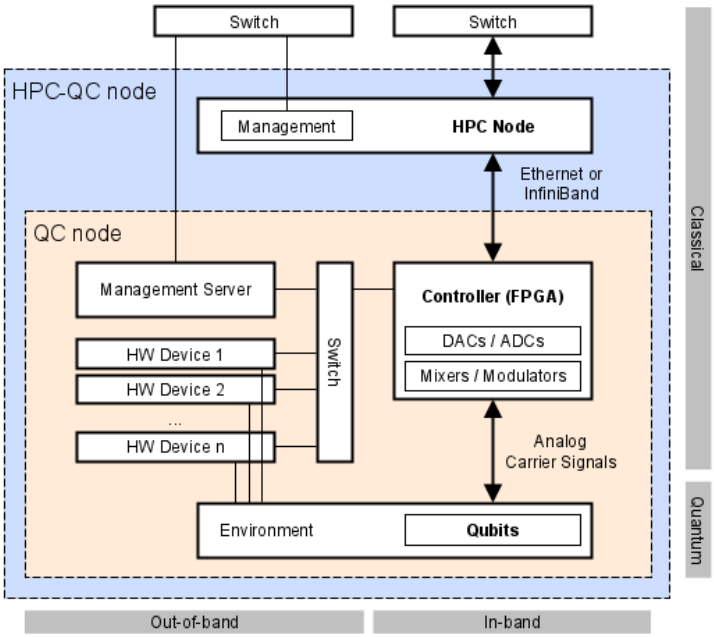

Fig. 1: This diagram shows the major components needed for the operation of a QPU within the HPC node infrastructure, from [1]. Individual components are grouped into so-called out-of-band and in-band scopes and are placed on the left-hand and right-hand side of the figure, respectively. The QPU, which contains the qubits and is capable of processing quantum information, is depicted at the lower part, whereas classical information processing components are shown in the upper part of the figure.

References

[1] T. S. Humble, A. McCaskey, D. I. Lyakh, M. Gowrishankar, A. Frisch and T. Monz, "Quantum Computers for High-Performance Computing", IEEE Micro, 41, 5, pp. 15-23 (2021)

Romeo and *Julia*: HPC, ranking, and the many-electron problem

Markus Wallerberger **and Karsten Held**

Institute of Solid State Physics, TU Wien, 1040 Vienna, Austria

Python and numpy are a bedrock upon which a vast scientific computing ecosystem is built. For this reason, scientific Python packages are widely deployed in supercomputers such as the VSC. In our computational material science group, we maintain an open-source Python/Fortran package for solving models of correlated electrons [1] as well as a novel Python/C package for the efficient solution of two-particle diagrammatic equations [2], among others.

Unfortunately, Python+X packages are supremely difficult to deploy on supercomputers: any build compliciations are compounded by the *DLL hell* of multiple versions of Python interpreter, numpy library, BLAS, other libraries, and MPI infrastructure, which all need to be made and kept consistent. Moreover, the vectorization needed to make numpy perform acceptably as well as the presence of external libraries provide a high barrier of entry to students and slow down prototyping efforts.

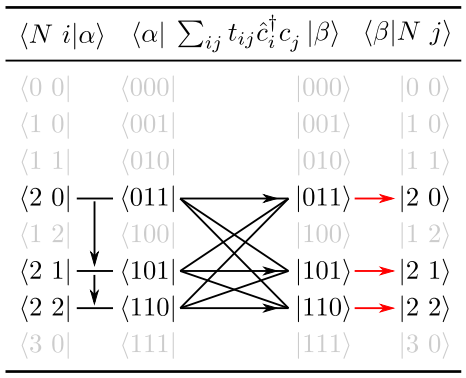

Fig. 1: Matrix–vector multiplication with an "implicit" form of a Hamiltonian for a quantum many-electron model with three sites.

The Julia programming language serves as an intriguing alternative for high-performance computing: firstly, it has a Python-like syntax, thereby easing any transition and entry. Secondly, it supports fast compiled loops, thus largely eliminating the need for supplementary native libraries. Thirdly, it vendors its own set of base libraries and maintains a content-addressable package repository, easing deployment considerably.

In this talk, we largely corroborate these advantages: specifically, we report on algorithmic improvements and an associated package for exact ground state calculations (ED/FullCI) for small many-electron systems [3]. The computational core of these calculations is the multiplication of the Hamiltonian H, a square matrix, with a quantum many-body state ψ, a vector of dimension $\sim 10^9$. H cannot be stored in any way, and is instead constructed on-the-fly in each multiplication with ψ (cf. Figure 1). Acceptable performance depends crucially on a fast map from a set of physical states (set of bit patterns) to a set of array indices, known as a "ranking".

Traditionally, ranking is done using a bisectioning search, which has poor cache performance on modern machines. We instead show that by leveraging *tries* (prefix search trees), we can achieve a two- to ten-fold speed-up in numerical experiments with only moderate memory overhead. For the important problem of ranking permutations, the corresponding tries can be compressed, cutting down the memory requirements [3]. These techniques would have been difficult to develop and benchmark in Python, since ranking cannot be vectorized and is difficult to modularize out into a native library. Using Julia, this was easily overcome.

References

[1] Wallerberger, M. *et al.*, *Comput. Phys. Commun.* **235**, 388–399 (2019)

[2] Wallerberger, M., Shinaoka, H., and Kauch, A., *Phys. Rev. Research* **3**, 033168 (2021)

[3] Wallerberger, M. and Held, K., *arXiv* **2203**:04158 (2022)

The Cluster Coffer: Teaching HPC on the Road

***Philipp Gschwandtner*, Alexander Hirsch, Peter Thoman, Peter Zangerl, Herbert Jordan, and Thomas Fahringer**

University of Innsbruck, Technikerstrasse 21a, 6020 Innsbruck, Austria

Teaching and public dissemination of HPC-related topics can be a difficult task. First, there is a number of sophisticated hardware and software components, each complex on their own and often showing non-intuitive interaction. Second, we consider education in HPC among the more difficult topics in computer science since larger distributed memory systems are ubiquitous yet inaccessible and intangible to students and the general public. By design, most of these systems are physically and logically encapsulated to provide a layer of abstraction that facilitates their use, simultaneously hindering on-hands education.

In this talk, we discuss our Cluster Coffer [1], a mobile miniature cluster computer that reduces the entry barrier to HPC in teaching and public outreach through direct access to all major components of a distributed memory system. It consists of an aluminium case with 16 NanoPI M4 ARM boards that form the compute nodes, while a NanoPC T4 board serves as the login node. All boards are equipped with an actively-cooled Rockchip RK3399 SoC with two higher-power Cortex-A72 and four lower-power Cortex-A53 cores and use Gigabit Ethernet as the interconnect. All components are low-cost and easy to maintain to minimize operational cost. Preliminary peak performance is 101 GFLOPS (HPL Rmax), total system power draw is less than 200 W.

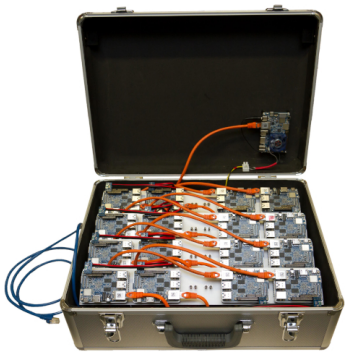

Fig. 1: The fully assembled Cluster Coffer. The blue Ethernet cable on the left connects to the outside world.

We employ the Cluster Coffer for two main use cases. First, we address teaching students in multi-objective optimization and Pareto optimality by comparing execution times of HPL or HPCG against energy consumption for varying cores or DVFS settings. Since the Cluster Coffer is equipped with high-resolution per-node power instrumentation, students can observe the trade-off of e.g. wall time and energy consumption for various points in the problem space.

The second use case demonstrates the possibilities of HPC to the public at the example of a space-weather prediction proxy app, simulating solar winds (originating from the H2020 project AllScale). However, experience has shown that high levels of audience engagement are crucial for maximizing public outreach success. Therefore, the proxy app was enriched with an interactive element by including live input from participating audience through a motion-detecting Kinect camera. It is attached to a host PC, which converts motion information into charged particles and transfers them to the Cluster Coffer, to be incorporated into the continuously running simulation, while simultaneously all current particle data is sent back to the host PC for visualization. This continuous loop allows the audience to directly influence the simulation and observe effects in real time. We have successfully demonstrated the Cluster Coffer at numerous outreach activities such as science fairs and open-day events. CAD blueprints, a bill of materials, and setup software is available online for reproduction [2].

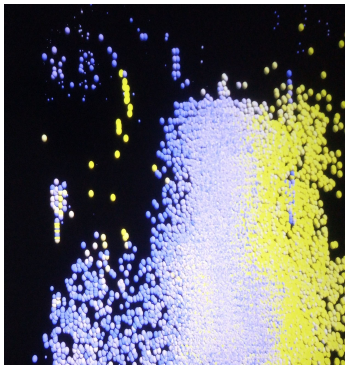

Fig.2: Visualization of charged particles caused by a person walking from left to right. Higher energy is shown in yellow, lower energy is shown in blue.

References

[1] Gschwandtner, Philipp, Hirsch, Alexander, Thoman, Peter, Zangerl, Peter, Jordan, Herbert, Fahringer, Thomas, Journal of Distributed and Parallel Computing **155**, 50 (2021)

[2] `https://github.com/uibk-dps-teaching/cluster-coffer`

Keynote Talk:

NERSC, Perlmutter and HPC in the U.S. Department of Energy Office of Science

Richard Gerber

National Energy Research Scientific Computing Center at Lawrence Berkeley National Laboratory

The National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkley National Laboratory is the mission High Performance Computing and Data (HPC) center for the U.S. Department of Energy Office of Science, the largest funder of fundamental research in the physical sciences in the U.S. NERSC and the leadership facilities at Oak Ridge and Argonne National Laboratories provide HPC at extreme scale and capability for open science research.

NERSC recently deployed *Perlmutter*, an HPE Cray Shasta system with AMD EPYC 7763 CPUs and NIVIDIA A100 GPUs that is currently ranked as the 5th most powerful supercomputer in the world. The Oak Ridge Leadership Computing Facility and Argonne Leadership Computing Facilities will soon deploy systems with exascale performance.

In this talk I will give an overview of NERSC and the HPC facilities landscape in the Advanced Scientific Computing Research program in the Office of Science. I'll describe Perlmutter and NERSC's efforts to enable a broad set of applications and users to productively use the compute capability offered by the system. I'll also describe how NERSC is adapting to changing workloads and how our Superfacility project connect NERSC to other experimental facilities.

Perlmutter is NERSC's first large GPU-accelerated system. To prepare NERSC's large diverse workload, the NERSC Exascale Science Application Program (NESAP) has been working with key application teams since 2019 to prepare their simulation, data analytics and learning/AI codes for GPUs. I'll give an update on performance and porting progress with NESAP applications in the areas of particles and grids, data, electronic structure, learning, molecular dynamics, and Lattice QCD.

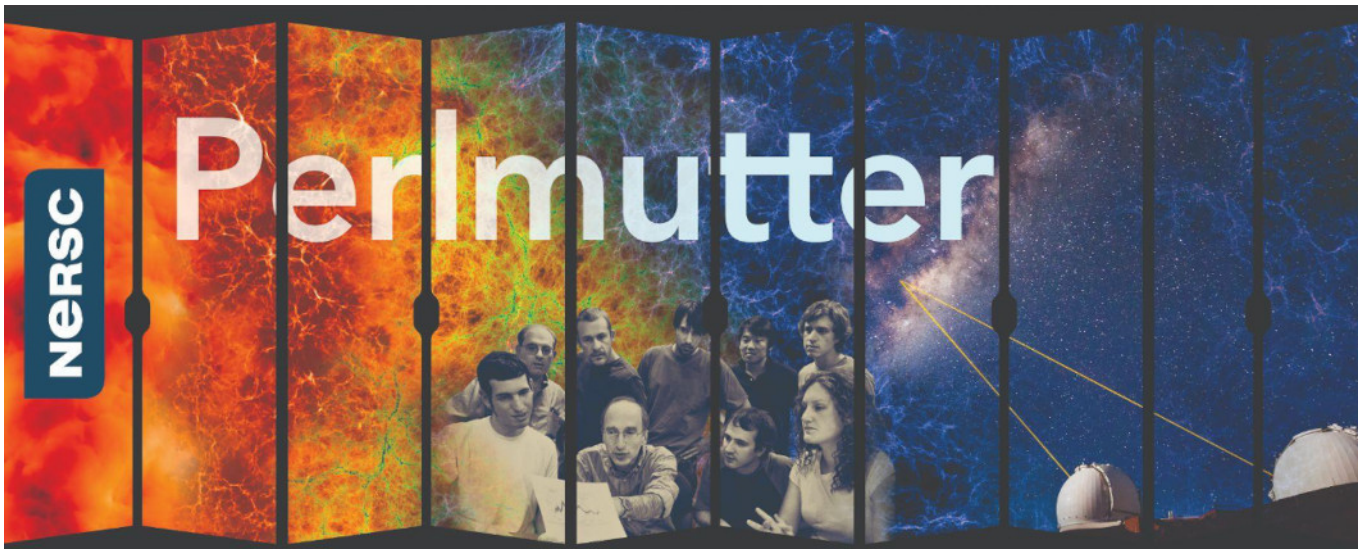

Fig. 1: Perlmutter

Poster session:

Case study: How to select the right Software mix for SMEs

László Ligeti

Governmental Agency for IT Development (KIFÜ) - Hungarian NCC

We strongly believe that HPC usage will increase among the industrial parties too. In Hungary we are focusing especially on the SME sector, as it plays an important role in our national economy.

In the summer of 2021, KIFÜ started market research among industrial players. In order to gain experience and build up relationships with industrial players – by for example assisting them in using HPC – we started with a detailed survey and market analysis for the mechanical engineering sector. This sector was selected for deeper research and analysis due to its high significance and contributions to the country's gross domestic product (GDP). The surveys were sent out to 37 different industrial entities and 24 detailed answers were received back. The results enabled us to identify the key driver of this sectors' stakeholders. We analysed their needs further through eight in-depth interviews. Since then we have regular meetings with small and medium-sized enterprises. One of the key observations is that most of the interviewed companies carry out Research and Development or innovation activities. Emerging computational tasks typically include mechanical, thermal, fluid simulations, measurements, testing and visual design.

Our mission is to support SMEs' economic growth with HPC solutions and softwares.

We experienced that several companies develop their own software to solve computationally intensive tasks, however, it is also typical that they use purchased software. Most of the cases they supplement their developments in areas where its functions are not sufficient to perform the task in depth.

Purchasing a software change might require a hardware investment also, so they will be considered carefully in many cases.

Based on the above mentioned we made the decisions to purchase HPC softwares by KIFÜ and make it available for the SME sector. It is important to underline that it is not enough for the centre to provide basic information and applications related to HPC, users also need personalised advice on how to optimise the use of resources. The aim is to build up a growing collection of software from which users can choose the software they need to carry out their research. This will be developed by the Competence Centre in the light of evolving needs. User training is also very important, and will be provided through training courses, workshops and online seminars.

The use case drives you through in details what we have done so far, and will present you the process of how we selected the right software mix.

Poster session (online):

Heterogeneous Workflows Scheduling in the Computing Continuum Systems

Vladislav Kashansky[a,c] **and Radu Prodan**[a]

[a] *Institute of Information Technology, University of Klagenfurt, Austria*
[c] *Eteronix GmbH, Austria*

Computing continuum systems [1] normally characterized by large amounts of data/compute-intensive tasks, utilizing heterogeneous computing and storage resources in non-stationary settings. Remarkable problem that still requires research effort is to enable efficient workflows scheduling in such complex environment. Multiple system temporal/spatial scales [1,2], heterogeneous workloads with complex inner structures and the arrival patterns make solution of the scheduling problem exponentially hard, requiring problem-specific heuristics. In this talk we will discuss how to approach this problem from simulation perspective, discuss important algorithmic and architectural aspects.

In particular, we will discuss some relaxations and two-phase heuristic for the minimum makespan scheduling of tasks and computing machines on large-scale computing systems, consisting of various matching phases with subsequent event-based MILP method for schedule generation. We evaluated the scalability of the heuristic using parallel MPI-variation of Constraint Integer Programing (SCIP) framework with various configurations based on data sets, provided by the MACS framework [3].

References

[1] Kashansky, Vladislav, Gleb Radchenko, and Radu Prodan. "Monte Carlo Approach to the Computational Capacities Analysis of the Computing Continuum." International Conference on Computational Science. Springer, Cham, 2021.

[2] Kashansky, Vladislav, et al. "The ADAPT Project: Adaptive and Autonomous Data Performance Connectivity and Decentralized Transport Network." Proceedings of the Conference on Information Technology for Social Good. 2021.

[3] MACS: Modular architecture for complex computing systems analysis. `http://www.edmware.org/macs/`, accessed: 2021-01-29

Acknowledgement This work has been supported by ADAPT Project funded by the Austrian Research Promotion Agency (FFG) under grant agreement No. 881703

Poster session:

arcControlTower for distributed HPC job management

Jakob Merljak[a], **David Cameron**[b], **and Andrej Filipčič**[a]

[a] *Jozef Stefan Institute, Jamova 39, 1000 Ljubljana, Slovenia*
[b] *University of Oslo, Problemveien 7, 0316 Oslo, Norway*

arcControlTower (aCT) is a generic job management framework capable of handling large volumes of jobs over many HPCs managed by the ARC middleware [1]. The advanced job management capabilities of aCT are useful for workflows of many users and projects. The ATLAS experiment, for example, uses aCT to run up to 0.5 million jobs per day on over 30 HPCs.

aCT was originally developed for the ATLAS experiment. To extend the functionality, a REST interface was added to aCT which allows the framework to be set up as a web service and provide remote job management capabilities to many clients [2]. A client program was implemented that integrates with the REST interface to provide a job management workflow similar to other middleware client tools, e.g. ARC Client Tools. The REST interface provides a suitable API for the implementation of a web browser application for job management.

Remote job management requires special handling of job input files. Two data management solutions were integrated into aCT. First, a data management endpoint was implemented in the REST interface which can be used to upload job input files to a directory local to the aCT service. The service can then pass them to the middleware. Second, the aCT client program can upload input files to a remotely accessible WebDAV directory that is provided by an external data management service, such as dCache.

Figure 1 shows how the aCT service is used in an HPC workflow. The service balances job distribution to all available HPCs to maximize resource usage available to users. The job input files can be provided to the HPCs through either the internal data management system of the aCT service or external WebDAV services.

A testing aCT service has been set up for the Vega supercomputer, ARNES HPC and IJS NSC HPC to evaluate workflows for different users. The final goal is to set up an aCT service that provides job management capabilities for the computing resources of the entire Slovenian national supercomputing network (SLING). Such a service would provide a more accessible and more efficient use of the Slovenian HPC infrastructure with transparent processing of data stored on national facilities.

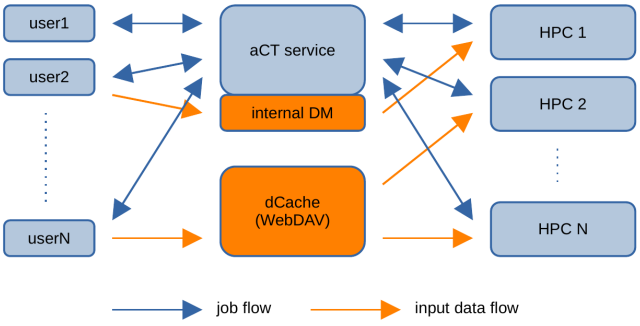

Fig. 1: A high level depiction of the interaction of the aCT service and its integrations with the HPC job and data workflows.

References

[1] Nilsen, J. K., Cameron, D., and Filipčič, A., Journal of Physics: Conference Series 664, 62042 (2015)

[2] Merljak, J., Interfaces for arcControlTower, (2018)

[3] SLING: Slovenian national supercomputing network, Available at: `https://www.sling.si/sling/en/`

Poster session (online):

HPC in design of devices for chiral nanotechnology

Renáta Rusková **and Dušan Račko***

Polymer Institute of the Slovak Academy of Sciences, Dúbravská cesta 9, 845 41 Bratislava, Slovakia
** dusan.racko@savba.sk*

In this contribution, we investigate the behaviour of knotted polymers confined in nanochannels with complex geometries. It is known that DNA is a natural source of molecular knots [1], but synthetic knots became possible thanks to the progress in macromolecular chemistry and synthesis.

Topologically, the majority of knots are chiral. However, not all knotted molecules are chiral, and not all chiral molecules are knotted. For a group of knots with increasing complexity characterized by a crossing number up to 11, there exists in total 801 prime knots, out of which 20 are achiral or amphichiral knots [2].

Polymer knots represent a group of chiral superstructures. The topological state of molecules affects physical, biophysical and biological properties and behavior of the molecules.

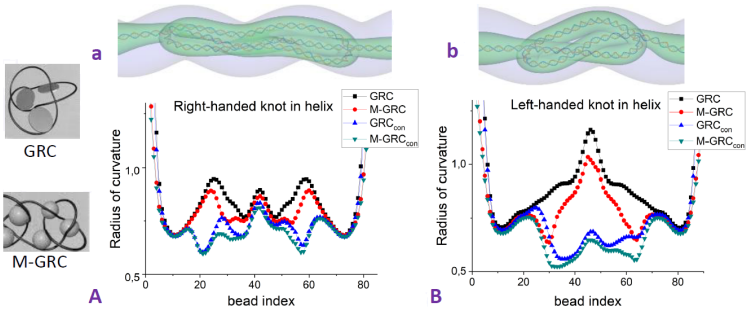

Fig. 1: Symmetry breaking observed for chiral knots in chiral confinement. Right-handed and left-handed knots in equichiral and antichiral confinements respectively.

Experimentally, the topological state of DNA molecules is detected by nanodevices that employ translocation of the polymers through nanochannels. In our contribution, we show that nanochannels with special geometry are able to detect knotted state together with its chirality [3].

HPC molecular simulations and computer experiments are an emerging approach that provides possible advantages for study and design of the nanochannels for topological sorting of molecular knots.

Acknowledgements: The work was supported by the Grant Agency of the Ministry of Education, Science, Research and Sport of the Slovak Republic No. VEGA 2/0102/20 and Slovak Development and Research Agency SRDA (Slovak-Austrian Billateral Project) SK-AT-20-0011.

References

[1] Frisch, H.L.; Wasserman, E. Chemical Topology1. J. Am. Chem. Soc. 1961, 83, 3789–3795. **53**(12), 4770 (2020)

[2] Vaughan, F.R. Jones A Polynomial Invariant for Knots via von Neumann Algebras. Bull. Am. Math. Soc. **12** 103 (1985)

[3] Rusková, R.; Račko, D.;Channels with Helical Modulation Display Stereospecific Sensitivity for Chiral Superstructures. Polymers **12** 103 (2021)

Poster session (online):

The free volume in PVME-water mixture as obtained from HPC simulations

Dušan Račko*

Polymer Institute of the Slovak Academy of Sciences, Dúbravská cesta 9, 845 41 Bratislava, Slovakia
**dusan.racko@savba.sk*

In this contribution we investigated the free volume microstructure in polymer-water mixture obtained by atomistic molecular dynamics simulations. The structures were obtained by simulated cooling, from 350 K down to 175 by decreasing temperature by 25 K. In each step, the density of the system was equilibrated during 20 ns, followed by another 10 ns production run in NVT ensemble.

The free volume was analysed by numerical volumetric integration by using voxellization of space and by semi-analytical method implementing Connoly's method for calculation of the solvent accessible volumes. Using these methods, the cavity volumes, with their numbers and distributions were obtained.

The obtained properties were further compared with the quantities predicted from the free volume theory, such as free volume fractions at glass temperature, volume expansivities, and free volume numbers and volumes from the Positronium Annihilation Spectroscopy.

The main results show, that the free volume has more complicated behavior than predicted by the classical free volume theory in Fox equation, while the free volume around specific parts of the polymer changes based on the presence of hydrofilic and hydrophobic groups.

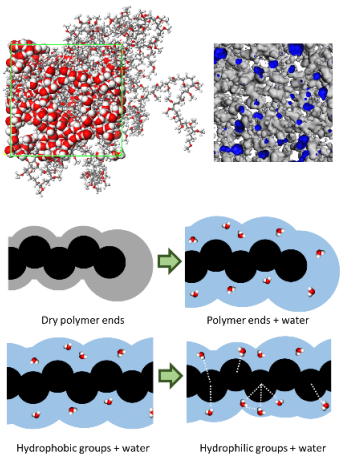

Fig. 1: Free volume in a polymer-water mixture.

The original research project was done in collaboration with Department of Molecular Simulations at Polymer Institute of the Slovak Academy of Sciences (Slovakia), Donostia International Physics Center and Department of Physics of Materials in San Sebastian (Spain), and Department of Industrial and Applied Genomics, IBM AI and Cognitive Software Organization, IBM Almaden Research Center, (USA).

The free volume simulations present a challenge for the computational experiments from two aspects. In the simulated cooling, the relaxation of the structure slows down exponentially with the cooling rate. Currently, computationally achievable cooling rates are still 9 orders of magnitude faster than their experimental counterparts, shifting the difference between the simulated and experimental glass transitions to tens of degrees higher temperatures. Secondly, the volumetric computations are computationally very expensive, while the complexity grows with power law N^3 of number of atoms for determination of the free volumes and N^4 for finding cavities.

Acknowledgements: The work was supported by the Grant Agency of the Ministry of Education, Science, Research and Sport of the Slovak Republic No. VEGA 2/0102/20.

References

[1] Capponi, S.; Alvarez, F.; Racko, D.; "Free volume in PVME-water mixture." Macromolecules **53**(12), 4770 (2020)

[2] Racko, D; Capponi, S.; Alvarez, F.; Colmenero, J.; "The free volume of poly(vinyl methylether) as computed in a wide temperature range." J. Chem. Phys. **134**(4), 044512 (2011)

Poster session:

Neural network potentials for fusion material research

Shokirbek Shermukhamedov **and Michael Probst**

Institute of Ion Physics and Applied Physics, University of Innsbruck, Technikerstrasse 25, 6020 Innsbruck, Austria

The interaction of plasma gases, for example argon (Ar), with the first wall and tungsten (W) divertor is one of the main problems for using tokamaks as fusion devices since heavy atoms can contaminate the plasma and cause disruptions. Despite a large amount of experimental and computational data on tungsten, sputtering yields of tungsten caused by low energetic Ar atoms are still poorly understood [1]. The best way to describe such systems is fitting very accurate atomic potential energy functions based on quantum chemical (DFT) calculations. One of the powerful tools for this purpose is the High Dimensional Neural Network Potential (HDNNP) [2].

In this work we trained a feedforward neural network with two hidden layers of 25 nodes each. The input nodes consist of weighted radial and angular Behler-Parrinello type symmetry functions [3]. The reference data set consists of 6230 configurations. 5655 total energies and 1 140 705 forces are used to train the final NNP. The remaining 10% of structures containing were used as a test set, to validate the interpolation capacity of the neural network potential energy function. With 50 training steps, the RMSE in the test set converged to 0.94 meV/atom for energies and 0.17 eV/Å for atomic forces, the corresponding values in the training set are 0.97 meV/atom and 0.19 eV/Å. There was no evidence for overfitting in the training process. The range of per atomic potential energies and forces present in the training data is quite large. The correlation of the NNP vs. the reference energies and the corresponding forces together with their distributions are shown in Figure 1. The trained HDNNP will be used for further sputtering simulations.

Training HDNNP requires a large amount of memory. The computational results have been obtained using the LEO4 cluster, a linux cluster (Intel Xeon Broadwell/Skylake processors) of the University of Innsbruck and the Vienna Scientific Cluster VSC-4 (Intel Xeon Skylake Platinum processors). The computational effort to train the final NNP was 1010 core hours with 28 cores in parallel, utilizing 30GB of RAM.

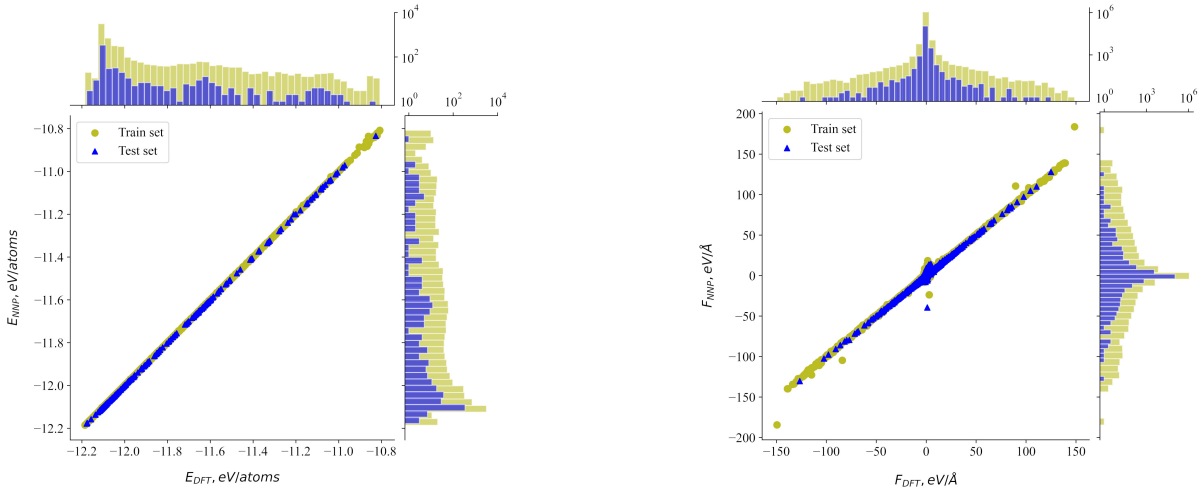

Fig. 1: Correlation between NNP – derived and DFT energies (left) and between the forces (right) for all configurations in the training and test set.

References

[1] Rieth, M., Doerner, R., Hasegawa, A., Ueda, Y. & Wirtz, M., J. Nucl. Mater. **519**, 334 (2019)

[2] Singraber, A., Behler, J. & Dellago, C., J. Chem. Theory Comput. **15**, 1827 (2019)

[3] Gastegger, M., Schwiedrzik, L., Bittermann, M., Berzsenyi, F. & Marquetand, P., J. Chem. Phys. **148**, 241709 (2018)

Poster session (online):

Computing the gravitational (electrostatic) potential on nested Cartesian meshes using the convolution method

Eduard Vorobyov[a], **James McKevitt**[a], **and Igor Kulikov**[b]

[a] *University of Vienna, Department of Astrophysics, Vienna, 1180, Austria*
[b] *Institute of Computational Mathematics and Mathematical Geophysics SB RAS, Lavrentieva ave., 6, Novosibirsk, 630090 Russia*

Gravity plays a key role in the dynamics of many astrophysical objects. Finding the gravity force of an arbitrary mass distribution requires the numerical solution of the Poisson equation.

$$\nabla \cdot \nabla \Phi = \triangle \Phi = 4\pi G \rho, \tag{1}$$

where Φ is the gravitational potential, $\triangle$ is the Laplace operator, G is the gravitational constant, and ρ is the mass volume density. Often such a solution must be obtained on non-homogeneous numerical meshes, which allow a finer numerical resolution in the regions with higher mass concentration. An example of such meshes are so-called nested grids (see Fig. 1). A contemporary solution of the Poisson equation on the nested grids involves moving from the coarsest toward the finest grid, using for the boundary conditions the potential at the coarser grid. However, such an outside-in solution is difficult to parallelize. It also requires the knowledge of the boundary potential at the outer coarsest grid, which is obtained through a time-consuming multipole expansion of the Laplace equation $\triangle \Phi = 0$ in spherical harmonics.

An alternative approach is to employ the convolution theorem for computing the gravity using its integral representation [1,2]

$$\Phi(\mathbf{r}) = -G \iiint \frac{\rho(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d^3\mathbf{r}', \tag{2}$$

where the integration is performed over the entire volume occupied by an object of interest. On a discrete three-dimensional computational mesh the integral takes the form of a triple sum. This method involves the fast Fourier transform and is usually not computationally efficient on non-homogeneous meshes (such as nested grids). Here, we present a modification to the convolution method, which can be efficiently applied to nested grids (or adaptive meshes in general). The modified method is easily parallelizable and does not require the costly multipole expansion to find the boundary conditions at the outermost grid. This method can also be applied equally well to finding the electrostatic forces on nested grids.

Figure 1 presents an example of the mesh decomposition which is used to compute the gravitational potential on nested grids. A nested grid with only two levels of refinement is shown for simplicity. The solution procedure consists of three steps: 1) calculating the potential $\Phi_{\rm m=2}$ on the finer $m = 2$ grid, 2) calculating the potential $\Phi_{\rm m=1}$ on the coarser $m = 1$ grid, and 3) calculating the potential $\Phi^{\rm hole}_{\rm m=1}$ on the coarser $m = 1$ grid, but with a central hole. The final step is a combination of the three values to form the potential on the nested mesh, including an interpolation of $\Phi^{\rm hole}_{\rm m=1}$ on the finer $m = 2$ grid. Each of the three steps is independent and can be done in parallel, unlike the outside-in solution procedure outlined above.

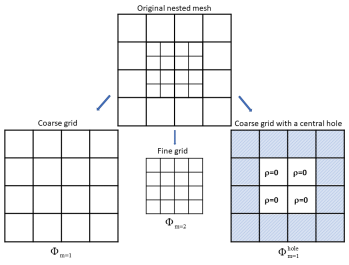

Fig. 1: Nested mesh decomposition. Note that the r.h.s. grid has a hole filled with zero densities.

References

[1] Binney, J., & Tremaine, S. 1987, Galactic Dynamics (Princeton, NJ: Princeton Univ. Press)

[2] Hockney, R. W., & Eastwood, J. W. 1988, Computer Simulation Using Particles (CRC Press)

Poster session:

Benchmarking using Relion GPU workloads

***Stefano Elefante*, Alois Schlögl, Andrei Hornoiu, and Stephan Stadlbauer**

Institute of Science and Technology Austria (ISTA)

Benchmarking is an important aspect for running HPC resources, to ensure quality controlled performance as well as efficient purchases. REgularised LIkelihood OptimisatioN (RELION) [1] is a tool for analysing cryo-electronmicroscopy (CryoEM) data, it is extensively utilised by several research groups at the ISTA HPC cluster [2] and it is very demanding in terms of CPUs and GPUs utilization as well as I/O operations.

Table 1: Performance measured in wallclock time (HH:MM:SS).

Node	**GPU Card**	**CPU**	**Power**	**Memory**	**HD disk**	**SSD disk**
Type	*Type*	*Type*	*Watt*	*GB*	*HH:MM:SS*	*HH:MM:SS*
A	GTX1080Ti	Xeon E5-2650 v4 @ 2.20GHz	250	12	6:37:50	6:36:41
B	RTX2080Ti	Xeon E5-2680 v4 @ 2.40GHz	250	12	5:26:11	5:15:32
C	RTX3090	EPYC 7502 @ 2.50GHz	350	24	4:41:37	4:18:46
D	A10	EPYC 7F72 @ 3.20GHz	150	24	3:37:39	3:35:10
E	A40	EPYC 7452 @ 2.35GHz	300	48	4:19:02	4:07:03

Benchmarks tests (3D classification [3]) have been performed using different node configurations (see table 1). Preliminary results are shown in figure 1. These test results suggest: i) the type of local storage (SSD vs HDD) is of very little importance, ii) newer GPUs are typically faster then those from earlier generations, iii) the CPU (e.g clock frequency) is also of critical importance as this will most likely explain why configuration E does not perform as well.

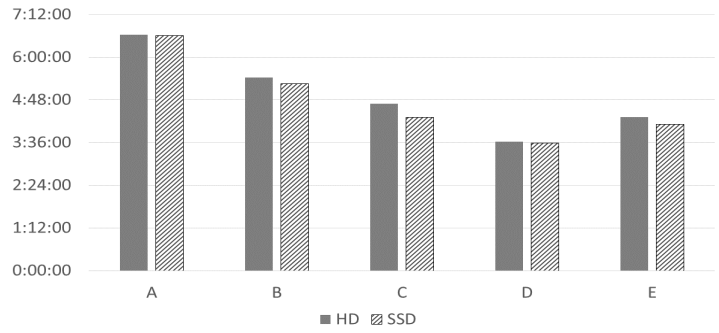

Figure 1: Elasped time: for different node configurations.

These results will be useful for the procurement of future systems. They show that the bottleneck for Relion is mostly the CPU and GPU card. Indeed, SSD or HDD disks do not have a significant impact on the processing time. The CPU and GPU configuration is therefore the major parameter to consider.

References

[1] Sjors H.W. Scheres. RELION: Implementation of a Bayesian approach to cryo-EM structure determination, J Struct Biol. 2012 Dec; 180(3): 519–530. doi: 10.1016/j.jsb.2012.09.006

[2] Schlögl, A., and Kiss, J. Scientific Computing at IST Austria. p.28, Booklet AHPC (2017)

[3] Relion Benchmarks, `https://www3.mrc-lmb.cam.ac.uk/relion`

Poster session (online):

Catch22 = Biomed2022

Boštjan Murovec[a], ***Leon Deutsch***[b], **and Blaž Stres**[b,c,d]

[a] *Laboratory for Machine Intelligence, Faculty of Electrical Engineering, University of Ljubljana, Ljubljana, Slovenia*
[b] *Department of Animal Science, Biotechnical Faculty, University of Ljubljana, Ljubljana, Slovenia*
[c] *Institute of Sanitary Engineering, Faculty of Civil and Geodetic Engineering, University of Ljubljana, Ljubljana, Slovenia*
[d] *Department for Automation, Biocybernetics and Robotics, Jožef Stefan Institute, Ljubljana, Slovenia*

The availability of human microbiome raw data deposited in biomedical databases and its production by high-throughput 'omics technologies are increasing at unprecedented rate. Extended cohorts comprising thousands of participants are increasingly being equipped with ample metadata, shaped by EU wide and national GDPR regulations. However, the actual conversion of raw sequencing outputs to data matrices amenable for extended machine learning is giving rise to scattered publications with low statistical power for effective separation of healthy human states from various disease types associated with human microbiome. Only few meta-analyses utilizing this rich data structure (taxonomy, functional genes, enzymatic reactions, metabolic pathways, metabolites, proteomes, transcriptomes, patient metadata) encompassing millions of variables have been conducted so far.

The availability of appropriate data is merely the beginning as the field has become extremely hungry for computational resources despite a number of hurdles: (i) many algorithms are CPU intensive and require submission of large jobs that are less frequently accepted by HPC schedulers; (ii) a single metagenomic analysis requires a large amount of data on a disk as DNA sequence information for a single disease can amount to several terabytes; (iii) the data need to be accompanied by large prebuilt databases of known DNA patterns to be mapped to the input dataset(s); (iv) the transfer and management of such large input, output files and underlying databases requires effective internet connections to sustain massive data transfer in real time as well as up-/down-loads from repositories; (v) the resulting datamatrices can encompass million variables with peculiar biomedical data characteristics. At the same time the biomedical field is (vi) noticeably less adapted to the established HPC protocols as reliance on GPU parallelism is severely limited, since it is not generally possible to divide input sets and databases into uniform chunks as in image processing and numerical calculations; (vii) processing of metagenomic data often cycles between CPU-intensive and disk-intensive time intervals. Hence, a large pool of CPUs together with a decent amount of RAM are required without guarantee of high and uniform CPU utilization; (vii) the obtained biological data vary in quality in an unpredictable way as large amount of input data is discarded during a quality control prescreening step, which substantially reduces computation time of later stages of analysis, while a high quality sample of equal size may require an order of magnitude more wall time than a typical or a low quality one.

Finally, as the character of biomedical studies is unlikely to adapt to the current prerequisites of HPC, the helping hand can only come from the side of HPC in the form of: (i) large disk arrays with fast connection to CPU nodes; (ii) nodes with large number of CPUs that are available for one computation for long time intervals; (iii) HPC policies should fairly accommodate less penalization of non-optimal CPU utilization, as well as they should accommodate the stated uncertainty in required wall time prediction.

Poster session:

Interactive Hands-on Introduction to Parallel Programming

Leon Kos

LECAD laboratory, Faculty of Mechanical Engineering, University of Ljubljana, Slovenia

Practical learning with notebook-style (Jupyter based) hands-on exercises has proven to be a useful approach for beginners and researchers focusing on getting results reproduced. Within *xeus-cling* Jupyter kernel for C++, FORTRAN, Python it is possible to use OpenMP, MPI and CUDA constructs as demonstrated in Fig. 1. Besides interactive C++ we can use notebook cell offloading with `%%file` and `!execute` directives.

The Introduction to Parallel Programming massive open online course (MOOC) [1] divided into 5 weeks and 20 steps each targets (parallel) programming beginners. Case studies and hands-on tutorials are done interactively on the cloud and therefore no computer is needed to learn the programming concepts. The notebooks are embedded in the MOOC platform [1] in two conceptual ways. For explanation of the concepts within video or article steps the notebooks are embedded at the end of the web page with the purpose of analysing the code snippet. Second concept is exercise where the problem is explained with a task to solve it. Advices are given, including quick-help for OpenMP and MPI constructs that appear as a cell (e.g. `?omp::parallel`) that splits the window and shows a "pragma" or a function from reference manuals. The notebooks are given in a language independent way so that one can compare C or Fortran example for OpenMP with addition of Python MPI using Mpi4Py examples. The cloud infrastructure "behind the scenes" is providing a cluster with a swarm of compute nodes that are seamlessly accessed through MOOC platform without the need of login. The MOOC platform provides Learning Tools Interoperability (LTI) to anonymously identify users and track the learners' progress on a run. Within each learning step we provide alternative links to to Binder (without GPU support) and Google Colab (without inline help) notebooks that opens our programming examples in a separate web page.

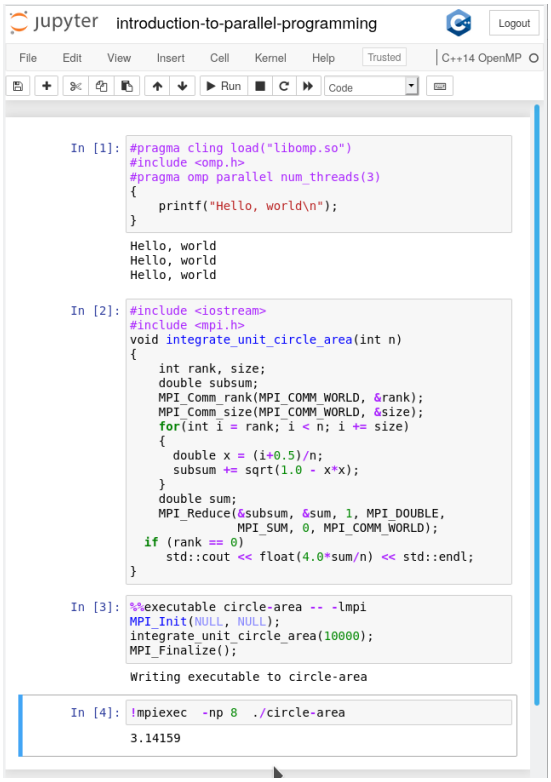

Fig. 1: Jupyter notebook in xeus-cling environment introducing OpenMP and MPI interactively.

Developed MOOC remote training platform [1] at the university provides up to 40 servers running as a Docker images under Windows WSL2 Linux with NVIDIA P400 GPUs. Single *HAproxy* access server provides Cross-Origin Resource Sharing (CORS) of embedded HTTPS pages and load balancing bases on LTI or origin cookies. Scaling for over 400 simultaneous user sessions is dynamic with pruning of inactive notebooks. After the notebook server is identified idle it goes to sleep and is woken up over the lan (WOL) when more users are connecting. The training platform developed shows capability to run other clustered paradigms (Kubernetes, Hadoop, Spark, ...) for HPC jobs with high availability and resilience built in.

References

[1] Kos L, Badovinac K.A., Bogdanović L and Brank M, MOOC Introduction to Parallel programming, (2022)
`https://www.futurelearn.com/courses/interactive-hands-on-introduction-to-parallel-programming`
`https://github.com/kosl/ihipp-examples`
`https://bitbucket.org/lecad-peg/remote-training-platform/`

Poster session:

Erasmus+ project: SCtrain

Pavel Tomšič

University of Ljubljana, Faculty for Mechanical Engineering, Aškerčeva 6, 1000 Ljubljana, Slovenia

The currently established curricula in HE do not offer a comprehensive inclusion of HPC knowledge for future professionals. We aim to establish international cooperation and exchange of knowledge between acknowledged professionals in the field of High-Performance Computing (HPC). Each state involved in this Erasmus+ call will build new HPC competence centers through the EuroHPC joint undertaking initiative. Basic educational programs are foreseen – HPC literacy programs, but further systematical approach (bottom-up) is needed to accelerate the level of knowledge. Throughout knowledge research and development by educators we are raising the level of competences for future science, technology, engineering and mathematics (STEM) professionals between Central European countries, including Italy, Slovenia, Austria and Czech Republic.

The project addresses multiple target groups; **For students:** Gaining skills that present a knowledge gap in current HE system, enabling future competitiveness; **For educators:** Gaining skills needed for conducting courses with HPC. Raising level of competences in theoretical, programming, mathematical and teaching skills; **For business experts:** Gaining skills needed for professional growth and competitiveness with introduction of HPC; **For future HPC HE courses:** Creation of a comprehensive set of learning and teaching material, which will embrace the field of HPC in Engineering and Data Science.

The project addresses the issue of raising the HPC knowledge by researching several topics in the field of HPC in Engineering and Data Science. Within each research objective, outputs include creation of a knowledge base in form of lectures (core knowledge) and exercises/tutorials (hands-on approach), tested through training events that benefit both the verification of knowledge transfer to students and the enhancement of the educators by gaining experience. The research topics include:

- Development of program on the topic of HPC in Engineering - with a focus on FEM (Finite Element Method) - finished in Autumn 2021,
- Development of program on the topic of HPC in Data Science - Introduction to Parallel Programming - finished in Spring 2022,
- Development of program on the topic of HPC in Engineering - with a focus on CFD (Computational Fluid Dynamics) - currently ongoing,
- Development of program on the topic of HPC in Data Science - with a focus on IOT and Big Data,
- A combined collection of knowledge on the topic of HPC in Engineering and Data Science.

Up to Spring 2022 we have successfully developed the first two topics and are moving forward on the topic of HPC in Engineering with a focus on CFD, that will be available in the Autumn of 2022. Results of our project are publicly available for all interested public through our website and we even organise a training week at the end of each topic research to test the gathered knowledge. For more information and participation at the next workshop you can visit the project website: `http://sctrain.eu/`

Project partners

• University of Ljubljana, Slovenia • VSB – Technical University of Ostrava, Czech Republic • CINECA Consorzio Interuniversitario, Italy • Technische Universitaet Wien, Austria

This project has been funded with support from the European Commission.

Keynote Talk:

Multiscale Methodologies for Electrolyte Characterization

Matteo Ambrosetti

Northvolt AB

When combined with clean electricity production, electric mobility has the potential to significantly reduce CO2 emissions of the road transportation sector, which currently accounts for about one fifth of total CO2 emissions in Europe.

Northvolt was founded in 2016 with the mission to build the greenest battery in the world with a minimal carbon footprint and the highest ambitions for recycling. To accelerate the transition to renewable energy, Northvolt is building Europe's largest li-ion battery factory, located in Skellefteå, in northern Sweden. The factory will supply sustainable, high-quality battery cells and systems.

In Västerås, one hour west of Stockholm, Northvolt Labs is instead used to facilitate the critical task of developing and maturing battery cells that provide the bedrock technology of new energy systems and products. The principal functions of Northvolt Labs include refining manufacturing processes and producing demonstration and sample cells for automobile manufacturers, industrial vehicle providers and other customers requiring battery systems.

One of the major projects currently ongoing at Northvolt Labs is the establishment of the world's first R&D campus covering the entire battery ecosystem, with the aim to become Europe's leading campus for battery technologies. A new R&D facility is already under construction, which will enable the development of novel battery cell materials and products.

The Advanced Materials Team (AMT), led by Hwamyung Jang, is involved in research and development of next-generation batteries. The team is setting up a pilot line for prismatic lithium metal cells, and therefore requires support from a committed and technologically outstanding supplier of Formation and Aging equipment. At the same time the AMT is working on the development of fast and reliable computational approaches able to accelerate the development of the next-generation batteries.

Several experimental and theoretical studies have focused on systems formed by pure and mixed solvents and lithium salt. However, various experimental methods are needed to measure their structural and transport properties. Therefore, it seems that the convenient choice would be to apply computational methods such as molecular dynamics (MD) simulation to obtain reliable predictions for the required properties.

In this talk I will present how High Parallel Computing (HPC) resources have been adopted by the AMT to leverage and accelerate the optimization of electrolytes for Li-ion batteries. A computational framework based on molecular dynamics simulations is proposed as a tool to explore the existing relation between molecular structure and ionic conductivity of the electrolyte with the final aim of developing a novel organic-based electrolyte able to enhance the performance of lithium-ion batteries (LIBs).

High-performance electricity consumption prediction

Tomaž Čegovnik[a], Andrej Dobrovoljc[b], *Janez Povh*[b,c], and Matic Rogar[c]

[a] *3Tav d.o.o., Podbreznik 15, 8000 Novo mesto, Slovenija*
[b] *Razvojni center Novo mesto, Podbreznik 15, 8000 Novo mesto, Slovenija*
[c] *University of Ljubljana, Faculty of mechanical engineering, Aškerčeva 6, 1000 Ljubljana, Slovenija*

The measurement of electricity consumption at 15-minute granularity is increasingly mandated in the EU, including for households, and this also allows, once sufficient data have been collected, the prediction of future consumption at the same time intervals. In this paper, we present preliminary results of the industry project that aimed to build AI models to predict the next day electricity consumption at 15-minute granularity. We identified the main influencing factors, developed scripts and databases to collect data about these features and the electricity consumption data for each 15-minute interval, and finally developed AI models to predict the next day electricity consumption, for each 15-minute interval and for each consumer.

The models that we generated are based on (i) similar days method combined with random forests and artificial neural networks, (ii) generalized linear models (GLM), and (iii) ARIMA (Autoregressive Integrated Moving Average) methodology. We implemented and tested these models on a dataset of approx. 2500 real customers, consisting of their 15-minute interval consumptions for the period of 2019-2021. The main influencing factor was the weather for which we collected (i) the data about the true weather: temperature, precipitation, pressure, humidity, wind speed, and power or solar radiation for each 30 minute time interval, for each weather station in Slovenia, and (ii) the data about the weather forecasts for the next three days, with the same weather variables, but with granularity of 4 hours, for every subregion in Slovenia having its own weather forecast.

Computing 2500 AI models in a given time window of few hours, where computation of each model demands approximately 5 minutes, is a computationally very intensive job.

We therefore decided to use an HPC and libraries for parallel computation. We used R with libraries *parallel* and *doParallel* and Python for parallel generation of models and predictions of consumptions. We also implemented automatic workflow for retrieving, cleaning and storing the data about the previous day consumptions, about the previous day weather and about the next 3 days weather forecasts into the MongoDB database. We report numerical comparison of predictions based on different models and demonstrate that scaling within one computing node is almost linear, while internode parallelization has not been tested so far. For a better user experience, a visualization of past consumptions and of predicted consumptions was provided based on Grafana tool.

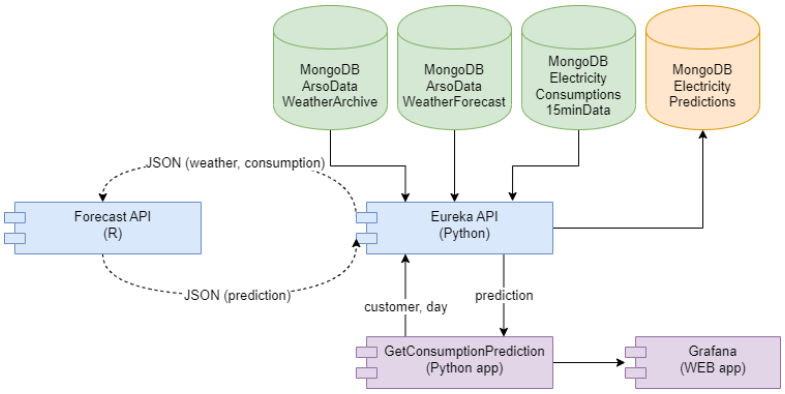

Fig. 1: The scheme of the whole workflow.

References

[1] Čegovnik, T., Dobrovoljc, A., Povh, J., Rogar, M., Electricity consumption prediction using artificial intelligence. In: **SOR '21 proceedings : the 16th International Symposium on Operational Research in Slovenia**: September 22 - 24, 2021, online. Ljubljana: Slovenian Society Informatika, Section for Operational Research, 2021

Speeding up Vectorized Benchmarking of Optimization Algorithms

Aleš Zamuda

University of Maribor, Faculty of Electrical Engineering and Computer Science, Koroška cesta 46, SI-2000 Maribor

This contribution focuses on speeding up of vectorized benchmarking, which includes optimization algorithms [1]. These algorithms are used within Machine Learning (ML) workloads together with benchmarking in order to compute the performance of instances of such algorithms on a whole benchmark [2]. Such performance provides a more general evaluation of an algorithm's applicability, i.e. an intelligence generality. However, the computational time for whole benchmark evaluation extends significantly compared to a single instance evaluation and hence speeding up might be required under time-constrained conditions, especially when e.g. used for robotic deep sea underwater missions [2].

A closer observation of execution times for workloads processed in [2] is provided in Fig. 1, where it is seen that the execution time (color of the patches) changes for different benchmark executions. Therefore, it is useful to consider speeding up of benchmarking through vectorization of the tasks that a benchmark is comprised of. These include e.g., parallel data cleaning part of an individual ML tile [1] or synchronization between tasks when executing parallel geospatial processing [3].

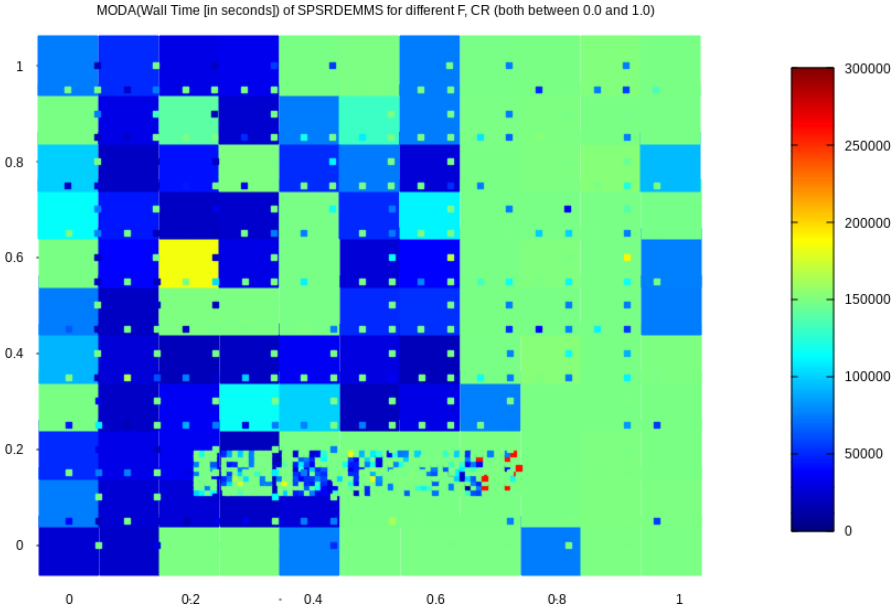

Fig. 1: Execution time of full benchmarks for different instances of optimization algorithms. Each patch presents one full benchmark execution to evaluate an optimization algorithm.

To enable the possibilities of data cleaning (preprocessing) as well as geospatial processing in parallell, such opportunities first need to be found or designed, if none yet exist for a problem tackled.

Therefore, this contribution will highlight some experiences with finding and designing parallell ML pipelines for vectorization and observe speedup gained from that. The speeding up focus will be on optimization algorithms within such ML pipelines, but some more future work possibilities will also be provided.

Acknowledgement: this work is supported by ARRS programme P2-0041 and EU project no. 957407.

References

[1] Zamuda, A., Lloret, E., Optimizing Data-Driven Models for Summarization as Parallel Tasks. Journal of Computational Science **42**, 101101 (2020)

[2] Zamuda, A, Brest, J., Self-adaptive control parameters' randomization frequency and propagations in differential evolution. Swarm and Evolutionary Computation **25C**, 72-99 (2015)

[3] Zamuda, A., Hernández Sosa, J. D., Success history applied to expert system for underwater glider path planning using differential evolution. Expert Systems with Applications **119**, 155-170 (2019)

Modelling-aided materials discovery at Montanuniversität Leoben

***Domink Gehringer*, David Holec, and Lorenz Romaner**

Department of Materials Science, Montanuniversität Leoben, Franz-Josef-Straße 18, A-8700 Leoben, Austria

Without any doubt, computer modelling has become an important part of modern materials science. For example, it allows for checking hypotheses originating from experimental observations by performing simulations with imposed artificial constraints; parallel computations open possibilities for high-throughput searches to narrow down and focus experimental exploration; or machine-learning algorithms combined with large data-sets are capable of unbiased predictions that can lead to unexpected discoveries.

In this talk, we will provide an overview of modelling activities of the Computational Materials Science (CMS) group at the Montanuniversität Leoben, which are enabled by the computational resources of the Vienna Scientific Cluster. The CMS group heavily employs *ab initio* quantum-mechanical calculations (primarily using the VASP and Wien2k codes) focusing on phase stability and phase transformations in various intermetallic alloys, diffusion in bulk as well as along extended defects, segregation to grain boundaries and dislocations, hydrogen production (via pyrolysis) and storage in carbon-based nanostructures, or exploration of electronic structure of functionalized quantum dots. This insight is further extended by atomistic studies using molecular dynamics which allows for models composed of millions of atoms and hence describing e.g. mechanical response of multilayered thin films including finite-temperature effects and defect generation (plastic flow).

The group is also active in methods development. We explore methods to combine experimental high-temperature and *ab initio* low-temperature data for CALPHAD prediction of phase diagrams. Another activity is related to linking quantum-mechanical (QM) calculations with molecular mechanics (MM) into a so-called QM/MM framework which is envisioned to allow for large scale simulations (e.g., extended defects or complicated microstructures) of matrix materials with simple chemistry containing regions with complex chemistry or requiring insights beyond MM (e.g. information about electronic structure). We will demonstrate this latter approach in detail and present our own implementation relying on a combination of the VASP code (for the QM region) and the LAMMPS simulator (for MM region). The method has been applied to predict segregation of ternary alloying elements to phase (α_2/γ) and grain (γ/γ) boundaries in intermetallic TiAl alloys. An example simulation cell, as well as the spatial partitioning, is presented in Fig. 2. To benchmark our approach we first predicted the preferred sublattice for the ternary substitution, which agreed well with former full-DFT predictions. Next, we calculated alloying impact on the phase/grain boundaries in terms of their elastic modulus as well as cleavage strength. For the phase boundary scenario, we considered segregation to the phase boundary, phase segregation as well as a potential change of the preferred sublattice all at the same time. Putting all this information together allows discussing alloying impact on mechanical properties of multi-phase TiAl alloys.

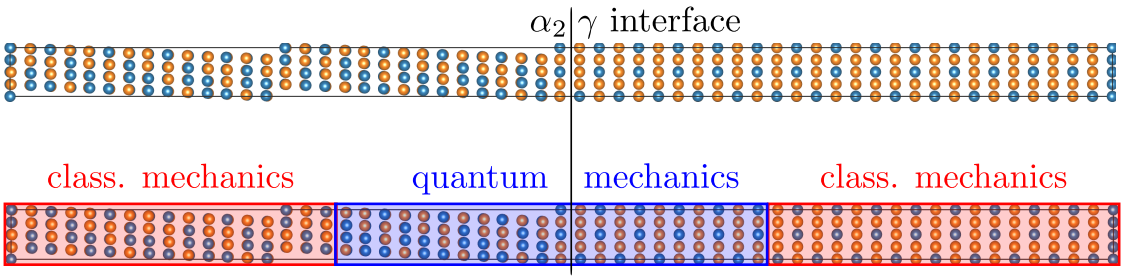

Figure 2: Spatial partitioning of a α_2/γ-TiAl simulation cell. In the QM/MM approach the region around the interface is treated with a quantum mechanical description, while outer bulk regions are treated with classical mechanics.

On the emergence of quasi-crystalline structures in a Wigner bilayer system

Benedikt Hartl[a], Marek Mihalkovič[b], Ladislav Šamaj[b], Martial Mazars[c], Emmanuel Trizac[c], and *Gerhard Kahl*[a]

[a] *Institute for Theoretical Physics, TU Wien, Vienna, Austria*
[b] *Institute of Physics, Slovak Academy of Sciences, Bratislava, Slovakia*
[c] *LPTMS, CNRS, Université Paris-Sud, Université Paris-Saclay, Orsay, France*

In a Wigner bilayer system classical point charges are confined between two parallel plates of opposite charge, maintaining overall charge neutrality; the system is characterized by the separation (η) and the ratio of the charge densities of the plates (A). Recent, detailed investigations [1] on the self-assembly scenarios of the charges on the two layers have revealed a rich plethora of zero-temperature ground state configurations. In well-defined regions of the parameter space, ordered structures (such as tilings with rectangle-triangle patterns on one plate and regular, discrete decorations on the other one or so-called snub-square type tilings) have been identified which can be considered as precursors of quasi-crystalline (QC) structures.

For this contribution we have focused on these regions and have explored the emergence of possibly stable QC structures. In particular we have increased the system size to up to 1142 particles per unit-cell via Stampfli inflation in our search for ground states in order to approximate QC tilings built up by equilateral triangles and squares. In an effort to perform the required structure optimization on these increasingly complex unit-cells, we implemented a replica exchange Monte-Carlo (REMC) method using "zipper" update-moves, which allows ergodic sampling of the tilings over a wide temperature range.

To compare newly discovered structures with the ground-state candidates known from the literature [1], we employ our REMC energy minimization procedure on an extended range of the system parameters A and η, comprising roughly 60.000 grid points for each unit-cell size. We utilize the *MPI4py*-framework [2] to distribute the independent Markov Chain Monte-Carlo "zipper"-workloads at different temperatures among several worker processes and apply the additional REMC tempering step synchronously in the main task: thus a single structural optimization procedure takes typically one hour on a single 48-core node of the VSC-4. Consequently we are left with substantial computational costs of (roughly) 487.296 CPU hours when optimizing the structures of 72 relevant (and computationally feasible) unit-cell sizes over the range in η.

In a region of the (A, η)-plane a novel ground-state phase featuring pseudo-dodecagonal clusters and super-clusters has been identified. Our investigations suggest that at small finite temperatures the system exhibits a random-rectangle–triangle phase with approximate dodecagonal symmetry, which provides evidence of the occurrence of a QC structure. The emergence of QC especially in soft-matter systems is usually related to two characteristic length-scales in the respective interparticle pair-potentials [3]. In contrary, the confined charges of our asymmetric Wigner bilayer system are subjected to a purely repulsive, featureless Coulomb pair-potential. To the best of our knowledge, the here studied Wigner bilayer system represents the first system based on long-range potentials where the formation of dodecagonal ordering has been observed.

References

[1] Antlanger, M., Kahl, G., Mazars, M., Šamaj, L., and Trizac, E., Phys. Rev. Lett. **117**, 118002 (2016) and J. Chem. Phys. **149**, 244904 (2018)

[2] Dalcín, L. D., Paz, R. R. and Stortiand, M., J. Parallel Distr. Com. **65** (2005)

[3] Dotera, T., Oshiro, T., and Ziherl, P., Nature **506**, 208 (2014)

Development of GPU accelerated molecular software (CmDock) for efficient high-throughput virtual screening

Davor Sluga[a], **Tine Erent**[a], **Marko Jukič**[b,d], **Črtomir Podlipnik**[c], **and Nejc Ilc**[a]

[a] *Faculty of Computer and information Science, University of Ljubljana, Slovenia*
[b] *Laboratory of Physical Chemistry and Chemical Thermodynamics, Faculty of Chemistry and Chemical Engineering, University of Maribor, Slovenia*
[c] *Faculty of Chemistry and Chemical Technology, University of Ljubljana, Slovenia*
[d] *Faculty of Mathematics, Natural Sciences and Information Technologies, University of Primorska, Slovenia*

In silico methods are becoming increasingly important in early drug discovery. Virtual screening to identify hits is the starting point for small molecule-based drug discovery and is central to both academia and the pharmaceutical industry. One of the most important steps in identifying new hits is screening compound libraries in silico (HTVS or virtual high throughput screening) as opposed to classical in vitro assays (HTS or high throughput screening) to obtain enriched compound libraries for further studies. In response to the limited chemical space in current chemical libraries, recent advances in compound library development and coupling with machine learning approaches have provided more efficient ways to sample chemical space and develop comprehensive chemical libraries. In addition, synthetic feasibility analysis has produced billions of sample libraries that are being evaluated for commercial availability, such as the Enamine REAL space [1]. The developments described above therefore place greater demands on downstream processing, and virtual screening software is needed to enable efficient evaluation of the extensive chemical libraries [2]. For this reason, we have started the development of CmDock, a molecular docking software based on rDock with the goal of optimising docking protocols, parallelization and, most importantly, the use of modern GPU hardware [3]. The latter is paramount in order to tackle the evermore complex chemical compounds.

Graphic processing units offer much more computing power than CPUs and can thus speed up the process considerably if efficient parallelization of the algorithms involved is achieved. Of note is also the fact that often, in modern HPC systems and personal computers, most of the computing power is concentrated in GPUs. Consequently, it is important to exploit this, in scientific software usually neglected, potential. Once the software (CmDock) matures it should enable researchers to tap into the vast amounts of computing power available in modern HPC systems and grid computing solutions, which rely on volunteer computing, such as BOINC.

References

[1] Hoffmann, T., Gastreich, M., The next level in chemical space navigation: going far beyond enumerable compound libraries, Drug discovery today, **24**, 1148 (2019)

[2] Ruiz-Carmona, S., Alvarez-Garcia, D., Foloppe, N., Garmendia-Doval, A. B., Juhos, S., Schmidtke, P., Barril, X., Hubbard, R. E., and Morley, S. D., rDock: a fast, versatile and open source program for docking ligands to proteins and nucleic acids. PLoS computational biology, PLoS computational biology **10**, e1003571 (2014)

[3] Jukič, M., Škrlj, B., Tomšič, G., Pleško, S., Podlipnik, Č. and Bren, U., Prioritisation of Compounds for 3CLpro Inhibitor Development on SARS-CoV-2 Variants, Molecules **26**, 3003 (2021)

High-Performance Computing with Relational Database Management Systems

Domen Verber

Faculty of Electrical Engineering and Computer Science, University of Maribor

Database management systems (DBMS) are optimized to process a vast number of transactions in a very short time. For this, DBMS can effectively utilize multiprocessor and multicore computer architectures, solid-state disks and large computer memories. If needed, the data transactions can be distributed to several computer nodes.

We will demonstrate how some data-intensive algorithms can be accelerated with modern relational DBMSs. This is especially useful if the data we are operating on is already contained in the database. The domain-specific Structured Query Language (SQL) is used to retrieve some information and manipulate the data in the relational database. The expressive power of SQL allows for a parallel operation on any number of inputs and can also produce (change) any number of outputs. On the other hand, the algorithms are usually presented as a list of instructions. For expressing the operation on several data elements simultaneously, loops are required. To resolve this, we introduced a translator that transforms instructions into a series of SQL statements that are compiled and optimized with DBMS during the execution. The translation to SQL code can be stored and reused until the basic algorithm changes.

The first practical problem for trying this approach was to shorten the payroll calculation for several larger companies in Slovenia [1]. Typically such calculations can be done in several minutes or even hours. For practical use, this was too slow. The payroll calculation can be expressed as a list of formulas that calculate specific payroll items. The formulas were prepared in advance. Each formula is (usually) a simple numerical expression that can reference previously calculated items. To cope with specifics of different employees, inline functions (like inline if) were used. Also, a Boolean predicate function can optionally be present, which determine if some formula is relevant for some employee or not. There is a one-to-one relation between payroll item definition and the equivalent SQL statements and can be stored in advance. Those statements are gathered and combined into one large SQL batch before execution. The method was implemented with MS SQL and Oracle databases, and it has been successfully used in practice for more than fifteen years.

The second problem we studied was quickly assessing if a student could enroll in a senior year. The conditions can be complex and differ between different study programmes. The conditions are described as a Boolean expression, translated into SQL condition expression and used for different purposes. The implementation of that is used in the Academic information system of the University of Maribor.

More recently, further studies of this approach were studied. In [2], the authors study how to translate general linear algebra structures and operations on relational DBMS. In our latest research, we are developing the translation of deep neural networks architectures into SQL structures to accelerate inference. We will present the result from the MNIST benchmark of recognizing handwritten digits with a simple deep learning convolutional network.

References

[1] Verber, D., Uporaba dinamično generirane programske kode za baze podatkov v poslovnih aplikacijah. Sodobne tehnologije in storitve: OTS 2010: zbornik petnajste konference, Maribor, 209 (2010)

[2] Shangyu, L., Zekai, J. G., Michael, G., Luis L. P., Dimitrije J. and Christopher J., Scalable Linear Algebra on a Relational Database System. Communications of the ACM **63**, 8 (2020)

Distributed computing for everyone

Giovanna Roda[a], Liana Akobian[b], and Dieter Kvasnicka[c]

[a] *TU Wien, TU.it, and EuroCC Austria, BOKU Wien; giovanna.roda@tuwien.ac.at*
[b] *TU Wien, TU.it; liana.akobian@tuwien.ac.at*
[c] *TU Wien, TU.it and VSC Research Center; dieter.kvasnicka@tuwien.ac.at*

Since the onset of the digital age, data volumes have been rising at a higher rate than processing power, making distributed computing an indispensable technology for data-intensive tasks. The Big Data ecosystem Hadoop with its main components MapReduce and Spark is one among the many available frameworks and libraries for data-intensive computing such as Dask for Python, Ray for Java/Python, distributed R, etc.

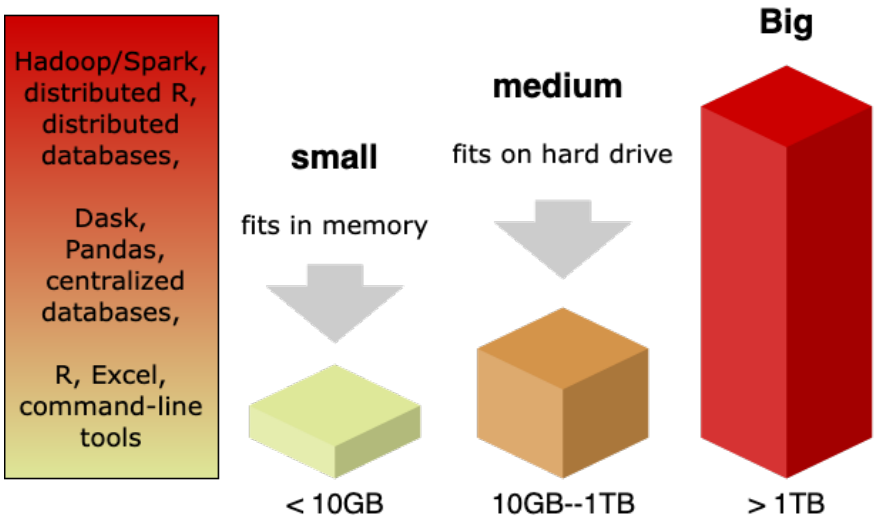

Big Data and the need for distributed computing.

By means of practical examples, we are going to illustrate how even a newcomer can leverage Hadoop by identifying embarrassingly parallel features in a problem's workload [1]. While data-intensive computations might require a high-performance cluster, with today's multi-core processors parallelism can be achieved at the cores' level on a single compute node.

R A straightforward way to increase code performance for an R user is to substitute the `lapply` function with `spark.lapply` to distribute computations across multiple nodes.

Python/Pandas Python users that are accustomed to Pandas can adopt *Koalas*, a library for writing Pandas code with a Spark backend.

SQL SQL developers can run SQL queries on distributed datasets with the Spark SQL API, using the Spark engine optimizations at the CPU and main memory level.

References

[1] P. H. Gelado and R. K. Pradhan, Convergence of HPC and Big Data: `https://youtu.be/EP1bWMjOGHA` (2021)

Keynote Talk:

HPC in the Exascale era and beyond

Jean-Pierre Panziera

Atos; jean-pierre.panziera@atos.net

Exascale now!
During the last decades, High-performance computing (HPC) simulations performance have steadily improved thanks to the introduction of new technologies and system architectures. The HPC users successively migrated from mainframes to vector supercomputers, to RISC parallel SMPs, then large clusters aggregating generic microprocessor servers, and now the first Exascale systems being deployed rely heavily on GPU accelerators. But, even though GPUs have become easier to program, porting and optimizing full scale HPC applications represent a sizeable effort.

The next challenge comes from the Exascale systems size. A full warehouse is now necessary to host such systems, and their electrical consumption amounts to tens of MegaWatts. This energy is all converted into heat and liquid cooling has become the only viable solution. Besides adapting to new hardware, HPC applications have evolved to integrate new techniques. (Big) Data Analysis helps refine computing models with the integration of real-life measurements. Recently, Artificial Intelligence (AI) started finding its way into HPC, it provides a software acceleration for example with the use of surrogate models defined with machine learning frameworks.

Beyond Exascale
To reach the next level of performance and scale up to Zettascale within a decade(?), current technologies must drastically evolve. The use of multiple specialized hardware components (CPUs, GPUs, FPGAs, IPUs, TPUs, neuromorphic, QPUs …) will boost performance within a reasonable power envelop. But to deliver the promised performance a tight hardware and software integration will be necessary, and AI will be key to facilitate the development and optimization of future HPC applications on such heterogeneous systems. Finally, for some specific workloads, Quantum Computing will open new possibilities, with the first prototypes coupling HPC and Quantum being expected to be deployed in 2023.

Index of contributors

DOI: https://doi.org/10.25365/phaidra.337

ISBN: 978-3-200-08499-5

Published by:
EuroCC Austria
c/o Universität Wien
Universitätsring 1
1010 Vienna, Austria
https://eurocc-austria.at/

Edited by:
Eduard Reiter, Research Area Scientific Computing, University of Innsbruck, Austria

Layout:
Irene Reichl and Claudia Blaas-Schenner, VSC Research Center, TU Wien, 2016