Newsletter 02/2024

Welcome to the newsletter of the dtec.bw project hpc.bw. If you want to subscribe to the newsletter, please send a message with subject line “Subscription hpc.bw Newsletter” to info-hpc-bw@hsu-hh.de.

Contents

Ready for the 2nd phase of dtec.bw

In the scope of the extension of the Digitalization and Technology Research Center of the Bundeswehr (dtec.bw), hpc.bw successfully applied for a project extension, which will run from 01.01.2025 to 31.12.2026. Besides continuing current work and research strands in the project, another focus will be put on the interaction with relevant groups from the federal armed forces and on the development of innovative multi-physics and multiscale simulations, for example with applications in fluid dynamics. In particular, hpc.bw merges for the latter purpose with the dtec.bw project “MaST: Macro/micro-Simulation of Phase Decomposition for Transcritical Fluids”, which becomes one particular strand of research of hpc.bw in the extension. Project work is to be complemented by activities on research software engineering, energy efficiency and internationalization.

Visit from the European Commission, 24.04.2024

Due to an audit of research projects funded by the European Union, hpc.bw was visited by delegates of the European Commission on the 24.04.2024. During the visit, Piet Jarmatz presented the scope, aims and structure of the hpc.bw project. He offered a tour of the CBRZ showing the operational hardware of the HPC cluster HSUper, and gave a short live demonstration showing how the supercomputer can be used and what its applications can be. The visitors from the Commission were fascinated, and were delighted to see our systems and to get some insight about a world that is otherwise often hidden in the background.

HPC Portal: First glance of the courses

The HPC Portal is finally and officially online and was presented in detail in the last newsletter(s). Furthermore, we went through the various webpages on-site with the participants of the second HPC Workshop, received reviews and feedback and adjusted the Portal based on internal evaluations. Long story short, the structure is set, and it is time to start with the planned core content of the HPC Portal: the web-based training modules. Have a look at the upcoming training offers and the structure of web-based training, which will be distributed via the HPC Portal.

Figure 1: Logo, HPC Portal

Web-based training or course?

First of all, let's have a look at the terms used in the HPC Portal and what the offers can be and what not.
The HPC Portal strives to support the acquisition of competences and skills related to High Performance Computing. For that, we want to create and distribute various formats of media in different formats such as cheat-sheets, videos, blogs, documentations or so-called self-learning-arrangements. The self-learning-arrangements are based on educational theory presentations with written and media-supported content. We don't wish to call them "courses", because the term comes with expectations similar to on-site educational events, such as those found in universities or adult education centers. It also comes with the expectation of having some kind of result-management (audits or tests) and a certificate if the course is completed successfully. We do not want to conform to these expectations with the self-learn-arrangements on the HPC Portal, and thus we chose the term "web-based training" instead.

Content of the first web-based trainings

As promised, the first prepared content in form of web-based training will be the topics covered by the HPC workshops.

1. A short introduction to HSUper and Linux Terminal
2. Data transfer and system resources
3. Module systems and Slurm jobs

All web-based trainings will have the same structural design to maintain some form of standardization throughout all topics. All trainings will start with a short introduction, legal information and learning goals for that specific subject. The learning goals are based on *Bloom's Taxonomy* so that every content creator can use terms referring to the same approach, which also makes the training coherent for learners.
After this preamble, every training module contains a page with information about prerequisites. In general, the training modules are not designed to be completed consecutively, but some training modules need prior knowledge in specific topics for successful completion. The actual content then starts on the next page.
Every training module also contains some kind of *tracking mechanism* shown as a small line on every page. It tracks content which has already been covered. Learners can use these to track their own progress and pause the course if necessary.

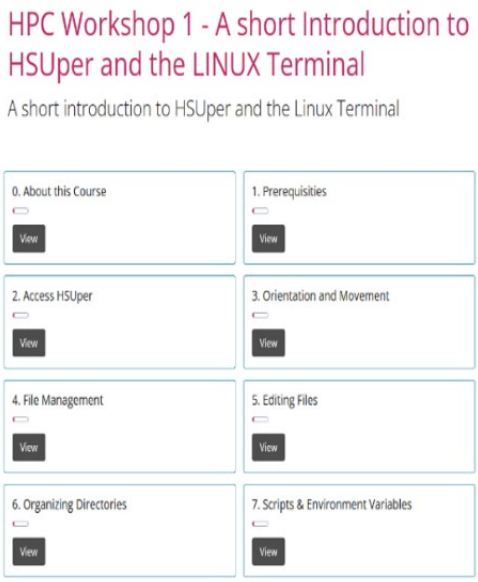

Figure 2: HPC Workshop 1, HPC Portal

Invitation: CBRZ User Meeting @17.10.2024

The next CBRZ user meeting will be held on Thursday, 17.10.2024 at 09:00 in seminar room 208 in building H11. The user meeting is open to every HSUper and ISCC user to share experiences. Feel free to join us! Please send an e-mail to info-cbrz@hsu-hh.de in case of any questions regarding the meeting.

hpc.bw @ISC High Performance 2024

hpc.bw presented a poster in the poster gallery of the ISC High Performance 2024 from Monday to Wednesday (13.05.-15.05.2024). During the poster session, many people were interested in the project results due to their experiences with similar problems. Contact details were exchanged with people working on projects who also intend to create HPC learning materials.

©Sylvain Maret

The project poster is available for download here:

hpc.bw @Day of Armed Forces, 08.06.2024

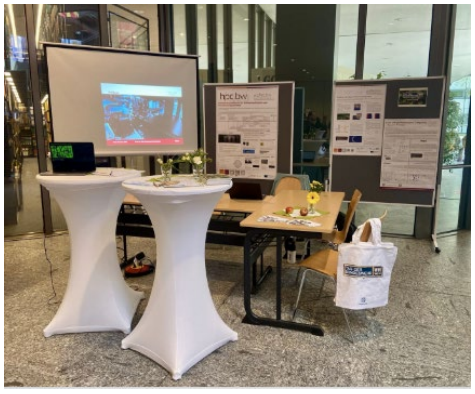

©Imane Bechelaoui

On the Day of Armed Forces, hpc.bw was represented with its own booth next to various exhibitors from both research areas and the Armed Forces. Interested visitors had the opportunity to find out more about our project - either in direct dialogue with our team or through the illustrative posters.

As a special highlight, three tours of HSUper were offered where visitors were able to gain an insight into HSUper from the inside and ask questions about the hardware. Visitors of all ages also had the opportunity to try out a game that we had programmed ourselves. This interactive activity was great fun for both young and old, and offered a playful approach for our research topics.

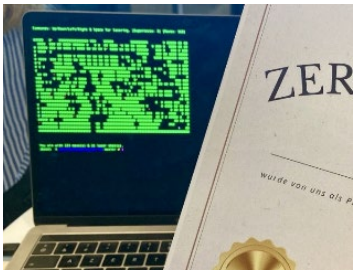

©Imane Bechelaoui

Here you can see our represented poster:

European Trilinos & Kokkos User Group Meeting (EuroTUG) @HSU, 24.06.-26.06.2024

©Marie Rathmann

The EuroTUG Meeting, hosted by hpc.bw, took place at HSU with a total of 23 participants from 24.06.-26.06.2024. The unique aspect of this year's EuroTUG was the merger of Trilinos and Kokkos into a single meeting.

On the first day, Jakob Bludau (TU Munich) gave an introduction to Kokkos, a Performance Portability Programming Ecosystem. The day ended with a talk by Chris Siefert (Sandia National Laboratories, online) on "From Epetra to Tpetra", to illustrate the migration of the code with Trilinos. The second day started with an introduction to Trilinos by Alexander Heinlein (TU Delft). Trilinos' package architectures enable the simultaneous development of many new functions in a federated system. This was followed by a hybrid (on-site/online) hackathon with Matthias Mayr (UniBw M), Alexander Heinlein, Chris Siefert, and other staff members from Sandia National Laboratories. The evening concluded with a joint dinner at Binnenalster in the heart of Hamburg. On the last day, Philipp Neumann (DESY/ Universität Hamburg) presented research on HPC at HSU and Piet Jarmatz (HSU) presented hardware and software information on HSUper. Afterwards, Hauke Preuss (HSU) led a guided tour to HSUper.
The EuroTUG Meeting was completed by various contributed talks.

A Roof for the CBRZ Building

Construction of the roof for the building which houses the CBRZ began in May. As a result, some cooling infrastructure was moved and reinstalled between May 13 and May 22. During this time, the CBRZ was offline, including HSUper and the ISCC. The roof construction should be completed before winter. The cooling systems will be moved to their final locations, requiring the CBRZ to be taken offline again for several days. This will be announced as soon as we have reliable time tables.

Project Update: Logistics and Optimization

One of the hpc.bw subprojects, namely "Logistics and Optimization", is dedicated to the investigation of mixed-integer mathematical optimization solvers available in the market, and how they can exploit parallel computing power off the shelf. In the plot, a comparison of the performed floating point operations per second (FLOPS) for the software Gurobi with default parameters for the benchmark ta064 is shown. The default parameters ensure that the number of threads is determined automatically (with a soft limit of 32), and ta064 refers to a special scenario of the time-dependent travelling salesman problem.

For the experiment, one node of HSUper (2x Intel Xeon Platinum 8360Y) is always used exclusively, and we limit the number of cores available to Gurobi via SLURM. The colors represent 72 cores (i.e., one node, 1533 seconds), 36 cores (i.e., one socket, 1786 seconds), 32 cores (i.e., amount of available cores which equals the soft limit for Gurobi threads, 1744 seconds), 18 cores (i.e., one NUMA domain and hyperthreading necessary if more than 18 threads are used, 2383 seconds), 16 cores (i.e., cores are overbooked with two threads per core, 2417 seconds), 8 cores (i.e., cores are overbooked with four threads per core, 2690 seconds). The

Gurobi solver is set up deterministically, i.e., the corresponding amount of reported work is identical (and the sum over all FLOPS is almost identical due to background noise such as from the operating system) for all traces. This shows that decreasing the number of cores often increases the runtime.
However, the runtimes for 36 and 32 cores, as well as for 18 and 16 cores, are almost identical whereas the increases in runtime are especially large if only one socket instead of two is used (approximately 200 seconds or 13 to 17 percent) and when hyperthreading (HT, i.e., less than 32 physical cores are used) is necessary (approximately 600 seconds or 33 to 37 percent). The reason for this can be that each socket has its own L3 cache and each physical core has its own L2 and L1 caches.

Hence, the available cache sizes and bandwidths are decreased by decreasing the number of cores/sockets used. Moreover, we repeated the computations and observed that the runs with HT have a significantly lower difference in the runtime (here up to approximately 11 seconds) than the runs without HT (here up to approximately 90 seconds) although the runtimes of HT runs are larger than the runtimes of runs without HT. These tests indicate that on hardware with many cores (and several sockets), it seems to be advisable to run multiple Gurobi problem-solving processes simultaneously instead of sequentially, since the increases in runtime for using one socket instead of two, as well as for overbooking cores, are modest.

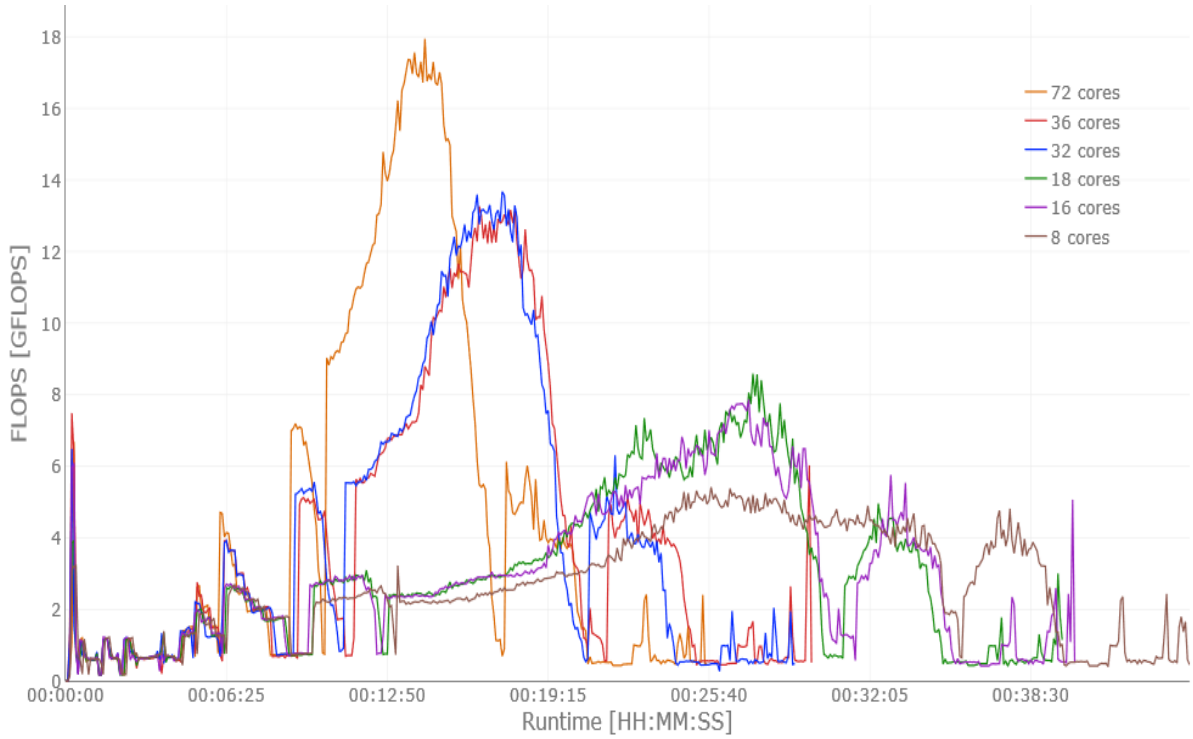

Figure 3: Floating point operations per second for solving ta064 with Gurobi. Amount of available cores controlled via SLURM.

Project Update: Enabling High-Throughput Studies of Reactive Materials

We completed the work aimed at implementing, benchmarking, and validating a high-throughput framework for the study of material properties at the atomic and molecular level. The VASP and LAMMPS simulation codes were wrapped in a python-based front-end for HSUper's SLURM scheduler, consisting of atomate, fireworks, pymatgen, and custodian. A validation was performed using the provided test cases, as well as a scan and a benchmark of parallelization-related parameters. Runtime penalties for unfavourable parameter combinations were identified (see Figure 4) and can be avoided in the future.

The resulting simulation framework enables the Chair of Computational Material Design to systematically explore an enormous combinatorial space of materials for their mechano-chemical properties. The benchmark results provide useful guidelines for parallelization settings in order to optimize the use of CPU

resources. The created templates on software usage of the HSUper scheduler and related advice served as an example for the SLURM segment of this year's HSUper workshop.

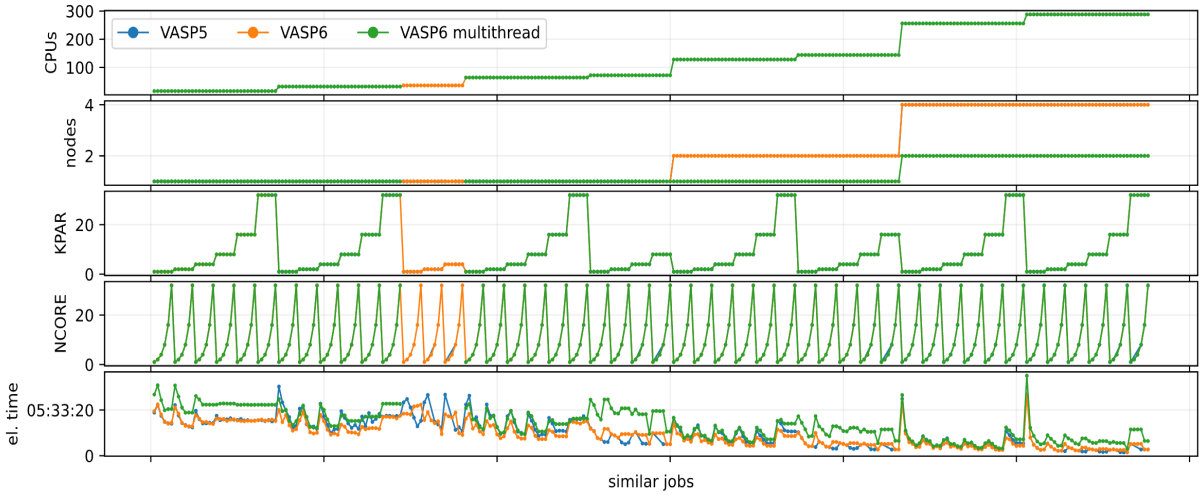

Figure 4: Data points from the run time scans over parameters of interest. The measurements suggest, amongst others, to avoid hyperthreading, and to use generally higher parallelization.