Newsletter 03/2023

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

Seminar Series Computation & Data in HT23

In the autumn trimester (HT 2023), we warmly invite you to the upcoming talks in our seminar series Computation & Data at the HSU and look forward to exciting discussions on the topic! In addition to attending in presence at the HSU, it is also possible to participate digitally via MS Teams.

Date	Room	Speaker	Title of lecture
25.10.2023; 16:00-18:00	Container-building C2/S2, room 113-115	Maria Maleshkova	Digital Twins – Building Blocks for Medical Data Scienc
		Thomas Carraro	The HiBRAIN project: Combined data- and model-based design of lithium-ion battery electrodes using artificial intelligence
29.11.2023; 16:00-18:00	Container-building C2/S2, room 113-115	Therese Rosemann & Yannis Schumann	Digital Competences and Digital Learning Behavior in Higher Education - Generation of individualized Feedback in a Longitudinal Study (DigiTaKS* and hpc.bw)
		Sergej Grednev	Prediction of Structure-Property Relationships in Cellular Materials
13.12.2023; 16:00-18:00	Container-building C2/S2, room 113-115	Maximilian Maigler	Coupled PIC-DSMC and Molecular Dynamics Modeling of Radio Frequency Gridded Ion Thruster Erosion
		Valentina Pessina	Modeling Rarefied and Continuous High Angle-of-Attack Hypersonic Reentry into Martian Atmosphere with Open-Source Software

To subscribe to the seminar mailing list, send an e-mail to info-hpc-bw@hsu-hh.de, subject line „Subscription Seminar Computation & Data". For more information, click here.

hpc.bw @ PinT, 17.07.–21.07.2023

The 12th Workshop on Parallel-in-time Integration, PinT2023, took place on 17–21 July 2023 at Hamburg University of Technology (TUHH). The conference was jointly organized by HSU and TUHH.

Parallel-in-Time (PinT) methods are a highly scalable way to leverage the computational power of distributed systems and achieve a speed-up in transient numerical simulations compared to conventional serial time stepping methods. The recent developments in exascale computing require a high degree of parallelization, so that advanced approaches like PinT are needed to allow scientists and engineers to run larger simulations

©David Ernst

of physical phenomena. Thus, research of PinT mathematical methods and algorithms is an important precursor of future HPC applications. The hpc.bw project was present with a project poster presentation on the conference as well as with a research poster by Piet Jarmatz, demonstrating PinT molecular-continuum flow simulation results obtained on our HPC cluster HSUper.
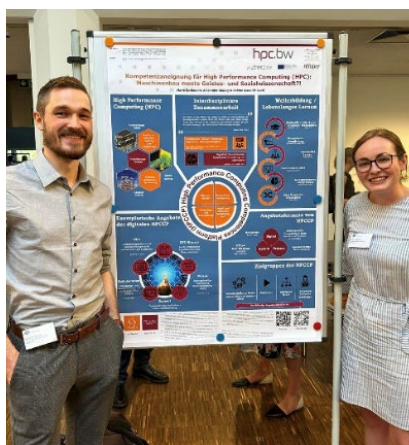
Here you can find the related article on our website and the presented poster:

hpc.bw @ Conference of Division for Adult Education (DGfE), 11.09.–13.09.2023

On 11–13 September 2023 the DGfE-Congress of the Division for Adult Education took place at the Ludwig Maximilian University of Munich. The event focused on "Sustainability" in further education and additionally invited to present scientific projects as a part of a poster presentation. The hpc.bw project was presented during that poster slam by Alexander Kolling and Marie Rathmann, focusing on the conceptual details of the learning environment provided by the High Performance Computing Competence Platform (HPCCP). Numerous interested participants of the conference approached the poster to learn more about hpc.bw, with a particular focus on the HPC competence platform.

Here you can see the associated article and the poster:

©Henrik Weitzel

hpc.bw @ NHR Conference, 18.09.–19.09.2023

On 18–19 September 2023, the first conference of the national high performance computing initiative (NHR) took place at the Zuse Institute Berlin. The conference focused this year on the three application areas atomistic simulation, life science and agent-based simulation. It brought together researchers to discuss current methods and trends that are under consideration at the NHR sites in Germany. From HSU, Philipp Neumann contributed a talk on efficient three-body calculations in molecular process engineering, including a short description of hpc.bw activities and the CBRZ infrastructure.

Find out more about the conference here:

hpc.bw @ Digital Total, 09.10.-10.10.2023

The recent event "Digital Total" at the University of Hamburg showcased the metropolitan region of Hamburg as a coming hub for digital innovation. Organized by the University of Hamburg's House of Computing & Data Science (HCDS) in collaboration with the scientific platform PIER PLUS and the Akademie der Wissenschaften, the Academy of Sciences, in Hamburg, the event underscored the significance of Data Science, AI, and Computation in the region.

The event, co-organized by the CDCS, emphasized the strong representation of these themes in the Hamburg metropolitan region with an in-depth discussion on "PIER Computing and Data Science." Additionally, the Akademie der Wissenschaften Hamburg hosted a cross-disciplinary forum on the future of digital transformation in the region, encouraging discussions on potential synergies across institutions. On day two, the focus shifted to the role of Data Science, AI, and Computation at the University of Hamburg. This day also marked the introduction of the Cross-Disciplinary Labs (CDLs) located within the HCDS.
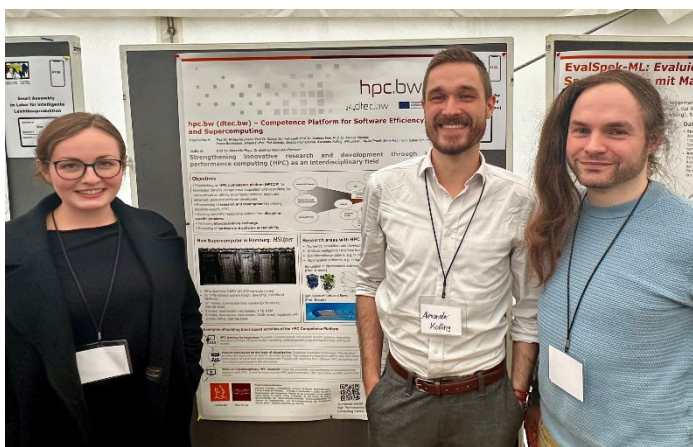
©Yannis Schumann

HSU contributed to Digital Total in several ways. Prof. Oliver Niggemann gave an overview talk on digitization at HSU, including HPC and dtec.bw activities. Several of the latter projects were also presented in the poster session of the event, which featured a total of almost 150 posters, including a poster of hpc.bw on the "Competences platform for supercomputers and efficiency." Feel free to check out the Book of Abstracts for more details on the sessions and contributions.

You can find the related article on our website as well as the poster here:

Interactive Scientific Computing Cloud (ISCC) Status Update

The Interactive Scientific Computing Cloud (ISCC) aims to fill the gap for users with specific needs that cannot be containerized, as well as users with workloads that cannot be satisfied by existing local machines and do not yet require the entire power of HSUper. The ISCC consists of 12 hosts (2 hosts with GPUs) for virtual machines. Resources can be divided as small as a single core, 4MB of RAM, and a quarter of a single NVidia A30 GPU. For example, this allows a GPU host to be sliced into 32 equal virtual machines using a quarter of a NVidia A30, 2 CPU cores and 32GB of RAM. Virtual machines can be sized according to the use case. All major operating systems are supported. Many Windows and Linux images are already available for immediate use. We welcome requests to add more OS images.

The alpha test phase has been a success. Currently, different use cases are being tested by different groups. One of the larger examples is the use of ISCC resources for special applications in lectures.

Please email us if you are interested in testing the new platform with us: info-hpc-bw@hsu-hh.de or visit the ISCC system documentation on our website for further information and a link to the ISCC resource request form:

HSUper Workshop Organized by hpc.bw, 28.09.2023

On 28.09.2023 HSU hosted a dynamic workshop centered on both the theoretical aspects and practical applications of HPC, with a specific emphasis on the use of the supercomputer HSUper. This event, presented by the hpc.bw-team, was a mixture of foundational and advanced topics tailored for both newcomers and experienced HPC enthusiasts.

©Alexander Kolling

The participants were introduced to a variety of subjects, ranging from the essentials of HPC, Linux, and data transfer to the complexities of software installation and GPU computing. The BeeGFS parallel file system was also a notable highlight, offering insights into next-generation data storage and access. To ensure a comprehensive understanding, practical demonstrations and manual tasks accompanied the theory, covering areas such as node allocation on HSUper and script writing.

Beyond the educational segments, attendees were also given the opportunity to visit the CBRZ with the HPC cluster HSUper and to interact during coffee breaks. In addition, we used this opportunity to gain more insight into the needs and professions of the participants through questionnaires. In our first evaluation, we were able to identify specific needs, which open up new possibilities for us to provide additional materials and offerings. Fortunately, with ca. 40 participants, the overall response was so great that we hope to be able to offer more workshops or events soon.

Scan the QR-Code or click here to see the related post on our website with the linked schedule of the workshop:


Infrastructure Data Analysis Supported by hpc.bw

In cooperation with the Chair of Engineering Materials and Building Preservation at the HSU, hpc.bw has successfully completed a Python code optimization for a software tool for infrastructure operators. The primary goal was to improve the software design of the data analysis script, to make it more maintainable. Besides, the possibility to export and import already trained models was introduced, so that they can be reused and inefficient repeated input processing and training can be avoided.

The example output plots from the data analysis tool show how the condition of bridge components relates to their age, how they can be clustered or categorized, and a SHAP analysis of a Random Forest model. This shows the impact of model inputs, i.e. properties of the bridge component, to the model output, i.e. the probability of failure of the respective component.

The pictures show plots from the data analysis script that resulted from this cooperation.
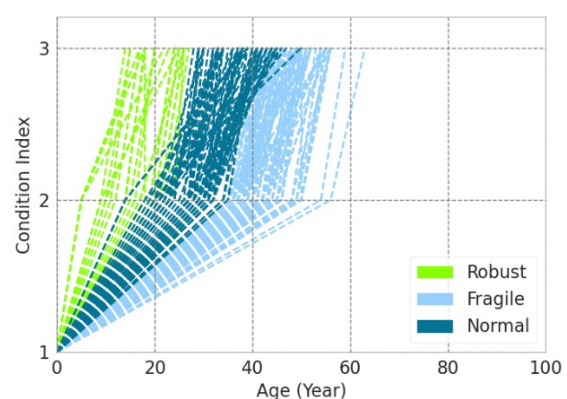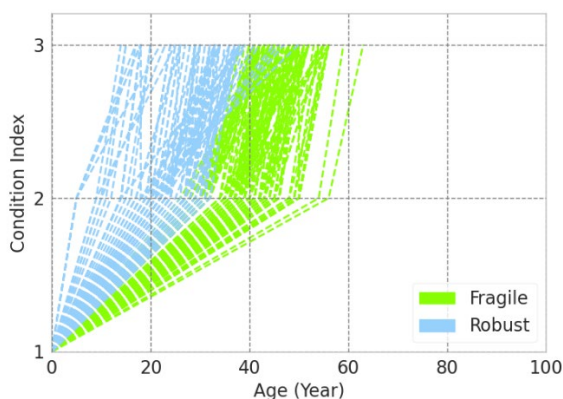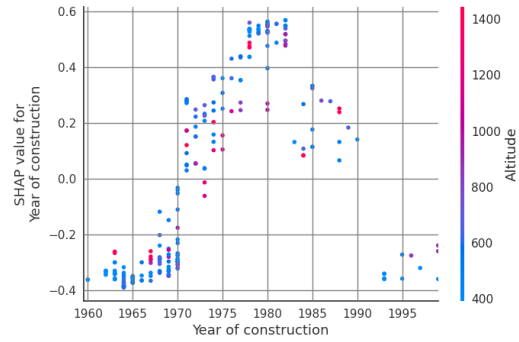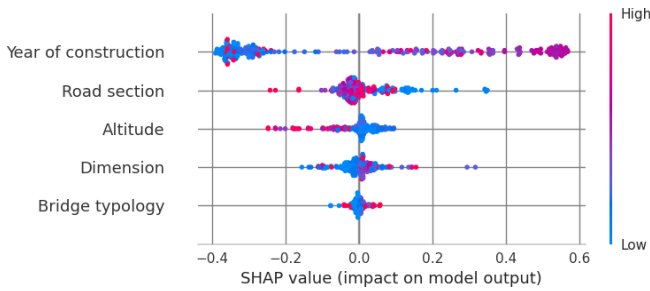
Project Update: Logistics and Optimization

In hpc.bw, one sub-project is dedicated to investigating how available mixed-integer mathematical optimization solvers are already able to exploit the parallel computing power of the shelf. Recently, we experimented with the software Gurobi on the ISCC (Interactive Scientific Computing Cloud).

In the figure, a performance experiment of the benchmark ta064, a time-dependent traveling salesman problem, is shown. This illustrates that the ISCC with HT (hyperthreading; i.e., up to 128 cores) and without HT (i.e., up to 64 cores) leads to almost identical computation times. However, HSUper (which runs in a different configuration such that up to 144 cores can be used in exchange for slightly smaller base and turbo frequencies) has for up to 32 threads higher computation times than ISCC (probably due to its smaller frequencies), whereas HSUper leads for 64 threads to smaller computation times than the ISCC (probably due to the overhead of its virtualization). Each marker represents the arithmetic mean of three computations. The differences in computation time between the markers are usually significantly larger than the differences between different runs of the same setup (e.g., the fastest setup with 16 threads has a difference between runs of the same setup of less than 2.6s (HSUper), 21.4s (ISCC_HT) and 3.3s (ISCC_noHT), whereas the slowest run on the ISCC is more than 6.8s (850s for the 1 thread case) faster than the fastest run on HSUper).
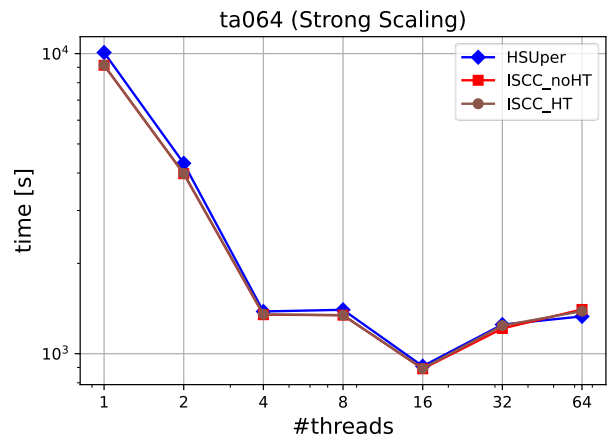
Update: HPC-Portal

We're pleased to announce that we're currently working on the debut release of the "HPCCP (High Performance Computing Competence Platform)", a website for self-learning, addressing HPC researchers and users. For that, we're navigating the challenges of data protection and IT security in cooperation with the HSU and adapting the content for dual use.

For those who recall the workshop on 28.09.2023, we're customizing its content to make it freely available on HPCCP. Plus, alongside the launch, we're introducing an online glossary to aid your learning journey.

Stay tuned for more updates!