hpc.bw: A Supercomputer with Competence Platform for the Universities of the Federal Armed Forces

Philipp Neumann*, Johann Duffek, Jessica Kleinschmidt, Willi Leinen, Michael Breuer, Sabine Schmidt-Lauff, Andreas Fink
Helmut Schmidt Universität/ Universität der Bundeswehr Hamburg
Hamburg, Gemany
*philipp.neumann@hsu-hh.de

Matthias Mayr, Max Firmbach, Alexander Popp
Universität der Bundeswehr München
Munich, Germany

Axel Auweter
MEGWARE Computer Vertrieb und Service GmbH
Chemnitz, Germany

***Abstract* – Computational and data science have evolved as major pillars in basically all research disciplines. When it comes to efficiently solving compute- or data-intensive tasks such as numerical simulations, optimization problems or machine learning, researchers are, however, often confronted with several challenges which result in a sub-optimal research productivity. In our project hpc.bw, we strive to increase research productivity by providing high performance computing (HPC) systems, skills and support. In a highly interdisciplinary approach, we develop for this purpose a HPC competence platform at the Universities of the Federal Armed Forces. In this contribution, we detail our anticipated approach to disseminate and to transfer HPC knowledge across all disciplines. This includes a multiplicator program, performance engineering activities to boost discipline-specific software performance, hardware and software sustainability considerations and the development of a building block-based concept for HPC competences. All these activities integrate with the establishment of a container-based high performance computing center at Helmut-Schmidt-Universität and its supercomputer HSUper.**

Keyword* – *High performance computing, supercomputer, competence platform, performance engineering, education

Nomenclature

CBRZ	Container-based high performance computing center[1]
HPC	High performance computing
HPCCP	HPC competence platform
HSU	Helmut-Schmidt-Universität/ Universität der Bundeswehr Hamburg
MPI	Message Passing Interface
OER	Open Educational Resources
UniBw M	Universität der Bundeswehr München

I. Motivation

Computational and data science have evolved as major pillars in basically all research disciplines, complementing analytical and empirical or experimental studies. However, researchers are easily confronted with severe challenges when it comes to relying on computational methodologies and technology, especially when these turn out to be very compute- or data-intensive: computational resources might be lacking; if these resources are available, e.g. in the form of a supercomputer, it is not straight-forward to use them; programming skills that are necessary to optimize and parallelize programs to be used on such high performance computing (HPC) systems need to be acquired first, and this is often not included in discipline-specific under- or postgraduate programs; time to apply these skills to the actual research task and software might be scarce; etc. All these hurdles, in conjunction with the fact that HPC resources become more and more complex (heterogeneous architectures, often comprised of hybrid systems with CPUs and accelerators such as GPUs; growing memory hierarchies), induce a sub-optimal research productivity.

Goal and motivation of our dtec.bw-funded project *hpc.bw: Competence Platform for Software Efficiency and Supercomputing* is to decrease this gap. The aim is to strengthen innovative cross-site HPC research at the Universities of the Federal Armed Forces and to accelerate the transfer of HPC knowledge to a wide range of disciplines to

- sustainably strengthen research and development in the respective disciplines,
- promote interdisciplinary exchange between HPC-related problems,
- derive new research questions for HPC from discipline-specific problems and answer them, and
- establish a joint HPC competence platform for users within and outside the federal armed forces.

These activities are supported by the establishment of a container-based high performance computing center. An overview of project activities is provided in FIGURE 1.

In this contribution, we give an overview of project activities. First, we focus on our efforts to build an effective

[1] German: **C**ontainer-**b**asiertes HPC **R**echen**z**entrum

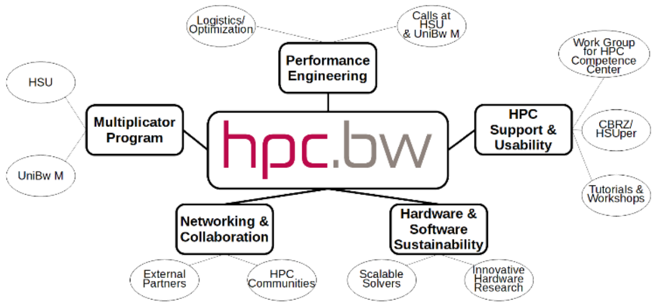

FIGURE 1: HPC.BW PROJECT OVERVIEW.

HPC community in Sec. II. Our approach to develop a HPC competence platform is explained in Sec. III. Hardware-aware HPC sustainability efforts with a focus on the establishment of a container-based high performance computing center including a PetaFLOP-capable supercomputer are detailed in Sec. IV. Research and support actions for HPC software are laid out in Sec. V. We close with an outlook to *hpc.bw* activities beyond the Universities of the Federal Armed Forces in Sec. VI and a conclusion in Sec. VII.

II. BUILDING A HPC COMMUNITY

Various measures have been taken to reach out to relevant user groups and to increase awareness for the project *hpc.bw* and its research and support actions, which shall in the long term feed into the formation of an active, highly interdisciplinary HPC community.

A. Multiplicator Program

An essential point from the very beginning of the project was the involvement of all potential HPC user groups. The term "user group" is interpreted in the widest sense: users may come from any discipline (social sciences, natural sciences, economics, engineering, …), they might already have HPC experiences at various levels, or they might be beginners, recognizing the potential that high performance computing may provide to their research field. To reach out to these groups in a top-down approach, a multiplicator program has been launched. One multiplicator of every involved organizational unit is a member of the project consortium. Due to current HPC demands, UniBw M and every department at HSU act as one organizational unit each. The multiplicators spread news on the project in the relevant university boards (e.g., information on *hpc.bw* newsletters) and are further involved in decision-making processes such as the evaluation of performance engineering project proposals, cf. Sec. V.A.

B. Newsletters

Complementing the multiplicator program, a newsletter series has been launched in the beginning of 2022 to disseminate *hpc.bw*-related activities to a wider audience [1], [2]. The quarterly newsletters are made available on the current project website[2] and on the portal OpenHSU[3].

C. Meeting the Community: Seminar Computation & Data

Besides, *hpc.bw* actively contributed to bringing together researchers that work in the fields of computational and data science in the widest sense. A monthly seminar series "Computation & Data" was launched in March 2022, where in each session 1-2 researchers from the Universities of the Federal Armed Forces, but also invited speakers and collaboration partners, present their work.

III. TOWARDS A HPC COMPETENCE PLATFORM

A specific aim in the overall *hpc.bw* project is to initiate, expand and establish the transfer of HPC knowledge via a joint HPC competence platform (HPCCP) for users within and outside the federal armed forces. The HPCCP represents an open concept with newly created and systematically organized learning resources for the use of HPC systems and various target groups, cf. Sec. II.A. The implementation of the adult education activities can be digital, on-site or hybrid. The learning content will be open access as OER where possible and appropriate.

By providing learning opportunities and adult education concepts, an interdisciplinary contribution will be realized, leading to the development of new areas of knowledge and research in HPC. The following examples are working areas of the HPCCP:

- Learning and networking opportunities for different user groups (e.g., seminar series, workshops, tool talks, poster presentations).
- Teaching/learning material for different target groups (e.g., wikis, online tutorials, tool talk)
- Illustrative material (also digitally prepared) for low-threshold, non-specialist access to various service areas of HPC (e.g., basic knowledge, application fields, implementation examples, benefit analyses)
- Development of digital OER
- Establishing a web platform for the provision of developed materials (e.g., reports, online tools, workshop concepts) for HPC competences
- Continuous improvement through piloting (e.g., feedback, evaluations, needs assessments)
- Continuity through thematic updates, adaptation to demand and collaborative further developments

IV. HSUPER AND HARDWARE SUSTAINABILITY

A. Container-based High Performance Computing Center and HPC cluster HSUper

At the core of high performance computing lies the use of supercomputers. The "from-scratch" design, provision and commissioning of these computational resources requires various considerations. At its heart, the design of the supercomputer must enable optimal performance for the

[2] https://www.hsu-hh.de/wb/hpc-bw

[3] https://openhsu.ub.hsu-hh.de/

FIGURE 2: CBRZ (TOP; FOUR CONTAINERS TO THE LOWER RIGHT, FRONT SIDE OF THE CONTAINER-BASED BUILDING) WITH HPC CLUSTER HSUPER (BOTTOM). PHOTOS COURTESY BY U. SCHRÖDER, HSU.

anticipated compute loads, which are, in this case, numerical simulations, but also data analysis and machine learning tasks. The massive power consumption of HPC systems results in the dissipation of energy in form of heat. Using this dissipated heat is highly desirable, so that the system is operated in an environment-friendly way ("green computing"). Infrastructural prerequisites need to be satisfied: besides aforementioned massive power supply, enough space for the compute system and its infrastructure (e.g., cooling system) must be available, which is a challenge for smaller universities and holds in particular for HSU, given the current, enormous build-ups on campus. Finally, a heterogeneous user base with foci on, amongst others, computational fluid dynamics simulations, material science investigations and machine learning, requests different kinds of compute systems.

Under these constraints, the concept of a container-based HPC center (CBRZ) including supercomputer HSUper was developed, cf. FIGURE 2. It was subsequently realized by MEGWARE Computer Vertrieb und Service GmbH. Four 20 foot containers, integrated in a container-assembled building, host the supercomputer, which currently consists of nine compute racks. To allow an as-compact-as-possible hosting of processors in the racks and to improve the energy-efficiency of the system, a liquid-cooling approach has been followed for most components of HSUper. The cooling system has been designed for the highest possible density and efficiency. Depending on the operational mode and the season, this allows to either emit the generated heat via roof-mounted fans or to reuse the heat for the heating system of the container building in the future.

The supercomputer HSUper consists of three partitions of compute nodes. The biggest partition consists of 571 dual socket Intel Icelake compute nodes, each node equipped with 256 GB RAM and 2x36 compute cores. This partition allows the execution of classical x86-based software (such as most commercial simulation software in science and engineering). A second partition comprises five nodes identical to the ones of the first partition, but each featuring 1 TB RAM, to handle extremely data-intensive workloads. A third partition comprises five nodes identical to the ones of the first partition, but each node featuring two additional NVIDIA A100 graphics cards. These nodes are optimal for, amongst others, machine learning workloads. All compute nodes of HSUper are interconnected by a low-latency HDR100 InfiniBand network. The computational components are complemented by two file systems: a 1PB-sized BeeGFS-based parallel file system for most use cases and a 1PB-sized Ceph storage, which shall be particularly leveraged in combination with a digital tomography center for materials science.

HSUper has entered the TOP500 list of the fastest supercomputers in the world in June 2022, ranked 339 with a Linpack performance of 2.13 PFLOPS [4] [3]. It is further ranked 70 in the Green500 list, which measures the computational performance per consumed Watt. After a beta user phase for HSUper since spring 2022 the system is in productive operation starting late summer 2022.

Computational fluid dynamics applications are one of the main use cases of HPC resources especially when transitional or turbulent flows and multiphysics simulations such fluid-structure interactions are the objective. Therefore, it is not surprising that fluid mechanics was one of the first application fields of the new system HSUper. FIGURE 3 shows an example of some basic research application on the transition from a laminar to a turbulent flow in a pipe at Reynolds number Re = 15,000 [4]. Originally, the flow is laminar as can be seen by the straight line of the tracer particles. Triggered by a wall-mounted ring-type obstacle, small disturbances develop at about the middle of the pipe. Further downstream these are increasing finally, leading to transition. That becomes obvious in FIGURE 3 from the chaotic mixing of the tracer particles which is typical for turbulent flows. The background of such direct numerical simulations is an in-house CFD solver which is parallelized by domain decomposition relying on the message-passing interface (MPI). Before running extensive application cases, we investigated the performance of the new system for varying numbers of compute nodes and cores. The test case chosen is based on 210 MPI processes which are split up in different manners. The extreme cases are that solely one core per node or all physically available 72 cores are applied. FIGURE 4 depicts the ratio of the wall-clock time of each case to the one core-per-node case using 210 nodes. Up to 8 cores per node, the ratio of the wall-clock times stays close to unity. However, for 12 and more cores per node a saturation of the memory bandwidth of the nodes sets in leading to significantly longer wall-clock times up to a factor of 3.8 in case all cores of the Intel Icelake compute nodes are used. Thus, a

[4] 1 PFLOPS=1 PetaFLOP/s=10^{15} **Fl**oating Point **Op**erations per **S**econd

compromise between required nodes and acceptable wall-clock times has to be found in practice.

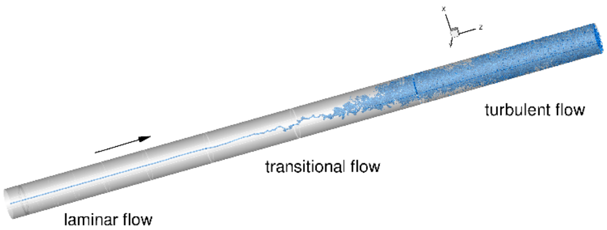

FIGURE 2: LAMINAR-TURBULENT TRANSITION IN A PIPEFLOW [4].

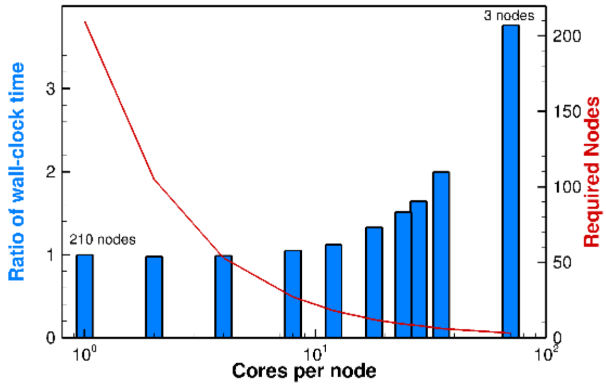

FIGURE 4: RESULTS OF PERFORMANCE TESTS ON HSUPER FOR THE COMPUTATIONAL FLUID DYNAMICS SHOW CASE.

B. Hardware Sustainability

HPC hardware develops rapidly, with novel attractive hardware released every year. It is essential to keep up with these hardware trends for strategic planning of the next-generation supercomputer. This also means that an understanding of the compute capabilities of potentially relevant novel hardware for user-specific application cases needs to be acquired and, based on this understanding, judge upon the hardware line's relevance in future procurements.

Against this background, selected novel HPC hardware is procured at small scale, installed and investigated in terms of user-defined benchmark cases by the project consortium of *hpc.bw*. Until now, two systems have been investigated:

- 4-node AMD EPYC[5] 7763 system, each node with 64 cores running at 2.25 GHz and with 256 GB RAM
- 8-node ARM-based Fujitsu A64FX[6] system, each node with 48 cores running at 2 GHz and with 32 GB high bandwidth memory

FIGURE 5 exemplarily shows scalability results for both pieces of hardware for a coupled flow simulation.

C. Outlook: Computing for Everyone – Interactive Scientific Computing Cloud

Supercomputers are efficient means for large-scale computational tasks. However, their use is more challenging and often requires advanced command line and shell scripting skills, use of HPC schedulers, and, to some extent, parallel architectures.

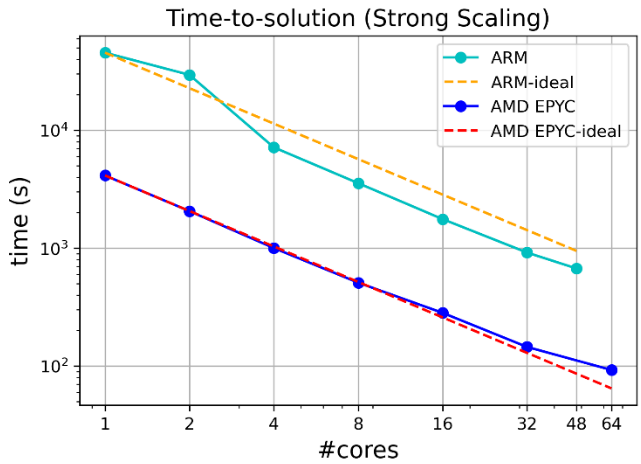

FIGURE 5: SCALABILITY, I.E. RUNTIME DEPENDING ON THE NUMBER OF COMPUTE CORES, OF A COUPLED MOLECULAR DYNAMICS-COMPUTATIONAL FLUID DYNAMICS APPLICATION, INVOKING A SIMPLE LATTICE BOLTZMANN FLOW SOLVER AND THE HIGHLY EFFICIENT MOLECULAR DYNAMICS SOFTWARE LS1 MARDYN/AUTOPAS [5]. COUPLING IS ACCOMPLISHED USING THE MACRO-MICRO COUPLING TOOL [6]. THE CONSIDERED HARDWARE SYSTEMS ARE ARM A64FX AND AMD EPYC 7763. ON BOTH PIECES OF HARDWARE, CLOSE-TO-OPTIMAL SCALABILITY IS OBSERVED, FOLLOWING THE IDEAL SCALING CURVES. OPTIMIZATION OF THE SOFTWARE FOR THE ARM ARCHITECTURE IS ONGOING WORK.

This also renders their use for (non-HPC) teaching unattractive. An interactive scientific computing on-premise cloud is planned

- to support users that have workstation-level (or slightly beyond) computational and data tasks,
- as platform for teaching simulation- and data-driven methods and their application fields,
- to simplify collaboration with other groups.

hpc.bw is supported in this endeavor by algorithmica technologies GmbH.

V. SOFTWARE EFFICIENCY AND SUSTAINABILITY

A. Performance Engineering Projects

At the core of *hpc.bw* lies the approach to boost research productivity by optimizing user-specific software on the one side and to derive new HPC research questions from this optimization process on the other side. This approach has been realized in terms of regular project calls for performance engineering projects. Every member of the Universities of the Federal Armed Forces can apply with his/ her own software. Projects are meant to last for a maximum of 12 months, but must not exceed six person months in total on HPC side for the optimization procedure. Submitted project proposals are reviewed by the *hpc.bw* project coordinator and the *hpc.bw* multiplicators, cf. Sec. II.A, to account for the interdisciplinary character of the projects.

[5] 5 out of the 10 fastest supercomputers (TOP500 list) leverage AMD EPYC processors at node-level, as of June 2022

[6] The supercomputer Fugaku, ranked 2 in the TOP500 list as of June 2022, leverages a similar node-level ARM technology

TABLE I: OVERVIEW OF PERFORMANCE ENGINEERING PROJECTS SUPPORTED BY HPC.BW IN 2022

Title	Research Field, University
benEFIT – Numerical simulation of non-destructive testing in concrete	Engineering Materials and Building Preservation, HSU
Enabling high-throughput studies of reactive materials	Computational Material Design, HSU
Monte Carlo simulations of real fluids	Thermodynamics, HSU
Optimization of an IGA code in Matlab for parallel computing	Mechanics and Statics, UniBw M

The first call was closed in February 2022. Four out of eight submissions have been selected for *hpc.bw* performance engineering support, cf. TABLE I. First results have been achieved in the project "Monte Carlo simulations of real fluids". Data buffering and SIMD pragmas for improved vectorization resulted in a speedup of 2 (i.e., halving the overall execution time) for examined use cases, cf. FIGURE 6. Upcoming calls are planned on an annual basis, with the next call anticipated for winter 2022/23.

B. *Supporting Software Sustainability Efforts*

To efficiently develop and maintain research and application codes and to allow developers to focus on their specific domain expertise (and to not have to re-write basic algorithms from scratch), today's software systems crucially rely on the existence of optimized and mature building blocks to be (re-)used by domain scientists. Within *hpc.bw*, we strive to provide such building blocks as enabling technology for other researchers and their HPC codes.

To support the simulation of coupled multi-physics phenomena in various disciplines, the *hpc.bw* project team at UniBw M develops and maintains a variety of multi-level block preconditioning methods to efficiently solve large linear systems of equations arising from the discretization of coupled problems. As illustrated in FIGURE 7, multi-level methods employ a hierarchy of coarse representations of the original fine level problem and reconstruct the fine level solution from information from the coarse levels [7]. This allows to solve problems with several million unknowns accurately and efficiently, while also obtaining *parallel scalability*, i.e. an optimal use of supercomputing hardware. We contribute our algorithms as open-source implementations to the MueLu package[7] of the Trilinos Project[8], such that they are elegantly embedded into a widely used software ecosystem for HPC and can be used by any researcher worldwide. As exemplary application, the *hpc.bw* team at UniBw M targets the embedding of thin fibers into solid bodies, as e.g. in reinforced concrete or composite materials.

To further increase the outreach into the scientific community, *hpc.bw* organizes the *European Trilinos User Group (EuroTUG) Meeting* to be held in Munich in 2022 and in Hamburg in 2024. The EuroTUG meeting series offers a platform for Europe-based users and developers of the Trilinos Project to learn about recent developments in Trilinos, report on their use cases and experiences with Trilinos, interact with the Trilinos leadership and core developers, and form a European network of Trilinos users and developers.

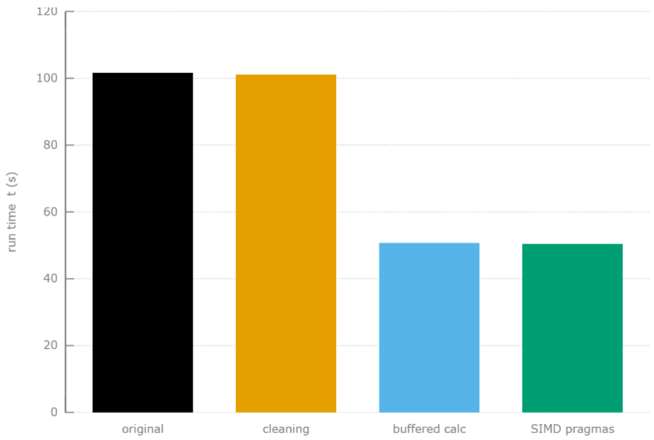

FIGURE 6: PRELIMINARY PERFORMANCE RESULTS FOR MONTE-CARLO SIMULATION AFTER SUBSEQUENT OPTIMIZATION STEPS FOR A TWO-BODY INTERACTION PROBLEM. ORIGINAL: RUNTIME BEFORE OPTIMIZATIONS. CLEANING: CODE REFACTORING TO SIMPLIFY PERFORMANCE ENGINEERING. BUFFERED CALC: BUFFERING OF PARTICLE DISTANCES TO ENABLE AUTO-VECTORIZATION. SIMD PRAGMAS: USE OF SIMD PRAGMAS TO LEVERAGE VECTORIZATION PROPERTIES OF X86 PROCESSORS.

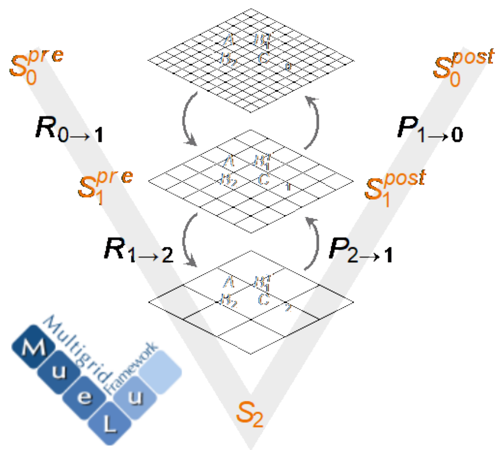

FIGURE 7: MULTI-LEVEL BLOCK PRECONDITIONING IN MUELU.

VI. OUTLOOK: HPC FOR LOGISTICS

While HPC is already well established in research and development fields such as engineering and science, there is still much room for solving optimization problems in business administration and economics [8]. Considering the application area of logistics, there are a multitude of diverse planning problems of high relevance notably with respect to economic efficiency and ecological sustainability (e.g., various kinds of scheduling and vehicle routing problems). Even simplified models of such problems from practice are often NP-hard,

[7] https://trilinos.github.io/muelu.html

[8] https://trilinos.github.io/

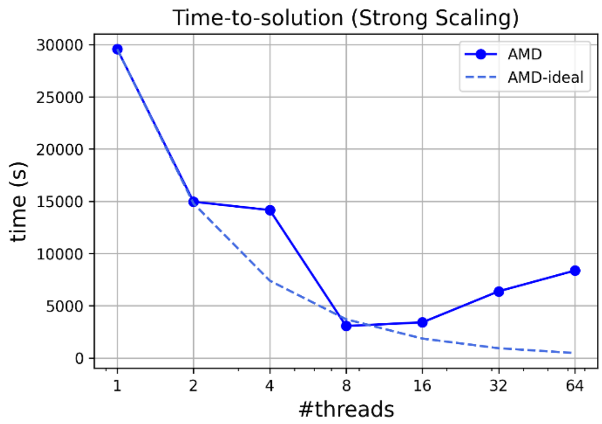

FIGURE 8: RUNTIME IN SECONDS OVER NUMBER OF USED COMPUTE CORES, I.E. USED THREADS TO COMPUTE AN OPTIMAL SOLUTION OF THE TIME-DEPENDENT TRAVELING SALESMAN PROBLEM TA061 FROM [9] WITH N=100 NODES USING IBM ILOG CPLEX OPTIMIZATION STUDIO, VERSION 22.1.0, WITH DEFAULT SETTINGS ON THE AMD SYSTEM

which means that to date many problems cannot be solved exactly even for relatively small problem sizes with ordinary computing power. Therefore, exploiting HPC power may extend the possibilities for tackling such computationally hard combinatorial optimization problems to a certain degree (regarding exactly solving somewhat larger problems or improving the approximation quality).

However, existing optimization methods often have a sequential flow logic at their core (with limited data-parallel concurrency), which is why such algorithms often must be adapted or even newly designed and implemented for HPC architectures with parallelization in mind [10]. In particular, hybrid optimization methods are being developed and investigated regarding parallelization (including mixed-integer mathematical optimization, constraint programming methods and metaheuristics/local search methods).

In a first step, we are investigating how available mixed-integer mathematical optimization solvers are already able to exploit parallel computing power of the shelf (considering both shared memory multi-core parallelization as well as distributed computing). FIGURE 8 shows first performance results for the time-dependent traveling salesman problem.

With the goal in mind to establish a joint HPC competence platform for users within, but also outside the federal armed forces, *hpc.bw* strives to collaborate with companies on HPC-related topics. The first collaboration has been initialized with Opheo Solutions GmbH, who develop a transport management software including an optimizer for tour planning. In the project, HPC experts from *hpc.bw* team up with the software developers of Opheo to improve the efficiency of the optimization process, researching parallel algorithms and efficient implementations thereof.

VII. CONCLUSION

With the project *hpc.bw*, the Universities of the Federal Armed Forces pave the way towards a competence platform around the wide field of high performance computing.

Establishing the container-based HPC center including the supercomputer HSUper was a challenge, as it required intense interactions of HPC experts, hardware vendor (including sub-contractors), infrastructure officers and architect. HSUper appears to be a very promising machine for the Universities of the Federal Armed Forces, and we look forward to exploring its performance and benefits for our HPC community.

While various view points from multiple user profiles have been taken into consideration to design the concept for the HPCCP, its actual implementation is the next step, which is planned for the upcoming months.

ACKNOWLEDGEMENT

The work described in this contribution has been funded by dtec.bw – Digitalization and Technology Research Center of the Bundeswehr which we gratefully acknowledge [project *hpc.bw: Competence Platform for Software Efficiency and Supercomputing*].

REFERENCES

[1] P. Neumann and J. Kleinschmidt, "hpc.bw Newsletter 01/2022," Hamburg, 2022.

[2] P. Neumann and J. Kleinschmidt, "hpc.bw Newsletter 02/2022," Hamburg, 2022.

[3] E. Strohmaier, J. Dongarra, H. Simon and M. Meuer, "TOP500," June 2022. [Online]. Available: https://www.top500.org/lists/top500/list/2022/06/?page=4. [Accessed 27 July 2022].

[4] F. Durst, M. Breuer, B. Ünsal and K. Haddad, "Laminar-to-turbulent transition of pipe flows triggered by wall-mounted ring-type obstacles," *J. of Turbulence,* vol. 23, no. 7, pp. 382-404, 2022.

[5] S. Seckler, F. Gratl, M. Heinen, J. Vrabec, H.-J. Bungartz and P. Neumann, "AutoPas in ls1 mardyn: Massively Parallel Particle Simulations with Node-Level Auto-Tuning," *Journal of Computational Science,* vol. 50, p. 101296, 2021.

[6] P. Neumann, H. Flohr, R. Arora, P. Jarmatz, N. Tchipev and H.-J. Bungartz, "MaMiCo: Software design for parallel molecular-continuum flow simulations," *Computer Physics Communications,* vol. 200, pp. 324-335, 2016.

[7] A. Brandt, "Multi-Level Adaptive Solutions to Boundary Value Problems," vol. 88, no. 138, pp. 333-390, 1977.

[8] G. Schryen, N. Kliewer and A. Fink, "High Performance Business Computing," *Business & Information Systems Engineering,* vol. 62, pp. 1-3, 2020.

[9] A. Fink and S. Voß, "Solving the continuous flow-shop scheduling problem by metaheuristics," *European Journal of Operational Research,* vol. 151, no. 2, pp. 400-414, 2003.

[10] G. Schryen, "Parallel computational optimization in operations research: A new integrative framework, literature review and research directions," *European Journal of Operational Research,* vol. 287, pp. 1-18, 2020.