

Towards an HPC cluster digital twin and scheduling framework for improved energy efficiency

Alexander Kammeyer^{*†}, Florian Burger^{*}, Daniel Lübbert^{*} and Katinka Wolter[†]

^{*}Physikalisch-Technische Bundesanstalt
 Abbestraße 2-12, 10587 Berlin, Germany

Email: {alexander.kammeyer, florian.burger, daniel.luebbert}@ptb.de

[†]Freie Universität Berlin,
 Takustraße 9
 14195 Berlin, Germany

Email: {a.kammeyer, katinka.wolter}@fu-berlin.de

Abstract—Demand for compute resources and thus energy demand for HPC is steadily increasing while the energy market transforms to renewable energy and is facing significant price increases. Optimizing energy efficiency of HPC clusters is therefore a major concern. Different possible optimization dimensions are discussed in this paper. This paper presents a digital twin design for analyzing and reducing energy consumption of a real-world HPC system. The digital twin is based on the HPC cluster at PTB. The digital twin receives information from multiple internal and external data sources to cover the different optimization opportunities. The digital twin also consists of a scheduling simulation framework that uses the data from the digital twin and real-world job traces to test the influence of the different parameters on the HPC cluster.

I. INTRODUCTION

THE Physikalisch-Technische Bundesanstalt (PTB) operates a compute cluster at its Berlin site that is used by various departments within PTB. The cluster is the backbone of numerous research activities. The HPC cluster currently consists of 60 CPU nodes with CPUs from two generations and two special GPU nodes with 10 GPUs. This amounts to an energy consumption of approximately 30 kW. PTB plans to extend the installed compute power by another 50 kW in Berlin and, perspective, to install a new cluster with over 200 kW installed power at its Brunswick site.

The PTB relies on HPC for many research activities in different departments including physics, mathematics and medicine. The new AI strategy of PTB will require more compute resources for AI applications and model training. Energy efficiency becomes a concern with this increased demand. While the energy efficiency of modern CPUs has improved over the years, the total energy consumption of HPC systems has increased at an even faster rate [1]. With this increasing energy usage, the total cost of ownership is impacted by increasing energy prices in a volatile market environment. Additionally, PTB, as a federal institution, is bound by the climate protection plan 2050 [2] and climate protection program 2030 [3] of the German federal government. New data centres, like the planned Brunswick site, need to meet the standards defined by the Blue Angel certificate for data centres [4].

Currently, the influence of different internal and external factors on the energy efficiency of the HPC cluster at PTB are not well understood. A digital twin can help in understanding those factors in detail and test changes to the system configuration without adverse effects to the production system and is commonly used for this purpose [5]. However, optimization goals and systems differ from every site and cluster. Thus, a specific, tailored solution is needed for this multi-factor optimization problem. Other HPC centres are also making efforts to improve their energy efficiency and are reporting on those efforts [6].

This paper introduces a design specification for digital twin of the HPC cluster at PTB. The digital twin shall help the operators of the cluster to improve the cluster operation with regards to several optimization goals. The digital twin collects data from various data sources for that purpose. In addition, the digital twin contains a simulation framework with a scheduling model. This simulation framework is used to test parameters and settings with regard to the combined optimization goals.

The next chapter presents the optimization dimensions that have to be taken into consideration. Chapter III presents the design specification of the digital twin and its different data sources. Finally, Chapter IV introduces the scheduling framework created to simulate the HPC cluster with data from the digital twin.

II. BACKGROUND

Increasing the energy efficiency and energy usage of a compute cluster is a multi-dimensional optimization problem. Energy consumption, CO₂ emissions, energy cost, cooling and energy limitations have been identified to be of particular interest of PTB [7]. Optimizations for each of these goals are possible. Each of these dimensions needs a way to be monitored and a representation within the digital twin is required. The rest of this chapter describes these dimensions in detail:

A. Energy consumption

The most obvious optimization is the overall reduction of energy consumption by the HPC cluster. Different mechanisms

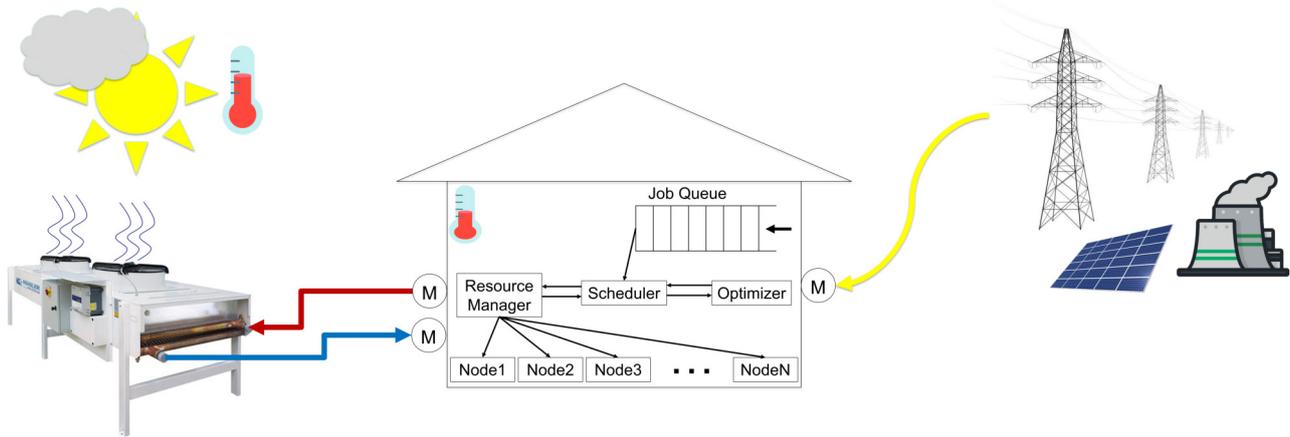


Figure 1. Schematic representation of the digital twin for PTBs HPC cluster

have been proposed like Dynamic Voltage and Frequency Scaling (DVFS) [8], turning off idle nodes [1], adapting jobs to the energy budget and running nodes at reduced frequencies [9]. The digital twin can be used to test different parameters and their effect on the cluster and monitor the long-term effect of configuration changes.

B. CO_2 emissions

The energy in the energy grid comes from different sources with different CO_2 emissions associated to them. Tracking the equivalent CO_2 emissions when using energy and moving jobs to times of low CO_2 emissions can help to reduce the CO_2 equivalent associated with the operation of the cluster. Since each energy source has different emissions, days with strong winds or low cloud coverage reduce emissions while sources like natural gas and coal used to cover base load in the energy grid increase emissions.

C. Energy cost

Energy prices for large consumers are dynamically determined at energy spot markets. These prices are volatile and shifting compute jobs to low-cost times can directly reduce energy costs. Optimizing for this goal can directly save energy costs. Using the data collected by the digital twin, it may be possible to identify a correlation between pricing and high energy availability, e.g. due to a lot of available renewable energy.

D. External influences

The weather has a direct influence on energy efficiency, hence most metrics, such as Power Usage Effectiveness [10], are averages over a year-long period to average certain weather and season related effects. PTB intends to move the cluster cooling to a free cooling system, which use temperature differences to cool the cluster. These machines are more efficient compared to classic compressor-based cooling machines but only work up to a certain outside temperature. Therefore, they might not provide sufficient cooling on extremely warm days.

The cluster needs to adapt to such conditions, e.g. by limiting the amount of active nodes. Monitoring system temperatures and integrating weather forecast data into the digital twin, the HPC system can be configured to stay within defined operational parameters, e.g. by reducing system load.

Another problem might be limited energy availability. Other industries with high energy usage have developed so called demand response mechanisms to reduce energy consumption in cooperation with energy service providers when not enough energy is available. With increasing power demands of HPC clusters such strategies might be required for HPC as well and some strategies have been proposed [11]. Similar to the limitations imposed by insufficient cooling capacities, the cluster can be adjusted to using less energy when not enough energy can be supplied by the energy provider.

III. DIGITAL TWIN FOR HPC

A digital twin is used to reason about a real-world object with data available digitally. In order to do so, real world data needs to be collected and models about the object have to be created. The previous section introduced the relevant dimensions for the optimization problem. These dimensions need to be represented in the digital twin with data and models.

The HPC cluster receives compute jobs from the users in a job queue. The scheduler decides which jobs are run on which nodes at a certain time. A resource manager module manages the available system resources. The compute nodes execute the jobs submitted to the system. Each node has its own CPU, RAM and network connection. Some nodes also have GPUs. The general operation of the cluster is simulated via a scheduling simulation. This simulation is key to understand the cluster operation. Chapter IV introduces the scheduling simulation. The nodes operation can be estimated with energy estimates. No actual jobs need to run to simulate the scheduling simulation.

The simulation results and data is used by the optimizer component that influences the scheduler and the resource manager to change the operation parameters of the HPC

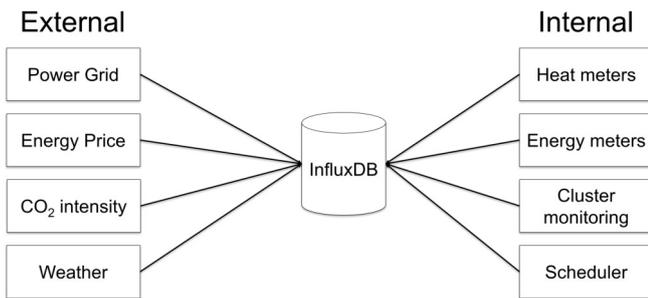


Figure 2. Software components of the digital twin with all data sources and models

cluster. This component is not yet functional but planned for future development with the results obtained by scheduling simulations from the simulation component.

The collected data generally falls into two categories: internal data from the HPC cluster and data centre, and external data about the outside world. The core component of the digital twin is an InfluxDB database. InfluxDB is a time-series database, optimized for processing time-dependent data. A collector program for each data source continuously writes the data to the Influx database. A Grafana dashboard provides access to and visualizations of the data for the user. All data collectors, Grafana and InfluxDB are managed via Docker containers. The scheduling simulator also uses the data from the InfluxDB directly.

The cluster status, including job and node information, is collected with ClusterCockpit [12]. Energy usage is collected via energy meters installed on site. This includes energy used by the cluster and the cooling infrastructure. Heat meters are used to monitor the amount of heat removed from the cluster by the cooling system.

All external data sources have been chosen specific to PTBs location and requirements as well as the optimization goals described in Chapter II. The external data sources for the digital twin include energy grid information, weather forecast data, energy prices and the CO₂ emission associated with the current energy mix in the energy grid. Information about the energy grid is obtained from Bundesnetzagentur [13]. This federal agency offers information about the German energy grid and the energy sources used at any given time. Energy prices are also retrieved from Bundesnetzagentur. CO₂ estimations are retrieved from Electricity Maps [14]. Finally, weather information are retrieved from the German Meteorological Service (DWD) [15] via the Bright Sky API [16].

Figure 1 shows the architecture described in this Chapter in a schematic representation. The individual data collectors described in this Chapter can be found in figure 2. They are grouped into internal and external data sources with InfluxDB being the core component.

IV. SCHEDULING SIMULATION FRAMEWORK

This section gives a brief overview of existing scheduling simulators and introduces a new scheduling simulator that uses

the digital twin data of the PTB HPC cluster as input for its simulations and models the operation of the cluster.

Energy aware scheduling has been of interest in the community with various goals, metrics and proposed solutions [17]. Scheduling simulators have been used to test different scheduling algorithms or additional parameters without changing production systems or the need of running actual jobs. Slurm [18] is a scheduler used in production HPC systems. Simakov et. al. developed multiple Slurm simulator versions [19], [20] that follow the steady Slurm development. All versions allow to test different Slurm parameters without altering the corresponding production system. Slurm can be combined with other schedulers like NQSV and digital twins [5]. However, since every setup is different, a custom solution is necessary. Yang et al. [21] proposed a scheduling simulator that offers two pricing options for scheduling but does not support dynamic pricing based on constantly changing energy prices.

The data of the digital twin is used for the scheduling simulations. The simulator connects to the database in order to get both the internal and external data. Because the digital twin is a purpose-build solution, no suitable simulator exists. Therefore a new solution was developed.

For development, testing and validation purposes the simulator component presented in this paper can generate synthetic job traces. For longer, realistic job traces the simulator can read the standard workload format (swf) from the Parallel Workloads Archive [22] and can use job traces from this archive as input. Additionally, job traces from the PTB compute cluster can also be used as input.

Furthermore, the data from the digital twin database as well as the job traces are time dependent. The scheduling simulation can run at arbitrary time points. The simulator was designed in such a way that it can run a given job trace at a chosen start time and match the data from the digital twin database accordingly. This allows to test varying internal and external factors on the same job trace or the same internal and external factors on different job traces.

The results of the simulations can be used to tune the system parameters to meet the optimization goals from Chapter II. An automatic optimizer can be implemented as a component of the HPC cluster for automatic tuning and adjustment of parameters.

V. CONCLUSION

With increasing compute demand and thus energy demand of HPC clusters, energy consumption and availability becomes a concern for HPC cluster operators. Possible goals for optimizations like energy consumption, CO₂ emissions, energy cost and external influences that affect energy and cooling have been discussed in this paper.

A design of a digital twin, a representation of a real-world HPC cluster at PTB, has been proposed. It has a time-series database at its core and is connected to data sources for the internal cluster aspects as well as external factors relevant for the optimization goals. The digital twin makes data about the cluster available to the administrators. This data helps to asses

and monitor the efforts towards meeting desired efficiency goals over time.

As part of the digital twin, a scheduling simulator has been developed that simulates the operation of the cluster. It uses the data from the digital twin and allows to test different job traces. Because of the time dependence of the data, the simulator can map different job traces with multiple data points to test traces with different data. The digital twin can be used by the administrators to test various configuration parameters and algorithms for their effect on the optimization goals.

So far, the digital twin is only used for simulations. It has not yet been integrated with the scheduler of the real-world HPC system. However, this shall be done in the future.

With the scheduling simulation framework, as part of the digital twin, further empirical studies of the optimization goals are planned. These simulations are a first step towards more efficient cluster operation at PTB and are the basis for future improvements of the real-world HPC cluster, especially toward an automatic optimizer component as part of the scheduler.

REFERENCES

- [1] O. Mämmelä, M. Majanen, R. Basmadjian, H. De Meer, A. Giesler, and W. Homberg, "Energy-aware job scheduler for high-performance computing," *Computer Science - Research and Development*, vol. 27, no. 4, p. 265–275, 2012. [Online]. Available: <https://doi.org/10.1007/s00450-011-0189-6>
- [2] Bundesministerium für Umwelt, Naturschutz und nukleare Sicherheit (BMU), "Klimaschutzplan 2050," <https://www.bmwk.de/Redaktion/DE/Publikationen/Industrie/klimaschutzplan-2050.pdf>, 2019.
- [3] —, "Klimaschutzprogramm 2030 der Bundesregierung zur Umsetzung des Klimaschutzplans 2050," <https://www.bundesregierung.de/resource/blob/974430/1679914/e01d6bd855f09bf05cf7498e06d0a3ff/2019-10-09-klima-massnahmen-data.pdf>, Oct. 2019.
- [4] R. UMWELT, *Energieeffizienter Rechenzentrumsbetrieb DE-UZ 161*, 2nd ed., <https://produktinfo.blauer-engel.de/uploads/criteriafile/de/DE-UZ%20161-201502-de%20Kriterien.pdf>, Fränkische Straße 7, 53229 Bonn, Feb. 2015.
- [5] T. Ohmura, Y. Shimomura, R. Egawa, and H. Takizawa, "Toward building a digital twin of job scheduling and power management on an hpc system," in *Job Scheduling Strategies for Parallel Processing*, D. Klusáček, C. Julita, and G. P. Rodrigo, Eds. Cham: Springer Nature Switzerland, 2023, p. 47–67.
- [6] M. Ott and D. Kranzlmüller, "Best practices in energy-efficient high performance computing," in *Workshops der INFORMATIK 2018 - Architekturen, Prozesse, Sicherheit und Nachhaltigkeit*. Bonn: Köllen Druck+Verlag GmbH, 2018, p. 167–176.
- [7] A. Kammeyer, F. Burger, D. Lübbert, and K. Wolter, "Optimization of energy efficiency of an hpc cluster: On metrics, monitoring and digital twins," in *Sensor and Measurement Science International*, ser. SMSI 2023. AMA Service GmbH, May 2023, p. 378–379. [Online]. Available: <https://doi.org/10.5162/SMSI2023/P51>
- [8] K. Ahmed, J. Liu, and X. Wu, "An Energy Efficient Demand-Response Model for High Performance Computing Systems," in *2017 IEEE 25th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS)*, 2017, p. 175–186.
- [9] A. Krzywaniak, J. Proficz, and P. Czarnul, "Analyzing Energy/Performance Trade-Offs with Power Capping for Parallel Applications On Modern Multi and Many Core Processors," in *2018 Federated Conference on Computer Science and Information Systems (FedCSIS)*, 2018, p. 339–346.
- [10] V. Avelar, D. Azevedo, A. French, and E. N. Power, "Pue: a comprehensive examination of the metric," *White paper*, vol. 49, 2012.
- [11] K. Ahmed, "Energy Demand Response for High-Performance Computing Systems," Ph.D. dissertation, Florida International University, Miami, 2018.
- [12] J. Eitzinger, T. Gruber, A. Afzal, T. Zeiser, and G. Wellein, "Clustercockpit — a web application for job-specific performance monitoring," in *2019 IEEE International Conference on Cluster Computing (CLUSTER)*, 2019, p. 1–7.
- [13] Bundesnetzagentur für Elektrizität, Gas, Telekommunikation, Post und Eisenbahnen (BNetzA), "SMARD - Strommarktdaten, Stromhandel und Stromerzeugung in Deutschland," <https://www.smard.de/home/marktdaten>, May 2023.
- [14] Electricity Maps ApS, "Electricity Maps," <https://www.electricitymaps.com/>, May 2023.
- [15] Deutscher Wetterdienst (DWD), "Open Data Server of the German Meteorological Service," <https://opendata.dwd.de/>, May 2023.
- [16] Bright Sky Developers, "Bright Sky JSON API for DWD's open weather data," <https://brightsky.dev/>, May 2023.
- [17] B. Kocot, P. Czarnul, and J. Proficz, "Energy-aware scheduling for high-performance computing systems: A survey," *Energies*, vol. 16, no. 2, 2023. [Online]. Available: <https://www.mdpi.com/1996-1073/16/2/890>
- [18] A. B. Yoo, M. A. Jette, and M. Grondona, "Slurm: Simple linux utility for resource management," in *Job Scheduling Strategies for Parallel Processing*, D. Feitelson, L. Rudolph, and U. Schwiegelshohn, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2003, p. 44–60.
- [19] N. A. Simakov, M. D. Innus, M. D. Jones, R. L. DeLeon, J. P. White, S. M. Gallo, A. K. Patra, and T. R. Furlani, "A slurm simulator: Implementation and parametric analysis," in *High Performance Computing Systems. Performance Modeling, Benchmarking, and Simulation*, S. Jarvis, S. Wright, and S. Hammond, Eds. Cham: Springer International Publishing, 2018, p. 197–217.
- [20] N. A. Simakov, R. L. DeLeon, Y. Lin, P. S. Hoffmann, and W. R. Mathias, "Developing accurate slurm simulator," in *Practice and Experience in Advanced Research Computing*, ser. PEARC '22. New York, NY, USA: Association for Computing Machinery, 2022. [Online]. Available: <https://doi.org/10.1145/3491418.3535178>
- [21] X. Yang, Z. Zhou, S. Wallace, Z. Lan, W. Tang, S. Coghlan, and M. Papka, "Integrating dynamic pricing of electricity into energy aware scheduling for HPC systems," in *International Conference for High Performance Computing, Networking, Storage and Analysis, SC*, 2013.
- [22] D. G. Feitelson, D. Tsafir, and D. Krakov, "Experience with using the parallel workloads archive," *Journal of Parallel and Distributed Computing*, vol. 74, no. 10, p. 2967–2982, 2014. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0743731514001154>