EE-HPC – A framework for energy efficient HPC system operation

Jan Eitzinger, Thomas Gruber
Erlangen National High Performance Computing Center
Erlangen, Germany

Christian Terboven, Radita Liem
IT Center, RWTH Aachen University
Aachen, Germany

Jose Gracia, Kingshuk Haldar
HLRS
University of Stuttgart
Stuttgart, Germany

Pay Giesselmann, Jan Frederik Engels
German Climate Computation Center (DKRZ)
Hamburg, Germany

Torsten Wilde, Marcel Marquardt, Jan Maeder, Christian Simmendinger, David Brayford
Hewlett Packard Enterprise (HPE)
Boeblingen, Germany

ACM Reference Format:
Jan Eitzinger, Thomas Gruber, Christian Terboven, Radita Liem, Jose Gracia, Kingshuk Haldar, Pay Giesselmann, Jan Frederik Engels, and Torsten Wilde, Marcel Marquardt, Jan Maeder, Christian Simmendinger, David Brayford. 2023. EE-HPC – A framework for energy efficient HPC system operation. In *Proceedings of International Conference for High Performance Computing, Networking, Storage, and Analysis (SC '23).* ACM, New York, NY, USA, 3 pages.

INTRODUCTION

The need for more compute resources is pushing the power envelope of data centers higher, leading to increased capital expenditure (CAPEX) and operating expenses (OPEX). Especially the cost of energy consumption in HPC data centers is becoming a decisive factor in the procurement and operation of the systems. For procurement's, offered systems are not only evaluated based on throughput benchmark results but also based on the energy efficiency of the quoted solution. Additionally, there is an increased public awareness for sustainability and carbon-neutral solutions. Both factors have led to intensified research on energy reduction while keeping system performance high. Newly deployed systems are equipped with meters and knobs to estimate and control the power consumption.

The implemented OS control system power management however, is working on a coarse level not being able to use knowledge of the actual workload running across multiple compute nodes. For a truly improved system energy efficiency this approach is inadequate.

The EE-HPC project aims to develop a solution with an application-aware, dynamic, user transparent, and policy-driven approach. The project started end of 2022 and builds on the existing software components ClusterCockpit [1] and LIKWID [2].

SC '23, November, 2023, Denver, CO

STATE OF THE ART

Multiple approaches to coordinate power usage in cluster systems like GEOPM (Global Extensible Open Power Manager) [3], EAR [2], READEX (Runtime Exploitation of Application Dynamism for Energy efficient eXascale computing) [4], the PowerStack Initiative [1], and REGALE [5] exist. GEOPM targets MPI and OpenMP applications and uses a hierarchical control system to implement dynamic power management among the compute nodes of a job. EAR manages the energy efficiency of HPC cluster systems by optimizing nodes locally, determines the job phases through intercepted MPI calls. These job phases are analyzed and compared with different models to map them to configurable policies. The EU-funded READEX project requires code annotations inside applications and libraries and applies auto-tuning to each phase. Since it relies on the annotations, it is of limited extent for production operations. The PowerStack Initiative proposes a dynamic power management software stack for HPC systems aiming to specify vendor-neutral requirements, interfaces, data exchange protocols as well as control knobs. REGALE is an open architecture project integrating with existing components of the HPC stack like resource management systems and MPI libraries. It tries to cover the whole end-to-end cycle of a supercomputing system with one part being energy management.

MANAGEMENT OF POWER AND ENERGY

In the EE-HPC project approach, job specific power distribution and energy consumption is guided by a central component. This central component (Energy Manager, EM) is aware of external power constraints, job policies, power domains, and node allocations per job. Interaction with the batch system is kept minimal. Compute node agents gather required metrics and enforce the decisions from the EM. Communication between node agents and EM is managed by a pub/sub interface to an in-memory database. Jobs are dynamically steered at runtime, based on job policies and application workload.

[1] https://www.clustercockpit.org/
[2] https://hpc.fau.de/research/tools/likwid/

Examples for job policies are “minimal energy to solution”, “maximal performance at minimal energy”, “high priority jobs”, “profiling queues” or “low priority job”. During a jobs execution, the compute behavior of the job is characterized dynamically, either, via a machine-learning approach (based on an application footprint derived from runtime measurements) and steered according to the detected footprint or directly steered via the optimization of a given metric such as the energy-delay product.

FINE-GRAINED ENERGY EFFICIENCY LIBRARY

To enable the characterization and fine-grained optimization of an application, the EE-HPC project is creating an instrumentation library. The library will tag the application's current runtime behavior (I/O-bound, compute-bound, or memory-bound), either via a direct user-guided instrumentation or via the MPI [3] or OMPT [4] profiling and tools interfaces. The node-agent forwards the information to the EM which then, according to job-policy and detected workload characteristics decides on a specific set of tunables: power/frequency and/or environment changes pertaining to the programming model. The set of tunables then is enforced by either the node-agent or in the case of environment changes to the programming models by the instrumentation library. The enforcing of tunables happens asynchronously. The EM registers an tag-specific action with the node-agent and the node-agent triggers the required action node-locally. In order to obtain application characteristics and the tagging of applications in an automated way the MPI profiling interface and OMPT interface are used. In this way, the well-studied [6] power characteristics of the programming models by opting for a, for example, more situation-appropriate collective algorithm for MPI, can be leveraged.

PERFORMANCE AND ENERGY MONITORING

Nowadays, many HPC centers monitor their systems for general health and reasonable usage but often do not provide access to this data to the user or even support personal. Moreover, users are only partly interested in the metrics the site operators look at. Therefore, the NHR@FAU started the development of a job-specific monitoring framework that provides relevant HPC metrics to help the users and support personal to asses the quality and optimization potential of jobs. The *ClusterCockpit* framework provides loosely coupled components for data collection (`cc-metric-collector` [5]), data storage and retrieval (`cc-metric-store` [6]) as well as a central controller with web frontend for visualization (`cc-backend` [7]) using standardized protocols [8]. Various throughput and energy related metrics are collected on each compute node including CPU metrics, GPU metrics, and OS statistics using different collection methods, for example, the Likwid library for CPU metrics, rocm-smi for AMD GPU metrics, or nvidia-smi for NVIDA GPU metrics. The job-specific dashboards provide general information, time series plots and a time-resolved Roofline model [7] that deemed most appropriate for the target group (HPC users and support personal).

[3]https://www.mpi-forum.org/docs/mpi-3.1/mpi31-report/node359.htm
[4]https://www.openmp.org/spec-html/5.0/openmpsu15.html
[5]https://github.com/ClusterCockpit/cc-metric-collector
[6]https://github.com/ClusterCockpit/cc-metric-store
[7]https://github.com/ClusterCockpit/cc-backend
[8]https://github.com/ClusterCockpit/cc-specifications

In the EE-HPC project, the ClusterCockpit framework will be extended by a new component *cc-node-control* which provides an interface to control knobs of the system as well as the job runtime for the Energy Manager. Furthermore, the cc-metric-collector is planned to read more energy related metrics from the nodes as well as PDUs and infrastructure to support the Energy Manager in its decisions. An communication channel with defined protocol provides information to and from the energy efficiency library for fine-grained control. Besides collecting and processing time series data, the EE-HPC project will add event based information throughout the software stack. This allows the web-frontend to display additional information to the users such as a timeline of the decision made by the EM and other event information such as the compute phase detection of an application.

TARGET VALIDATION - ICON

The ICON modeling framework is a joint project between the German Weather Service (DWD), the Max Planck Institute for Meteorology (MPI-M), the German Climate Computing Center (DKRZ), and the Karlsruhe Institute of Technology (KIT). It aims at developing a unified next-generation global numerical weather prediction and climate modeling system. The ICON model has been introduced into DWD's operational forecast system in January 2015 and is also used by many other weather services in Europe. The ICON atmosphere and ocean part is based on an icosahedral grid for the horizontal layout and uses a rectilinear grid for the vertical setup. It allows for high resolution climate and weather experiments, as well as regional simulations and nested setups. ICON also includes modules for land and chemical processes.

The model is currently being extended to run on the first exascale HPC systems in Europe, which will require better scalability and use of GPU accelerators. These systems will be needed to run the most ambitious simulations, namely a global climate prediction on a uniform grid of just 1km in size. At this fine resolution, atmospheric convection in particular, and hence the evolution of clouds and precipitation processes, can be simulated almost explicitly, without the need for error-prone parameterisations of these processes as at coarser resolutions. Small-scale features such as ocean eddies are also now resolved and can be studied in more detail.

ICON is used by all partners in the EE-HPC project to test and validate the energy management and application library on each partner's HPC system. Typical climate projections are adjusted based on reference experiments to achieve the most energy efficient design. The scientific output of the model must not be changed by the approaches in EE-HPC, but for example small increases in runtime may be acceptable if significant energy savings can be achieved. Finding the right balance is particularly important for operational weather forecasting, where simulations must be completed within fixed time windows. Long-term climate projections are somewhat more relaxed in this respect.

ACKNOWLEDGEMENTS

The initial development of the ClusterCockpit framework was funded by the NHR initiative. The EE-HPC project is funded by the German Federal Ministry of Education and Research in the “GreenHPC” call.

REFERENCES

[1] Christopher Cantalupo, Jonathan Eastep, Siddhartha Jana, Masaaki Kondo, Matthias Maiterth, Aniruddha Marathe, Tapasya Patki, Barry Rountree, Ryuichi Sakamoto, Martin Schulz, and Carsten Trinitis. [n. d.]. A Strawman for an HPC PowerStack. ([n. d.]). https://doi.org/10.2172/1466153

[2] Julita Corbalan, Lluis Alonso, Jordi Aneas, and Luigi Brochard. 2020. Energy Optimization and Analysis with EAR. In *2020 IEEE International Conference on Cluster Computing (CLUSTER)*. 464–472. https://doi.org/10.1109/CLUSTER49012.2020.00067

[3] Jonathan Eastep, Steve Sylvester, Christopher Cantalupo, Brad Geltz, Federico Ardanaz, Asma Al-Rawi, Kelly Livingston, Fuat Keceli, Matthias Maiterth, and Siddhartha Jana. 2017. Global Extensible Open Power Manager: A Vehicle for HPC Community Collaboration on Co-Designed Energy Management Solutions. In *High Performance Computing*, Julian M. Kunkel, Rio Yokota, Pavan Balaji, and David Keyes (Eds.). Springer International Publishing, Cham, 394–412.

[4] Joseph Schuchart, Daniel Hackenberg, Robert Schöne, Thomas Ilsche, Ramkumar Nagappan, and Michael K Patterson. 2016. The shift from processor power consumption to performance variations: fundamental implications at scale. *Computer science-research and development* 31 (2016), 197–205.

[5] Eishi Arima TUM, Bruno Raffin UGA, Yiannis Georgiou, Daniele Cesarini, Nikela Papadopoulou, Mohsen Seyedkazemi, Andrea Bartolini, Pierre-François Dutot UGA, and Martin Schulz. 2021. REGALE: An open architecture to equip next generation HPC applications with exascale capabilities. (2021).

[6] Akshay Venkatesh, Krishna Kandalla, and Dhabaleswar K. Panda. 2013. Evaluation of Energy Characteristics of MPI Communication Primitives with RAPL. In *2013 IEEE International Symposium on Parallel Distributed Processing, Workshops and Phd Forum*. 938–945. https://doi.org/10.1109/IPDPSW.2013.243

[7] Samuel Williams, Andrew Waterman, and David Patterson. 2009. Roofline: An insightful visual performance model for multicore architectures. *Commun. ACM* 52, 4 (2009), 65–76. https://doi.org/10.1145/1498765.1498785