Towards energy-efficient data center infrastructure – a holistic approach based on software for modeling, simulation, and (re)configuration of the energy network

Torsten Wilde*, Tanja Clees†, Horst Schwichtenberg†, Hayk Shoukourian*‡, Detlef Labrenz*, Inna Torgovitskaia†, Michael Schnell†, Nils Hornung†, Bernhard Klaaßen†, Eric Lluch Alvarez†,
*Leibniz Supercomputing Centre (LRZ) of the Bavarian Academy of Sciences and Humanities, Garching bei München, Germany
Email: torsten.wilde@lrz.de
†Fraunhofer SCAI, Sankt Augustin, Germany
Email: tanja.clees@scai.fraunhofer.de
‡Technical University Munich (TUM), Munich, Germany

I. Extended Abstract

According to a data center survey done by DataCenterDynamics Intelligence [1], in the past decade the power consumption of data centers has increased substantially from 18GW in 2007 to an estimated 40GW in 2013 making power consumption an important constrained for new data centers. Also, according to the annual Data Center Industry Survey 2013 by the Uptime Institute [2] the data center average Power Usage Effectivness (PUE) improved from over 2.7 in 2007 to 1.65 in 2013. This indicates that around 16GW of the total data center power consumption was spent on the infrastructure needed to run and cool the IT systems. Therefore, increasing the energy efficiency of the data center infrastructure and evaluating the possible reuse of waste heat are becoming increasingly important issues in the design and operation of large data centers. In the past this was not a strong focus since all data centers used air cooling. Today, many High Performance Computing (HPC) data centers use a combination of different cooling technologies (chilled water, chiller-less cooling, air cooling, immersion cooling etc.), but there exists a lack of knowledge regarding the energy efficient operation of the different cooling infrastructures which are still designed and optimized for the maximum possible load, like traditional air cooling. However, this peak load is rarely seen in production. For example, SuperMUC (Number 14 on Top500 list from 11/2014) installed at the Leibniz Supercomputing Centre (LRZ) has an idle power consumption of 800kw, an average power consumption of 2.2 -2.4 MW in normal operation, and a maximum power consumption of 3.6MW achieved during the compute intensive Linpack benchmark. The combination of this load variability and the use of new cooling technologies in the data center provides an opportunity to increase the energy efficiency of the data center.

SIMOPEK [3] is one project funded by the German Ministry of Education and Research (BMBF) trying to improve the energy efficiency of modern data centers. It uses a wholistic approach based on the "4 Pillar Framework for energy efficient HPC data centers [4]". The main idea is to create a model of the data center cooling infrastructure in order to be able to simulate it's behavior and to automatically find possible optimizations. For creating a model and to verify and validate the results one needs to be able to collect and use real operational data from both the data center infrastructure (cooling and electrical circuit) and the installed super computer (power consumption and job information). To that end, the Power Data Aggregation Monitor (PowerDAM) software was developed by LRZ. PowerDAM [5] is able to collect data from the LRZ building automation system (Johnson Controls), the power management system (SIEMENS WinCC), from decisive HPC systems (CooLMUC, SuperMUC and SuperMIG) as well as from two batch scheduling systems (Slurm, LoadLeveler). It provides analysis and reporting capabilities, such as, Energy-to-Solution per application, energy budget information on a per user basis, and reports of data center key performance indicators such as PUE and cooling systems Coefficient of Performance (COP). In SIMOPEK, PowerDAM is used to transform and export collected data for model development, and for verification and validation of the simulation and optimization results provided by MYNTS.

MYNTS ([6], [7]), developed by Fraunhofer SCAI, is a software for modelling, simulating, analyzing, and optimizing energy networks (e.g. gas transport, electricity, water). In SIMOPEK, a model of the data center cooling loops, which cool CooLMUC and SuperMUC, was built based on PowerDAM data including a schematic representation of the network (created by SCAI's AutoCAD parser) as well as characteristic fields for involved pumps and other devices. Several scenarios reflecting important operating conditions of CooLMUC and SuperMUC are considered as a basis for simulation, statistical analysis (parameter ensembles) and optimization. The result of the optimization should help to improve the energy efficiency of the LRZ chiller-less cooling infrastructure. Along with decreasing the data center power overhead, created by the cooling infrastructure, the re-use of waste heat could also be beneficial. One possible way to do this is to use adsorption cooling to create cold from the waste heat. Therefore, part of the MYNTS

development is the creation of models describing existing SorTech adsorption chillers and any future developments. This will allow MYNTS to find an optimal adsorption chiller design for data centers.

Overall, the SIMOPEK project is developing methods and software components for modeling, simulating, and optimizing the cooling infrastructure of a data center. The models take into account both the highly dynamic load behavior of the HPC system as well as new technological components (high-temperature liquid cooling) and concepts for recycling the generated waste heat (adsorption cooling by SorTech). In this way, a virtual reconfiguration of the cooling circuits can be performed and studied prior to physically rebuilding the system with the goal to efficiently use and re-use energy.

Acknowledgment

The work presented here has been carried out within the SIMOPEK project, which has received funding from the German Federal Ministry for Education and Research under grant number 01IH13007A.

References

[1] DCD Intelligence, "Is the industry getting better at using power?" *Data Center Dynamics FOCUS 33, January/February 2014*, vol. 33, pp. 16 – 17, 2014. [Online]. Available: http://content.yudu.com/Library/A2nvau/FocusVolume3issue33/resources/index.htm?referrerUrl=

[2] M. Stansberry, "2013 uptime institute annual data center industry survey report and full results." [Online]. Available: http://www.data-central.org/resource/collection/BC649AE0-4223-4EDE-92C7-29A659EF0900/uptime-institute-2013-data-center-survey.pdf

[3] SIMOPEK, http://www.simopek.de/.

[4] T. Wilde, A. Auweter, and H. Shoukourian, "The 4 pillar framework for energy efficient hpc data centers," *Computer Science - Research and Development*, pp. 1–11, 2013. [Online]. Available: http://dx.doi.org/10.1007/s00450-013-0244-6

[5] H. Shoukourian, T. Wilde, A. Auweter, and A. Bode, "Monitoring power data: A first step towards a unified energy efficiency evaluation toolset for HPC data centers," *Environmental Modelling And Software*, vol. 56, no. 0, pp. 13 – 26, 2014, thematic issue on Modelling and evaluating the sustainability of smart solutions. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1364815213002934

[6] S. Grundel, N. Hornung, B. Klaassen, P. Benner, and T. Clees, "Computing surrogates for gas network simulation using model order reduction," in *Surrogate-Based Modeling and Optimization*, S. Koziel and L. Leifsson, Eds. Springer New York, 2013, pp. 189–212. [Online]. Available: http://dx.doi.org/10.1007/978-1-4614-7551-4_9

[7] J. Lorenz, E. Bar, T. Clees, R. Jancke, C. Salzig, and S. Selberherr, "Hierarchical simulation of process variations and their impact on circuits and systems: Methodology," *Electron Devices, IEEE Transactions on*, vol. 58, no. 8, pp. 2218–2226, Aug 2011.