学際大規模情報基盤共同利用・共同研究拠点 公募型共同研究 2024年度採択課題
16th Symposium
Joint Usage / Research Center for Interdisciplinary Large-scale Information Infrastructures
jh240074

Toshihiro Hanawa (The University of Tokyo)

Energy Efficient Operation for Supercomputer Systems

Introduction

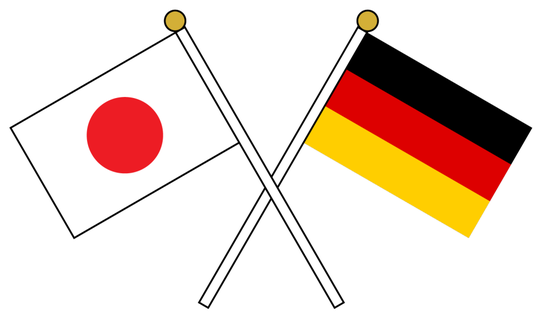

- Higher operation cost is severe issue for Supercomputer centers.
- In near future, Carbon footprint must be reduced and realized Net-0.
- The performance of the supercomputer systems in Japan will be able to be compared with that of German supercomputer centers, including the difference in climate and the performance profile of the advanced high-temperature water cooling system.

Purpose

- Explore optimal energy-efficient operation methods for the supercomputer system
- Supercomputer systems in operation are measured using benchmarks and real applications, and energy consumption and carbon footprint are modeled according to those results.
- Clock frequency should be appropriately controlled for the processors (CPU & GPU) at frequency suitable for processing at the time.

Resource

Japan	
U. Tokyo	Wisteria/BDEC-01 (Odyssey: A64FX, Aquarius: ICX+A100 GPU), Miyabi (OFP-II, 2025/1-)
Tokyo Tech	TSUBAME 4.0 (Gemini+H100 GPU)
Kyoto U.	Camphor 3 (SPR HBM)
Osaka U.	SQUID (ICX+A100 GPU, SX-Aurora Tsubasa)

Germany (warm water cooling)	
RWTH Aachen	SPR+H100 GPU (PUE monitoring)
Paderborn U.	Noctua-1, **Noctua-2** (A100, FPGA), **NHR-1** (GPU)
LRZ	**SuperMUC-NG** (Intel SPR+PVC)
TU Dresden	Under planning

Theme

1. Study on the usage of low-precision calculation and for energy reduction (Sakamoto, Hanawa, Nakajima, Fukaya, Katagiri, Ohshima, Wellein)-

- The power efficiency improvement by aggressively using low-precision arithmetic for HPC applications will be studied
- In collaboration with **jh240021 project** for target applications, we will investigate energy profiles using various architectures on both Japanese and German computing resources.

2. Study on analysis corresponding among cooling, energy consumption, and performance (Hanawa, Sumimoto, Miki, Ohara, Nomura, Nonaka, Miura, Shoji, Terai, Plessl, Simon, Rohde, Ilsche, Huber, Müller)

- To lower the operating costs of supercomputing systems ➔ Higher temperature cooling water should be used and ideally only cooling tower without chiller system
- Higher operating temperature may impact the CPU power consumption due to the leakage current, and it may degrade performance due to DVFS (Dynamic Voltage-Frequency Scaling) mechanism activated to maintain the power and temperature within TDP (Thermal Design Power) limit.
- We will directly compare such systems using different cooling methods in Germany and Japan.

3. Study on parameterizing optimal configuration based on application performance (Yamazaki, Shimokawabe, Hanawa, Takizawa, Endo, Kawai, Fukazawa, Nanri, Miwa, Yoshida, Kusaba, Honda, Eitzinger, Müller, Terboven)

- Variations on performance and power consumption for each processors can be observed even in the same cluster [2].
  - Variations are strongly affected by both the kernels that are running and the process generation [5].
- According to the algorithms, some processor cores could be set to sleep to save energy and allow better use of the power and thermal budget [7]. In collaboration with **jh240072 project**, optimal parameters for parallelization can be determined dynamically.
- Evaluate kernel characteristics in benchmarks and applications using various kinds of architectures, analyze to classify them into CPU-bound, memory-bound, communication-bound, IO-bound, etc., and parameterize to determine optimal configuration for energy efficiency.

Reference

[1] R. Sakamoto, M. Kondo, K. Fujita, T. Ichimura, and K. Nakajima, “The Effectiveness of Low-Precision Floating Arithmetic on Numerical Codes: A Case Study on Power Consumption,” HPC Asia 2020, pp. 199-206, Jan. 2020.
[2] J. Nonaka, T. Hanawa, and F. Shoji, “Analysis of Cooling Water Temperature Impact on Computing Performance and Energy Consumption,” IEEE CLUSTER, pp. 169-175, Sep. 2020.B. Wang,
[3] A. Kozhokanova, C. Terboven, M. S. Müller, “RLP: Power Management Based on a Latency-Aware Roofline Model,” IEEE Intl. Symp. on Parallel and Distributed Processing (IPDPS), pp. 446-456, May 2023.
[4] H. Shoukourian, T. Wilde, H. Huber, and A. Bode, “Analysis of the efficiency characteristics of the first High-Temperature Direct Liquid Cooled Petascale supercomputer and its cooling infrastructure,” Journal of Parallel and Distributed Compt., Vol. 107, pp. 87-100, 2017.
[5] K. Yoshida, R. Sageyama, S. Miwa, H. Yamaki, and H. Honda, “Analyzing Performance and Power-Efficiency Variations among NVIDIA GPUs,” the 51st International Conference on Parallel Processing (ICPP), No. 65, pp. 1-12, Nov. 2022.
[6] Y. Miki and T. Hanawa, “Performance evaluation of N-body computation code on NVIDIA H100 PCIe and AMD MI210,” IPSJ SIGHPC Tech. Report, Vol. 2023-HPC-190, No. 28, pp. 1-13, Jul. 2023.
[7] A. Klawonn, M. Lanser, O. Rheinbach, G. Wellein, and M. Wittmann, “Energy efficiency of nonlinear domain decomposition methods,” Intl. Journal of High Performance Computing Applications, vol. 35, No. 3, pp. 237-253, 2021.