Development of Monitoring and Management Methods for Efficient Large HPC System Use

Joint research project between JAXA and Morgenrot Inc., Japan.

Hiroshi Ito[*1], Masamichi Nakamura[*1], Takuro Shoji[*1], Shunji Uno[*2], Minori Hirose[*2], and Naoyuki Fujita[*2]

[*1] Morgenrot Inc., Japan [*2] Japan Aerospace Exploration Agency (JAXA)

Introduction

- The job filling rate is reported as about 95% in JAXA Supercomputing System, JSS3[1]. To improve effective and healthy uses, we have continued the research on analysis and monitoring methods about its usage[3]. However, that research has been limited to test analysis, and has not been extended to other metrics or implemented in concrete tools yet considering realistic management.
- In this study, we deepened detailed analysis and monitoring methods for the large HPC system. The detailed usages are analyzed, and new metrics detections are demonstrated by using EAR[3], a system software for energy management and accounting. Also, those are applied to development for new monitoring tool.

Conclusion

- **Analysis of JSS3 jobs**
  Varieties of CPU utilization are observed in job-allocated states. That enables us to grasp status change in the system instantly.
- **Detections of detailed metrics**
  Flops, I/O, and bandwidth benchmarking are performed via EAR[2]. Those show similar results as each software and it is a suitable tool for metrics detections.
- **Development of monitoring tool**
  The monitoring tool is developed for displaying analysis results and metrics. It has several dashboard figures to grasp system status easily.

Analysis of the current system JSS3

Overview

- JSS3, JAXA Supercomputing System, mainly consists of TOKI-SORA, the main computing platform, and TOKI-RURI, the general-purpose computing system. TOKI-SORA and TOKI-RURI systems have the peak-performances of 19.4 and 1.24 PFLOPS, respectively. There are about 400,000 jobs from the operation start in 2020 to Jun 2025.

**1 Data for 2025 is covered until middle of June. **2 Some data has no data in the starting time.

Number of JSS3 jobs through Dec 2020 – Jun 2025

Year	TOKI-SORA	TOKI-RURI
2020	17792	116798
2021	1497979	14819784
2022	1209599	5690189
2023	430647	7228169
2024	494027	3416708
2025**1	218104	1766088
NA**2	259797	225209
Sum	4127945	33262945

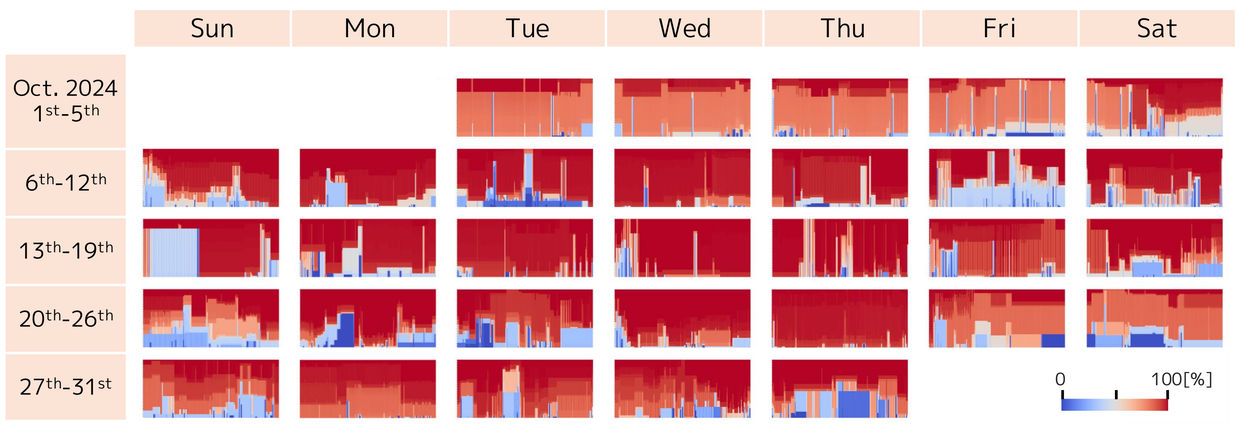

CPU utilizations (User mode) in sample 48 nodes for Oct. 2024[3]

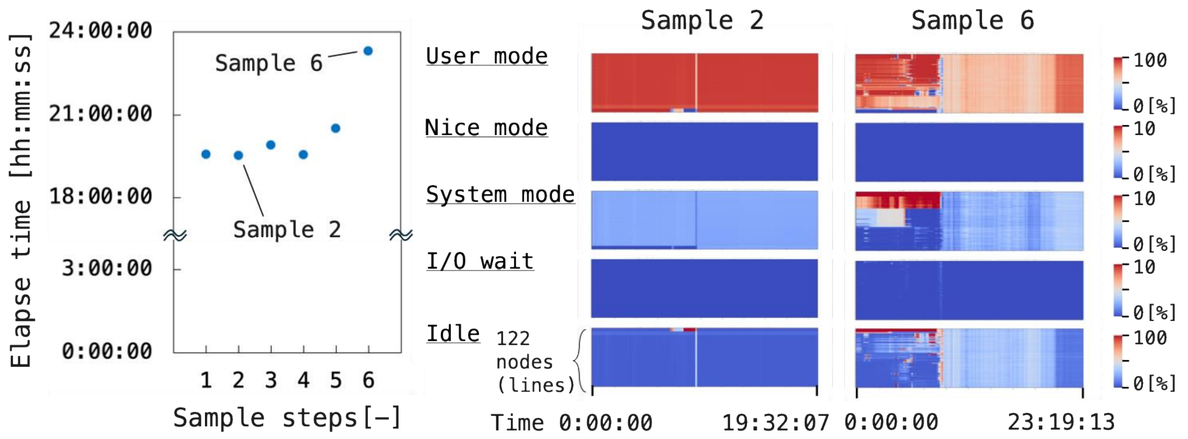

Performance observation by detailed CPU utilization

Calendar style monitoring in CPU utilization

- It enable us to check the daily and weekly usage tendencies. In the sample case of Oct 2024, the usage did not depend on a specific time and a day of week.
- This monitoring style is useful to grasp the status and usage of system easily.

Performance checking for jobs

- In the sample step job, the sample 6 has lower user mode value than that of the sample 2 especially in the latter half of the job.
- This example enable us to grasp status change in the system instantly.

Detailed metrics detections via EAR[2]

FLOPS

- EAR shows flops data for each benchmark conditions easily. It will contribute to management system in the performance detections.

**3 Test calculations were carried out using volatile variables and SIMD intrinsics (_mm_add_ps, _mm256_add_ps, _mm_add_pd, _mm256_add_pd) for float and double data types.

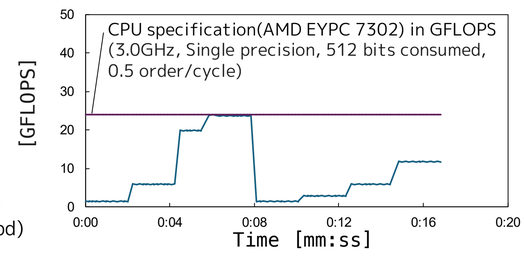

Flops during sample jobs[3]**3

I/O

- Considering the output intervals of 10 s, EAR shows I/O rates appropriately matched with the IOR[4], benchmark soft for I/O. This will be useful to monitor the system conditions.

**4 IOR benchmark conditions: write process only, processes = 2, blockSize = "128g", segments = 1, transferSize = "4m", iterations = 5

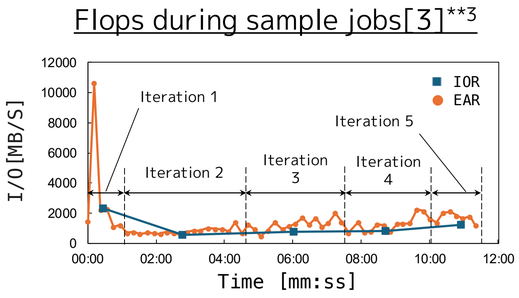

I/O during IOR sample runs[3]**4

Bandwidth

- EAR shows GBS value appropriately matched with the STREAM[5], benchmark soft for bandwidth.

**5 Benchmark is performed on a test environment of AMD EPYC 7302(CPU).

Software	Index	Condition and results	
STREAM	Function		Copy
STREAM	Array size		20,000,000
STREAM	Cycles		10,000
STREAM	Best time[s]	0.034498	–
STREAM	Ave time[s]	–	0.052915
STREAM	[GB/s]	9.276	6.047
EAR*	GBS		6.27

GBS during STREAM sample run**5

Development of the monitoring method/tool

Overview of the tool

- CPU utilizations and other metrics examined in this study are monitored on our monitoring tool, MORGENROT Arthur[6].
- Outline of MORGENROT Arthur[6]
  - Reading logs of job schedulers (Slurm, OpenPBS, etc.) and EAR, and displaying statistical data on the dashboard
  - On-premises and Web versions (without/with cloud service in the backend system, respectively)

Example system of cluster and monitoring tool, MORGENROT Arthur[6]

References

[1] Japan Aerospace Exploration Agency., JAXA Supercomputer System Annual Report (February 2024-January 2025) [2] Energy Aware Solutions S.L., Energy Aware Runtime, https://www.eas4dc.com/ [3] Morgenrot Inc., Presentations in Exhibit at ISC High Performance 2025. [4] Hongzhang Shan et al., "Using IOR to analyze the I/O Performance for HPC Platforms." Lawrence Berkeley National Laboratory, June 8, 2007. UNT Digital Library, ark:/67531/metadc898900. [5] John D. McCalpin, STREAM, https://www.cs.virginia.edu/stream/ [6] Morgenrot Inc., MORGENROT Arthur, https://morgenrot.net/en/arthur

JAXA Annual Report[1] MORGENROT Arthur[6]

Acknowledgements

We are grateful to Mr. Raphael Brochard, Dr. Julita Corbalan Gonzalez, and Dr. Luigi Brochard, Energy Aware Solutions S.L., for supports in usage and discussions on results, related to EAR (Energy Aware Runtime).