A View from the Facility Operations Side on the Water/Air Cooling System of the K Computer

Jorji Nonaka*
HPC Usability Dev. Unit
RIKEN R-CCS
Kobe, Japan

Keiji Yamamoto
System Operations & Dev. Unit
RIKEN R-CCS
Kobe, Japan

Akiyoshi Kuroda
Application Tuning Dev. Unit
RIKEN R-CCS
Kobe, Japan

Toshiyuki Tsukamoto
Facility Operations & Dev. Unit
RIKEN R-CCS
Kobe, Japan

Kazuki Koiso
Grad. School of System Informatics
Kobe University
Kobe, Japan

Naohisa Sakamoto
Grad. School of System Informatics
Kobe University
Kobe, Japan

ABSTRACT

Eight years have been passed since the first appearance of the K computer in the TOP500 list as the most powerful system at that time. Currently, the count down for the final shutdown has already started, and the preparations for the substitution are also underway. The Operations and Computer Technologies Division at the RIKEN R-CCS is responsible for the management and operations of the entire Facility, which includes the auxiliary subsystems such as the power supply, and water/air cooling systems. It is worth noting that part of these subsystems will be reused in the next supercomputer (Fugaku), thus a better understanding of the operational behavior as well as the potential impacts especially on the hardware failure and energy consumption would be greatly beneficial. In this poster, we will present some preliminary impressions of the impact of the cooling system on the K computer, from the Facility Operations side point of view, focusing on the potential benefits of the use of low water/air temperature respectively for the CPU (15^oC) and DRAM memory modules (17^oC) produced by the cooling system. We expect that the obtained knowledge will be helpful for the decision support and/or operation planning of the next supercomputer.

CCS CONCEPTS

• **Information systems** → **Data analytics**; *Data centers*; • **Mathematics of computing** → *Time series analysis*; • **Hardware** → *Power and energy*.

KEYWORDS

Facility operation, chilled water cooling, hardware failure, energy consumption

*Corresponding Author: jorji@riken.jp.

SC19 Poster, November 2019, Denver, CO

ACM ISBN 978-1-4503-9999-9/18/06…$15.00
https://doi.org/10.1145/1122445.1122456

ACM Reference Format:
Jorji Nonaka, Keiji Yamamoto, Akiyoshi Kuroda, Toshiyuki Tsukamoto, Kazuki Koiso, and Naohisa Sakamoto. 2019. A View from the Facility Operations Side on the Water/Air Cooling System of the K Computer. In *Proceedings of IEEE/ACM Supercomputing Conference (SC19 Poster).* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/1122445.1122456

1 INTRODUCTION

Almost five years have been passed from the first detailed operation analysis [6], and four years from the detailed hardware failure analysis [5] of the K computer system. There was no surprise in the results, and just confirmed the high reliability and availability explained by Miyazaki et al. [4] about the adopted design principles of the hardware system. A worthy note was the extremely low CPU failure rate (< 0.01%) during the regular usage, and with some peaks during some heavy usage, such as LINPACK and some Grand Challenge applications by using the entire system, but not surpassing 0.02%. Temperature is considered one of the key factors that can accelerate the transistor aging [1], which can degrade the performance and reduce the expected lifetime. It is worth noting that there is also the leakage current problem, and in [7] it is explained that the leakage current of the SPARC64 VIIIfx CPU was significantly decreased by adopting long gate transistors, and by lowering the junction temperature to 30^oC via chilled water cooling. For this purpose, the CPU water cooling temperature was maintained around 15 ± 1^oC since the beginning of the K computer's operation.

Except during the large-scale maintenance period, the set of chiller systems has produced cold water at around 10^oC to be used for the Heat Exchangers (HEX) for cooling down the CPU cooling water, and for the Air Handling Units for cooling down the air for the memory modules (DRAM), power supply, and other

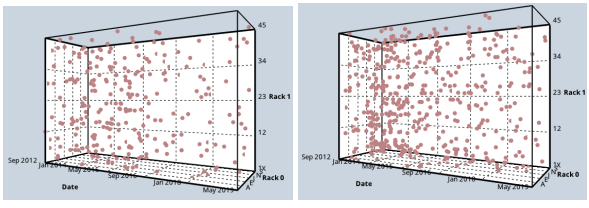

Figure 1: Spatio-temporal distribution of the critical failures of the CPU (Left) and DRAM (Right) per compute rack.

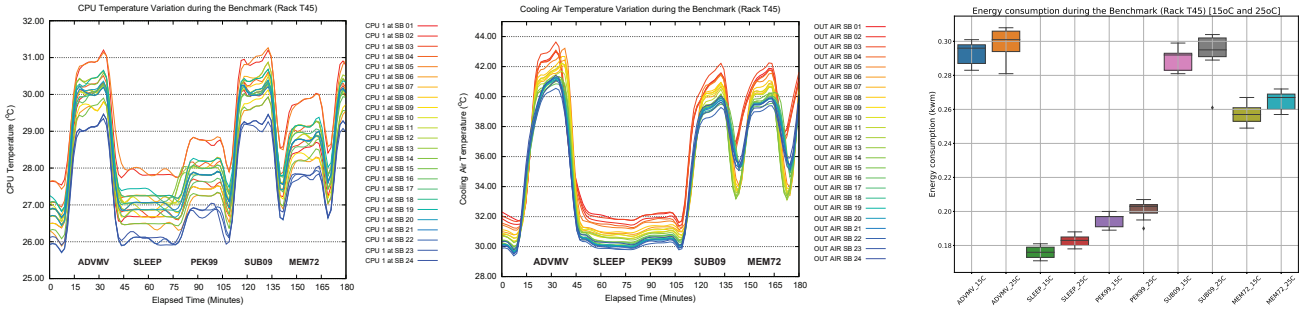

Figure 2: Temperature effects on the energy consumption. CPU temperature (Left), cooling air temperature (Center), and energy consumption (Right) when running the benchmark applications of ADVMV, SLEEP, PEK99, SUB09, and MEM72.

components inside the racks. A detailed schematic design of the K computer cooling system is presented in [3], and the Facility Operations and Development Unit has been operating the entire cooling system by using different combinations of chillers and controlling the flow of the cooling water and air. To better understand the possible impact of the temperature on the hardware failures and energy consumption, we applied post-hoc visual log data analysis for the former, and also a set of benchmark applications on a single compute racks by using increased CPU water cooling temperatures for the latter. Some findings and impressions are presented in the following subsections.

2 HARDWARE FAILURES

Figure 1 shows the spatio-temporal distribution of the accumulated failures of the CPU and DRAM per compute racks [2]. We only took into consideration the hardware failures that required maintenance intervention and hardware component substitution. The maximum number of accumulated CPU failures in the racks was three, and for the DRAM was five. It is worth noting that inside a single compute rack, there are 96 CPUs and 768 DRAM modules, and we could not verify a clear correlation between the CPU and DRAM failures. As shown in the Figure 2, the CPU temperature has increased no more than 5^oC even when running CPU intensive applications, and the airflow used to remove the heat from DRAM and other components on the SB (System Board) increase no more than 12^oC even when running memory intensive applications. It is also worth noting that there also exist some compute racks that there were not affected by CPU or DRAM failures during the entire period of the operation.

3 ENERGY CONSUMPTION

We utilized a single compute rack (96 nodes) of the K computer to verify the impact of the CPU cooling water temperature on the energy consumption. We used the low-priority "Micro" class job, with a maximum of 30 minute running time, which is used to maximize the usage of idle compute racks. We utilized a compute rack (T45) with an attached power monitoring and logger device, to run a set of benchmark applications with different behaviors: SLEEP (Do nothing); PEK99 (CPU intensive); MEM72 (Memory intensive); SUB09 (CPU/Memory balanced usage); and ADVMV (Kernel from a production grade application). Figure 2 shows the CPU temperature (Left) and the cooling air temperature (Center) inside the compute rack T45 during the job running when using the CPU cooling temperature at 25^oC. The cooling water temperature was changed by managing the primary loop water flow with a chilled water at 10^oC. The graph at the right side shows the boxplot graph of the power consumption distribution of these five benchmark applications, when using CPU cooling temperature at 15^oC and 25^oC. We could verify that the increase in energy consumption was less than 5%. We could also observe that even increasing the CPU cooling water temperature in 10^oC, the CPU running temperature rarely surpasses 30^oC, and a low impact on the energy consumption. As a result, we can think that a wider fluctuation on the CPU cooling water temperature could be allowed without much consequences during the operation.

Application	15^oC	25^oC	Increase
SLEEP	0.176	0.182	3.79%
PEK99	0.194	0.201	3.32%
SUB09	0.288	0.295	2.14%
MEM72	0.257	0.265	2.90%
ADVMV	0.292	0.299	2.50%

4 CONCLUSION

Although more deep evaluation is required, we could observe in practice that the theoretical benefits of using low cooling water temperature (15 ± 1^oC) when running the K computer. From this small set of experiments, we could observe that even increasing the CPU cooling water temperature in 10^oC, it may still allow the hardware to operate within specification with limited impact on the energy consumption and hardware failure rate. We expect that the obtained knowledge will be helpful for the decision support and operation planning of the new supercomputer Fugaku.

ACKNOWLEDGMENTS

Part of the results was obtained by using the K computer at the RIKEN R-CCS. We thank the colleagues at the RIKEN R-CCS, and also the staffs from Fujitsu for their supportive assistance and helpful discussions. We also thank some ex-colleagues who worked on the benchmark application codes utilized in the experiments.

REFERENCES

[1] J. Keane and C. H. Kim. 2011. An odomoeter for CPUs. *IEEE Spectrum* 48, 5 (2011), 28–33.

[2] Kazuki Koiso, Naohisa Sakamoto, Jorji Nonaka, and Fumiyoshi Shoji. 2018. A Transfer Entropy Based Visual Analytics System for Identifying Causality of Critical Hardware Failures Case Study: CPU Failures in the K Computer. In *Proceedings of the 18th Asia Simulation Conference (AsiaSim 2018)*.

[3] Motohiko Matsuda, Hiroya Matsuba, Jorji Nonaka, Keiji Yamamoto, Hiroshi Shibata, and Toshiyuki Tsukamoto. 2019. Modeling the Existing Cooling System to Learn its Behavior for Post-K Supercomputer at RIKEN R-CCS. In *Proceedings of the EE HPC State of the Practice Workshop 2019*. 1–10.

[4] Hiroyuki Miyazaki, Yoshihiro Kusano, Naoki Shinjou, Fumiyoshi Shoji, Mitsuo Yokokawa, and Tadashi Watanabe. 2012. Overview of the K computer system. *Fujitsu Scientific and Technical Journal* 48, 3 (2012), 255–265.

[5] Fumiyoshi Shoji, Shuji Matsui, Mitsuo Okamoto, Fumichika Sueyasu, Toshiyuki Tsukamoto, Atsuya Uno, and Keiji Yamamoto. 2015. Long term failure analysis of 10 Petascale supercomputer, ISC 2015: HPC in Asia (Poster).

[6] Keiji Yamamoto, Atsuya Uno, Hitoshi Murai, Toshiyuki Tsukamoto, Fumiyoshi Shoji, Shuji Matsui, Ryuichi Sekizawa, Fumichika Sueyasu, Hiroshi Uchiyama, Mitsuo Okamoto, Nobuo Ohgushi, Katsutoshi Takashina, Daisuke Wakabayashi, Yuki Taguchi, and Mitsuo Yokokawa. 2014. The K computer Operations: Experiences and Statistics. *Procedia Computer Science* 29 (2014), 576 – 585. 2014 International Conference on Computational Science.

[7] Toshio Yoshida, Mikio Hondo, Ryuji Kan, and Go Sugizaki. 2012. SPARC64 VIIIfx: CPU for the K computer. *Fujitsu Scientific and Technical Journal* 48, 3 (2012), 274–279.