

High Performance Computing (HPC), Supercomputing  
Intel® Xeon® CPU Max Series, Intel® Xeon® Scalable Processors

## Kyoto University enhances HPC performance to advance research and development

**Academic Center for Computing and Media (ACCMS) maximizes the Intel® Xeon® Max CPU series to enhance the speed and efficiency of scientific investigations, resulting in accelerated research outcomes.**

“We required a user-friendly CPU for applications in the Kyoto University system, meaning high B/F value, x86 CPU with DDR5, and large memory x86 system. And based on our research, there were no CPUs apart from Intel® Xeon® CPU Max series that meet our requirements.”

—Keiichiro Fukazawa  
Associate Professor, Computing  
Research Department, ACCMS,  
Kyoto University



京都大学  
KYOTO UNIVERSITY



Working at the forefront of technological innovation and academic excellence, Kyoto University is home to the Academic Center for Computing and Media Studies (ACCMS), a dynamic hub for advanced computation and media research. Over the years, the ACCMS has been a focal point for advancing groundbreaking scientific research and development through its state-of-the-art computational resources. This ongoing evolution reflects a commitment to staying at the forefront of technological advancements, with the aim of pushing the boundaries of computational research.

However, amidst these advancements, it becomes apparent that certain challenges persist. Many of the intricate simulation codes, integral to the research processes, face constraints associated with memory bandwidth within the existing High-Performance Computing (HPC) resources. In simpler terms, the speed at which these codes can read from or write to memory emerges as a limiting factor, impacting their overall performance.

This memory-bandwidth limitation presents an ongoing challenge for computational researchers. To maximize the performance of their codes, they were looking to seek ways to optimize them to work within these constraints. This could involve using memory more efficiently, adjusting code to better match the memory hierarchy of the HPC systems, or even developing new algorithms and techniques that are less reliant on memory bandwidth.

It is this continuous quest for optimization that led Kyoto University to collaborate with Intel in updating their supercomputing systems. Built with the latest Intel® Xeon® CPU Max series, these new systems are designed to meet user requirements for exceptional high-performance memory bandwidth, expansive memory capacity, and optimal parallel performance within a well-balanced High-Performance Computing (HPC) infrastructure.

### The Need for Enhanced Computing Resources

For effective scientific research, accelerating the production of results is a constant objective. Beyond the requirement for quicker run times, researchers grapple with a host of key challenges. These hurdles span a broad spectrum, from data management and analysis complexities to the need for advanced computational resources. Understanding and addressing these challenges is crucial for researchers to expedite scientific discoveries and accelerate innovations.

The rapid advances in computing have led to an increasing demand for faster and larger calculations. For Kyoto University, as the complexity and computational demands grew, the need for more powerful resources became apparent. These advanced resources were essential in providing the computational power needed to deliver results in shorter times and accelerate research productivity. But, apart from the requirement for faster run times, another key challenge is the need for more memory bandwidth per node as the workloads increase and the need to handle tasks at larger scales arises.

“In the pursuit of advancing research outcomes, researchers consistently strive for expedited and more extensive program execution, explains Keiichiro Fukazawa, Associate Professor, Computing Research Department, ACCMS, Kyoto University. “Their specific requirement revolves around the necessity for a generous allocation of memory per node,” he adds.

## Finding the Solution

The emphasis on faster processing and larger memory capacities to obtain scientific results quickly reflects the continuous drive within the research community. This is important to enhance the efficiency and capabilities of computational systems, to tackle increasingly intricate and data-intensive scientific challenges. This was exactly what Kyoto University was looking for when they started the process for updating ACCMS HPC systems.

Professor Fukazawa explains, “A few years earlier, our main HPC system was configured around the Intel® Xeon® Phi 7250. This setup included 16GB of MCDRAM, and despite its peak performance of 3TFlops, the bandwidth was approximately 400GB/s per unit, resulting in a B/F value of 0.1333. This value represented a higher bandwidth compared to the DDR4 memory at that time. However, roughly five years post the implementation of the previous system, advancements in computer technology have led to an increased demand for swifter and larger computations. Furthermore, on the Intel® Xeon® Phi, we also observed performance degradation with non-vectorized applications, which seemed to stem from an issue with the CPU core.”

In the context of HPC and supercomputing, where large-scale simulations and complex computations are common, optimizing the performance of a system involves considering both computational power and data transfer efficiency. The Bytes/Flop value (B/F value) provides insights into how well a system utilizes its computational resources by indicating the amount of data movement required for each floating-point operation. Taking this into consideration, Professor Fukazawa emphasized the fact that the need of the hour was to enhance their current systems with a CPU that comes with high B/F value



So, under the guidance of Professor Fukazawa, ACCMS undertook the task of designing new systems that incorporate cutting-edge technologies to address the vital need for higher memory bandwidth performance in high-performance computing (HPC). The core design principle of these systems revolved around a three-configuration system. “We initiated the three-configuration system three generations ago, focusing on the power of many-core processing with high memory bandwidth for our code, general-purpose use, and large memory requirements”, Professor Fukazawa states.




Figure 1: ACCMS three-configuration system

He further elaborates on the goal behind this strategy, “Out of the three system types (Figure 1) - System A (Camphor 3), System B (Laurel 3), and System C (Cinnamon 3) – Camphor 3 is predominantly used by a majority of users. These users are primarily engaged in research that involves the use of custom-built applications for scientific computations. Indeed, many users prefer to use applications created during this period in their original form. In essence, more than 80% of applications running on Camphor 3 necessitate a high B/F value. Therefore, a CPU that could cater to these demands was required.”

### Maximizing the Power of Intel® Xeon® Processors for Optimal Performance

To arrive at the new designs, Professor Fukazawa and his team researched the latest technologies and ran benchmarks of the latest Intel® Xeon® CPU Max series. This was crucial as Intel® Xeon® processors provide max value, particularly in CPU performance. The Intel® Xeon® CPU Max Series supercharges Intel® Xeon® Scalable processors with high bandwidth memory (HBM) and is architected to unlock performance and speed discoveries in data-intensive workloads, such as modeling, artificial intelligence, deep learning, high performance computing (HPC) and data analytics.

Intel® Xeon® CPU Max series is optimized to leverage the broad range of software ecosystem that includes compilers, math libraries, open-source applications, and more. Another key advantage is that Intel® Xeon® CPU Max series provides a seamless experience and enables the best performance on a variety of the workloads. Apart from the performance benefits, Intel® Xeon® CPU Max based systems come with HBM support for enhancing the overall performance and accelerating the research journey. This means that researchers can focus on their true research and don’t need to spend much time for coding and optimization.

Explaining the reason behind choosing Intel® Xeon® CPU Max series, Professor Fukazawa states, “We required a user-friendly CPU for applications in the Kyoto University

system, meaning high B/F value, x86 CPU with DDR5, and large memory x86 system. And based on our research, there were no CPUs apart from Intel’s that meet our requirements.” He further adds, “And, when renewing the system, it was essential to choose a configuration with the highest possible B/F value. At that point, the choice had to be a CPU equipped with HBM memory, and the choices were inevitably narrowed down to two options: Intel® Xeon® CPU Max series and another CPU. However, when considering the computational performance, even if we used a CPU that was a different option, the computational performance would only be about half of the Intel® Xeon® CPU Max, making the Intel® Xeon® CPU Max the ideal choice for our requirements.”

“On the other hand, although Laurel 3 does not have as high a need for a B/F value as Camphor 3, there were still a train of thoughts that emphasized on the need for a wider memory bandwidth. Therefore, we started considering the possibility of using DDR5, but at the time of consideration there were few CPU options that officially supported DDR5 and could meet the required performance, so we benchmarked from several candidates and found the results. In the end, 4th Gen Intel® Xeon® Scalable processors were chosen, taking into consideration the timing of procurement and other factors.” Professor Fukazawa adds.

### Delivering Impactful Results

With the deployment of the new systems, Kyoto University has started seeing significant performance gains. According to Professor Fukazawa, “Based on comparisons made with the previous generation systems, Camphor 3 is already achieving an average speed increase of 4.7 times<sup>1</sup>, and Laurel 3 is achieving an average speed increase of 3.7 times<sup>1</sup>.” (Figure 2 and Figure 3)

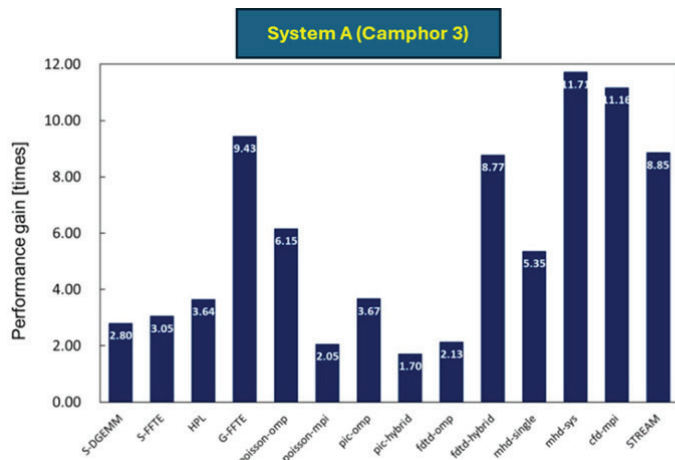


Figure 2: Camphor SPR+HBM vs KNL system comparison

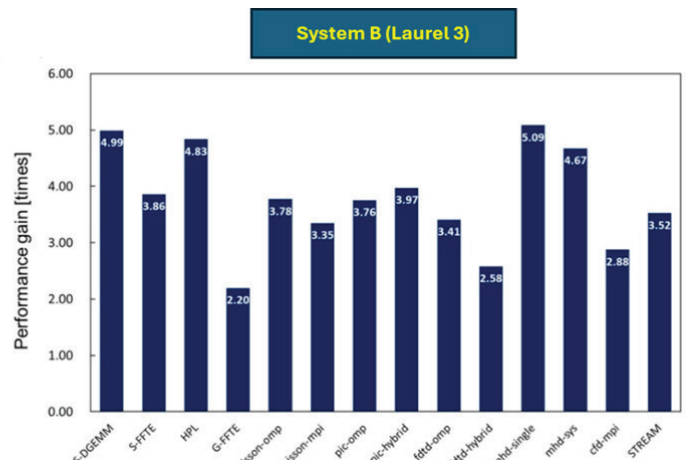


Figure 3: Laurel SPR+DDR vs Broadwell system comparison

Emphasizing on how Intel® Xeon® CPU Max played a key role in driving this performance advantage, Professor Fukazawa says, “At the center, we are conducting research called program coding support joint research. This involves receiving the user’s code, optimizing it over a period of about a year, and then returning it to the user.

In the case of Intel® Xeon Phi, it is often necessary to optimize the application accordingly in order to bring out the performance. However, for the Intel® Xeon® CPU Max series, Intel’s compiler and Math Kernel Library used as is, performance can be easily extracted without any special optimization.”

To drive this point further, Professor Fukazawa provides a couple of examples of how the Intel® Xeon® CPU Max series plays a crucial role in advancing research and development, “In one of the projects that I recently took up, I was working on an application that involves solving the global Jovian magnetosphere with 3D MHD simulation. This is a high B/F value application, and due to the huge size of the magnetosphere with small grid spacing, it takes more than a year to observe the time evolution. However, with Intel® Xeon® CPU Max series, I was able obtain results at more than twice the speed. Another example is the general circulation model (GCM) that is used to study the effects of global warming by running simulations with various parameters. GCM is also a high B/F application, allowing researchers to benefit from Intel® Xeon® CPU Max series. The application helps in performing simulations with multiple parameters and can identify parameters that may mitigate global warming.”

## Getting Ready for the Future Ahead

Kyoto University’s ACCMS stands as a beacon of academic excellence and innovation, committed to forging the path for Japan’s leadership in advanced research. With an unwavering dedication to the pursuit of knowledge and breakthroughs, ACCMS is poised to continue its impactful academic research across various disciplines.

In alignment with this commitment, Intel, a global technology leader, is positioned to play a crucial role. By extending robust technological support to the ACCMS, Intel aims to amplify the impact of academic research, facilitating the exploration of new frontiers in science, technology, and

various fields. This collaborative endeavor signifies not only a commitment to the advancement of academic pursuits at Kyoto University but also a broader dedication to contributing to the overall development and progress of diverse fields within Japan and beyond. Through this collaboration, Intel is committed to being a catalyst for transformative advancements, fostering innovation, and shaping the future of academic research in multiple domains.

Reiterating the need to drive this collaboration further and achieve the milestones set by Kyoto University, Katsumi Yazawa, Director of HPC Business Development, Industry Business Unit, Intel Japan signs off, “We fully understand that the B/F value is important in the HPC market. However, solutions that implement HBM and achieve high memory bandwidth will inevitably be expensive. Therefore, at Intel, we understand the requirements and are considering various new memory technologies. In the near future, we hope to improve the B/F value by offering MCR-DIMMs that have the same form factor as DDR but can achieve nearly twice the memory bandwidth. As a trusted advisor to Kyoto University, Intel is always looking to enhance our collaboration and look forward to providing the roadmap for meeting the requirements of HBM solution on HPC/AI in our long-term relation.”



<sup>1</sup>Source: Information Processing Society of Japan Research Report Vol.2023-HPC-192 No.12.

Performance varies by use, configuration and other factors. Learn more at [www.Intel.com/PerformanceIndex](http://www.Intel.com/PerformanceIndex).

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

For workloads and configurations visit [www.Intel.com/PerformanceIndex](http://www.Intel.com/PerformanceIndex). Results may vary. Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy. Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.