# A Custom Change

AMD OPTERON™ PROCESSORS BRING JAPAN'S HPC MARKET INTO THE X86 FOLD.

**BY PAUL KALLENDER**

When Dr. Satoshi Matsuoka, director of the Tokyo Institute of Technology (Tokyo Tech) Global Scientific Information and Computing Center, wanted to replace the university's NEC SX-5 and SGI Origin 2000 supercomputers in the fall of 2004, he had a tall order to fill. Tokyo Tech needed the computing power to handle the demands of 10,000 staff and students as well as academic and industrial collaborators around the globe.

Matsuoka planned to replace several legacy systems, including an Origin 2000 with 128 R12000 nodes and an HP GS320 with 64 Alpha nodes. The technical specifications for the new-generation system were demanding: It had to handle peak performance above 40 teraflops and run at an annual cost of $8 million to $9 million within a 1-megawatt power budget.

"We were also looking for a system that could fit into fewer than 80 racks in less than 350 square meters—the available space in our center—as well as offer strong continuity with PC platforms, x86 and Linux. And of course, memory and bandwidth needed to be as high as possible," he says.

## AMD OPTERON™ PROCESSORS GET THE CALL

Had Matsuoka required this type of system at the beginning of the millennium, he probably would have had to commission a series of custom processors, which was the common practice in the Japanese high-performance computing (HPC) market through the mid-2000s. By 2006, however, Matsuoka was able to meet his specification requirements by employing a solution based on x86 AMD Opteron processor technology.

The result is the TSUBAME supercomputer (named after the Japanese word for the swallow, the institute's mascot). The system comprises 10,480 AMD Opteron processor cores in Sun Fire x64 servers (specifically, 655 nodes of Sun Fire X4600 with 2.4-GHz Dual-Core AMD Opteron™ processor Model 880 × 8 sockets and 16 cores per node), ClearSpeed coprocessors, InfiniBand interconnects and more than 21 terabytes of memory.

PHOTOGRAPH: KATJA DE BRUIJN, ICONOGENIC

The TSUBAME became the fastest supercomputer in the Asia-Pacific region in June 2006, weighing in at 85 teraflops at its peak. After beating the Earth Simulator supercomputer at the Earth Simulator Center in Kanazawa, Yokohama, the TSUBAME was acknowledged as the seventh-fastest supercomputer in the world in the 27th Top 500 project list in July 2006. The system continues to serve as the core computing resource on the Tokyo Tech campus.

The combination of the AMD Opteron™ processor's performance features provided a compelling solution, says Matsuoka. He particularly noted the combination of the memory bandwidth and performance, AMD Smart Fetch technology, the AMD Wide Floating Point Accelerator and the AMD Balanced Smart Cache.

"It was the combination of the coherent HyperTransport™ Technology and the embedded memory controller, allowing for high memory bandwidth and fat-node socket scalability, that was the most compelling," says Matsuoka. The value of the TSUBAME is not only that it is one of the world's leading supercomputers but also that it allows for the coexistence of capacity and capability, he adds.

Matsuoka's experience with the Dual-Core AMD Opteron processor has been repeated at many other leading institutions in Japan's academic HPC community as they have abandoned custom solutions in favor of 65-nanometer Quad-Core AMD Opteron processor-based x86 systems, says Hiro Yamano, senior manager, Enterprise Product Marketing, at AMD Japan Ltd.

## CUSTOM SOLUTIONS FALL FROM FAVOR

Since 2005, with the advent of the Quad-Core AMD Opteron processor, there has been a shift away from RISC-based and custom-tuned proprietary solutions to x86-based cluster computing underpinned by AMD Opteron processors in both the mid- and high-end HPC segments. Within them, the AMD Opteron processor family has scored major successes, including the adoption of a 512-core PC cluster system by Doshisha, a leading private university in Kyoto, Japan, and the T2K Open Supercomputer by three of Japan's leading institutions: the University of Tsukuba, University of Tokyo and Kyoto University.

"We have excelled in the Japanese HPC market, and the key words are *price, performance* and *performance-per-watt*," Yamano says.

Professor Tomoyuki Hiroyasu, deputy director of the Department of Medical Information, Faculty of Life and Medical Sciences at Doshisha University, agrees. In opening the Faculty of Life and Medical Sciences near Kyoto, in April 2008, the university needed an HPC environment that could cope with a multitude of medical science and basic biology research needs and crunch an enormous amount of data on biophenomena.

Doshisha was experienced with PC cluster systems, and the university had already been using AMD Opteron processors in its 512-CPU-based, 1.169-teraflop "Super Nova" cluster since 2003. For its replacement, university officials opted for 64 Dell PowerEdge SC1435 servers housing two Quad-Core AMD Opteron processors running at 2.3 GHz for its new "MISC" PC cluster system.

"The major reasons we chose this solution were, first and foremost, the memory bandwidth and the system's ease of use," Hiroyasu says.

In addition to the superior memory

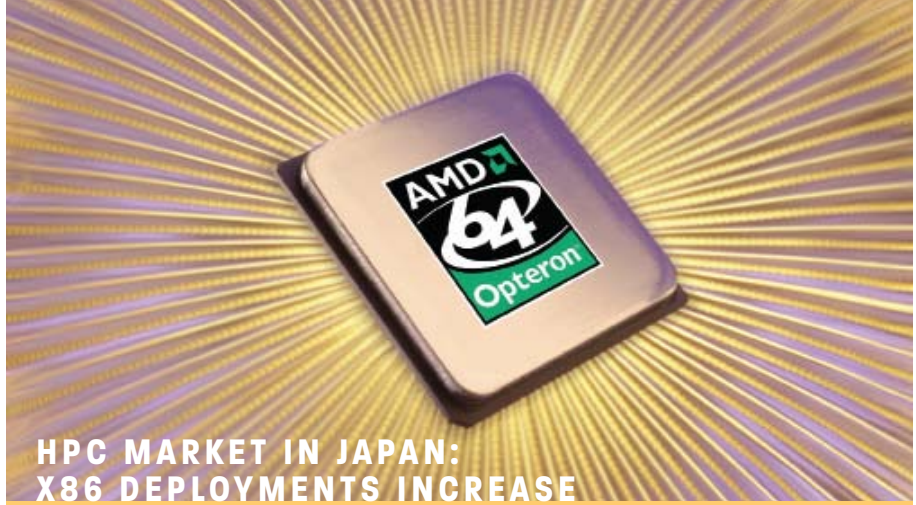


## HPC MARKET IN JAPAN: X86 DEPLOYMENTS INCREASE

AMD Opteron™ processors captured 49.1 percent of Japan's x86 HPC server market in 2006, according to Ikuo Takafuji, research analyst, Servers, at IDC, underscoring the fact that x86-based solutions are rivaling the more traditional, custom-based solutions that are popular in Japan's HPC market.

Japan's HPC market differs from the U.S. market, in that it is slightly more limited in scope, says Takafuji. In the U.S., HPC use has diffused into, for example, the financial and commercial bioscience, geosciences and geoengineering sectors. In Japan, on the other hand, demand is focused mainly on academic research, electronic design automation and computer-aided engineering. Takafuji continues, "In the U.S., x86 is more dominant in the HPC market than it is in Japan, although the Japanese market has been changing in that the RISC-based market has been shrinking. And in Japan, the government and academic sector has been embracing x86, so this sector provides 50 percent, 60 percent or even 70 percent of the x86 HPC market in Japan.

"The reason for the increase in x86-based systems and users in the Japanese HPC market is very simple—the systems' superior cost performance over custom solutions," Takafuji says. "Basically, x86 is very open and the price is very reasonable," he adds. According to IDC, the x86 server market share for HPC in Japan rose from 3.7 percent in 2002 to 36.6 percent in 2008, to a projected 42.5 percent in 2009.

"What has been so appealing to the HPC market is the superior memory bandwidth of quad-core. HPC computing has been faced by memory bandwidth bottlenecks," Takafuji says. "In particular for the TSUBAME, the Sun Microsystems servers based on AMD Opteron processors are very strong. The most important thing for selecting HPC servers is the IPC [instructions per clock] for CPUs," he concludes.

AMD ARCHIVE

bandwidth and ease of use, other considerations were the processing performance supplied by AMD Direct Connect Architecture, supported by the latest HyperTransport Technology, and the integrated memory controller, allowing for high-speed computation irrespective of factors such as complex memory distribution in a program.

MISC is theoretically capable of 4.71 teraflops, even with the AMD Opteron™ processors using up to 1 amp per node less than other processor types.

"The difference in power and performance, combined with the reduced power consumption and the space savings we achieved in a two-rack solution, was compelling," he says.

## T2K OPEN SUPERCOMPUTER

The T2K system comprises three PC clusters based on Quad-Core AMD Opteron processors formerly code-named "Barcelona," with specifications developed jointly by the University of Tsukuba's Center for Computational Sciences, the University of Tokyo's Information Technology Center and Kyoto University's Academic Center for Computing and Media Studies.

In July 2006, the three universities began jointly developing common specifications so the next-generation supercomputer would be able to use an open source hardware architecture and open source system software to make the supercomputers widely available among the three institutions. The Quad-Core AMD Opteron processor-based solutions met the specifications.

Kyoto University's 50.5-teraflop (LINPACK) system is built around a subsystem of 416 Fujitsu HX600 HPC server nodes, each with four AMD Opteron 8356 Series processors. The system, says Professor Hiroshi Nakashima of Kyoto University's Academic Center for Computing and Media Studies, has more than lived up to expectations for the more than 80 user groups conducting research with it. The research includes oceanography, quantum electrodynamics, molecular chemistry, colloid physics, plasma physics, electronics, fundamental and structural mechanics, computational neurology and natural language processing.

"What impressed us most at Kyoto was that the Quad-Core AMD Opteron processor offered a genuine quad-core solution with 4-flop-per-cycle throughput for each core—that is to say excellent peak performance per socket," he says.

The Quad-Core AMD Opteron processor-based solution met the university's needs, Nakashima says, because it provided high overall performance at a reasonable cost.

In actual use, Nakashima found the power consumption of the Quad-Core AMD Opteron processor-based solution to be 400 kilowatts, or about 20 percent below what the university required, saving it up to $100,000 a year in electricity bills.

"The other notable fact was that we got many users from outside of traditional numerical computing. For example, my heavy user list includes natural-language-processing people who have never used conventional supercomputers but now find that our T2K Open Supercomputer is a powerful tool for their research work," he says.

In the case of the University of Tokyo, the university replaced its Hitachi SR8000/MPP supercomputer with 952-node Hitachi HA8000-tc/RS425 technical servers with four Quad-Core AMD Opteron processors per node, says Yutaka Ishikawa, a professor at the Computer Science Department/Information Technology Center at The University of Tokyo. With a theoretical peak performance of approximately 140 teraflops, it had the potential to be the fastest in Japan when it launched, on June 1, 2008. The published LINPACK performance is 82.98 teraflops using 512 + 256 nodes. But in fact, "the performance is what we expected, the system is more stable than we expected and the electricity cost is less than we expected," he says.

## POWERFUL COMPUTING FOR ALL?

Beyond narrow performance criteria and the AMD Opteron processor's advantages, the question remains: Can huge clusters of commercial processors help deliver the simulations, projections, models and breakthroughs their communities of users need? Can one size fit all?

The answer is yes, says Dr. Taisuke Boku of the University of Tsukuba's Department of Computer Science Graduate School of Systems and Information Engineering. Boku's research into lattice quantum chromodynamics (the study of the interactions of quarks and gluons) requires an extremely large system—one with more than 10,000 cores.

"On our site, the large system size with more than 10,000 cores greatly helps several heavy users. These successful results lead the users in our center and other nationwide institutes to the next step of research," he says.

Matsuoka concurs, saying that Tokyo Tech's TSUBAME works for novices and experts, for traditional as well as non-traditional applications, something that underlines AMD's credo of doing what's right for all technology user communities and customers.

Fusing technology with might? You bet! ■

**"WHAT IMPRESSED US MOST AT KYOTO UNIVERSITY WAS THAT THE AMD OPTERON™ PROCESSOR OFFERED A GENUINE QUAD-CORE SOLUTION WITH 4-FLOP-PER-CYCLE THROUGHPUT FOR EACH CORE—THAT IS TO SAY EXCELLENT PEAK PERFORMANCE PER SOCKET."**

—PROFESSOR HIROSHI NAKASHIMA
KYOTO UNIVERSITY'S ACADEMIC CENTER FOR COMPUTING AND MEDIA STUDIES

related stories

**HPC in the Mainstream**

**Lawrence Livermore's Extensive Computing Needs Results in Unique Collaborative Program**