Case study
Intel® Xeon Phi™ Coprocessor
High-Performance Computing
Education

Look Inside.™

Building a supercomputer with Intel® Xeon Phi™ coprocessors

University of Tsukuba builds a supercomputer capable of 1 PFlops with Intel® Xeon® processor E5 v2 family and Intel Xeon Phi coprocessor

筑波大学
University of Tsukuba

Center for Computational Sciences, University of Tsukuba

Location: Tsukuba Campus, University of Tsukuba, 1-1-1, Tennodai, Tsukuba-city, Ibaraki Prefecture, Japan
Details of Business: Promotion of interdisciplinary computational science, development of ultra-high-speed computer systems and network technology, ultra-high-speed simulation and large-scale data analysis, applied research in information technology, training human resources responsible for computational science, and managing the Interdisciplinary Joint-Usage program
http://www.ccs.tsukuba.ac.jp/

Challenges

- Expanding its high-speed system to support interdisciplinary computational science, a fusion of science and computer science.

Solution

- An Intel® technology-based supercomputer called COMA (PACS)-IX with 768 Intel® Xeon® processors E5 v2 family, Intel® Xeon Phi™ coprocessors, and Intel® Cluster Studio XE for MPI developers and C++/ Fortran programmers

Benefits

- Supercomputer COMA (PACS-IX) achieved 1.001 PFlops peak performance with Intel® architecture
- Optimized code that performs better on the many-core coprocessor
- Promote research focusing on the latest hybrid computing architecture

The second PFlops supercomputer at the Center for Computational Sciences, University of Tsukuba

The Center for Computational Sciences at the University of Tsukuba promotes interdisciplinary computational science, a fusion of science and computer science. In April of 2014, it introduced COMA (PACS-IX), a supercomputer with peak computing performance of 1.001 PFlops. The ultra-parallel cluster is composed of 393 nodes. Each node contains two Intel Xeon processors E5 v2 family and two Intel Xeon Phi coprocessors. It is a matter of pride for the center to achieve the highest domestic performance (as of June 2014), with peak computing performance over 1 PFlops.

With COMA and the existing HA-PACS, the Center for Computational Science at the University of Tsukuba now has two supercomputers for educating undergraduate and graduate students in the field of high-performance computing.

Scientific research running on the supercomputer promotes supercomputer development

With the Center for Computational Physics that was founded in 1992 as its predecessor, the current Center for Computational Sciences restructured and expanded in 2004. In 2010, it was certified by the Joint Usage/Research Center, Ministry of Education, as AISCI (Advanced Interdisciplinary Computational Science Collaboration Initiative). Currently, under the Multidisciplinary Joint-Use Program that strengthens the collaboration of computer science and science, it acts as the computer center for research institutes across the country, supports workshops for promoting interdisciplinary computational science, and offers resources for researchers and students for developing extra-scale supercomputers.

Taisuke Boku, deputy director and chairman of computer system operation, explained, “The role of the Center for Computational Sciences is not just as the information center where researchers can

The many-core Intel® Xeon Phi™ coprocessor has contributed to the development of interdisciplinary computation science by combining science and computational science

"A large number of CPU cores are mounted on a single chip, and are attached to the CPU via the PCI Express* bus in the Intel® Xeon Phi™ coprocessor to achieve optimum cost performance, power performance and space performance on the supercomputer."

- Professor Taisuke Boku
Computing Systems Application Committee Head/
Assistant Center Head
Center for Computational Sciences, University of Tsukuba

access supercomputers. Being a unique national university, it serves many teachers and researchers from various subjects like computational science, application development, and scientific research. This sharing promotes the development of the supercomputer itself, something unique among university computer centers. These initiatives, defined as interdisciplinary computational sciences, are a major feature of the Center for Computational Sciences. It is a place for broad exploration in computational sciences."

The Center for Computational Sciences also promotes the use of information technology in basic science, high-speed simulation, large-scale data analysis, and applied research. With HA-PACS and COMA bringing PFlops computing to reality, it is looking forward to breakthroughs in research and development of computational sciences in the fields of particle physics, space, nuclear physics, material science, life, and global environment.

Evaluating the high cost performance and the power performance and adopting of Intel® Xeon Phi™ coprocessor

In recent years, CPU performance has helped increase the performance of massively parallel supercomputing clusters. However, the conventional way of increasing performance by adding nodes and stacking servers takes up limited data center space and power. At the Computational Science Research Center, a course showed how low-power, high-performance supercomputers can be accelerated by arithmetic coprocessors, focusing on many-core processors (MIC) in which multiple cores are integrated in a single chip. The research center evaluated the MIC with Intel, and is continuously researching how to improve tuning and performance. As a result, it decided to build COMA immediately after the Intel Xeon Phi coprocessor was launched.

With up to 61 cores on a single chip, Intel Xeon Phi coprocessor offers up to a Tflop of double-precision peak performance, which delivers higher parallel performance per watt than Intel Xeon processors.[1,2,3]

Another benefit of the Intel Xeon Phi coprocessor, based on the same architecture as the Intel Xeon processor, is its high level of programmability. Boku said, "There comes a sense of security when applications can be developed without wasting the existing codes and without having to learn new things. Just like writing applications for a general-purpose processor, the Intel Xeon Phi coprocessor can be programmed by using FORTRAN, C++ and Open MP."

In 2013, the Computational Science Research Center designed the next supercomputer system in collaboration with the Information Technology Center of the University of Tokyo and has created Joint Center for Advanced High Performance Computing (JCAHPC), an organization for cooperation and management. As the goal for the financial year 2015, JCAHPC is planning to introduce a class of very large supercomputers with scores of PFlops at the Kashiwa campus of the Information Technology Center, University of Tokyo. COMA also plays a role as an experimental system for carrying out the research and development of tuning and coding of the next-generation supercomputer.

Ensuring high-speed performance with a biased MIC board of two units with one CPUMIC and CPU

COMA, introduced by the Center for Computational Sciences, stands for cluster of many-core architecture processor. The term is also derived from the English name Coma Berenices, a typical cluster of galaxies - a collection of stars is a galaxy (many core), and a collection of galaxies (a cluster) is a galaxy cluster.

COMA is a parallel system equipped with 393 compute nodes where one compute node has two Intel Xeon processors E5-2670 v2 (2.50GHz) with 10 cores per CPU and two Intel Xeon Phi coprocessors 7110P with 61 cores. Theoretical peak computing performance of a single node is 0.4 Tflops for CPU, 2.147 Tflops for MIC, and 2.547 Tflops combined. Theoretical peak computing performance of the entire system is 157.2 Tflops for CPU, 843.8 Tflops for MIC, and 1.001 PFlops combined.

In addition, all compute nodes are connected by a mutual high-speed network (InfiniBand* FDR), with a communication performance 10 times that of Ethernet.

Center for Computational Sciences, University of Tsukuba
Professor Taisuke Boku
Computing Systems Application Committee Head/Assistant Center Head

Parallel processing is uniform between compute nodes. A Lustre* file server with a total capacity of 1.5 petabyte by RAID-6 is used for storage. It is freely accessible from all compute nodes via the InfiniBand FDR network.

"Compute nodes consist of two MIC boards equipped with the Intel Xeon Phi coprocessor 7110P and a network board of InfiniBand FDR, next to one Intel Xeon processor E5-2670 v2, ensuring high-speed communication capability between the coprocessor and the network board. Just one Intel Xeon processor E5-2670 v2 can maintain high-speed performance," said Boku.

Different operating modes to support different computing needs

COMA operations can handle three different operating modes at the node level: CPU only, MIC only, and a combination of CPU and MIC.

CPU only

CPU only mode is used to support the multi-core-oriented programs of the previous-generation supercomputer 2K-Tsukuba, which reached its end of life in February 2014. This mode uses 16 of the total of 20 cores in the Intel Xeon processor E5-2670 v2.

MIC only

Under this mode, all computing resources on the Intel Xeon processors E5 v2 and the Intel Xeon Phi coprocessors are used. It runs hybrid programs in the MIC Native or Symmetric mode.

Besides the flexibility of multiple execution models, the architecture of the Intel Xeon Phi coprocessor brought additional advantages. Unlike other GPU architectures, the Intel Xeon Phi coprocessor is compatible with Intel architecture. There is no need to rewrite applications or master new programming models. With a single programming model, parallel optimization techniques simultaneously deliver performance for both the Intel Xeon processor and Intel Xeon Phi coprocessor. Researchers familiar with existing programming tools can immediately take advantage of the performance offered by the coprocessor.

Intel Cluster Studio XE is used as the development tool in the programming environment. Researchers use the Intel® FORTRAN and C++ compilers and the Intel® MPI library as needed to create the programs. Boku said, "Since a free choice of compilers can also be effective for some programs, we have developed an application system in which we do not specify

System Block Diagram

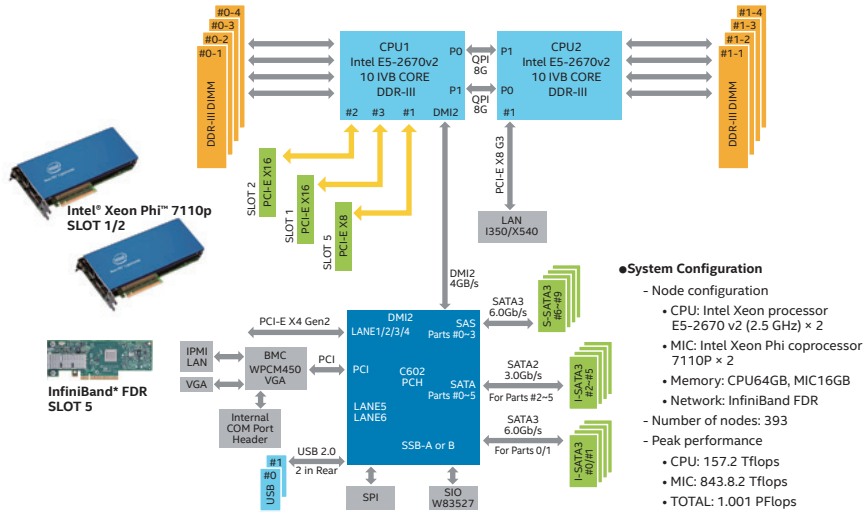

The Center for Computational Sciences, University of Tsukuba, and the Information Technology Center, University of Tokyo, have set up the Joint Center for Advanced High Performance Computing (JCAHPC)

The Center for Computational Sciences, University of Tsukuba, and the Information Technology Center, University of Tokyo, have set up a supercomputer system to be designed primarily by the teaching staff of both institutions in the Tokyo University Information Technology Center located in the Kashiwa Campus of the University of Tokyo. This is the first experiment of its kind in Japan, setting up and jointly operating and managing a supercomputer.

The Center for Computational Sciences, University of Tsukuba, and the Information Technology Center, University of Tokyo, have implemented projects in the past to set joint specifications and procure supercomputers for each university through the T2K Open Supercomputer Alliance, consisting of three centers of the Academic Information Media Center of the University of Tokyo. The specifications of the T2K Open Supercomputer Alliance are intended to reduce procurement costs and improve system performance through freedom from vendor dependence. They are formulated based on openness of hardware and openness from user dependence. The establishment of JCAHPC is one step in this direction, and the design and development as well as the operation and management of the high-performance supercomputers are jointly carried out by the information centers of both universities with the objective of promoting state-of-the-art computational science.

The policy framework established for the design and development of supercomputing systems does not follow existing supercomputing products. Instead, it uses many-core processors to design a state-of-the-art system, also linking the system to the OS, programming language, numerical computing library, and other systems that form the core technology for the software during the design and development process.

a particular product for the compiler or library. Instead, we allow the researchers to freely choose a compiler or library."

Used for development of computational science applications in fields such as particle theory and life science

Researchers across the country can use COMA free of charge under the Interdisciplinary Joint Utilization Program to support advanced scientific research in computational science and computational engineering. Arrangements have also been made for its free use under the High Performance Computing Interconnections (HPCI) Program promoted by the Ministry of Education, Culture, Sports, Science and Technology, and under the Large-Scale General Use Program, which researchers across the country can use free of cost.

Boku said, "There are no limitations on research using the supercomputers. Researchers can choose as they like and, after careful examination of the subject matter of the research, make use of the Interdisciplinary Joint Utilization Program. We hope this computer will be used in every field including elementary particles, space, life science, and global environmental research."

Lectures on high-performance computing give graduate students competence to carry out computations

COMA also helps educate students in the Center for Computational Sciences. According to Boku, "Today, Japan is running overwhelmingly short of researchers in the scientific field who can write their own programs, when compared to the situation abroad. To enable our researchers to write their own computational programs in various fields of research, we are conducting high-performance computation classes for graduate students. We are also running intensive courses introducing high-performance computing to 40 to 50 graduate school students every year and continuing with our initiatives to develop qualified researchers with knowledge of high-performance computing.

"And even in our Global 30 (G30) study program for overseas students, who can only attend classes and obtain credits in English, we are planning to conduct classes on supercomputers such as COMA and HA-PACS directed at these overseas students in order to offer advanced educational services," Boku said.

"In the Center for Computational Sciences, we plan to continue to use COMA in various fields of research and study the performance of many-core processors to gain know-how and knowledge to help achieve high-performance computing," Boku continued. "In addition, we are planning to continue with research and performance evaluation to achieve the planned launch in 2105 of a supercomputer using a many-core processor, to be jointly operated by the Center for Computational Sciences and the Information Technology Center, University of Tokyo."

Intel will continue to work with the Center for Computational Sciences to support the development of interdisciplinary computational science based on the performance of the Intel Xeon processor and Intel Xeon Phi coprocessor.

Find the solution that's right for your organization. Contact your Intel representative, visit Intel's **Business Success Stories for IT Managers,** or explore the **Intel.com IT Center.**

JPN/1501/PDF/SE/ESS/FN 330732-001 US