

Advancing academic research in Japan with Intel® Xeon Phi™ coprocessor

Intel® Xeon Phi™ Coprocessor
Intel® Xeon® Processor E5 v3 Family
High Performance Computing



Okinawa Institute of Science and Technology Graduate University (OIST)

Campus: 1919-1 Tancha, Onna-son, Okinawa, Japan

Established: 1 November 2011

Faculty: 50 (September 2014)

Researchers: 337, PhD Students: 79 (September 2014)

Overview: World-class research and education, the advancement of science that contributes to Okinawa's economic development, and English-medium education and research

<http://www.oist.jp/>

Challenges

- **Insufficient HPC Environment:** An increase in the number of researchers, students, and laboratories made the existing High Performance Computing (HPC) environment unable to meet demand.
- **Decrease in Performance:** As the demand for high throughput, parallel computing increased, the storage I/O speed in the existing environment became a bottleneck.

Solution

- **A new Intel® architecture-based HPC system:** The 426 node HPC system is made up of the Intel® Xeon® processor E5 v3 family, Intel® Xeon Phi™ coprocessor, Intel® SSD, and Lustre*.

Impacts

- **Increased Performance:** Achieved four times better performance in peak operation performance of entire system, and two times better I/O performance.
- **Increased Use in Research that Uses Parallel Computing:** Intel Xeon Phi coprocessor speeds up the HPC performance meeting the needs of researchers and allows more researchers to have access to HPC for their researches.

HPC-Based Large-Scale Computing Environment That is Vital to Cutting Edge Research

The Okinawa Institute of Science and Technology Graduate University (OIST) is located an hour drive from Naha Airport in Onna, Kunigami District, in the center of the Okinawa main island. The headquarter is alongside Japan National Route 58, which runs along the coast of the East China Sea. It was established in 2011 with the goal of contributing to Okinawa's autono-

mous development and the advancement of science and technology in the world through outstanding education and research related to science and technology. It has invited world-class faculty from the fields of neuroscience, environmental and ecological science, molecular, cellular, and developmental biology, mathematical and computational science, physics, and chemistry to conduct world-class research in 50 research units. Over half of the faculty and students are foreigners who all



Intel® Xeon Phi™ coprocessor accelerates programs developed for Intel® Xeon® family-based platforms with x86 compatibility

“The performance of parallel workloads has been greatly improved by the Intel® Xeon® processor E5-2600 v3 product family, with new optimizations and expanded functions such as the Intel® Advanced Vector Extensions.”

– Okinawa Institute of Science and Technology Graduate University
Chief Information Officer
Tim Dyce

teach and research in English. Since beginning to accept students in September 2012, 20 students from around the world have joined every year, and there were 50 faculty members, 337 researchers, and 79 students as of September 2014.

A large-scale computing environment based on high performance computing (HPC) is essential for OIST to promote cutting edge research. The HPC system comprised of 384 servers (increased to 480 in 2013) was implemented as the graduate university’s second generation HPC cluster in the fall of 2011. Since the number of students, researchers and laboratories has grown significantly, the system was unable to handle the analysis demands. An HPC system equipped with Intel® Xeon® processor E5-2600 v3 product family and Intel Xeon Phi coprocessor was deployed then implemented and named “Sango” (Coral). Operating know-how from the previous system was put to use to achieve a variety of improvements in Sango.

Chief Information Officer, Tim Dyce, who supports research through maintenance of the HPC system, explains, “Learning lessons from our previous system, we greatly improved I/O performance from the previous environment. The previous system was a hybrid of Ethernet and Infiniband*. In order

to optimize system costs, due to the rapidly increasing demand for parallel computing, the new environment was unified under InfiniBand FDR. This unified networking environment has led to dramatic increases in performance and simplified management. The Intel Xeon processor E5-2600 v3 product family was used in the new environment, resulting in substantial performance improvements. The increased processing performance has come not only from the increased core counts of the Intel Xeon processor E5-2600 v3 product family, but also from the new extensions, such as Intel® Advanced Vector Extensions 2 (Intel® AVX2). Many of our research applications are well placed to capitalize on these new extensions, and we expect increasing utilization over time.”

Also, in order to accelerate storage access, Mr. Dyce stated that the new environment unifies the previous General Parallel File System (GPFS*) and Lustre systems into a single high performance Lustre system. Furthermore, Intel SSDs have been adopted for nodes with built-in large capacity memory to provide high performance, low-cost storage, especially for applications with very large numbers of small files.

The Intel Xeon Phi coprocessor has also been newly adopted for this HPC system. The Intel Xeon Phi coprocessor achieves ground-breaking performance in massively parallel applications by supplementing the Intel Xeon processor in the system. Because accelerated parallel processing can be achieved by reusing the source code developed for the widely used Intel® architecture programming model, it is an ideal environment to run more complicated simulations. HPC and Research Computing Engineer Eddy Taillefer says, “There



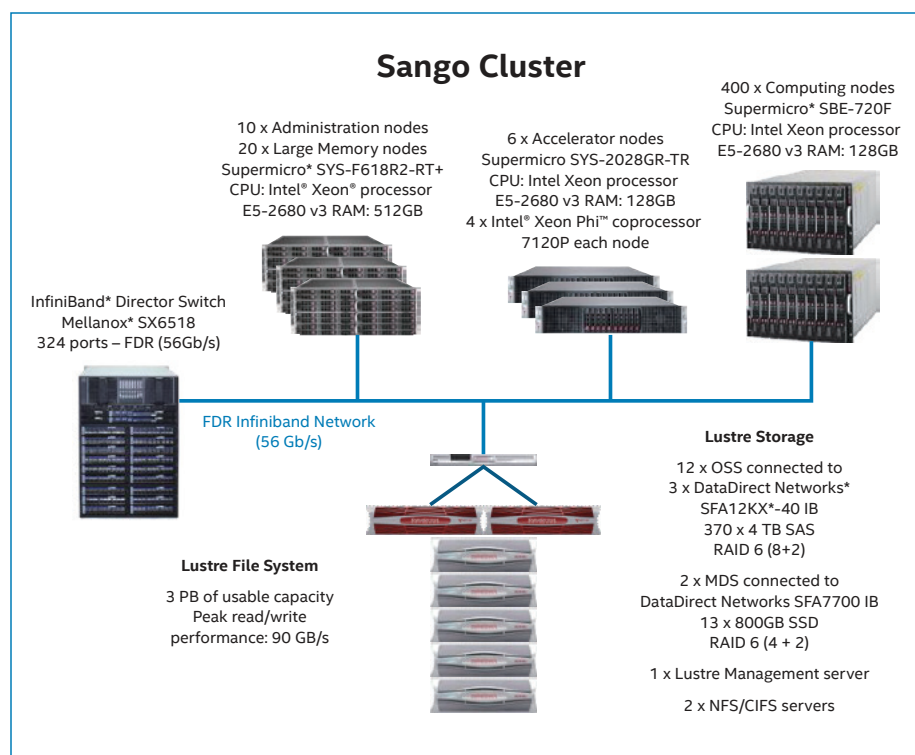
are currently research projects that make use of specialized GPGPU programs to solve large scale problems or to significantly reduce the computation time. However, recently more and more researchers want to achieve similar computation scalability and time reduction while keeping their codes written in C/C++ or Fortran. By taking advantage of OpenMP* automatic parallelization and Intel compiler vectorization acceleration, the highly versatile Intel Xeon Phi coprocessor was able to meet their needs. We also think that the information that we are collecting from actual research computation on Sango will benefit to the integration of future Intel Xeon Phi coprocessor.”

Comprised of 400 Computation Nodes Peak Operations Performance: 250TFLOPS

Sango, which OIST began using in April 2015, is comprised of 400 computation nodes, 20 large capacity memory nodes, 3 Intel Xeon Phi coprocessor nodes, and 3 acceleration nodes. Built into each node is an Intel Xeon processor E5-2680 v3 that has 12 cores. HPC and Research Computing Engineer Francesca Tartaglione says, “Sango’s peak performance across the entire cluster is more than 250 TFLOPS. This is substantial improvement, more than 4x, compared to the previous cluster Tombo which could reach only 60 TFLOPS as peak performance”. When a parallel computation simulation was experimentally run in an environment with increased I/O speed between nodes thanks to InfiniBand* FDR, the speed improvement was confirmed to be a minimum of 1.5 to 2 times the original speed.

Utilized for Simulations of Ocean Currents in Marine Science and For the Modeling of Viscous Fluids

The HPC implemented by OIST is now being utilized in a variety of research fields. Researchers and students can request use of the large-scale comput-



Structure chart for “Sango (Coral)”, the HPC system built by OIST

ing environment at any time and carry out necessary data analysis and simulations. Among the 50 research units at OIST, the Marine Biophysics Unit which specializes in marine science is one of the research units that performs parallel computing most often.

Shohei Nakada, member of Marine Biophysics Unit, explained, “We simulate ocean currents with HPC to learn the patterns of objects like eggs floating in Okinawa’s coral reefs and oceans flow through the ocean. Specifically speaking, the sea area is divided into 960 subareas, and the computations for each subarea are allotted to a HPC computation node for parallel processing. Furthermore, by sharing information like water temperature between each of the nodes, the next flow can be predicted. Two generations ago, early stage research utilizing high performance computing was programmed with GFortran. As programmers move from GFortran to using Intel® Fortran we have seen increases in performance, allowing for simulations that took one

week to now be finished in two or three days. In addition to the use of Intel Fortran, I/O speeds are further accelerated with the new Sango system, so a further increase in performance is expected.”

The Mathematical Soft Matter Unit, which researches matter comprised of multiple atoms and molecules, utilizes the HPC system for large-scale computational science processing. Unit member Ryohei Seto explained, “We are researching not only simple objects and fluids, but also the physical properties of biological semiconductors and soft matter. We are currently conducting research that investigates the principles of viscous fluid movement under boundary conditions. But because this field does not have well-established models, we have to run simulations while improving existing models and establishing new models ourselves. We utilize HPC’s multithread processing to model the flow of viscous fluids while running large-scale computations by trial and error.”

The HPC system at OIST, making use of Okinawa Prefecture's frequent typhoons, is also being utilized to predict the destructive power of approaching typhoons. Tapan Sabuwala from the Continuum Physics Unit stated that they are doing research which involves collecting over a long period the magnitudes of battering strong winds that arise from typhoons, using mathematical models to analyze them, and predict the magnitude of typhoons.

Examining the Utilization of Knights Landing To Optimize HPC


Dyce discussed the prospects, "The number of students and researchers contin-

ues to increase at OIST. By implementing Sango we succeeded in obtaining an HPC system which meets our projected demands and allows for a flexible increase in capacity if required. This environment supports our goal to actively support researchers in accelerating their analysis and simulation activities through optimizing their code and increasing the utilization of parallel computing. Looking forward, we are considering how we can use the next generation Intel Xeon Phi coprocessors to further accelerate and support research activities at OIST."

OIST is a place where world class intellect gather to research and educate cutting edge science and technology. The

HPC system equipped with technology including the Intel Xeon processor and Intel Xeon Phi coprocessor will continue to make contributions as the core computational foundation of this graduate university.

Find the solution that's right for your organization. Contact your Intel representative, visit [Intel's Business Success Stories for IT Managers](#), and check out [IT Center](#), Intel's resource for the IT industry.



(From left to right)

Scientific Computing & Data Analysis Section
HPC and Research Computing Engineer
Eddy Taillefer

Scientific Computing & Data Analysis Section
HPC and Research Computing Engineer
Francesca Tartaglione

Chief Information Officer
Tim Dyce

Continuum Physics Unit
Research Specialist
Tapan Sabuwala

Marine Biophysics Unit
Technician
Shohei Nakada

Mathematical Soft Matter Unit
Staff Scientist
Ryohei Seto



Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to <http://www.intel.com/performance>.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Intel, the Intel logo, Intel Inside, the Intel Inside logo, Xeon, and Intel Xeon Phi are trademarks of Intel Corporation in the United States of America and other countries.

Copyright © 2015 Intel Corporation. All rights reserved.

* Other company names, product names, etc. are generally representations, trademarks, or registered trademarks of these companies.

Printed in USA JPN/1509/PDF/SE/ISG/FN 332672-001US