Case Study
Intel® Xeon® Processor E5 Family
Intel® Xeon Phi™ Coprocessor
High-Performance Computing
Government/Public Sector

High-Performance Computing Drives Scientific Research

Jefferson Lab more than doubles application performance running optimized code on the Intel® Xeon Phi™ coprocessor

Researchers from across the United States and around the world use resources at the Thomas Jefferson National Accelerator Facility (Jefferson Lab) to conduct nuclear physics experiments that enhance our understanding of matter. As part of a continuous effort to improve the performance of scientific applications, the lab implemented a new cluster that integrates Intel® Xeon Phi™ coprocessors alongside the Intel® Xeon® processor E5 family. By optimizing software for the Intel Xeon Phi coprocessor, developers have increased performance by more than two times over the most optimized code on Intel Xeon processors, enabling researchers to accelerate projects and address complex scientific questions.

Challenges

- **Accelerate application performance.** Adopt technologies that will help improve performance of research applications.
- **Maximize the benefit of optimization work.** Make the most of programming work by enabling developers to use the same software development models for both coprocessors and industry-standard x86 processors.

Solution

- **Seneca* servers based on the Intel Xeon processor E5 family and Intel Xeon Phi coprocessors.** Jefferson Lab installed 16 nodes equipped with Intel Xeon processors E5-2650 and Intel Xeon Phi 5110P coprocessors running on the Red Hat Enterprise Linux* operating system.

Technology Results

- **Improved performance.** By running optimized code on an Intel Xeon Phi coprocessor, the Jefferson Lab team can run some applications approximately 2.2 times faster than with the Intel Xeon processor E5 family. Researchers are looking to achieve additional gains through further software optimization.

Business Value

- **Faster results.** Running optimized code on Intel Xeon Phi coprocessors, researchers could speed completion of complex scientific calculations. The experiments help researchers gain a better understanding of the building blocks of matter and how these particles build our world and universe.
- **Better science.** Faster processing enables researchers to enhance the precision of calculations and explore more complex problems, which would have been too costly to explore in the past.

"With optimized code, researchers will be able to generate results faster than before....[T]hey can explore additional, more complex scientific questions."

– Chip Watson, Manager, High-Performance Computing, Jefferson Lab

One of 17 national labs funded by the U.S. Department of Energy, Jefferson Lab enables researchers to use its Continuous Electron Beam Accelerator Facility to conduct basic research into the nuclei of atoms. The lab's Scientific Computing Group provides the computing resources to perform calculations that support those experiments in nuclear and particle physics while also engaging in its own research in high-performance computing (HPC).

The Scientific Computing Group at Jefferson Lab continuously explores new technologies that might help researchers address more complex problems and generate results faster.

For example, the Scientific Computing Group has been experimenting with clusters equipped with graphics processing units (GPUs) and clusters equipped with Intel Xeon Phi coprocessors for several years. "We have had good success improving the performance of research codes on both architectures," says Chip Watson, manager of high-performance computing at Jefferson Lab.

For Jefferson Lab, the introduction of Intel Xeon Phi coprocessors provided an interesting alternative to GPUs. "At the time of our last purchase, the Intel Xeon Phi coprocessors offered a better price/performance ratio

Intel® Xeon Phi™ coprocessors help boost speed for solving complex calculations

than GPUs," says Watson. "In addition, they can help reduce programming work by allowing developers to use the same software development models as they would for industry-standard x86 processors."

Exploring Intel Xeon Phi Coprocessors

To evaluate the potential of Intel coprocessors for scientific research, Jefferson Lab implemented a new cluster that features Intel Xeon Phi coprocessors. The cluster was named 12M, which designates the year the cluster was commissioned (2012) and its use of the Intel Xeon Phi coprocessor's many-integrated core architecture (M). The 12M cluster includes 16 nodes, each equipped with the Intel Xeon Phi 5110P coprocessor alongside the Intel Xeon processor E5-2650. The cluster runs Lattice Quantum Chromodynamics* workloads on a Linux operating system.

While some organizations take the approach of offloading only some application code to a coprocessor or GPU, the team at Jefferson Lab is taking an alternate path. "We are strongly interested in running entire applications on the Intel Xeon Phi coprocessor instead of using them in an offload mode," says Watson. "With GPUs, we were reaching the limits of the performance that could be achieved when only part of the code was running on the accelerator. The Intel Xeon Phi coprocessor provides a potential opportunity to have much more of our code running on the high-performing architectures."

Streamlining Application Optimization

Since making the decision to implement a cluster equipped with Intel Xeon Phi coprocessors, developers at Jefferson Lab have been working to optimize research code for the new architecture. Jefferson Lab developers are working closely with members of the Intel® Parallel Computing Lab. "The Intel team has helped us overcome several key challenges, including finding a way to reorganize data to take full advantage of the many-core architecture," says Bálint Joó, a Jefferson Lab computer scientist focusing on that software development work. "We are also working with Intel to develop a package of solvers to provide a high-performance library."

Using the Intel Xeon Phi coprocessor helps streamline development work. "While the high-performance kernels took some development effort, porting the rest of the code has been straightforward because the Intel Xeon Phi coprocessor enables us to use regular programming models," says Joó.

Code optimizations made for the Intel Xeon Phi coprocessors should also help improve application performance when that code is running on other Intel Xeon processors. "The fact that they are both fundamentally x86 architectures allows us to make the most of our efforts on one to the benefit of the other," says Watson.

Accelerating Results and Improving Precision

The lab has seen significant performance improvements by running optimized code on the Intel Xeon Phi coprocessors. "Our optimized code runs 2.2 times faster on a single Intel Xeon Phi coprocessor than on two processors from the Intel Xeon processor E5 family," says Joó.

Better application performance offers multiple benefits to researchers. "With optimized code, researchers will be able to generate results faster than before," says Watson. "As a result, they can explore additional, more complex scientific questions that they wouldn't have been able to afford to explore before. At the same time, they can improve the precision of calculations so results are more accurate. Ultimately, researchers will produce not just more science, but better science."

Lessons Learned

While the initial porting of the code for the Intel® Xeon Phi™ coprocessor has been straightforward, the Jefferson Lab team still highlights the importance of software development and optimization. "To achieve the best performance, you will always have to invest time and effort in the optimization process," says Bálint Joó, Jefferson Lab computer scientist. "On the plus side, we have found that optimizing the code for the Intel Xeon Phi coprocessors has resulted in up to a 3.6 times increase in performance on Intel® Xeon® processors compared with our previous baseline. Multiplied by the 2.2 times increase from moving from Intel Xeon processors to Intel Xeon Phi coprocessors, the new architecture and code together bring a total performance increase of close to 8 times over our initial baseline code on Intel Xeon processors."

Sharing the Benefits

Looking ahead, the Jefferson Lab developers will continue to optimize applications for the Intel Xeon Phi coprocessor so researchers across multiple projects can benefit from performance gains. The lab will consider implementing a production cluster designated for research within the next two years.

In the meantime, additional researchers will likely benefit from the development work done at Jefferson Lab. The organization is incorporating optimized code as part of its extensible Chroma* C++ framework, which researchers can use as a building block for their research applications. "Our hope is to offer the software that we have developed with Intel to researchers around the world as open source," says Watson.

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.**

 Printed in USA 0619/LJ/TDA/XX/PDF Please Recycle 330378-001US