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

Look Inside.™

Accelerating the Performance of Science and Engineering Applications

The National Institute for Computational Sciences helps researchers optimize applications with a cluster based on Intel® Xeon Phi™ coprocessors

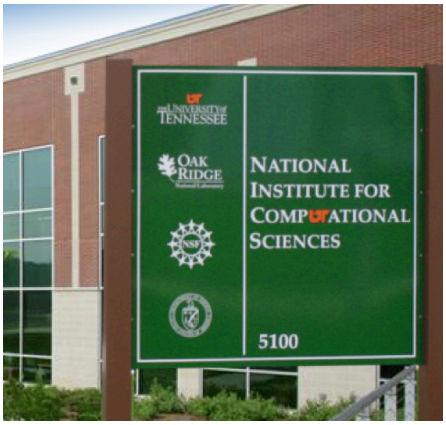

"With one optimized computational fluid dynamics code, we achieved a little more than 2.25 times the performance on an Intel® Xeon Phi™ coprocessor compared with two Intel® Xeon® processors E5-2670... [R]esearchers can build clusters that use Intel Xeon Phi coprocessors to boost performance while reducing costs."

– Glenn Brook, Chief Technology Officer, Joint Institute for Computational Sciences, University of Tennessee

With funding from the National Science Foundation (NSF) and the University of Tennessee (UT), along with assistance from Intel, the National Institute for Computational Sciences (NICS) built Beacon—a cluster designed to help researchers optimize science and engineering applications for the Intel® Many Integrated Core Architecture (Intel® MIC Architecture). Nodes are equipped with the Intel® Xeon® processor E5 family and Intel® Xeon Phi™ coprocessors. By optimizing code, the NICS team and associated researchers found they could achieve more than 2.25 times better performance using a single Intel Xeon Phi coprocessor compared with two processors from the Intel Xeon processor E5 family. Building similar clusters will enable researchers to solve larger, more complex problems while controlling costs.

Challenges

- **Optimize performance.** Provide a high-performance computing (HPC) cluster that would enable researchers to optimize their code for the Intel MIC Architecture.
- **Control costs.** Find ways to build more efficient clusters that can help researchers solve larger, more complex problems without significantly increasing hardware acquisition expenses, energy consumption, or software development costs.

Solutions

- **Cray CS300-AC cluster supercomputer* with the Intel Xeon processor E5 family and Intel Xeon Phi coprocessors.** Beacon comprises 48 compute nodes and 6 I/O nodes with a total of 768 conventional cores and 11,520 accelerator cores. Compute nodes are equipped with Intel Xeon processors E5-2670 and Intel Xeon Phi coprocessors 5110P. Intel® Solid-State Drives (Intel® SSDs) are integrated into the storage environment.
- **Intel® Software Development Tools.** Software developers used the Intel® Cluster Studio XE suite to optimize code.

Technology Results

- **Increased performance.** Optimized code running on an Intel Xeon Phi coprocessor can deliver more than 2.25 times better performance than with two Intel Xeon processors E5-2670.

Business Value

- **Solving larger problems while reducing costs.** Using Intel Xeon Phi coprocessors can help researchers address new challenges while keeping infrastructure and software optimization costs in check.

"In the past, adopting newer, faster processors has been the primary way scientists and engineers have achieved better application performance," says Glenn Brook, CTO at the Joint Institute for Computational Sciences (JICS) at the University of Tennessee. "Today, they need to parallelize code and capitalize on the rapidly increasing numbers of processing cores for significant performance gains."

While some software developers have begun to modify code for graphics processing units (GPUs) or similar workload-offloading technologies, that software development approach can be very labor-intensive. "Taking advantage of GPUs often involves rewriting code in a proprietary, vendor-specific language," says Gregory Peterson, associate professor of electrical engineering and computer science at the University of Tennessee and project director at NICS. "It would be much less labor-intensive to use a familiar programming model."

Collaborating with Intel on a New Approach

In 2011, NICS began working with Intel to explore ways researchers could take advantage of the Intel MIC Architecture that was in development. Intel provided preproduction versions of what would later be named the Intel Xeon Phi coprocessor along with Intel

Intel® Xeon Phi™ coprocessors help boost application performance while controlling costs

Software Development Tools. "Unlike other acceleration technologies, the Intel Xeon Phi coprocessor supports a variety of programming models," says Brook. "Our goal was to evaluate whether optimizing code for the Intel Xeon Phi coprocessor could improve the performance of parallelized scientific and engineering codes while minimizing software development work compared with proprietary accelerator solutions."

Building Beacon

With funding from the NSF, NICS built Beacon, a Cray CS300-AC cluster supercomputer with 48 compute nodes and 6 I/O nodes joined with InfiniBand* connectivity. "NICS has a strong relationship with Cray," says Brook. "We knew that Cray could provide solid hardware, offer excellent support, and deliver systems on our accelerated timeline."

The compute nodes are equipped with both Intel Xeon processors E5-2670 and Intel Xeon Phi coprocessors 5110P. In all, the cluster includes 768 conventional cores and 11,520 accelerator cores. NICS also uses Intel SSDs for the cluster's local storage environment to provide higher input/output operations per second (IOPs) and lower latency than traditional spinning disks.

The Intel Xeon processor E5 family provides the flexibility NICS needs to support an array of scientific and engineering applications. "We serve a diverse user base," says Peterson. "While researchers running simulations might need raw compute power, bioinformatics users require large amounts of memory. The Intel Xeon processor E5 family can accommodate a very wide variety of workloads."

Bringing together Intel Xeon processors and coprocessors opens new possibilities for software development and infrastructure testing. "This hybrid environment allows us to explore a variety of programming and processing scenarios," says Brook. "At the same time, the environment is designed to help us examine energy efficiency, data movement, and other variables. We hope to find new ways to maximize performance, minimize energy consumption, and reduce costs."

Beacon was ranked first on the Green500 list in November 2012, delivering nearly 2.5 billion floating-point operations per second (gigaFLOP/s) per watt. The cluster has also been ranked on the TOP500 list in November 2012 and June 2013, having demonstrated more than 210 TFLOP/s of combined computational performance.

Streamlining Optimization with Intel® Software Development Tools

Intel provided early versions of Intel Cluster Studio XE tools to help the NICS team members and researchers optimize codes for the new architecture. Using familiar Intel tools helped streamline the optimization work and multiply potential benefits. "With the Intel Software Development Tools, optimizing for the Intel Xeon Phi coprocessor is not substantially different than optimizing for the Intel Xeon processor E5 family," says Brook. "In addition, the optimization work done for the Intel MIC Architecture will also benefit codes if they are run on standard Intel Xeon processors."

Boosting Application Performance by Up to 2.25 Times

Using Intel tools, NICS developers quickly saw the potential of the Intel Xeon Phi coprocessors. "With a reasonable degree of optimization, you can achieve the same level of performance on an Intel Xeon Phi coprocessor as on a pair of Intel Xeon processors from the E5 family, at the same power level," says Brook. "If you go farther and more fully expose the parallelism, you can realize more significant gains."

NICS testing with optimized code showed just how significant those performance improvements can be. "With one optimized computational fluid dynamics code, we achieved a little more than 2.25 times the performance on an Intel Xeon Phi coprocessor compared with two Intel Xeon processors E5-2670," says Brook. "Those results indicate that researchers can build clusters that use Intel Xeon Phi coprocessors to boost performance while reducing costs."

Controlling Hardware Acquisition and Software Development Costs

The price/performance advantages of incorporating the Intel MIC Architecture into future clusters have become clear from the Beacon project. "Using Intel Xeon Phi coprocessors, organizations can build smaller clusters with fewer nodes and achieve the same performance as larger clusters, saving hardware acquisition and energy costs," says Brook. "Of course, many organizations will choose to capitalize on cost advantages to achieve even greater performance while keeping costs flat."

Lessons Learned

"The Intel® Many Integrated Core Architecture (Intel® MIC Architecture) enables a flexible approach to programming and computing," says Glenn Brook, CTO at the Joint Institute for Computational Sciences (JICS) at the University of Tennessee. "You can write applications to run directly on the coprocessor or to offload work from the processor to the coprocessor. For organizations that have already written code for GPUs, that flexibility is very helpful in letting them capitalize on the Intel® Xeon Phi™ coprocessors."

Intel Xeon Phi coprocessors can also help reduce software development costs. "With the Intel MIC Architecture, programmers do not have to rewrite code in another language—they can use the same programming model they used for Intel Xeon processors," says Peterson. "That saves time and money."

Working Toward Sustainable Computing

Through the Beacon project, the NICS team has been able to envision a more sustainable computing paradigm—one that can help answer increasingly complex research questions without adding resources. "We hope to expand the Beacon project, creating more of a production environment that is available for science and engineering research," says Brook. "At the same time, we will continue to evaluate the ways Intel MIC Architecture can help reduce energy consumption and control the demand for human resources in software development. We believe that this technology is an important step in fostering more sustainable computing in the future."

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 0614/LJ/TDA/XX/PDF ♻ Please Recycle 329980-001US