

ARRA-enabled 'Barracuda' computing cluster allows scientists to team up on larger problems

March 20 2012



EMSL scientist Dr. Karol Kowalski and colleagues use Barracuda to develop novel software solutions for high-impact science.

Within the Department of Energy's (DOE) EMSL, new highperformance computing breakthroughs often are the result of combining the best of two worlds. Experimental and computational tools are integrated; suites of leading-edge hardware and software are developed in tandem; and, perhaps more than ever, scientists from different disciplines combine expertise. The addition of the Barracuda computing cluster, funded by the American Recovery and Reinvestment Act, has brought a new level of collaboration between teams of domain scientists



(such as chemists) and computer scientists.

"If there wasn't already a walking path there, we would have worn one into the grass," said Dr. Karol Kowalski, while Dr. Sriram Krishnamoorthy agreed jokingly, referring to the expanse of lawn between EMSL and CSF, the Computational Sciences Facility at Pacific Northwest National Laboratory (PNNL).

"These frequent collaborations help us get the most out of Barracuda as we prepare a major update to NWChem [EMSL's widely used opensource computational chemistry software application]," Kowalski continued. "The ultimate goal, of course, is to further optimize how we use computations to predict the properties of matter."

FASTER SOLUTIONS, FEWER RESOURCES

Predicting the properties of matter has intrigued curious minds for thousands of years. Today, computational chemists like Kowalski and computer scientists like Krishnamoorthy are pushing the boundaries of such predictions using Barracuda, which is part of EMSL's Molecular Science Computing capability. In particular, the focus is on calculating the properties and structures of molecules involved in the most societally important chemical reactions—those related to energy innovations, environmental protection, national security and human health. For example, more efficient solar panels can result from highly advanced simulations of how electrons reorganize themselves when exposed to light.

The problem with scientifically impactful calculations is they tend to be very costly in terms of time-to-solution.

"We have a code for highly advanced theoretical formalisms for approximate solving of the Schrödinger equation that describe the



properties of molecules," Kowalski said. "But, it requires a high investment of computational resources to achieve reliable answers. Barracuda uses GPUs, or graphics processing units, a new type of architecture that can get to the solution faster. We're working with people like Sriram, Wenjing Ma, and Oreste Villa [PNNL highperformance computing experts] to translate NWChem for use on this architecture.

"This is the first documented attempt to apply GPU-based technology to the most advanced theoretical methods."

For researchers creating simulations and models for complex scientific problems, their implementations will translate to scientifically significant answers for larger systems, with a lower investment of resources.

GPUS: FROM GAMING TO GAME-CHANGING

A paradigm shift is happening within the highperformance computing world: the move from homogeneous to heterogeneous computer architectures. In other words, rather than using multiple identical cores in parallel to solve problems, new systems are using multiple types of cores. In Barracuda's case, it uses both CPUs (central processing units) and GPUs. With 60 nodes, each node of Barracuda consists of two quadcore Intel Xeon X5560 CPUs with 8 MB L2 cache running at 2.80 GHz.

The strategy is part of this decade's "holy grail" quest to achieve exascale computing—a thousandfold increase in performance over today's fastest supercomputers.

GPUs originated in the late 1990s as an innovation for the video game and computer graphics industries. Surprisingly, they have risen to prominence in broader computing applications.



"At first, the motivation was quickly manipulating a screen full of pixels," said Krishnamoorthy. "Traditional CPUs weren't very good at it, so a new architecture was built to move lots of data from memory to the monitor. Over time, people realized that GPUs are not only fast, but they offer significant improvements over CPUs in memory bandwidth and power efficiency. Now, these advantages are being applied far beyond graphics."

In fact, GPU computing at its present state can bring about significant increases in overall speed, using its advantages to handle the most computationally intensive tasks and remove data bottlenecks during calculations.

FLEXIBILITY FOR IMPACT

Of course, a new generation of GPU-enhanced hardware is only as good as the software developed to run on it. As new supercomputers around the world are being built with GPUs, including Titan at DOE's Oak Ridge National Laboratory, EMSL and PNNL scientists are preparing for the shift by using Barracuda to help them develop the GPU extension of NWChem—an overhaul that adds such functionality to the majority of the software application. Specifically, it improves highly accurate methods accounting for instantaneous interactions between electrons and methods designed to treat very large systems (plainwave density functional theory methods).

"We're taking NWChem to the next level, to harness the power of GPUs for studying molecular systems," Krishnamoorthy said. "But, the key is to keep it flexible enough to work on a large variety of heterogeneous architectures. Nvidia's CUDA, or compute unified device architecture, is just one GPU computing engine, but there are others such as the [Khronos Group's] optimize applications for GPU architectures, PNNL is the first DOE national laboratory to be recognized by Nvidia as a



CUDA Research Center.

More information: Ma W, et al. In press. "Optimizing tensor contraction expressions for hybrid CPU-GPU execution." Cluster Computing, *Online First*. DOI: 10.1007/s10586-011-0179-2

Ma W, et al. 2011. "GPU-Based Implementations of the Noniterative Regularized-CCSD(T) Corrections: Applications to Strongly Correlated Systems." *Journal of Chemical Theory and Computation* 7(5):1316-1327. DOI: 10.1021/ct1007247

Ma W, et al. 2010. "Acceleration of Streamed Tensor Contraction Expressions on PGPU-Based Clusters." In Cluster Computing (CLUSTER), Proceedings of the 2010 IEEE International Conference on *Cluster Computing*, pp.207-216. September 20-24, 2010, Heraklion, Crete. Institute of Electrical and Electronic Engineers, Piscataway, N.J. DOI: 10.1109/CLUSTER.2010.26

Provided by Environmental Molecular Sciences Laboratory

Citation: ARRA-enabled 'Barracuda' computing cluster allows scientists to team up on larger problems (2012, March 20) retrieved 27 April 2024 from <u>https://phys.org/news/2012-03-arra-enabled-barracuda-cluster-scientists-team.html</u>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.