Oak Ridge Leadership Computing Facility Snapshot
The Week of March 15, 2010

OLCF Releases "Science at the Petascale 2009: Pioneering Applications Point the Way"

Highlights from the first petascale research conducted on Jaguar

"Science at the Petascale 2009: Pioneering Applications Point the Way," a document detailing the first ultrascale research conducted on the Cray XT5 known as Jaguar at the Oak Ridge Leadership Computing Facility (OLCF) is now available online at http://www.nccs.gov/media-center/nccs-reports/.

"The projects featured in this document used pioneering applications to facilitate understanding in areas of research that are critically important to problems of national importance," said Doug Kothe, director of science at the OLCF. The term petascale refers to a computer system capable of performing one quadrillion floating point calculations—or a petaflop—per second. Jaguar became the world's first petascale system for unclassified research when it came online in November 2008 with a peak performance speed of 1.64 petaflops.

In the first half of 2009, the OLCF invited 28 of the world's leading computational research teams to participate in a six-month program of early petascale research on the newly upgraded Jaguar. Using more than 360 million combined processor hours, these research teams tackled some of the most pressing contemporary issues in climate science, chemistry, materials science, nuclear energy, physics, bioenergy, astrophysics, geosciences, fusion and combustion. The OLCF's primary goals for this phase of petascale early science were to not only deliver pioneering science results, but also to engage a broad community of users capable of hardening the nascent system for the nearly 40 projects allocated time on Jaguar in 2009 by the Department of Energy's Innovative and Novel Impact on Computational Theory and Experiment (INCITE) program. Several research teams using Jaguar during this early petascale phase ran the largest calculations ever performed in their field, and three codes achieved sustained production performance of over one petaflop.

"Science at the Petascale 2009" highlights 20 of the 28 research teams that participated in the early petascale program.

Petascale early science projects in climate included climate models of unprecedented resolution. One project carried out century-long simulations with a coupled atmospheric-oceanic climate model to explore the possibility of predicting climate on a decadal scale. The ultra-high resolution of the simulated climate system allowed researchers to study fine-scale phenomena such as ocean eddies and storm patterns and determine the effect of a warming climate on tropical storm frequency and intensity.

The list of accomplishments also included the largest known simulations of a nuclear reactor. This work provided researchers with an accurate depiction of the distribution of power within a modern fission reactor, paving the way for the design of next-generation reactors.

Other projects included studies of high-temperature superconductors, cost-effective means for producing cellulosic ethanol from biomass, and calculations of the influx of uranium pollution from the Hanford 300 waste disposal site into the Columbia River basin.

"With powerful applications and huge computational allocations on Jaguar, these researchers were able to punch through long-standing barriers and create new and useful knowledge," said Kothe. "We are confident that their achievements will stand over time as substantial contributions to science and technology."

<u>Document Maps the Road to Exascale</u>

Center releases "Preparing for Exascale:OLCF Application Requirements and Strategy"

Scientists have proven many times over that they can use supercomputers to simulate systems—from entire galaxies to the universe's smallest particles—with breathtaking accuracy. Together with theory and experiment, simulation has become a pillar of scientific inquiry.

Nevertheless, there are many critical questions these researchers still cannot answer, simply because they do not have sufficient computing power. To meet our computing needs in climate research, alternative energy development, and other critical areas, we will need exascale systems, capable of more than a million trillion calculations each second and a thousand times more powerful than today's most advanced supercomputers.

We cannot create exascale supercomputers simply by replicating existing leadership systems and making them a thousand times bigger. Years of coordinated effort will be required from hardware manufacturers, supercomputing centers, and researchers. In response to this challenge, the OLCF has taken the lead in preparing for the exascale age by encompassing the expertise of its staff and user community in a 136-page blueprint entitled *Preparing for Exascale: ORNL Leadership Computing Facility Application Requirements and Strategy*.

"It's becoming clear that getting to exascale will require paradigm shifts in how we do high-performance computing," said OLCF staff scientist Wayne Joubert, who produced the document with staff scientist Hai Ah Nam and OLCF Director of Science Douglas Kothe. "We need to begin preparing now for the hardware and software changes required for exascale computing."

The specific challenges addressed in the document are drawn from four general research areas pursued by Department of Energy's (DOE's) Office of Science researchers, including energy security, nuclear security, scientific discovery and innovation, and environmental responsibility. Specific research goals include such things as

- developing high-resolution models capable of analyzing climate change at the regional level;
- developing realistic simulations of lignocellulose (used in producing biofuels) and next-generation materials, efforts that will require a fortyfold increase in the number of atoms and a tenfold shrinking in the units of time being simulated; and

- developing accurate simulations of a full nuclear reactor core and working from the fundamental laws of nature to calculate nuclear fission.

Supercomputers able to handle these problems will be a collaboration among many communities. Manufacturers will combine different, but complementary, types of computer chips to get the highest possible performance. Developers of programming languages and software tools will give application programmers the resources they need to take advantage of new, mixed architectures. And programmers themselves will need to adapt and, in some cases, rewrite codes that are already enormously complex. In many areas, these communities will need to develop creative new technologies and new ways of thinking to overcome looming roadblocks.

“We’ve reached the limit of getting gains by increasing the clock speed of processors,” Joubert said. “Many people believe we will need some sort of accelerator technology to get to exascale.”

At the OLCF, for instance, the accelerator technology will come from graphics processing units (GPUs), originally created for video applications. Also known as graphics accelerators, GPUs are parallel processing on steroids, with each chip containing hundreds of small cores able to race through simple instructions. They will be combined in the next OLCF system with traditional processors—known as central processing units—to create a system ten times faster than today’s most powerful supercomputers.

Tools have already been developed for this new computing environment, and they are continually being improved. Graphics processor manufacturer NVIDIA has created a programming environment called CUDA (for Compute Unified Device Architecture) that allows programmers to make the most of these mixed systems while still writing in traditional languages such as C and Fortran. A similar, open-source environment is known as OpenCL.

Even with these tools, however, it will be a massive undertaking for developers to adapt applications that typically contain tens of thousands of lines of code. Nevertheless, Joubert said research teams have been open to the idea, especially given the boost in computing power they can expect.

“There’s been a remarkable openness among scientists. They are seeing that this is going to be needed to get to the next level.”

He noted also that CUDA is becoming a common tool among young scientists, thanks to the efforts of universities.

“There are hundreds of universities teaching CUDA,” he said. “It’s new to high-performance computing, but it’s very dominant in some of the other software development communities.”

By working with the OLCF user community, Joubert, Kothe and Nam were able to incorporate views from many of the world’s most advanced scientific computing teams. That perspective, combined with the experience of OLCF staff, make this document unique.

For instance, both Joubert and Nam work with the OLCF's Jaguar system, a Cray XT5 capable of more than 2 petaflops (2 thousand trillion calculations a second) and currently ranked the world's fastest supercomputer. Nam is a developer for a nuclear coupled-cluster application called NUCCOR, which, among other things, explores the forces holding together atomic nuclei. Joubert is a developer of an application called Denovo that, among other things, simulates nuclear reactors.

"We have an extremely concrete view of things because we're in the trenches with the codes," said Joubert. "These are some of the only major application codes that have scaled to 200,000 processors. We have a perspective that's informed by real, practical experience with these applications."

Computational Powerhouses Hold Workshop

Users learn ins and outs of fastest hardware for science

Two ORNL-based computing facilities joined the National Energy Research Scientific Computing Center (NERSC) to sponsor a Cray XT5 workshop February 1-3 at the University of California, Berkeley.

Staff from the Department of Energy's OLCF and the University of Tennessee's National Institute for Computational Sciences (NICS) attended the Joint NERSC/OLCF/NICS Cray XT5 Workshop to train high-performance computing and scientific communities in the use of the world's largest and most powerful leadership systems. Two of these—Jaguar and Kraken—are housed at ORNL.

The processing cores for Jaguar, the fastest supercomputer in the world, and Kraken, the fastest academic machine, were upgraded in October from four to six new AMD Istanbul Opteron processor cores per node. This and other new hardware features present novel potential for applications. Staff and representatives from each facility and Cray were on site to lecture and provide assistance during hands-on sessions with the latest hardware, helping users effectively utilize the new XT5 architectures.

Since Jaguar and Kraken's users are stationed across the country, and sometimes the world, the February workshop at UC Berkeley closely mirrored another held at ORNL in December.

"The OLCF hosted the Hex-core Workshop at Berkley in conjunction with NERSC to make it easier on our West Coast users to attend this important workshop," said Ashley Barker, the OLCF's User and Assistance Outreach group leader. "We recognize that our users' time is very valuable and it is our hope that by hosting the workshop on both sides of the country we made the training more accessible to all of our users."

Topics included programming effectively for the XT5 and proper use of the new six-core CPU architecture.

This marks the first time the three computing facilities have collaborated to host a workshop.