Oak Ridge Leadership Computing Facility Snapshot
The Week of December 14, 2009

Denovo Uses Jaguar to Test Its Muscle for Reactor Problems

Radiation transport code runs largest known simulations of a nuclear reactor

Denovo, a new parallel solver for radiation transport developed at Oak Ridge National Laboratory (ORNL), is contributing to the design of radiation-tracking tools to deal with deadly threats to national security.

Now the Denovo code, developed 2 years ago by nuclear engineer Tom Evans, is evolving to tackle bigger, more complex tasks—ones that require a state-of-the-art supercomputer such as the Jaguar system at the Oak Ridge Leadership Computing Facility (OLCF). Evans received a 5-million-processor-hour allocation during Jaguar's 2009 Early Science Program, which devoted the system to especially promising research projects after it was upgraded to a peak speed of more than a quadrillion calculations a second or a petaflop.

That allocation gave Evans and ORNL collaborators Scott Mosher and Kevin Clarno an opportunity to test whether Denovo can develop into a scalable code that can calculate such big science problems as the multiple uncertainties in radiation transport in next-generation fission reactors and hot plasmas in fusion prototype reactors.

The team successfully ran the largest-known simulations to date of a nuclear reactor on Jaguar's 150,000-core XT5 partition. The code solved a variety of large-scale problems in nuclear engineering and technology. Calculations were done to determine when a reactor reaches and sustains a chain reaction and to determine the optimal safeguards. They did calculations in shielding, to determine the protective perimeter that is needed around large nuclear complexes and radioactive sites, as well as for ITER and individual pressurized-water reactors. They used the code to gauge what dose is absorbed by containers that transport or store radioactive materials.

Denovo solves the linear Boltzmann transport equation for the interaction between radiation and matter. The equation relates the production of radiation particles to their loss: on one side of the equation is the generation of ionized particles, on the other is how they are dispersed. Evans' code successfully characterized a pressurized-water reactor—the most common design in the United States—with 8.164 trillion degrees of freedom or unknowns. Next, it characterized a pressurized-water reactor core with 26 billion degrees of freedom, and then a boiling-water reactor with 126 billion degrees of freedom.

The simulations gave researchers a more accurate picture of the distribution of fission power, allowing them to resolve what had been only approximations of the neutrons' individual energies and directions of movement. This is critically important in both the design of next-generation reactors and improvement of existing ones.

The power of the Jaguar leadership system was needed because numerical solutions to the Boltzmann equation gobble up computational resources. The calculations solve the density of the radiation as it interacts with matter at any point, multiplied by the speed it travels. They are necessarily three dimensional and solve all points in phase space over seven independent variables: three in space, two in angle, and one each in energy and time. When the researchers separate out all such variables in three dimensions, the computational demand is staggering. Besides the complexity of calculating the transport of ionized neutrons from fission in the core of a reactor, there is the challenge of calculating the interaction of the neutrons with the other parts of the reactor's architecture. Reactors are made of massive arrays of fuel rods, coolant channels, and control rods, as well as reflectors and shielding that are riddled with ducts and other geometric anomalies. Without the power of Jaguar, researchers could not begin to simulate such transport.

Jaguar's petascale speed and balance were such that the researchers could run full-facility simulations, dividing the reactor space into a billion zones and generating detailed equations that calculate the trajectories of the neutrons as they scatter in fission. The researchers were able to turn such problems around in an hour.

Evans ran Denovo on transport problems spread over 20,000 to 40,000 cores on Jaguar, with one big run on 60,000 cores. The code showed excellent weak scalability—as the problem and machine sizes grew, the time to solution remained constant—over problems as much as eight times the size of the original problem. For fusion they ran heating calculations on 1,600 to 14,400 cores, and Denovo scaled perfectly along those ranges. The code handled a whole new class of problems in nuclear reactor analysis, which is an area traditionally hampered by unreliable, low-order approximations.

Evans has shown that Denovo can harness Jaguar to solve high-fidelity radiation transport calculations for a variety of very large nuclear energy and technology problems. "This is a new code, and I designed it to use the latest parallel algorithms, even though the majority of users don't actually take advantage of those. That allowed us to go out on Jaguar and see what kind of problems we can run."

The researchers next want to calculate transport in large reactor systems and then compare their findings with operating parameters in existing reactors. "If we can do that, that would be a very meaningful result. Then we can say 'Okay, these calculations have merit.' I think with the early science project, we proved we can run problems at that scale," said Evans.

Workshop Prepares Users for Upgraded Cray XT5 Machines

Users get a chance to test-drive two of the world's fastest supercomputers

A workshop held at ORNL December 7-9 helped users and developers acclimate to an upgrade from four- to six-core AMD processors in both the Cray XT5 machines known as Jaguar and Kraken, respectively owned by the OLCF and the National Science Foundation- (NSF) supported National Institute for Computational Sciences (NICS).

“We have these workshops periodically to educate and re-educate the users on the best way to make use of the resources that we have,” said Ricky Kendall, group leader for the scientific computing group at OLCF. The center aims to advance the state of high-performance computing by providing scientists with the most powerful computational tools available including not only Jaguar, but also state-of-the-art data analysis and visualization technology, and a highly-trained staff that can work closely with scientific teams to ensure maximum results.

“In these workshops, we cover all components of the machine,” explained Kendall. “If we can do optimizations that get [the users] a 5X improvement in the run time of their code, then there’s that much more science they can get from their allocations.”

Representatives from over 20 organizations attended the workshop which included lectures by staff members of the OLCF, NICS, Cray, and AMD. The lectures focused on XT5 architecture, effective programming on the XT5, AMD six-core processors, and issues involving the NUMA compute nodes located on each processor. The workshop also featured hands-on sessions that allowed participants to access Jaguar using their own codes and work one-on-one with OLCF and NICS staff members to resolve any issues. Some users have never run their codes on a machine with as many processors as Jaguar or Kraken, and that can pose a problem.

“Things that you do on a 1,000 core cluster that you don’t see as a problem often bite you at 100,000 cores,” explained Kendall. “In the process of developing the leadership computing facility, we’ve gone from 25 teraflops to 54 teraflops to 119 teraflops to 264 teraflops to 1 petaflop and now to 2 petaflops and at every jump, codes that were stellar performers had to change things to be able to compute at the next scale.”

Jaguar and Kraken are mainly employed by research teams from laboratories and universities the world over, but they also serve industries who are interested in using computing and computational science in their research and development processes. Major industries such as GE Global Research, Whirlpool Corporation, and BMI Corporation as well as smaller companies had representatives present at the workshop. These companies are either currently using ORNL computing resources through the Industry Partnership program or are anticipating applying for time.

The American Recovery and Reinvestment Act-funded upgrade has placed two of the top five supercomputers in the world at ORNL, moving Jaguar from the number two spot into the lead position on the Top500 list of the world’s fastest supercomputers with a peak performance speed of 2.33 petaflops—or 2.33 thousand trillion calculations per second. Managed by the University of Tennessee for NSF, Kraken moved from sixth to third on the list with a peak performance speed of 1.03 petaflops, also making it the first academic supercomputer to breach the petascale.

Through hands-on workshops, state-of-the-art technology and a skilled technical support staff, the OLCF will continue to bring together leaders in science and technology to provide the best environment for research teams from business, academia and industry to tackle issues of global importance.

CPUs to GPUs

Researcher talks about next-generation of supercomputing at Oak Ridge seminar

On Dec. 15, 2009, Pacific Northwest National Laboratory researcher Oreste Villa spoke about effective methods of utilizing graphic processing units (GPUs) in high-performance computing (HPC). Villa's talk was part of the OLCF Seminar Series. The series hosts monthly presentations covering areas of interest in high-performance computing, with a focus on petascale computing.

The HPC community anticipates future computer architectures to make substantial use of GPUs, which crunch data much faster than central processing units (CPUs). But GPUs must first receive information from CPUs, and this communication has been routinely slow. Villa presented attendees with new methods that increase the speed of communication and, subsequently, the calculated results.

"These kinds of investigations are important because they explore and help delineate the capabilities of GPUs for specific application areas in high-performance computing," says Donald Frederick, a member of the User Assistance and Outreach Department at the OLCF.

BOF Seeks User Assistance Excellence

HPC Centers meeting evolving

For the fourth time, the HPC Centers birds-of-a-feather (BOF) gathered at SC09 to brainstorm best practices for user assistance departments at some of the country's top high-performance computing centers.

57 participants, representing user assistance personnel from the Department of Defense, the Department of Energy, NASA, the National Science Foundation, the National Center for Atmospheric Research, the University Corporation for Atmospheric Research, and several universities, lent their ears and minds to the open forum discussion on user assistance issues.

Organized by Bobby Whitten, an HPC user support specialist at the OLCF, the meeting is slowly changing from a BOF to a regular working group.

Beginning in January 2010, the BOF will go from its semiannual physical meeting to monthly teleconferences, though they will continue to meet physically at SC and possibly other conferences. "What started as a birds-of-a-feather has now reached critical mass and this collaboration promises to bring user assistance to the next level," said Whitten.

According to participant Julie Mullen of MIT's Lincoln Laboratory, all of these labs share many of the same issues. "If you can leverage what someone else has already done, it's helpful . . . many heads are always better than one." And besides the brainstorming, the meeting provides a great networking opportunity in a narrow field, said Mullen.