

National Center for Computational Sciences Snapshot

The Week of October 27, 2008

Oak Ridge Delivers Several Breakthroughs

Jaguar makes mark on ASCR document

A recently released document showcasing 10 scientific computing milestones includes five projects conducted at the National Center for Computational Sciences (NCCS) at Oak Ridge National Laboratory (ORNL). The document, *Breakthroughs 2008*, chronicles major advances in simulation over the past 18 months under the auspices of the Department of Energy's (DOE's) Office of Advanced Scientific Computing Research (ASCR).

Among the praised ORNL-based research were one of the largest simulations ever produced of plasma confinement in a fusion reactor, which will potentially pave the way for energy production that emits no carbon dioxide into the atmosphere; a billion-particle simulation of the dark matter halo of the Milky Way galaxy, in which researchers performed the largest simulation to date of the dark matter cloud holding our galaxy together; and combustion simulations that dissected how flames stabilize, extinguish, and reunite, possibly leading to cleaner, more efficient diesel-engine designs. The ORNL accomplishments, which represent half of the total list, took place on the laboratory's Cray XT4 known as Jaguar.

The list of breakthroughs was compiled by a distinguished panel of computational scientists, applied mathematicians, and computer scientists made up of representatives from ASCR-associated labs (including Oak Ridge) and participating universities. Each of the cited accomplishments was supported in a broad sense by ASCR through the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program, the Scientific Discovery through Advanced Computing (SciDAC) program, and/or its base program. The report can be viewed at http://www.nccs.gov/wp-content/media/nccs_reports/Breakthroughs_2008.pdf.

Meet the Performance Police

Pat Worley's team fixes problems that threaten simulations

When scientists want more accurate or more detailed simulations, they turn to modeling experts and software engineers who upgrade the capabilities of simulation models such as the Community Climate System Model (CCSM), a megamodel coupling four independent models whose codes describe Earth's atmosphere, oceans, lands, and sea ice. But when the software engineers need help, they turn to Pat Worley at ORNL. He leads a DOE SciDAC project with Arthur Mirin of Lawrence Livermore National Laboratory and Raymond Loy of Argonne National Laboratory to scale up climate codes, enabling them to solve larger problems by using more processors, and to evaluate software and new high-performance computing platforms such as the Cray XT4 and IBM Blue Gene/P supercomputers.

“An important practical aspect of climate science is figuring how much science you can get in the model and still get the simulations done in time,” said Worley, whose team works with

researchers and manufacturers to identify bugs in CCSM codes, performance bottlenecks in the algorithms used in the CCSM, and glitches in a machine's software. "Our contribution is getting the component models to run as efficiently as possible. The software-engineering aspects of the code are always changing, and often the new code has unexpected performance issues. We monitor things. We're kind of the performance police."

Worley and his colleagues push codes to their limits. If a code runs slowly on 1,000 processors but quickly on 2,000, they might assign more processors to work on a problem. If, owing to algorithmic restrictions, the code can't use more than 1,000 processors, changing algorithms may be the only option to improve performance. Different science also imposes different performance requirements. Ocean scientists may choose to run a high-resolution ocean model coupled to a low-resolution atmosphere model, whereas atmospheric scientists may pick the converse. Changes to the codes to improve performance for one scenario must not slow down the code for another or hurt performance on a different (or future) platform. Often the performance team introduces algorithm or implementation options that scientists can choose to optimize performance for a given simulation run or on a particular computer system.

On Cray and IBM systems, the group has improved performance through both algorithmic and implementation efforts. Recent work improved performance 2.5-fold on benchmark problems on ORNL's Cray XT4. "With the improvements to the scalability of the CCSM software by Pat and his colleagues, along with the dramatic growth in the performance of Jaguar, the CCSM developers are seriously considering model resolutions and advanced physical processes that were not on the table before," said Trey White, who is an NCCS liaison to the CCSM project helping the scientists make the most of the machines.

"Pat Worley's group has provided critical support in improving the scalability and performance of the CCSM across a wide range of architectures," said Mariana Vertenstein of the National Center for Atmospheric Research (NCAR), head of the engineering group responsible for CCSM's software development, support, and periodic community releases. "The CCSM project played a major role in the fourth assessment report of the Intergovernmental Panel on Climate Change, or IPCC AR4, through an extensive series of modeling experiments and in fact resulted in the most extensive ensemble of any of the international global coupled models run for the IPCC AR4. This accomplishment could not have occurred without Pat's contributions."

Worley's team is currently working with a large multilab SciDAC project led by John Drake, chief computational scientist for the Climate End Station at ORNL, with Phil Jones of Los Alamos National Laboratory to build a first-generation Earth system model, which will extend the physical climate model by including chemical and ecological processes. The computer allocations are provided through the Climate Science Computational End Station, an INCITE program award led by NCAR's Warren Washington on Jaguar in the NCCS.

"For DOE, which is very concerned with the carbon cycle and with the impact of climate change on ecology and ecosystem services, this kind of Earth system model is really called for," Drake said. "We're trying to do whatever we can to get there as quickly as possible."

Phoenix Makes Way for Petascale Age

Venerable vector system retired as next generation comes on line

ORNL's Phoenix supercomputer, still one of the fastest vector systems in the world, was taken out of service at the beginning of October to make way for petascale systems capable of 1,000 trillion calculations a second (1 petaflop).

Installed in 2003 with a peak performance of 3.2 trillion calculations a second (3.2 teraflops), Phoenix was ORNL's most powerful system when the lab's Leadership Computing Facility was established in 2004. It has been upgraded in the years since and in its latest configuration had more than 1,000 multistreaming vector processors and a peak performance of more than 18 teraflops. In its 5½ years at ORNL, Phoenix rose as high as No. 17 on the TOP500 list of the world's fastest supercomputers.

At the time it was installed, supercomputing was still dominated by Japan's Earth Simulator, but both systems have long been surpassed. In fact, Phoenix is being removed from ORNL's computer room to make way for two systems that are each more than 300 times as powerful as that 3.2-teraflop configuration. The NCCS's Cray XT5 Jaguar system will boast a peak performance of 1 petaflop, while the National Institute for Computational Sciences' Cray XT5 Kraken system will peak at just under that level.

In its time Phoenix was critically important to progress in fields such as computational fluid dynamics, climate science, fusion studies, astrophysics, and materials science. "Many of our users loved working on Phoenix," said NCCS Director of Science Doug Kothe. "It was a fantastic machine."

ORNL's Anthony Mezzacappa and colleagues used the system to advance the world's understanding of core-collapse supernovae. The team first discovered that the shockwave created by a star's collapsing iron core becomes unstable and wobbly before it blows most of the star into space. Later it showed that this instability is very possibly responsible for the pulsar that is all that remains of the star.

"My heart sank a little when Phoenix was finally turned off, and not just for sentimental reasons," said Bronson Messer of the NCCS, a member of Mezzacappa's team who first used Phoenix while he was a postdoctoral researcher at the University of Chicago. "The machine was remarkably capable even at the end of its tenure and allowed us to perform supernova simulations—ranging from the birth of pulsars to the signature of galactic supernovae in terrestrial neutrino detectors—with amazing efficiency."

Climate scientists moved to Phoenix and found it to be a great advance over earlier systems used to provide data for the IPCC AR4.

"We saw a factor-of-15 increase in speed over current production," noted John Drake, chief computational scientist for the Climate End Station at ORNL. "It was a great machine for climate simulation."

ORNL's Thomas Schulthess and colleagues used Phoenix to make a major breakthrough in materials science, showing that a model known as the 2D Hubbard model correctly describes high-temperature superconductors where more conventional density functional theory fails. The team's discovery moved forward a line of research that may eventually prove revolutionary in areas such as electric-power generation, electronics, and transportation.

"We did important work on Phoenix," said ORNL's Thomas Maier, a member of Schulthess's team.

Kothe noted that while Phoenix's passing marks the end of an era, the system is moving out for the best of reasons.

"We hate to see this great vector hardware and software leave," he acknowledged, "but we expect these petascale systems to enable unprecedented advances in computational research, and we'll very likely see more vector hardware in next-generation systems."

New System Maintaining Software for NCCS Users

"NCCS Software Environment" streamlines researchers' third-party preferences

Running a Leadership Computing Facility presents plenty of challenges, all of which the NCCS is tackling head-on. Take third-party software for example. Different researchers need or prefer different outside libraries, editors, management tools, and the like—all of which are stored on NCCS systems. Needless to say, with numerous researchers using a host of favorite utilities, managing and maintaining this external software inventory presents plenty of organizational headaches and consumes precious time and, of course, money.

The NCCS currently supports more than 70 different (nonvendor) libraries, tools, and applications for its Cray XT4 system alone. And for most applications and libraries, the NCCS supports multiple versions of the software and multiple builds of each version to support three different compilers and different compilations. Somehow, some way, it has to be organized, an unenviable task by any measure.

As the NCCS has expanded over the years, the list of third-party software and associated organizational challenges has also grown. To address and improve the situation, NCCS staff member Mark Fahey has created a new suite of tools and policies for installing and maintaining users' preferred software, including naming schemes, template scripts for building and testing software, and scripts to police the entire installation "area." While the new tools, dubbed the NCCS Software Environment, were originally designed for the center's Cray XT4, they now manage all third-party installations for the Blue Gene/P and two AMD-based quad-core clusters (Lens and Smoky) at the NCCS as well.

The installations are primarily managed by a core team of 3 staff members, but the infrastructure is designed so that many staff members can contribute. For example, 20 different staff members contributed XT4 installations when the new file system went on line, and they are occasionally called upon when issues arise. Fahey says the new third-party area

is staying organized and working well, a testament to the new scheme's design and implementation.

The NCCS Software Environment package was unveiled internally in March, and all the Jaguar software was reinstalled by Jaguar's upgrade acceptance in May. Following acceptance, Fahey was able to rebuild and test approximately 30 applications in a day's time, a huge improvement from the past when this work was all done manually. Furthermore, the package is now largely automated, meaning that after a machine is upgraded with new hardware, a new operating system, or new compilers, the old software builds can easily be tested and if necessary relinked or rebuilt, quickly making them research ready—a stark contrast to the time when reaching this point would take many days of tedious work.

Software web pages are easier to maintain now, too, thanks to a feature that allows the automated generation of web pages for each third-party application. The installer of the package writes a “basic” description file describing the package and how to use it on the machine, and a script in the NCCS Software Environment generates the web pages for each machine with a web page for each package. “Now software web pages are a nonissue,” said Fahey. Furthermore, he added, the web pages are available in “list” and “category” views, making it easier for researchers to more quickly find what they're looking for. For example, in the category view researchers can browse for libraries, performance, tools, and such in the individual categories, largely taking the “search” out of research.

The NCCS runs the developer-built tests for the new packages as well and rates their verification for users on its website, which also shows up on the application web page. “Users will now have some expectation of how well the package works through the verification posted on the NCCS website,” said Fahey. Finally, Fahey has added a little something extra for the users: a set of date-sensitive scripts that alerts the NCCS via email when new versions of third-party software are released. “We are proactively ahead of users on new versions and packages,” he said.

Fahey has also developed a related software-tracking system that enables NCCS administrators to view which users are using which software. Because software management and maintenance consume resources, it is helpful for Fahey and others to know which packages the center should work to update and maintain. By tracking application link lines, Fahey can see what libraries, compilers, and such are being used and how often. If users have abandoned a particular software package, then continuing to support it is a waste of time and, perhaps most importantly, money.

The tracking project complements the Software Environment package and vice versa, said Fahey, adding that “there are so many ways we can target information now.” So far the tracking system operates on only the NCCS's Cray XT4 test system, but Fahey hopes to have it on the NCCS's upgraded Cray XT petascale system this fall and on the remaining systems soon thereafter.

Besides the tools, there are also new rules governing how new software is to be installed by staff members. Fahey has developed a set of scripts that checks daily to make sure staff are

following protocol. He calls it “policing,” adding that it has been very useful. “The policing activities are what allow us to spread many installations across many staff members, with just a few staff members making sure every installation has met our design goals, with the entire tree of software then managed by these few staff members.”

The new NCCS Software Environment package adds up to big savings in money through big savings in time for the NCCS. It can now provide the software the users need as well as keep up to date on latest production versions and new compilers more effectively than ever. And when the tracking goes live, the organization will be able to even further refine its focus. “We now have control over the software and are policing it regularly,” said Fahey. The rules are straightforward, but Fahey is quick to add that he is still improving and tweaking certain aspects of the system.

But it’s not about the NCCS, of course, it’s about the users. The more organized the NCCS is, the better it can serve researchers and ultimately do big science with minimal headaches, continuing its role as one of the top supercomputing facilities in the world.