

# **Oak Ridge Leadership Computing Facility Snapshot**

## **The Week of November 30, 2009**

### **Jaguar: The World's Most Powerful Supercomputer. For Science!**

*Oak Ridge supercomputer takes number 1 on Top500 list*

An upgrade to a Cray XT5 high-performance computing system deployed by the Department of Energy (DOE) has made the “Jaguar” supercomputer the world’s fastest. Located at Oak Ridge National Laboratory (ORNL), Jaguar is the scientific research community’s most powerful computational tool for exploring solutions to some of today’s most difficult problems. The upgrade, funded with \$19.9 million under the Recovery Act, will enable scientific simulations for exploring solutions to climate change and the development of new energy technologies.

“Supercomputer modeling and simulation is changing the face of science and sharpening America’s competitive edge,” said Secretary of Energy Steven Chu. “Oak Ridge and other DOE national laboratories are helping address major energy and climate challenges and lead America toward a clean energy future.”

To net the number-one spot on the TOP500 list of the world’s fastest supercomputers, Jaguar’s Cray XT5 component was upgraded this fall from four-core to six-core processors and ran a benchmark program called High-Performance Linpack (HPL) at a speed of 1.759 petaflop/s (quadrillion floating point operations, or calculations, per second). The rankings were announced November 15 in Portland at SC09, an international supercomputing conference. The HPL benchmark is one of four applications on Jaguar that have exceeded 1 petaflop of performance.

In 2004, DOE’s Office of Science set out to create a user facility that would provide scientists with world-leading computational research tools. One result was the Oak Ridge Leadership Computing Facility (OLCF), which supports national science priorities through the deployment and operation of the most advanced supercomputers available to the scientific community.

“Our computational center works closely with the science teams to effectively use a computer system of this size and capability,” said James Hack, director of the National Center for Computational Sciences that houses Jaguar in the OLCF.

Jaguar began service in 2005 with a peak speed of 26-teraflop/s (trillion calculations per second) and through a series of upgrades in the ensuing years gained 100 times the computational performance. The upgrade of Jaguar XT5 to 37,376 six-core AMD Istanbul processors in 2009 increased performance 70 percent over that of its quad-core predecessor.

Researchers anticipate that this unprecedented growth in computing capacity may help facilitate improved climate predictions, fuel-efficient engine designs, better understandings of

the origin of the universe and the underpinnings of health and disease, and creation of advanced materials for energy production, transmission, and storage.

The Oak Ridge computing complex is home to two petascale machines. In addition to DOE's Jaguar system, the National Institute for Computational Sciences, a partnership between the University of Tennessee and ORNL, operates another petascale Cray XT5 system known as Kraken, which was ranked 3rd on the November Top500 list at a speed of 831.7 teraflops.

“The purpose of these machines is to enable the scientific community to tackle problems of such complexity that they demand a well tuned combination of the best hardware, optimized software, and a community of researchers dedicated to revealing new phenomena through modeling and simulations,” said ORNL Director Thom Mason. “Oak Ridge is proud to help the Department of Energy address some of the world's most daunting scientific challenges.”

Simulations on Jaguar have primarily focused on energy technologies and climate change resulting from global energy use. Scientists have explored the causes and impacts of climate change, the enzymatic breakdown of cellulose to improve biofuels production, coal gasification processes to help industry design near-zero-emission plants, fuel combustion to aid development of engines that are clean and efficient, and radio waves that heat and control fuel in a fusion reactor.

“The early petascale results indicate that Jaguar will continue to accelerate the Department of Energy's mission of breakthrough science,” said Jeff Nichols, ORNL's associate laboratory director for computing and computational sciences. “With increased computational capability, the scientific research community is able to obtain results faster, understand better the complexities involved, and provide critical information to policy-makers.”

### **ORNL-Led Team Takes Prize for World's Fastest Science App**

*Second ORNL-led team also finalist for Gordon Bell Prize*

A team led by ORNL's Markus Eisenbach was named winner Thursday of the 2009 ACM Gordon Bell Prize, which honors the world's highest-performing scientific computing applications. Another team led by ORNL's Edo Aprà was also among nine finalists for the prize.

Results of the contest were announced in Portland, Oregon, during the SC09 international supercomputing conference. The prize is supported by high-performance computing pioneer Gordon Bell and is administered by the Association for Computing Machinery.

Eisenbach and colleagues from ORNL, Florida State University, and the Institute for Theoretical Physics and Swiss National Supercomputing Center achieved 1.84 thousand trillion calculations per second—or 1.84 petaflops—using an application that analyzes magnetic systems and, in particular, the effect of temperature on these systems. By accurately revealing the magnetic properties of specific materials—even materials that have not yet been produced—the project promises to boost the search for stronger, more stable magnets, thereby

contributing to advances in such areas as magnetic storage and the development of lighter, stronger motors for electric vehicles.

The application—known as WL-LSMS—achieved this performance on ORNL’s Cray XT5 Jaguar system, making use of more than 223,000 of Jaguar’s 224,000-plus available processing cores and reaching nearly 80 percent of Jaguar’s peak performance of 2.33 petaflops. Earlier in the week Jaguar was named number one on the TOP500 list of the world’s fastest computers. The system was recently upgraded from four-core processors to six-core processors, boosting its peak performance to 2.33 petaflops.

WL-LSMS allows researchers to directly and accurately calculate the temperature above which a material loses its magnetism—known as the Curie temperature. The team’s approach differs from earlier efforts because it sets aside empirical models and their attendant approximations to tackle the system through first-principles calculations.

“What we can do is calculate the Curie temperature for materials with high accuracy without external parameters,” Eisenbach explained. “These first-principles calculations are orders of magnitude more computationally demanding than previous models; it’s only with a petascale system such as Jaguar that calculations like this become feasible.”

WL-LSMS combines two methods to achieve its goal. The first—known as locally self-consistent multiple scattering, or LSMS—applies density functional theory to solve the Dirac equation, a relativistic wave equation for electron behavior. The code has a robust history, having been the first code to run at a sustained trillion calculations per second, and earned its developers the prestigious 1998 Gordon Bell Prize. This approach, though, describes a system in its ground state at a temperature of absolute zero, or nearly  $-460^{\circ}\text{F}$ . By incorporating a Monte Carlo method known as Wang-Landau, which guides the LSMS application, Eisenbach and his colleagues are able to explore technologically relevant temperatures ranges.

The work improves on previous advances in magnetic materials, Eisenbach said. He noted that materials research has led in the past century to more than a 50-fold increase in the magnetic strength of materials per volume and in the last decade to more than a 100-fold increase in the density of magnetic data storage. Other efforts that may benefit from the research include the design of lighter, more resilient steel and the development of future refrigerators that use magnetic cooling.

Aprà’s team—the other finalist led by an ORNL researcher—achieved 1.39 petaflops on Jaguar in a first principles, quantum mechanical exploration of the energy contained in clusters of water molecules. The team, comprising members from ORNL, Australian National University, Pacific Northwest National Laboratory (PNNL), and Cray Inc., used a computational chemistry application known as NWChem, which was developed at PNNL.

The application used 223,200 processing cores to accurately study the electronic structure of water by means of a first-principles quantum chemistry technique known as coupled cluster. The team will make its results available to other researchers, who will be able to use this highly accurate data as inputs to their own simulations.

## **Tennessee Supercomputing Titans Triumph at HPC Challenge Awards**

*'Jaguar' takes three gold medals and a bronze while 'Kraken' scores two silvers*

Two powerful Cray XT5 systems at the ORNL computing complex outmuscled competitors to win half of this year's High-Performance Computing (HPC) Challenge awards. Results of the "Best Performance" awards, which measure excellence in handling computing workloads, were announced Nov. 17 at SC09, an international gathering of supercomputing professionals. The Department of Energy's "Jaguar" supercomputer took home the lion's share of the honors, with three "gold medals" and one "bronze." "Kraken," an academic supercomputer funded by the National Science Foundation (NSF) through a partnership with the University of Tennessee, showed with two "silver medals" that it too is a contender.

"The HPC Challenge benchmarks examine the performance of HPC architectures using kernels with more challenging memory access patterns than just the High Performance Linpack (HPL) benchmark used in the TOP500 list," said Jack Dongarra of University of Tennessee-Knoxville and Oak Ridge National Laboratory. Jaguar ranks first on the TOP500 list of the world's fastest supercomputers; Kraken ranks third.

Jaguar won first place for speed in solving a dense matrix of linear algebra equations by running the HPL software code at 1,533 teraflop/s (trillion floating point operations per second). Kraken, the world's fastest academic computer, took second by running HPL at 736 teraflop/s.

The fastest cat in the HPC jungle also ranked first for sustainable memory bandwidth by running the STREAM code at 398 terabytes per second. STREAM measures how fast a node can fetch and store information.

Jaguar's third "gold" was for executing the Fast Fourier Transformation (FFT), a common algorithm used in many scientific applications, at 11 teraflop/s. Kraken took second with a speed of 8 teraflop/s.

Edged out by IBM Blue Gene machines at Lawrence Livermore and Argonne national laboratories, Jaguar took third place for running the RandomAccess measure of the rate of integer updates to random locations in a large global memory array.

Lawrence Livermore National Laboratory's Blue Gene/L machine took third in HPL and second in STREAM competitions, and the Japan Agency for Marine-Earth Science and Technology placed third in both STREAM and FFT contests.

"It is very gratifying that Jaguar has been recognized as a very powerful machine," said Buddy Bland, project director for ORNL's Leadership Computing Facility, which hosts Jaguar. "The HPC Challenge benchmarks are designed to give a better view of the entire system's performance. Jaguar was designed to be the most powerful system for scientific applications, and these results reflect that design and implementation. It's no surprise that

Kraken, using the same architecture and also designed for high-performance scientific computing, is now the second most powerful machine in the world for that purpose.”

To support development of the hardware and software needed to use supercomputers capable of executing quadrillions of calculations each second, the Defense Advanced Research Projects Agency, International Data Corporation, DOE, NSF, and the Center for Information Technology Research sponsor the HPC Challenge.

### **ORNL Computing Garner Awards from Online Computing Publications**

*HPCwire and insideHPC recognize Oak Ridge for computing leadership*

ORNL has been hand-picked by insideHPC readers to receive the publication’s first-ever HPC Community Leadership Award. A U.S. Department of Energy laboratory, ORNL has routinely contributed to scientific computing in numerous fields including astrophysics, fusion energy, materials science, and biophysics. As home to the world’s most powerful computing complex, ORNL houses Jaguar, the fastest supercomputer in the world and the first petascale machine available for open science.

“ORNL has blazed a trail at the very high end of supercomputing in recent years,” said John West, editor of insideHPC (<http://www.insidehpc.com>). “Bringing together the expertise, funding, and organizational resources to build a record of sustained accomplishment at this level is a truly remarkable achievement.” West presented the award to Jeff Nichols, ORNL’s associate laboratory director for computing and computational sciences, November 16 at the 2009 International Conference for High Performance Computing, Networking, Storage and Analysis (SC09) in Portland, Oregon. “insideHPC’s readers have highlighted the confidence that the supercomputing community has not only in what ORNL has already accomplished, but in the leadership they will provide in the future,” West said.

Jaguar and ORNL also received an HPCwire Editors’ Choice Award for “Top Supercomputing Achievement.” The award was presented November 19 to Nichols at SC09 by Tomas Tabor, publisher of HPCwire (<http://www.HPCwire.com>)

“This award, which represents a partnership between the HPCwire global readership and our publishing team, is a salute from the global HPC community. Being selected as an award recipient means that you are at the top of mind of HPCwire readers, editors, and luminaries in the field,” said Tabor, whose publication covers the ecosystem of computationally- and data-intensive computing, including software, middleware, hardware, networking, storage, tools, and applications.

### **ADIOS Ignites Combustion Simulations**

*ORNL-led research enhances scaling of codes*

Despite the muscle of today's premier supercomputing systems, scaling and I/O often prevent leading software packages from taking full advantage of the latest hardware's potential power. These ancillary tasks take precious time and resources away from studying the fundamental science that impacts our everyday lives.

To address these issues, a team of researchers from ORNL, Georgia Tech, and Rutgers University developed the ADaptable I/O file System (ADIOS). ADIOS is an I/O middleware package that has shown great promise with leading fusion codes, scaling up to 140,000 cores for XGC-1, GTC, and GTS, and in astrophysics with the CHIMERA code. Recently, ADIOS made its mark in the field of combustion.

Researchers at Sandia National Laboratory (Ray Grout, Jackie Chen, and Chun Sang Yoo) and Andrea Gruber of SINTEF, the largest independent research organization in Scandinavia, are using the leading combustion code S3D to perform the first direct numerical simulations of reacting jets in cross flow. These transverse jets are a class of flows used in practical applications in which high mixing rates are desirable -- for example, in fuel injection nozzles in stationary gas turbines for power generation or in aero-gas turbines. The gas turbine industry is faced with numerous R&D challenges in adapting conventional hydrocarbon burner designs to operate safely and cleanly with hydrogen-rich syngas.

Though the S3D team was able to fully scale the code to ORNL's entire Cray XT5 supercomputer, known as Jaguar, the scaling of the I/O proved to be difficult. Enter ADIOS. Besides its ability to scale, ADIOS's BP file format is resilient to failures in the compute nodes and the file system, an attribute that hinted to researchers that S3D's analysis routines would pair well with the ADIOS ecosystem. After consultations with ORNL's Qing Liu and Scott Klasky (a leading ADIOS developer) and Georgia Tech's Jay Lofstead, Chen's post-doc Ray Grout decided to integrate ADIOS as an alternative I/O mechanism for S3D. ADIOS quickly surpassed the networking limitations imposed by the previous I/O stack.

"I appreciate the clean yet capable interface and that from my perspective, it just works," said Grout. "It would have been fairly difficult for us to get this work finished in time for our targeted paper deadline using our previous I/O solution, but now I'm confident that we'll make it thanks to ADIOS and the help of its team."

S3D's newfound scalability and flexible architecture has led a member of the S3D team, ORNL's Ramanan Sankaran, to explore creating a custom I/O transport method to perform special manipulations to the I/O and the data. With this new transport method, no source code changes would be required.

Thanks to the seamless integration of ADIOS into S3D, the middleware's 1.0 release may have already achieved the team's goal of satisfying 90 percent of ADIOS users "out of the box," said Klasky, adding, "It's really exciting to be able to help the top scientists such as Ray and Jackie obtain results more efficiently. These results make a difference in the world and will hopefully make the world a better place to live."

The ADIOS team thanks all of the scientists who have helped make this work possible, including Hasan Abbasi, Julian Cummings, Divya Dinkar, Ciprian Docan, Stephane Ethier, Garth Gibson, Steven Hodson, Scott Klasky, Zhihong Lin, Qing Liu, Jay Lofstead, Xiaosong Ma, Ron Oldfield, Manish Parashar, Norbert Podhorszki, Milo Polte, Alex Romosan,

Ramanan Sankaran, Karsten Schwan, Arie Shoshani, Mladen Vouk, Matthew Wolf, Yong Xiao, Weikuan Yu, and Fang Zhang.