# Oak Ridge Leadership Computing Facility Snapshot
# The Week of July 20, 2009

**<u>Oak Ridge Supercomputers Provide First Simulation of Abrupt Climate Change</u>**
*Researchers use 'Phoenix' and 'Jaguar' to study climate's past and future*

At the Department of Energy's (DOE's) Oak Ridge National Laboratory (ORNL), the world's fastest supercomputer for unclassified research is simulating abrupt climate change and shedding light on an enigmatic period of natural global warming in Earth's relatively recent history. The work, led by scientists at the University of Wisconsin and the National Center for Atmospheric Research (NCAR), is featured in the July 17 issue of the journal Science and provides valuable new data about the causes and effects of global climate change.

This research is funded by the Office of Biological and Environmental Research within DOE's Office of Science and by the National Science Foundation (NSF) through its paleoclimate program and support of NCAR.

In Earth's 4.5-billion-year history, its climate has oscillated between hot and cold. Today our world is relatively cool, resting between ice ages. Variations in planetary orbit, solar output, and volcanic eruptions all change Earth's temperature. Since the Industrial Revolution, however, humans have probably warmed the world faster than nature has. The greenhouse gases we generate by burning fossil fuels and forests will raise the average global temperature 2 to 12 degrees Fahrenheit (1 to 6 degrees Celsius) this century, the Intergovernmental Panel on Climate Change (IPCC) estimates.

Most natural climate change has taken place over thousands or even millions of years. But an episode of abrupt climate change occurred over centuries—possibly decades—during Earth's most recent period of natural global warming, called the Bolling-Allerod warming. Approximately 19,000 years ago, ice sheets started melting in North America and Eurasia. By 17,000 years ago, the melting glaciers had dumped so much freshwater into the North Atlantic that it stopped the overturning ocean circulation, which is driven by density gradients caused by influxes of freshwater and surface heat. This occurrence led to a cooling in Greenland called the Heinrich event 1. The freshwater flux continued on and off until about 14,500 years ago, when it virtually stopped. Greenland's temperature then rose by 27 degrees Fahrenheit (15 degrees Celsius) in several centuries, and the sea level rose about 16 feet (5 meters). The cause of this dramatic Bolling-Allerod warming has remained a mystery and source of intense debate.

"Now we are able to simulate these transient events for the first time," says Zhengyu Liu, a University of Wisconsin professor of atmospheric and oceanic sciences and environmental studies whose team simulated the abrupt climate changes using DOE supercomputers at ORNL. The Oak Ridge Leadership Computing Facility allocated supercomputing time through DOE's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. "It represents so far the most serious validation test of our model capability for simulating large, abrupt climate changes, and this validation is critical for us to assess the model's projection of abrupt changes in the future," according to Liu.

The Oak Ridge Leadership Computing Facility is funded by the Office of Advanced Scientific Computing Research in DOE's Office of Science.

Liu, director of the University of Wisconsin's Center for Climatic Research, and his collaborator Bette Otto-Bliesner, an atmospheric scientist and climate modeler at NCAR, lead an interdisciplinary, multi-institution research group attempting the world's first continuous simulation of 21,000 years of Earth's climate history, from the last glacial maximum to the present, in a state-of-the-art climate model. The group will also extend the simulation 200 years into the future to forecast climate. The findings could provide great insight into the fate of ocean circulation in light of continued glacial melting in Greenland and Antarctica.

**Three parts to abrupt change**

Most climate simulations in comprehensive climate models so far are discontinuous, amounting to snapshots of century-sized time slices taken every 1,000 years or so. Such simulations are incapable of simulating abrupt transitions occurring on centennial or millennial timescales. Liu and Otto-Bliesner employ petascale supercomputers, capable of a quadrillion calculations each second, to stitch together a continuous stream of global climate snapshots and recover the virtual history of global climate in a motion picture. They use the Community Climate System Model (CCSM), a global climate model that includes coupled interactions between atmosphere, oceans, lands, and sea ice developed with funding from the NSF, DOE, and NASA.

Based on insights gleaned from their continuous simulation, Liu and his colleagues propose a novel mechanism to explain the Bolling-Allerod warming observed in Greenland ice cores. The three-part mechanism they suggest matches the climate record.

First, one-third of the warming, or 9 degrees Fahrenheit (5 degrees Celsius), resulted from a 45 parts-per-million increase in the atmospheric concentration of carbon dioxide, the scientists posit. The cause of the carbon dioxide increase, however, is still a topic of active research, Liu says.

Second, another one-third of the warming was due to recovery of oceanic heat transport. When fresh meltwater flowed off the ice sheet, it stopped the overturning ocean current and in turn the warm surface current from low latitudes, leading to a cooling in the North Atlantic and nearby region. When the melting ice sheet was no longer dumping freshwater into the North Atlantic, the region began to heat up.

The last one-third of the temperature rise resulted from an overshoot of the overturning circulation. "Once the glacial melt stopped, the enormous subsurface heat that had accumulated for 3,000 years erupted like a volcano and popped out over decades," Liu hypothesizes. "This huge heat flux melted the sea ice and warmed up Greenland."

Liu and Otto-Bliesner's collaborators include Feng He, a doctoral student at the University of Wisconsin–Madison who is mainly responsible for the deglaciation modeling, as well as ocean modeler Esther Brady (NCAR), atmospheric scientist Robert Tomas (NCAR), glaciologists Peter Clark (Oregon State University) and Anders Carlson (University of Wisconsin–Madison), paleoceanographers Jean Lynch-Stieglitz (Georgia Institute of Technology) and William Curry (Woods Hole Oceanographic Institution), geochemist Edward Brook (Oregon State University), atmospheric modeler David Erickson (ORNL), computing expert Robert Jacob (Argonne National Laboratory), and climate modelers John Kutzbach (University of Wisconsin–Madison) and Jun Cheng (Nanjing University of Information Science and Technology). "This interdisciplinary team, each member contributing to a different aspect of the project, ranging from a proxy data interpretation to supercomputing coding, has been essential for the success of this project," says Liu.

The 2008 simulations ran on a Cray X1E supercomputer named Phoenix and an even faster Cray XT system called Jaguar. The scientists used nearly a million processor hours in 2008 to run one-third of their simulation, from 21,000 years ago—the most recent glacial maximum—to 14,000 years ago—the planet's most recent major period of natural global warming. With 4 million INCITE processor hours allocated on Jaguar for 2009, 2010, and 2011, they will complete the simulation, capturing climate from 14,000 years ago to the present and projecting it 200 years into the future.

"This has been a dream run of both of ours for a long time," says Otto-Bliesner. "This was an opportunity to take advantage of the CCSM, the computing facility at Oak Ridge, and the INCITE call for proposals." No other research group has successfully simulated such a long period in a comprehensive climate model.

**Science-based forecasts**

More accurately depicting the past means clearer insights into climate's outlook. "The current forecast predicts the ocean overturning current is likely to weaken but not stop over the next century," Liu says. "However, it remains highly uncertain whether abrupt changes will occur in the next century because of our lack of confidence in the model's capability in simulating abrupt changes. Our simulation is an important step in assessing the likelihood of predicted abrupt climate changes in the future because it provides a rigorous test of our model against the major abrupt changes observed in the recent past."

In 2004 and 2005, climate simulations on DOE supercomputers contributed data to a repository that scientists worldwide accessed to write approximately 300 journal articles. The published articles were cited in the Fourth Assessment Report of the IPCC, which concluded that global warming is definitely happening and humans have probably caused most of the warming since the mid-20th century.

Liu and Otto-Bliesner's simulations may soon find their way into IPCC's data repository and reports as other groups succeed in continuous simulation of past abrupt climate changes and demonstrate the results are reproducible. Meanwhile, Earth's climate continues to prove that change is an eternal constant. Understanding how we affect the rate of change is a grand challenge of our generation. Petascale computing may accelerate answers that in turn inform our policies and guide our actions.

**Fusion Gets Faster**

*New optimizations, enhanced I/O increase speed of GTC by more than 100 percent*

When it comes to scientific computing, the amount of science reaped from a simulation is largely determined by the speed and scalability of the software. Likewise, a code's speed is often at the mercy of its I/O performance. The more efficient the I/O, the faster the code and the more simulations can be run over a period of time.

Few codes require faster I/O or scale better than today's fusion particle codes. GTC and XGC-1, for instance, are running on more than 120,000 cores on the Oak Ridge Leadership Computing Facility's (OLCF's) Jaguar Cray XT5 supercomputer, the fastest system in the world for open science with a peak performance of 1.6 petaflops.

"These are the largest runs with the largest datasets," said Scott Klasky of the OLCF and the SciDAC scientific data management center. "And they are at the extreme bleeding edge of scalability and I/O."

Thanks to Klasky and a diverse team of collaborators, GTC recently became twice as fast. This number, said Klasky, was not reached only for an ideal benchmark case but for an actual production simulation. This impressive performance is the result of cross-discipline collaborations that have led to significant software and middleware improvements.

These advances are the result of software enhancements by Cray Inc. and a combined team effort of physicists (Y. Xiao and Z. Lin of the University of California–Irvine and S. Ethier of Princeton Plasma Physics Laboratory), vendors (N. Wichmann of Cray and M. Booth of Sun Microsystems), and computational scientists (S. Hodson, S. Klasky, Q. Liu, and N. Podhorszki of Oak Ridge National Laboratory [ORNL]; H. Abbasi, J. Lofstead, K. Schwan, M. Wolf, and F. Zheng of Georgia Tech; and C. Docan and M. Parashar of Rutgers).

"In order to advance the science, collaboration is essential," said Lin. "High-performance computing is more than benchmark numbers; it is about advancing scientific breakthroughs and that is accomplished by achieving high performance from both the code and the computing system [Jaguar]."

The various technical improvements include a new Cray compiler, optimizations to the code itself, and further I/O enhancements to ADIOS, an I/O middleware package created by Klasky and collaborators at Georgia Tech and Rutgers. From core physicists to programmers to hardware vendors, this group effort cut across organizational and disciplinary lines. "Working with some of the top computational scientists in the world, such as Parashar and Schwan, allows us to bring in new ideas that help enable more science in these codes," said Klasky.

While other members of the collaboration worked on enhancements in their respective areas, the ADIOS team was busy improving the I/O of some of the most scalable codes run at

ORNL. In the past, said Klasky, I/O wasn't a major issue simply because simulations had not reached the enormous scales seen on today's most powerful high-performance computing systems. Now, however, fusion simulations generate up to 100 terabytes of data per day.

"Researchers want easy-to-use, fast, scalable, and portable I/O," said Klasky, adding that the team is currently making additional updates to the ADIOS package for improved analysis capabilities. Today's supercomputers can make I/O performance difficult, thus the need for ADIOS, an I/O componentization layer that requires the users to add only a few lines of code to their applications to gain substantial I/O performance.

In part, it works by allowing users to switch between several best-practice I/O methods (ADIOS-BP with MPI I/O, POSIX, parallel HDF5, parallel NetCDF4, or even in situ visualization) without ever having to fundamentally change their codes or recompile. The new version will greatly ease the I/O burden on users, said Klasky.

"We are getting close to the peak performance of the I/O system when it comes to reading data in GTC," he said, adding that the GTC code read 100 terabytes of analysis data per hour on Jaguar when reading from 512 cores in a recent simulation. Previously, fusion codes used multiple different file formats, none of which worked well for both small- and large-scale data streams. This was the motivating factor behind ADIOS, which defines a novel, metadata rich, binary-packed "BP" file format capable of writing out GTC data at 80 gigabytes per second on Jaguar's XT5 component.

"ADIOS implements a new file format that was developed specifically to work well with parallel file systems from the ground up," said Klasky. Rutgers' Parashar agreed. "File formats for today's parallel file systems need to be redesigned to get bleeding-edge performance for both reading and writing," he said. "The old-school view that contiguous file formats are best for I/O is being revisited in the context of parallel file systems and is being challenged by new file formats."

"It's all about the science, and the best way to help the scientists is to work with them as a team to develop new and innovative software," said Klasky. These advancements in GTC will eventually find their way into other codes as well, he said, further allowing researchers to probe the complex properties of Mother Nature and tackle today's greatest scientific challenges.

ADIOS is open source and can be obtained from Scott Klasky at the OLCF (email: klasky@ornl.gov).

## Supercomputing Charts Unfamiliar Waters

*Simulations explore the air–water interface*

Between water and air lies a unique world of unknowns, called the interface. The molecules and ions in this region—located at the surface of your glass of water—behave differently than they do in either water or air.

Christopher Mundy, a physical chemist at Pacific Northwest National Laboratory (PNNL), and Doug Tobias, a chemistry professor at the University of California–Irvine, plan to explore the mysteries of aqueous interfaces. And they will have help from the Cray XT Jaguar supercomputer at Oak Ridge National Laboratory, on which they received 4 million processor hours in 2009 from the Department of Energy (DOE). Such DOE allocations give top researchers time on supercomputers capable of a quadrillion calculations per second, or a petaflop.

"Life as we know it takes place at interfaces," Mundy said. "On Jaguar we will look at water's ions in the vicinity of the air–water interface. We want to characterize the system both structurally and thermodynamically and compare it to experimental measurements. We can then begin to explain, with greater confidence, specific reaction mechanisms that are important for many chemical processes used in industry."

The commonly known properties of water usually apply to bulk water. These rules are tossed aside at the interface, which is known to have a dramatically different electronic structure.

"The forces on the molecules or ions at the interface may be different from those experienced in the middle of your glass, where they are completely surrounded by water," Mundy said.

The typical pH or acidity scale, for example, was determined in bulk liquid water. At the interface, though, water dissociates differently into its components: an acidic hydronium ion (H3O+) and a basic hydroxide ion (OH-).

Mundy and Tobias will simulate the air–water interface on the world's fastest computer for open scientific research, enabling them to see the breakup of water molecules into their ions as well as the interaction of water molecules with each other.

**Mapping a new world**

Using the CP2K computer code, designed to perform atomistic and molecular simulations of solids or liquids, the researchers have calculated the free energy to transfer a hydronium ion from the bulk to the interface and into the air. The resulting potential of mean force, or average of all forces, can be used as a starting point for a deeper understanding of the molecular interactions that could be responsible for the novel chemistry taking place in aqueous systems in the vicinity of interfaces.

With computer time from the DOE's Innovative and Novel Computational Impact on Theory and Experiment (known as INCITE) program and the National Science Foundation's Early PetaApps program, the researchers are performing state-of-the-art calculations to compare the structure and thermodynamics of the air–water interface to bulk water.

Another question the team is attempting to answer is the location of the hydroxide anion at the air–water interface. Having used about half of their processor hours, the researchers have found that the ions are absorbed back into water molecules—another interface reaction that differs from its bulk-water equivalent. With their remaining computer time, the team members hope to create a consistent picture of the chemistry and thermodynamics present in aqueous interfacial systems.

A more complete understanding of the interface also fits into the long-term goals of not only Mundy and Tobias, but also physical chemists in general because they want to know the limitations of simpler molecular models of ions in aqueous solutions. Mundy and his colleagues at PNNL are working toward creating more efficient molecular models that are transferable to different environments such as the air–water interface. He will try to establish a relationship between models run on the petascale supercomputer and less detailed models.

"Our research will provide benchmark calculations and should stimulate the development of new, cheaper molecular models because we can't all run Cray XT5s for months on end to get answers to complex problems," Mundy said. "We need to learn from these high-fidelity calculations what essential physics needs to be present in the construction of simpler and hopefully more accurate molecular models. In order to tackle the complex problems of chemical physics, you cannot tie yourself to a single model."

This project, like so many others, is part of a quest for a deeper understanding of the chemical processes around us. Even something as ordinary as tap water has mysteries yet to be explained.