OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

U.S. DEPARTMENT OF ENERGY | Office of Science

Twenty-five Years of Leadership Science at the Oak Ridge Leadership Computing Facility

Outreach manager – Katie Bethea

Writers – Eric Gedenk, Jonathan Hines, Katie Jones, and Rachel Harken

Designer – Jason Smith

Editor – Wendy Hames and Deborah Counce

Photography – ORNL Creative Services

Stock images – iStockphoto™

Oak Ridge Leadership Computing Facility
Oak Ridge National Laboratory
P.O. Box 2008, Oak Ridge, TN 37831-6161

Phone: 865-241-6536
Email: help@olcf.ornl.gov
Website: https://www.olcf.ornl.gov

Facebook: @oakridgeleadershipcomputingfacility

Twitter: @OLCFGOV

The research detailed in this publication made use of the Oak Ridge Leadership Computing Facility, a US Department of Energy Office of Science User Facility located at DOE's Oak Ridge National Laboratory. The Office of Science is the single largest supporter of basic research in the physical sciences in the United States and is working to address some of the most pressing challenges of our time.

For more information, please visit **science.energy.gov**.

Contents

OLCF

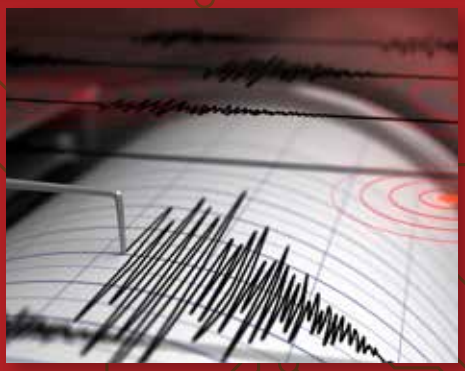

25
years

Twenty-five years ago, high-performance computing (HPC) in the United States stood at a crossroads. Established computing architectures were approaching their limits in performance and competiveness, while the country's need for computing power to solve challenging problems in science, energy, and national security continued to grow.

Out of this period of technological transition, a new force in scientific computing emerged in East Tennessee, a force that continues to shape the HPC landscape today.

In 2017, the Oak Ridge Leadership Computing Facility (OLCF) is celebrating 25 years of leadership in HPC. Since its founding as the Oak Ridge National Laboratory (ORNL) Center for Computational Sciences (CCS) in May 1992, the center has consistently delivered supercomputers of unprecedented capability to the scientific community on behalf of the US Department of Energy (DOE). Scientists, in turn, have used these versatile systems to solve critical problems in areas as diverse as biology, advanced materials, climate, and nuclear physics.

From the beginning, the OLCF has contributed to a rapid evolution in scientific computing that has produced a millionfold increase in computing power. This rise has included the launch of the first teraflop system (IBM Power3 Eagle) for open science, the science community's first petaflop system (Cray XT5 Jaguar), and two top-ranked machines on the TOP500 list, including the OLCF's current leadership-class machine, Titan. Additionally, the next chapter in the OLCF's legacy is set to begin in 2018 with the deployment of Summit, a pre-exascale system capable of 5 to 10 times the performance of Titan.

Using OLCF systems, scientists have expanded the scale and scope of their research, solved complex problems in less time, and filled critical gaps in scientific knowledge. Today, simulation is considered on par with experiment and theory as an essential standard of modern science.

At the start of the 1990s, however, predicting the OLCF's elevated HPC role would have seemed far-fetched. At the time, most of DOE's supercomputers were found in laboratories dedicated to national security and defense. Though ORNL's scientific computing group had been testing experimental architectures since the '80s, the only HPC systems in Oak Ridge were a Cray X-MP with a single CPU and a novel shared-memory system called a Kendall Square.

An opportunity to leap to the front of the scientific computing pack presented itself in 1991 when Congress signed the High-Performance Computing and Communication Act, legislation that called for government funding of the technology and infrastructure that would come to be known as the Internet. Additionally, the act directed DOE's Office of Energy Research, now the Office of Science, to issue a call for proposals for the creation of new HPC research centers to serve scientists from national laboratories, universities, and private industry.

At ORNL a team assembled that could put forth a compelling proposal for an HPC center to be located at Oak Ridge. The resulting document, titled "Partnerships in Computational Science" (PICS) and submitted in collaboration with three other national laboratories and seven universities, called for the construction of a large parallel computer to address three grand challenges: groundwater remediation, materials modeling, and quantum structure. The proposal also requested funds to develop new computational techniques, software, and tools to support this work and educational programs to develop interest in supercomputing among students in kindergarten through college.

When DOE awarded ORNL the bid on May 24, 1992, the announcement rocked the global scientific computing community. A new HPC powerhouse had arrived.

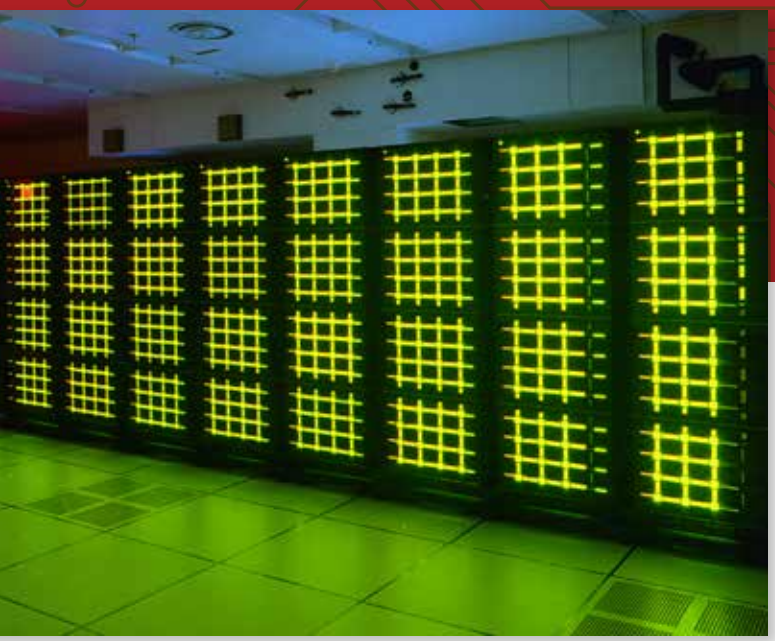

Technology

Upon its establishment, CCS immediately got to work with partner Intel to stand up a relatively new type of machine: a massively parallel supercomputer. This type of system solved large problems by dividing the work among many processors, allowing calculations to be carried out simultaneously. The approach represented a shift from older architectures, which relied on less flexible, increasingly expensive vector processors.

In 1992, CCS unveiled the Intel Paragon, a 66-core system capable of 5 gigaflops (5 billion calculations per second). A series of ever-more-powerful parallel computers at ORNL culminated in the 150-gigaflop Paragon XP/S 150, the fastest supercomputer in the world upon installation.

Working with scientific users, CCS staff helped showcase the benefits of parallel computing for research by getting scientific applications to run faster and faster. World-class science results followed, with projects providing insights into problems such as how solids melt, how combustion occurs in engines, and how air flows over a plane's wing.

The Paragon was replaced in 1999 by an IBM Power3 system called Eagle, the first teraflop machine in DOE's Office of Science. Other machines followed, including an IBM Power4 (Cheetah), a Compaq AlphaServer (Falcon), and a Cray X1 (Phoenix). The arrival of each supercomputer presented new opportunities and challenges for CCS staff, who documented achievements in technical papers and conference proceedings to benefit the HPC community.

In the early 2000s, ORNL cemented supercomputing's place at the laboratory by setting aside space in a new office building for a 40,000-square-foot datacenter. The investment paid off in 2004 when Congress passed the DOE High End Computing Revitalization Act, a response to bolster the nation's position in HPC in the wake of Japan's Earth Simulator, a machine as powerful as 14 of the fastest supercomputers in the US combined.

After a competitive proposal process, DOE awarded ORNL the first leadership computing facility, later named the Oak Ridge Leadership Computing Facility, to provide computing resources 100 times larger than those commonly available at the time. In addition to working with industry, universities, and other labs, the center was also charged with partnering with other federal agencies, including the National Science Foundation, the National Nuclear Security Administration, the National Institutes of Health, and the National Oceanographic and Atmospheric Administration, among others.

The establishment of the OLCF led to a new Cray XT3 system called Jaguar, which would undergo a series of upgrades—single-core processors to dual-core processors; then quad-core processors to six-core processors. In 2008, Jaguar, now a Cray XT5, became the first HPC system to run a scientific application at a

sustained petaflop. The achievement, carried out using a quantum cluster code called DCA++, won an ORNL-led team the prestigious Gordon Bell Prize the same year. In 2009, Jaguar reached No. 1 on the TOP500 list after running the Linpack benchmark application at 1.75 petaflops.

At its apex in 2012, Jaguar boasted nearly 300,000 cores and a theoretical peak of 3.3 petaflops, allowing researchers to investigate phenomena like high-temperature superconductivity in incredible detail. As one of the first machines to demonstrate petascale performance with multiple applications, Jaguar was widely recognized as much for its scientific productivity as for its technical prowess.

The OLCF's next move, however, would transform Jaguar into something else entirely. The combination of 16-core central processing units (CPUs) and graphics processing units (GPUs) marked the birth of Titan, a hybrid machine that took parallelism to the next level. Central to this achievement is the GPU, a highly parallel, highly efficient processor that can take on the primary calculations of computationally demanding scientific applications, freeing up CPUs for more complex operations.

Conceived by the OLCF and Cray as the answer to two seemingly incompatible objectives—a significant increase in computational power for a minimal increase in energy consumption—Titan debuted at No. 1 in November 2012 on the TOP500 list. Specifically, Titan, with a theoretical peak of 27 petaflops, realized 10 times the performance of Jaguar with only a 20 percent increase in electricity usage. Using Titan, researchers have contributed to the development of better biofuels, more efficient utility-scale gas turbines, and hazard maps for earthquake-prone regions. This new research has led to thousands of published scientific journal articles enabled by OLCF systems and, more importantly, has contributed to a safer, more energy-secure future.

Twenty-five years in, the OLCF continues to plan for the future. Summit will expand on Titan's groundbreaking hybrid architecture, leveraging advances in processor power, bandwidth, and data movement to deliver several times the computational power of the 27-petaflop Titan. Beyond Summit, plans are already being made to launch an exascale system, a machine capable of a billion billion calculations per second.

As the bar continues to rise, the OLCF remains committed to an unwavering mission: partnering with HPC vendors, scientific users, and others to deliver unprecedented tools for scientific discovery and global competiveness.

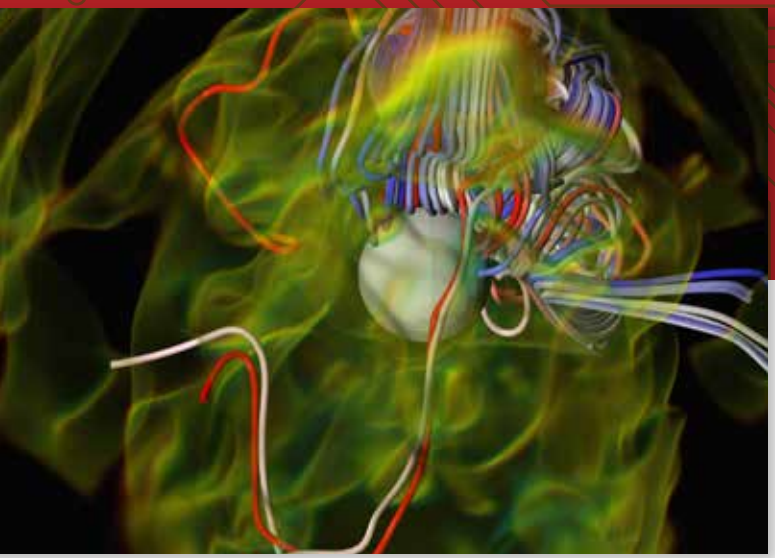

Science

Supercomputers enable watershed moments in science by driving new discoveries through modeling and simulation, connecting what we know from experiments and observations with our underlying theories about how the world works.

For the past 25 years, the OLCF has been a global leader in high-performance parallel computing and the development of computational applications and tools for scientific research. Each year, the facility provides systems 10 to 100 times more powerful than most computers available for research to international teams of scientists who make breakthroughs in biology, chemistry, materials science, cosmology, industrial modeling, and more. Publication productivity for the OLCF has grown to more than 460 publications per year in 2016, with a large fraction in leading journals and conferences. This report shares just a few influential success stories of OLCF user projects.

When CCS was founded in 1992, researchers studying environmental modeling and remediation, material properties, and the quantum structure and dynamics of materials were calling for major computing resources. Within 3 years of deploying supercomputers at CCS, ORNL researchers developed an award-winning parallel computing code for materials modeling (p. 24 and 32), and visiting researchers from Rice University developed parallel computing algorithms for modeling the flow of groundwater mixed with contaminants through rocks and soil (p. 20).

By 2003, scientists were taking supercomputing to the cosmos with unprecedented simulations of the collapse of massive stars, resulting in one of the first discoveries to be predicted entirely through modeling and simulation—the existence of a shock wave that contributes to stellar explosion (p. 14). Once the OLCF stood up the Cray XT Jaguar series, moving from teraflops to petaflops, research accelerated in combustion (p. 12), biofuels (p. 10), cosmology (p. 28), and climate (p. 26). In recent years, OLCF users have made huge strides in simulating thousands of unknown nuclei (p. 22) and more realistic earthquake frequencies for high-risk regions like Southern California (p. 18).

More than a dozen companies, from innovative small businesses to Fortune 500 giants, use OLCF resources each year to find new ways to get an edge on the competition. Industry partners are increasingly looking to supercomputers to accelerate research and development—such as SmartTruck using Jaguar to develop an add-on for long-haul trucks to optimize airflow, improving fuel mileage by 7 to 10 percent (p. 16), or Procter & Gamble using molecular dynamics simulations to extend the shelf life and stability of consumer products (p. 30).

With the deployment of Summit, users will have a machine designed not only for fast and complex modeling and simulation but also for powerful data analytics and machine learning. On Titan, researchers are already mining meaningful connections in cancer data and plant genetics, preparing for a new decade of data science and discovery on Summit and exascale machines. The next 25 years could prove to be the most exciting yet.

Education and Outreach

In the early 1990s, the PICS proposal emphasized the importance of establishing outreach and education programs at CCS to provide training in the computational sciences. Rather than focusing on a single pathway, the education and outreach effort of CCS involved multiple programs that aimed to build a pipeline from elementary to graduate-level students—and ultimately to a proficient computer science community. Within the education effort, CCS placed a high value on diversity and sought to target minorities and disadvantaged students from its very beginning.

Throughout the 1990s, CCS promoted a hands-on program called Adventures in Supercomputing that taught critical thinking in a multidisciplinary, multisubject setting. In collaboration with faculty members from ORNL, Ames Laboratory, and Sandia National Laboratories, teachers from more than 70 high schools integrated HPC concepts into their teaching materials.

In 1991, DOE sponsored the Computational Science Education Project, which relayed overviews of computational science tools to undergraduate and graduate students in the life and physical sciences via a comprehensive e-book, featuring clickable elements, summaries, and visualization packages. Dozens of researchers from universities as well as CCS and other government agencies contributed to the e-book. CCS also created the Undergraduate Computational and Engineering Sciences project, which created a broad range of educational materials for secondary-school and undergraduate students.

DOE established its flagship education program in 1991—the Computational Science Graduate Fellowship (CSGF), which gives doctoral students the opportunity to use computing resources while completing computer science summer practicums. The OLCF remains an active participant in the program to this day, hosting CSGF fellows for their practicums, facilitating workshops at CSGF Annual Review meetings, and hiring program alumni. Since 1991, CSGF alumni have built the foundation for the computer science community and will continue to pass the torch to the next generation of scientists.

With the advent of new computing resources such as Titan, the center recognized the need for an accessible education and outreach tool. Tiny Titan, a demonstration parallel computer, was built in 2014 to provide a portable tool for science education. The machine, featuring eight Raspberry Pi processors, has been shared with thousands of students at the lab and other events—including National Science Bowls, science fairs, and festivals. Building on Tiny Titan's popularity, the OLCF is creating a similar demonstration system for the Summit supercomputer.

The lasting legacy of the center's initial education programs lies in the past, present, and future contributions of the people touched by these programs; their contributions make the center's education programs both possible and successful. It also lies in the large amount of information and tutorials the OLCF provides through its website, teleconferences, and in-person workshops. These resources continue to make HPC accessible to audiences in STEM fields and beyond. Improving computer literacy and scientific knowledge at the OLCF and in the computer science community as a whole remains a major objective of the center as it continues to strive for excellence through education, training, curriculum materials, and outreach activities to create the next generation of leading scientists.

"This detailed knowledge of lignin behavior can guide genetic engineering of enzymes that bind less to lignin and therefore produce bioethanol more efficiently."
—Loukas Petridis,
Oak Ridge National Laboratory

Breaking Down Biomass

Since 2007, a team led by ORNL's Jeremy Smith has used OLCF systems to provide molecular-level insights to help biofuel researchers overcome hurdles to producing cost-competitive cellulosic ethanol fuel and other high-value products from woody plants and waste biomass.

Advanced biofuels hold the promise of reducing US greenhouse gas emissions and fossil fuel dependence with a sustainable energy source that doesn't compete with food production and contributes to rural development. Realizing this potential, however, requires more efficient processes for converting plant matter into fuel.

Through a series of simulations—ranging from thousands of atoms to more than 20 million—Smith's team captured key interactions among major components of the plant cell wall, along with the enzymes and solvents used in biofuel production. These results, obtained using the Cray XT5 Jaguar and Cray XK7 Titan, help explain and point toward strategies to mitigate a tenacious obstacle to ethanol production: lignin, a ubiquitous structural component of plants that blocks the enzymatic breakdown of biomass.

The Science

In nature, lignin helps provide the scaffolding for plant cells, reinforcing the cell wall's slender cellulosic fibers—the primary raw ingredient of cellulosic ethanol—and serving as a protective barrier against disease and predators. This hardy polymer, however, also makes plants stubbornly resistant to breakdown, or hydrolysis. Therefore, to hydrolyse biomass and make it available for fermentation by microbes, it must first be subjected to large amounts of acid, water, and heat, adding cost and complexity to the process.

Understanding the structure of lignin is critical to devising rational strategies to overcome its adverse effects. So to better understand the lignin problem, Smith's team first resolved the structure of lignin down

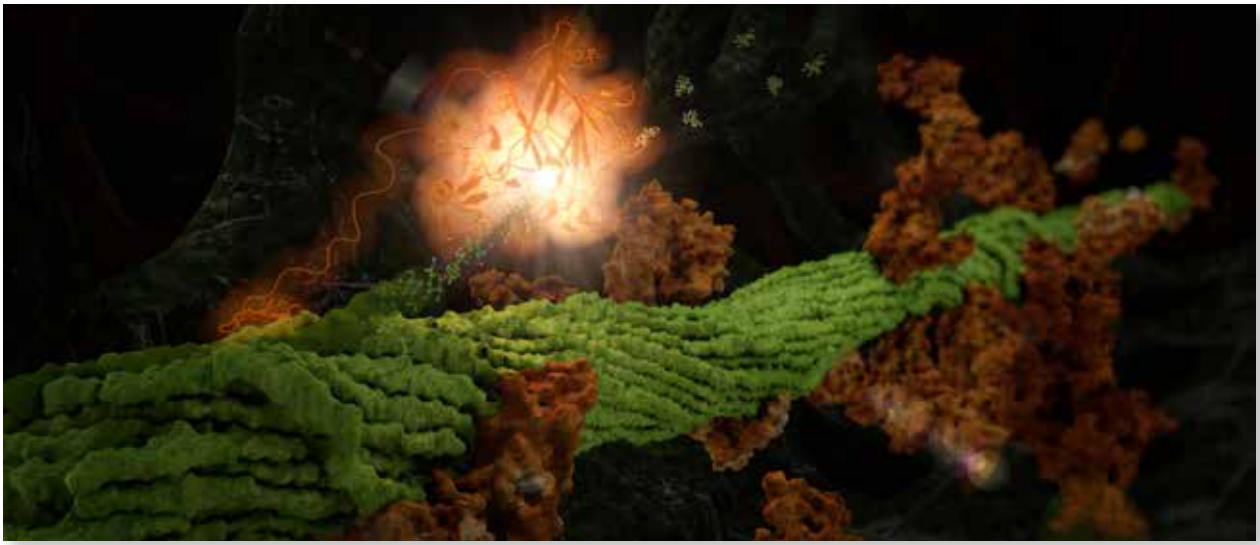

A team at ORNL performed its largest biological simulation to explain why lignin is so potent in blocking the enzymes that break down cellulose. Here, an enzyme (orange) hydrolyzes cellulose (green) despite the presence of lignin (brown). Image credit: ORNL/scistyle.com

to 1 angstrom (equal to 10 billionths of a meter, or smaller than the width of a carbon atom) with the help of Jaguar in 2011. The project revealed the rough and folded surface of the polymer in unprecedented detail.

The team transitioned its research to Titan, on which its atomistic simulations expanded into a representative model of pretreated biomass made up of lignin, cellulose, and enzymes. The simulations in 2014 demonstrated exactly how lignin binds to cellulose and enzymes to block hydrolysis. Furthermore, the team identified a pathway to circumvent the problem, finding that amorphous, or less-ordered, cellulose fibers interact less with lignin and therefore are more accessible to enzymes.

In addition to simulating conventional biomass, Smith's team used supercomputing resources to provide context to biofuel experiments funded by DOE's BioEnergy Science Center. In support of one such experiment, the team simulated a promising genetic modification to the lignin of an aspen tree and thus explained why it worked. In another instance, the team subjected an experimental solvent to the OLCF's "computational microscope" to explain why pretreatment with the solvent was more than three times as effective as conventional dilute acid when applied to maple wood. Both cases highlighted strategies for removing lignin with increased efficiency.

The Legacy

As the first group to perform supercomputer simulations on biomass, Smith's team demonstrated the value of identifying molecular mechanisms to explain problems in multicomponent molecular systems. Leveraging these insights, biofuel researchers have accelerated progress toward the production of economically viable cellulosic ethanol, a vision that predates the world's first mass-produced automobile.

The team's work also spearheaded a shift in computational biophysics from individual molecules and proteins toward more complex, multicomponent systems. This transition was largely enabled by petascale resources, beginning with the OLCF's Jaguar, that continue to grow, allowing for the simulation of biomolecular systems of even greater size and complexity.

Mostofian, B., et al. (2016), "Local Phase Separation of Co-solvents Enhances Pretreatment of Biomass for Bioenergy Applications." Journal of the American Chemical Society, Volume: 138, no. 34. DOI: 10869–10878.

"If low-temperature compression ignition concepts are widely adopted in next-generation autos, fuel efficiency could increase by as much as 25 to 50 percent."
—Jacqueline Chen,
Sandia National Laboratories

Detailing Combustion

A team led by Jacqueline Chen at Sandia National Laboratories (Sandia) has used OLCF resources since 2005 to conduct simulations of turbulent combustion, which can help engineers develop better predictive models for fuel-efficient engines.

Using the OLCF's Cray XT4 and XT5 Jaguar and Cray XK7 Titan, Chen's team has performed direct numerical simulations (DNS) of combustion phenomena, including fuel-air mixing, autoignition, and flame dynamics at multiple scales. Chen is particularly interested in combustion processes that more efficiently distribute heat and lower nitrogen oxide emissions.

By detailing the important mechanisms that can contribute to more efficient combustion processes, the team is helping pave the way for automobiles that could use 25 to 50 percent less fuel than those on the road today.

The Science

During diesel engine combustion, air and fuel mix violently and react, generating heat that causes the mixture to spontaneously ignite (autoignition). Finding the right conditions for efficient combustion is tricky, and simulating these conditions on multiple scales is impossible without a leadership-class resource. Multiple factors—such as how much fuel and air is added, how much mixing takes place, and whether combustion occurs as spontaneous ignition or propagating flames—contribute to this complex process.

Using the DNS code S3D, which solves for compressible, reacting flows using Navier–Stokes equations, Chen's team explored these factors, capturing length scales ranging from micrometers—the size of some molecular interactions—to centimeters, which allows for the nuances of fuel-air mixing and chemical reactions to be studied with unprecedented fidelity.

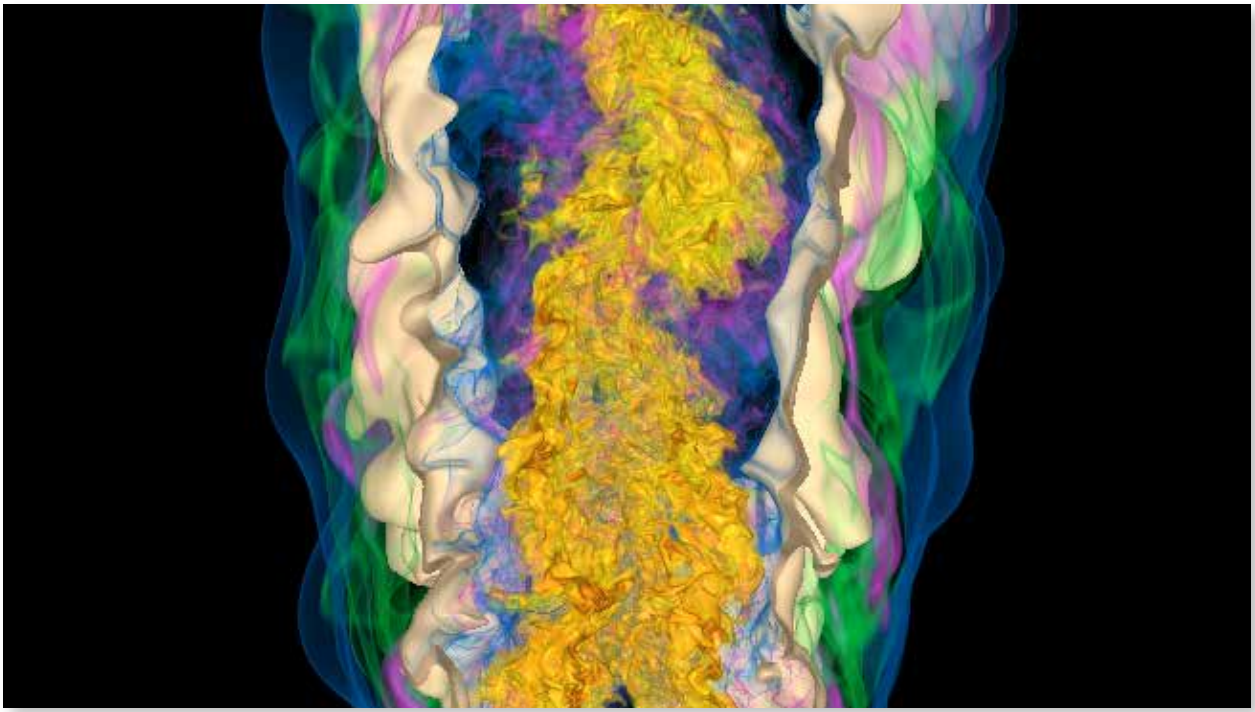

A visualization of the mixture fraction isosurface (gold) and HO_2, which shows autoignition occurring in fuel-lean mixtures at regions with high temperatures, low mixing rates, and short ignition delay times. Image credit: Hongfeng Yu, University of Nebraska; and Kwan-Liu Ma, University of California, Davis

In 2006 Chen's team used Jaguar to simulate a 3D turbulent lean methane–air premixed flame. Lean premixed combustion is of interest to the engineering community because lean burning distributes heat efficiently and generates lower nitrogen oxide emissions. The team discovered that intense small turbulent eddies can cause a broadening of the flame's preheat zone, an area that continuously heats a flame so that it does not die once ignited. This same turbulence, though, does not disrupt the thin areas where chemical reactions are occurring. The DNS data from this study were used to assess models of the degree of flame wrinkling, which affects the overall burning rates.

Since autoignition can occur simultaneously with turbulent mixing of cold gaseous fuel and heated oxidant, the team also simulated a 3D lifted flame to study how a flame stabilizes above a burner. By capturing the full range of turbulence scales from the largest to the smallest eddies, the team provided the first fully-resolved simulation of a lifted turbulent hydrogen flame stabilized by autoignition, concluding that autoignition was the key mechanism responsible for flame stabilization.

The Legacy

The simulation of a 3D lifted hydrogen flame provided a large dataset for engineers to use as they develop the predictive models that will optimize the internal combustion engines and gas turbine designs that could reduce emissions and increase efficiency. The team's most recent work continues to uncover new details about combustion processes, thanks to the wide range of scales—both in space and time—made available by leadership-class computing.

Chen is the founding director of the Center for Exascale Simulation of Combustion in Turbulence—an interdisciplinary center focused on codesigning the software and tools necessary for the future of combustion modeling and simulation in the exascale era. Her work is setting the stage for combustion simulation on exascale computing systems.

Sankaran, et al. (2007), Structure of a Spatially Developing Turbulent Lean Methane–Air Bunsen Flame, Proceedings of the Combustion Institute, Volume: 31, no. 1. DOI: 10.1016/j.proci.2006.08.025.

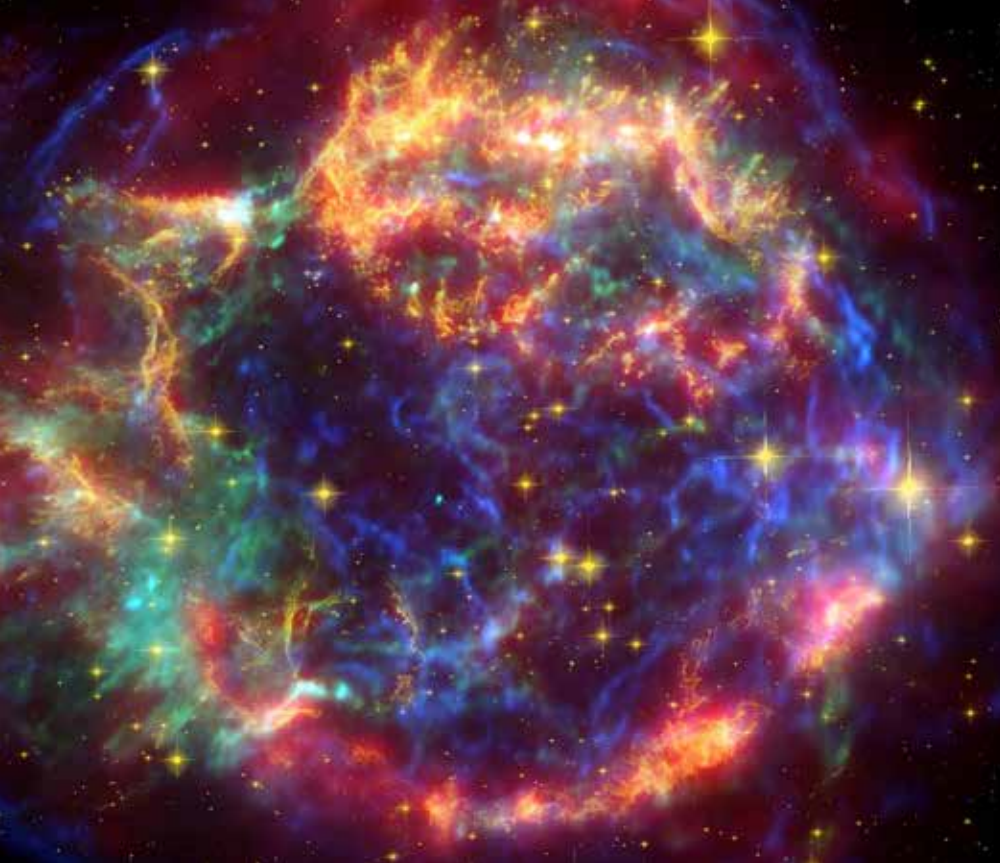

"Once we discovered the SASI, it became very much a part of core-collapse supernova theory. People feel it is an important missing ingredient."
—Tony Mezzacappa, University of Tennessee–Oak Ridge National Laboratory

Revealing Supernova Secrets

In 2003, a group of computational astrophysicists led by ORNL's Anthony Mezzacappa proved supercomputers can be more than tools for understanding. They can also be tools for discovery.

The team's simulations of the death of a massive star, known as a core-collapse supernova, revealed a previously unknown phenomenon—a shock-wave-distorting feature that emerges in the early stages of a star's demise and contributes to its eventual explosion. The team named the discovery the "stationary accretion shock instability," or SASI. After identifying the feature in 2D through work partly conducted on the OLCF's IBM Power3 Eagle, Mezzacappa and company went on to study SASI in 3D detail using systems such as the Cray X1 Phoenix and Cray XT3 Jaguar.

In 2014, a separate team of astronomers reported observational evidence from high-energy x-ray telescope data that supported the SASI model. The corroboration, concluded from the radiation signature of the Cassiopeia A supernova remnant gathered by the National Aeronautics and Space Administration's (NASA's) NuSTAR telescope, is a rare example of simulation predicting a physical phenomenon before it is observed in nature.

The SASI discovery supplied a key piece of missing information in supernova theory, bringing scientists one step closer to understanding stellar explosions and, by extension, the source of most of the elements that make up the world around us, including life on Earth.

The Science

After the energy-producing processes of nuclear fusion have ceased in a massive star, the star's iron core collapses to densities greater than those found in an atomic nucleus. The core "bounces" in response, resulting in a rebounding shock wave that pushes outward. Eventually, the shock wave blows most of the star's mass violently into space. In the process, new elements are produced and scattered throughout the cosmos, seeding the universe with the raw material of new stars, planets, and the life itself.

Between the bounce and the bang, however, lies a mystery of extreme physics scientists have been endeavoring to solve using the most powerful tools available. Only the most advanced supercomputers can supply the processing power necessary to calculate essential supernova physics in a realistic time frame, with some of the team's simulations consuming tens of millions of processor hours.

Through work conducted on OLCF systems, Mezzacappa's team contributed two significant pieces to the core-collapse puzzle. The team's initial discovery demonstrated that the shock wave—stalled by material falling inward toward the star's core—is unstable, causing it to become distorted. This distortion results in two rotating flows—one aligned with but directly below the shock wave and an inner flow that travels in the opposite direction. The result is a coherent, cigar-shaped slosh.

Additionally, the SASI phenomenon supplied a physics-based explanation for the spins of some pulsars, spinning neutron stars left behind in a supernova's wake. Before SASI, astronomers had assumed the spin of a pulsar was a result of the spin of the original star. If that were true, however, a pulsar would spin much faster than its larger forerunner, because of the same forces that cause an ice skater to spin faster by pulling her arms in.

SASI provided an alternative theory for pulsars with millisecond-scale spin periods, suggesting that the settling of SASI's inner flow onto the dying star's core could serve as the spinning mechanism.

The Legacy

Today, SASI is a fundamental component of high-fidelity supernova simulations and a widely accepted concept of core-collapse supernova theory. Its computational discovery stands as a testament to supercomputing's ability to probe some of the most challenging problems in science.

The core-collapse supernova simulations conducted on OLCF systems and subsequent confirmation by astronomers validated the efforts of the computational astrophysics community and its ability to capture the complex physical mechanisms of exploding stars.

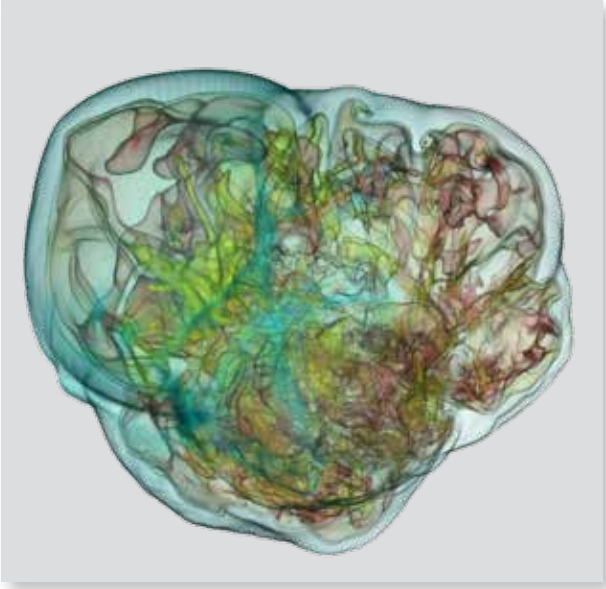

Visualization of the entropy of the inner 250 kilometers of a 15 solar-mass star during a 3D simulation of a core-collapse supernova using the CHIMERA code. Large-scale distortion of the supernova shock can be seen, along with smaller-scale convection. Image credit: Michael Matheson, ORNL

Since then, supernova models have built on SASI and other successes to incorporate greater realism by leveraging subsequent leaps in computing power, such as the GPU-accelerated architecture of the OLCF's Cray XK7 Titan. This work has resulted in improved nuclear reaction networks—the range of elements tracked throughout a simulation—and understanding of neutrino physics, which are thought to play a central role in the star's explosion.

Ultimately, the community is working toward a complete explanation of how the collapse of a star's core leads to the explosion that ejects most of its layers. If history is any guide, scientists may find this answer on Earth before seeing it in space.

Blondin, John M.; Mezzacappa, A.; DeMarino, C. (2003), "Stability of Standing Accretion Shocks, with an Eye toward Core-collapse Supernovae," The Astrophysical Journal, Volume: 584, no. 2. DOI: 10.1086/345812.

"Using Jaguar and FUN3D has provided us with many 'aha moments' and this was one of them."
—John Anastos, SmartTruck

Building a SmartTruck

To increase the fuel efficiency of Class 8 long-haul trucks small business SmartTruck Systems (SmartTruck) used the OLCF's Jaguar supercomputer in 2010 to better understand the complex airflow around these vehicles so the company could create unique add-on components that dramatically reduce fuel-sapping drag.

At that time, these so-called "18-wheelers" averaged only 6 mpg or less and the million-plus in operation consumed more than 1.5 million barrels of fuel per day. New fuel mileage standards were looming, emphasizing the need for technology that would increase truck fuel efficiency so that semis could remain compliant with the regulations.

In 2010, the SmartTruck company (at the time, a start-up), undertook a project led by then-CEO Mike Henderson and current SmartTruck engineer John Anastos to successfully employ Jaguar to develop a detailed and highly accurate long-haul truck model and simulate the airflow around it, revealing the truck's complex aerodynamics. Using this new understanding, the team designed add-on parts that redirected the air around the vehicle to dramatically decrease drag. Trucks retrofitted with the parts operated at an increased highway fuel efficiency of up to 12 percent, allowing them to meet the new regulations as well as achieve an estimated annual fuel savings of $5,000 per truck.

The Science

When the California Air Resource Board (CARB) announced aggressive new fuel efficiency regulations for long-haul trucks, Henderson saw a business opportunity to make 18-wheelers more aerodynamic. The same had already been done for automobiles and planes via modeling and simulation, and Henderson and his team felt confident that similar computational techniques could be applied to long-haul trucks.

Using the OLCF's Jaguar and NASA's Fully Unstructured Navier–Stokes code called FUN3D, a computational fluid dynamics (CFD) software suite that typically models aircraft aerodynamics, the team

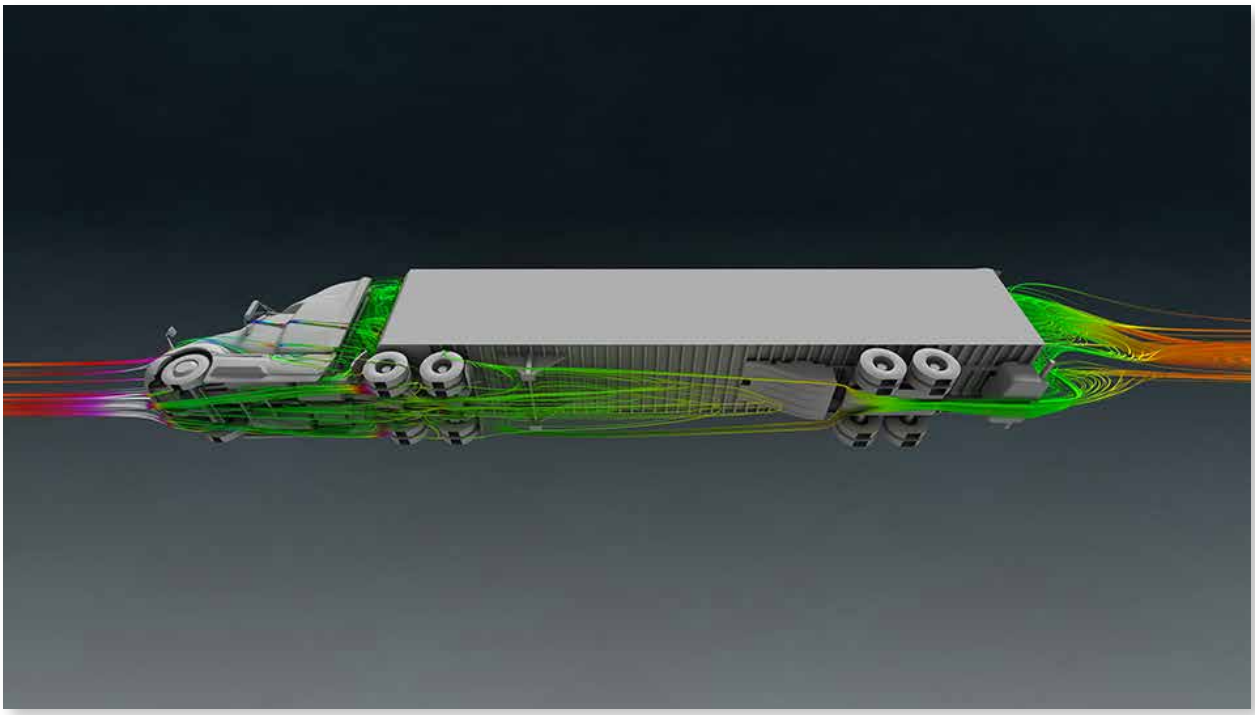

Supercomputing simulations at ORNL enabled SmartTruck engineers to develop the UnderTray System, some components of which are shown here. The system dramatically reduces drag—and increases fuel mileage—in long-haul trucks. Image credit: Michael Matheson, ORNL

developed one of the most detailed (and therefore the most accurate) numerical models of a Class 8 tractor and trailer to date.

The team broke the truck into hundreds of pieces and examined the drag on each piece to determine how they interacted across the entire system. The simulations revealed the details needed to create an optimized "UnderTray System" of fuel-saving add-on parts, such as aerodynamic wheel fairings and a special airflow-diverting sled that attaches to the axels on the underside of the vehicle. The system's biggest component, the rear diffuser, was a major factor in optimizing airflow and boosting fuel efficiency. The researchers also redesigned the rain gutter to recapture air, create a healthy airflow boundary layer, and direct the flow down into the truck's trailing wake to improve operating efficiency.

Access to Jaguar permitted the SmartTruck team to run models approximately 10 times faster than was possible on their internal systems, dramatically accelerating their understanding of truck aerodynamics and reducing the time to create the new parts. The physical tests, results were within one percent of the CFD calculations, validating the models and affirming the predictive accuracy of simulations and the operational efficiency of the new add-on component designs based on those simulations. This reduced the need to develop a time-consuming and costly series of physical prototypes, enabling SmartTruck to reduce the time from concept to a manufacture-ready design by 50 percent—from 3 years to 18 months.

The Legacy

SmartTruck's products were introduced successfully to the market and were named in Heavy Duty Trucking Magazine's list of Top 20 products in 2011, 2015, and 2017. By providing SmartTruck Systems with advanced computational resources, the OLCF demonstrated how its industrial partnerships can advance understanding of important science and boost the competitiveness of US businesses.

Henderson, Michael Lorenzo and Stephen Alfred Wulff. "Devices and Methods for Reducing Vehicle Drag." SmartTruck Systems, LLC, assignee. Patent US 8342595 B2 (1 Jan 2013).

"We are at a point now where computers can do these calculations using physics and improve our ability to do the type of analysis necessary to create a safe environment for society."
—Thomas Jordan,
Southern California Earthquake Center & University of Southern California

Shaking up Earthquake Preparedness

Since 2010, a multi-institutional team led by Thomas Jordan of the Southern California Earthquake Center (SCEC) has been using OLCF supercomputers to simulate earthquakes at higher wave frequencies and to assess the hazards posed by seismic activity to cities and critical infrastructure.

Beginning with the Cray XT5 Jaguar and continuing on the Cray XK7 Titan, Jordan's team carried out large-scale simulations to predict the intensity of ground shaking at specific sites in earthquake-prone Southern California. The results from individual earthquake simulations help inform and improve SCEC's CyberShake platform, the first physics-based probabilistic seismic hazard model, which requires large ensembles of earthquake simulations. City planners, structural engineers, and emergency preparedness officials, in turn, are considering CyberShake results for decision-making processes to update infrastructure, building codes, and emergency response plans.

SCEC's advanced earthquake system science projects can help mitigate the effects of large earthquakes (including those on the next well-known San Andreas Fault), saving lives and curtailing the cost of damage by providing more accurate seismic hazard information for California. Furthermore, the modeling tools created by Jordan's team can be applied to other at-risk regions around the world.

The Science

When an earthquake strikes, its destructive effects on nearby areas can vary greatly based on a range of factors, including the quake's strength, its depth of origin, and the terrain through which it travels. Scientists and engineers measure the frequency of seismic waves, the energetic vibrations that ripple through the Earth, to determine how the shaking will affect buildings and other societal infrastructure.

Seismic wave frequency, which is measured in hertz (cycles per second), is important to engineers who are designing buildings, bridges, and other infrastructure to withstand earthquake damage. Low-frequency waves, which cycle less than once per second (1 hertz), are easier to model. Tall buildings (20 stories or higher)

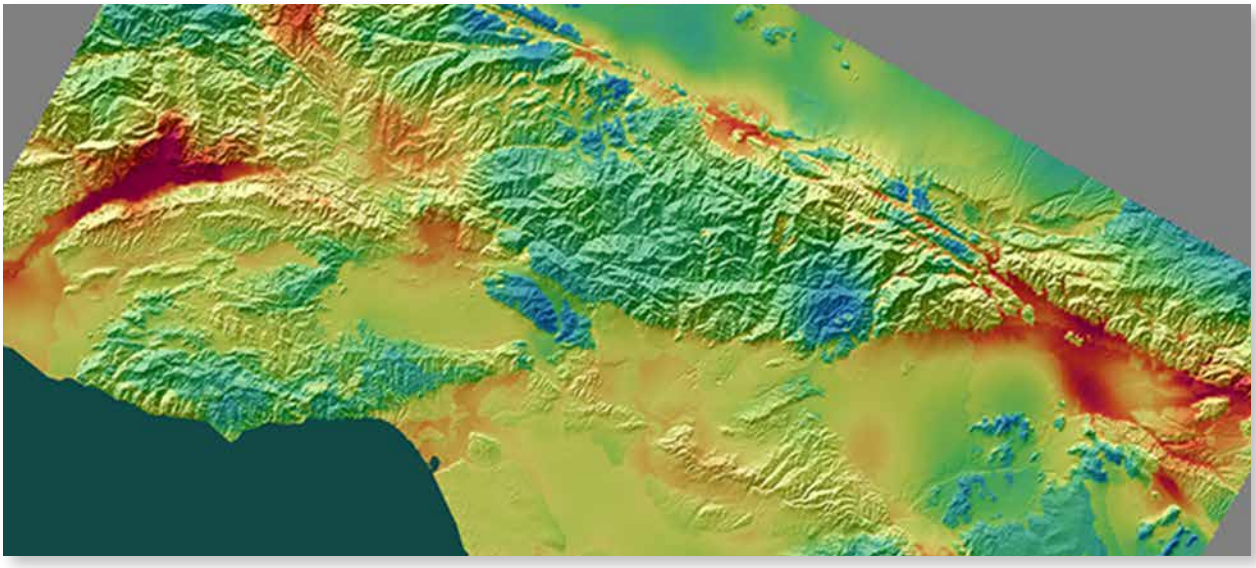

The CyberShake seismic hazard map shows the magnitude for the Los Angeles region, defined by the amount of change of a surface or structure in a 2-second period. The map provides engineers with vital information needed to design more seismically safe structures. Image Credit: Scott Callaghan, Kevin Milner, and Thomas H. Jordan; Southern California Earthquake Center

are most affected by shaking at these frequencies. High-frequency waves (in the 2–10 hertz range), however, are more difficult to simulate, requiring extremely detailed models and significantly more computing resources. These types of quakes affect smaller structures like hospitals, highways, and homes.

To predict the effects of high-frequency quakes in Southern California, Jordan's team has calculated the propagation and ground motion of waves radiating from the San Andreas Fault using a 3D model of the Earth's crust. These simulations require highly realistic earth models and wave propagation software which account for factors such as topography, gradual energy loss, and small-scale heterogeneities.

In 2010, the SCEC team used the 1.75-petaflop Jaguar to simulate a magnitude 8 earthquake along the San Andreas Fault. Those simulations peaked at 2 hertz. At the time, the Jaguar simulations were conducted, doubling the wave frequency would have required a 16-fold increase in computational power. But on Titan in 2013, the team was able to run simulations of a magnitude 7.2 earthquake up to its goal of 10 hertz. By modifying its CPU-based code, AWP-ODC, for Titan's GPU-accelerated architecture, the team significantly improved speedup, producing higher-resolution results in a fraction of the time.

In 2015, Jordan's team released an updated CyberShake map for the Southern California region that incorporated simulations for 336 separate locations in the region and doubled the maximum simulated frequency from 0.5 to 1 hertz. The map marked a significant step forward in supplying policymakers with an unprecedented tool for weighing the risks of seismic activity. In 2017, the team used Titan to calculate the first 1 hertz CyberShake physics-based seismic hazard model for Central California.

The Legacy

As more people move to cities in seismically active regions, the potential costs of a devastating earthquake will continue to increase. The techniques being pioneered by Jordan's team are improving scientists' understanding of earthquakes, helping civic leaders prepare for rare, disruptive events, and giving earthquake engineers new tools for predicting strong ground motions and their effects on the built environment.

Impact studies following SCEC's scenario simulations became a motivation for the City Of Los Angeles to develop its new city-wide earthquake resilience and preparedness plan, Resilience by Design. And, although earthquake prediction is still in its infancy, the simulation data being produced by Jordan's team is being used to train machine-learning algorithms as part of California's Earthquake Early Warning system, potentially increasing the time citizens have to respond to the imminent threat of a quake.

Jordan, T. H. (2015), Societal Impact of Earthquake Simulations at Extreme Scale, in Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, ACM.

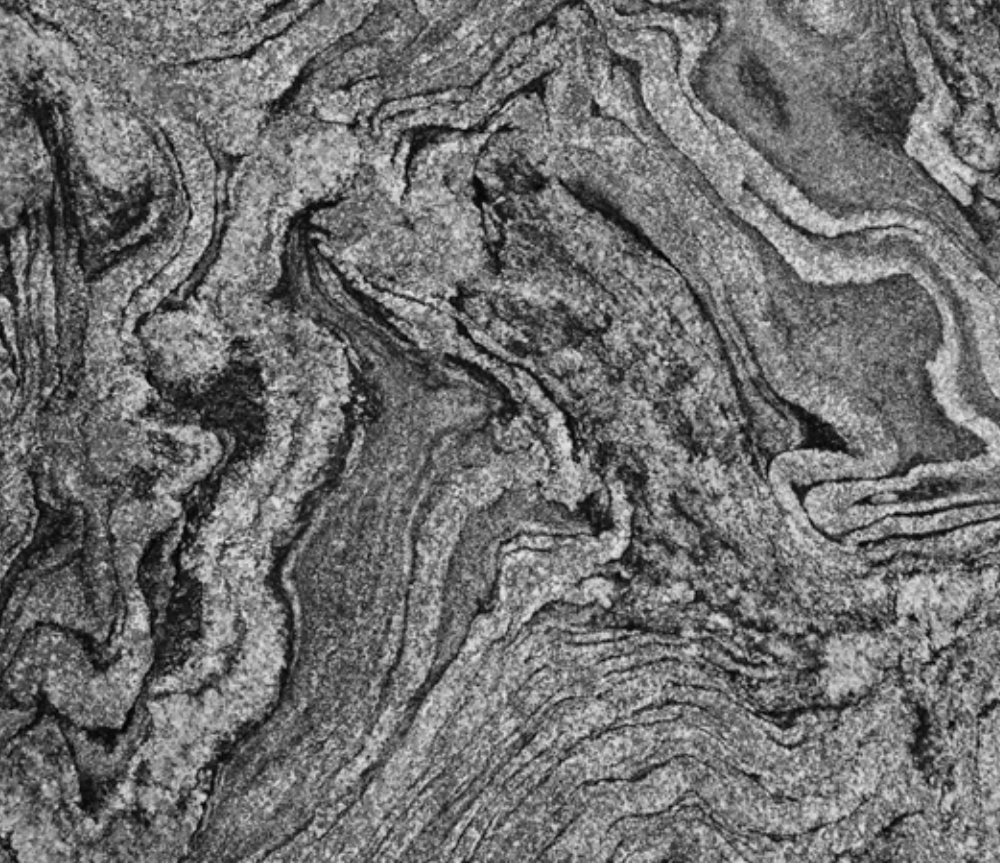

"The two-phase flow codes our team needed to begin solving these problems weren't available, but we created methods and algorithms for groundwater flow and transport that are still used today to help scientists answer important groundwater questions."
—Mary Wheeler, University of Texas

Unearthing the Subsurface

In the 1990s, a multi-institutional team used CCS supercomputers to develop the first parallel algorithms for understanding how different substances interact underground. Specifically, the team was interested in groundwater flow and transport in porous media like soil and rocks.

At the time, modeling how groundwater mixed with and transported contaminants below the surface was essential to correcting problems caused by radioactive waste left over from nuclear weapons research conducted during World War II. The problem was so important that ORNL selected it as one of three grand challenges to be addressed in its proposal for an HPC center. The proposal outlined the creation of new models that would help scientists assess risks at contamination sites and evaluate remediation methods.

When ORNL's CCS was founded in 1992, a current faculty member at the University of Texas–Austin (UT–Austin), Mary Wheeler, and her team pioneered the effort to create the algorithms necessary for simulating groundwater processes, giving scientists a first look at some of the complex interactions that occur during groundwater contamination and remediation.

Wheeler's team used the IBM SP2 and Intel Paragon systems at CCS to develop groundwater models and codes of unprecedented complexity. The research laid the groundwork for studying groundwater contamination problems on parallel computing systems. Scientists still use these codes today as they search for contaminant removal strategies and prevention plans for some of the world's most complex groundwater issues.

The Science

From 1950 to 1980, radioactive waste was disposed of in unlined trenches at ORNL under the assumption that geochemical processes would slow any leakage. However, radionuclides—atoms that emit harmful nuclear energy—flowed from disposal sites into local streams, contaminating the environment.

Wheeler and her collaborators created algorithms that could model the kinds of factors that led to the leakage, forming the basis for the first parallel groundwater simulation codes. In 1995, Wheeler helped develop

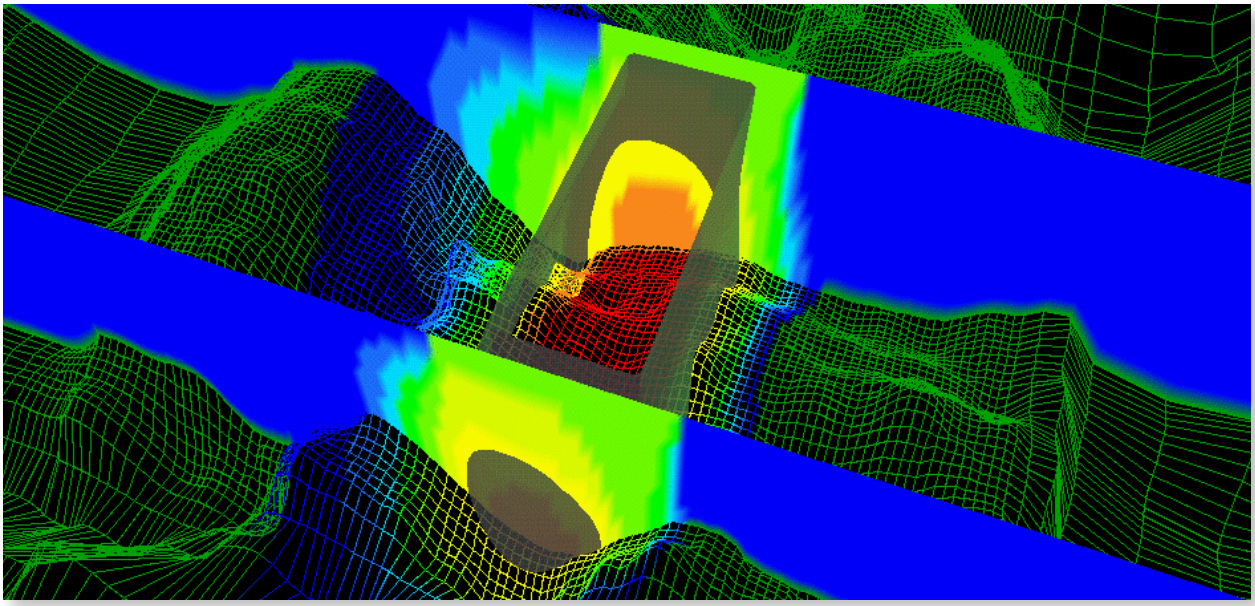

The PICS Groundwater Contaminant Code GCT 1.0 was released in 1994 to simulate saturated/unsaturated flow and reactive transport using massively parallel architectures. Here, a GCT simulation showing migration of a high-density contaminant in a clayey aquifer. Image credit: Texas A&M University

a parallel algorithm for two-phase (air and water) systems of groundwater contaminant transport. The algorithm used the Godunov method and mixed finite element methods, which can accurately account for the mass of a system at a given point in space and time, as well as model grids that change over time, as is the case in porous media transport. The methods used in this algorithm were then implemented into groundwater codes designed in part by the Center for Subsurface Modeling at UT–Austin.

Two of these codes—ParSSim, the Parallel Subsurface Simulator, and GCT, the Groundwater Contaminant Transport simulator—were tested and run on the massively parallel IBM SP2 and the Intel Paragon XP/S 35 supercomputers at CCS, with the help of former ORNL researcher Laura Toran. The ability to generate 3D models of subsurface radioactive decay and the transport of radioactive materials via groundwater provided scientists with new tools to study groundwater patterns and remediation strategies.

In 1998, Wheeler became part of a simulation project based on the radioactive contaminants found at the ORNL site. The team used ParSSim and another geochemistry code called PHREEQC to model how concentrations of strontium, a common waste product of nuclear research, change over time under different levels of acidity in a groundwater system. The team ran ParSSim on the Intel Paragon XP/S 5—a supercomputer donated to Rice University by CCS—using a 1D 100-cell model. The project resulted in the creation of the first model that used a two-tiered approach of simple and complex methods to explain interactions between contaminants and porous media in groundwater transport.

The Legacy

The algorithms designed by Wheeler and her collaborators ultimately became valuable tools for simulating the complex interactions in groundwater contaminant transport processes. Today, the ParSSim code is used at the Center for Subsurface Modeling at UT–Austin to model single-phase flows and the reactions of multiple chemical substances in groundwater systems.

Wheeler and her colleagues set the precedent for modeling flow and transport in porous media using Godunov and mixed finite element methods, bringing scientists a step closer to understanding and solving some of the world's most difficult groundwater contamination problems.

In 2013 Wheeler received the International Society for Porous Media, Interpore's Honorary Lifetime Membership Award for her immense contributions to the field of geosciences.

Arbogast, T., Dawson, C. N., and Wheeler, M. F. (1995), A Parallel Algorithm for Two Phase Multicomponent Contaminant Transport, Applications of Mathematics, Volume: 40, no. 3.

"Over the years we've tried to improve the models of the nucleus to include more and more knowledge and insights. We are building a nuclear model based on the best theoretical input guided by the best experimental data."
—Witek Nazarewicz,
Michigan State University

Mapping the Nuclear Landscape

In 2012, an ORNL and University of Tennessee team used the Cray XK6 Jaguar supercomputer to calculate the number of isotopes allowed by the laws of physics.

The team, led by Witold Nazarewicz, applied a quantum approach known as nuclear density functional theory (DFT) to six leading models to determine that there are about 7,000 possible combinations of protons and neutrons allowed in bound nuclei with up to 120 protons.

Most of these nuclei have never been observed experimentally. Of the total, about 3,000 have been found in nature or produced in a lab. Many others are expected to be created in massive stars and violent stellar explosions.

By carrying out large-scale computations, the team provided nuclear physicists with a more detailed roadmap that marks the so-called "drip lines," the boundaries of nuclear existence. This guide lends insight to researchers exploring exotic, weakly bound nuclei that directly impact the way elements are produced in stars. Additionally, the calculations point toward unexplored regions of the nuclear chart that may contain isotopes with properties useful in health care, material science, and nuclear power, among other applications.

The Science

For each number of protons in a nucleus, there is a limit to how many neutrons are allowed. For example, a helium nucleus, which contains two protons, can hold no more than six neutrons. If another neutron is added, it will simply "drip" off. Likewise, there is a limit to the number of protons that can be added to a nucleus with a given number of neutrons. The closer the isotope is to one of the drip lines, the faster it decays into more stable forms.

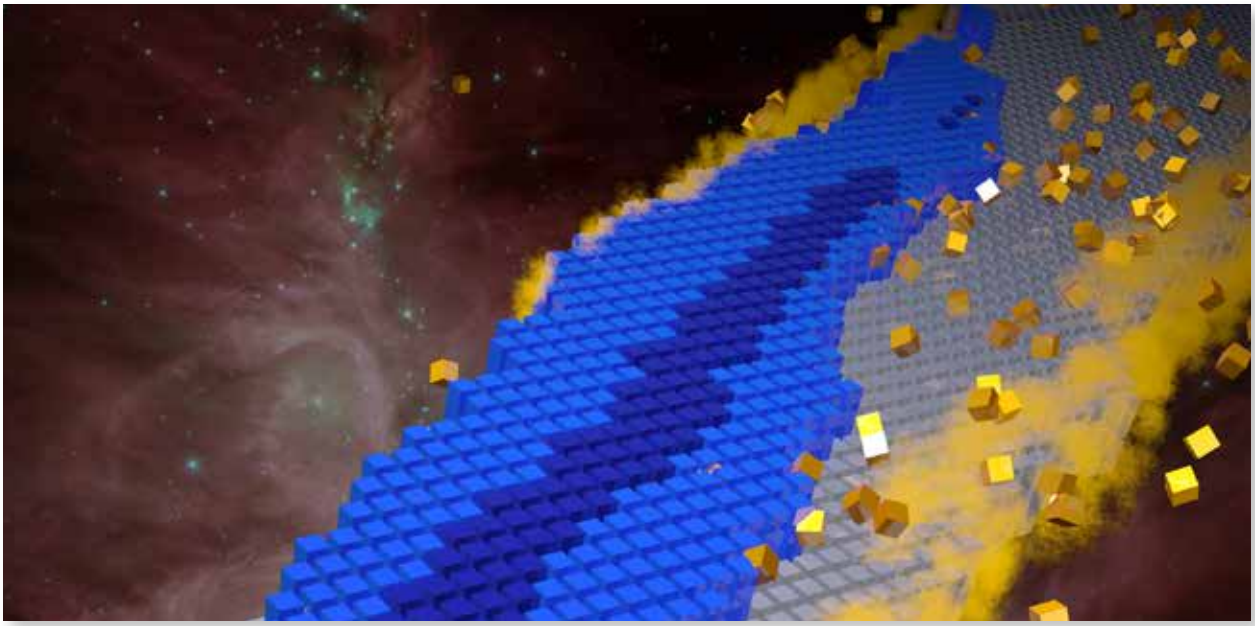

An illustration of the nuclear landscape, with isotopes arranged by an increasing number of protons (up) and neutrons (right). The dark blue blocks represent stable isotopes. The lighter blue blocks are unstable isotopes that have been observed. The gray blocks are bound isotopes that have not been observed. Nuclear existence ends at the drip lines (orange clouds), where there is no longer enough binding energy to prevent the last nucleons from dripping off (floating blocks). Image credit: Andy Sproles, ORNL

Identification of the drip lines for heavier elements is based on theoretical predictions well beyond the reach of experiment and is, therefore, uncertain. Nazarewicz's team carried out calculations based on the microscopic forces that cause neutrons and protons to cluster into nuclei to predict where exotic nuclei might exist. The calculations themselves were massive, with each set of nuclei taking about 2 hours on the 244,256-processor Jaguar system.

Applying DFT to six nuclear energy density functionals, the team found its results to be surprisingly consistent across models. By using several models, theorists were able to quantify uncertainties for the predicted drip lines for the first time, a measurement that gauges the accuracy of the predictions. The accomplishment served as one of the capstones of a 5 year multi-institutional collaboration under the DOE Office of Science's Scientific Discovery through Advanced Computing (SciDAC) program to develop a predictive theory to determine the properties of atomic nuclei using petascale systems.

The Legacy

The team's nuclear landscape study represented a milestone in the nuclear physics community with regard to both its scale and its scope. Since 2012, other groups have followed in the Nazarewicz team's footsteps, producing similar findings. Additionally, the team pioneered the practice in nuclear physics of providing theoretical predictions with an estimated error range, an idea that is becoming standard in the field.

As more powerful particle accelerators come online, the simulation results will continue to guide scientists' efforts to confirm the existence of exotic isotopes through experiment. This coupling hastens the discovery process and could lead to new "designer nuclei" with advantageous properties for humans. Examples of such nuclei existing today include terbium-149, which has shown an ability to attach to antibodies and irradiate cancer cells without affecting healthy cells, and radium-225, which could help scientists understand why there is more matter than antimatter in the universe.

Erler, J., et al. (2012), The Limits of the Nuclear Landscape, Nature. Volume: 486, no. 7404.

"The discovery of high-contrast tunneling magnetoresistance by Bill Butler's team was a clear example of the predictive power of computational science. From there, two courageous experimental groups, one in the US and one in Japan, were able to confirm these computational results."
—Jack Wells,
Oak Ridge National Laboratory

Tunneling Electrons

At the turn of the 21st century, materials scientists at ORNL were working on one of the biggest technical challenges of the new century: information storage. Since the OLCF was founded in 1992, the amount of digital information stored around the world each year had increased a thousandfold, from billions to trillions of gigabytes.

The ability to increase storage capacity is not only important for preserving information but also for developing more powerful and efficient computers. That's why a research team led by ORNL's William Butler used OLCF resources, including an IBM Power system, to model the material properties of layered magnetic films—to expand the limits of computing and storage capacity.

The Science

When a computer "reads" data in binary code, expressed as a series of ones and zeroes, the read-out head on the disk drive is sensing the magnetic fields from the magnetic films on a hard disk and converting them into electrical signals.

In 1988, physicists Albert Fert and Peter Grünberg independently discovered that electrical resistance may decrease when an external magnetic field aligns the magnetization of two layers of film, enabling the more efficient writing and erasing of data—a discovery that later won Fert and Grünberg the 2007 Nobel Prize in Physics. This effect, known as giant magnetoresistance (GMR), became a popular subject of materials research for computing and devices.

A decade later at ORNL, Butler's team was modeling the GMR effect in thin film systems composed of transition metals, which can have useful electromagnetic properties for writing and storing digital information. ORNL researchers with expertise in x-ray diffraction, microscopy, high-temperature materials, and HPC were trying to identify the best

transition metals for creating a high magnetoresistance effect to increase storage density. The team was also interested in materials that could withstand power interruptions without losing data, thus increasing reliability.

The thin films like those used on disk drives include two types of magnetic materials: ferromagnetic (in which the magnetism of the atoms aligns in the same direction) and antiferromagnetic (in which magnetism aligns in different directions).

Butler's team was using computationally intensive, first-principles calculations on ORNL supercomputers to model a phenomenon related to GMR known as tunneling magnetoresistance (TMR)—in which a magnetic tunnel junction is created by "sandwiching" an antiferromagnetic (insulating) barrier between two ferromagnetic (conductive) layers.

In classical physics theory, an electron cannot cross the insulating barrier. However, the team was demonstrating through the laws of quantum mechanics that, with the right materials, a magnetic tunnel junction could be created, allowing electrons to "tunnel" from layer to layer when their magnetism was aligned in the same direction. At a high TMR ratio, or contrast, a material exhibits a high probability of tunneling when magnetism is aligned across the layers, and a low probability for tunneling when magnetism is anti-aligned across the layers. A high TMR contrast means the size of information could shrink with less distortion or loss. At first, in laboratory observations of some common transition metals, the ORNL team and others were recording about a 50 to 70 percent TMR ratio.

The team used a first-principles electronic structure code to calculate the magnetoresistance of an Fe|MgO|Fe "sandwich," which used three layers of two ferromagnetic layers (iron) on either side of an insulating, antiferromagnetic layer of magnesium oxide (MgO). TMR had been experimentally demonstrated for iron and MgO but had not been calculated. What the ORNL team found was a staggering contrast: about 500 percent.

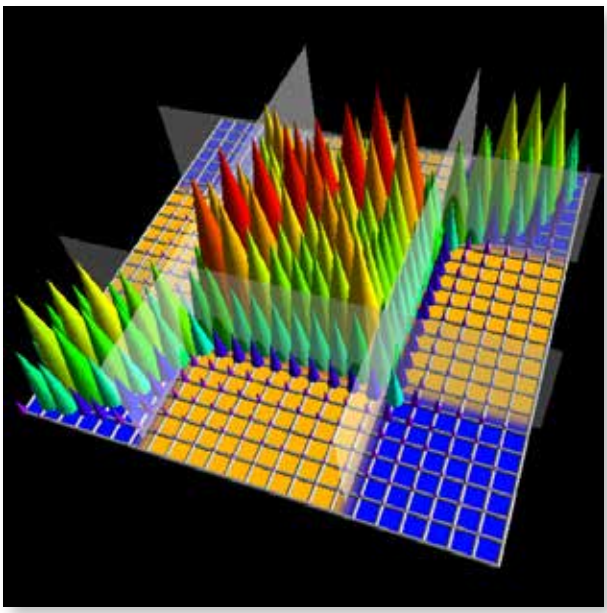

Giant magnetoresistance (GMR), a large change in a magnetic material's electrical resistivity caused by an applied magnetic field, was a popular subject of materials research for computing and devices following the phenomenon's discovery in 1988. A computer illustration of GMR, shown here, relied on massively parallel computing and the expertise of Williams Butler's team of computational and materials scientists at ORNL in the 1990s.

The Legacy

Such a high contrast enables the magnetic read-out head of a disk drive to capture more sensitive signals sent by smaller magnetic regions, thereby increasing storage capacity without losing quality. First discovered in computer modeling at ORNL, the tunnel junction between the Fe and MgO layers was confirmed by experimentalists a few years later; and even greater contrasts were observed between crystalline MgO insulating layers and ferromagnetic materials other than iron. Today, most devices in magnetic read-out heads rely on MgO-based tunnel junctions, as does a type of resilient memory known as magnetoresistive random-access memory (MRAM) that has long been of interest to the computing community for its potential use as fast, reliable, and energy-efficient permanent memory.

Butler, W. H., Zhang, X.-G., and Schulthess, T. C. (2001), Spin-dependent Tunneling Conductance of Fe|MgO|Fe Sandwiches, Physical Review B, Volume: 63.

"Now we are able to simulate these transient events for the first time. It represents so far the most serious validation test for our model's capacity 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."
—Zhengyu Liu,
University of Wisconsin

Resurrecting Earth's Paleoclimate

Beginning in 2008, a multi-institutional team led by University of Wisconsin–Madison professor Zhengyu Liu used two flagship OLCF supercomputers to perform the first simulation of an abrupt climate change event that occurred thousands of years ago.

Starting with the Cray X1E Phoenix and continuing its work on the Cray XT Jaguar supercomputers, the research team ran a continuous simulation of 21,000 years of the Earth's climate history, and simulated 200 years into the future to help forecast climate.

Liu indicated that the team's simulations were the most serious validation test for its model. He also noted that the validation of the model served as a critical step in assessing how well the model could project abrupt climate changes in the future.

In 2008, the team used Phoenix to simulate the first third of its simulation—from 21,000 years ago to 14,000 years ago—then transitioned to Jaguar to finish the simulation.

The Science

Climate scientists know that most shifts in the Earth's climate have happened over many millennia, but the Liu team focused on an era roughly 19,000 years ago when the climate shifted much faster. This period was called the Bolling-Allerod warming. In the span of a few millennia, the sea level rose 16 feet and Greenland's temperature rose by 27 degrees Fahrenheit.

These rapid shifts were largely driven by the interaction of oceans and glaciers. In a period of relative warming, glaciers began to melt and release large amounts of water into oceans. The melted freshwater largely stopped ocean circulation, and circulation continued to fluctuate for the next 5,500 years. When this fluctuation ceased, Greenland and other parts of the artic experienced the massive warming period and the subsequent rapid sea level rise.

The research team used simulation to understand this rapid climate shift. Using the Community Climate System Model (CCSM), the team created a coupled simulation that took into account atmospheric conditions, sea ice and ocean interactions, and land.

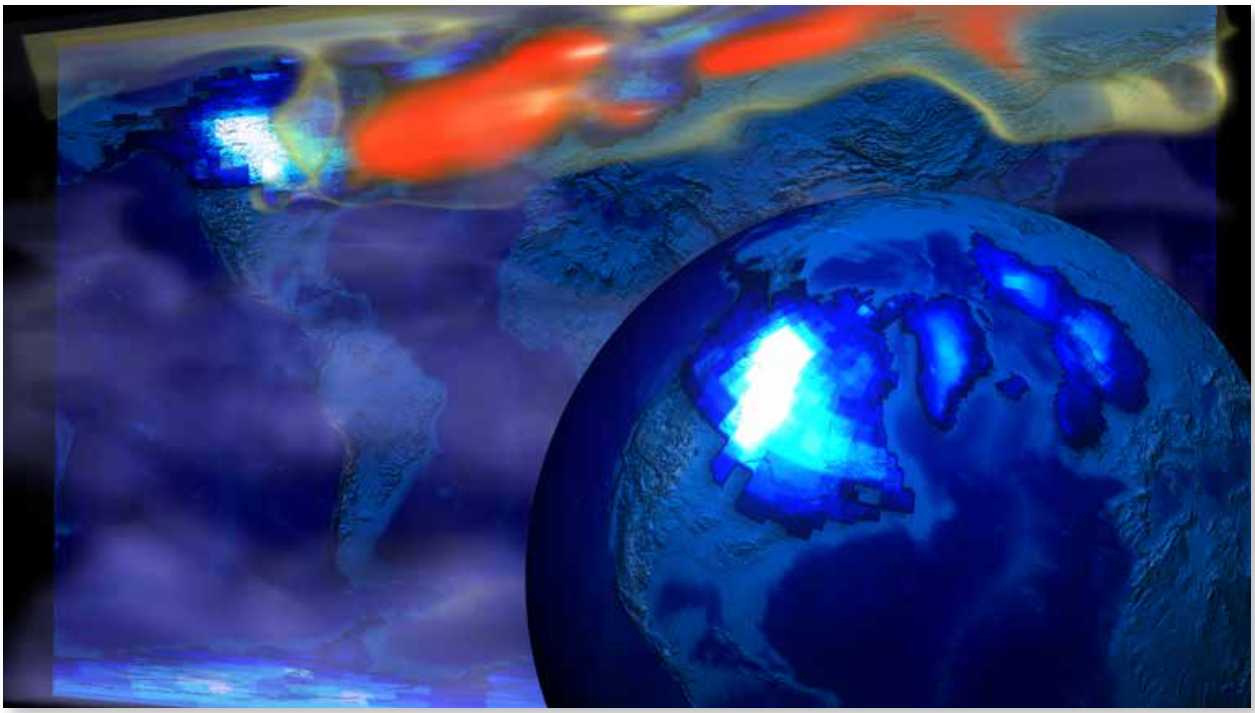

Simulations show deglaciation during the Bølling-Allerød, Earth's most recent period of natural global warming. Image credit: Jamison Daniel, ORNL

After running the simulation, the team discovered that three aspects contributed to drastic climate shifts during the Bolling-Allerod warming. First, the team indicated that one-third of warming during the period could be traced to a drastic increase in the carbon dioxide concentration in the air.

The second aspect arose from how oceans transport heat. With large amounts of glacial melt stopping ocean circulation, the water temperature in the North Atlantic cooled. Once ocean circulation resumed, warmer water currents returned to the North Atlantic and warmed the water there.

The third aspect also dealt with the ocean's circulation, but it occurred in a more surprising fashion. Once the warmer water currents returned to the North Atlantic, they were even stronger than before they were stopped by the glacial melts—the team referred to this increase of overturning circulation as an "overshoot." It accounted for one third of the warming during the Bolling-Allerod warming event.

The Legacy

The team's simulation served as an essential benchmark for researchers trying to expand climate models' scope and accuracy. The work done during this project, and its subsequent verification in 2012, led to the model's being included in the Intergovernmental Panel on Climate Change CCSM version 3—considered the premier climate modeling tool at the time.

In 2013, the team used calculations performed during this project to further study what processes contributed to the sudden increase in global temperatures. Once again using OLCF computing resources, the team found a combination of factors, including changes in the Earth's orbit, ocean circulation, and insolation, or the amount of solar radiation reaching Earth.

Liu, Z.; Otto-Bliesner, B. L.; He, F.; et al. (2009), Transient Simulation of Last Deglaciation with a New Mechanism for Bolling-Allerod Warming, Science, Volume: 17, no. 325. DOI: 10.1126/science.1171041.

"We expected a hierarchy of structure in cold dark matter. What we did not know is what sort of structure would survive the assembly because as these subclumps come together they are subject to tidal forces and can be stripped and destroyed. What we find is the survival fraction is quite high."
—Piero Madau,
University of California,
Santa Cruz

Illuminating Dark Matter

In 2007, researchers shone a galactic-sized light on the dark matter that holds our galaxy together by performing a large-scale simulation.

Using the Cray XT4 Jaguar supercomputer, a team led by astrophysicist Piero Madau of the University of California, Santa Cruz carried out the largest simulation to date of the Milky Way's dark matter —the invisible material that provides most of the universe's mass—and its evolution over 13 billion years. The work revealed new details regarding the likely distribution of the invisible substance throughout the galaxy, especially in the Milky Way's dense inner reaches, and gave astronomers a valuable tool in their search for dark matter.

As ongoing experiments and high-powered telescopes, such as NASA's Fermi Gamma-ray Space Telescope, continue to search the universe to increase understanding of dark matter, large-scale simulation could provide the critical information that tells scientists what to look for.

The Science

Since the 1930s, scientists have realized that visible matter supplies far too little gravitational force to explain the observed orbits of stars and galaxies. To account for the behavior of stars traveling around galaxy centers and galaxies traveling around each other, they concluded that all galaxies must be dominated by matter invisible to the eye. In fact, researchers have concluded that what we see makes up less than a fifth of the matter in the universe.

Fortunately, researchers do not need to see dark matter to simulate it. All they need to know is that dark matter interacts with ordinary matter through gravity and is cold, meaning the matter is made up of particles that move slowly when galaxies and clusters begin to form.

One lingering question facing scientists in 2007 was whether small clumps of dark matter remain identifiable or smooth out as they join other clumps to form halos large enough to host galaxies. Using initial conditions provided by observations of the cosmic microwave background, Madau and his collaborators

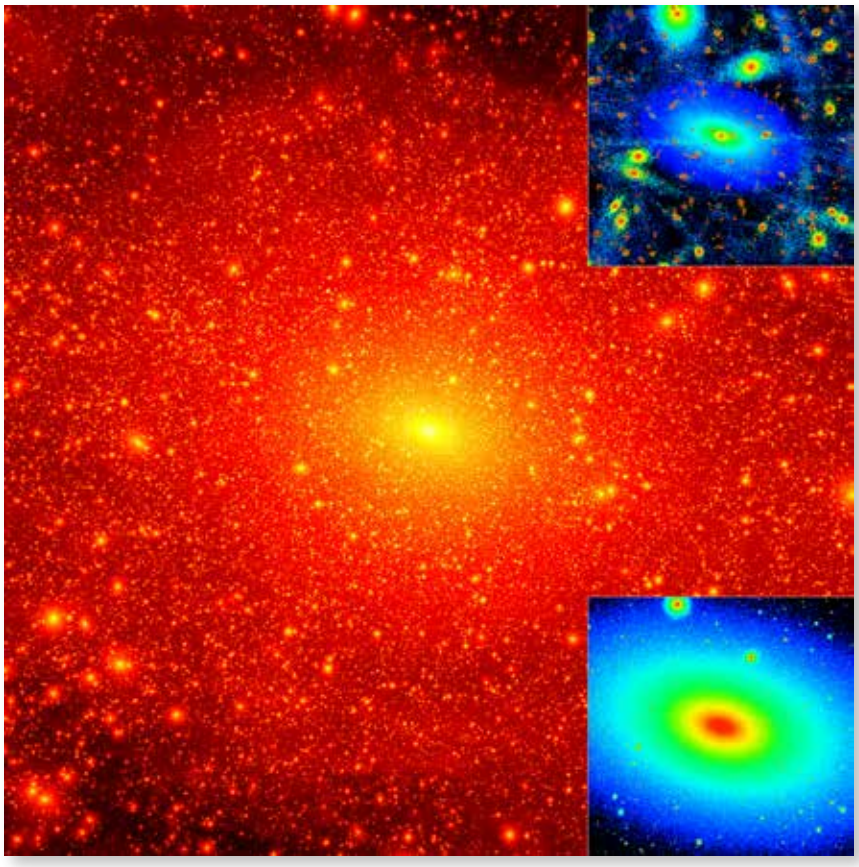

A cube of 800 kiloparsecs (kpc) per side is shown. The insets focus on an inner cube of 40 kpc per side (outlined in white), and show local density (bottom inset) and local phase-space density (top inset). Image Credit: Piero Madau, University of California, Santa Cruz

simulated a galaxy's worth of dark matter through nearly the entire history of the universe. The team used a computer application called PKDGRAV2 that ignored visible matter and focused entirely on the gravitational interaction among a billion dark matter particles. The simulations used about 1 million processor hours on the Jaguar system, which was capable of 120 trillion calculations a second at the time.

While earlier simulations did not possess the resolution to resolve any unevenness in the dark matter distribution—especially in the galaxy's dense inner reaches—the Madau team's billion-cell simulation verified that the earliest forms of dark matter do indeed retain a distinct form.

The Legacy

The resistance of dark matter clumps to indistinguishable homogeneity predicted by Madau's team reshaped the discussion about how our universe is put together. By capturing the substructure of dark matter in the neighborhood of our solar system, the team paved the way for future studies and observations.

Astrophysicists have developed and are testing new techniques, including gravitational lensing and the study of cold stellar streams in the Milky Way, to detect dark matter clumps predicted by Madau's team. If dark matter exists and is cold, it is expected that these elusive clumps will soon be unveiled by ongoing and planned galaxy surveys.

Diemand, J.; et al. (2008), Clumps and Streams in the Local Dark Matter Distribution, Nature, Volume: 454, no. 7205.

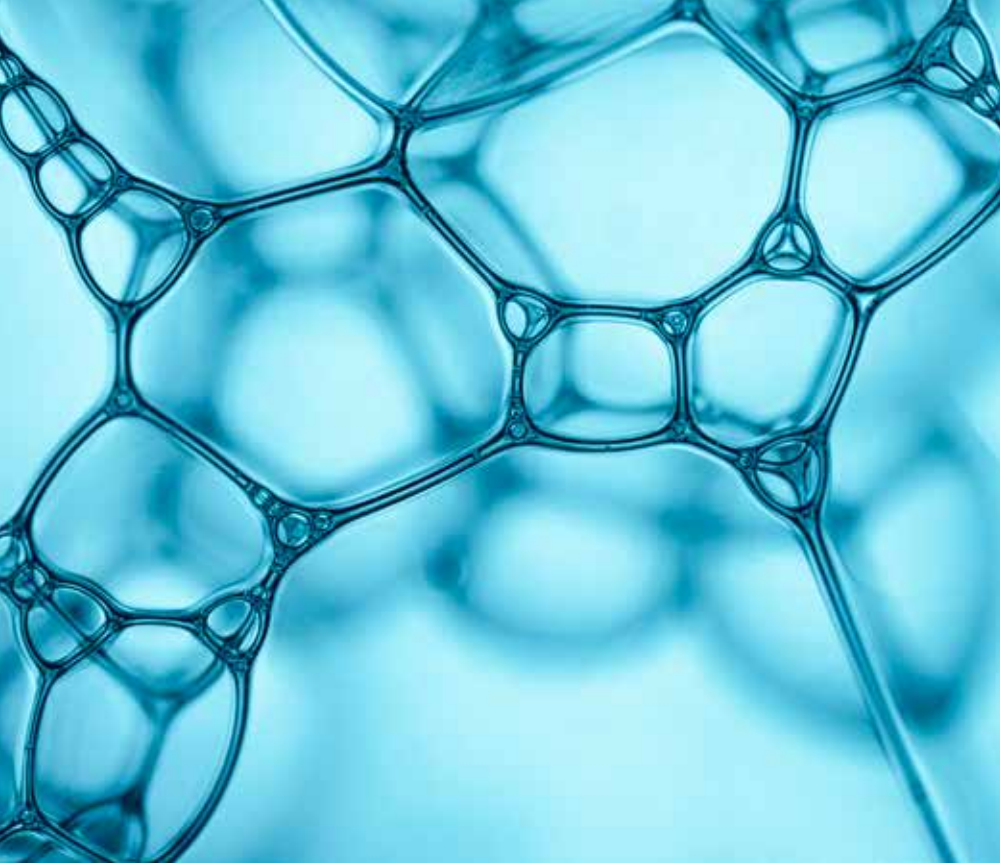

"These simulations allowed us to investigate the very early events associated with the initiation of vesicle fusion. If you want a complete mechanistic understanding of a process, these details and the capability to study them are extremely important."
—Russell DeVane, Procter & Gamble

Improving Everyday Products

Consumer-products giant Procter & Gamble (P&G), a Fortune 500 firm founded in 1837, came to ORNL in 2011 to better understand microscopic processes that can threaten product performance and stability.

Many of P&G's products for the body and the laundry are comprised of systems of fat-soluble molecules called lipids. These molecules dictate the flow, thickness, performance, and stability of P&G products.

In 2011, researchers at P&G with colleagues at Temple University turned to ORNL's Jaguar supercomputer, one of the fastest supercomputers in the world at the time, in an attempt to model and simulate microstructural arrangements of lipid vesicles.

The Science

Lipid vesicles form when lipid bilayers "roll" into spheres. Vesicles also can encircle each other like Russian nesting dolls. In the body, vesicles transport insulin, neurotransmitters, and other vital substances in their interiors and then fuse with cell membranes to deliver the contents. In manufactured products such as lotions and fabric softeners, man-made vesicles called liposomes encapsulate and suspend perfumes, dyes, and active ingredients.

Over time vesicles can fuse with each other, causing phases to separate and otherwise threatening consistent product performance. For P&G, it is crucial to understand vesicle fusion in order to extend the shelf lives of such products as fabric softeners, body washes, shampoos, and lotions.

To speed discoveries about why vesicle systems have stability problems, the researchers knew they would need to simulate the reorganization of vesicles over time. Before vesicles fuse, their surfaces strain and break, exposing water-avoiding parts of molecules to an aqueous environment. Following every molecule in a system of fusing vesicles as they contort and combine would be a herculean feat. P&G researchers hoped such simulations would help to predict performance attributes before a product is manufactured, so they would then be able to deliver high-quality products faster than competitors and earn greater

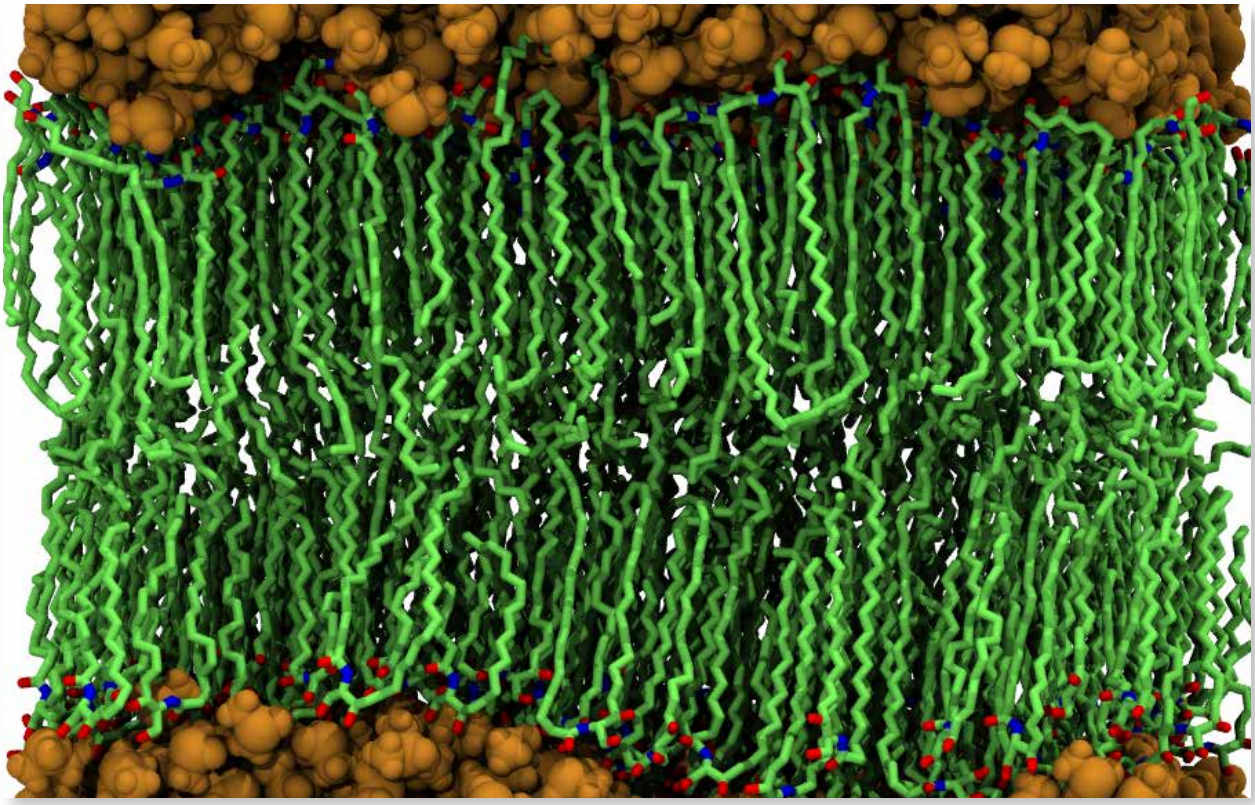

Understanding lipid–additive interactions gives P&G the opportunity to better design products and anticipate stability issues that may affect product performance and consumer satisfaction. Image credit: P&G

revenues in the marketplace. But the magnitude of the challenge exceeded the capabilities of even P&G's powerful supercomputers. They needed extremely long simulations of extremely big lipid systems made possible by OLCF resources.

With the help of 69 million core hours on Jaguar over two years, they were able to set the stage for simulations of large, complex systems of lipid assemblies. Until then, no one had ever been able to simulate vesicle fusion in realistic systems, such as biological vesicles whose fusion is triggered by proteins. With Jaguar's computational power, the researchers succeeded in artificially triggering the deformations of a single vesicle that happen before it fuses with another vesicle. The mathematical model applied an electrical field oscillating at 10 megahertz to one vesicle in a simulation that tracked the movements of the 5,000 lipid molecules in that vesicle.

These simulations allowed the team to investigate the very early events associated with the initiation of vesicle fusion. This benchmark simulation allowed comparisons to experimental data that verified the simulation and the interpretation of the experiment.

The Legacy

P&G has long understood the value of modeling and simulation, applying computational tools to significantly reduce the number of experiments involved in their process design for products like Pampers. The team sought out to achieve a similar goal when designing formulated products like shampoos and fabric enhancers. Since 2011, this team has continued using OLCF resources in developing codes and improving simulations capable of showing how microscopic properties of vesicles influence macroscopic properties, such as phase separation as a product ages. Gaining that type of knowledge has given untold long-term benefit as it helps the company wash its hands of potential poor performers and more quickly embrace formulations that are more likely to grow its fortune.

Fiorin, G.; M. L. Klein; R. DeVane; and W. Shinoda. "Computer Simulation of Self-assembling Macromolecules." In Hierarchical Macromolecular Structures: 60 Years after the Staudinger Nobel Prize I, 93 -107. Springer International Publishing, 2013.

"In order to do these kinds of models, you need a first-principles electronic structure method that will scale to a few thousand atoms, minimum. LSMS has that capability."
—Malcolm Stocks,
Oak Ridge National Laboratory

Revealing the Quantum World of Materials

Metals and alloys are ubiquitous, so we may barely notice them. However, new discoveries and optimizations of these materials are essential to modern life, contributing to energy, transportation, medical, manufacturing, and information technologies.

With the introduction of massively parallel computing at ORNL, a team led by Materials Theory Group leader Malcolm Stocks saw an opportunity to greatly increase our understanding of material properties. Over the last quarter-century, the team has leveraged supercomputing to improve predictive materials modeling, enabling them to model larger systems of many thousands of atoms, disordered structures that are difficult to predict yet important to performance, and the finite temperature properties of magnets.

The Science

When CCS was founded at ORNL in 1992, scientists were already using theory and computation to predict the properties of materials based on their atomic and electronic structure, but their calculations were limited by computational power. Researchers could compute only a small number of atoms at a time and were restricted to materials with neatly ordered structure, making it difficult to model the range of metals and alloys used for industry and technology applications.

One of the "Grand Challenges" recognized by the teams of researchers who wrote the 1991 proposal that led to increased funding for HPC and the creation of CCS was the need to simulate increasingly complex materials from first principles. Based on the fundamental laws of quantum mechanics that govern the behavior of electrons and atomic nuclei, first-principles calculations require an immense amount of computing power. However, physicists and materials scientists needed more than just a bigger, faster computer to advance materials structure modeling. The ORNL team quickly realized that the predictive methods that had been developed to compute the underlying electronic structure of metals and alloys did not complement parallelism. At every increase in system size—from tens to hundreds or hundreds to thousands of atoms—they knew the computational effort required would balloon. The parallelism actually worked against the computational methods developed

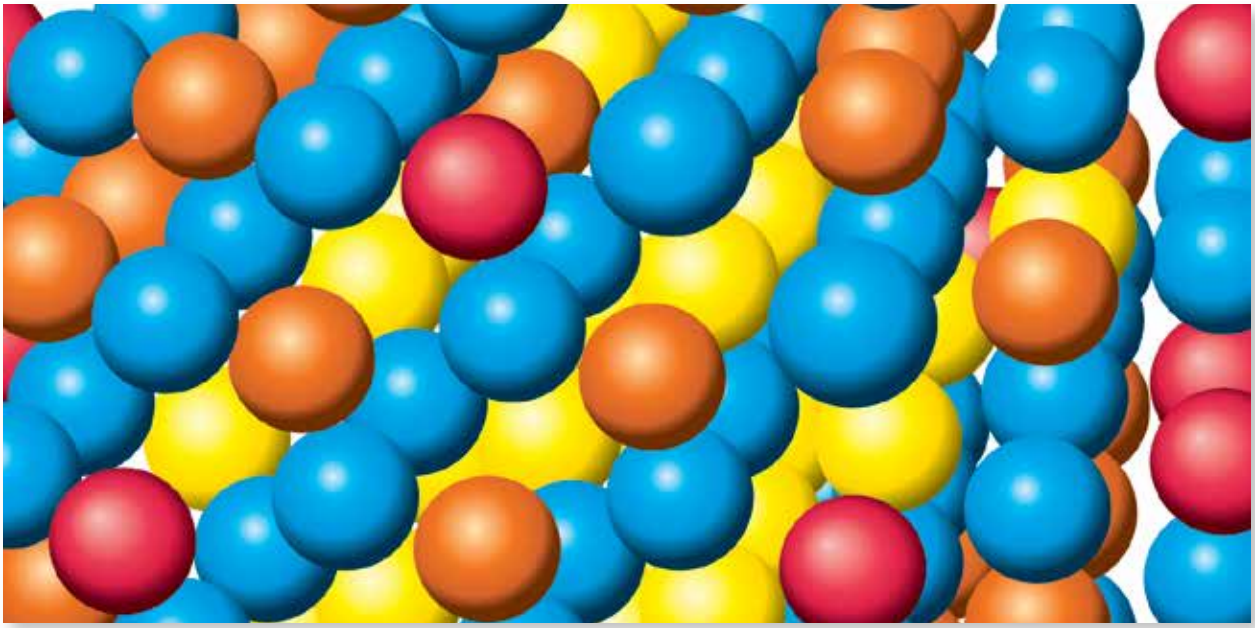

First described in a 1995 Physical Review Letters paper written by ORNL researchers, the LSMS code for modeling the electronic structure of materials has since been advanced for new generations of supercomputers. LSMS is now the basis of the WL-LSMS code, which has been used on Jaguar and Titan to study material disorder and magnetic properties.

to that point. So ORNL researchers took a different approach.

Using real space multiple scattering theory, which describes the propagation of electrons through solids, ORNL researchers developed the Locally Self-consistent Multiple Scattering (LSMS) electronic structure code specifically for large numbers of parallel processors. The LSMS code solves the runaway problem of calculating the electronic structure for large numbers of atoms by assigning each individual atom to a single computer node. The effect of neighboring atoms on each atom is then computed on each node, capturing the intent of parallelism. Computing local electronic structure in parallel reduced the computational burden while increasing system sizes and maintaining accuracy.

The Legacy

First presented in a Physical Review Letters article in 1995 based on computations carried out on the 1,024-node Intel MP Paragon XP/S-150 MPP supercomputer, LSMS was used to calculate the total energy of copper by modeling as many as 1,024 atoms—a system size previously inaccessible to first-principles methods. The ORNL team demonstrated not only that the code could produce accurate energies but also that the computing time scaled linearly with system size. Subsequent calculations that exploited the full power of the XP/S-150 produced insights into many materials systems, including disordered alloys, metallic magnets, and magnetic interfaces.

ORNL researchers continue to develop LSMS for ever more powerful supercomputer architectures. In 2000, the code was the first science application to perform at over one teraflop on a Cray T3E outside the laboratory. On the OLCF's Cray XT5 Jaguar in 2008, LSMS became the second application to run at over one petaflop and was later adapted for GPUs in preparation for the 27-petaflop, hybrid CPU–GPU Titan supercomputer that came online in 2012. Team members working on LSMS development have received two ACM Gordon Bell Prizes for peak performance in parallel computing: first in 1998 and again in 2009, when the application was combined with advanced statistical mechanics techniques to simulate the stability of alloys and magnets at finite temperature.

By leading the way into new territory time and again, LSMS has revealed the behavior of magnetic and electronic systems at unprecedented scales and provided an accurate framework for materials models beyond the HPC community and into mainstream modeling for research and industry.

Wang, Y.; et al. (1995), Order-N Multiple Scattering Approach to Electronic Structure Calculations, Physical Review Letters, Volume: 75.

Acronyms

CARB	California Air Resource Board
CFD	computational fluid dynamics
CCS	Center for Computational Sciences
CCSM	Community Climate System Model
CPU	central processing unit
CSGF	Computational Science Graduate Fellowship
DFT	density functional theory
DNS	direct numerical simulations
DOE	US Department of Energy
FUN3D	Fully Unstructured Navier–Stokes 3D code
GMR	giant magnetoresistance
GPU	graphics processing unit
HPC	high-performance computing
LSMS	Locally Self-consistent Multiple Scattering
MRAM	magnetoresistive random-access memory
NASA	National Aeronautics and Space Administration
OLCF	Oak Ridge Leadership Computing Facility
ORNL	Oak Ridge National Laboratory
P&G	Procter & Gamble
PICS	Partnerships in Computational Science
SASI	stationary accretion shock instability
SCEC	Southern California Earthquake Center
SciDAC	Scientific Discovery through Advanced Computing
TMR	tunneling magnetoresistance
UT–Austin	University of Texas–Austin