OAK RIDGE NATIONAL LABORATORY
REVIEW
• MANAGED BY UT-BATTELLE FOR THE DEPARTMENT OF ENERGY •

Vol. 45 • No. 3 • 2012
www.ornl.gov/ornlreview

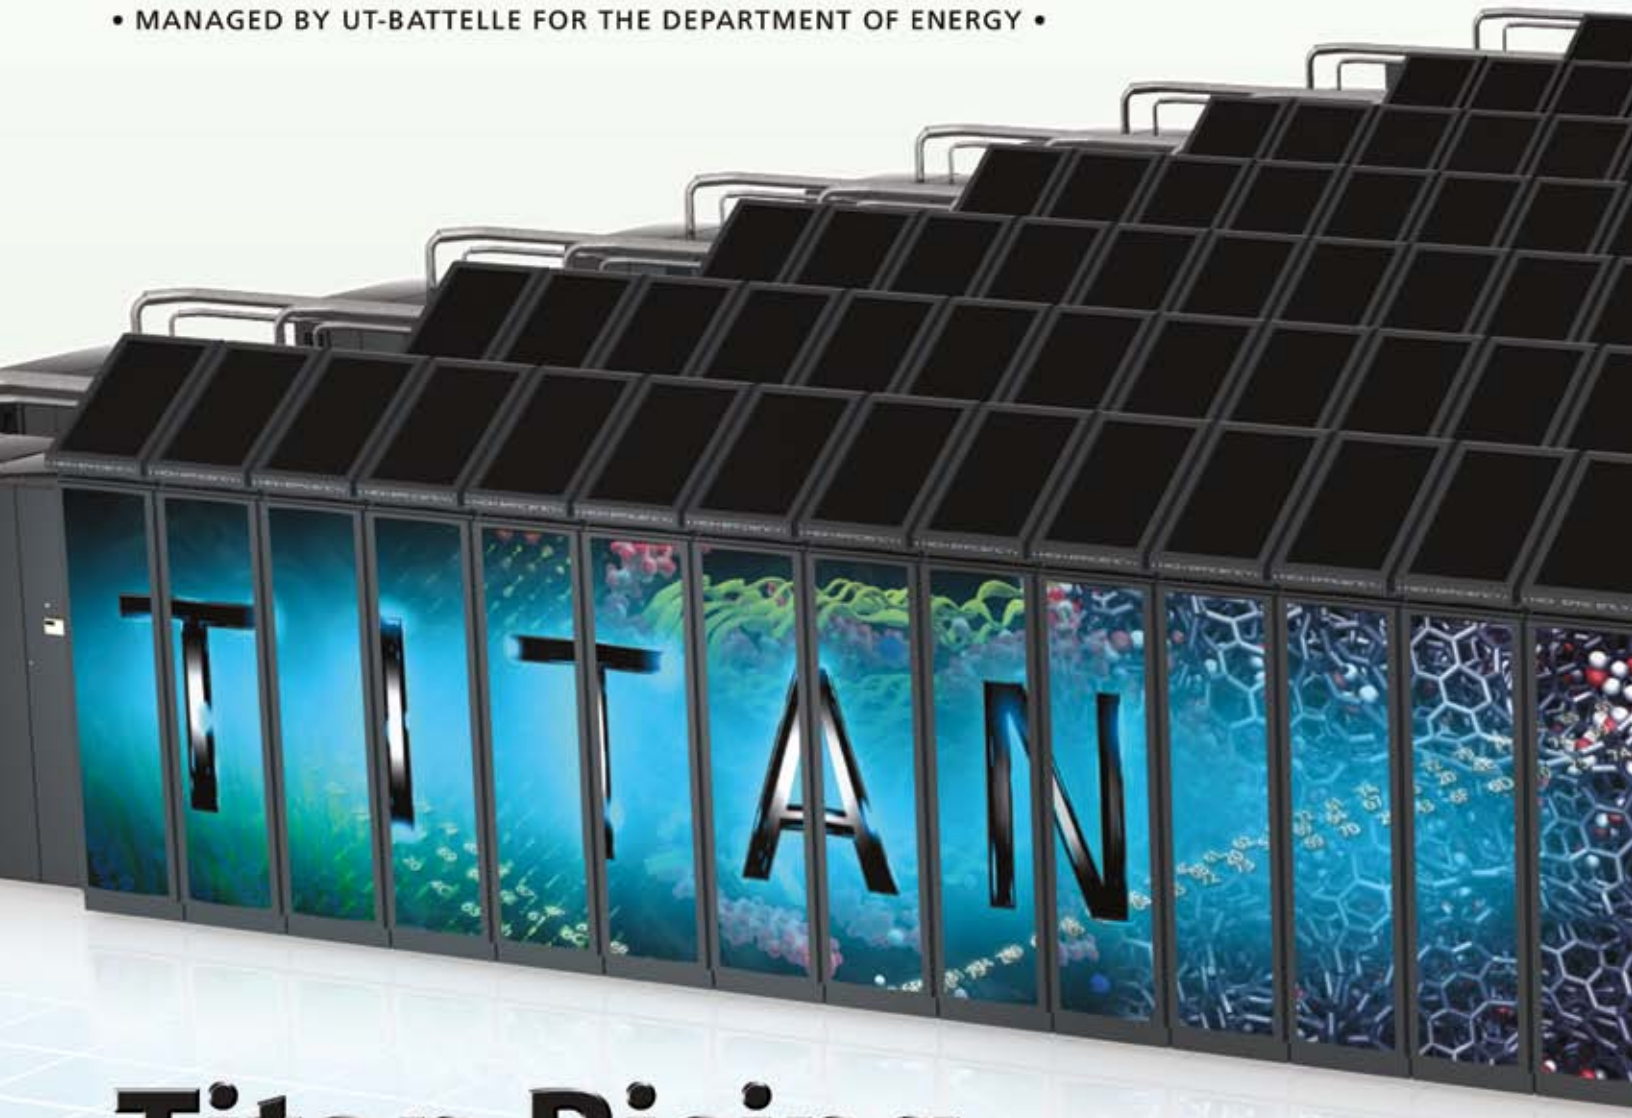

Titan Rising

The world's fastest computer for science

Expanding the realm of the possible

Ready for research on day one

OAK RIDGE NATIONAL LABORATORY
REVIEW
• MANAGED BY UT-BATTELLE FOR THE DEPARTMENT OF ENERGY •
Vol. 45, No. 3, 2012

contents

on the cover

ORNL's Titan supercomputer will extend the boundaries of computational science. Its peak performance of 20 petaflops, or 20,000 trillion calculations each second, will make it more powerful than any computer now in existence and six or more times more powerful than ORNL's current Jaguar system.

Image: Andy Sproles

TITAN: *Driving R&D through supercomputing*

When Titan is available to users in 2013, it will likely be the fourth time ORNL has hosted the world's fastest computer. The first was the Oak Ridge Automatic Computer and Logical Engine (ORACLE), created in the 1950s in partnership with Argonne National Laboratory to tackle problems in nuclear physics, radiation effects, and reactor shielding.

In 1995, ORNL acquired the Intel Paragon XP/S 150. At 150 GFlops (billion floating point calculations per second), it used 3,000 processors to deliver world-record speed in research areas including climate and groundwater modeling.

In 2000, ORNL became the first DOE laboratory to break the teraflop barrier (one trillion calculations per second), and sprinted through generations of increasingly powerful machines culminating in Jaguar, a petaflop (quadrillion calculations per second) computer in 2010. Jaguar, today with 300,000 processors, has generated world-class R&D across multiple disciplines including materials science, combustion, nuclear physics and astrophysics, biology, climate, and energy.

Of course, computing speed does not translate into great science without software—algorithms to perform the calculations and applications that convert scientific problems into computer simulations—and ORNL has an impressive history of software development, too.

In the 1960s, ORNL theorists Dean Oen and Mark Robinson used computer modeling to discover ion channeling. The discovery was critical to understanding ion implantation in a variety of applications including integrated circuit fabrication, and it provides a compelling example of scientific discovery through computer simulation.

In the 1980s, ORNL made numerous breakthroughs in parallel computing, research that anticipated the future of high performance computing. ORNL's Parallel Virtual Machine (PVM) software, with more than 400,000 users, became the worldwide standard for clustering computers into a virtual supercomputer. In the 1990s, ORNL teamed with other national labs and IBM to develop an ultrafast data storage system, now the standard for supercomputers across the nation. That same decade, Jack Dongarra of UT/ORNL led the development of the Message Passing Interface as well as linear algebra algorithms now used on virtually all supercomputers.

Since 2000, progress at ORNL has continued at a blistering pace. Titan will be more than 10,000 times more powerful in hardware, and more than a thousand times more powerful in software—a million-fold increase in capability in little more than a decade. This spectacular advance in performance is unprecedented, and is likely to be followed by another factor of a thousand or more over the next decade.

So what does this mean for science?

This issue of the *Review* provides a glimpse of that future, from the enabling hardware role of Titan's graphics processing units (GPUs), to crucial software advances at the ORNL Center for Accelerated Applications Readiness, to transformative applications including virtual reactors, extreme nuclei, climate modeling, biofuels and materials physics.

Titan will drive R&D in two ways. First, Titan will provide access to a set of problems that are impractical to attack with current supercomputing resources. Second, Titan will accommodate increased complexity in existing computer models, increasing accuracy and fidelity. As a result, we anticipate advances across broad reaches of science and technology including climate, seismology, combustion, materials, nuclear physics, biology, fluid dynamics, and fusion and fission energy.

These advances foresee an era where computer simulations stand alongside theory and experiment as a third leg of science and engineering, accelerating both discovery and design. ORNL is providing leadership for this revolution by fielding the most capable machines, developing the enabling software and applying supercomputing across a broad spectrum of science and technology challenges.

Jim Roberto

Associate Director for Science and Technology Partnerships

Oak Ridge National Laboratory

Solvable problems

Titan expands the realm of the possible

Supercomputers like Jaguar and Titan allow researchers to simulate complex processes such as the combustion of fuels in next-generation engines. Image: Sandia National Laboratories

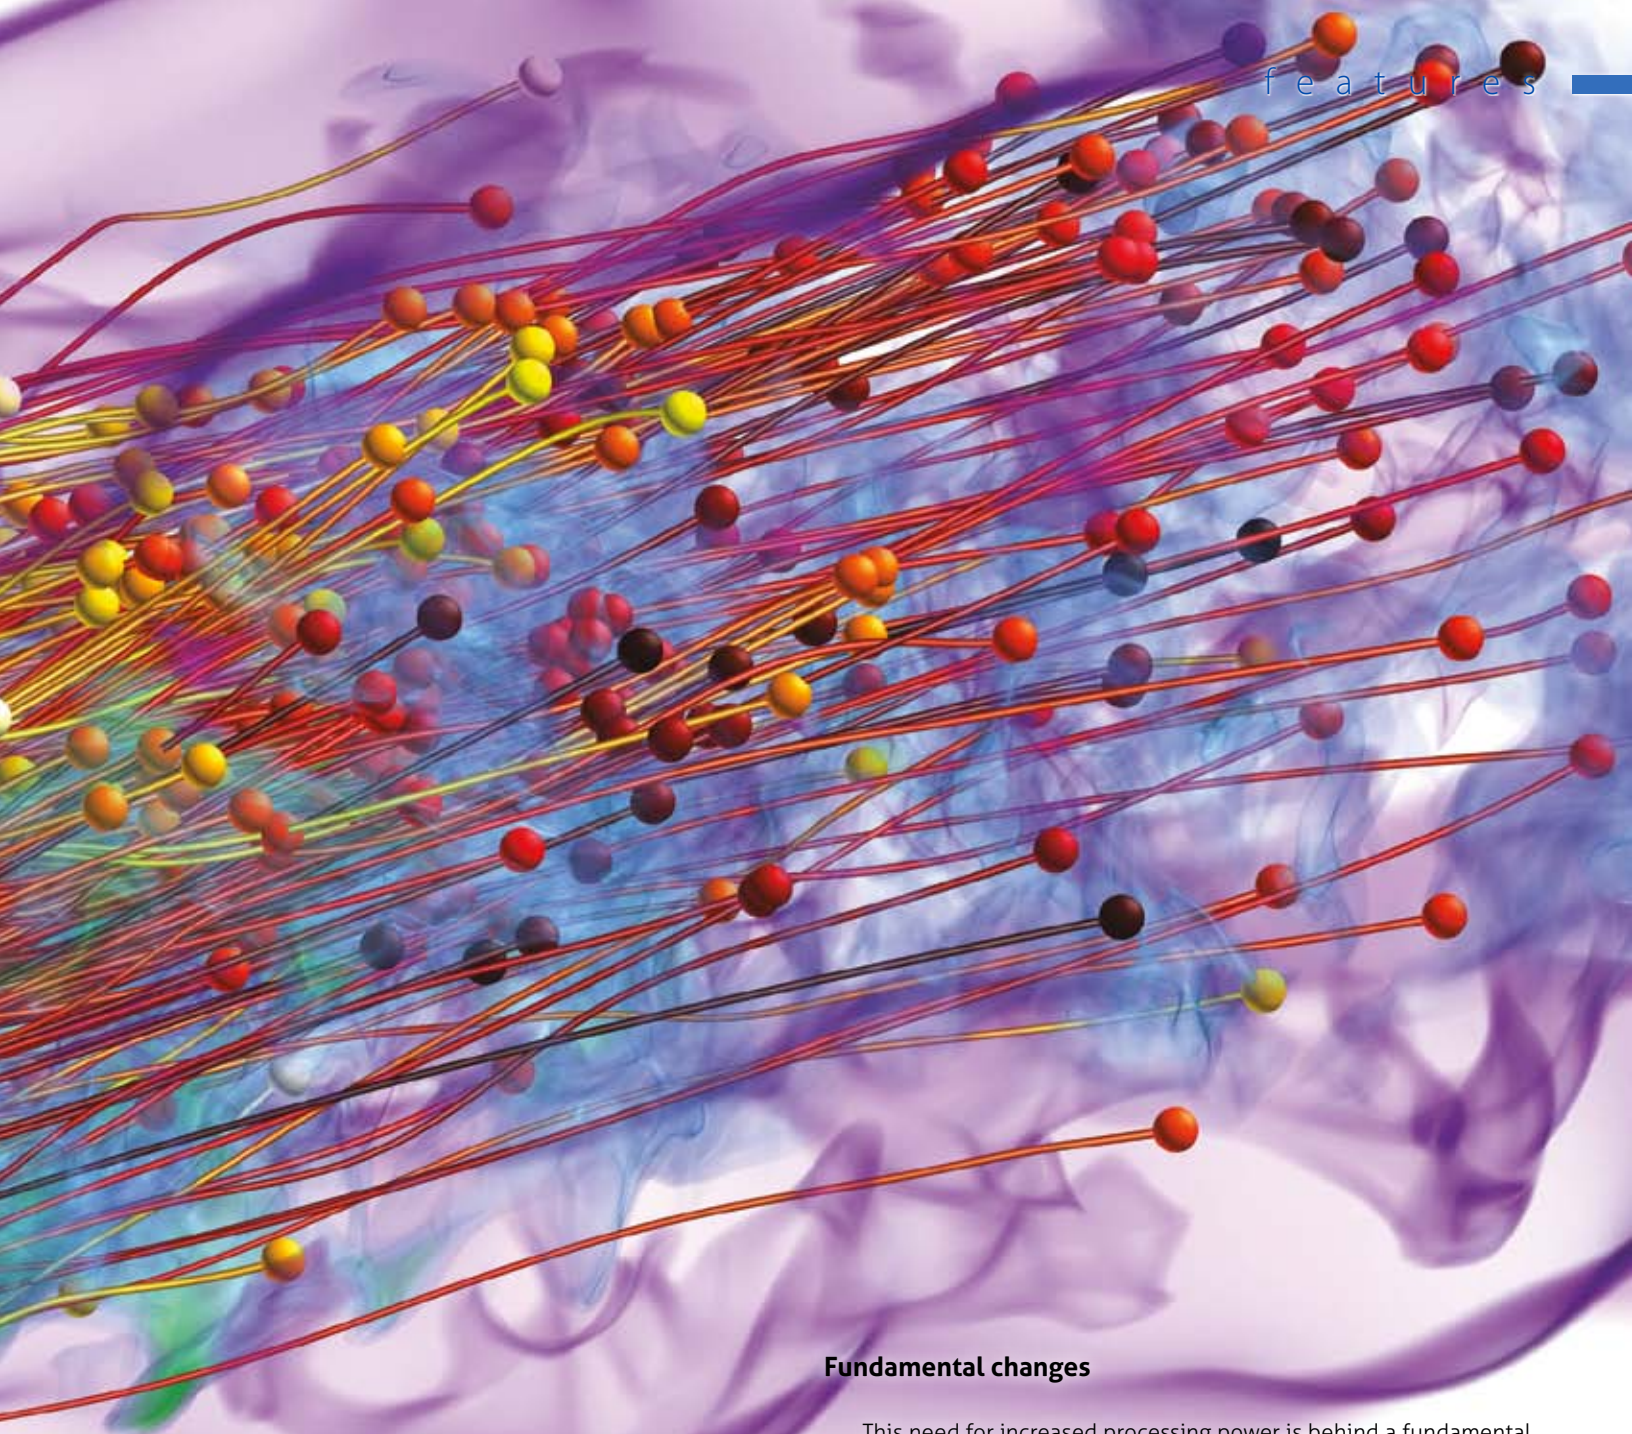

When the US Department of Energy asked researchers across a range of disciplines what they would do with a thousand-fold increase in computing power, ORNL computational scientist Jim Hack recalls, they all had the same need: the ability to create more detailed simulations.

"Increasing computational resources provides scientists with the opportunity to simulate whatever phenomenon or problem they're looking at with considerably more fidelity," says Hack, who directs the laboratory's National Center for Computational Sciences. "This increased accuracy can come in the form of greater resolution, or greater complexity. For example, when combustion researchers investigate how to burn different types of fuel more efficiently, they rely on computer models that involve a lot of complex chemistry and turbulent motion of gases. The more computational power they have, the more realistic the simulation will be."

Fundamental changes

This need for increased processing power is behind a fundamental change in supercomputer design that's taking place at ORNL and other leading centers of computational research.

"We have ridden the current architecture about as far as we can," Hack says. "Our Jaguar computer has 300,000 processor cores in 200 cabinets connected by a high-performance network. If we wanted to create a machine that is 10 times more powerful than Jaguar using the same architecture, we would have to put 2,000 cabinets on the floor, consume 10 times the power, et cetera."

The solution to this dilemma, at least for the NCCS, has been to design a machine that's different from Jaguar in two important ways.

The first step in Jaguar's transformation was to boost the performance of the computer's central processing units by increasing the number of cores per node from 12 to 16.

"Over the last few years, especially, CPUs have gotten faster because vendors have incorporated more processing cores," Hack notes. "The processing cores aren't getting any faster, so the whole strategy for improving performance in the same footprint has been to

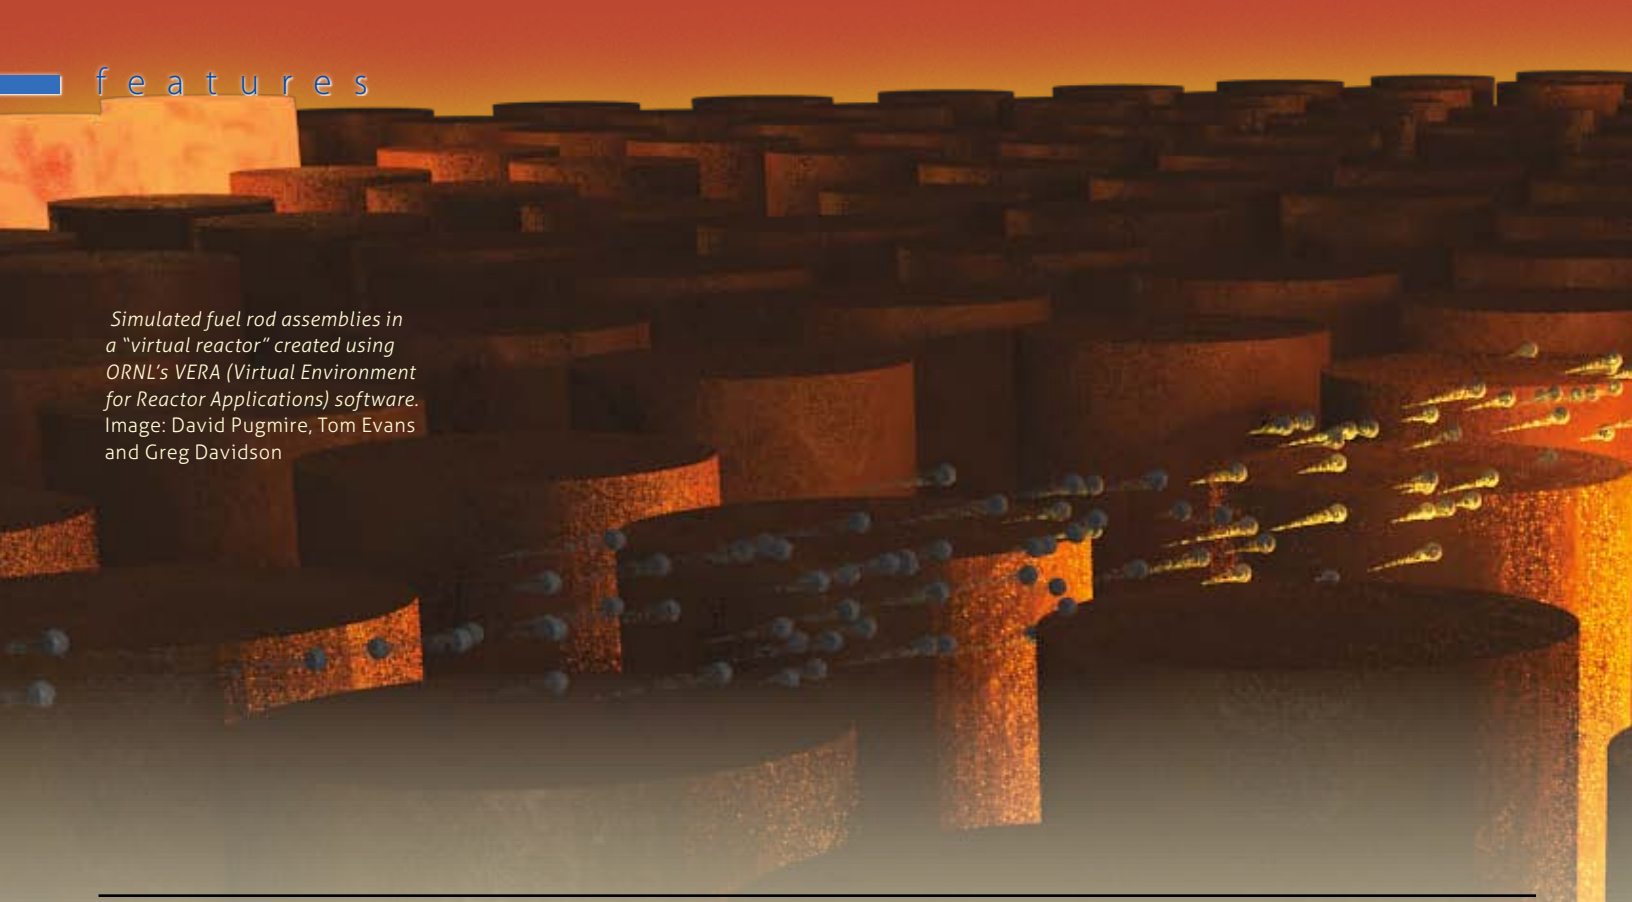

Simulated fuel rod assemblies in a "virtual reactor" created using ORNL's VERA (Virtual Environment for Reactor Applications) software. Image: David Pugmire, Tom Evans and Greg Davidson

add more processing cores. However, we're running out of room to go much further in that direction."

The remedy for this space constraint is the second part of the upgrade solution being put into practice in the new Titan supercomputer. In addition to amped-up CPUs, Titan will incorporate a feature found in many PCs—graphics processing units. On a laptop, GPUs are used to speed up video rendering, particularly for computer gaming apps. Unlike CPUs, which orchestrate the flow of information among various parts of a computer program, GPUs have a relatively simple task: run the computationally intensive parts of the program really, really fast.

New, portable code

Despite these upgrade modifications, when Titan is deployed, the basic footprint for the machine will be identical to Jaguar's. But instead of two 6-core processors, each of the 18,000 or so compute nodes will have a 16-core CPU paired with a GPU. The GPU has the potential to provide 10 times the performance of the CPU.

"The GPU is not as sophisticated as the CPU," Hack explains. "It waits for the CPU to say, 'I want you to do this to that data stream.' Then the data flows through the processor and an answer comes out the other side. This happens in thousands of parallel streams, so we can make better use of data locality—rather than moving data from node to node to node. With the right programming, GPUs work through the data in a very structured way."

In addition to processor-related improvements, Titan will employ an upgraded "interconnect," the internal network that enables compute nodes to talk to one another. The new network can handle a much higher volume of data and has a greater capacity to quickly re-route data when a data path is unexpectedly occupied or otherwise unavailable.

The linchpins holding these hardware improvements together are popular computer codes retooled by ORNL to take full advantage of Titan's rigorously parallel architecture. Much of this work is being done by the laboratory's Center for Accelerated Application Readiness. CAAR includes not only NCCS staff but also experts from various scientific disciplines and hardware and software vendors.

"When Titan comes online this fall, we expect it to offer researchers a full-fledged computing environment," Hack says. "Creating this environment has involved working with a wide range of software technologies and wasn't accomplished overnight. But it's easier now that we have better tools. Our collaboration with software vendors has helped us to evolve tools that make working on Titan easier for the application programmers."

Titan's ability to bring greater power to bear on any given challenge automatically expands the realm of solvable problems.

CAAR's initial goal was to facilitate research in astrophysics, biology, bioenergy, chemistry, combustion and energy storage by identifying widely used application codes in these areas. When the group analyzed the results, they found that the underlying algorithms in six major codes accounted for a significant fraction of total machine usage at the time. So those codes were the first to be migrated from the current architecture to the new Titan architecture.

"The CAAR group is showing the way," Hack says. "They're the trailblazers who are demonstrating not only how to optimize code for Titan, but also how to do it a way that's performance-portable—meaning that when you're done writing your code you can run it on any machine, regardless of its architecture. At the same time, they are establishing best practices for other programmers to follow when migrating their own codes."

A step toward exascale

In addition to being a world-class scientific computing resource, Titan is a testbed for concepts that its designers hope will move computer technology toward the goal of "exascale" computing—about 50 to 100 times faster than Titan's expected top speed.

"Our main concerns in this area included limiting both Titan's footprint and its power consumption," Hack says. "It's important to note that a cabinet costs about the same to populate and operate whether it's on Jaguar or Titan. With the hybrid architecture, for about the same price and energy consumption, we will get about 10 times the performance out of Titan. This is a significant step toward making exascale computing practical and affordable."

Of course the boost in performance was achieved largely through the incorporation of GPU technology into Titan's compute nodes. The power of the GPU comes from its simplicity and raw processing power. The price for this computational advantage is the additional responsibility placed on programmers or application scientists to figure out how to structure their data so they can provide the processor with input that seamlessly streams through it.

CAAR is also working to provide programmers with the tools they'll need to pass information efficiently to and from the GPUs. Historically, GPU software has been somewhat proprietary in the sense that each of the video accelerator manufacturers has had its own instruction set that tells the GPU what to do. Titan's GPUs are manufactured by NVIDIA and use an instruction language called CUDA (Compute Unified Device Architecture).

"The need to program the GPU using CUDA presented a problem to our programmers, who were already a couple million lines deep in code written in FORTRAN or C++," Hack explains. "It wouldn't be practical for them to rewrite that code in CUDA. So our strategy was to prepare the community to migrate to this kind of architecture by working with NVIDIA and other hardware and software vendors to establish a common set of directives that allow software written in common scientific computing programming language to tell the GPU what set of CUDA operations it needs to perform on the data.

"Code that has been re-factored in this way to run efficiently on Titan's GPUs has the added advantage of running faster on other processors that can take advantage of the fine-grained parallelism in the code as well," Hack says. "My message is that you can't lose with this kind of investment—the code allows any processor to take advantage of opportunities for concurrency and helps it to avoid moving data around unnecessarily."

Solvable problems

Hack observes that, if a scientist wants to tackle a problem with a computer simulation that would take a year's worth of computation to solve, then for practical purposes, it's an intractable problem. Solvable, problems are often a function of how much computational power the researcher has access to, so Titan's ability to bring greater power to bear on any given challenge automatically expands the realm of solvable problems.

"As technology gets better, we are able to tackle much more complex problems," Hack says. "More powerful machines provide scientists with the opportunity to get high-fidelity answers. The opportunity to use Titan will enable researchers in a range of disciplines to build simulations that address important scientific questions faster, with more detail and with greater accuracy."

— Jim Pearce

Moving to Titan

ORNL prepares for a leap in computing power

ORNL's Titan will extend the boundaries of computational science, increasing the complexity and confidence of climate prediction, boosting the efficiency of traditional engines and biofuel production, and pushing forward the full spectrum of computational science from astrophysics to nuclear structure theory.

When Titan is available to users in the first half of 2013, it will have a peak performance of 20 petaflops, or 20,000 trillion calculations each second, making it more powerful than any computer now in existence and six or more times more powerful than ORNL's current Jaguar system.

This leap in power will come from a new approach to computing: the use of accelerators. Titan will use thousands of graphics processing units, or GPUs, which have been developed since the 1980s to improve the video performance of computers and the look and feel of computer gaming. The addition of GPUs in supercomputing provides a major boost in computing power without a corresponding boost in the electricity required to run the system.

Gordon Moore and computer speed

The introduction of accelerators is at least the third strategy the computing world has used to get more out of a system.

For decades, advances came from shrinking transistors to fit more onto a sliver of silicon. Chip-makers did an admirable job of this; indeed, when Intel co-founder Gordon Moore famously predicted in 1965 that the number of transistors on a processor would double every two years, he was looking forward only a decade. Clearly he underestimated the industry he helped to found.

Nevertheless, advances have slowed in recent years.

"Since 2004, microprocessors have not gotten any faster on single-core performance," notes Buddy Bland, project director for the Oak Ridge Leadership Computing Facility (OLCF). "But Moore's Law, which people interpret as saying performance will double every 24 months, is still alive.

"So how do you reconcile these two statements? Moore's Law has gone into adding more transistors, and those additional transistors have been used to add more cores. You get more performance on the processor through more parallelism, not through increased single-thread performance."

In home computers, this approach can be seen in dual-core or quad-core chips. In the largest supercomputers, it can be seen in systems that incorporate hundreds of thousands of computing cores. Jaguar, for instance, incorporates 18,688, 16-core processors to reach a total of 299,008 separate processor cores.

Acceleration

Titan has been developed with the recognition that even this level of parallelism has limits. Traditional computer chips, known as central processing units, are fast but power-hungry. If supercomputers are to move beyond the petascale to the next major milestone—exascale computers capable of a million trillion calculations per second—they must find a more energy-frugal way to do it.

"The traditional microprocessors are optimized to make single threads go really fast," Bland explains, "but at a cost of using more power per calculation than necessary. There are several ways being developed to get more parallelism without increasing the power consumption of the processors dramatically, and one of those ways is using GPGPUs, or general purpose graphics processing units."

Researchers moving from single-core to many-core computers had to break their calculations into smaller problems that could be parceled out separately to the different processing cores. This approach was known as parallel computing. Accelerators are pushing parallelism by allowing researchers to divide those smaller problems even further.

"These processors have many, many, many threads of execution," Bland explains. "Each one runs more slowly than a traditional CPU thread of execution, but you have so many of them that it allows you in aggregate to get much higher performance at a similar power consumption."

Upgrade in two phases

Titan is transitioning from Jaguar in two phases.

The first was completed in February 2012. In it, Jaguar's AMD Opteron processors were upgraded to the company's newest 6200 series, and the number of processing cores was increased by a third, from 224,256 to 299,008. In the process, two six-core processors were removed from each of Jaguar's 18,688 nodes, or connection points. At the same time, the system's interconnect was updated and its memory was doubled to 600 terabytes.

In addition, 960 of Jaguar's 18,688 nodes received an NVIDIA GPU. This portion of the system, known as TitanDev, gives OLCF staff and selected researchers a platform for testing approaches to GPU computing (see sidebar: "Preparing users for Titan").

The upgrade increased Jaguar's peak performance from 2.3 to 3.3 petaflops.

The second phase of the upgrade will begin in October 2012. At that time the 960 accelerators will be removed, and most of Titan's nodes will get one of NVIDIA's next-generation Kepler GPU processors added to the existing AMD processor. Each of the Kepler chips will be capable of more than a trillion calculations each second, or 1 teraflop, and the new system as a whole will have a peak performance of at least 20 petaflops.

During this process the OLCF will also replace its file system. Bland says the target for the new system is to move a terabyte of data each second, making it 8,000 times as fast as a typical business Ethernet connection.

As the countdown to Titan winds down, facility staff are working overtime to ensure that the new system will be delivering valuable results from day one. TitanDev has proved important in this effort, giving the facility's staff scientists and its most sophisticated users a chance to try out approaches to the new architecture.

Bland notes that the first step—identifying more effective ways to divide large problems and spread them out among the available resources—would pay off even without the accelerators.

"We've done a lot of work to help those applications expose more parallelism. In doing that, we have typically doubled the performance of the application before you even use the accelerator, and when you add the accelerator we're typically seeing another doubling of performance," Bland says.

"When we get the next generation of GPUs in the fall, we expect it will be even faster."—*Leo Williams*

Titan's leap in power will result from the addition of thousands of graphics processing units. GPUs provide a major boost in computing power without a corresponding increase in the electricity required to run the system. Photo: Jason Richards

Preparing users for Titan

In a sense, the morphing of ORNL's flagship supercomputer from Jaguar to Titan resembles earlier installations. Being at the forefront of scientific supercomputing, the Oak Ridge Leadership Computing Facility regularly upgrades its capabilities to give researchers access to the most powerful resources possible.

But Titan is different. Although Jaguar has gone through six upgrades over the past seven years, the challenge was always one of scale. The system got newer processors, and there were more of them, but the type of processor remained the same.

Titan, on the other hand, will combine two very different types of processors—CPUs (traditional computer chips) and GPUs (chips originally developed to speed video performance)—and ask users to adapt their applications accordingly.

"GPUs are fast as lightning," explains ORNL's Bronson Messer, "but you can't ask them to make a decision. All they can do is take a big chunk of data and do the same thing on it."

Because CPUs and GPUs are good at different things, the new configuration asks researchers to rethink their problems. To help, the OLCF created the Center for Accelerated Application Readiness, or CAAR, a collaboration among application developers, Titan manufacturer Cray, GPU manufacturer NVIDIA, and the OLCF's scientific computing experts.

CAAR has been working for nearly two years to establish best practices for code writers. The center is divided into six teams working with six of the OLCF's most advanced and representative applications.

"Applications make use of supercomputers by finding effective ways to spread out the work over the available resources," Messer explains. "On Jaguar, for instance, the first level was to spread it across the nodes. Each node had two six-core AMD processors, so the work at hand would be spread across those.

"With Titan, each node will have not only a 16-core AMD processor but also an NVIDIA GPU capable of tackling hundreds of tasks simultaneously. Therefore, an application using Titan to the utmost must also find a way to keep the GPU busy, remembering all the while that the GPU is fast, but less flexible than the CPU."

The process must also allow researchers to move freely between supercomputing centers, notes OLCF Director Buddy Bland. The solution, he says, is in the creative use of programs known as compilers.

Application developers don't typically write code in languages that can be understood by Titan—or any other computer, for that matter. Rather, they write in languages such as FORTRAN, C, and C++ that must be translated into instructions that a computer can understand. The translator is a program known as a compiler.

"We've been working with the compiler community and the vendor community for a couple of years to help develop accelerator directives," Bland explains. "These are directives that you can put into your FORTRAN or C++ code that describe the parallelism in the code, but you let the compiler figure out how to express that parallelism in the machine language of whatever machine it's running on.

"In our case, it would figure out how to translate those instructions to run on the GPU. But they can also figure out how to run them most efficiently on the myriad of other processors that are out there."—*Leo Williams*

A Cray technician prepares to upgrade nodes on the Jaguar supercomputer.
Photo: Jason Richards

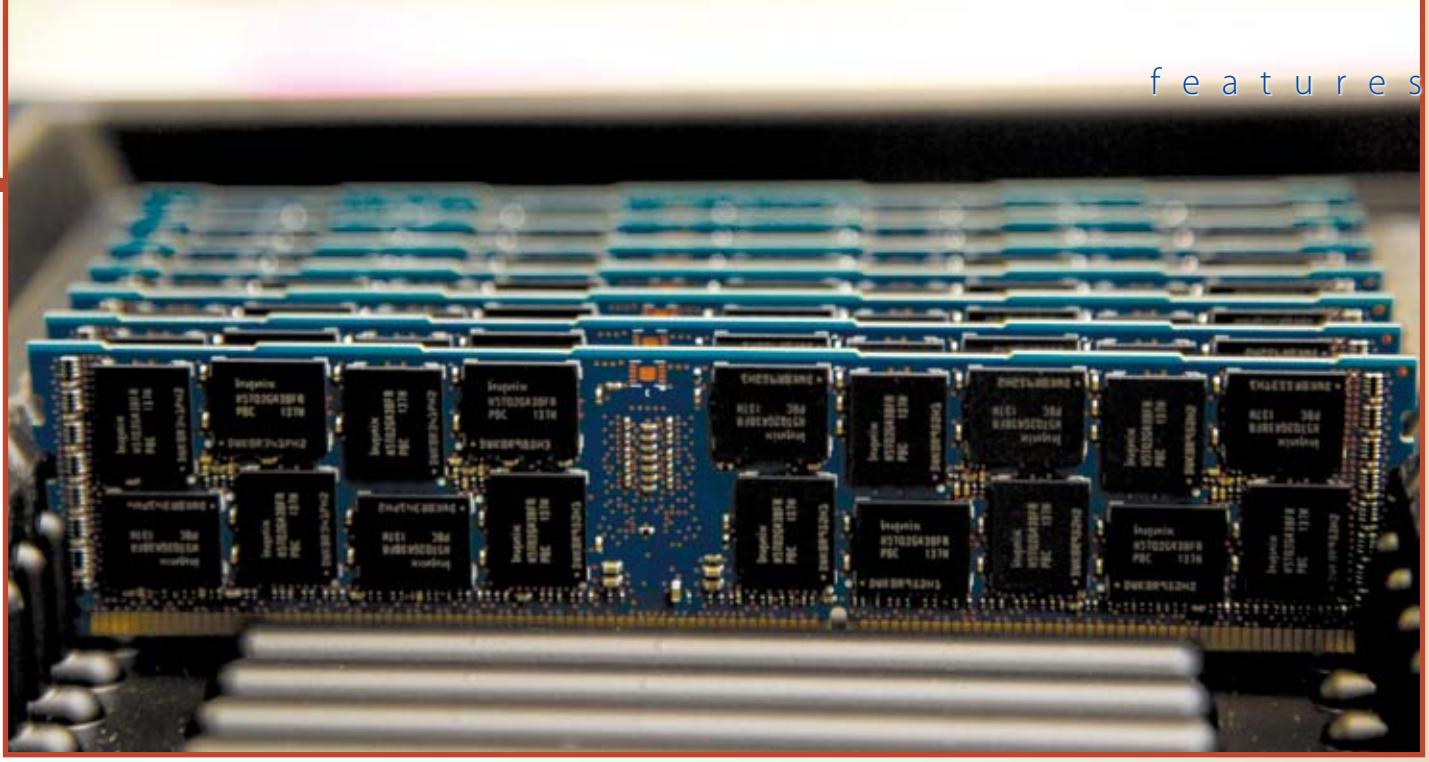

Science at 20 petaflops

The Center for Accelerated Application Readiness is working with six world-class applications that are representative of the Oak Ridge Leadership Computing Facility computing workload to show how GPUs can revolutionize computational science when they are combined with CPUs.

Combustion with S3D – Three-quarters of the fossil fuel used in the United States goes to cars and trucks, which produce a quarter of the country's greenhouse gases. Advances in fuel efficiency and alternative-fuel engines benefit national security and the environment. The S3D application simulates fuel burning under highly turbulent conditions. At 20 petaflops, it will move beyond simple fuels to tackle complex, larger-molecule hydrocarbon fuels such as isooctane (a surrogate for gasoline), commercially important oxygenated alcohols such as ethanol and butanol, and biofuel surrogates (blends of methyl butanoate, methyl decanoate and n-heptane.)

Magnetic systems with WL-LSMS – Magnetism at the atomic scale plays an important role in many industrial materials, including steels and iron-nickel alloys; and lightweight yet strong permanent magnets are important components in highly efficient electric motors and generators. Small improvements in the performance of these materials will result in more competitive industries and greater energy efficiency. WL-LSMS combines two methods, known as locally self-consistent multiple scattering and Wang-Landau. It analyzes magnetic materials at the nanoscale, allowing researchers to directly and accurately calculate, for example, the temperature above which a material loses its magnetism. At 20 petaflops, WL-LSMS will improve calculations of a material's thermodynamics or calculate the underlying magnetic states with greatly reduced margins of error.

Biophysical science with LAMMPS – Biofuels are among the most promising approaches to alternative energy production, but the process of turning woody plants into fuel is laborious and expensive. LAMMPS—or Large-scale Atomic/Molecular Massively Parallel Simulator—explores bioenergy using molecular dynamics, modeling problems such as membrane fusion, large biomolecular simulations for proteins and lignocellulose for biofuels. At 20 petaflops, it will be able to overcome size limitations on systems with charged particles, expanding to millions of atoms.

Nuclear reactors with Denovo – Nuclear power provides abundant, reliable and emission-free electricity. To be effective, though, safety must be guaranteed, and the volume of radioactive waste must be reduced (e.g., by burning the fuel longer in the reactor). Denovo is a powerful tool for ensuring the safe and efficient operation of today's nuclear power plants. At 20 petaflops it will take only about 13 hours to simulate a fuel rod through one round of use in a reactor core. The same simulation took 60 hours on Jaguar.

Climate change with CAM-SE – Improved atmospheric modeling will help climate researchers better understand future air quality, as well as the effect of particles suspended in the air—a large source of uncertainty regarding the climate's response to natural and human effects. The Community Atmosphere Model–Spectral Element simulates long-term global climate to inform public policy and improve scientific understanding of climate changes. At 20 petaflops it will be able to increase the simulation speed to between one and five years per computing day. The increase in speed is needed to make ultra-high-resolution, full-chemistry simulations feasible over decades and centuries and would allow researchers to quantify uncertainties by running multiple simulations.

Radiation transport and advanced algorithms with NRDF – The Non-Equilibrium Radiation Diffusion application models the journey of noncharged particles. It is also being used to develop advanced computing techniques that allow applications to solve larger problems by focusing computing power only on critical areas of a simulated system. NRDF has applications in areas such as astrophysics, nuclear fusion, and atmospheric radiation; the algorithms being developed for it should prove valuable in many other areas, such as fluid dynamics, radiation transport, groundwater transport, nuclear reactors and energy storage.—*Leo Williams*

Programming Titan

Hybrid architecture points the way to the future of supercomputing

The lab's new Titan supercomputer promises research that is literally global in scale.

A group of researchers from Princeton University is proposing to use Titan—the beefier successor to the lab's Jaguar supercomputer—to create seismological simulations of the entire Earth.

Not just a single fault line or continent. The world.

"This group wants to use seismic data from around the globe to image the substructure of the planet—as if they were giving the Earth an ultrasound," explains Jack Wells, science director for ORNL's National Center for Computational Sciences (NCCS).

"If they tried to run that simulation on Jaguar, it would have required the whole machine for half a year. However, if they optimize their code and win a big allocation of time on Titan, they just might be able to do it."

Titan comes online later this year and may be able to handle this feat of computational strength because of a fundamental difference in the way it approaches complex calculations. Rather than increasing the speed of its individual processor cores, Titan's computational prowess rests on tens of thousands of next-generation processors, each of which contains hundreds of cores designed to rip through calculations.

"They're not getting faster, they're getting wider," Wells says.

These "wider" processors can handle hundreds of parallel threads of data and are changing the way programmers and software designers work in the high-performance computing arena.

Speed and efficiency

For decades, computers got faster by increasing their central processor unit's "clock rate"—how often a CPU, or "core," cycles through its assigned tasks. However, fast processors produce a lot of heat. Twice the speed results in a chip that's eight times as hot. So by about 2004, manufacturers decided it made more sense to simplify the circuitry on their chips and use the extra space for more, slower cores running in parallel, rather than finding new ways to cool progressively faster monolithic CPUs. This led to the dual-core and quad-core processors that dominate the personal computer market today. Even an early version of ORNL's Jaguar supercomputer was based on dual-core processors.

"Parallelism at the chip level increases with the number of cores on the chip," Wells explains. "We went to dual core, then quad core, then hex core. The traditional processors we have placed in Titan have 16 cores."

These multicore CPU chips are flexible, powerful and designed to orchestrate complex computational activities. What they're not is energy efficient. "If they were a building, the heat and lights would be on all the time," Wells says.

Jaguar, with its 300,000 cores, consumes about 8 megawatts of power annually. To reach the center's goal of creating an "exascale" computer (about 400 times faster than Jaguar) with the same technology would require 100 to 200 MW of power. "That's not a practical option," Wells says. "We need a way to increase performance with much higher energy efficiency."

It just so happens that there is a rapidly emerging technology that does exactly that: the graphics processing unit, or "accelerator." Rather than following Jaguar's lead and using two CPU chips on each of the nodes on its communication network, Titan will replace one of the CPUs with a GPU. While consuming slightly more power than two CPUs, this hybrid arrangement allows the central processor chip to hand off the most computationally intensive and time-consuming activities to the accelerator, which divides them among hundreds of streamlined GPU cores and returns the results to the CPU. By radically boosting throughput, the hybrid configuration provides a big increase in processing speed.

As its name implies, the GPU was originally developed for accelerating graphics processing. The demand for increasingly realistic visual content in video games and other applications compelled GPU designers to endow their chips with the ability to apply complex physical equations to the task of producing convincing renderings of real-world phenomena like explosions, smoke and fire. Eventually game programmers and GPU manufacturers realized that such "game physics" could also have applications in the field of scientific computing and simulation.

"At that point, NVIDIA a leading GPU manufacturer, began to develop a business strategy for producing GPUs for scientific

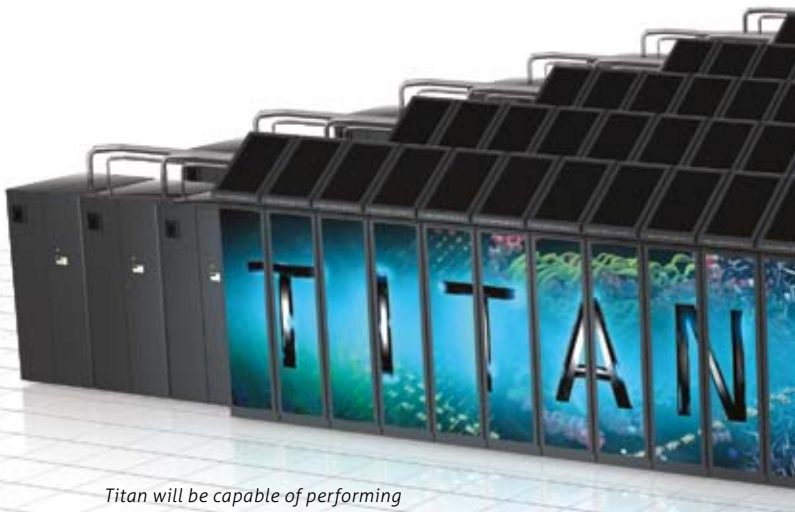

Titan will be capable of performing 20,000 trillion calculations every second, making it six times more powerful than ORNL's current Jaguar system. Image: Andy Sproles

computing," Wells says. "Our partnership with them on Titan is part of that strategy. NVIDIA recognized that, although scientific computing is a smaller market than video gaming, it has considerable impact, and it is growing. So now they have a line of products for scientific computing."

New challenges

The use of these highly parallel GPUs shifts some of high-performance computing's complexity from its hardware to its software, providing new challenges for software developers.

"To take full advantage of this hybrid architecture, a programmer needs to concentrate on revealing all of the available parallelism in the computer code to the processor," Wells explains. "Any tasks that can be done in parallel, need to be made available to the hierarchy of GPUs, CPUs, the computer's communication network, and the memory structure that goes with it."

Wells says that, in anticipation of the move to Titan, the laboratory's Center for Accelerated Application Readiness has been retooling key computer codes in a range of research areas to exploit opportunities for parallel processing. As a result, most of them are now also twice as fast when they're running on traditional CPU-based systems.

One of these applications is LSMS, a code used for modeling magnetic materials that was initially developed at ORNL 15 years ago. It employs a computational technique called matrix-matrix multiplication to simulate the interactions between electrons and atoms in magnetic materials. Because the code was developed for early parallel computers, its programmers anticipated many of the nuances of "message passing," the ability to efficiently use parallel streams of data that will be critical to operating in Titan's highly parallel environment.

"The message-passing capabilities and data structure of LSMS allow the program to stride through memory in a regular fashion," Wells says. "That means the processor doesn't have to wait for data to arrive. Time isn't wasted waiting on resources."

As a result, while the program will require some work to bring it up to date, it is a good candidate for being adapted for use on highly parallel computers.

Wells explains that the accelerators give researchers new tools to apply to their research problems. "The question they need to consider is how they will structure their data in order to use these resources efficiently.

"Think of a GPU as a factory with materials moving around the factory floor," he suggests. "When you have big, expensive machines available, you want the necessary materials to be there when you need them, so the machine operators are running the machines, not waiting around drinking coffee.

"An important part of our Titan project is partnering with software vendors to make sure that appropriate programming tools such as compilers and debuggers are available, along with the necessary libraries, to help make this programming task more manageable.

"Again, restructuring the data to take advantage of all available parallelism is the basic programming task for all foreseeable architectures. Accomplishing this is the first thing on which the programmer must focus. We have observed that the performance of the restructured codes, in general, is twice as fast as the old code running on the same, traditional CPU-based hardware. That is before we offload work to the GPUs."

Next-generation GPUs

The application of this hybrid architecture to building Titan began in earnest earlier this year and will occur in two phases. Phase 1

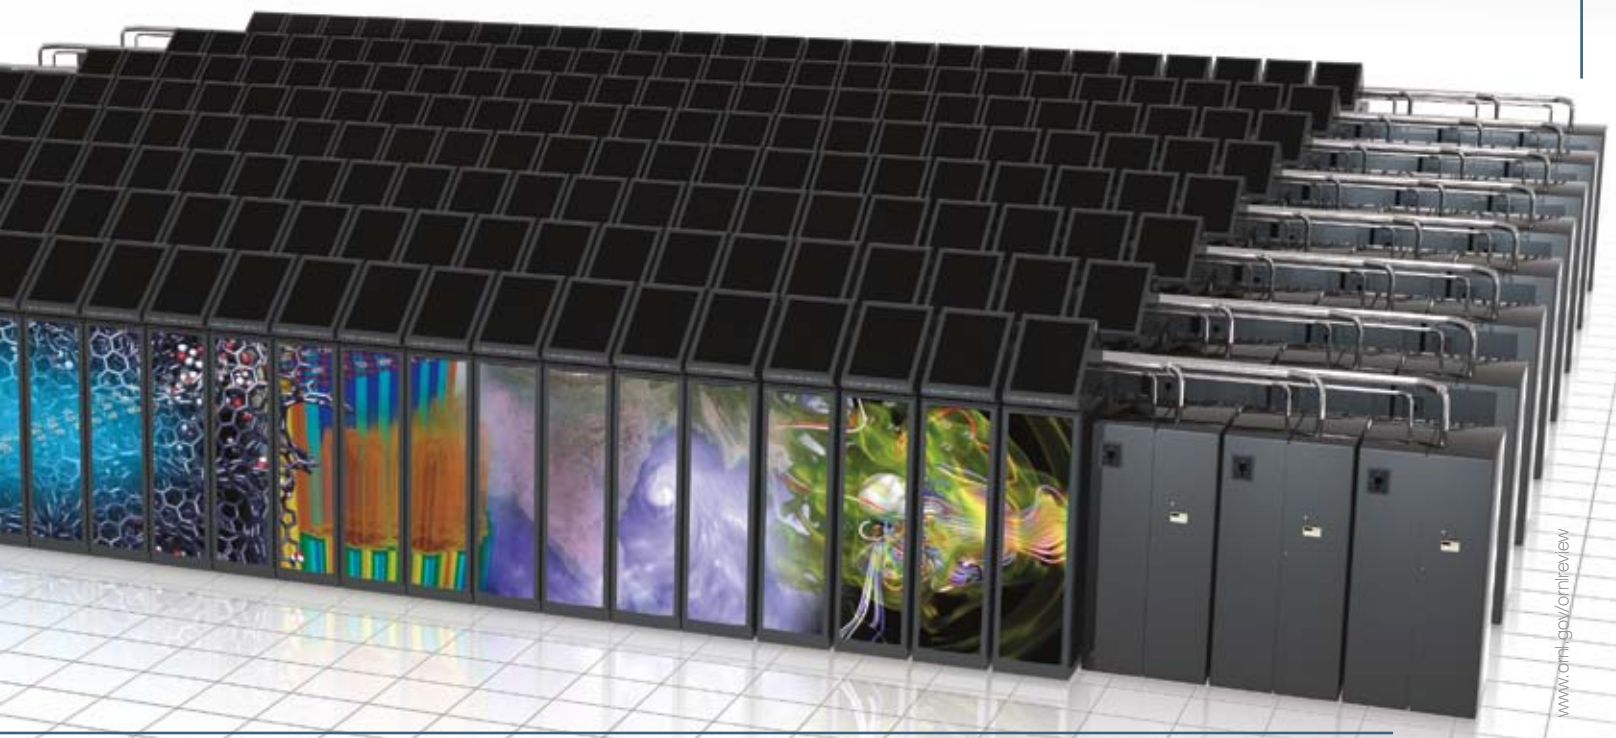

involved replacing all the circuit boards in Jaguar with new ones that came with new CPUs and installing twice as much memory, a slot for the yet-to-be-installed GPU, and a new network interconnect to speed communication among nodes. The upgraded interconnect is also "hot-swappable," so if something goes wrong with the node (if a processor fails, for example) it can be replaced while the machine is still running.

Phase 2, which is scheduled to begin in October, will fill the empty accelerator slots of the new nodes with NVIDIA's next-generation "Kepler" GPU.

While they wait for Kepler to make its appearance, NCCS researchers have partitioned off 10 of Jaguar's 200 cabinets to try out the new hybrid architecture using NVIDIA's current-generation GPU, called Fermi.

Wells notes that this developmental mini-Titan has already yielded a good comparison of traditional CPU-only architecture, like that used in Jaguar, with the hybrid arrangement that will be used in Titan.

"Depending on the code we're running, with the Fermi chip, the hybrid node is usually a factor of 1.5 to 3 times faster than a node containing two multicore CPUs," he says. "On a few applications it has been 4 to 6 times faster. We made these comparisons after rewriting the code to take better advantage of the parallelism in the Fermi accelerator. So it's an apples-to-apples comparison.

Of course the big question is, once the Jaguar-to-Titan transformation is complete, how fast will Titan be? There are algorithms for extrapolating from Fermi's performance to Kepler's, and those suggest that Titan could be 10 times faster than Jaguar. The crucial variable, however, is how many of Jaguar's 200 cabinets will be upgraded to include Kepler GPUs. That has yet to be determined. So a reliable prediction of Titan's power is still just out of reach.

Opportunity for innovation

When Titan goes online, researchers will be able to create some of the simulations they have been dreaming about, but were just too big or too detailed for other systems.

For climate scientists that might mean generating 10 times as much detail in their climate models. For example, climate simulations often assume that air pressure is relatively constant—the same at the ground as it is at the jet stream. They make this compromise with reality because calculating variable air pressure would add a lot of computation time, and researchers have a limited number of hours on the computer to get their calculations done.

"There's a lot of interesting stuff that happens at different altitudes," Wells says. "Titan will now be able to explore that."

Other users plan to apply Titan to studying the fundamentals of combustion—how materials burn—in greater detail than ever before. These studies will give researchers a better understanding of the basic processes that underpin the use of both fossil fuels and complex renewable fuels, such as ethanol or biodiesel, for transportation and electricity generation.

"We expect this research to result in the development of more energy-efficient combustion processes for these fuels," Wells says.

Another group will use highly accurate molecular dynamics techniques to simulate the fusion of biological membranes. This process is fundamental to cell division and is related to disease processes, such as cancer.

"This basic biophysical phenomenon is poorly understood," Wells says. "Titan's computing power will enable highly accurate simulations and new insights.

Expanding the base

"These are the kinds of problems we want scientists to consider addressing with Titan. We're asking the scientists, if you had a big allocation of time on Titan a year from now what would you do with it?"

The response to Wells' question has been enthusiastic—and maybe a little surprising. "As expected, most of the scientists and engineers we have traditionally worked with are interested in working with Titan's hybrid programming model," Wells says. "What we didn't anticipate was the number of users working in the GPU computing space—researchers doing accelerated computing on their workstations, rather than on supercomputers—who have been inspired to think of bigger problems that could be done on Titan. We were worried that Titan's hybrid architecture would alienate some of our traditional users; instead, it is actually attracting a new community of users to Titan.

"This is a very healthy development. A big part of our mission is reaching out to new users and new communities and encouraging them to take advantage of our unique resources for scientific computing."—*Jim Pearce*

Rather than following Jaguar's lead, Titan will replace one of the CPUs on each of its nodes with a GPU. By radically boosting throughput, this hybrid configuration provides a big increase in processing speed.
Photo: Jason Richards

Virtual reactor

VERA analyzes nuclear reactor designs in unprecedented detail

Some people would call ORNL computer scientist Tom Evans an optimist. He's building a mathematical model of a nuclear reactor so detailed that the only computer it can run on hasn't been built yet.

Most of the time, reactor models are used by utility companies to optimize and troubleshoot power plant operations. Usually, these simulations run on small, networked groups of computers, called "clusters," or on garden-variety desktop and laptop computers—which makes Evans' gargantuan programming effort all the more mysterious. That is, until you start thinking like a computer scientist.

"The fact of the matter is that today's supercomputer is tomorrow's cluster—and the next day's laptop," he explains. "If we don't push the envelope in the area of supercomputing now, we won't have anything to offer users working on desktops and clusters five years from now."

Evans works for the laboratory's Consortium for Advanced Simulation of Light Water Reactors, and over the next five years, he and his colleagues will be focused on producing a "virtual reactor" called VERA. This software tool will enable users in the utility industry to simulate and analyze almost every aspect of the performance of existing or proposed reactor designs in unprecedented detail.

Development of VERA, which began on the laboratory's Jaguar supercomputer, will eventually coalesce on Titan, Jaguar's brawnier successor expected to be available this fall.

Virtual blueprint

Currently, because most nuclear utilities have limited access to high-end computational resources, like Jaguar or Titan, the industry standard for reactor simulation consists primarily of low-resolution models that do a good job of capturing basic reactor behavior but don't provide a great deal of detail or flexibility. CASL's goal is to use VERA's higher-resolution view of reactor operations to get a better understanding of problems that these models don't anticipate, such as corrosion, interactions among fuel rod components, and various safety concerns, and then to share VERA's capabilities with utility companies.

VERA provides users with a comprehensive view of reactor core operations by integrating a range of specialized modeling tools designed to simulate what goes on inside a reactor and providing a framework that enables them to interact with one another.

"We take a number of factors into consideration," Evans says, "including thermal hydraulics, neutron transport, fuel depletion, chemistry and a variety of other processes that affect the behavior of the core. There is a lot of physics at play here. Each piece of VERA has a set of model equations that represent the physics involved."

The upshot of VERA's wide-ranging perspective is that its users have the ability to monitor the performance of any part of the reactor core at any point in its fuel cycle. An engineer can ask the model about the temperature of a particular fuel pin or the power being produced by a specific part of the core and get a data-based snapshot of what is happening in each location.

Evans and his colleagues constantly benchmark the data produced by VERA with data from real-world reactors under identical conditions. These comparisons provide them with a high level of confidence that measurements made on the model will match up with those made on operating reactors at any point in time.

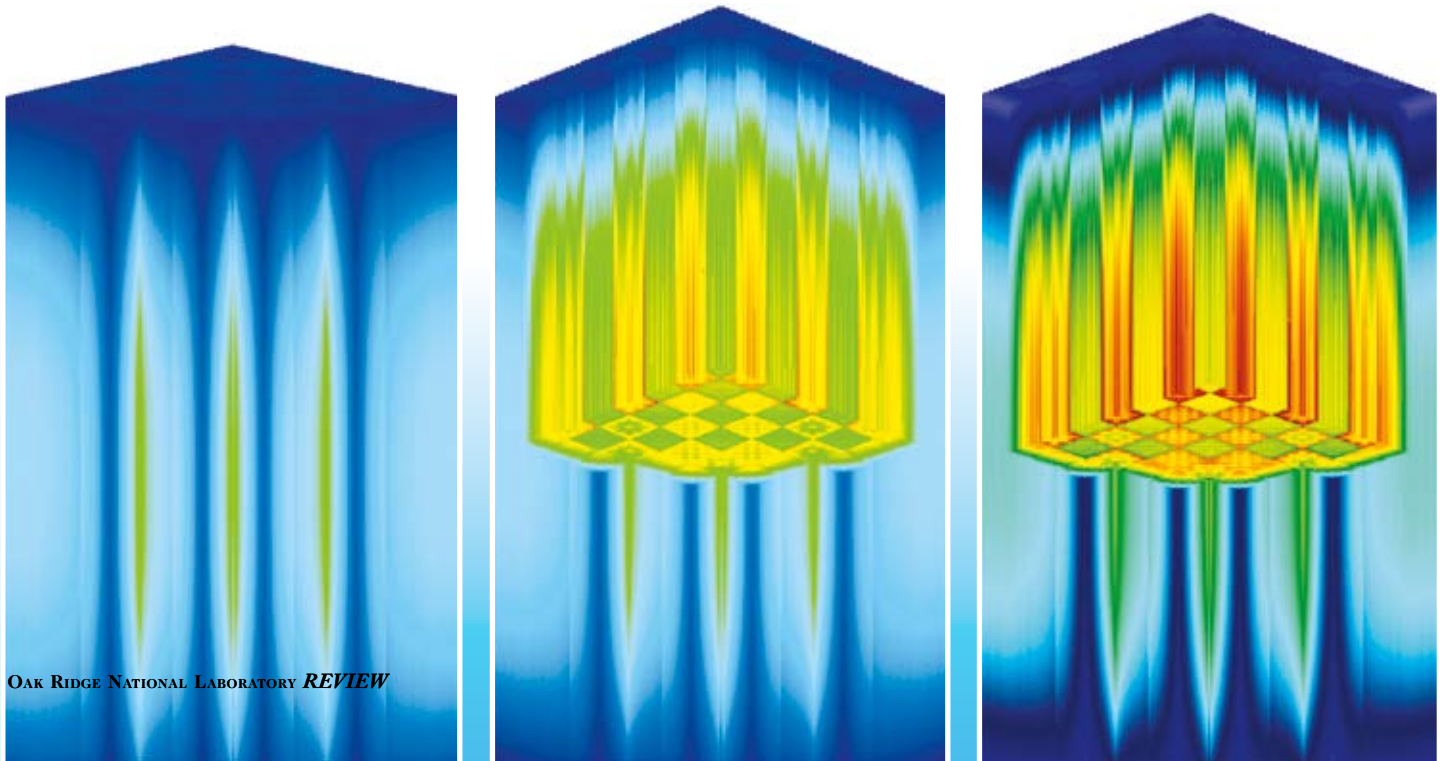

VERA's users can monitor the performance of any part of the reactor core at any point in its fuel cycle, receiving a data-based snapshot of what is happening in each location. Image: Andrew Godfrey, Tom Evans, Greg Davidson and Josh Jarrell

Hybrid path forward

The technology that will enable Titan to support VERA's computational demands is a novel hybrid computer architecture that combines the sophistication of traditional central processing units with the raw data-handling power of graphics processing units. This combination is expected to provide Titan with 10 times the power of its predecessor.

However, with that surge in speed comes a similar increase in the complexity of the software needed to make the most of the GPUs' ability to process thousands of streams of data simultaneously.

"I like to compare moving to a hybrid computer architecture to the transition we made from serial to parallel computers several years ago," Evans says. "It wasn't a drastic change to go from 1 to 4 to 8 to many more processor cores. Obviously things happen when you're using 100,000 cores that don't happen with 1,000 cores, but fundamentally it wasn't a tremendous shift.

"What we're doing now to adapt to hybrid architecture is much more dramatic. We had years to scale from one processor to thousands of processors, so each step up felt like a rolling hill. Getting ready to take advantage of the new hybrid architecture is like going from one processor to 300,000 cores overnight."

Optimizing for Titan

Fortunately, Evans and his colleagues have been able to draw on work that has been done in the laboratory's Center for Accelerated Application Readiness. CAAR has brought together dozens of researchers from national labs, computer vendors and universities to ensure that critical computer codes in a range of research areas will be ready to take full advantage of Titan when it comes online later this year.

"The reality is that GPUs are vectorized machines," Evans says, "which means they are designed to break problems into many threads of operation. So we have to think about programming in a different way. For example, if I want to multiply 5 × 4 on a GPU, it will work faster to do it many times on multiple threads even though the same operation is repeated. GPUs do not perform single operations efficiently. It's an entirely different way of thinking about programming, so it's a fairly steep learning curve."

Evans notes that the driving force behind the move to hybrid supercomputers is not just the need for greater computational power, but the difficulty of sustaining both the current architecture and the current level of power consumption.

"We can't keep scaling the number of chips up indefinitely," he says, "and in order to get more computational power per megawatt, we need to take advantage of new architectures. That's the only way we're going to move forward."

Once Titan is up and running, Evans and his colleagues hope it will be able to complete a simulation of an entire reactor fuel cycle in a week. The data provided by this kind of detailed model would be invaluable to power plant operators, and the one-week timeframe would mesh well with the engineering workflows at most nuclear power facilities.

Evans and his colleagues have tested pieces of VERA on a developmental version of Titan, and preliminary results indicate it's three to five times too slow to meet the one-week goal. However, the next-generation GPUs that will be available when the full version of Titan goes online are expected to make up some or all of the shortfall, providing a 3X to 4X boost in speed.

"We're leveraging our expertise with numerics and algorithms to optimize our code and make the GPUs work harder for us," Evans says, "so we may be able to do this. Creating a model with this level of detail was inconceivable not too long ago, and now it's within the realm of possibility. Will we capture every little piece of physics that happens in the reactor? No. Will we capture a more comprehensive model of the reactor core? I think we will. Will the performance metrics from the development machine really scale up to the full machine? We will find out.

"We've seen a lot of signs that suggest we're on the right path," Evans says.

—Jim Pearce

Joining forces for biofuels

Computing and neutron science overcome bioenergy roadblocks

Computational biophysicist Jeremy Smith specializes in performing science at the intersection of multiple disciplines. Smith, who directs ORNL's Center for Molecular Biophysics and holds a Governor's Chair at the University of Tennessee, is bringing together the laboratory's strengths in supercomputing and neutron science to overcome challenges in interdisciplinary areas such as bioenergy.

Collaborating with scientists from ORNL's Bioenergy Science Center, Smith focuses on a particularly troublesome molecule called lignin. Lignin is a major component of plant cell walls, where it intertwines with plant sugars called cellulose and hemicellulose. Lignin's hardiness provides benefits to plants in the wild but has been a major frustration to bioenergy researchers. During biofuel production, in which plant biomass is converted into alcohol, lignin inhibits the release of sugars necessary for fermentation.

This recalcitrance is a major bottleneck in the production of biofuels from non-food crops such as switchgrass and poplar trees. Smith's team of researchers has been trying to understand how and why lignin behaves as it does, which has proven to be no easy task.

"One of the problems with lignin is that it doesn't seem to have a sequence," Smith says. "If you think of DNA, it has a code, a sequence of letters that makes up your genome, what you are. Lignin is a mixture of chemicals, but in a seemingly random order; there's no sequence to them.

"There is less known about lignin than almost any other biopolymer because it's only found in plants, which makes it less interesting to medicine. It also has a complex random structure, making it difficult to apply many techniques that you normally would to understand polymers and biopolymers."

Tricky translation

To figure out the physical basis of lignin's recalcitrance, Smith combined neutron-scattering experiments at ORNL's neutron facilities with high-performance simulations on lab supercomputers. Using the two resources in tandem required a large amount of theoretical research and technique development to allow experiments to be translated into simulations.

"Neutron experiments tend to produce relatively simple and smooth-looking signals, as they only 'see' a molecule's motions at low resolution," Smith says. "In contrast, data from a supercomputer simulation are complex and difficult to analyze, as the atoms move around in the simulation in a multitude of jumps, wiggles and jiggles. Reconciling these different views of the same phenomenon is a tricky problem."

The complementary techniques of simulation on ORNL's Jaguar supercomputer and neutron scattering at the lab's High Flux Isotope Reactor enabled Smith's team, working with experimentalists from the ORNL Biofuels Science Focus Area, to resolve the structure of lignin aggregates down to 1 angstrom—that's one 10-billionth of a meter, smaller than the width of a carbon atom. These aggregates, known as "clumps," cause problems in biofuel production because they latch on to enzymes used in pretreatment to release sugars from cellulosic biomass.

The combination of experiments and simulation yielded a molecular model of lignin's surface structure at scales ranging from 1 to 1,000 angstroms, giving the research team a broad yet highly detailed view of the molecule. Additional supercomputing simulations revealed how lignin's structure changes during the high-temperature pretreatment process. Smith's team determined that lignin collapses to form problematic clumps even at relatively hot temperatures, rather than only during the cool-down phase, as previously believed.

"Looking at the simulations and experiments together, we've been able to figure out when lignin scrunches up into a ball and when it likes to float around in an extended form," Smith says. "That's very important to know, and it looks like we have a consistent picture. When lignin collapses and when it extends, under what circumstances, and how it interacts, what drives it to interact with cellulose or otherwise—this basic understanding is needed for improving the biofuel production process. It requires both the neutrons and the simulations."

Smith is looking forward to more advanced simulations of bioenergy-related systems.

Toward Titan

As ORNL's flagship supercomputer, Jaguar, morphs into the more powerful Titan, Smith is looking forward to more advanced simulations of bioenergy-related systems beyond lignin.

"Titan allows us to perform simulations of bigger systems," Smith says. "Now we'll be thinking not just of simulations of lignin and cellulose, but we'll be including enzymes and even the microbes—the things that eat up the biomass. We'll be thinking of the interface between the plant surface and microbe surface and maybe even doing simulations of the whole interface at atomic detail."

Smith also anticipates the need for software and methods that must be developed to keep up with the increase in supercomputing speed.

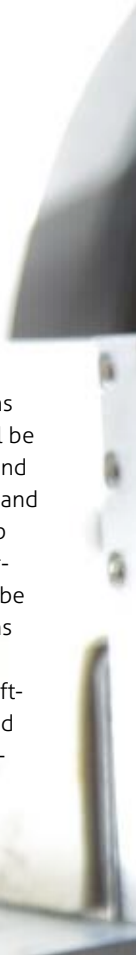

"As supercomputers increased in power, until about 5 years ago, it was fairly simple to just take your program and run it on the faster computer," he explains. "And it would run faster. But that's all changing now. Because the architecture of these computers is becoming so complex, it leads to huge challenges in getting these programs to run quickly.

"When we get to exascale, in principle, we'll be able to simulate a whole living cell at atomic detail—every atom in a cell will be in our simulation. In practice, that's not going to happen without tremendous methods development."

Future footsteps

The full benefits of these pioneering advances in supercomputing may not be realized for years, Smith says, referencing the first molecular simulations, which were made in the late 1970s on computing systems a hundred million times less powerful than those currently available. "Those first simulations were not necessarily very scientifically useful at the time, but they showed you could do it," he says.

"Then later everyone else caught up, and you had this massive amount of information available. In a way, you're putting a marker down; you're saying that we're going here, in terms of system size or level of detail, and proving you can do it. Someone else may come along afterwards and extract even more interesting science, using our methods or following in our footsteps."

Although Smith's research team has made great strides in joining the power of supercomputers with neutron experiments, there is still work to be done. Smith envisions a future in which a scientist who visits ORNL to conduct experiments at the Spallation Neutron Source will have simultaneous access to supercomputing facilities and assistance in using them. Part of this process involves training neutron researchers to design their experiments in such a way that the results can be more easily interpreted using simulations.

"In the future, this approach is going to extend all over the sciences—materials science, bioenergy, condensed matter physics—most of the types of science that SNS deals with," Smith says. "They already use high-performance simulation at SNS on a regular basis to interpret results. But the question for the future is whether Oak Ridge is going to put together a framework for making this nexus between supercomputing and facilities at SNS such that the users have a really unique experience."

—Morgan McCorkle

Neutron-scattering experiments combined with computer simulations help scientists understand roadblocks to the production of ethanol. Photo: Jason Richards

Behind the data

Knowledge discovery provides context for world events

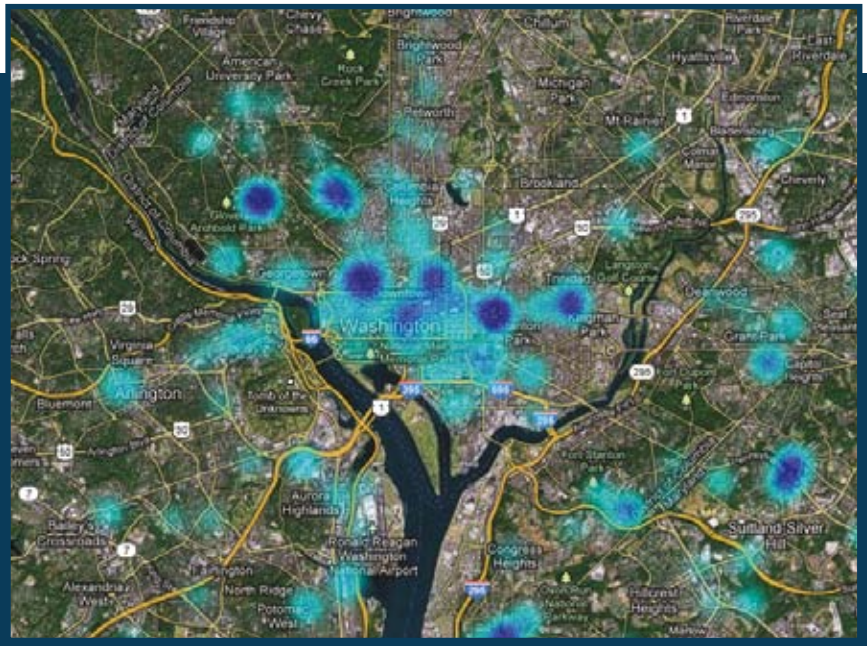

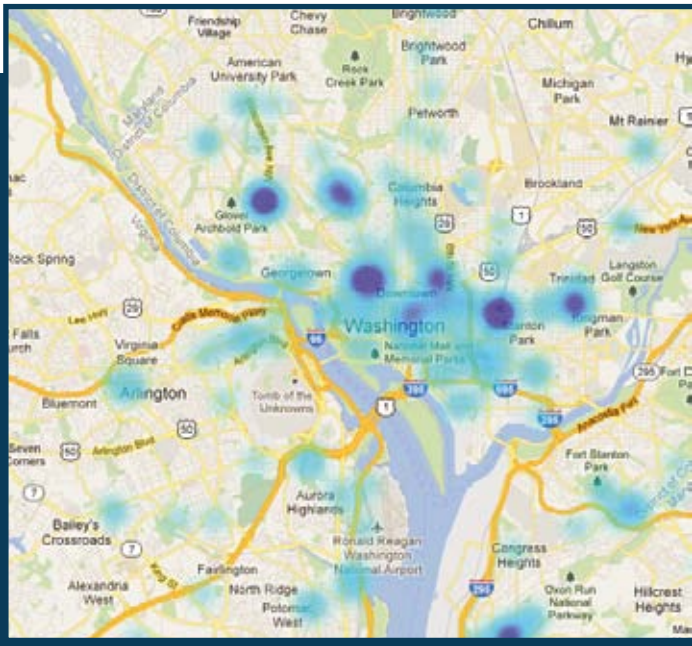

Thanks to the digital revolution, the information available on any given topic has morphed from a sporadic stream of text and photos into a nonstop maelstrom of multimedia shrapnel. Traditional news reports and transcripts are now terabytes of pages, posts and tweets. A billion smartphones upload pix and vids to the net 24/7, and GPS data ties people to places, adding a spatial component to the puzzle.

The need to parse, understand and develop predictions from this barrage of information, particularly for national security purposes, is what the field of knowledge discovery (KD) is all about.

"Knowledge discovery involves working our way from individual pieces of information toward an understanding of what's going on," says Robert Patton, a computer scientist on ORNL's Intelligent Computing Research team. "We try to build a context around events."

That context is provided by analyzing various types of data that provide indirect evidence of other events.

"For example," Patton says, "when the US military found pictures, letters and other documents on Osama bin Laden's hard drives, that data provided part of the context for his life in Pakistan."

Expanding that context may have required finding out: Who wrote the letters? To whom? Who took the photos? What was in the photos? Where were they taken? KD analysts answer questions like these by applying a variety of techniques to understand the relationships among people, places and information and to recreate the story behind the data.

"Usually our projects start when an organization comes to us with a large set of data they want to use to answer certain questions," Patton says. "They're pretty sure the answers are in there, but they don't know how to find them, where to start, or even what parts of the information are relevant."

Plotting patterns

One of the main ways Patton and his ICR colleagues extract context from data is to look for patterns. For example, email traffic among members of a group is used to construct network diagrams. The relationships found in the diagram can help identify key people within the group.

"We might look at the diagram and realize that one guy has sent tons of e-mail, but he never gets any responses," Patton says. "This might suggest he's a leader because no one is questioning his authority."

GPS data is especially helpful in KD because it enables analysts to link information to a location. For example, if there were an explosion at a particular location at a particular time, analysts would want to gather other location-specific data surrounding the event—mobile phone calls, security camera video, eyewitness accounts—and display it on a map to see if it reveals a pattern. The ability to discern consistent patterns in and around events is a step toward being able calculate the likelihood of similar events—or even to predict when they will occur.

Accuracy and speed

Two aspects of knowledge discovery that are critical for national security applications are accuracy and speed.

"If I want to know what the weather will be like next week, but all I know is what the weather is like today, I won't be able to predict a whole lot," Patton explains. "But

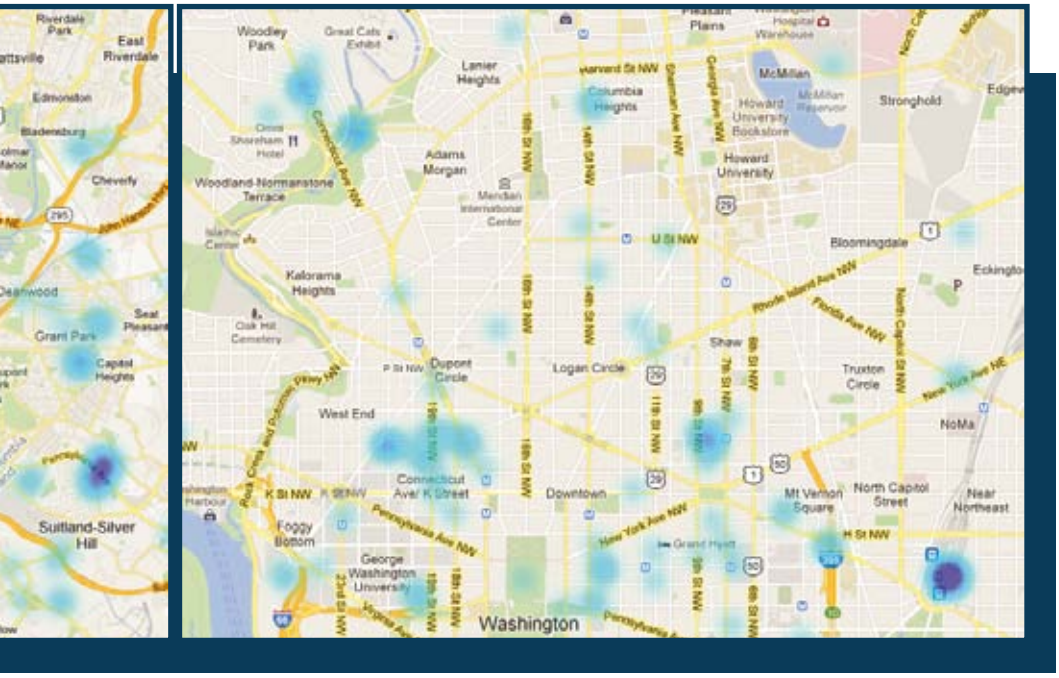

Social media posts often contain location information. Here, tweets generated in the Washington, DC, area are clustered around subway stations. Image: Chad Steed and Chris Maness

let's say I have data for the whole eastern seaboard. That gives me more data, and I can improve my prediction. Generally speaking, the more data you accumulate, the more accurate your predictions become."

However, accumulating data can be a two-edged sword. The more data you have, the harder it is to process it quickly. Patton points to the role played by social media in the recent "Arab Spring" uprisings across the Middle East.

"Social media outlets were the channels through which masses of people were communicating and coordinating their actions," Patton says. "We need analytical systems that can keep up with this volume of data if we want to understand events and be able to respond in time to protect our national interests—or to have the opportunity to influence events in some way. Applying high-performance computing to KD enables us to meet this need."

Speed isn't always the primary consideration for KD, but it often is. Sometimes ICR's national security customers have a very short turnaround for questions that involve analyzing new tactical data. In these cases meeting the deadline is more important than being extremely accurate, so analysts have to consider as much data as they can within the timeframe and give their best answers. Other customers have strategic concerns that operate on longer timelines, so they have the luxury of considering all available data.

"For example," Patton says, "some of our customers who have concerns about social and political changes in another country want to be able to look at the terabytes of data that document the history of the country or the social and political movements involved, consider how similar situations transpired, and be able to project the likelihood of certain outcomes—as well as evaluating things we could do to influence or respond to these outcomes."

Looking forward

In the course of their work on ORNL's Jaguar supercomputer, Patton and his colleagues are breaking new ground in applying KD tools and techniques to unprecedentedly large data sets. Because most of their customers don't have access to computers like Jaguar, Patton likens the customers' relationships with ICR to that between an auto manufacturer and its racing division. The research his group does for customers' specialized, high-performance KD applications today will be applied to improve the KD software used on desktop computers a few years down the road.

"Working on Jaguar today—and on its successor, Titan, that will be available this fall—provides us with an opportunity to see how KD algorithms and techniques perform at scales that are way beyond what anyone else is doing right now," Patton notes.

"For example, last year we tried pushing a huge number of documents through our document-clustering tool called Piranha. It broke at half a million. Once we looked into the problem, It was easy to understand and easy to fix, but because we had never tried to process that many documents before, we never knew the problem was there. If we learn to fix these glitches now, then when our customers need to process that volume of data, the software will be ready."

Getting ahead of the game

Having access to Titan will provide ICR with an environment that can support the massive uptick in the quantity of online data.

"Many of our current algorithms were not designed for handling information on that scale," Patton says. "So when they try to process petabytes of data, they start breaking down. Titan will help us develop new ways to handle that volume of information."

Patton notes that the volume of online data is sometimes so large that it can't be stored now and analyzed later. As a result, next-generation KD applications will not only have to process huge amounts of information, but also do it "on the fly," making the most of the one look they'll get as the data stream passes.

"Titan is going to enable us to keep up with the flow of online information and push our current applications to the breaking point and beyond," Patton says. "It's critical that we get ahead of the game by learning how to handle data on this scale now, so we'll have that capability ready for our national security customers when they need it tomorrow."

—Jim Pearce

Climate models for a changing world

Simulations predict possible climate outcomes

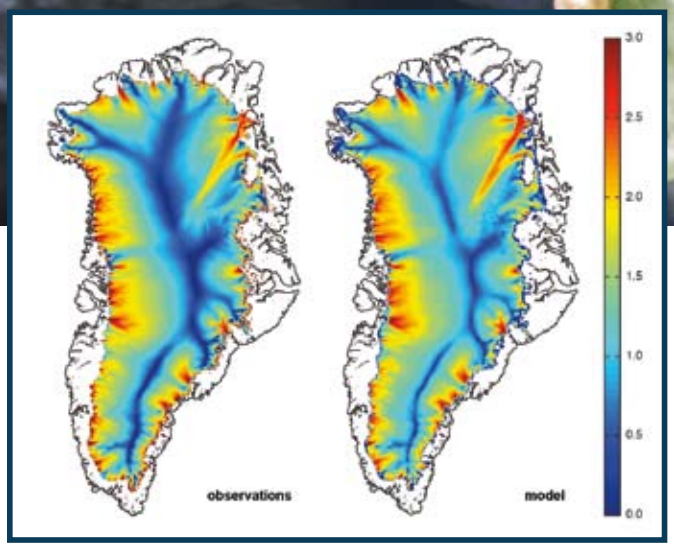

Recent advances in modeling have improved prediction of the movement of ice sheets. The simulated (right) versus the observed (left) flow velocities for Greenland ice sheets match well. Simulation data: Steve Price, Los Alamos National Laboratory; Observational data: Jonathan Bamber and colleagues in Journal of Glaciology, 46, 2000.

Using ORNL's Jaguar supercomputer to capture glimpses of Earth's future, researchers are trying to bridge the gap between science and public policy. Upgrades to Jaguar may make that gap smaller.

Ongoing upgrades to Jaguar, one of the fastest supercomputers in the world, have allowed internationally recognized climate scientist Warren Washington and his colleagues to see Earth's future climate in the long term and with much more clarity.

Washington and Jerry Meehl, climate researchers at the National Center for Atmospheric Research (NCAR), have paired up with Jaguar to shed light on the confounding and often controversial relationship between greenhouse gas emissions and their effects on climate change.

They head a group of scientists and computer experts that has carried out extensive climate change research for the 20th and 21st centuries and beyond for the Intergovernmental Panel on Climate Change assessment reports, including the 2007 report that concluded the Earth is definitely warming and human activity probably contributed to the warming. A subsequent assessment that will take advantage of model improvements is expected in 2014.

"The US Department of Energy and other parts of the federal government want to know how different energy strategies affect the environment," Washington says. "Tools like the computer climate models allow them to ask conditional, 'what if,' questions."

The Community Earth System Model, a virtual time machine supported by DOE and the National Science Foundation, is a set of coupled computer programs used on Jaguar to simulate Earth's past, present and future climates. Predicting what the future has in store requires a mixture of model development subprojects representing numerous Earth system climate components and lengthy, complex simulations of climate variability and change.

"We have to make some assumptions about how society is going to deal with greenhouse gases," Washington says. "For example, if the world predicts a 'business as usual' scenario, meaning the world continues to emit greenhouse gases at the present rate or even faster, our CESM shows future projections of a much warmer climate by the end of the century."

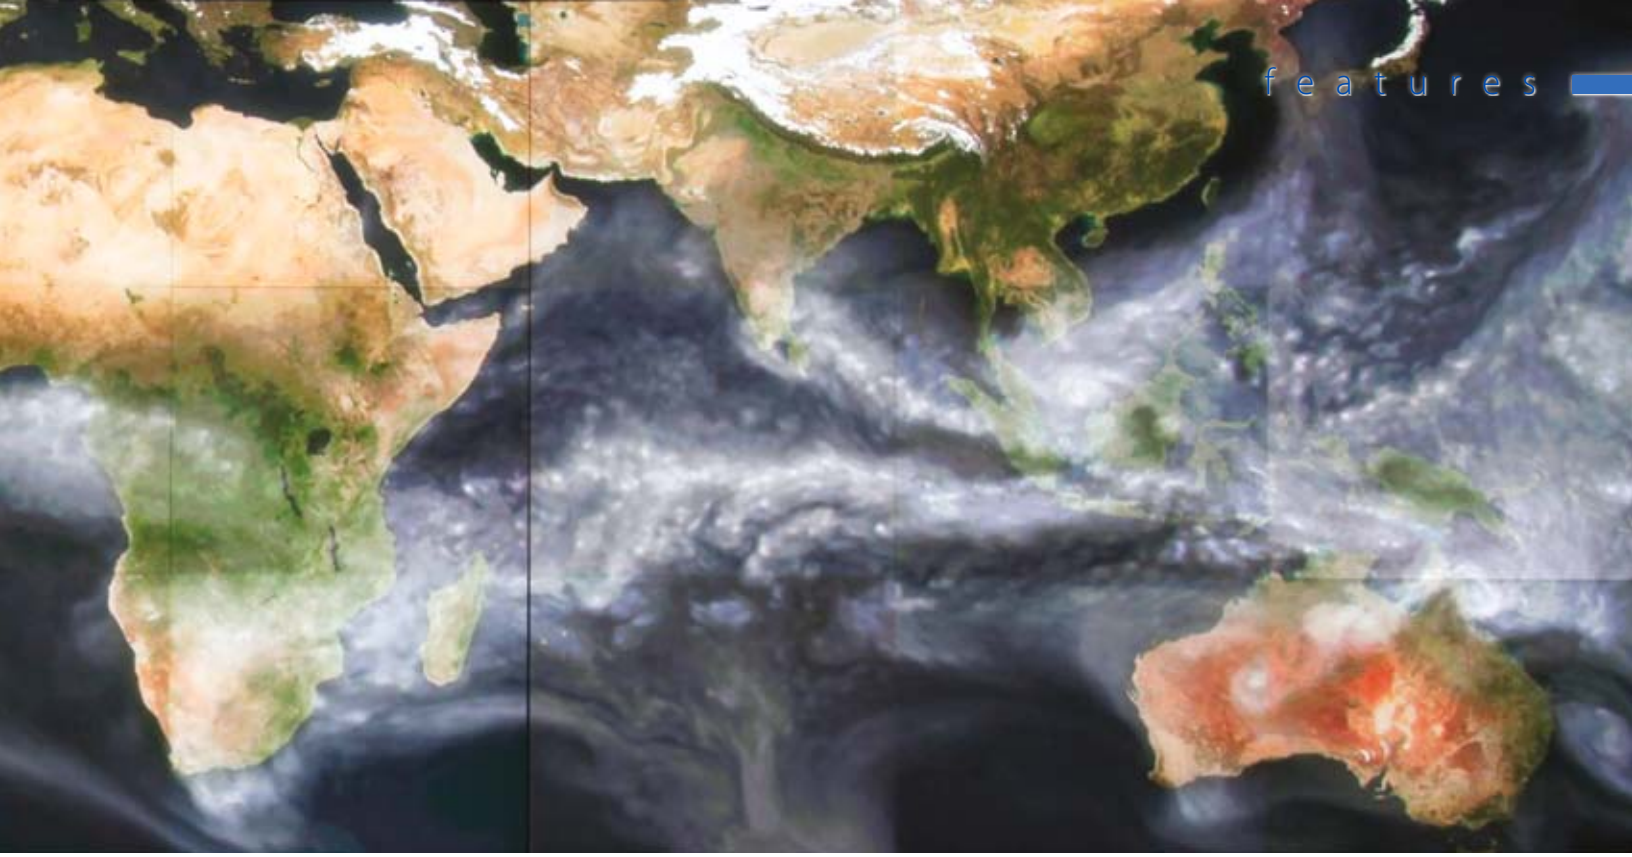

Jaguar can also produce ensembles of experiments that separate natural climate fluctuations from those caused by man. These experiments show what will happen when there are greenhouse gas mitigation efforts. If emissions of greenhouse gases are cut back significantly, the model shows a much more moderate increase in global warming.

"Our research tells the public and the policy makers what future climate change will look like, given different scenarios," Washington says.

In 2012, DOE granted Washington's team and their project, the Climate End Station, a total of 86 million processor hours through the Innovative and Novel Computational Impact on Theory and Experiment program. The team has 56 million processor hours on Jaguar and 30 million processor hours on Argonne National Laboratory's supercomputer to generate climate simulations. This is equivalent to the power of 28 million dual-core laptops for one hour. However, unlike millions of separate laptops, Jaguar's massive array of parallel processors are interconnected, allowing them to perform millions of calculations simultaneously and making more complex simulations possible.

The CESM can use up to 100,000 of Jaguar's almost 300,000 processors at a time. "We carry out the calculation by breaking the globe up into areas to which we assign a certain number of processors," Washington says. Depending on the resolution of the model, the processors may run for hours or days.

These models give Washington and his colleagues a much more detailed picture of potential climate change, including information regarding the warming of the polar regions, the extent and frequency of droughts and whether heat waves will increase in the future.

The team's history of working with Jaguar has allowed them to examine the relationship between carbon dioxide emissions and the potential for global climate change. In the coming months, upgrades to Jaguar will incorporate high-performance graphics processing units, or accelerator chips. The newly installed, highly specialized GPUs will offer a faster way to process information than relying on central processing units alone. When the upgrades are complete, the system will be renamed Titan.

While Titan will have many advantages over Jaguar, Washington's computer code in its current state is unable to take full advantage of the supercomputer's improved power and speed. However, ORNL's and NCAR's computing experts are already working with the CESM codes to take advantage of the added capabilities.

"GPUs are meant to accelerate really heavy parts of the computation, while CPUs handle the lighter parts," said Matthew Norman, a computational climate scientist within the Scientific Computing Group in ORNL's National Center for Computational Sciences.

The CESM incorporates interacting components: Earth's atmosphere, land surface, oceans, sea ice and land ice. Computations involving the atmosphere take up most of the time and supercomputing resources when the team runs the higher-resolution models. This is where Titan's GPUs, which effectively power through dense calculations, will be most useful.

"We will take certain dominant parts of the atmosphere simulation and put those calculations on GPUs where they will run 75 percent more efficiently," Norman says. "Through the use of GPUs, Washington's colleagues will use fewer computing resources to achieve their groundbreaking simulations."

Its use of GPUs will place Titan in a unique position, allowing it to run more sophisticated climate models effectively and efficiently. Researchers will be able to test more phenomena, helping the climate models more faithfully depict the future.

"We are greatly indebted to DOE, and to ORNL in particular, for providing computer time so we can carry out these types of experiments," Washington said. "The use of Jaguar has been vital to pursue this very important research objective for the nation and the world."—*Jennifer Brouner*

Dhiraj Catoor uses a high-resolution x-ray diffractometer to study defect interactions in structural materials.
Photo: Jason Richards

Modeling material defects

Theory and experiment improve our understanding of materials

New technology promises to improve the efficiency of nuclear reactors and to reduce waste, but it also requires careful consideration of how higher temperatures and intense radiation fields affect the materials with which reactors are built.

"Metallurgy is a very old subject, so you might think that by this point we would understand everything about it. However, that's not the case," says Malcolm Stocks, director of ORNL's Center for Defect Physics. The CDP is conducting groundbreaking research into the structure and performance of materials used in nuclear reactors and other applications where radiation damage is a concern.

Radiation damage begins when a material, steel for instance, is struck by a high-energy neutron. The core of a nuclear reactor is full of very energetic neutrons. When one of these strikes an atom in the reactor vessel, it knocks the atom out of its normal position in the lattice-like atomic structure that is characteristic of the steel alloys used in high-radiation applications.

"That atom hits another one, which hits another, and so on," Stocks says. "Pretty soon you have thousands of them knocked out of place. Most of the time, this is a very transient effect, and these atoms very quickly go back to their original positions. The few that don't are the defects that will remain and will eventually interact with other defects to form larger flaws that cause the material to become brittle, to crack, or to swell.

"There are features of the radiation damage process which no one has ever measured," Stocks explains. "They happen extremely quickly—inside chunks of material—so it's very difficult to see them."

Because no measurements have ever been made of the process of defect formation, researchers' understanding of the phenomenon is based on simulations rooted in classical molecular dynamics, which describes the basic physical interactions among atoms. The CDP proposes to achieve a more thorough accounting of the process by making the first-ever measurements of several aspects of defect formation and then creating a much more accurate and comprehensive simulation that will be used to help understand and interpret the experimental results.

One-of-a-kind experiments

Stocks explains that all materials have naturally occurring defects and dislocations in their atomic structure. CDP scientists are working to illuminate the types and numbers of defects that result from irradiation, as well as how these induced defects interact with naturally occurring dislocations and microstructural features, such as "grain boundaries" (the areas between grains of metal where defects can gain a foothold) to change the structure of a material and affect properties such as strength and ductility. These

phenomena are studied using use x-ray and electron microscopy, as well as nanoscale tests of material's mechanical properties.

"Experimentally, we have two ways to see dislocations," Stocks says. "we are using electron microscopy to provide direct imaging of both single dislocations and clusters of defects. We also have researchers who travel to Argonne National Laboratory to use the tightly focused x-ray beams of their Advanced Photon Source to capture the details of the interaction between a single dislocation and the grain boundary. The idea is that we can use very sophisticated experimental techniques to take snapshots of the process of defect formation at very short time intervals and then use that data to develop a simulation of what we measured. We want to do the same thing in the studies of the interactions between dislocated atoms and grain boundaries."

Insights gained through these studies are used to build computer simulations of the structure of these materials in order to identify the "fundamental events" of both defect production and interactions among dislocations and defects.

"This pairing of experimental and theoretical research on the same length and time scales—or as close as possible—is something that was previously not possible," Stocks says. "We are pushing experiment to ever shorter length and time scales while extending theory and modeling to longer scales—eventually to the point where the two overlap. This allows the most direct comparison between theory and experiment."

Ideally, the computer models will reflect the results of the experiments as closely as possible while providing explanations of the process based on molecular dynamics and quantum theory.

Both of these experimental approaches involve doing things that have never been done before, and neither would be understandable without simulation to back them up. High-performance computers like ORNL's Jaguar and its soon-to-be-online successor, Titan, provide the computational power that enables CDP scientists to try to make a link between complex experimental data and simulations.

High-performance pedigree

To accurately simulate the defect formation process, the CDP will need to develop models that comprise very large numbers of atoms—from tens of thousands to millions. That's a big jump from the hundreds of atoms that have traditionally been possible in calculations based on quantum theory, due to the amount of computing power required to carry them out.

"For most simulation methods that use quantum electronic structure techniques, the amount of computational work required is proportional to the third power of the number of atoms," Stocks explains. "However, beginning in the mid-1990s, we developed methods that scale as the first power of the number of atoms—so N rather than N^3."

Without supercomputers like Titan and Jaguar, we wouldn't even consider doing this.

The combination of the CDP's modeling methods and the extra computational muscle provided by Titan brings the center's goal of "tens of thousands to millions" of atoms just within reach.

"The complexity of these massive simulations is the reason we need high-performance computers," Stocks says.

Roughly speaking, each of the computer's processors handles the calculations for some number of atoms in the simulated system. The more computer cores there are, the more data can be processed in parallel, and the more complex the simulation can be.

"Our method naturally maps onto parallel computers because, when we developed this method in the 1990s, we anticipated massively parallel computer architecture," Stocks says. "Our LSMS (Locally-Self-Consistent Multiple-Scattering) software that models interactions among electrons and atoms in magnetic materials is the first documented example of code that ran on a parallel computer at one teraflop—a trillion calculations per second. A derivative of LSMS, called Wang-Landau, was one of the first codes to get to one petaflop—a thousand trillion calculations per second."

CDP scientists have been working on adapting these and other codes to the rigorously parallel architecture of Titan in anticipation of creating more meticulous defect models than were previously possible.

"When we have the opportunity to apply more computing power to a problem, we have two basic options," Stocks says. "We can do more calculations about a small number of atoms—including calculations of the properties of materials, their stability, and their magnetic structure. Or we can do fewer calculations with many more atoms. For some of the questions we want to answer, the minimum number of atoms we need in order to get an insight into the process is many hundreds of thousands or millions.

Faith in science

To make the most of these experimental insights, over the next few years, Stocks expects that the CDP will apply Titan to the task of creating unprecedentedly detailed simulations of the million or so atoms involved in a radiation damage cascade, at very short time intervals—on the order of trillionths of a second. The simulations will take into account not only the physical movement of atoms within the system, but also the magnetic interactions among the atoms and their myriad electrons.

"Without supercomputers like Titan and Jaguar, we wouldn't even consider doing that," Stocks says.

While acknowledging the potential value of the CDP's findings for the next generation of nuclear power plants, Stocks emphasizes that the goal of the CDP isn't to design a better alloy for use in reactor vessels, but to use experimentation and simulation to understand the defects that underpin radiation damage and the ultimate strength or weakness of materials.

"If we're successful," Stocks says, "the CDP will provide theoretical models that have been validated against experiment and then can be used to guide the development of new materials.

"We have this great faith in science; we believe that if we understand defects in these materials at the most fundamental level, that knowledge will help us to determine what happens in more complex real-world situations. Today radiation-resistant materials are developed largely by intuition, trial and error, and experimentation. We're trying to improve on that by developing a better basic understanding of the defect-formation process."—*Jim Pearce*

ORNL's associate laboratory director for Computing and Computational Sciences leads the laboratory's high-performance computing efforts in areas such as climate change, fusion energy, nanotechnology, and biotechnology. This includes managing the Oak Ridge Leadership Computing Facility, which was established at ORNL in 2004 and now hosts several of the world's most powerful computer systems. We asked him to share his thoughts on how supercomputing impacts both R&D and day-to-day living.

JEFF NICHOLS

How does supercomputing touch our daily lives?

Supercomputers affect almost everything we do, often in ways that we don't realize.

Weather forecasters, for example, use supercomputers. Manufacturers depend on supercomputers to fast-track product designs and increase their global competitiveness. Environmental scientists use computer simulations to model the flow of groundwater and protect the drinking water supply. Computing technology touches our lives all the time. We just don't think of it as supercomputing.

After NASA sent astronauts to the moon, we saw a lot of innovations based on discoveries made in the space program. Supercomputing works the same way. Innovations in the field are rapidly spun off into consumer products. Smart phones can operate without burning a

hole in your pocket because the computing and cooling technologies used in these devices were first developed for high-performance computers.

Computer simulations are now used in almost every scientific discipline. What's the attraction?

The attraction is that it allows scientists to predict things with a high degree of accuracy. One example is drug discovery. For the last several years, the pharmaceutical industry has been using supercomputers to help predict which chemical compounds are most likely to be effective for a particular purpose, based on computer models of their chemical structure. The results of this research have enabled drug companies to reduce the time it takes to develop a new drug from about 10 years to two.

The same techniques could be applied to many other areas, like advanced manufacturing, the development of new materials, nanoscience, and molecular machines. New products in all these areas can be designed and then tested with supercomputers—which is pretty cool.

Jaguar has often been called the most powerful scientific computer in the world. How will Titan expand that legacy?

Titan will take advantage of new, faster processors and will also replace many of its standard processors with GPUs (graphics processing units) that are 10 times more powerful and designed specifically to handle computationally -intensive operations. These upgrades will result in a 10X performance boost over Jaguar.

The result will not only be the world's most powerful computer, but like Jaguar, Titan will be a machine that can be used for research across a wide range of scientific disciplines—nuclear energy, climate modeling, etc. Making sure that Titan is as accessible to scientific researchers as Jaguar is a tough job because the upgrades shift the machine's complexity from its hardware to its software to accommodate differences in how GPUs handle information. A 10X improvement in hardware performance will have been achieved, but the software guys will have to write new code that takes full advantage of that improvement by breaking down research problems into tens of millions of parallel elements—that is a very tough assignment.

Titan's contribution to Jaguar's legacy will be to generate 10 times the raw performance from the same amount of energy—while delivering a similar increase in speed for its scientific applications.

You've said that having an understanding of the scientific applications that run on supercomputers is just as important as having the fastest supercomputer. Why is that?

Part of our strategy has been to ensure that we have groups at the laboratory that are doing application development, not only in computational materials, but also in computational chemistry, computational biology, computational fusion science, computational nuclear science, and computational astrophysics. We have computational science groups in each of those areas developing applications to run on our systems. The reason is that those groups then understand the concept of scalability and extensibility and have the ability to generate or develop the next-generation applications that will be able to take advantage of Titan's capabilities and do science on day one.

We are unique, in some sense, compared to the other national labs in that we have computational science skills across all of those domains, and we can apply those folks to developing next-generation applications. I think that is what sets us apart from other national labs. They all have good people and good computational scientists, but I think we have the breadth of applications and the ability to solve science problems across the board that some of the other labs can't.

You came to the laboratory 10 years ago. In the area of computing, what's the biggest change you've seen over that time?

When I came to ORNL 10 years ago, the concept of scalable infrastructure was not as well understood as it is today. It was understood that we had to have a physical presence, and Thomas Zacharia (ORNL's former deputy lab director for Science and Technology) had the vision to build a 40,000-square-foot computer center. At the time, there was no computer to put in it, so Thomas was a visionary when it came to the scalable infrastructure. He also anticipated that we could take advantage of the space by supporting multiple organizations and agencies.

Thomas hired me from Pacific Northwest National Laboratory because I was doing scalable computational chemistry. He wanted the same thing for ORNL—not just in chemistry, but in materials, climate, biology, and all the other research areas.

Today we are extending that vision. In 2000, Thomas delivered the laboratory's first teraflops capability, and eight years later Jaguar delivered the first petaflops performance. Now with Titan, we are building on those achievements and working toward the goal of delivering the first exaflops supercomputer.

ORNL also hosts supercomputers for the University of Tennessee, the National Science Foundation, and the National Oceanic and Atmospheric Administration. What makes the laboratory fertile ground for supercomputing?

The main reason is that we have 600 plus people who take the business of high-performance computing very seriously. We also have unique capabilities in terms of scalable infrastructure—not just power, space and cooling, but all of the resources that surround computing operations. We have the file systems, the archives, the visualization theater—all of the things that enable us to field a well-balanced system and allow us to understand the science that is being produced. That's why these organizations want to locate their supercomputing operations here—because we know how to do it, and we've been doing it very well for the last 10 years. When an organization works with us, they are doing science on a scale that gives them a competitive advantage across the country and around the world.

Oxygen-23 loses its halo

Jaguar enables researchers to choose between contradictory experiments

The oxygen-23 isotope is rare and ephemeral.

This is not the oxygen that keeps your body running. It exists at the edge of the nuclear landscape, where isotopes are always fleeting and most commonly found within exploding stars.

Nevertheless, if we are to understand how the universe is put together, we must understand exotic isotopes such as oxygen-23. A research team from ORNL, the University of Tennessee and the University of Oslo in Norway recently contributed to this understanding with intense calculations of the oxygen-23 nucleus performed on ORNL's Jaguar supercomputer. In doing so the researchers also demonstrated that supercomputer simulation has become an indispensable tool for scientific discovery, on par with physical experiment. Their work is discussed in a recent edition of the journal *Physical Review C*.

The isotope that makes up nearly 100 percent of naturally occurring oxygen is oxygen-16, whose nucleus has eight positively charged protons and eight uncharged neutrons. It doesn't decay and is especially important to sustaining life, both as the stuff that keeps us breathing and as the heavier of water's two elements. When you step on a scale, oxygen-16 is nearly two-thirds of the weight that stares back up at you.

With eight protons and 15 neutrons, oxygen-23 does decay—and quickly. The oxygen-23 nucleus has a half-life of 82 milliseconds, meaning that if you have 10,000 atoms now you'll be down to two or three within a second. Its neighbor, oxygen-24, is believed to be the heaviest an oxygen isotope can get; beyond it lies the so-called neutron drip line, where neutrons will no longer attach to a nucleus. While they may be rare, these and other exotic isotopes are important, at least in part because they challenge current theories of how a nucleus—and therefore the universe—is constructed.

"Our goal is to explain the origin of the elements," explained Gaute Hagen of ORNL, who did the calculations with Øyvind Jensen of the University of Oslo and Thomas Papenbrock of ORNL and UT. "The only way to go beyond oxygen on the nuclear chart is to go via paths that are along the very drip line of the nuclear chart. On a bigger scale, astrophysics is also nuclear physics. What happened in the first milliseconds after the Big Bang, that's nuclear physics as well."

Nuclear physicists have worked for more than 60 years within an approach known as the nuclear shell model, akin to the atomic shell model that governs electron orbits. The nuclear model notes that as you add protons or neutrons to a nucleus, it becomes especially stable at certain numbers such as 2, 8, and 20. These are known as magic numbers. Calculated for protons and neutrons separately, they indicate that a "shell" within the nucleus has been filled. In addition, there are "subshells" in between these numbers where the nucleus is relatively stable, but less so than with a full shell.

Magic isotopes

This model does well describing stable isotopes, but it becomes problematic as you move toward unstable, exotic nuclei. For instance, the model would suggest that oxygen-28, with eight protons and 20 neutrons, is magic, yet no such isotope has been observed and researchers believe oxygen-24 is the heaviest possible oxygen isotope. In fact, a recent paper from Hagen and colleagues in the journal *Physical Review Letters* supports the contention that oxygen-24 itself is magic, although existing theory would say otherwise.

By exploring exceptions to the shell model as it has been understood, researchers seek a deeper understanding of all nuclei, stable and unstable alike.

"The idea is that these naïve shell model pictures of the nucleus do not hold when you go to the very extreme of the nuclear chart," Hagen noted, "where you have very neutron-rich or unstable or fragile systems."

Enter oxygen-23. The oxygen nucleus reaches a relatively stable subshell with 14 neutrons at oxygen-22. The twenty-third neutron is essentially left over. Experimental data from a decade ago suggested that the twenty-third neutron did not even touch the others, but rather hovered over the nucleus as a halo. That conclusion grew from observation that the nucleus had an especially large cross section; in other words, it was very wide.

Data from more recent experiments by Rituparna Kanungo of Saint Mary's University in Halifax, Nova Scotia, disagreed. Working at Germany's GSI Helmholtz Centre for Heavy Ion Research, Kanungo concluded that the cross section of oxygen-23 was far smaller, meaning that it did not have a neutron halo.

The question, then: Which experiment was right?

Kanungo invited Hagen to explore the question computationally, which he did with a first-principles approach known as coupled cluster theory. To ensure the most accurate possible answer, he did the calculations for several oxygen isotopes, from oxygen-21 with 13 neutrons to oxygen-24 with 16. Each calculation ran on about 100,000 processors and used several million processor-hours.

Supercomputer simulation
indispensable tool for scie

A tough calculation

Oxygen-23 is a very neutron-rich nucleus, which until very recently could not be accurately analyzed microscopically with existing supercomputers. Hagen noted that the calculations necessary to handle 23 strongly interacting particles were very complex and required a system of Jaguar's power.

"This work could not have been done without Jaguar," he said. "It would not be possible. It's putting us at the forefront of our field just being at this location."

The calculations showed that Kanungo's conclusion was right and oxygen-23 indeed does not have a halo.

"It's not a halo nucleus, our calculations verified," Hagen said. "We computed the ground state, the radius of the system, and the density profile of the neutrons to see where the neutrons really are. By using those results and the reaction calculation, Kanungo was able within the uncertainty of the experiment to put our calculations in there, and there was very nice agreement between theory and experiment."

The study of this isotope is a piece in the puzzle of nuclear physics. As supercomputers grow in power, they will tackle ever-heavier isotopes and, possibly, help us understand how matter is put together.

"It affects our understanding of the universe, its evolution, and why we're here," Hagen said. "All those questions relate to what we do, and they are really not located at the stable part of the nuclear chart. They're at the extremes."—*Leo Williams*

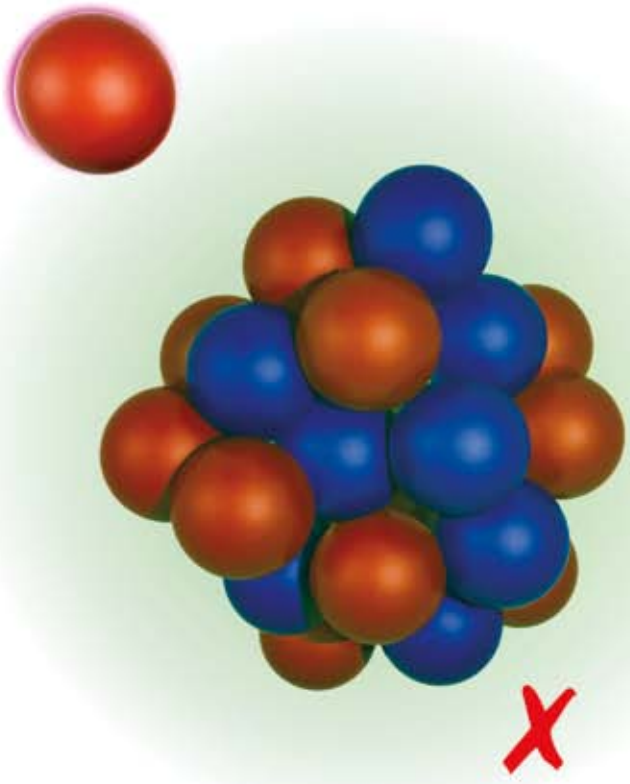

A halo resides outside the nucleus. Oxygen-23 does not have such a nucleus. Image: Andy Sproles

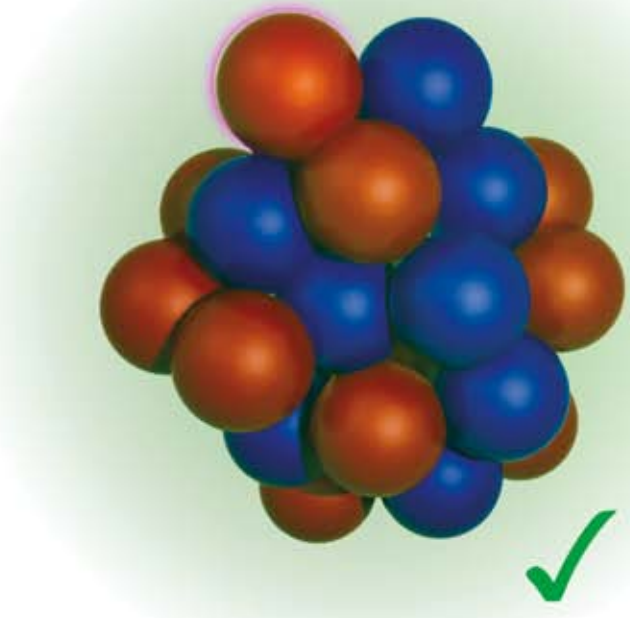

n has become an ntific discovery.

ORNL/UTK team maps the nuclear landscape

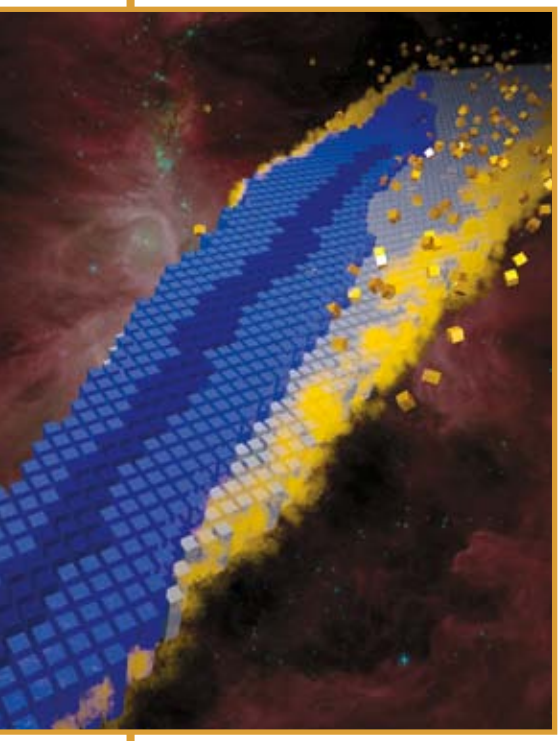

This image represents the nuclear landscape, with isotopes arranged by an increasing number of protons (up) and neutrons (right). The dark blue blocks represent stable isotopes. The lighter blue blocks are unstable isotopes that have been observed. The gray blocks are bound isotopes that have not been 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: Andy Sproles

An ORNL and University of Tennessee team has used the Department of Energy's Jaguar supercomputer to calculate the number of isotopes allowed by the laws of physics.

The team, led by Witek Nazarewicz, used a quantum approach known as density functional theory, applying it independently to six leading models of the nuclear interaction to determine that there are about 7,000 possible combinations of protons and neutrons allowed in bound nuclei with up to 120 protons (a hypothetical element called "unbinilium"). The team's results are presented in the June 28 issue of the journal *Nature*.

Most of these nuclei have not been observed experimentally.

"They are bound, meaning they do not spit out protons or neutrons," Nazarewicz explains. "But they are radioactive—they are short-lived, because there are other processes, such as beta decay, that can give rise to transmutations."

Of the total, about 3,000 have been seen in nature or produced in nuclear physics laboratories. The others are created in massive stars or in violent stellar explosions.

The computations allowed the team to identify the nuclear drip lines that mark the borders of nuclear existence. 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 to the nucleus, 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. Placement of the drip lines for heavier elements is based on theoretical predictions extrapolated far from experimental data and is, therefore, uncertain.

The closer an isotope is to one of these drip lines, the faster it decays into more stable forms. Particle accelerators have been unable to identify most of these exotic isotopes, especially those approaching the neutron drip line, because they are impossible to produce using current combinations of beams and targets. In fact, says Nazarewicz, all radioactive isotopes decay until they are transformed into one of 288 isotopes that form the so-called "valley of stability." These stable isotopes have half-lives longer than the expected lifetime of the solar system (about 4.6 billion years).

Earlier estimates of the nuclear landscape varied from as few as 5,000 to as many as 12,000 possible nuclei, Nazarewicz notes. He said his team's calculations were based on the microscopic forces that cause neutrons and protons to cluster into nuclei, adding that results from the six separate models were surprisingly consistent. By using several models, theorists were able for the first time to quantify uncertainties of predicted drip lines.

Because most of these nuclei are beyond our experimental reach, he explains, models must conform to known nuclei in a way that allows researchers to extrapolate results for exotic nuclei. Insight on the nature of most exotic nuclei must be extrapolated from models, he said.

"This is not a young field," Nazarewicz notes. "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."

The calculations themselves were massive, with each set of nuclei taking about two hours to calculate on the 244,256-processor Jaguar system. Nazarewicz notes that each of these runs needed to include about 250,000 possible nuclear configurations.

"Such calculation would not be possible two to three years ago," he says. "Jaguar has provided a unique opportunity for nuclear theory."

Nazarewicz notes that this work, supported by DOE's Office of Science—which also supports the Jaguar supercomputer—and by the Academy of Finland, has both existential value, helping us to get a better understanding of the evolution of the universe, and potential practical applications.

"We are not doing nuclear physics just to see whether you can get 7,000 species," he explains. "There are various nuclei that we can use to our advantage, eventually. Those we call 'designer nuclei.'"

Among these valuable nuclei are iron-45, a collection of 26 protons and 19 neutrons, which may help us understand superconductivity between protons; a pear-shaped radium-225, with 88 protons and 137 neutrons, which will help us understand why there is more matter than antimatter in the universe; and terbium-149, with 65 protons and 84 neutrons.

Terbium-149 has shown an ability to attach to antibodies and irradiate cancer cells without affecting healthy cells.

"They have done experiments on mice and now humans in which they would look at the effectiveness of this treatment," Nazarewicz said. "This treatment is called an 'alpha knife.'

"Applications will certainly follow from the basic knowledge."—*Leo Williams*

next issue

Disruptive materials

on the web

www.ornl.gov/ornlreview

OAK RIDGE NATIONAL LABORATORY REVIEW
P.O. Box 2008, Oak Ridge, Tennessee 37831-6266

ORNL *Review*

Editor—Jim Pearce
Writers—Jennifer Brouner, Morgan McCorkle, Jim Pearce, Leo Williams
Designer—LeJean Hardin
Graphic Artist—Andy Sproles
Technical editors—Deborah Counce, Walter Koncinski
Photographers—Jason Richards
Web developer—Cindy Latham
Stock images—iStockphoto™

Phone: (+1) 865.574.4160
Fax: (+1) 865.574.0595
E-mail: ornlreview@ornl.gov
Internet: www.ornl.gov/ornlreview

Oak Ridge National Laboratory is managed by
UT-Battelle, LLC, for the U.S. Department of Energy
under contract DE-AC05-00OR22725

ISSN 0048-1262

UT-BATTELLE

OAK RIDGE National Laboratory

U.S. DEPARTMENT OF ENERGY