CENTER FOR COMPUTATIONAL SCIENCES

2005 ANNUAL REPORT

Center for Computational Sciences

Located at Oak Ridge National Laboratory
P.O. Box 2008, Oak Ridge, TN 37831-6161
TEL 865-241-7202
FAX 865-241-2850
EMAIL help@nccs.gov
URL www.nccs.gov

The research and activities described in this report were performed using resources of the Center for Computational Sciences at Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC0500OR22725.

Contents

1

CCS Overview

CCS Overview

The year 2005 was groundbreaking for CCS. It was the organizational year: we installed the two primary supercomputers, the Phoenix and the Jaguar, and completed the acceptance process for them; awarded the first time allocations on those computers to five "grand challenge" science projects; and increased staff by 50%. In addition, we established the Cray Supercomputing Center of Excellence that will reside at CCS and work with Cray users all over the country. It was a year of huge new efforts and successes.

Jeff Nichols, Director of Computer Science and Mathematics, led the CCS throughout 2005 up to the inception of the LCF Project in 2006

In one year, CCS has become the largest open scientific computing facility in the nation. ORNL was selected to establish the U.S. Department of Energy's National Leadership Computing Facility in 2004. The task in 2005 was to execute the plan that won us the award. In 2006, CCS is seeing the results of breakthrough science in significant papers and journal articles and awards for high-impact scientific discoveries through simulation.

Organizationally, CCS became a separate ORNL division in 2005 to focus the resources of the CCS user facility on users and enable them to accomplish breakthrough science. To carry out that vision, we formed four new groups, all focused on providing support to facility users: User Assistance and Outreach, Scientific Computing, High-Performance Computing Operations, and Technology Integration. CCS also formed a Future Technologies group within its sister organization—Computer Science and Mathematics—to keep it in the forefront of new computing technologies being developed.

CCS hired Ph.D. group leaders from Lawrence Berkeley and Pacific Northwest National Laboratories, Ames Laboratory, and Northrop Grumman Corporation and hired many new staff members from other top-notch research institutions. Bringing new employees on board from other leading computing centers allows CCS to benefit from their experience with what has worked in those places.

CCS is not your typical supercomputing center: It is dedicated exclusively to leadership computing and high-impact science. Other supercomputing centers must provide for any researcher who needs their resources; they may support thousands of users, each using a limited amount of resources. CCS supports only a few dozen teams using large allocations of computer time to pursue potentially groundbreaking research. It focuses on researchers who need high-end resources and can take advantage of them to do breakthrough science. CCS wants to offer those users 100 times the computing resources they can get at other supercomputing centers, plus the support services to use them.

Research projects are selected for CCS through annual calls for proposals. In 2005, when the

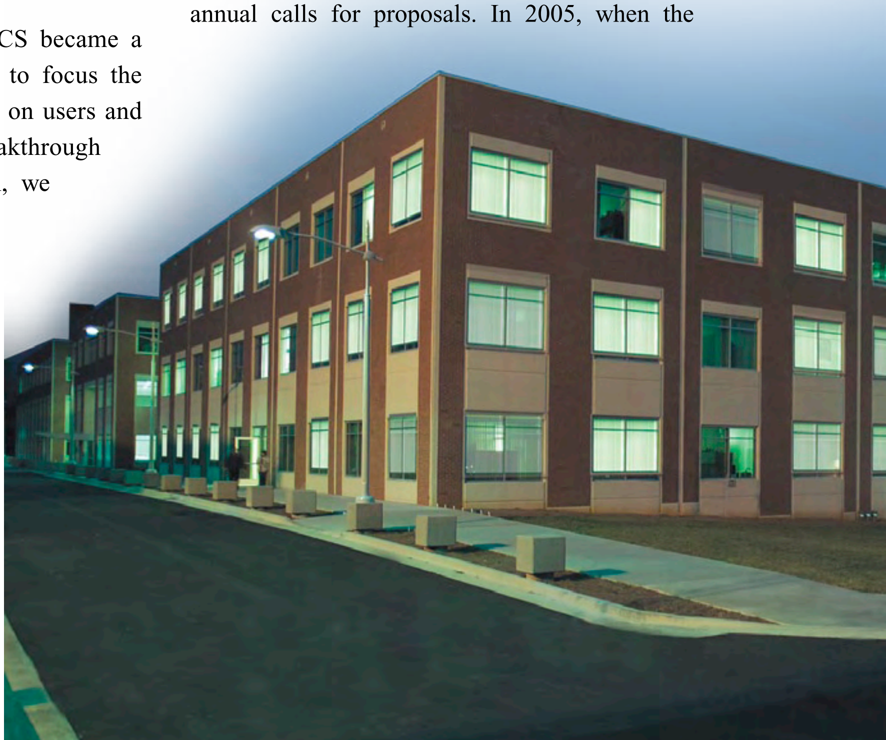

center was still in the process of ramping up its capabilities, five allocations were awarded to grand challenge projects in chemistry, combustion, astrophysics, accelerator physics, and fusion simulation. We are already aware of breakthrough achievements in some of those projects.

The call for proposals for 2006 resulted in 22 research awards, each allocating a total 3 to 4 weeks of dedicated computing time on thousands of processors.

During 2005, CCS went from a 6.4-TF Cray X1 to the 18.5-TF Phoenix and the 25-TF Cray Jaguar. We have set a very aggressive path toward upgrading hardware. Jaguar will be taken to 100 TF, a fourfold increase, in 2006 and 250 TF in 2007. We will add a petaflops system before the end of the decade.

Computing capacity is growing exponentially. There will not be a time when we're caught up and can rest on our accomplishments. CCS must provide numerous and ever-increasing capabilities, and it must provide 100 times as much of those capabilities as other computing centers. As the leadership computing institution, CCS must lead in providing computing resources.

The exponential increases in computing capacity bring expanding needs for data storage, visualization, and networking. CCS has unmatched network bandwidth capability. It has multiple 10-gigabits connections to all major networks: ESNet, UltraScience Net, Teragrid, Internet2, Cheetah, and National Lambda Rail. Researchers located anywhere in the United States can access CCS, and the data generated by its computers can be moved to any other site.

The CCS goal is to provide leadership computing for the nation, regardless of the agency or affiliation of the researchers. Researchers from government, universities, and industry are encouraged to apply for allocations. CCS is the only large computing center in the Southeast, and we want to tap into the unique resources available here, such as the enormous amounts of power available from the Tennessee Valley Authority.

CCS is about doing groundbreaking science in a synergistic way with theory and experimentation. Our goal is to integrate simulation with theory and experimentation as an equal partner and make it a force that can lead theory and experiment. For example, we plan to give researchers the capability to predict through simulation what they will see in their experiments on instruments such as the Spallation Neutron Source and thus better understand the experimental results.

In 2006, CCS transitioned its operations to focus on the Leadership Computing Facility (LCF) Project, led by Buddy Bland. While it continues to provide CCS with high-performance resources, the LCF Project is concentrating its energy on the development and deployment of the first open-access petaflops computer system—100 times more powerful than current leadership-class computers.

There are many reasons CCS is poised to be successful. It's up to CCS to implement the plan that carries out our vision.

CCS building at night

"OUR GOAL IS TO INTEGRATE SIMULATION WITH THEORY AND EXPERIMENTATION AS AN EQUAL PARTNER AND MAKE IT A FORCE THAT CAN LEAD THEORY AND EXPERIMENT."

—Jeff Nichols
2005 CCS Director

2

Science Perspectives

Science Perspectives

CCS is about reshaping the process of scientific discovery. For centuries, the scientific method has relied upon two "legs": theory and physical experimentation. Computational science (using computers to analyze scientific problems) has firmly established itself over the past several years as the third leg of the scientific method. With terascale (and soon petascale) parallel scientific simulation on the leadership computing platforms at the CCS, computational science can truly unite theory and experiment in a "numerical laboratory."

There are several DOE and NSF high-performance computing centers in the United States that are invaluable resources for researchers. But no other center provides the level of sheer computing power, the collaborative model, or the dedication to groundbreaking science that make CCS the national leadership computing facility. CCS exists to enable big science: research with the potential for breakthrough discoveries that will define our future. It provides researchers enough computing capability (thousands of fast processors) and enough allocation time (typically millions of CPU hours) to run the largest simulations ever performed in their respective fields. CCS computers are terascale-class systems: they can perform trillions of floating point operations per second. To put that into perspective, a supercomputer operating at a trillion calculations per second could read information equivalent to the 5 million volumes in the New York Public Library in about 5 seconds.

In the pages that follow, you will read about the projects awarded allocations at CCS during 2005: understanding how massive stars explode into supernovae, how turbulence dissipates energy both in conventional combustion devices and in fusion reactors, exact solutions to quantum problems associated with chemical catalyst design, designing an accelerating cavity for a linear collider, and why high-temperature superconducting materials can conduct electricity without resistance. Only massive computational capability can move science like this forward because the calculations require an enormous number of operations that only a terascale computer can complete in a reasonable turn-around time (e.g., days as opposed to years).

It's appropriate to say that the resources being provided at CCS and the talent coming together here constitute a revolution in scientific computing, because we are actually creating an environ-

Jamison Daniel of the Scientific Computing group manipulates an image on the EVEREST PowerWall ment where the likelihood of discovery happening is much greater than ever available previously. CCS is fertile ground for discovery. Access to leadership-class computers allows researchers to ask bigger, more complex questions than they can explore at other computer centers.

CCS contributes more than big, fast computers to the revolution; it provides a fully integrated partnership between CCS staff and user teams. Our staff do not merely support—they collaborate. CCS has staff members who serve as liaisons between CCS and the research teams using the computers. They work directly with the scientists to help them use CCS resources efficiently, for example, helping them port and tune applications and trouble-shoot codes. These staff members possess advanced science degrees and have research experience and portfolios of their own. By bringing their expertise and experience to the table, they help make breakthrough science more achievable.

Another CCS innovation is the end-station model, long-term allocations awarded to specific user communities who coordinate breakthrough research in their fields with model and code development and optimization. End station allocations are awarded only to projects in which the simulation tools are recognized formally by peers and have matured sufficiently to represent a computational laboratory. Just like an experimental lab, the tools can accept new researchers coming in with new ideas and questions. The end station is a working mode of research in some fields—climate is an example—in which the community accepts simulation and computational science as the principal research arm in the field, has rallied around a suite of tools, and has contributed collectively to the tools so they are essentially community property.

The areas on which CCS resources are focused have great potential to increase our fundamental understanding in ways that will benefit society. In computational chemistry, for example, to be able to design chemical catalysts at the nano level is a breakthrough that would allow much more efficient, cost-effective design of catalysts for the pharmaceutical and oil industries. (Incredibly, in that area, the most advanced computers are still probably a factor of 1000 away from the computing power we need to fully solve this problem.)

In fusion research, simulation is being used as the primary design tool for an experimental fusion reactor. The impact of such a reactor would be huge—a long step toward a virtually inexhaustible energy source that creates no hazardous waste or greenhouse gas emissions. Finally, climate researchers are using CCS resourses to predict atmospheric greenhouse gas concentrations and their impact on global warming for international panels that set environmental policies worldwide. This research will affect us in ways we don't even think about—for example, better, longer-range hurricane prediction.

Doug Kothe
Director of Science

Simulation is not only the third leg of the scientific method "stool"; in many cases it is the strongest leg or replaces another one. Astrophysics is a good example—you can't set up a supernova core collapse experiment in a laboratory. Simulation is never divorced from experiment and theory, but in many cases it guides and validates theory and helps design experiments. It enables scientists to explore "what ifs" and bound the possibilities of a hypothesis.

Users of the CCS resources feel privileged to have access to this tremendous asset—it is a unique resource. They understand that here they have a chance to really go after breakthrough science.

Big Stars End with a Bang

Great stars don't go gently. They die by collapsing into themselves and then exploding, flinging matter far into space. It is these explosions that provide the stuff of the world as we know it: they open up the stellar furnaces where, over the eons, all the elements heavier than oxygen have been forged. Our Sun consists of the debris from their demise. Our planet's atmosphere, its skeleton of rock and skin of soil, its blanket of fields and forests, and the elements in our own bones and blood all have their origin in the death throes of ancient stars.

Stars are powered by the fusion of lighter into heavier elements. As heavier elements form, they naturally sink toward the center of the star. Eventually these elements fuse into iron, which is too tightly packed to fuse further, and the star begins to accumulate an iron core. In a massive star, this core will grow until its own gravity causes it to implode. As it does, enormous pressure compresses the inner core into a super-dense mass that bounces a shock wave back to meet the collapsing outer core. A key consequence of this collapse and bounce is that copious amounts of weakly interacting particles called neutrinos are produced, as electrons are driven into the increasingly heavy nuclei in the core. Because their interaction with matter is so tenuous, neutrinos are the only agents capable of transporting energy out of the extremely dense inner core.

What happens when that shock wave meets the infalling matter is one of the mysteries being explored at CCS by the Terascale Supernova Initiative (TSI) of the DOE Office of Science. Tony Mezzacappa of Oak Ridge National Laboratory, TSI project leader, and John Blondin of North Carolina State are two of about two dozen researchers from 11 research institutions using CCS to investigate the processes by which a massive star (at least eight times the mass of our Sun) explodes in a supernova. In addition to ascertaining the core collapse supernova mechanisms, Mezzacappa's team is trying to understand supernova phenomenology such as element synthesis, neutrinos, gravitational waves, and gamma ray signatures and provide a theoretical foundation in support of DOE Office of Science experimental facilities. Using the Phoenix supercomputer, they have produced a 3-dimensional simulation of a core-collapse supernova that provides fundamental insights into the physics of the event.

Previous simulations indicated the shock wave loses energy and stalls as it hits the collapsing outer core. However, this stalled shock wave (called a standing accretion shock) re-energizes within

Volume rendered sequence of the development of the stationary accretion shock instability in a 3D simulation

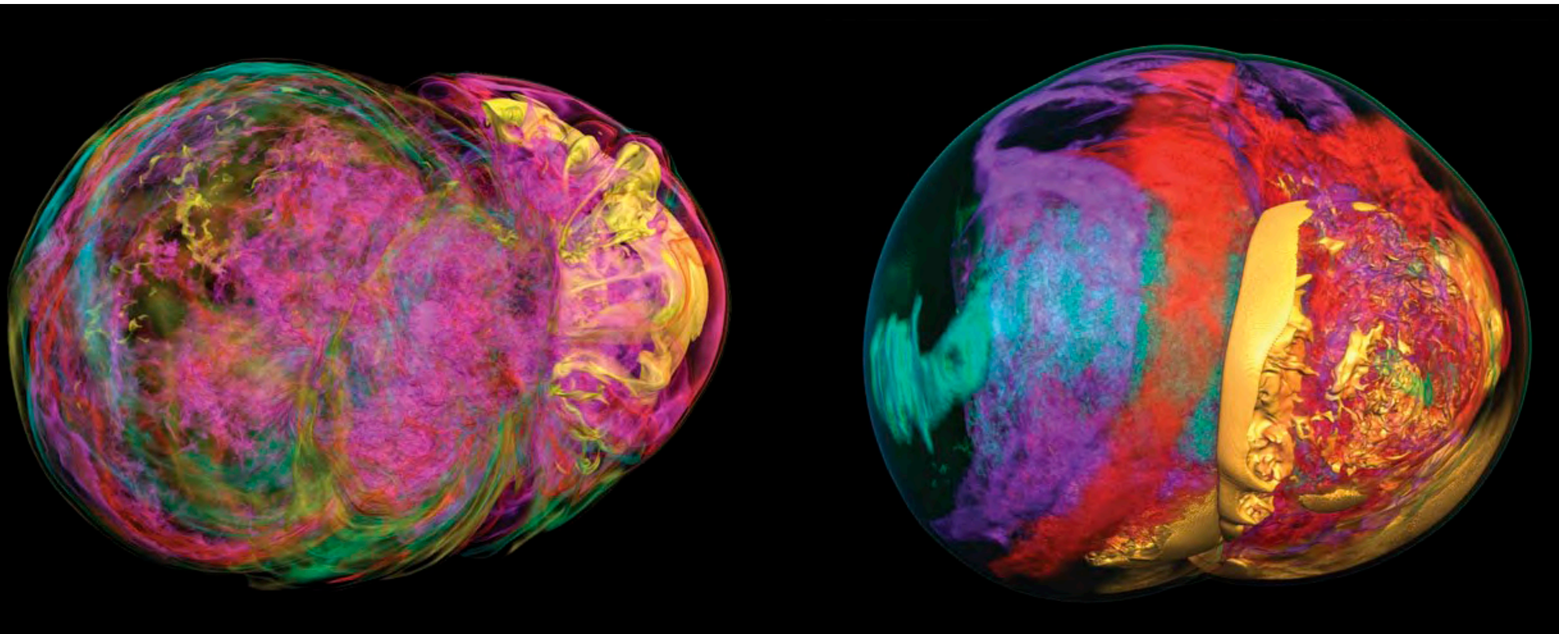

milliseconds to blast the outer layers of the star into space. How it restarts is a key question. From earlier simulations, TSI astrophysicists postulated that incoming matter generates pressure and sound waves that bounce against the shock wave violently enough to jar it back into motion—a standing shock accretion instability (SASI). They were eager to see if the 3D simulations on Phoenix would confirm the SASI hypothesis.

Not only did the SASI show up in the Phoenix simulations—to the surprise of the researchers, it grew and evolved into a rotation, a spin, disrupting the symmetric nature of the core collapse.

"The SASI induces counter-rotating flows of stellar matter on the inside of the star," Mezzacappa explains. "As matter spins and accretes on the central object of the simulated star, it deposits angular momentum on the central object, spinning it up. We started with no spin, and our simulation generated an object that spins at tens of milliseconds. This was an exciting discovery."

That discovery elucidates another mystery: As it explodes, a supernova's core becomes a neutron star, and most neutron stars spin violently for a time. (Astronomers know these stars spin because they emit beams of radio waves that seem to pulse dozens of times a second—thus they are called "pulsars.") The TSI simulation provides a plausible mechanism for how a supernova morphs into a newborn, fast-spinning pulsar.

The 3D simulations of supernovae achieved at CCS are far more realistic than 2D and 1D models. "In our multidimensional simulations, we also take into account other parameters, such as neutrino direction, neutrino energy, and time." Mezzacappa says. "We predict that the SASI, along with neutrino transport and the magnetic fields from within the star, affect how the shock wave generates the explosion. These findings are important contributions to astrophysics theory."

Not until terascale computing became available was it possible to conduct the realistic multidimensional simulations essential to determine how a supernova explosion occurs and explore the phenomena that accompany it. The fast vector processors of Phoenix make these 3D simulations possible. Each run produces tens of terabytes of data, about 100 TB over the course of the project.

This new capability to begin analyzing the processes that drive supernovae is intrinsically important in gaining insight into how the universe behaves and came to be the way it is, says Bronson Messer,

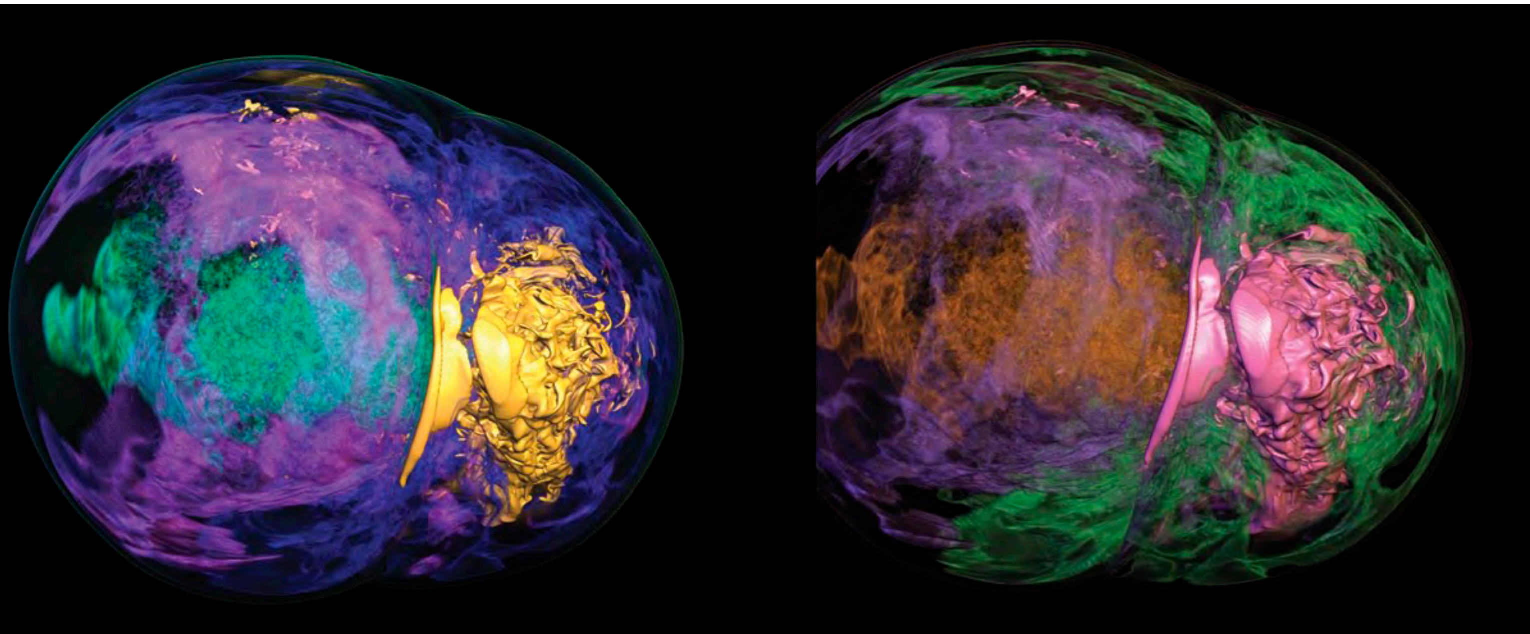

a CCS staff member and liason. Exploding stars crunch a handful of light elements into all the other elements that make up our world, and they account for some of the most bizarre of all physics phenomena: neutrinos, black holes, and gravitational waves. We have no instruments that can physically observe and analyze, millisecond by millisecond, the progress of a supernova far outside our galaxy. Thus supernova simulation is the only tool we have to reveal what these events tell us about the behavior of matter.

With petascale computers available at CCS, the PSI teams will incorporate more realistic 3D physical models, such as general relativity (as opposed to Newtonian), sophisticated neutrino transport, no imposed symmetry, and better treatments of turbulence and magnetic field effects. They may also simulate aspects of hypernovae, a newly discovered type of supernova in which the shock dies and the massive star collapses into a black hole that ultimately produces gamma-ray bursts.

Mezzacappa says his team may be able to do ten parallel 3D simulations in a year on a 1-petaflops supercomputer (CCS is on track to develop a petascale computer by 2008). Such a computer, 1000 times more powerful than a terascale machine, will be able to make a quadrillion calculations per second and will have hundreds of terabytes of memory.

Astronomers observe supernovae throughout the universe almost daily, but it has been more than 400 years since the German scholar Johannes Kepler observed a supernova in our own Milky Way. "The best way to validate our code simulations is to compare the results of our 3D petascale calculations with actual data from a real supernova in our galaxy," Mezzacappa says. "If I could play God, I would not give humans a galactic supernova until 2010. By then we will have far more sophisticated models with great predictive ability. These models will be more receptive to being proved correct or incorrect by detailed observations." There is no predicting when the universe will grant us a supernova in our own galactic neighborhood, but when it happens, Mezzacappa and his collaborators plan to have the tools to take advantage of the opportunity.

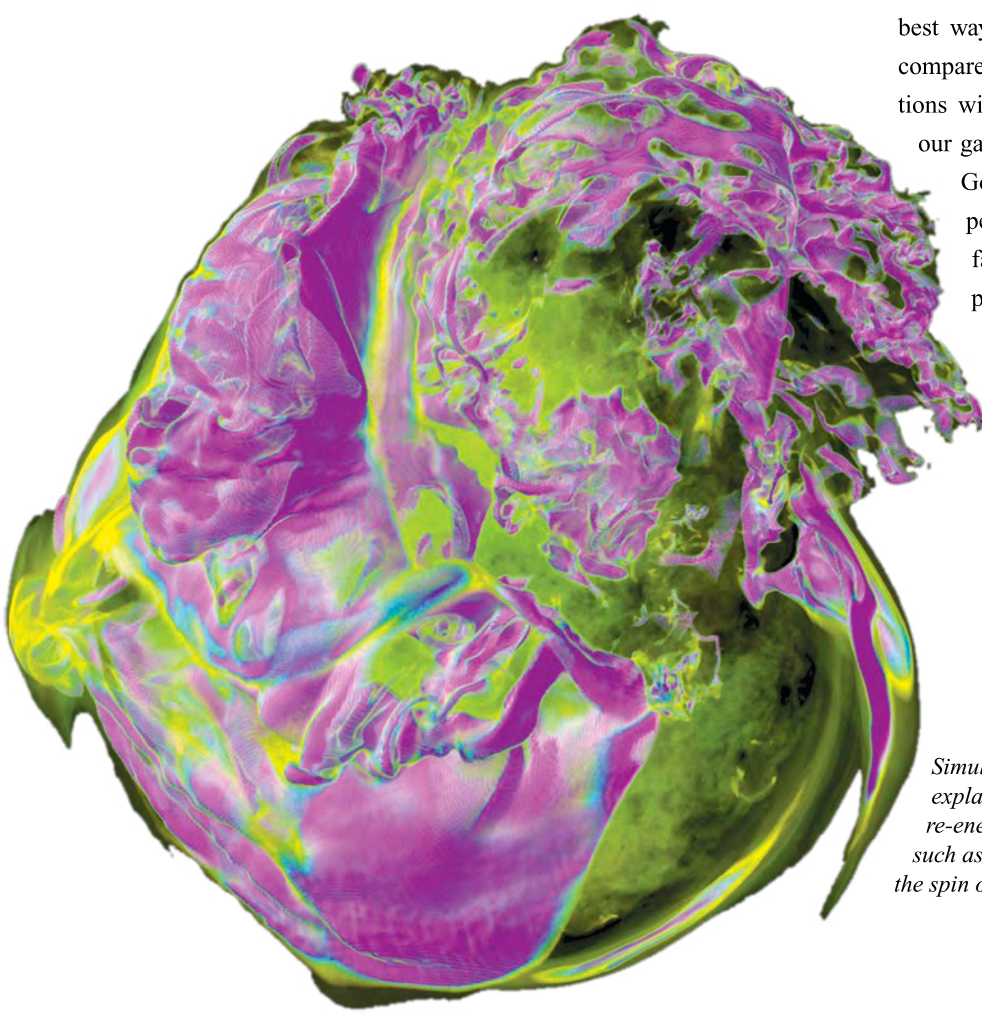

Simulations are helping to explain how the shock wave is re-energized and observables such as neutron star "kicks" and the spin of newly-born pulsars

This remnant of supernova 1987A, the collapse of a star in the Large Magellanic Cloud galaxy. Clouds of gas surround the remnant and its inner and outer rings of material (slightly above center of image) (Courtesy of the Hubble Heritage Team, NASA)

Studying Flames in the Wind

Combustion powers the machinery of modern life. We may aspire to a world powered by fuel cells and fusion reactors; but for the foreseeable future, our electricity, transportation, manufacturing, and heating/cooling all depend overwhelmingly on the burning of hydrocarbon fuels.

At the same time, finite, costly supplies of those fuels and harmful emissions produced by hydrocarbon combustion threaten the very way of life they support. So developing cleaner-burning, more efficient devices for combustion is essential to building a sustainable energy infrastructure for the near and mid-term.

Improving the design of combustion devices (e.g., power turbines, vehicle engines, furnaces) has traditionally been slow and incremental because it requires multiple iterations of making small design changes and then building hardware to test the results. Combustion device manufacturers want to eliminate the need for most of the cumbersome, expensive hardware testing by using numerical simulations to design and test new generations of combustors. As yet, the engineering models they need do not exist. But highly detailed computational simulations are paving the way for them by building the knowledge base needed to optimize and design combustion devices.

"The potential impact is huge in terms of improved fuel efficiencies of devices," says Jackie Chen of Sandia National Laboratories, leader of the combustion research team using the resources of the CCS. For example, she points out, a 50% increase in automobile fuel efficiency due to advanced engine designs could translate into 3 million barrels of oil saved per day, or a 21% reduction in oil used for transportation.

To develop useful combustion models, researchers must account for a mindbending array of parameters that affect the combustion process. The length scales involved, for example, range from the molecular scale where chemical reactions take place to combustors measuring several cubic meters. Similarly, time scales range from nanoseconds to hours. Hundreds of chemical species and reactions are involved, as well as dozens of other variables. The only approach feasible for such complex problems is to simulate directly a smaller range of scales and use the data to generate models that can be used in higher-level simulations.

Chen's team is focusing its work on one of the fundamental issues that must be understood for model development—the behavior of flames in a turbulent environment. Better understanding of the details of the turbulent combustion process is needed, including chemistry–turbulence interactions that affect the efficiency and emissions characteristics of devices. Data at this level of precision and completeness are far beyond the capabilities of physical experiments; massively parallel simulation is the only way to obtain them.

The operating conditions inside combustion devices are too extreme and complex to characterize fully by experimentation (e.g., high temperatures and pressure, and scores of chemical reactions), says Chen. Multiple underlying coupled processes occur in a combustion chamber, including turbulent mixing, spray evaporation, autoignition, flame combustion, and emissions generation. So computation, in conjunction with experiment, is the only way to fully characterize the conditions and to understand the intricate coupling. During FY 2005, the combustion research team used the Phoenix and Jaguar computers at CCS to conduct

Dynamics in a turbulent jet flame

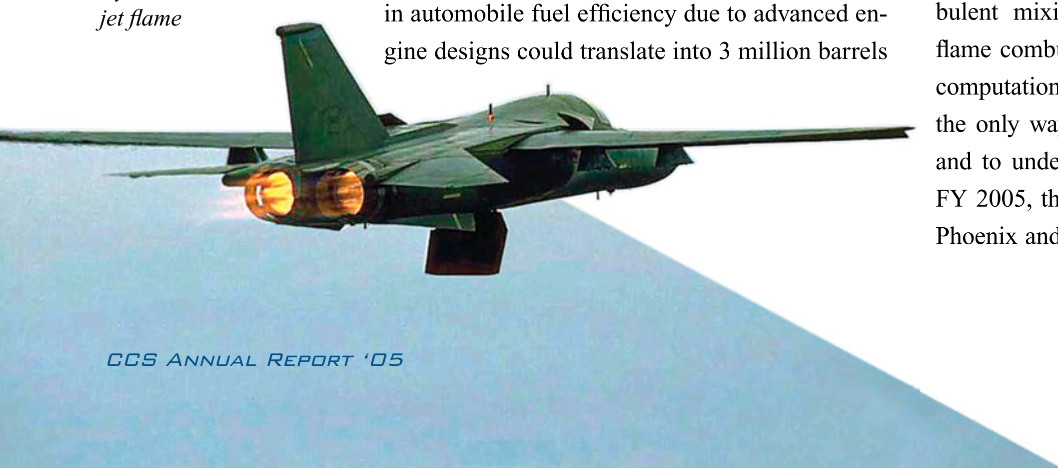

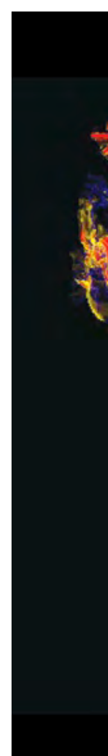

Researchers perform detailed simulations of 3D turbulence and complex chemistry

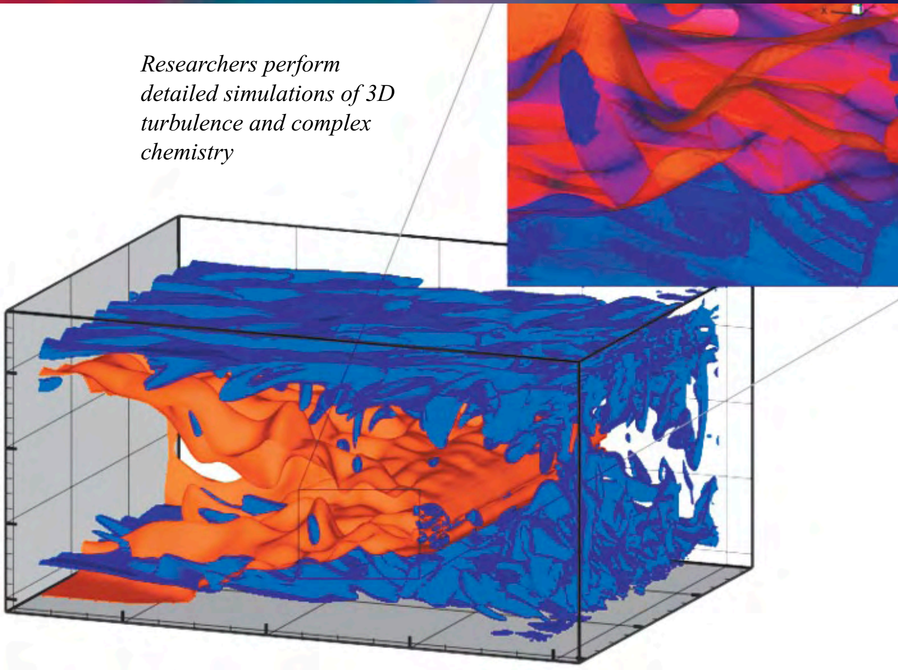

the first high-resolution, 3-dimensional (3D) direct numerical simulation (DNS) of non-premixed flames with detailed chemistry. DNS differs from the usual numerical simulation in that the turbulence is fully resolved numerically, rather than described by a model.

The calculations used approximately 300,000 computational hours on Jaguar and 500,000 on Phoenix in a series of runs. The largest run on Phoenix required 4 weeks of computational time on half of its 1024 processors, and the largest Jaguar run required 1 week on 40% of its 5212 processors. The DNS code S3D, used extensively in analyzing turbulence–chemistry interactions, was the simulation software employed by Chen and her collaborators.

The tremendous resources provided by CCS are essential to this work because DNS of turbulent combustion requires a "huge" number of grid points and scales, Chen says. The simulation must run for long periods to reach a point where the statistics extracted are stationary for model development and validation, and many variables must be included to represent even the simplest hydrocarbon fuels. "Without LCF (a leadership computing facility) we would be simulating incomplete physics: either 2D unsteady flows with detailed chemistry, or 3D turbulent flow with global 1-step chemistry. LCF enables us to include both 3D turbulence and complex chemistry and therefore lets us study turbulence-chemistry interactions in combustion directly."

As "early users" in 2005, combustion researchers performed DNS of turbulent non-premixed CO/H_2 jet flames to study extinction and re-ignition. Rapid mixing of fuel and air in the combustion chamber promotes efficient, clean combustion. However, if the mixing rates are so rapid that the chemical reactions cannot keep up, portions of the flame may be extinguished, leading to reduced efficiency and higher emissions. These simulations provide a better understanding of these fundamental processes and high-fidelity numerical benchmark data for model validation.

In separate runs, the research team also simulated turbulent lean methane–air Bunsen flames to better understand how intense turbulence can affect flame structure and propagation. Lean combustion is important in gas turbines used for stationary power generation because it promotes high thermal efficiency and low emissions of nitrogen oxides due

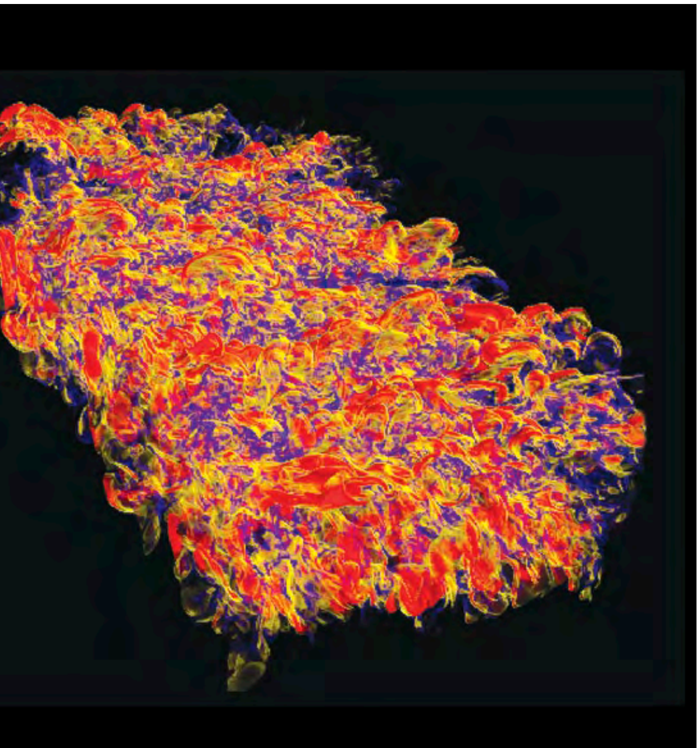

Simulated planar jet flame, colored by the rate of molecular mixing

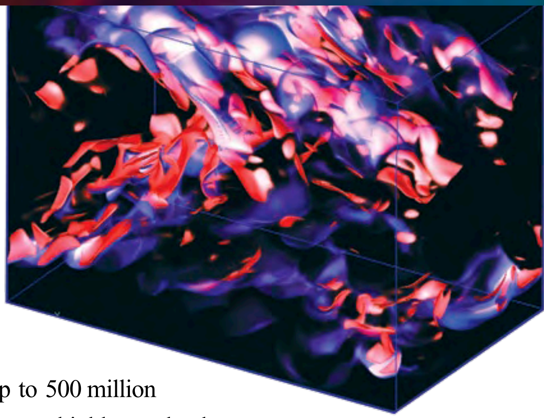

to lower flame temperatures. However, combustion at lean flammability limit conditions risks local extinction, emissions of unburned hydrocarbons, and large-amplitude oscillations in pressure that can result in poor combustion efficiency, toxic emissions of CO and unburned hydrocarbons, and even mechanical damage to the turbo-machinery used in power production. A fundamental understanding of the dynamics of premixed flame propagation and structure at this limit is required to advance predictive models.

Direct numerical simulation of a methane-air turbulent Bunsen flame

Fanning The Flames

Understanding the mechanisms governing turbulent mixing and flame extinction and re-ignition in a turbulent environment is key to developing predictive combustion models. In many common combustors, fuel and air are injected separately into the combustion chamber rather than being premixed. Fundamental questions about the combustion process revolve around the rate at which the fuel and air mix in the chamber. Rapid mixing produces rapid energy release, allowing the use of smaller combustion chambers and reducing emissions. However, above a critical level, rapid mixing and the associated turbulence can extinguish combustion in areas of the flame or even destabilize the entire flame. Extinguished fuel–air pockets that fail to reignite quickly are exhausted from the combustor, and abundant extinguished pockets that do not re-ignite can halt combustion altogether. Thus extinction adversely affects energy efficiency, emissions, and safety.

The CO/H_2 flame simulations were performed on Jaguar and Phoenix with up to 500 million grid points, the most highly resolved simulations ever conducted, and generated 30 TB of raw data. The methane–air flame simulations were conducted on Phoenix. The data are being analyzed to gain insights into how turbulent air–fuel mixing interacts with chemical reactions, and the dynamics of flame extinction and re-ignition. The data eventually will be made available to an international community of researchers working to advance basic understanding of combustion processes.

For FY 2006, the combustion research team was awarded 3.6 million computational hours on Jaguar and Phoenix to perform further simulations of turbulent combustion and to study flame stabilization mechanisms. The goal is a more complete understanding of several combustion phenomena: in addition to continuing the study of flame extinction and re-ignition, the team also is considering flame stabilization, soot formation, flame propagation, and auto-ignition.

Direct simulation of actual operating combustors is far beyond the capability of even the largest, fastest terascale computers or the petascale machines that will follow them. Rather, the knowledge gained will guide theory and experimentation in the field and build the science base. A step at a time, computational combustion researchers are moving toward predictive models that will enable engineers to design combustion devices that are more efficient, environmentally friendly, and safe.

Better Materials through Computation

Human history is measured literally in the advent of new materials: Stone Age, Bronze Age, Iron Age—the names memorialize simple but civilization-changing additions to the materials portfolio, separated by thousands of years. In our time, conversely, new materials with painstakingly engineered properties appear by the dozens every year, many of them filling such specialized niches that few people ever hear of them.

Continued progress in creating new materials is the key to advancement in every technological field. Thomas Schulthess and his collaborators are performing high-power computational simulations at CCS to aid the development of classes of advanced materials that will affect daily life in significant ways, even if they never become household words. Their simulations of high-temperature superconductors (HTSCs) already have answered key questions about materials that have the potential to revolutionize our electrical power system, and they set the stage for the development of better superconductors that will help make widespread use of this technology a reality.

Superconducting materials do not resist the flow of electrical current as other conductors do. Low-temperature superconductors (which superconduct at temperatures close to absolute zero) have limited utility because keeping them cooled consumes so much power. HTSCs, though, conduct electricity without loss at temperatures obtainable at much lower cost.

HTSCs are crucial to improving energy efficiency because they eliminate distribution losses in power cables and transformers. A significant portion of the electricity distributed through the existing power grid is lost because of resistance, and much of the complexity in the system results from methods to keep these losses small. The availability of superconducting materials operating at room temperature could therefore greatly enhance the efficiency and the stability of the power distribution system. Efforts have been under way for years to develop practical HTSCs for this and other applications. To support this work, computer models are being used to explore the fundamental nature of superconductors.

The model most widely used to study the physics of HTSCs is the Hubbard model. One of the main challenges in superconductivity research was determining the mechanism that underlies high-temperature superconductivity. Schulthess and his collaborators—Thomas Maier at ORNL, Paul Kent of the University of Tennessee, and Mark Jarrell of the University of Cincinnati—used CCS's Phoenix supercomputer in FY 2005 to achieve the first credible solution of the Hubbard model. The results show that the electron pairing responsible for superconductivity in HTSCs (called "Cooper pairing") can result from strong electronic correlations. This achievement—which settles a key physics question about HTSCs—was made possible by the fast vector processors and high memory

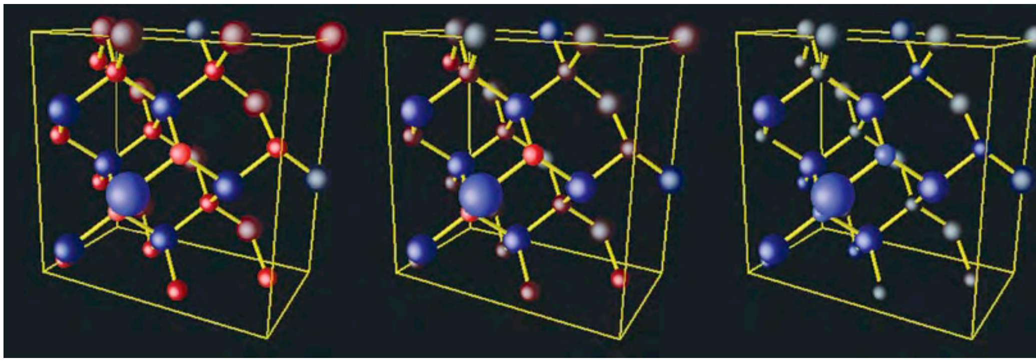

Magnesium (large blue sphere) in three different semiconductor materials (from left), GaAS, GaP, and GaN bandwidth of Phoenix and a new algorithm that takes a different approach from previous methods.

There are two logical next steps that build on the work done in 2005, said Schulthess. First, now that a reliable model exists for superconductivity, the team will use the resources of CCS in 2006 to study the mechanism that leads electrons to pair into Cooper pairs. "The question of understanding the pairing mechanism is one of the biggest problems in physics," he said.

Second, the researchers will try to relate the model directly to specific materials and connect it to first-principles calculations. "That's where we are saying we have copper and oxygen and lanthanum—specific elements—in the materials. The Hubbard model is very generic. But now the big question is how to connect the first-principles calculations to the model. When that has been achieved—and this is probably a multi-year project—I think that's when you can claim a major breakthrough. Once you establish this connection between the model and specific materials, you're in the business of designing materials, making better superconductors."

Smaller Is Better

High-density data storage relies on altering the magnetic state of a material. Iron-platinum (FePt) materials have exceptionally high magnetocrystalline anisotropy energy (MAE), meaning once the magnetic state is changed, a great deal of energy is required to reverse it. In low-MAE materials, temperature increases can cause random scrambling of previously set states.

Processes are being developed to create nanoparticles of uniform size that can be engineered to dramatically increase their capacity to store information compared with bulk materials. A major challenge is designing nanoparticles for high MAE, since MAE is proportional to particle volume. Modeling is helping unravel the relationship between chemical, structural, and magnetic properties of these promising materials.

It is possible to synthesize FePt nanoparticles with sufficient MAE for use in magnetic storage media that store tens of terabits per square inch. The magnetization at room temperature would be stable, and information stored in a particle could be retained for the lifetime of the storage medium. But there is a roadblock: in such tiny particles, the magnetization can't be reversed with conventional writing techniques. Information can be stored, but it can't be written! So new methods are needed to switch the magnetic moment in these nanoparticles with a nanometer-sized write head at rates appropriate for use in hard drives. Researchers at CCS are using the electronic structure code LSMS and an extension of the Wang-Landau algorithm to simulate magnetic systems with several thousand atoms at non-zero temperatures.

The simulations have achieved over 81% of theoretical peak performance for particles containing up to 2662 atoms. They have shown variation in moment size and orientation as a function of position within a nanoparticle and nanoparticle size and composition.

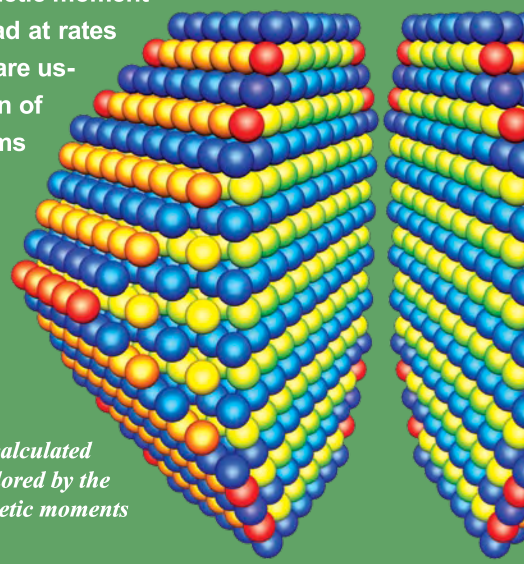

Magnetic structure of FePt nanoparticle calculated using CCS computers. Atoms are colored by the calculated magnetic moments

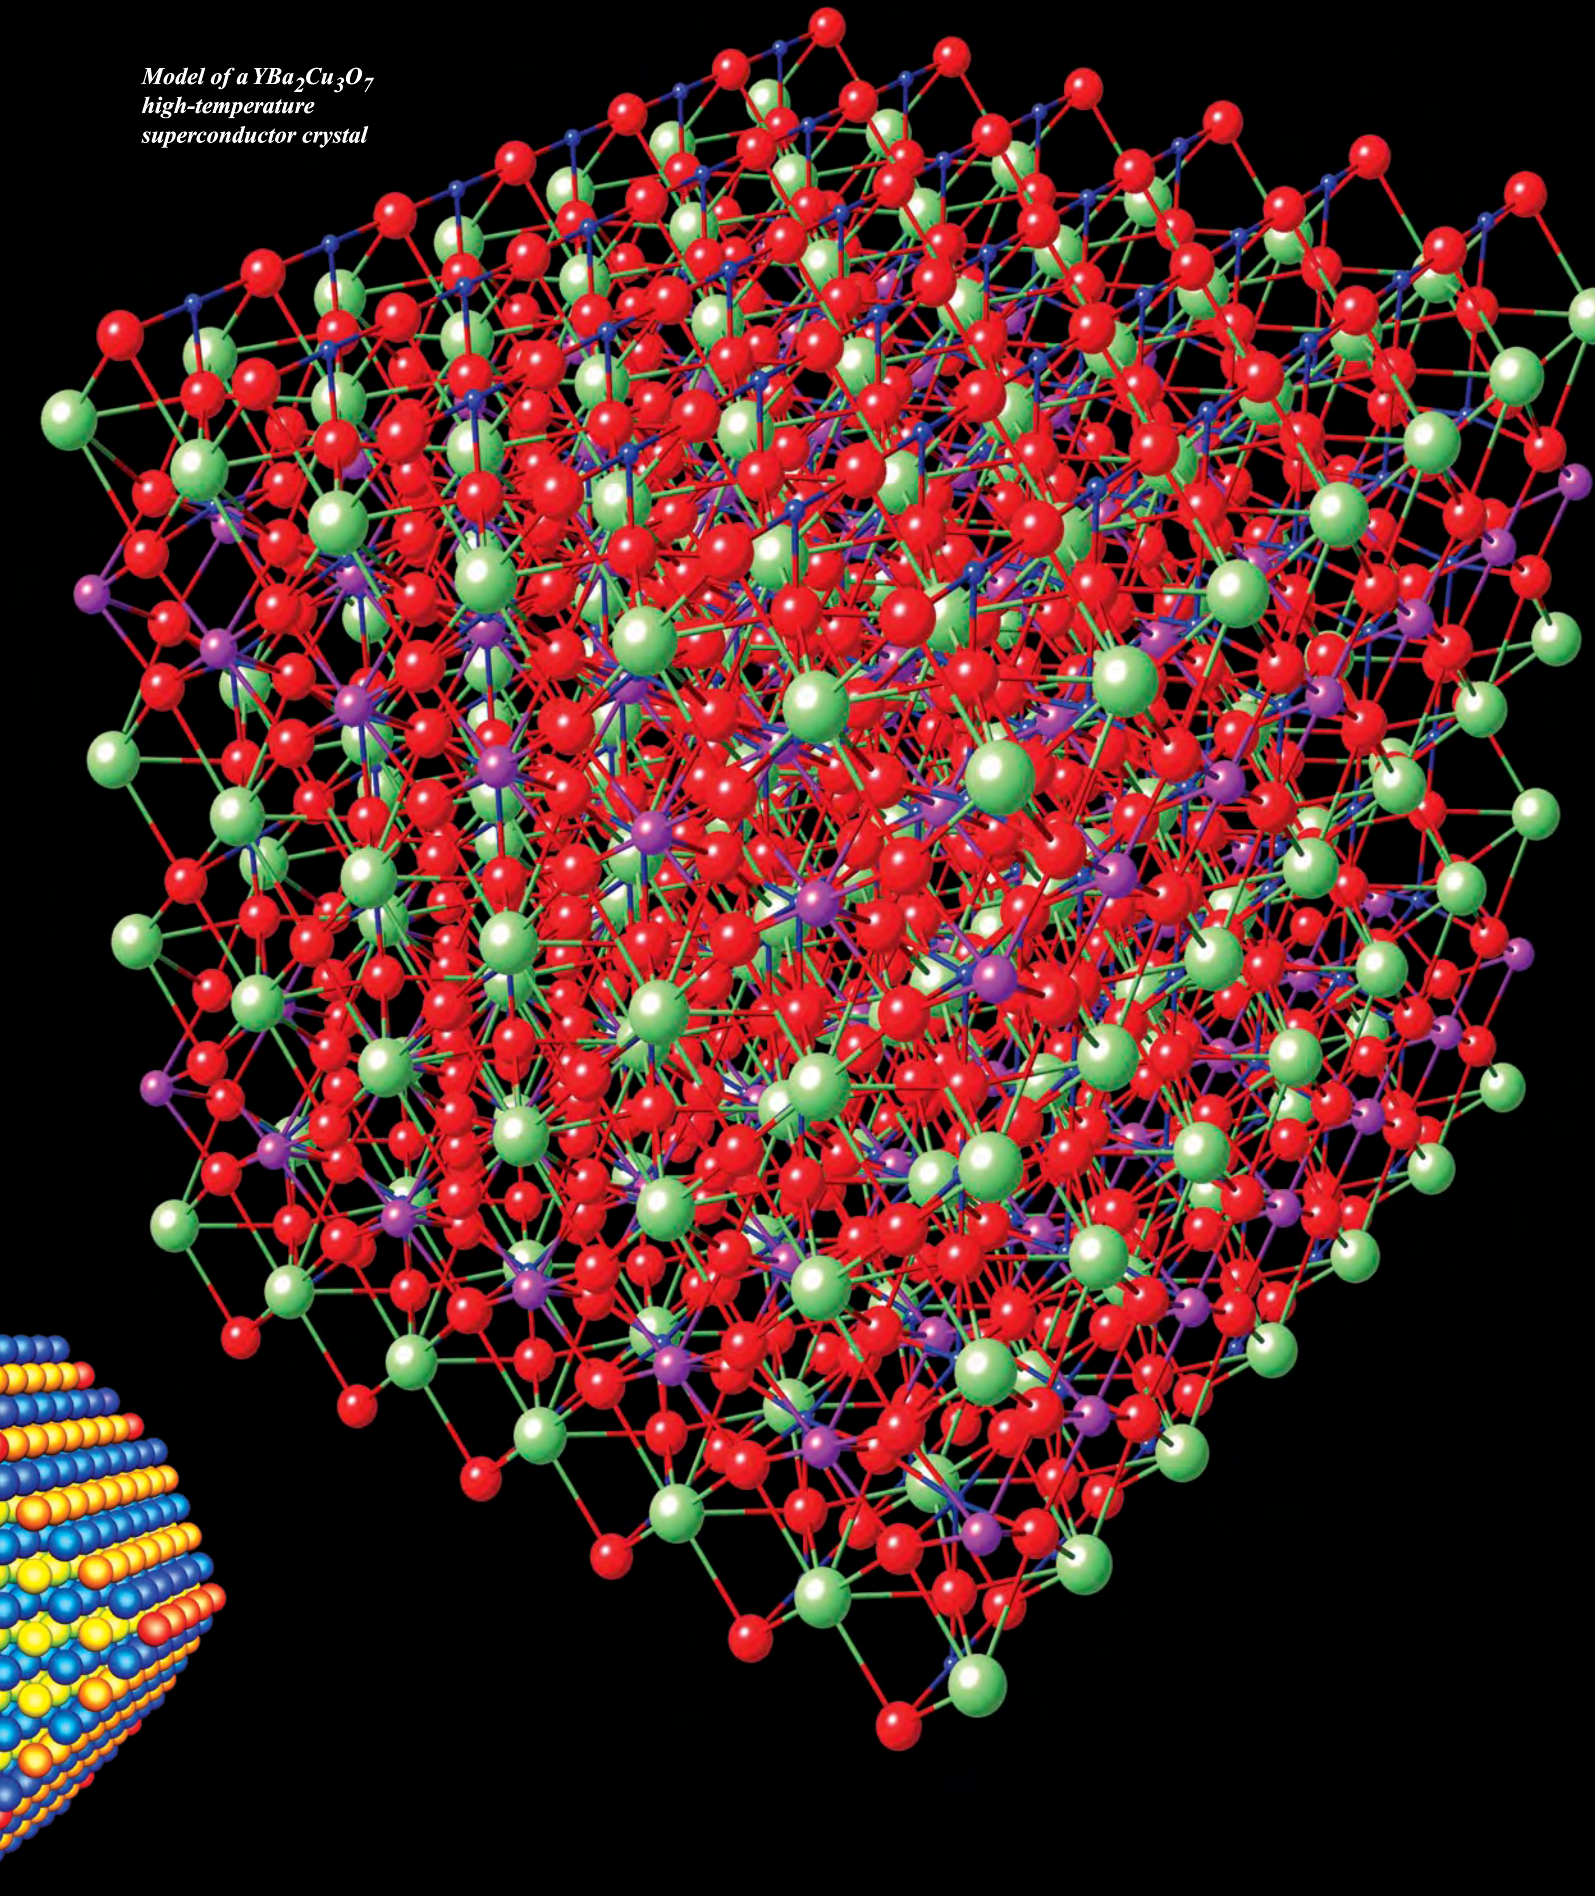

Model of a $YBa_2Cu_3O_7$ high-temperature superconductor crystal

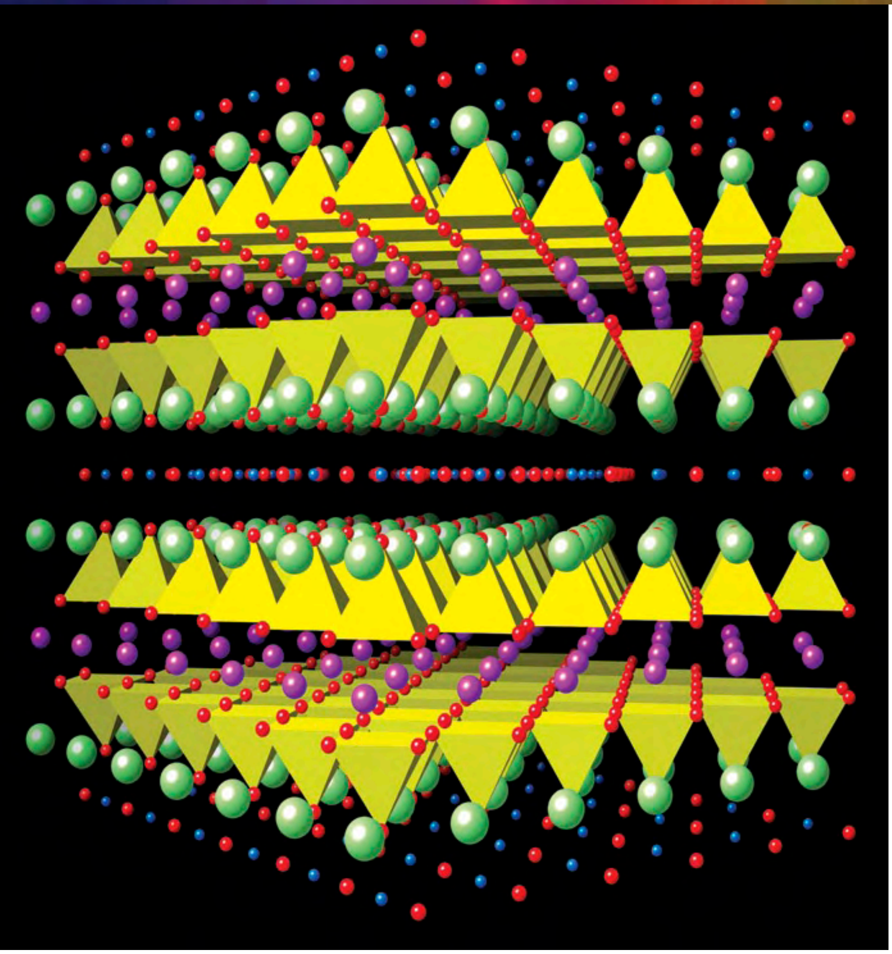

Superconducting materials conduct electricity without loss caused by resistance

In addition to the HTSC research, the team will be conducting two other major projects at CCS in 2006. Gonzalo Alvarez of ORNL and Elbio Dagotto of the University of Tennessee are working on a "spintronics" project that uses spin fermion models to simulate magnetic semiconductors and colossally magnetoresistive oxides, both of which have potential for use in future generations of magnetic data storage and other information technologies.

"They look very promising from a physics point of view, but the progress with these materials has been slow because of a lack of theoretical understanding," Schulthess said. "We hope we can make a change here. The models have been around for a long time, but it has not been possible to solve them accurately. Jaguar is making a difference because of its sheer power. Having a factor of 100 or 1000 more power gives us the ability to solve these in relevant parameter ranges." The calculations can run in a week or two on Jaguar, compared with a year or two on smaller computers. "If you have to run them for 2 years, you just can't do the work."

The researchers are hoping to understand the colossal magnetoresistive effect in a realistic model. Like the Hubbard model, the spin fermion model is inspired by real materials. If it can be solved with effects measured in real materials, it can be used to design new materials. "What you can do now that you could not do without these machines is solve it for realistic models that include the effects of chemical disorder," Schulthess said. "Experimentally, it is known that chemical disorder is important, but nobody previously was able to incorporate this in a calculation. We can now do these computational experiments with very realistic systems."

A third materials project will use Jaguar to simulate iron–platinum alloys, which hold promise for magnetic recording applications. Part of this project could develop into a petascale computing problem. It is expected to run at 50–90% of peak and to scale to 200,000 processors.

The ongoing computational research will directly impact the discovery and design of new materials. The combination of these theoretical and computational capabilities with new synthesis techniques developed at the Center for Nanophase Materials Sciences (CNMS), the world's highest-resolution electron microscopes, and the Spallation Neutron Source is a powerful one that will move materials science forward, Schulthess said.

The CCS computing capability is not isolated; it is part of a bigger ORNL set of tools, he noted. "With computing we don't discover. With computing, we only make predictions. It's when your predictions are verified experimentally that you call it a discovery. That's why it's so important that our models are connected to real materials. What happens in the real world is what counts in science."

Seeking Revelations about Space and Time

The fundamental nature of matter can be exposed by studying the behavior of subatomic particles, revealing the very makeup of space and time. Such matter can be studied by smashing together subatomic particles at very high energies. Doing so requires large, sophisticated, and precise instruments. One such device is the proposed $10-billion+ International Linear Collider (ILC)—the highest-priority future accelerator project in high-energy physics.

A group of scientists from Europe, Asia, and North America is designing the ILC to create high-energy particle collisions between electrons and positrons that would reach the scale of a TeV—tera electron volt, or one trillion electron volts—and would open many new possibilities for discovery. (By comparison, molecular bonds are on the order of a just a few electron volts.) This energy level can be realized only if the tens of thousands of accelerator cavities can maintain a stable high-energy beam. As the beam travels the many-kilometer length of the accelerator, wakes—not unlike those trailing ships in the ocean—can be generated and must be controlled to minimize loss of beam stability.

A low-loss accelerator cavity design is being evaluated as an alternative to the standard design, because it has a lower operating cost while delivering a higher performance. Designing such a cavity is the goal of the simulation project collaboration between the Stanford Linear Accelerator Center (SLAC) and the Center for Computational Sciences (CCS).

Second model of KEK L-band low-loss 9-cell (ICHIRO) cavity (Courtesy of KEK, Japan)

Computer rendering of electronic fields inside a super-conducting accelerator cavity (Courtesy of DESY, Germany)

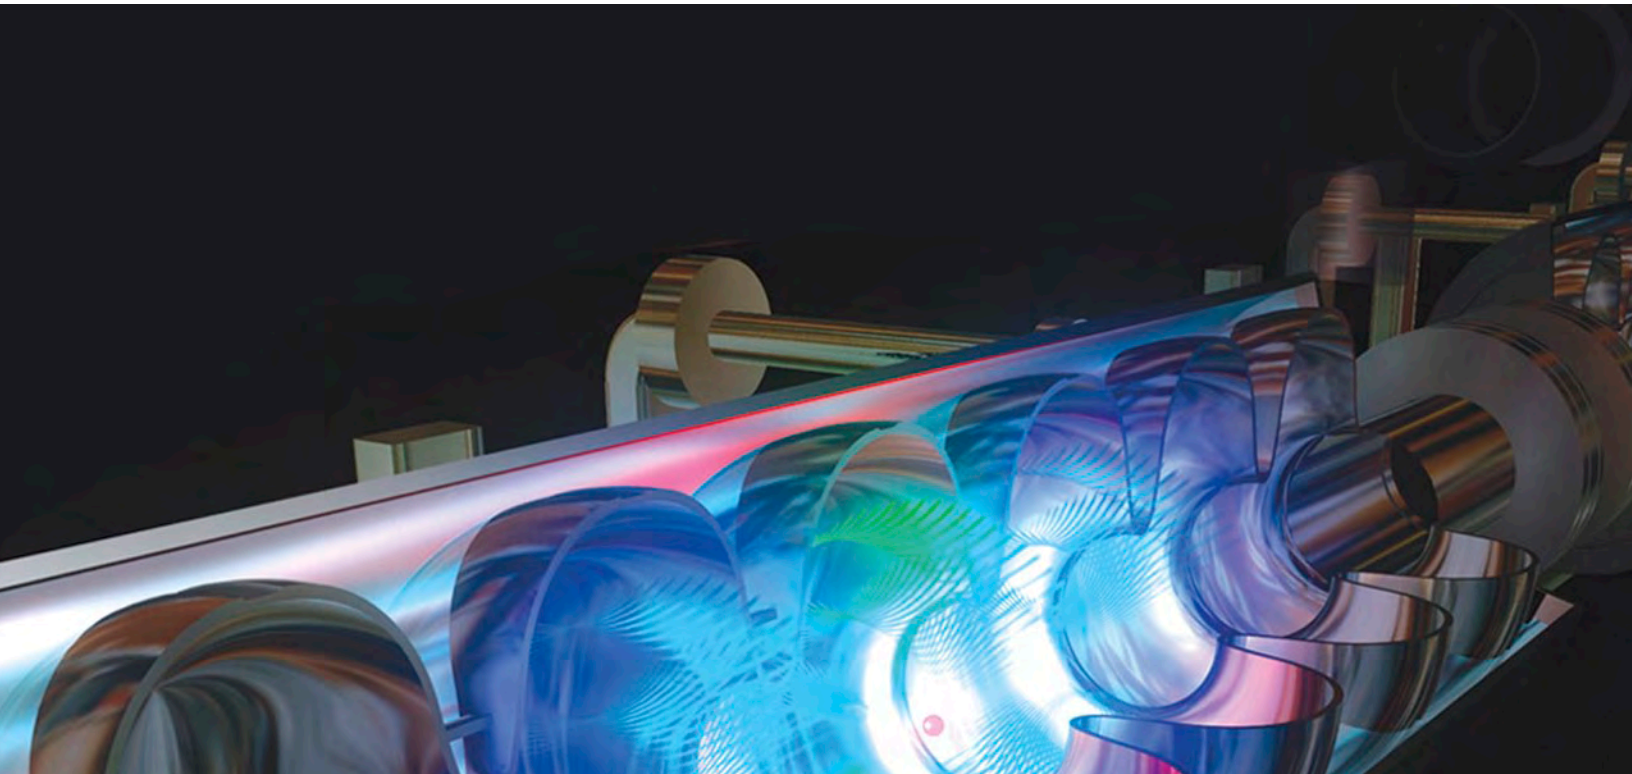

Kwok Ko of SLAC leads the effort to conduct CCS computer simulations of wakefield suppression in a new low-loss cavity design with less energy dissipation that will help provide input for determining the ILC baseline design. The goal is to optimize the shape of the cavity so that the disruptive wakefields generated by the accelerating beam are suppressed below an acceptable level without compromising the cavity's performance.

Collaboration with researchers in the DOE SciDAC program's Integrated Software Infrastructure Center and Scientific Application Pilot Program has led to significant advances in numerical models and simulation capabilities used by Ko and his collaborators. Effective use of these complex 3-dimensional models and advanced software requires computer simulation on the scale of that found at the CCS. In particular, the large memory and faster processing units available on Phoenix, the world's largest open-access Cray X1E, enable researchers to carry out many large-scale computation experiments required for the optimization of the ILC cavity design. Physical experiments of this nature and scale would be cost- and time-prohibitive.

"Utilization of the CCS facilities is crucial to our work," Ko says. "Simulation currently provides the only means of studying and understanding wakefields in the low-loss cavity, because the prototypes are still under fabrication. We also are greatly appreciative of the support we have received from many CCS staff members during this project."

Ko says SLAC scientists are using the new software to aid in the design of the low-loss cavity for the ILC. "The most critical component of the ILC is the accelerating cavity, which imparts energy to the beam and constitutes a significant fraction of the machine cost," Ko explains.

This work is also in collaboration with the KEK (National Laboratory for High Energy Physics) in Japan, DESY (Deutsches Elektronen-Synchrotron) in Germany, TJNAF (Thomas Jefferson National Accelerator Facility), and FNAL (Fermi National Accelerator Laboratory).

Ko and his team—along with other scientists in our nation and abroad—will continue their efforts to provide solid scientific research that will impact the way the ILC is designed and constructed so that, ultimately, the collision of the smallest imaginable particles will provide unparalleled insight into our understanding of the universe.

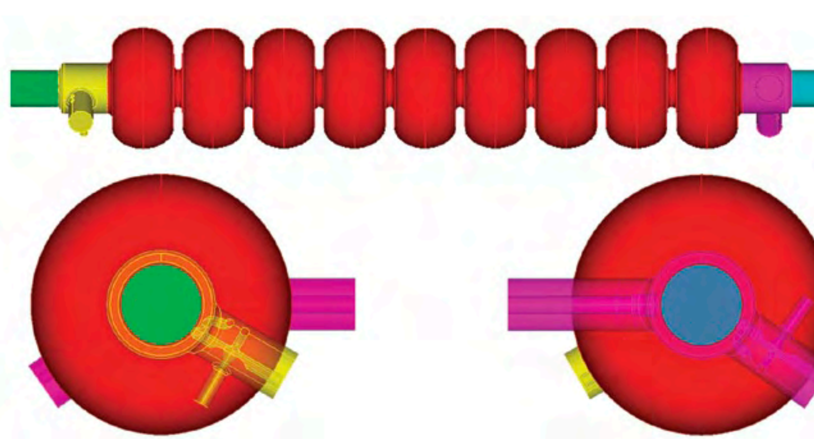

Researchers are using 3-dimensional electromagnetic modeling on Jaguar to design a new low-loss accelerating cavity for the ILC (Courtesy of SLAC)

Tracking Plasma Turbulence

As worldwide demand for energy accelerates, the need for energy sources other than fossil fuels is becoming acute. Since the 1950s, researchers have been exploring the possibility of harnessing nuclear fusion, the process that releases the immense energy of the Sun and other stars, as a terrestrial energy source. An agreement in recent years among several nations to build the International Thermonuclear Experimental Reactor (ITER), an experimental magnetic fusion reactor, was a big step toward demonstrating the feasibility of fusion energy.

Fusion occurs when two isotopes of hydrogen nuclei, deuterium and tritium, fuse with each other at extremely high temperature to form an isotope of helium, releasing a great amount of energy in the process. (Fusing a kilogram of hydrogen releases as much energy as burning 10 million kg of coal.) Igniting a sustained fusion reaction also requires an enormous input of energy: the planned operating temperature of ITER is 100 million degrees centigrade. At such temperatures, the atoms form a plasma—an ionized gas of charged atomic particles. In magnetic fusion, the dominant approach used, this high-energy plasma is confined under intense pressure by surrounding the cloud of charged particles with a powerful magnetic field. At present, the most promising type of apparatus for producing fusion power is the tokomak, a donut-shaped magnetic chamber.

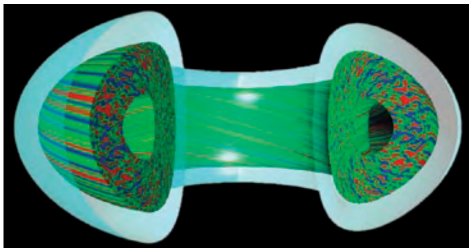

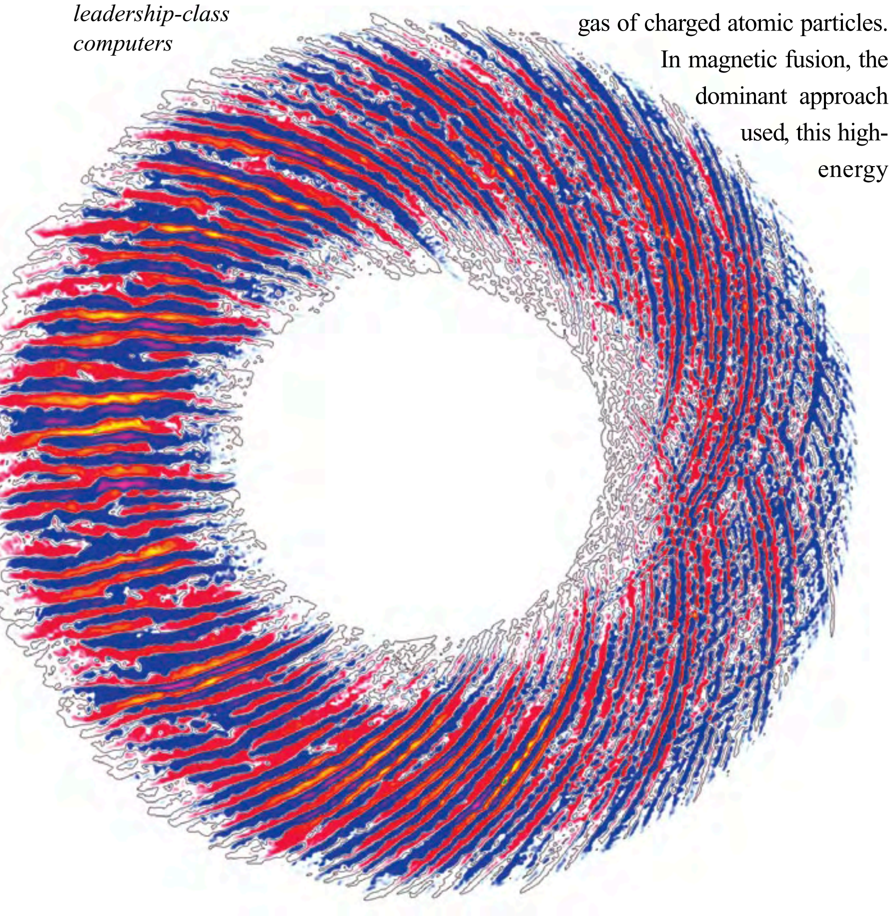

Simulation of plasma microturbulence requires leadership-class computers

Although fusion has been successfully demonstrated on a small scale, reactor technology is in its infancy and fundamental questions still must be resolved. To conduct experiments successfully in ITER, scientists must have far more sophisticated fusion simulation tools than the ones available now. They will provide the data needed to develop models that can predict how plasma will behave under the conditions that will exist inside the reactor. A race is on to develop those tools by the time ITER comes on line (projected for 2015).

A fusion research team led by Wei-li Lee of the Princeton Plasma Physics Laboratory is using the Phoenix and Jaguar supercomputers at CCS to simulate plasma turbulent fluctuations that cause particles and energy to travel from the center of the plasma and flow toward the edge. Turbulence causes the plasma to lose the heat that is essential to maintaining the fusion reaction, so it must be controlled to enable a fusion reactor to operate successfully. Reliable modeling of turbulence processes is an indispensable step toward formulating control strategies.

The project studies the flow of charged particles in a plasma and the associated evolution of turbulence over an extended period, and shows what is happening to particles as turbulence occurs. The Gyrokinetic Toroidal Code (GTC), a particle-in-cell code for simulating complex microrturbulence properties in fusion-grade plasmas, is the principal code being used for these calculations.

GTC has achieved a sustained 3.8 teraflops on 97% of Jaguar's 5212 processors and 2.2 teraflops on about 94% of Phoenix's 1024 processors. The largest calculation was 28 billion particles on 4800 of Jaguar's processors. A typical run takes about 80 to 100 hours.

Data from the CCS simulations are contributing to advances in understanding the degradation of the confinement of energy and particles in fusion plasmas caused by turbulence associated with small-scale plasma instabilities driven by plasma pressure gradients. However, the detailed physics of how these instabilities grow and reach their upper limits, how they impact plasma confinement, and how turbulence might be controlled remain major scientific challenges. Eventually, the researchers plan to include more physics in GTC so they can actually simulate ITER-like plasmas in terms of complexity and physical size. Plans for work in 2006 include using Jaguar to conduct simulations with increased fidelity for "shaped plasma" in which the geometry is closer to that of an actual experimental tokamak such as ITER.

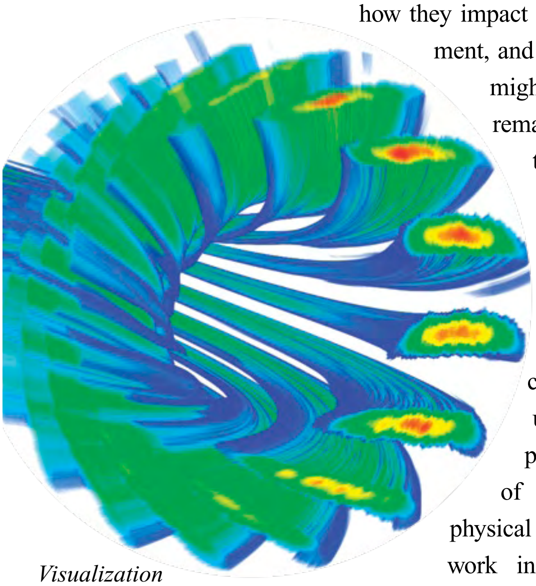

Visualization of turbulent eddies in plasma flow

As they refine their understanding of how turbulence operates, researchers will be able to see which reactor scenarios promote stability and which ones increase turbulence. Once models can reliably identify situations that are favorable to efficient reactor operation, engineers can use the resulting data to determine how to design equipment to create those scenarios.

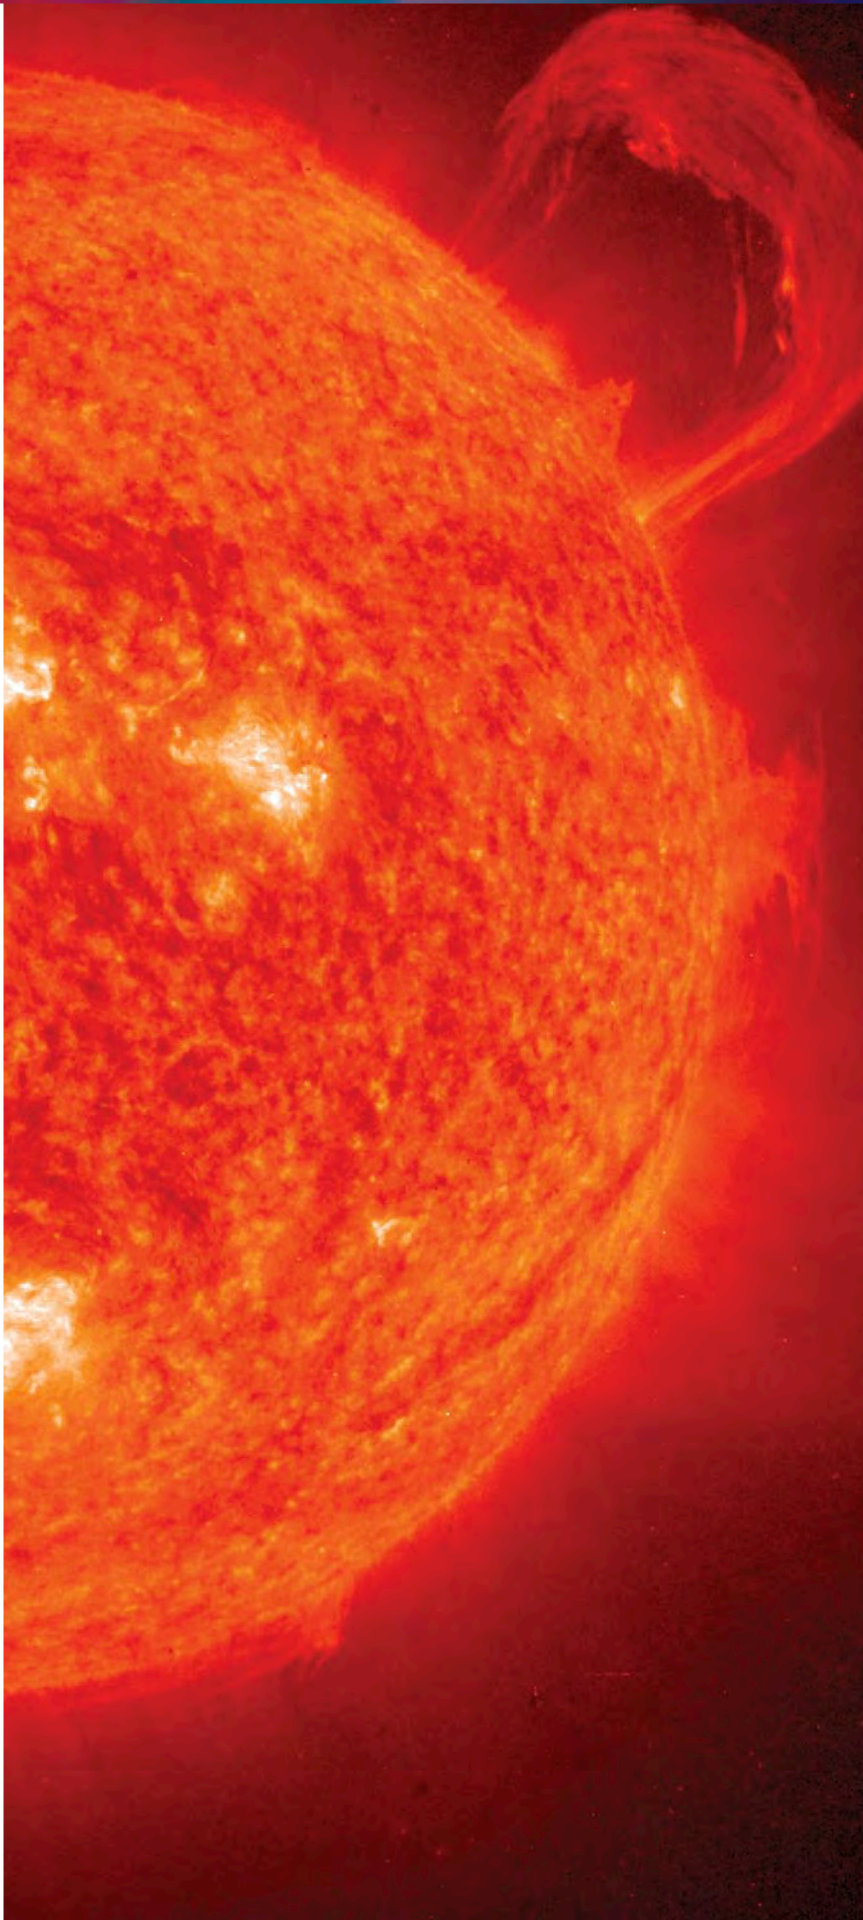

Courtesy of NASA

Benchmarking Small Molecules

Chemists spend their lives studying the structure and interactions of molecules because understanding molecules, and the smaller atoms that make them up, is essential to deciphering a wide range of phenomena, from the fate of contaminants in the environment to the treatment of genetic diseases. Thanks to advances in theoretical, computational, and experimental capabilities, chemists are now able to characterize matter at increasingly detailed atomic and molecular levels. Computational modeling is an essential part of this endeavor.

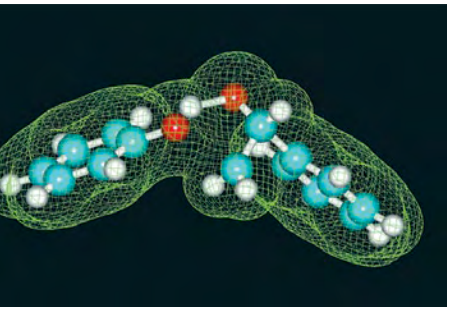

Results of electron interaction are represented here by the green mesh surrounding constituent atoms in this molecule

Atoms combine into more complex molecular systems, ranging from simple two-atom clusters like the sodium and chloride combination in common table salt to the intricate, winding-patterned DNA that holds the genetic code. The way atoms bind to form molecules depends in large part on how the electrons within the atoms interact. The interactions of these subatomic particles, referred to as "electronic structure," determine how the larger molecule behaves. Configuration Interaction (CI) is a method for modeling the electronic structure of both atoms and molecules.

Full CI, the most accurate quantum chemistry method, can provide results that allow researchers to distinguish between various sources of error in approximate methods. This "benchmarking" exercise requires performing computationally intensive CI calculations on small molecules and is essential in assessing the reliability and precision of calculations on large molecules. Full CI is especially important for benchmarking molecular excited states caused by adding energy to a molecular system.

Mathematically exact calculations provide benchmarks for small molecules and the foundation for more precise simulations of large systems

Because many calculations must be performed in the benchmarking process, it is essential to have a fast, efficient computer program. The new algorithm employed by Robert Harrison, joint University of Tennessee–ORNL faculty researcher, stores just two vectors of CI coefficients in memory; and, in the largest calculation on the CCS Phoenix to date, it included 65 billion coefficients in a vector.

The CCS computers also allow the use of larger one-electron basis sets (an important parameter for the accuracy of the calculations), which are more accurate and balanced than the small basis sets employed in previous studies. Researchers can now focus on systems previously not accessible to full CI benchmarking, notably molecules or electronic states with unpaired electrons (open-shell systems), which challenge conventional approximate methods. Open-shell systems are very important in chemical reactions, and excited states are central to topics such as photochemistry.

Large, fast computational power is needed to advance from approximate to exact models of molecules, especially for complex open-shell systems and excited states. These benchmark calculations will enable researchers to calibrate various approximate models that can then be used to study much larger molecules. When fully functional, the CCS Jaguar will enable calculations with 300 billion coefficients.

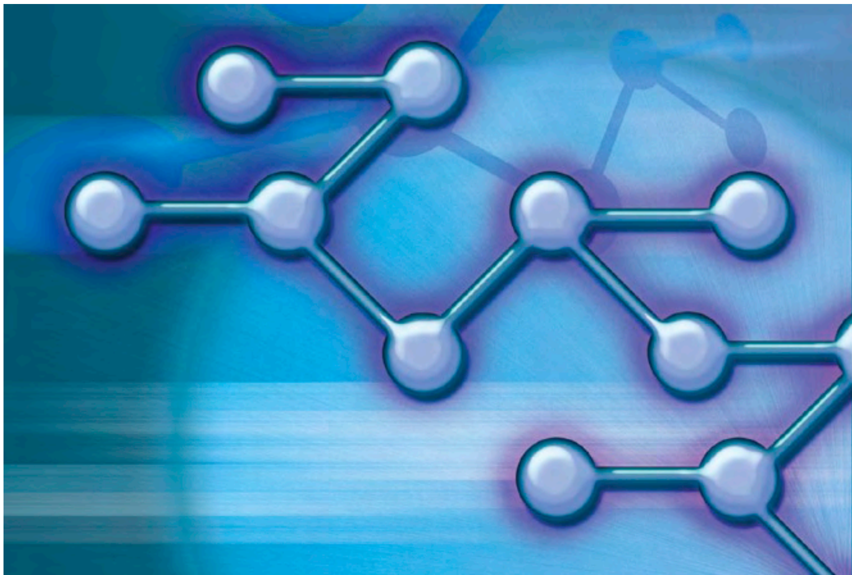

3

CCS Proposal Process

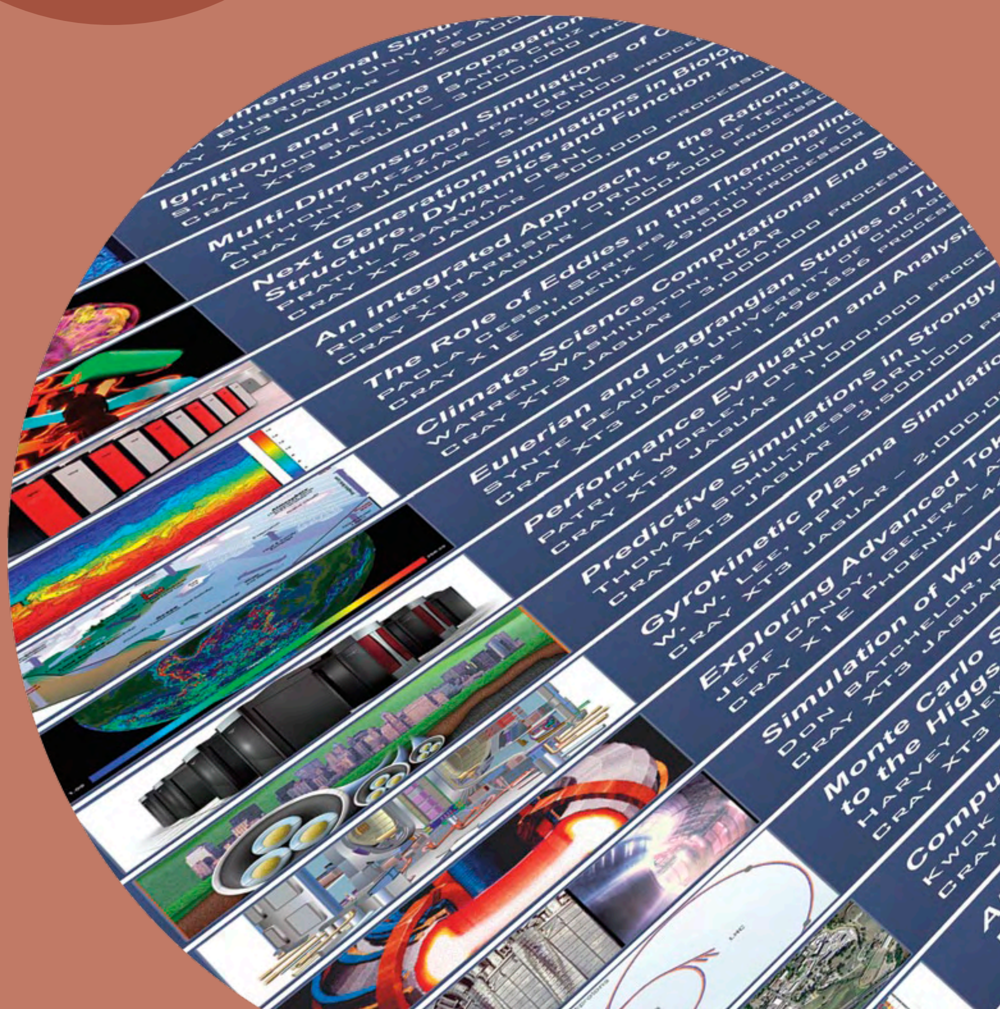

Proposing Leadership Computing Projects

CCS resources are dedicated to a few computationally intensive projects with the potential for breakthrough discoveries in research areas of great importance to the United States. The basis for selecting these projects is a call for proposals followed by a peer-reviewed selection process.

DOE issues an annual call for proposals for CCS, inviting researchers around the country to submit individual or team proposals. Proposals typically fall into three categories:

Research activities in porting, testing, and tuning applications on the CCS machines prepare a project for larger resource allocations. These test or "pilot" projects are usually short in duration and involve only a few tens of thousands of processors. The project team can work with members of CCS's Scientific Computing group to optimize their applications.

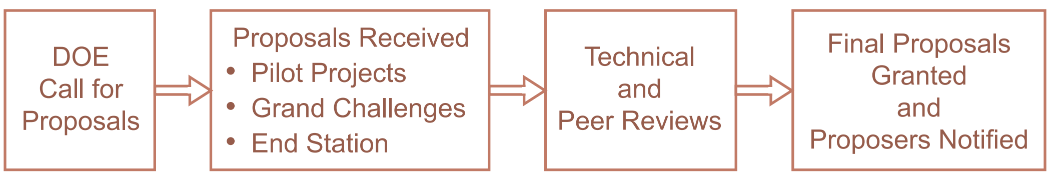

Large, "grand challenge" scale research efforts, usually multi-year in scope, have the potential to lead to scientific breakthroughs. These typically demand millions of computational hours. Examples are the fusion, combustion, astrophysics, catalysis, and accelerator simulations that began running on the Jaguar and Phoenix in 2005.

End station projects could be characterized as "grand challenge plus." They usually involve a large user community in a specific research domain. These projects are aimed at scientific discovery, but they also include a significant effort to add to the functionality and improve the performance of software and processes and develop next-generation supercomputing applications for the research community.

Researchers submit their proposals electronically. Proposals receive a technical review and a peer review. "We accepted 22 of 47 proposals for 2006," said Julia White of CCS User Assistance and Outreach. "Our investment of resources is in big science that uses large fractions of the center's capability."

In the proposal, researchers describe the project, its goals, and the theoretical and computational methods it uses; its place in the context of research on the topic; and how using the CCS computers will enable progress in the research. The proposal includes technical details such as the number of hours requested on the CCS computers, needs for data storage and transfer and visualization, methods to be applied in the simulation codes; and the development necessary to prepare the simulation to run on the CCS systems.

Proposers must show that their codes can scale effectively to a large fraction of the CCS machine's thousands of processors. "If a code can only run on 100 processors, it's probably not appropriate for the CCS," White said.

Use of CCS is not limited to DOE-funded research. Through the proposal process, researchers from academia, national laboratories and industry may all request access to CCS resources. A rigorous review identifies the top projects, and a select number are granted the extraordinarily large allocations necessary to achieve science goals otherwise unobtainable.

CCS Projects for 2006

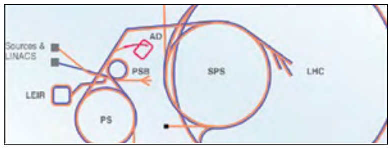

Monte Carlo Simulation and Reconstruction of CompHEP–Produced Hadronic Backgrounds to the Higgs Boson Diphoton Decay in Weak-Boson Fusion Production Mode
Harvey Newman (California Institute of Technology)
30,000 processor-hours: Jaguar

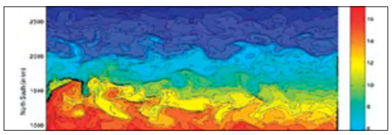

The Role of Eddies in the Thermohaline Circulation
Paola Cessi (Scripps Institution of Oceanography, University of California, San Diego)
29,000 processor-hours: Phoenix

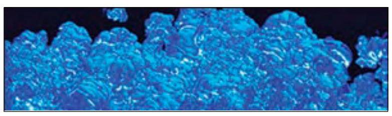

Ignition and Flame Propagation in Type Ia Supernovae
Stan Woosley (University of California, Santa Cruz)
3,000,000 processor-hours: Jaguar

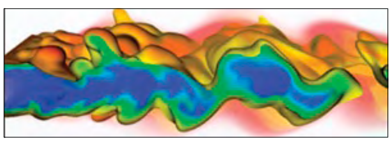

High-Fidelity Numerical Simulations of Turbulent Combustion—Fundamental Science Toward Predictive Models
Jackie Chen (Sandia National Laboratories)
3,000,000 processor-hours: Jaguar / 600,000 processor-hours: Phoenix

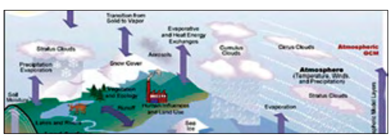

Climate-Science Computational End Station Development and Grand Challenge Team
Warren Washington (National Center for Atmospheric Research)
3,000,000 processor-hours: Jaguar / 2,000,000 processor-hours: Phoenix

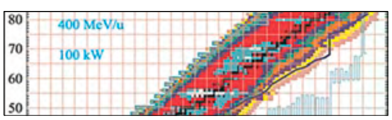

Ab-initio Nuclear Structure Computations
David J. Dean (Oak Ridge National Laboratory)
1,000,000 processor-hours: Jaguar

Performance Evaluation and Analysis Consortium (PEAC) End Station
Patrick H. Worley (Oak Ridge National Laboratory)
1,000,000 processor-hours: Jaguar / 200,000 processor-hours: Phoenix

Computational Design of the Low-loss Accelerating Cavity for the ILC
Kwok Ko (Stanford Linear Accelerator Center)
500,000 processor-hours: Phoenix

Gyrokinetic Plasma Simulation
W. W. Lee (Princeton Plasma Physics Laboratory)
2,000,000 processor-hours: Jaguar / 225,000 processor-hours: Phoenix

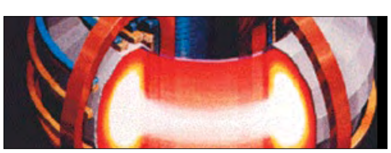

Exploring Advanced Tokamak Operating Regimes Using Comprehensive GYRO Gyrokinetic Simulations
Jeff Candy (General Atomics)
440,240 processor-hours: Phoenix

Multi-dimensional Simulations of Core-Collapse Supernovae
Anthony Mezzacappa (Oak Ridge National Laboratory)
3,550,000 processor-hours: Jaguar / 700,000 processor-hours: Phoenix

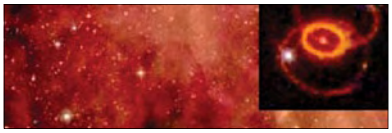

Multi-dimensional Simulations of Core-Collapse Supernovae
Adam Burrows (University of Arizona)
1,250,000 processor-hours: Jaguar

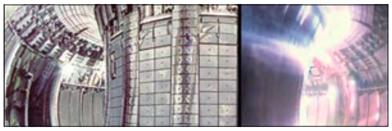

Simulation of Wave-Plasma Interaction and Extended MHD in Fusion Systems
D. B. Batchelor (Oak Ridge National Laboratory)
3,000,000 processor-hours: Jaguar

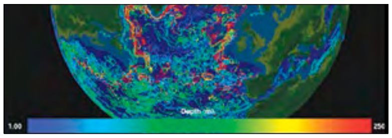

Eulerian and Lagrangian Studies of Turbulent Transport in the Global Ocean
Synte Peacock (University of Chicago)
1,496,856 processor-hours: Jaguar

An Integrated Approach to the Rational Design of Chemical Catalysts
Robert Harrison (Oak Ridge National Laboratory and University of Tennessee)
1,000,000 processor-hours: Jaguar / 300,000 processor-hours: Phoenix

Next Generation Simulations in Biology: Investigating Biomolecular Structure, Dynamics and Function through Multi-scale Modeling
Pratul K. Agarwal (Oak Ridge National Laboratory)
500,000 processor-hours: Jaguar

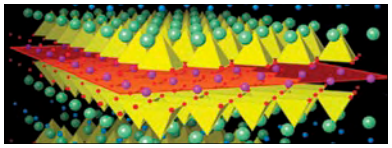

Predictive Simulations in Strongly Correlated Electron Systems and Functional Nanostructures
Thomas Schulthess (Oak Ridge National Laboratory)
3,500,000 processor hours: Jaguar / 300,000 processor-hours: Phoenix

INCITE Projects for 2006

Molecular Dynamics Simulation of Molecular Motors
Martin Karplus (Harvard University)
1,484,800 processor-hours: Jaguar

Real-Time Ray Tracing
Evan Smyth (Dreamworks)
950,000 processor-hours: Jaguar

Development and Correlations of Large-Scale Computational Tools for Flight Vehicles
Moeljo Hong (The Boeing Company)
200,000 processor-hours: Phoenix

Direct Numerical Simulation of Fracture, Fragmentation, and Localization in Brittle and Ductile Materials
Michael Ortiz (California Institute of Technology)
500,000 processor-hours: Jaguar

Interaction of ETG and ITG/TEM Gyrokinetic Turbulence
Ronald Waltz (General Atomics)
400,000 processor-hours: Phoenix

4

Operations Perspectives

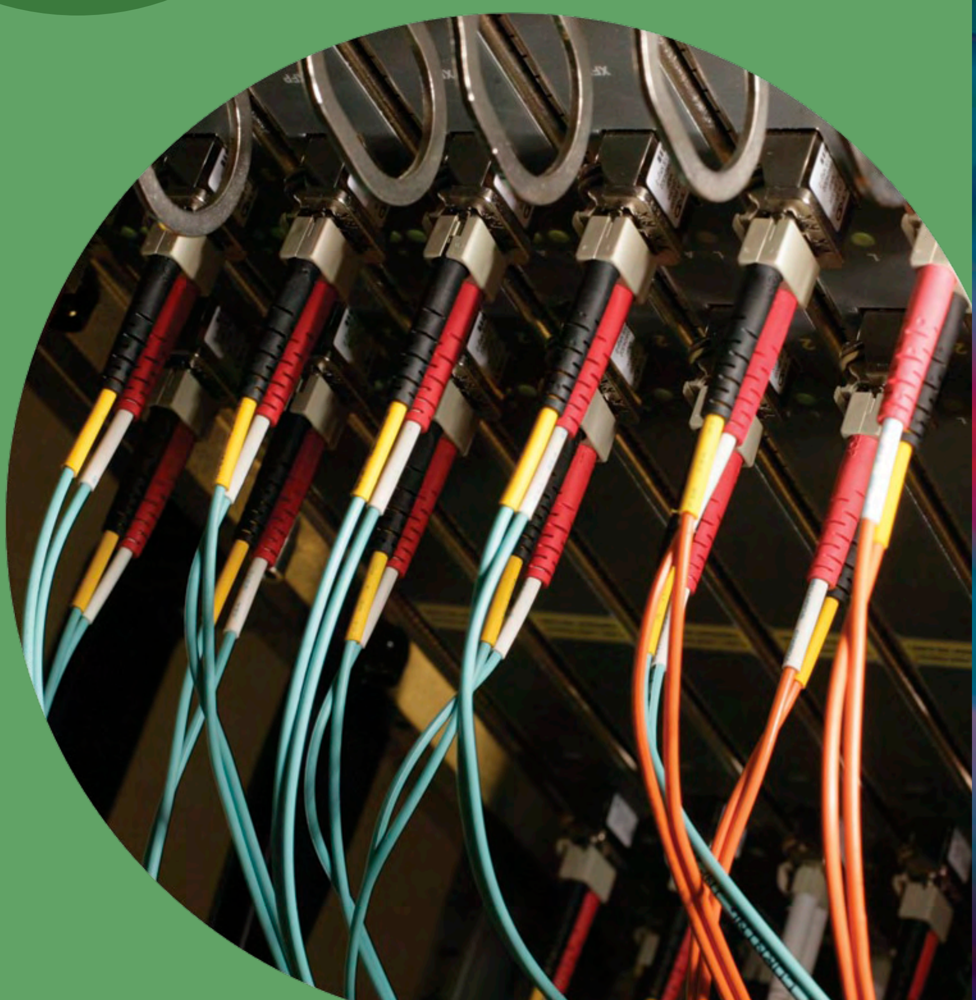

Operations Perspectives

The year 2005 was a busy time for the operations team at CCS. Expansion to meet the escalating demand for scientific computing resources drove a wave of new installations and upgrades.

This has been the biggest year ever for changes to the computer system in the CCS. Every new supercomputer system is rapidly challenged by one even more powerful. Thus maintaining “leadership” status demands not just perpetual but ever-accelerating motion.

The most dramatic addition to CCS was the 6-month phased delivery of the 25-TF Cray XT3 (Jaguar)—installation, stabilization, and integration into the CCS environment. This newest supercomputer in the CCS collection contains 5212 processors and 10.7 TB of memory. Acceptance of the system was completed on September 30; the machine was officially open for production then, and the first set of time allocations was assigned.

Equally important to researchers who need its specialized capabilities was the upgrade of the Cray X1 to the 18.5-TF X1E (Phoenix). With 1024 multistreaming vector processors, Phoenix is the largest vector processing supercomputer available for open research in the United States, and its high-performance processors are essential for some types of simulations.

During 2006, CCS plans an ambitious set of upgrades designed to enable scientific discovery for the scientists and engineers who use the center. Work is under way to quadruple the size of Jaguar to 100 TF by 2006, making it the most powerful unclassified scientific computer in the world. The CCS goal is always to have the most powerful open scientific machines in the world.

Care is being taken to accomplish this upgrade with minimal interruption of the existing system. The upgrade to 100 TF is a two-step process: replacing the existing processors with dual-core processors, and then doubling the number of cabinets.

The CCS computer room

Another major advance for 2006 is a centralized file system, dubbed "Spider," spanning all the production systems plus data analysis system and visualization. This is going to be huge. It will mean users won't have to worry about moving data between systems; they can put it in one place and access it from everywhere. It will be high-performance—we're aiming for many gigabytes per second. A prototype is running and being tested for compatibility with all the major systems. Spider is built on the Lustre technology, which is already in use on Jaguar.

We have begun planning for the next large ultra-scale system, a 1-petaflops system to be installed in 2008. This system is likely to test the limits of our existing facility. CCS was designed to be upgradable for power and cooling; it will probably be upgraded beyond the original design point.

The computers of the future will use dramatically more power. In the past, performance improvements always outstripped increases in power requirements; now power demand is beginning to increase in parallel with performance. Jaguar uses 2 MW. The 1-PF machine will draw 8–10 MW. CCS is developing plans to increase the power and cooling available to support this class of system.

Another addition in 2006 will be establishing a CCS users group, comprising all researchers with accounts on the systems, as an integral part of the QA process. It will meet two or three times a year to provide information on what is and is not working well to guide plans for the next year. Input from the user group will feed into a requirements document, which will go to a technology council that selects technologies to meet user requirements.

CCS is always dealing with tomorrow's needs at the same time it is dealing with current needs.

Buddy Bland
LCF Project Director

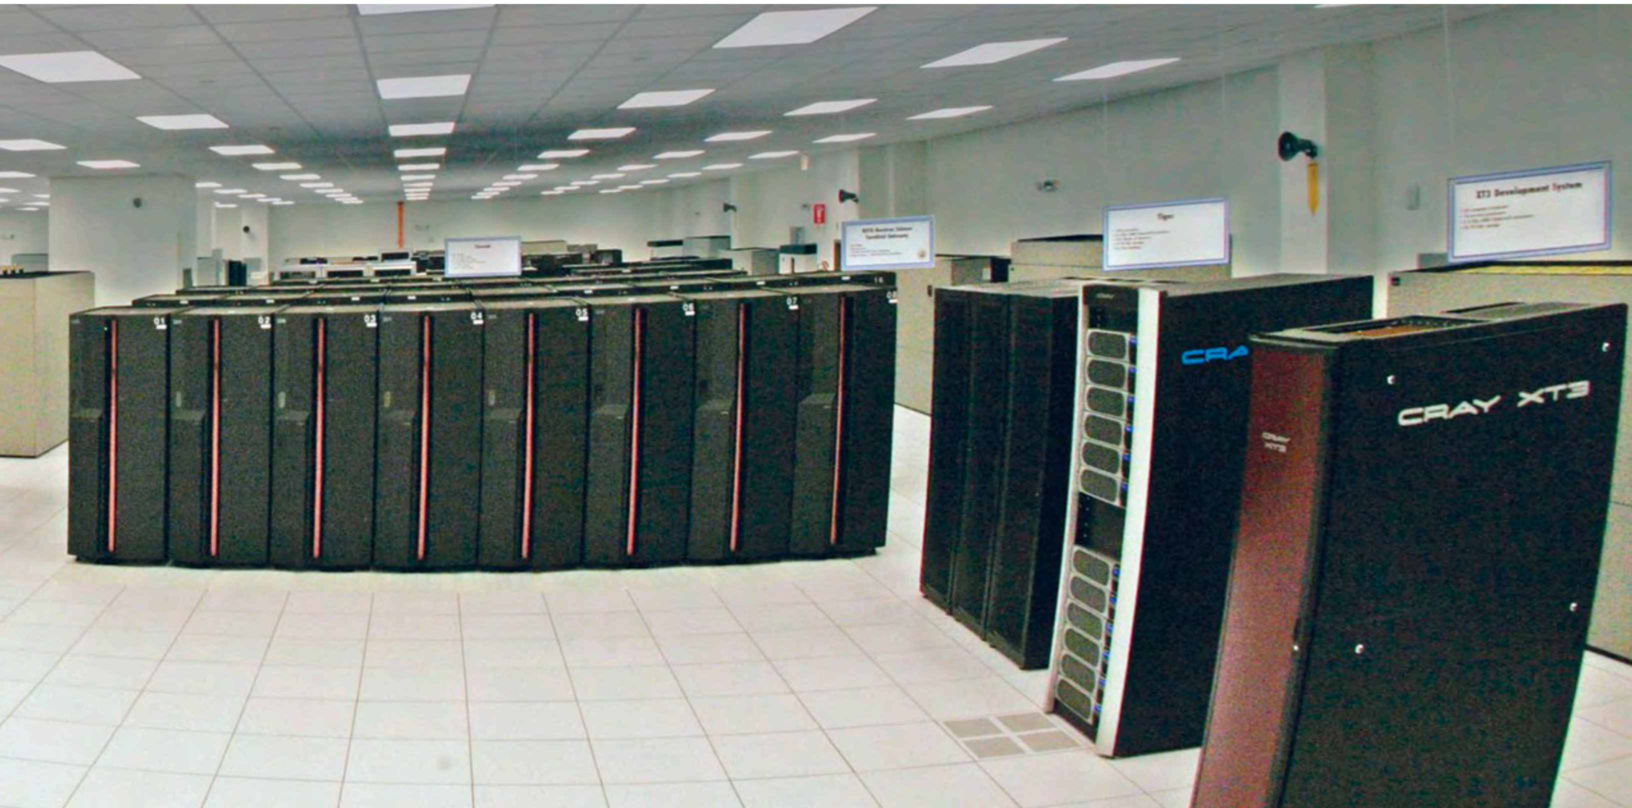

CCS Infrastructure

CCS houses not only the country's highest-capacity computers but also ultramodern infrastructure and a highly experienced professional staff to pave the way for breakthrough scientific computing.

The Computing Systems

The primary CCS computing systems are the Cray Jaguar and Phoenix computers. Two smaller systems serve as utility resources.

JAGUAR is a 5212-processor Cray XT3 system providing a peak performance of over 25 teraflops and over 10 TB of memory. There are planned upgrades of Jaguar to 100 teraflops in 2006 and eventually to 250 teraflops. Jaguar's features include

- scalable processing elements, each of which has its own high-performance AMD Opteron processors and memory
- high-bandwidth, low-latency interconnect
- MPP-optimized operating system
- standards-based programming environment
- sophisticated RAS and system management features
- high-speed, highly reliable I/O system

Each Jaguar processing element includes an Opteron 2.4-GHz, 64-bit processor, dedicated memory, and a HyperTransport link to a dedicated Cray SeaStar communication chip. The design ensures uniform performance across distributed processes to enable the use of scalable algorithms. Each processor has an on-chip 1-MB cache that can issue as many as nine instructions simultaneously. The integrated memory controller is particularly appropriate for algorithms that require irregular memory access because it enables access to local memory in under 60 nanoseconds. The 128-bit memory controller provides 6.4 GB/second local memory bandwidth per processor, more

Jaguar (Cray XT3 system)

CCS Roadmap

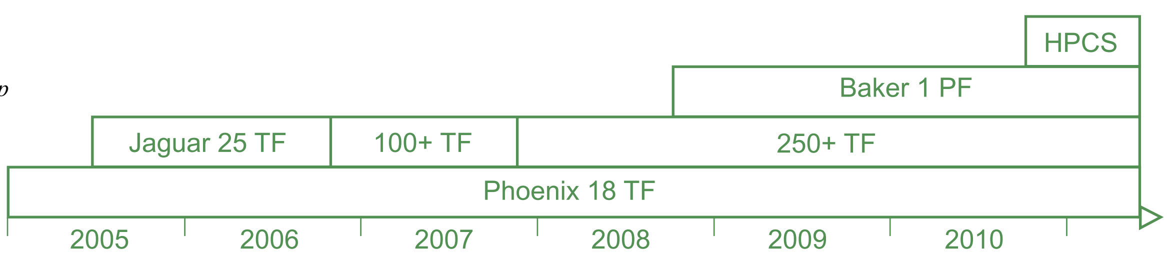

than one byte per floating operation; this is an advantage for algorithms that stress local memory bandwidth. HyperTransport enables a 6.4-GB/second direct connection between each processor and the system interconnect. Each processing element has 2 GB of memory. All memory is protected and highly reliable even in systems with tens of thousands of memory modules.

The architecture incorporates a high-bandwidth, low-latency interconnect that directly connects all the processing elements in a 3-dimensional torus topology, eliminating the need for external switches. It improves reliability and allows systems to economically scale to tens of thousands of nodes.

The operating system, UNICOS/lc, is designed to run large, complex applications and scale efficiently to 30,000 processors.

PHOENIX is a Cray X1E with 1024 multistreaming vector processors (MSPs) and 2 TB of globally addressable memory. The peak performance of Phoenix is 18.5 teraflops. Its features include

- an advanced processor architecture that combines vectorization with hardware-enabled processor coupling
- a scalable system architecture with thousands of processors able to share memory with one another
- a true single-system-image operating system
- a mature programming environment, including vectorizing compilers and support for a variety of optimized parallel programming models

Phoenix (Cray X1E system)

Each MSP has 2 MB of cache. Four MSPs form a node with 8 GB of shared memory. Each MSP contains four single-streaming vector processors with hardware features that allow them to be operated collectively with a second, low-latency level of parallelism called streaming. The compilers can recognize and schedule standard Fortran, C, or C++ code to take advantage of both vectorization and stream parallelism.

Memory bandwidth is very high, roughly half the cache bandwidth. The interconnect functions as an extension of the memory system, offering each node direct access to memory on other nodes at high bandwidth and low latency. The MSPs on each node can directly address memory on any other node by requesting data over the interconnect, bypassing the local cache. This mechanism is more scalable than traditional shared memory. Each processor can have up to 2048 outstanding memory references, allowing applications to tolerate global network latencies.

The operating system, UNICOS/mp, is a true single-system-image operating system. UNICOS/mp takes advantage of the distributed shared memory architecture to simplify administration, the I/O architecture, and provide a single login. Administrators can manage Phoenix as if it were a single node.

RAM is a 256-processor SGI Altix with 2 TB of shared memory. Each processor is the Intel Itanium 2 1.5-GHz processor. The full system runs a single Linux image, and the large shared memory facilitates analysis of very large data sets. The peak performance of Ram is 1.5 teraflops. Ram is available to users of Jaguar and Phoenix for pre- and post-processing work.

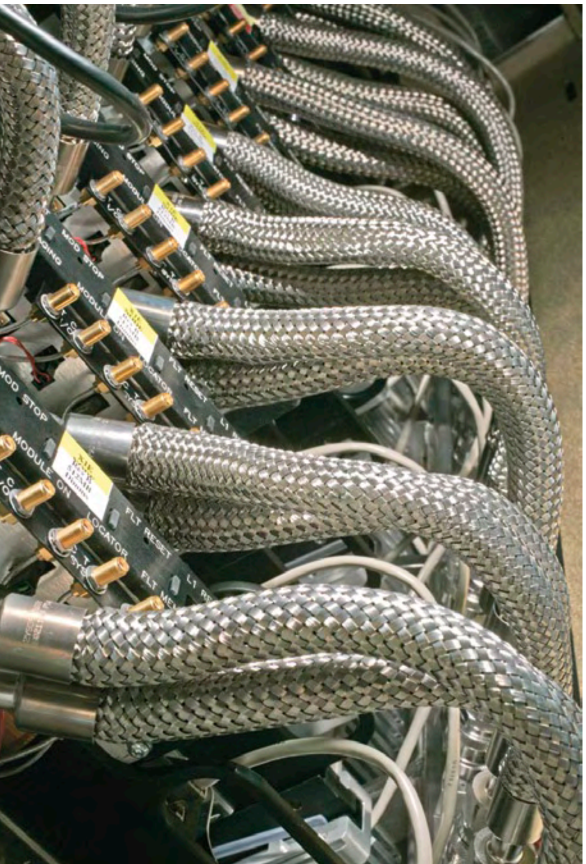

CHEETAH is a 27-node IBM Power-4 system. Each Power-4 node has thirty-two 1.3-GHz Power4 processors. Twenty nodes have 32 GB of memory, five nodes have 64 GB of memory, and two nodes have 128 GB of memory. The peak performance of Cheetah is 4.5 teraflops.

George Phipps (L) and Bronson Messer examine Jaguar.

Power and Cooling

The Tennessee Valley Authority (TVA) supplies power for the facility, which currently can provide up to 12 MW of power and 3600 tons of cooling. TVA is installing a new ORNL substation, to be operational in November 2006, to expand power supplies to CCS and other facilities. It will have two independent 161-kV supply sources with a capacity of 150 MW.

Cooling for CCS—3600 tons of capacity—is provided by chilled water that is directly connected to water-cooled systems and connected through computer room air-handling units to air-cooled systems. A redundant chiller enables continued operation in the event of chiller failure or during maintenance. The chiller capacity is being upgraded to accommodate up to 40 MW of power and to provide greater redundancy by connecting with the laboratory-wide chilled water system. The combination of power and cooling upgrades will allow CCS to house as many as three 10- to 12-MW petascale computer systems simultaneously.

Network Connectivity

Exponential growth in computing speed and datasets demands matching expansion of data storage, I/O, and networking capabilities. CCS undertook a major upgrade to a 10-gigabits (Gb) Ethernet infrastructure in 2005. It provides 10-Gb links among the CCS production and storage systems and a 10-Gb path to external networks that connect CCS with research institutions around the country. ORNL is connected to every major research network via optical networking equipment running over leased fiber optic cable. Only a fraction

What Is the HPSS?

The High-Performance Storage System (HPSS) is a CCS system that allows researchers to store and rapidly access data from their simulations. HPSS uses storage area network technology to provide highly scalable hierarchical storage management of disk and tape data.

CCS can store petabytes of data in its robotic tape library. Each tape cartridge holds from 20 to 200 GB of uncompressed data. New tape technology soon will allow storage of 500 GB of uncompressed data on one cartridge and 750 GB or more with compression.

In 1997, CCS adopted StorageTek Powderhorn libraries to position itself to effectively handle petabytes of data. The CCS tape silos (libraries) can each hold around 5000 cartridges. Four silos house eighteen 9840A drives (20-GB cartridges, uncompressed), two 9940A drives (60-GB cartridges, uncompressed), and sixteen 9940B drives (200-GB cartridges, uncompressed). The 9840A and 9940A drives read and write uncompressed data at 10 MB/second; the 9940B drives read and write at 30 MB/second. The 9840 tape technology provides fast seek time for small file access. The 9940 drives can store more data on each tape for large data sets; these are the capacity drives.

Rob Silva views one of the HPSS tape libraries

The tape library can locate a tape, take it from the slot, insert it into a tape drive, and have it begin transferring data on an average of 12 seconds for the 9840A and 59 seconds for 9940B tape drives.

HPSS is the result of a collaboration among Oak Ridge, Lawrence Livermore, Los Alamos, Sandia, and Lawrence Berkeley national laboratories, with significant contributions by universities and other laboratories worldwide. IBM is the commercial partner.

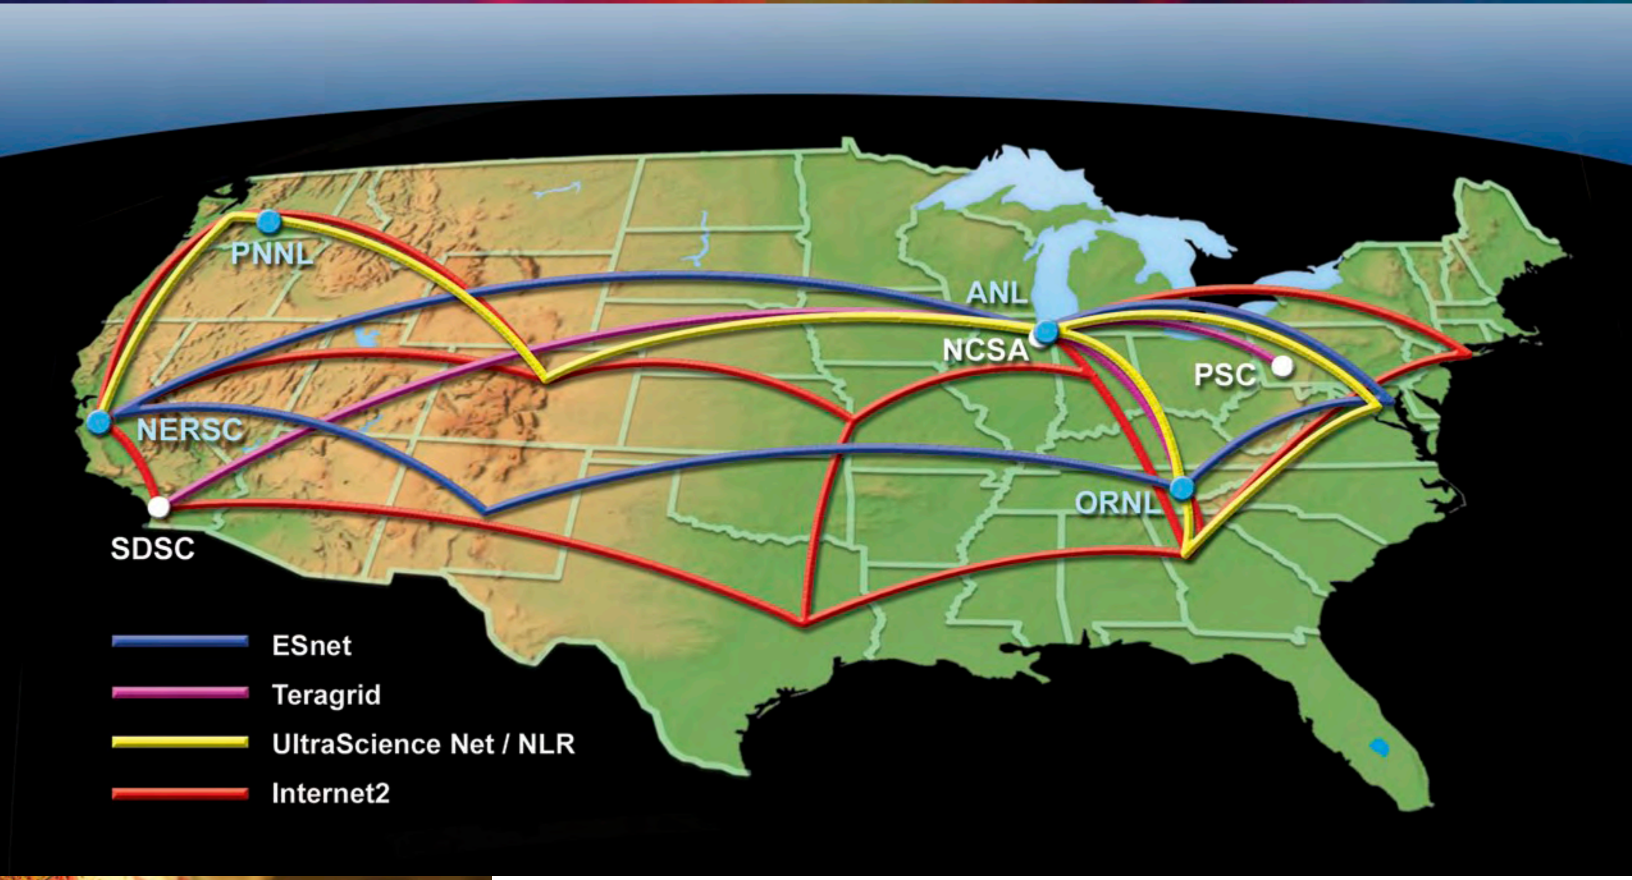

of the available 10-Gb circuits are currently in use, allowing for almost unlimited expansion of networking capacity. Currently, the connections into ORNL include TeraGrid, Internet2, ESnet, and Cheetah at 10 Gb/second, as well as UltraScienceNet and National Lambda Rail at 20 Gb/second.

Inside CCS, the development of "SuperCNS" by the Technology Integration group significantly improved the network performance of the Cray X1E by increasing bandwidth to users to more than 800 Mb/s over a 1000-Mb/s link. Bandwidth is also being added for the High Performance Storage System; its I/O bandwidth will increase in 2006 to about 56 Gb/s, compared with less than 1 Gb/s previously. Reworking maintenance schedules also eliminated 75% of the downtime for HPSS.

Cybersecurity, always a top priority, is complicated by the explosion in computing capacity and speed. Since few cybersecurity tools exist for operation at 10 Gb/s, CCS is working with vendors to ensure that security is maintained. During 2005, a one-time password infrastructure based on RSA SecurID electronic passwords was put into place. It will help eliminate system downtime due to attacks from hackers. The system will be fully in place for the new FY 2006 allocations.

Visualization and Collaboration

Scientific visualization at CCS expanded dramatically in 2005. The Exploratory Visualization Environment for REsearch in Science and Technology (EVEREST) is a valuable scientific tool for CCS users, allowing them to view collections of data that could not be rendered visually in the past.

EVEREST is a 30×8 ft PowerWall for data exploration and analysis. It displays the output from 27 projectors arranged in a 9×3 array, each providing 3500 lumens for a bright display. The projections are almost seamlessly edge-matched for an aggregate resolution of more than 11,000 by 3,000 pixels, or a total of 35 million pixels of visual detail. There is no need for users to pan or zoom to view their data—large datasets can be viewed on the wall in their entirety.

EVEREST has a 600-ft^2 projection area and a 1000-ft^2 viewing area that can accommodate up to 25 people. The PowerWall Toolkit is a GUI environment that enables groups to use the PowerWall as a large desktop pixel space for static images, movies, and interactive 3-dimensional visualizations. Other visualization capabilities include LCD arrays and a reconfigurable CAVE.

The projection environment is driven by a 64-node rendering and analysis cluster made up of dual-processor Opteron workstations. The cluster is networked to the other CCS resources and performs additional visualization-related functions including computation, pre-analysis, and pre-rendering. The rendering environment uses 64-bit Suse Linux, Chromium, Distributed Multi-Head X (DMX), and state-of-the-art graphics cards with pixel shader support.

Visualization resources at CCS can be used on site or accessed remotely. Plans for 2006 are to make visualization an integral part of the simulation process and easier to use, making it simpler for scientists to do visualization on their own either in EVEREST or on their own desktops.

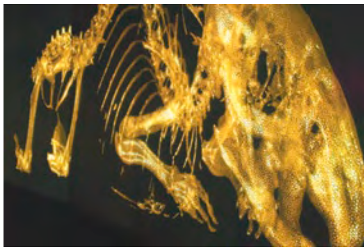

Sean Ahern of the Scientific Computing group at the PowerWall in EVEREST

CCS Staff

CCS staff more than doubled during 2005 as the center ramped up operation. At the same time they were racing to put new systems in production, CCS staff were developing their roles in a new organizational structure. Four new groups were established during 2005 to help operate CCS, two focusing on computers and technology and two on users and the science they are conducting.

The Technology Integration group

The High-Performance Computing Operations group

Technology Integration works with the chief technology officer and High-Performance Computing Operations to develop and enhance the unifying infrastructure that supports the CCS systems and provides systems-level expertise. This includes evaluating and integrating emerging technologies in areas such as archival storage, file systems, networks, and cyber security. The group also provides system programming to integrate new systems into the CCS infrastructure and guides HPSS development.

High-Performance Computing Operations is responsible for installation and configuration, systems administration, and cyber security for the CCS computers and networks as well as the Teragrid cluster. In addition to supporting the computational platforms, Operations supports the storage needs of the CCS, ARM, Teragrid and other projects with the High-Performance Storage System. The group provides around-the-clock operations coverage of the CCS systems, as well as configuration management, system performance monitoring, account software management, disk space management, web services, problem reporting, and software licensing.

User Assistance and Outreach is the "face" of CCS for new users. It sets up new user accounts, answers questions about supercomputer operations, helps users run or compile code and access their accounts, prioritizes service requests, and maintains access policies and procedures. The group provide phone response to user queries 24/7

The User Assistance and Outreach group

and operate the User Assistance Center from 9:00 to 5:00, Monday through Friday.

The **Scientific Computing** group are research scientists who provide a liaison between the computer users and CCS. Through direct collaboration, they augment and extend project computational and domain-specific expertise. They also represent users in CCS planning exercises in day-to-day facility operation and serve on the User Council, one mechanism for requirements gathering. Members of the group have extensive experience in problem solving and in porting, tuning, and developing application software on CCS resources. They use their experience to provide in-depth support through active participation in domain sciences, applications, algorithms, libraries, tools, visualization, data movement and workflow.

Members of these four operational and engineering groups staff CCS 24 hours a day, 365 days a year to provide for continuous operation of the center and immediate problem resolution.

Cray Supercomputing Center for Excellence

In addition to those groups, CCS is housing the newly established Cray Supercomputing Center for Excellence, which is composed entirely of Cray employees. This group provides system expertise to facilitate breakthrough science on Cray architectures, i.e., application targeting, porting, optimization, library development, tool development, and training. In addition to their own expertise, Cray Center staff can draw directly from Cray to address performance problems quickly. When not engaged in addressing specific problems, the staff help port and optimize code and train ORNL staff and researchers in effective use of the Cray systems.

The Scientific Computing group

5

Looking Ahead

2006 and Beyond

Thomas Zacharia
Associate Laboratory Director for Computing and Computational Sciences

In May 2004, DOE announced that ORNL would lead a U.S. effort to reclaim this country's historical position as the world leader in scientific computing. In response to the challenge, by September 2005 CCS had installed and commissioned the most powerful supercomputers for unclassified scientific research in the country. But the path forward has no rest stops: the next two years will see even greater advances in the capacity of CCS to support groundbreaking computational science.

Exciting results are emerging from simulations conducted on the CCS computers. Materials scientists working at CCS achieved the first credible solution of the Hubbard model for describing superconductivity in materials—a key physics problem. Plasma physicists are using CCS to conduct the fastest, most detailed simulations of fusion plasma turbulence ever achieved. Both of these research areas have enormous implications for our energy future.

Early this year, ORNL entered into a contract with Cray to install a petaflops machine at CCS by 2008—the first contract ever for acquiring a petaflops computer. Installing a machine capable of petaflops speed is a remarkable milestone, but the most important accomplishment will be the science these new computers will make possible. Researchers are using the ORNL computers for the most advanced studies ever conducted in areas such as astrophysics, combustion, chemistry, global climate change, fusion energy, accelerator design, and materials.

Computation is now synonymous with scientific discovery, and pre-eminence in high-performance computing is indispensable to maintaining U.S. leadership in science and technology. CCS is excited to be instrumental in pushing computational science in the United States to new levels.

ORNL's role as the center for high-performance computing in the United States complements a portfolio of other accomplishments such as the Spallation Neutron Source, the Center for Nanophase Materials Sciences, the Advanced Microscopy Laboratory, and being named the headquarters for the International Thermonuclear Experimental Reactor. ORNL is the place where superb science will be possible in the coming years because of the unique, powerful research tools available here.

An image of a doped carbon nanotube displayed on the EVEREST PowerWall

Buddy Bland
LCF Project Director

During the past year, ORNL and DOE announced that the leadership computing facility at ORNL would switch to a project basis to focus squarely on the development and installation of a petaflops-speed supercomputer at CCS. The Leadership Computing Facility Project is tasked with upgrading the Jaguar supercomputer to 100 teraflops by the end of 2006 and 250 teraflops in 2007 and then installing a petaflops supercomputer by the end of 2008.

This is an aggressive roadmap that requires careful planning and execution by both ORNL and Cray, which is supplying the computer systems. A great deal of rigor and careful planning are necessary to ensure the plans are well thought out and feasible to implement.

It is CCS's desire and goal to minimize disruption to the user community as much as possible while bringing the upgraded system into production. As the planned upgrades are made to the existing Jaguar computer, however, there will of necessity be some hopefully brief interruptions as the computer is upgraded and the acceptance tests completed at each phase.

The new computer will be power-hungry, using as much as 10 MW of electricity. To support this and other power needs of the Laboratory, ORNL has undertaken to install a new 70-MW power station on the campus and to upgrade its chilled water capability to help cool the computer systems.

These upgrades are necessary to meet the increasing need of scientists for faster, more powerful tools to run the simulations that have become indispensable to modern research. The breakthrough discoveries these machines will enable are the ultimate goal of our continuing efforts.

6

Appendices

CCS Ancillary Organizations

CCS receives advice and direction from a number of affiliate groups, including an external advisory committee, an operations council, and the Joint Institute for Computational Sciences.

External Advisory Committee

The CCS Advisory Committee is composed of 12 to 18 distinguished scientists from academia, national laboratories, industry, and other research institutions across the nation. The committee provides advice to the CCS director in the broader areas of computational science, computer science, applied mathematics, operation of a national user facility, and interagency communication and coordination. The committee reports to the ORNL ALD for Computing and Computational Sciences. Responsibilities of the committee include

- Providing advice on priorities and strategies to effectively execute the mission of CCS
- Providing scientific advice, for example, which domains may be ready to achieve "breakthrough science," and potential scientific directions for the CCS user program

In addition, the Advisory Committee should advocate and promote effective communication between ORNL leadership, DOE, other federal agencies, and the user community to help facilitate mutual understanding in support of achieving maximum impact by CCS users.

The JICS building on the ORNL campus

Operations Council

CCS is one of 18 major user facilities at ORNL. The mission of the Operations (Ops) Council is to ensure that the center operates in a safe, secure, and effective manner. The Ops Council is chaired by the deputy for operations and meets weekly to discuss current operational status, concerns, activities, and future direction. The Ops council is composed of representatives of each of the operational elements needed to provide the underlying CCS infrastructure: High-Performance Computing Operations, Technology Integration, User Assistance and Outreach, Scientific Computing, Networking, Visualization, and Facility Management. Representatives for Human Resources, Recruiting, Cyber Security, Quality Assurance, Environmental, Safety and Health, Finance, and Procurement meet monthly with the Ops Council.

Joint Institute for Computational Sciences

The Joint Institute for Computational Sciences (JICS) members are drawn from ORNL and university partners to create major new modeling and simulation capabilities for terascale—and beyond—computers. JICS also promotes the training of researchers to use these tools to investigate fundamental systems and the educating of the next generation of computational scientists and engineers.

Research Alliance in Math and Science

The long-term goal of the Research Alliance in Math and Science (RAMS) is to increase the number of underrepresented individuals with advanced degrees in science, mathematics, engineering, and technology in the workforce. This program is sponsored by the Mathematical, Information, and Computational Sciences Division of the Office of Advanced Scientific Computing Research, U. S. Department of Energy.

Statistics

In this section are charts illustrating statistics on the usage of the CCS resources during our first year as the Leadership Computing Facility for the United States. They reflect the wide range of research areas that depend on leadership-class computers to move forward with science that can lead to breakthrough discoveries.

The charts on this page illustrate the percentage of computational hours used by DOE program office and by discipline. The program usage chart shows that Basic Energy Sciences and Fusion Energy Sciences were the largest users, accounting for 26% and 25%, respectively, of the hours logged.

The breakdown by discipline shows that fusion and materials research accounted for almost half of the computational hours used during 2005, with 25% and 23%, respectively.

As illustrated, CCS supports a broad portfolio of research projects critical to the nation.

Percentage of computational hours on CCS machines used for research sponsored by the various DOE program offices (top) and by researchers in various scientific disciplines (bottom)

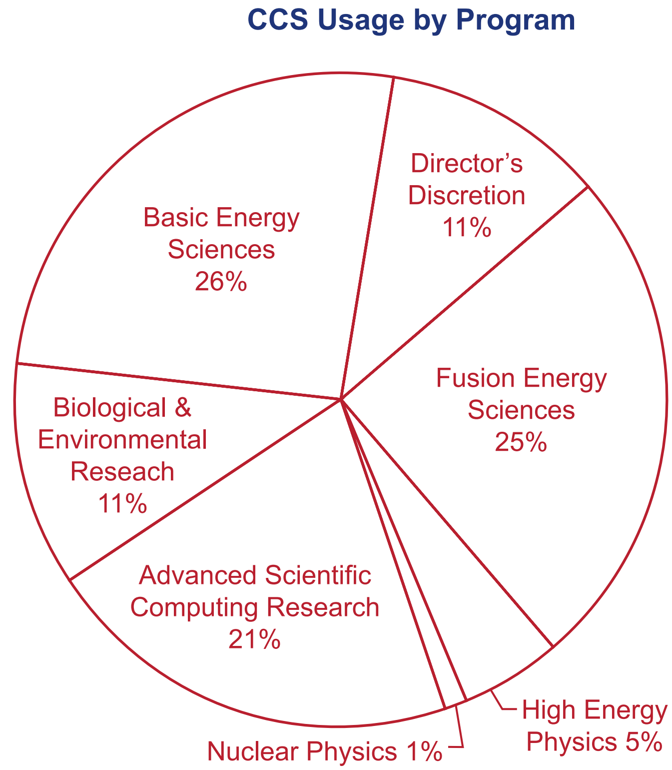

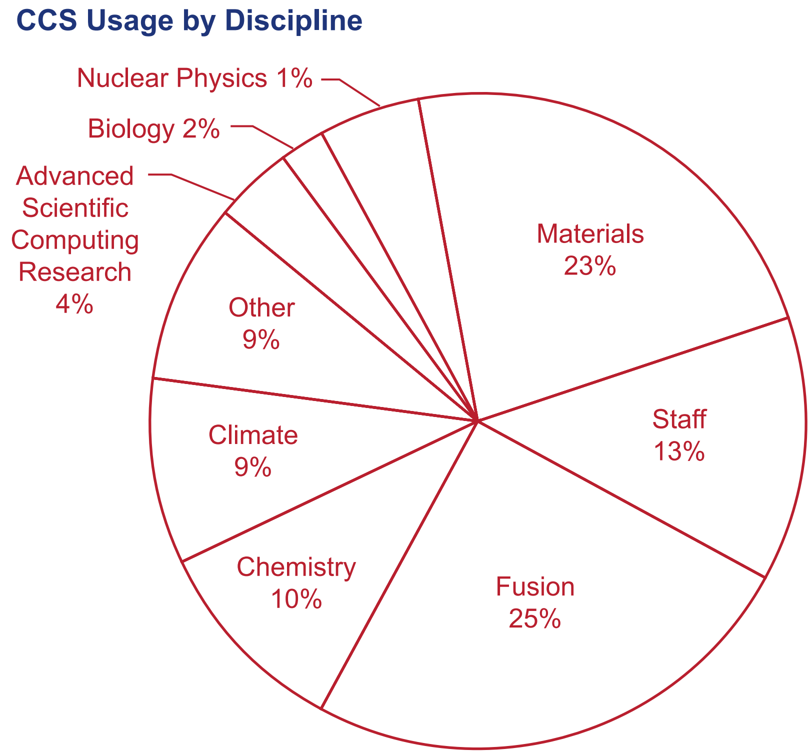

Active Users by Sponsor

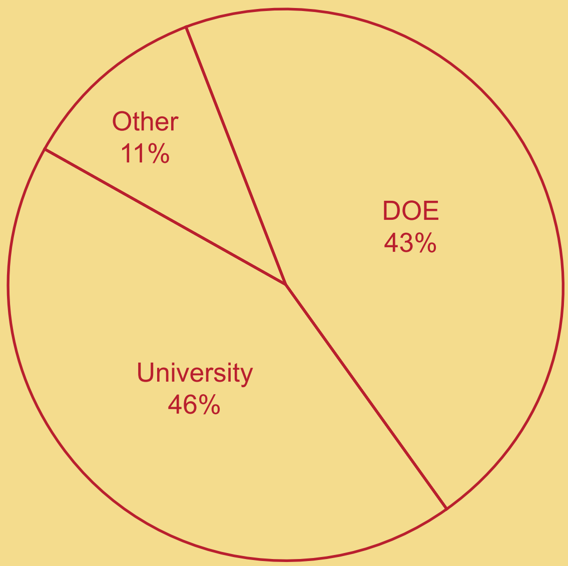

Active Users by Discipline

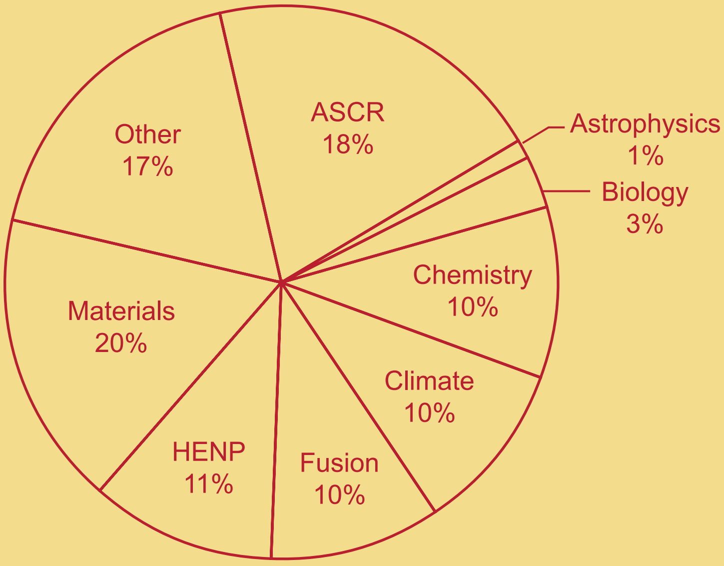

These charts illustrate the breakdown of active users of CCS systems by sponsor and by discipline. The chart on the left shows researchers from DOE programs and universities are by far the primary users: active users from universities constitute 46% of the total and researchers for DOE programs 43%. Users representing other research sponsors accounted for only 11%. According to the chart on the right, materials was the most active discipline, making up 20% of the active users. Researchers in Advanced Scientific Computing Research made up 18% of users. High-energy nuclear physics users made up 11%, and fusion, climate, and chemistry researchers each constituted 10%. Biological researchers were 3% of active users and astrophysics 1%; active users from various other disciplines made up 17% of the total.

HPSS Archival Storage

The usage chart for the High-Performance Storage System graphs trends in storage of data from 1998 to 2005. It indicates slow but steady growth until about 2002, when the amount of data stored began to climb sharply. The chart shows an increase of over 250 terabytes in data archived from 2004 to 2005.

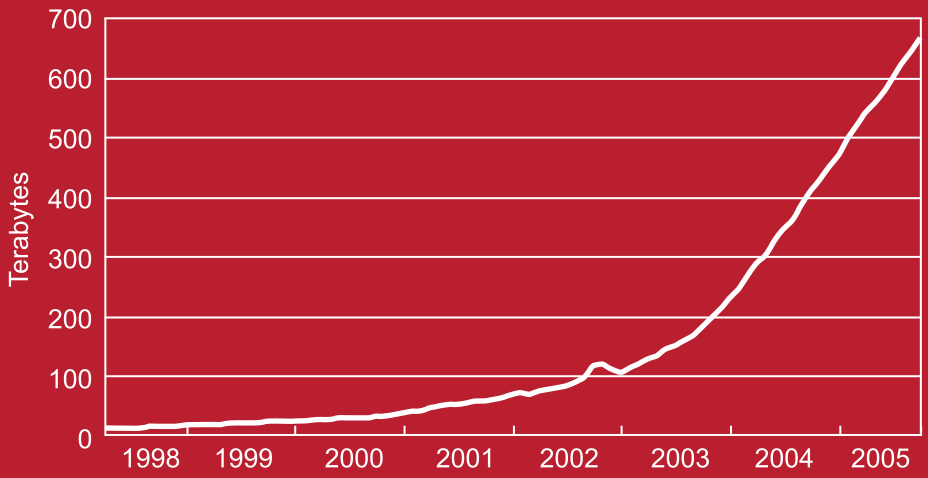

Meetings, Workshops, and Tours

CCS Users Meeting

CCS sponsored a workshop for project teams with FY 2006 allocations on Jaguar and Phoenix to allow researchers to get acquainted with each other and with CCS. Participants were scientists who had competed for and won time on the CCS comput-

CCS provides end-to-end solutions that enable users to take maximum advantage of their computing allocations

ers and support systems through the peer-reviewed proposal process. Those attending represented 17 leadership computing and 5 Innovative and Novel Computational Impact on Theory and Experiment (INCITE) projects in scientific domains such as climate, fusion, materials, astrophysics, chemistry, and biology and topics important to U.S. industry such as airplane design and animation.

The first 2 days of the workshop included overviews of CCS, Jaguar and Phoenix architecture and software, and support services available to researchers. Users toured the facilities and saw live demos of the equipment. Representatives of the projects presented overviews of their research goals. The group spent an afternoon organizing a User Council and Tech Council to facilitate communication among projects and ensure that all system and software requirements are accommodated.

Day 3 was devoted primarily to hands-on tutorials to answer porting and optimization questions. Attendees also learned about the CCS High Performance Storage System that is available to them for archiving data sets. Cray representatives presented sessions on Jaguar and Phoenix.

Users heard a detailed presentation on visualization using the EVEREST facility for data exploration and analysis, which is being used to push the limits on scientific visualization.

Workshops

Cray Rainier S/W Implementation–Roadmapping Workshop

February 1–3, Mendota Heights, MN

This workshop was organized to enable principal investigators in the DOE Mathematical, Information, and Computational Sciences research program to contribute to delivery of a robust Cray Rainier system software environment to help ensure that CCS will deliver the maximum amount of science to the nation. In this first workshop, the focus was on the lower layers of the software stack: operating system, programming models/libraries and I/O. Future workshops will focus on other layers of the S/W stack.

Triad Meeting: Multiscale Simulation: Atomistic to Continuum

April 3–5, ORNL

The Multiscale Simulation Workshop was an invitation-only event hosted jointly by ORNL, Imperial College, and Georgia Tech. The program included leading researchers from all three institutions in the computational science areas of climate, fusion, astrophysics, chemistry, materials, and biology.

SciDAC 2005

June 26–30, San Francisco

The Scientific Discovery through Advanced Computing Program (SciDAC) operated by DOE's Office of Science sponsors the annual SciDAC workshop. More than 300 researchers participated in the 2005 workshop, organized and managed by ORNL. The program included more than 100 presentations. In addition to plenary talks and technical talks each day, poster sessions were held on two evenings and there were panel discussions by applications community members, computer scientists discussing infrastructure needs, and a vendor panel discussing architecture plans. In addition, planning began for the next phase of SciDAC, which begins in 2006.

XT3 Workshop

June 14–17, ORNL

CCS and Cray held a preliminary information workshop to discuss the Cray Jaguar supercomputer coming online at CCS. Because the machine was in pre-production mode, attendance was limited to the CCS Scientific Computing group, representatives of application domains at ORNL, and Cray staff. Workshops presenters included staff members of CCS, the Cray Center of Excellence located at ORNL, and the Portland Group (the company that provides the compiler software). Four sessions were working sessions for users to port existing codes to Jaguar and obtain performance profiles for code optimization. More workshops with expanded attendance are planned for porting codes to Jaguar.

National Leadership Computing Facility Computational Chemistry Workshop

August 1–5, ORNL

This workshop featured a mix of invited and contributed talks covering nearly all aspects of computational chemistry with a common theme of enabling new science through large-scale computation. The final day of the workshop included discussion and planning of coordinated activities, collaborations, and future funding opportunities.

Workshop on Enabling Petascale Science and Engineering Applications

December 9, Georgia Tech

The focus of this workshop, organized by ORNL, Georgia Tech, North Carolina State University, and the University of North Carolina, was to characterize science and engineering applications that may require petascale computing platforms. There was a focus on the three emerging computational areas of nanotechnology, biology and biomedical applications, and environmental science. Participants were encouraged to submit benchmark codes to assist in the design and configuration of high-performance computing resources.

Many people of all ages toured the CCS EVEREST Visualization Laboratory during Community Day 2005

Viewing the ORNL booth at Supercomputing 2005 are (from left) Walt Polansky, DOE Office of Advanced Scientific Computing Research; Barbara Helland, DOE Mathematical, Information, and Computational Sciences; and John Drake, ORNL

Community Day '05

August 27, ORNL

ORNL Community Day, a public open house for visitors from the surrounding area, included public tours of CCS, the Visualization Laboratory, and the Joint Institute for Computational Sciences.

Supercomputing '05

November 12-18, Seattle

CCS staff were key participants in Supercomputing 2005, the premier international conference on high-performance computing, networking, and storage. The annual supercomputing conference convenes scientists, teachers, programmers, executive, and other representatives from the world's leading computing facilities and companies to showcase the role of high-performance computing in research, business, and education. Several dozen CCS/ORNL staff presented papers, workshops and other sessions, and posters.

Tours

CCS welcomes visitors who wish to tour its facilities. In 2005 it hosted at least 149 tours, involving more than 800 participants. These included tours of CCS alone and tours of CCS as a part of a visit to ORNL as a whole. Visitors to CCS cover a broad spectrum, including officials of DOE and other federal agencies, legislators, officials, representatives of other research institutions, Tennessee state officials, business and industry groups, student groups, and delegations from other countries.

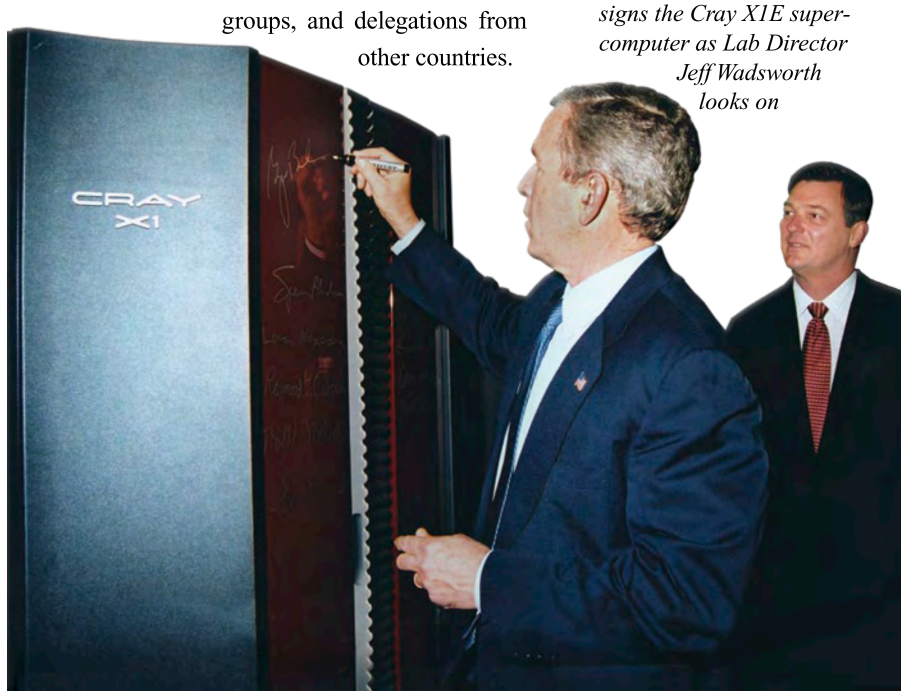

President George W. Bush signs the Cray X1E supercomputer as Lab Director Jeff Wadsworth looks on

Government officials tour the facility to view firsthand the results of allocated funding and to learn about the latest breakthroughs in computational science for future funding decisions. Representatives of other research institutions come to compare ongoing projects, facility design, and operation and to discuss ideas about how to better manage ongoing scientific research. They might also discuss possible collaborations with CCS for projects in the immediate future. Local public officials often visit to gain a firsthand perspective of the research being conducted in their state and increase the visibility of the facility for larger government bodies. Representatives of business and industry come to the facility to see the remarkably industry-friendly infrastructure of the Oak Ridge area and to discuss possible collaborations between CCS and their enterprises. Student groups that visit CCS are given a glimpse into the cutting edge of mathematics and science, hopefully encouraging the next generation of America's leaders to pursue these fields in their future careers.

Samuel Bodman, left, U.S. Secretary of Energy is shown the Cray X1E supercomputer by Thomas Zacharia, Associate Laboratory Director for Computing and Computational Sciences

Finally, because America is leading the world in the future of technology, international delegations are given a close-up look at the country's premier computational science facility. CCS hopes this knowledge will set an example and up the ante for research throughout the rest of the world. CCS has hosted tours for everyone from the Boy Scouts to top government officials and the Center's advances will no doubt continue to draw more people in the future. Tours are tailored to meet the needs of specific groups and are led by an experienced tour guide who will discuss the technical aspects of the facility and explain how it is helping to solve real-world problems.

Carl Kohrt, President and CEO of Battelle Memorial Institute, signs the Cray XT3 supercomputer during a visit to the Laboratory

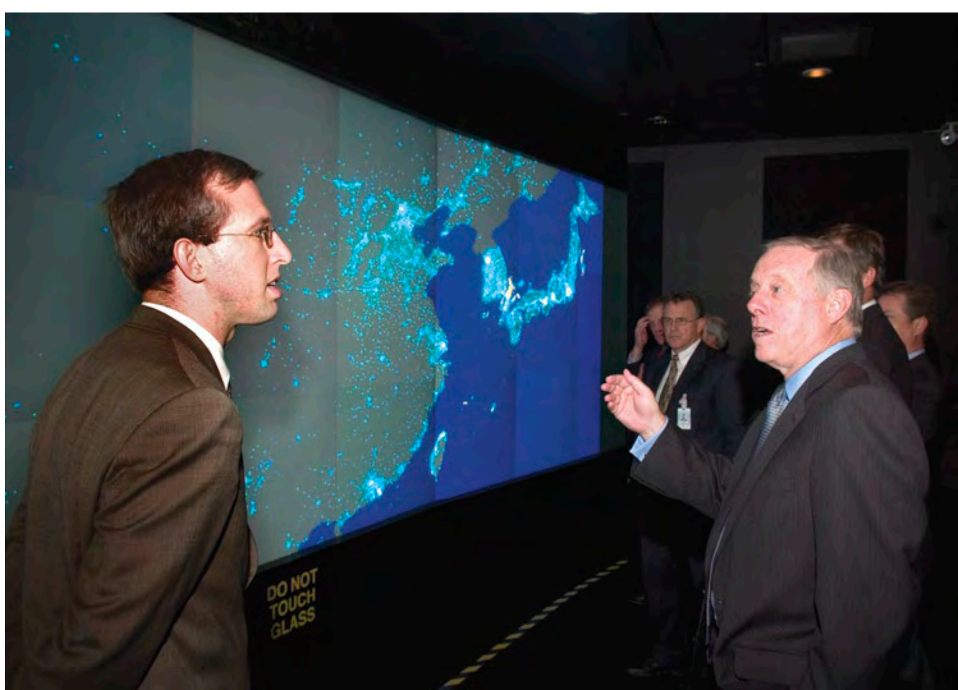

Sean Ahern of CCS, left, gives Tennessee Governor Phil Bredesen a tour of the EVEREST visualization laboratory

Selected User Publications

Numerous scientific breakthroughs have ocurred as a result of research conducted at CCS. Listed below is a small sampling of the nearly 100 publications from 2005, grouped by related discipline, that highlight a portion of the work being achieved through the combination of talented researchers and CCS resources. For a complete listing of 2005 publications, please refer to the enclosed CD.

ASTROPHYSICS

Altun, Z., A. Yumak, N. R. Badnell, J. Colgan, and M. S. Pindzola. 2005. "Dielectronic recombination data for dynamic finite-density plasmas, VI: The boron isoelectronic sequence (vol. 420, p. 775, 2004)." *Astronomy & Astrophysics* **433**, no. 1, 395.

Blondin, J. M., and A. Mezzacappa. 2006. "The spherical accretion shock instability in the linear regime." *Astrophysical Journal* **642**, 401–409.

Walder, R., A. Burrows, C. D. Ott, E. Livne, I. Lichtenstadt, and M. Jarrah. 2005. "Anisotropies in the neutrino fluxes and heating profiles in two-dimensional, time-dependent, multi-group radiation hydrodynamics simulations of rotating core-collapse supernovae." *Astrophysical Journal* **626**, 317-332.

CHEMISTRY

Gohda, Y., and S. T. Pantelides. 2005. "Charging of molecules during transport." *Nano Letters* **5**, 1217-1220.

Schulthess, T. C., W. M. Temmerman, Z. Szotek, W. H. Butler, and G. M. Stocks. 2005. "Electronic structure and exchange coupling of Mn impurities in III–V semiconductors." *Nature Materials* **4**, 838–844 (Nov.).

Sumpter, B. G., P. Kumar, A. Mehta, M. D. Barnes, W. A. Shelton, and R. J. Harrison. 2005. "Computational study of the structure, dynamics, and photophysical properties of conjugated polymers and oligomers under nanoscale confinement." *Journal of Physical Chemistry B: Condensed Matter, Materials, Surfaces, Interfaces and Biophysical* **109**, no. 16, 7671–7685.

CLIMATE

Guo, D. X., and J. B. Drake. 2005. "A global semi-Lagrangian spectral model of the shallow water equations with variable resolution." *Journal of Computational Physics* **206**, no. 2, 559–577.

COMBUSTION

Hawkes, E., R. Sankaran, P. Pébay, and J. Chen. 2006. "Direct numerical simulation of ignition front propagation in a constant volume with temperature inhomogeneities, Part II: Parametric study." *Combustion and Flame* **145**, 145–159.

FUSION

Estrada-Mila, C., J. Candy, and R. E. Waltz. 2005. "Gyrokinetic simulations of ion and impurity transport." *Physics of Plasmas* **12**, 022305 (Feb.).

LATTICE QCD

Aubin, C., et al. 2005. "Semileptonic decays of *D* mesons in three-flavor lattice QCD." *Physical Review Letters* **94**, 011601 (Jan.).

MATERIALS

Klie, R., J. Buban, M. Varela, A. Franceschetti, C. Jooss, Y. Zhu, S. Pantelides, and S. Pennycook. 2005. "Enhanced current transport at grain boundaries in high-T_c superconductors." *Nature* **435**, 475-478.

Maier, T., M. Jarrell, T. C. Schulthess, P. R.C. Kent, and J. B. White. 2005. "Systematic study of *d*-wave superconductivity in the 2D repulsive Hubbard model." *Physical Review Letters* **95**, 237001 (Nov. 29).

Smirnov, A. V., W. A. Shelton, and D. D. Johnson. 2005. "Importance of thermal disorder on the properties of alloys: origin of paramagnetism and structural anomalies in bcc-based $Fe_{1-x}Al_x$." *Physical Review B* **71**, 064408 (Feb.).


CRAY
X1E
CRAY
X1E
CRAY
X1E