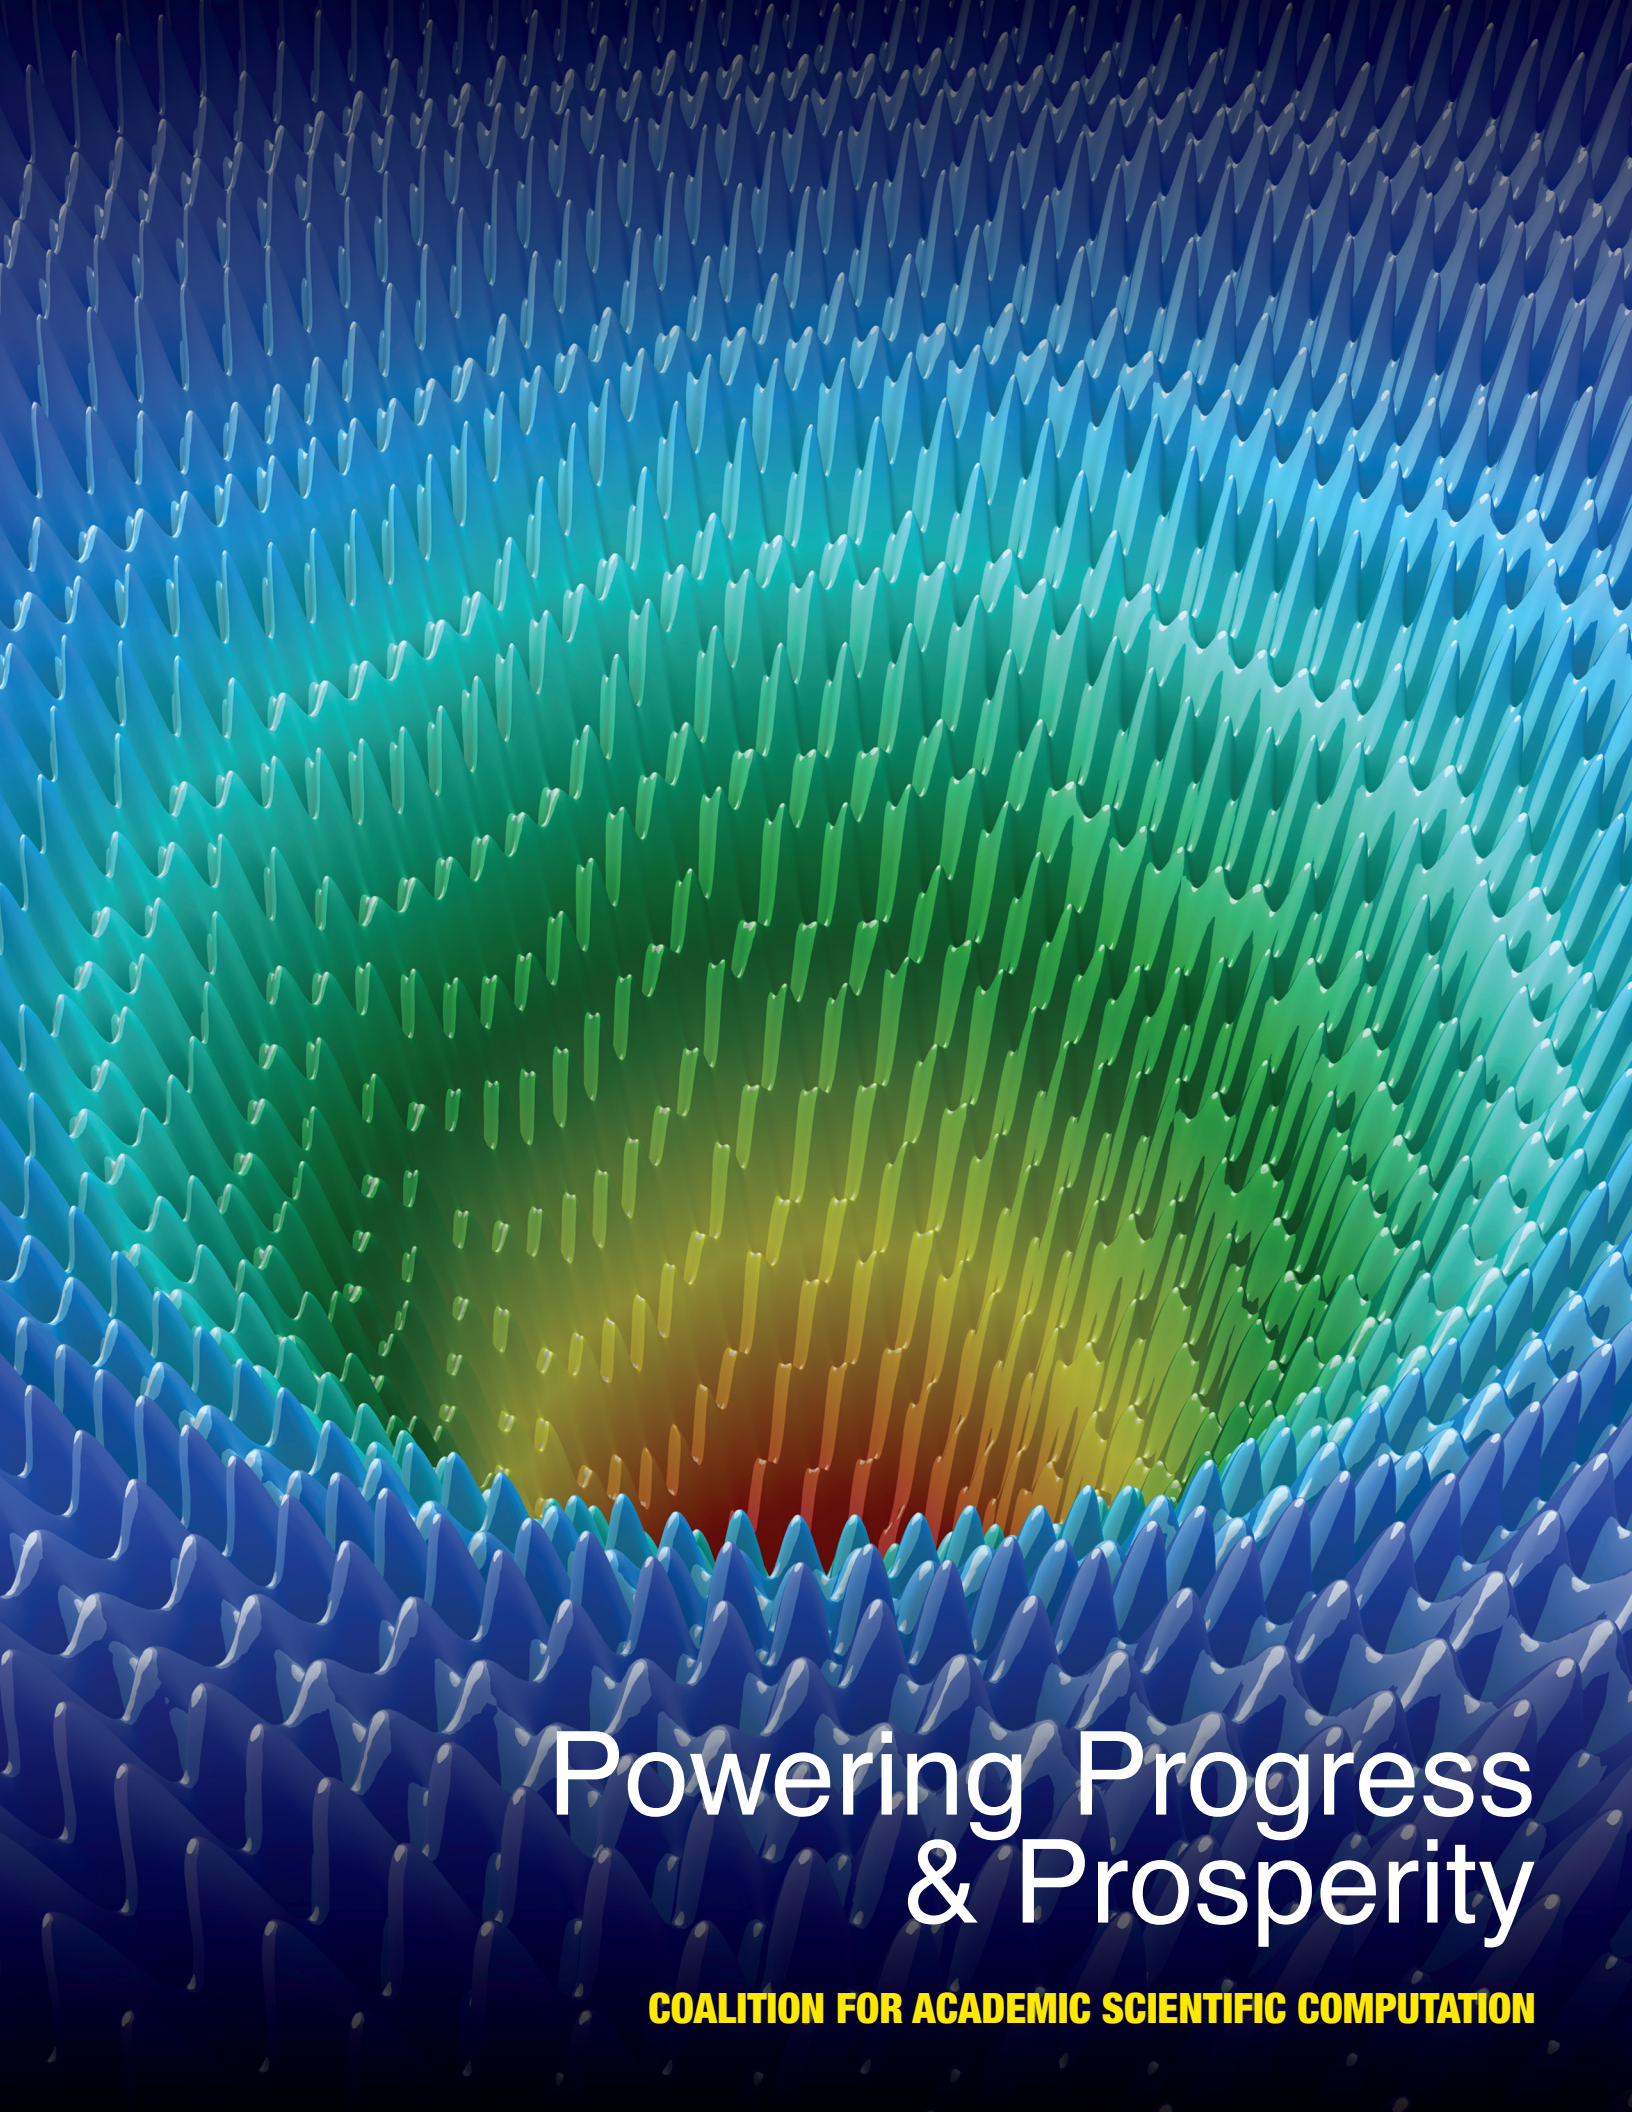

Powering Progress & Prosperity

COALITION FOR ACADEMIC SCIENTIFIC COMPUTATION

CONTENTS

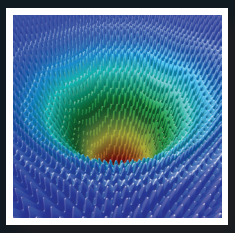

ABOUT THE COVER
What looks like a colorful whirlpool is actually a three-dimensional "cost-surface" containing many peaks and valleys. Valley locations outside the red-shaded base of the whirlpool represent "artificial minima," a mathematical phenomenon that bedevils engineers charged with cleaning up solvent-contaminated groundwater.

"We use computer models to predict the behavior of alternative cleanup strategies. We link these computer models with computerized search algorithms to try and find the least expensive design. But the search algorithms are prone to finding artificial minima instead of the truly best region of the design space," said Dr. Loren Shawn Matott of the University at Buffalo's Center for Computational Research. "So we need to develop better search algorithms that can avoid these artificial minima, not an easy task." Applied nationwide, these enhanced search algorithms could lead to better and cheaper cleanup strategies, **potentially reducing cleanup costs by as much as $100 million**.

Matott and a team of undergraduates participating in the National Science Foundation-sponsored "URGE to Compute" program are currently applying a suite of enhanced search algorithms to a 580-acre EPA Superfund site on the outskirts of Billings, Montana. They are using the algorithms to efficiently identify least-cost "pump-and-treat" strategies for two large plumes of chlorinated solvents slowly migrating through the subsurface towards the Yellowstone River.

COLLABORATORS FROM LEFT TO RIGHT
L. Shawn Matott, University of Buffalo Center for Computational Research; Shirmin Aziz, Buffalo State College; Beth Hymiak, Buffalo State College; Camden Reslink, University at Buffalo; Christine Baxter, University at Buffalo

The **Coalition for Academic Scientific Computation** is an alliance of 70 of America's most forward thinking research universities, national labs, and computing centers. We are dedicated to providing researchers with one of their most important scientific tools: Computation.

Scientists use CASC member facilities every day to sift millions of molecules for promising drug candidates; slow down nanosecond events happening inside atoms, molecules, cells, and stars; peer into the heart of a flame; rearrange atoms like tinker toys; and travel to the edge of the universe and beginning of time.

The research enabled by CASC member institutions not only broadens our scientific horizons, but also powers the progress and prosperity of our nation by

- Securing our **energy** future,
- Advancing **medicine & health**,
- Innovating for **industry**,
- Meeting challenges to our **environment & natural resources**,
- Exploring **matter & the universe**, and
- Encouraging and **educating** new generations of scientists and technologists.

With smartphones in every pocket, chips in every appliance, and more of our daily lives moving to the cloud, computing dedicated to scientific research may seem anachronistic. It's not.

Scientific breakthroughs depend on specialized computing resources, including powerful "parallel" computers, specialized scientific software, and the mathematical and technical expertise to assist scientists in using these new tools of discovery.

In fact, a 2010 study published in the *Journal of Information Technology Impact* found that "consistent investments in HPC [high-performance computing] at even modest levels are strongly correlated to research competitiveness." This publication offers a glimpse of how those investments are paying off and will continue to do so in future.

For more information, please visit our web site at **http://www.casc.org**

Energy

Computational methods have become key to tackling every aspect of U.S. energy challenges, from stretching the fossil fuels we primarily rely on today, to developing new energy sources and improving the alternatives we already employ, such as solar and wind power.

To bring down the costs of producing biofuels, researchers need to understand how to break down the tough fibers of plant cellulose into sugars. Molecular-level simulations are helping them understand how enzymes—present in fungi and other organisms—already perform this trick.

These kinds of simulations are also used to understand photosynthesis to build better, more efficient solar cells.

Even with progress in cleaner energy sources, fossil fuel-fired power plants will continue to generate 69 percent of American electricity in 2035, according to the U.S. Energy Information Administration.

Computation is key to developing cleaner fossil-fuel combustion, such as coal gasification and high-efficiency, low pollutant industrial burners.

Atom-by-atom simulations, freely available models, and open databases speed the search for the new materials we need for efficient fuel cells, longer-lasting batteries, and high-tech membranes that can separate carbon dioxide and other harmful gases out of fuel plant exhaust.

In 2012, the President emphasized the importance of scientific computing, databases, and new algorithms to the race for these new materials when he announced the $100 million Materials Genome Initiative.

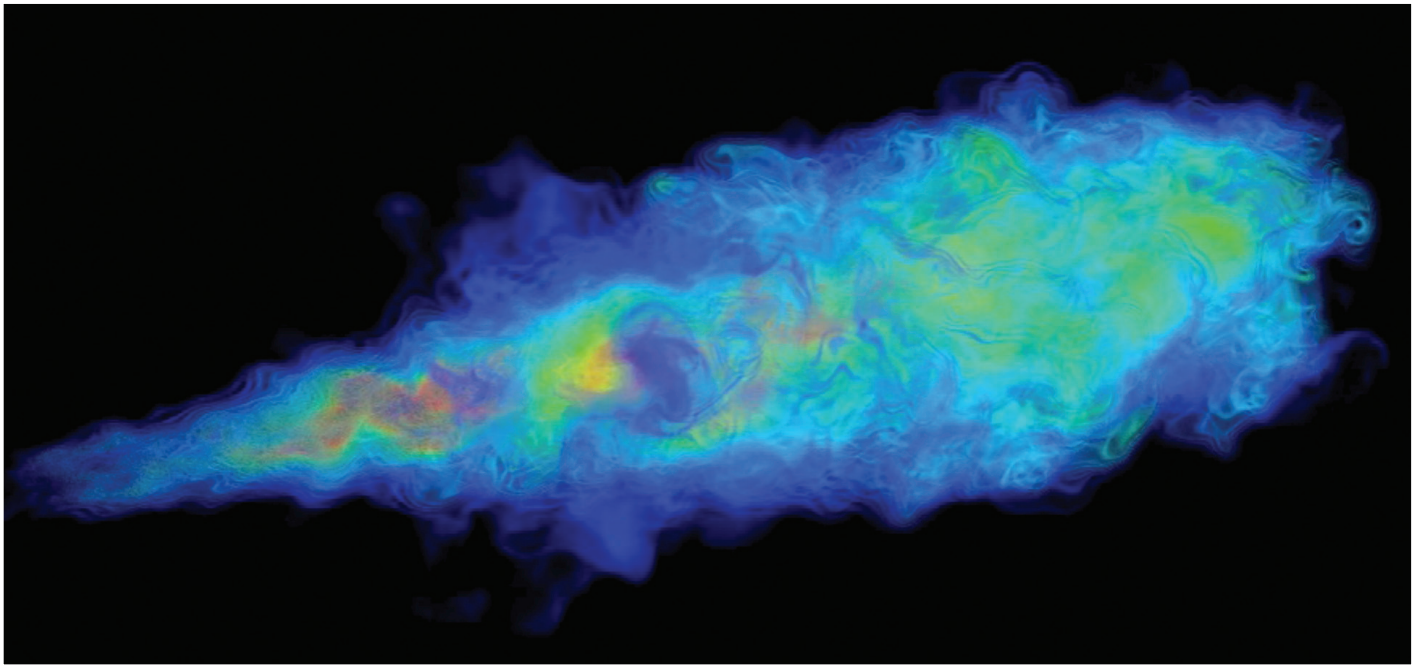

With coal still providing half of the nation's electricity today, researchers are investigating cleaner methods for extracting energy from fossil fuels, like gasification.

Unlike a traditional coal-fired boiler, a gasifier uses a little oxygen and a lot of pressure to kick-start a chemical chain reaction that results in "syngas." Once stripped of pollutants, syngas is primarily composed of hydrogen and carbon monoxide, which can be cleanly combusted in existing gas turbines to generate electricity. Scientists at the University of Utah have developed tools to model the complex processes of coal gasification, as well as methods to validate these models. A detailed understanding of coal gasification will help engineers improve existing plants' efficiency and improve the design of future plants.

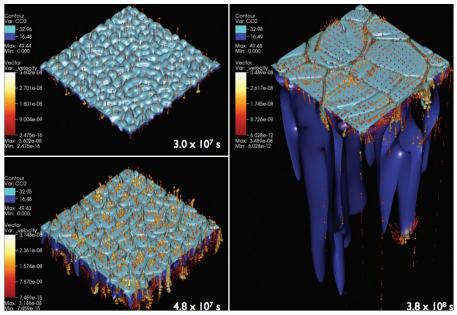

One idea for keeping excess carbon dioxide (CO^2) out of the air is to capture it from fossil-fuel burning plants and pump it into saltwater reservoirs deep underground.

To understand the long-term fate of CO^2 in such "geologic sequestration" sites, researchers are turning to simulations like this one, created at Lawrence Berkeley National Laboratory. This 3D model shows how some CO^2 mixes with brine over time creating long fingers that begin sinking. Unlike earlier models, this one uses a modeling technique that focuses computing power on the most active portions of the simulation, yielding more detailed results.

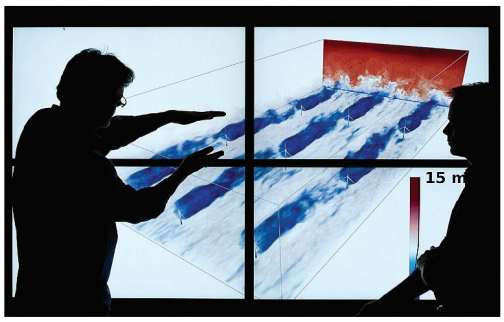

New computer simulations are able to predict the performance of large wind power plants with greater accuracy.

Design tools are generally effective for basic optimization of wind farm layout, but they can't simulate with consistent accuracy how wakes propagate or how wind turbines interact with one another to impact efficiency. Researchers at the National Renewable Energy Laboratory are creating more sophisticated models that look at a wind farm as a total system rather than just a collection of wind turbines.

"American economic competitiveness, environmental stewardship, and enhanced security depend on picking up the pace of energy technology innovation in this decade."

— The President's Council of Advisors on Science and Technology, *Accelerating the Pace of Change in Energy Technologies Through an Integrated Federal Energy Policy*, 2010

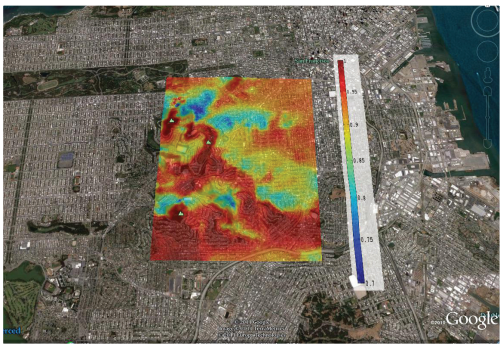

Utility companies and homeowners need to get a realistic picture of the solar power output they can typically expect and how that's affected by hills and valleys.

Researchers at the University of California, San Diego, used measurements from NASA's Shuttle Radar Topography Mission to predict how changes in elevation impact power output in California's solar grid. Jan Kleissl and Juan Luis Bosch built the model at the request of the California Public Utilities Commission. The results were integrated into commercial software for estimating site-specific solar energy potential.

Health & Medicine

"[T]he major game-changing advances in the biomedical sciences and drug development will not occur on a short time scale without the extreme use of supercomputing."

— Colin Hill, CEO and president of GNS Healthcare quoted in Council on Competitiveness' Case Study, *Bringing the Power of HPC to Drug Discovery and the Delivery of 'Smarter' Health Care*

As computer simulations and models grow faster and more accurate, so does our ability to understand and treat disease and injuries.

From drug discovery to surgery, to unraveling the mysteries of Parkinson's and Alzheimer's diseases, many advances in medicine would not be possible without the sophisticated modeling, visualization, and data analysis enabled by scientific computation.

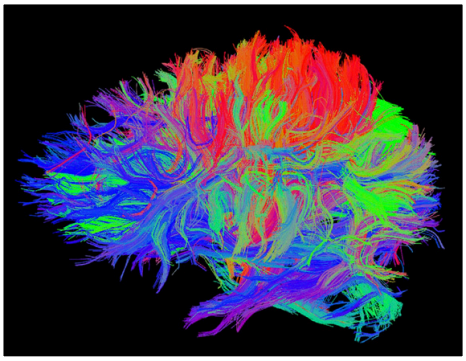

A computer model could help confirm a diagnosis of dementia and allow neurologists to predict a patient's cognitive state at any given point in the future.

"They could tell whether and when the patient will develop speech impediments, memory loss, behavioral peculiarities, and so on," said Ashish Raj of Cornell University. His model builds on recent evidence that diseased proteins involved in neurodegenerative illnesses travel along connected neural fiber tracts, rather than being transmitted by proximity like most diseases. Using the same kinds of calculations applied to the diffusion of gases, Raj was able to correctly predict the patterns of degeneration that result in a number of different forms of dementia. In this image, colors represent the fibers involved in different neuronal tracts.

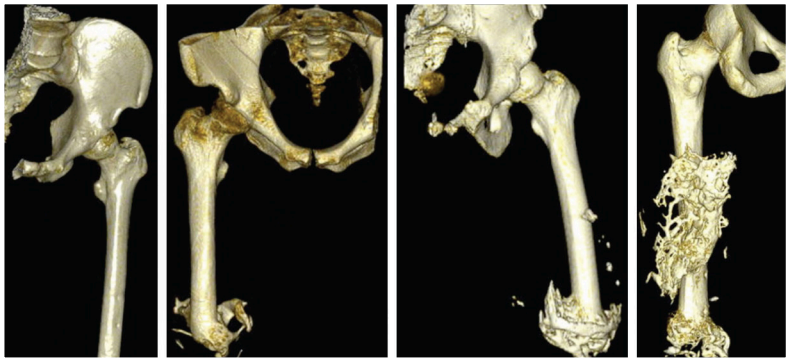

For many service members, the trauma of losing a leg is compounded by bone growths that make wearing a traditional prosthetic too painful, if not impossible.

At the University of Utah's Scientific Computing and Imaging (SCI) Institute, researchers are collaborating with biomedical engineers and orthopedic specialists to speed the healing time for implantable prosthetic legs. "Osseointegration" involves attaching a rod to an amputee's remaining bone as a base for a prosthetic. The technique, already used in Europe, is rare because it takes 18 months for bone to grow into and stabilize the rod. Doctors wondered if electrical stimulation, used to promote healing of bone fractures, could also speed the attachment process. Starting with CT scans, Brad Isaacson of the VA Medical Center, Salt Lake City, used SCI Institute software to create individual, 3D models of the residual legs of 11 servicemen: bones, muscles, fat and skin. Researchers used the models to best place an electrical stimulation device and to understand how extra-skeletal bone growths affected its use. Once osseointegration is approved for use in the U.S., patient-specific computer models like these developed at the SCI Institute could guide its use.

Researchers now routinely screen millions of potential drug molecules using computer models and algorithms to identify the most promising candidates for further testing. In fact, the first HIV drug approved to attack a key enzyme called integrase got its start *in silico* (Latin for "in silicon" versus *in vivo*, for "in life").

Simulations test drug molecules looking for those that "fit" an active region in a target protein. Given enough computing power, the same approach could also screen out molecules strongly attracted to the wrong targets, which can be toxic. It's akin to trying millions of keys to find the ones that can open a lock, but without setting off a booby trap.

Yet, seemingly good *in silico* candidates can fail in the lab. As supercomputers grow more powerful, simulations can become more detailed and accurate.

At the University of Utah's Scientific Computing and Imaging (SCI) Institute, researchers are creating patient-specific computerized models of human bodies. Their work is being used to diagnose and treat heart disease and to explore replacing the missing legs of wounded service members with implanted prosthetics.

Likewise, "molecular dynamics" simulations are revealing the complex biological processes that underpin drug-resistant influenza, antibiotic resistance, drug toxicity and devastating diseases such as Alzheimer's, Parkinson's, and schizophrenia.

Combined with enough computing power and genetic information, these breakthroughs will one day form the basis for individualized medical care.

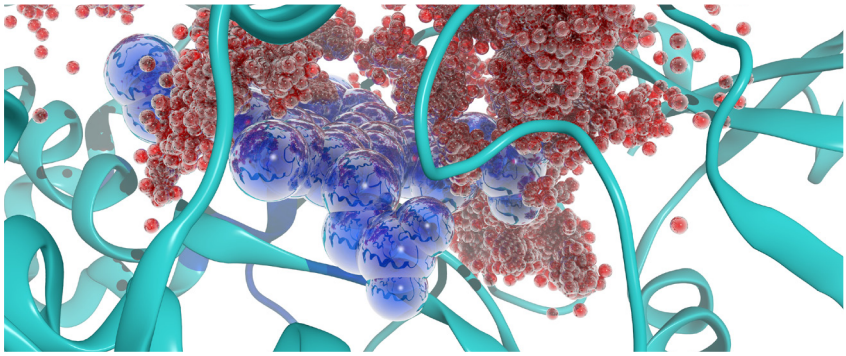

Why do people react differently to the same drug?

Supercomputer simulations are giving scientists unprecedented access to a key class of proteins involved in drug detoxification, and a window into one mechanism that may help answer that question. Jerome Baudry and Yinglong Miao, both jointly affiliated with the University of Tennessee and Oak Ridge National Laboratory, simulated the motions of water molecules in a class of enzymes called P450s. One of these enzymes is responsible for detoxifying many drugs in humans. Simulating 0.3 microseconds, their work reconciled a long-standing discrepancy between theory and observation and contributed to a broader understanding of drug processing by these enzymes. It's possible that mutations in P450 could effect water movement, and contribute to the differences in drug processing between individuals. Shown here are water molecules (red) moving into and out of the active site (blue) of a P450 enzyme.

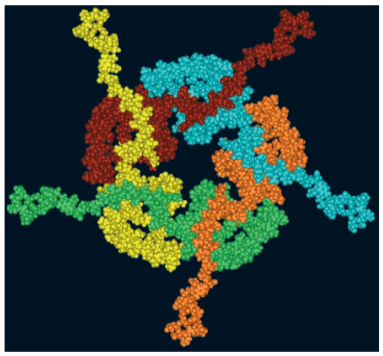

New evidence has lent support to a once-marginalized theory about the cause of Parkinson's.

University of California, San Diego, researchers used 3D molecular-level simulations with lab work to provide a step-by-step explanation of how alpha-synuclein, a "protein-run-amok," forms ring-like structure that punch holes in the membranes of brain neurons causing the symptoms of Parkinson's disease. "We think we can create drugs to give us an anti-Parkinson's effect by slowing the formation and growth of these ring structures," said Igor Tsigelny, lead author of the study and a research scientist at UCSD's San Diego Supercomputer Center and Department of Neurosciences. The new results conflict with an older theory that long strings of misfolded proteins called amyloids cause neurogdegenerative diseases, including Parkinson's and Alzheimer's.

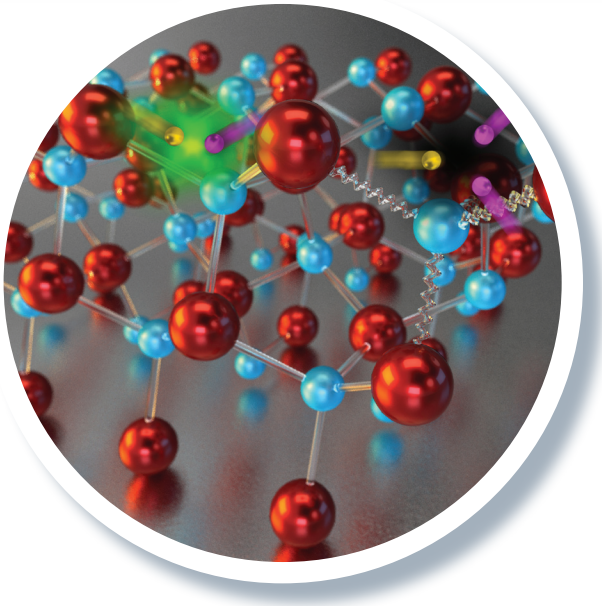

Despite being cool, ultra-efficient, and long lasting, the light-emitting diode (LED) has yet to conquer the general lighting market because of "efficiency droop."

At low power, the blue LEDs used in white lighting convert a big fraction of electricity into photons of light, producing almost no waste heat. But when researchers pump up the power to levels that could light a room, efficiency takes a nosedive. To help understand the cause of this "droop," researchers from the University of California, Santa Barbara, simulated the physics of the silicon wafer at the heart of this kind of LED. Carried out at the National Energy Research Scientific Computer Center (NERSC) at Lawrence Berkeley National Laboratory, the simulations showed that at higher power, energy was "leaking" out of the chip through a process known as auger recombination (shown as a dark spot on the crystal lattice in an artist's illustration at left). Their work is helping engineers develop a new generation of high-performance, energy-efficient lighting that could save the nation 1488 terawatt-hours of electricity (worth about $120 billion) and spare an estimated 246 million metric tons of greenhouse gas emissions over 20 years.

"Engineering innovation in almost every discipline has been revolutionized through the use of virtual models…leading to dramatic cost reductions, …reduced design cycles times, and…more competitive designs."

– National Science Foundation Advisory Committee for Cyberinfrastructure Task Force on Grand Challenges, *Final Report*, 2011

Industry & Innovation

Just as innovation is the lifeblood of U.S. industry, so is computation increasingly vital to innovation. Computer simulations, data analysis, and modeling are speeding up the design and adoption of new materials, machines, and manufacturing processes.

Many CASC members are engaged in public-private partnerships to advance U.S. industrial innovation and competitiveness. And even those who may not have ties to specific companies support industry with research that leads to new products and services.

Simulation has been used to analyze the performance and design of aircraft components since the early days of computing. Today, computational simulation is an essential tool in the complete design process not only for aircraft, but also for automobiles, manufacturing robots, and other machines.

Computer-aided design drastically reduces the need to construct and test prototypes. The Boeing 777, for example, was the first jetliner to be completely digitally designed. The airplane was preassembled by computer, eliminating the need for costly full-scale mock-ups.

Recently, a group of 60 universities and companies led by the University of California, Los Angeles, formed the Smart Manufacturing Leadership Coalition, a public-private partnership to advance U.S. manufacturing. While individual systems within factories are smarter, more instrumented, and more integrated, this group is working to marry systems more closely together with simulation and modeling to make American manufacturing smarter, more efficient, and more competitive.

The simulation of complex phenomena, such as breaking waves and spray sheets, play a key role in the design and operation of Naval combat ships.

This simulation shows the resulting turbulence as water flows past a NACA 0024 airfoil that approximates a ship's hull. A ship moving through water forms steep breaking waves, shedding spray along the crests and also near the ship's bow where thin sheets of water form. The Navy uses such simulations to learn how water interacts with naval vessels, not only to understand the propulsive power needed to overcome water resistance, but also to study the signature that the ship leaves in its wake, and to model how waves, which can break anywhere along the hull, affect ship stability. Paul Navratil and Bill Barth of the Texas Advanced Computing Center at The University of Texas at Austin, and Hank Childs of Lawrence Berkeley National Laboratory created these visualizations from simulations performed by Doug G. Dommermuth of Science Applications International Corporation.

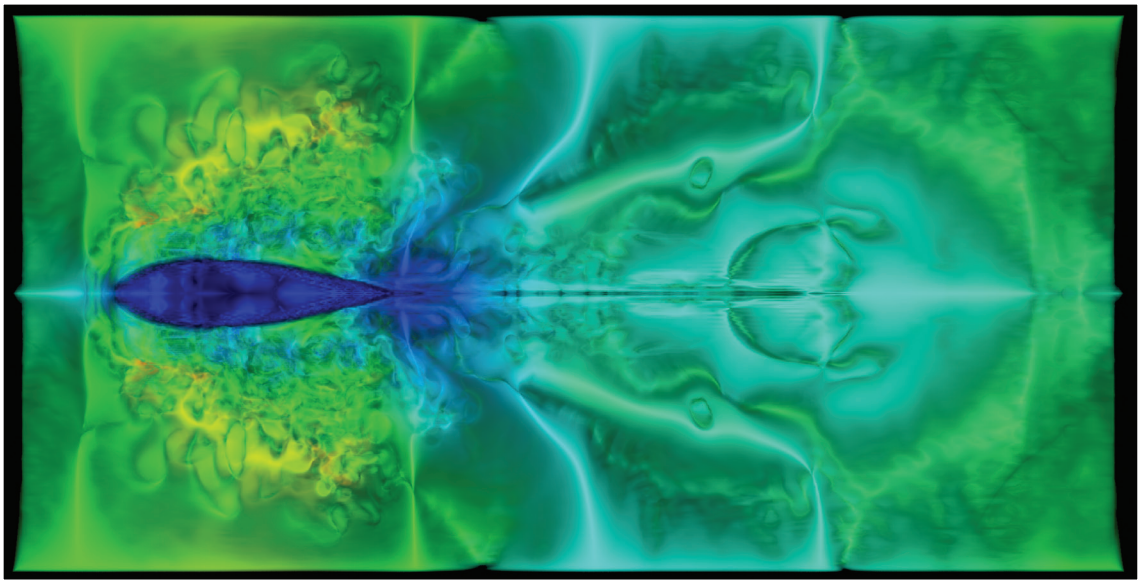

Understanding turbulent mixing noise sources for jet exhaust nozzles is critical to delivering the next generation of greener, lower-noise jet engines.

A complex large eddy simulation offers incredibly detailed, direct computation of turbulence, noise sources, and heat transfer. This work will enable the design of quieter, more efficient airfoils for wind turbines as well as nozzles for jet engines. Reducing wind turbine noise by even one decibel could translate into two to three percent more yield per turbine. Joseph A. Insley of Argonne National Laboratory and the University of Chicago Computational Institute created the visualization from simulations performed at Argonne by Umesh Paliath and Anurag Gupta of General Electric Global Research.

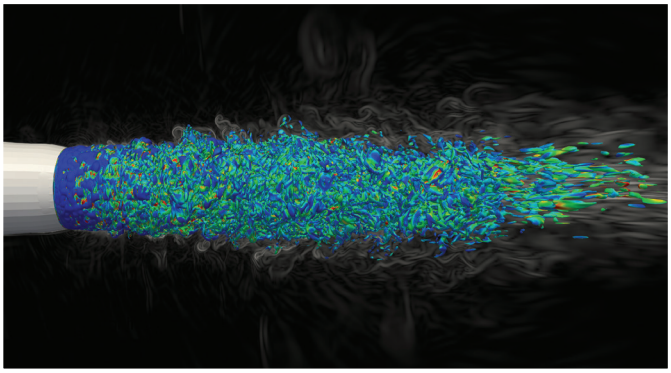

Environment & Resources

In the coming decades, changing climate and an ever-growing population will put unprecedented strains on our nation's natural resources, agriculture, and environment. Predicting and preparing for shifting temperatures, changing growing seasons and rain patterns, extreme weather events, increased drought, flooding, wildfires, and changes in plant and animal populations are among the most important roles of scientific computing.

The Center for Robust Decision Making on Climate and Energy Policy (RDCEP) at the Chicago Computation Institute integrates climate, agricultural, and economic models to help inform policy decisions.

For example, researchers there recently modeled the effects that climate change—such as the hotter, drier summers experienced in 2012—will have on U.S. corn crops.

The iPlant project at The University of Texas at Austin brings researchers together with vast databases of plant genomic information and tools that may help develop and adapt agricultural crops, among other uses.

At the University of California, Berkeley, researchers are using robotic floaters to transmit information about the Sacramento River delta into real time models that may someday help resource planners balance increasing demands and competing pressures on this waterway.

Louisiana State University's Center for Computation and Technology and UT Austin researchers are modeling the destructive water rises (storm surges) caused by tropical Atlantic tempests and passing that information on to emergency officials who rely on the information to direct evacuations and other responses.

And, of course, climate models are growing more sophisticated and more important in understanding and predicting the changes underlying all this work.

"Computer science research and information technology innovation will be critical to meeting sustainability challenges."

—National Research Council, *Computing Research for Sustainability*, 2012

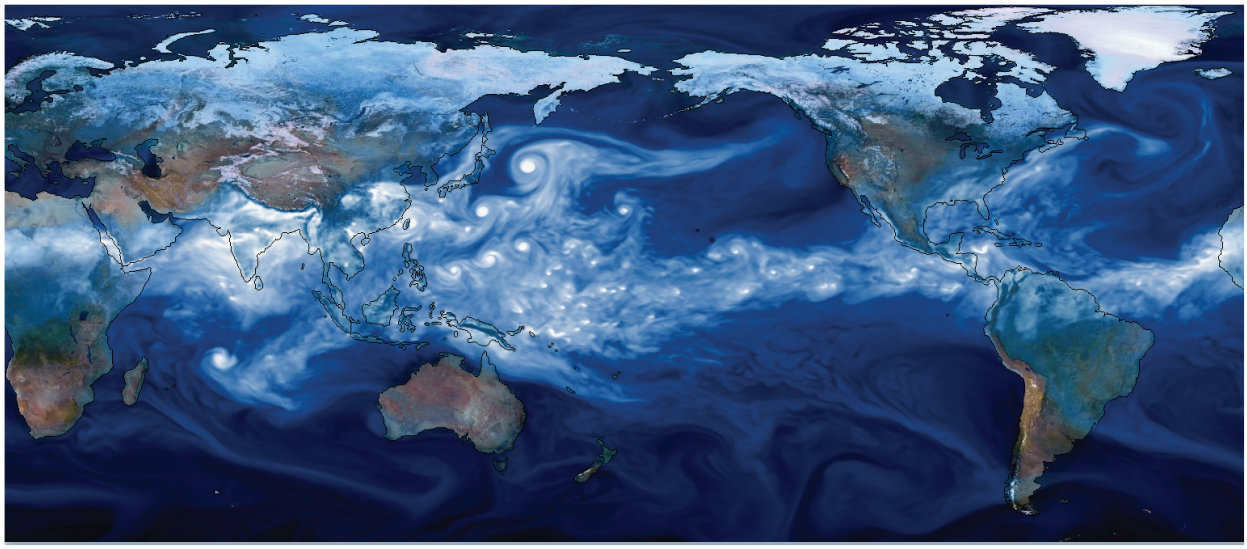

Climate models are becoming more detailed and more accurate.

This simulation was created to test a new version of climate modeling software (the National Center for Atmospheric Research's Community Atmospheric Model). Using historical weather data, Lawrence Berkeley National Laboratory researchers Michael Wehner and Prahbat found the model spontaneously spawned tropical cyclones in line with the actual storms from that period.

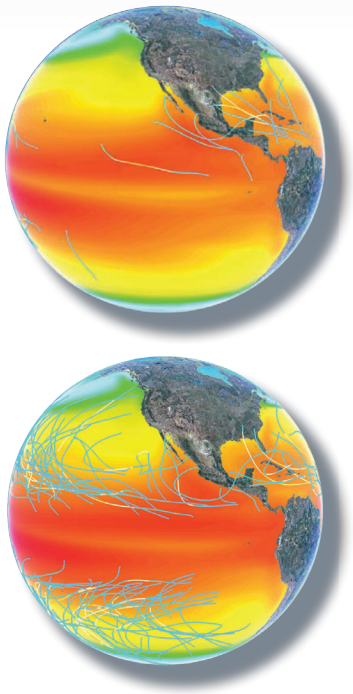

Researchers found that more frequent cyclones in Earth's ancient past fed into warmer temperatures.

During the Pliocene Epoch (3 to 5 million years ago) conditions were roughly the same as those of modern day Earth, including the amount of carbon dioxide in the air. Yet, global mean temperatures were up to 8ºF higher than today. Yale University and Massachusetts Institute of Technology researchers studying Pliocene climate models found that more frequent cyclones fed into a loop that helped keep temperatures high. How it started and whether we may experience the same isn't yet known. In this image, simulated cyclone tracks (blue) are shown overlaying modern climate (top) and Pliocene (bottom). The coloring represents sea surface temperatures where blue is coldest and red hottest.

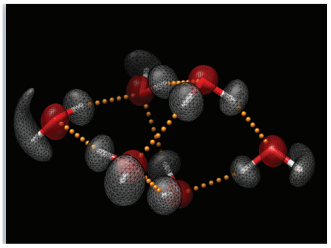

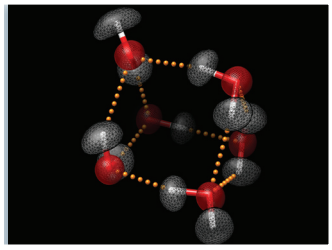

Water is vital to life on the planet, yet some of its most basic properties are just now being uncovered.

Simulations carried out at the San Diego Supercomputer Center at the University of California, San Diego, and experiments by researchers at UC San Diego and Emory University have uncovered previously unknown details about the hexamer structures that make up the tiniest droplets of water. These six molecule formations offer clues to what underlies water's wide range of properties. The research provides a new interpretation for experimental measurements as well as a vital test for future studies of our most precious resource. Understanding the properties of water at the molecular level can ultimately impact climate change research and drug development.

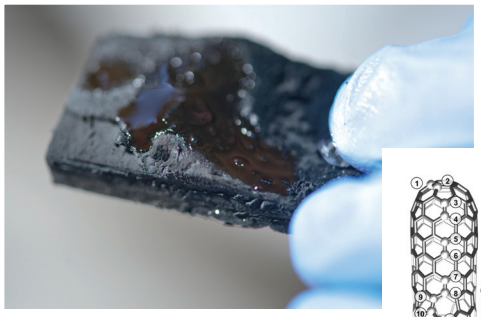

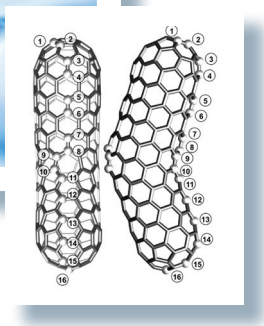

An oil-absorbing carbon nanotube sponge offers a promising material for cleaning up ocean oil spills.

Researchers used simulations to test their idea of adding boron to the process of creating carbon nanotubes, causing them to not only develop flexible bends, but to also fuse into a sponge-like material. The synthesized material worked as models predicted. It absorbs up to 100 times its weight in oil, repels water, and can be reused repeatedly. It can also be manipulated with magnets, allowing it to be steered into oil and retrieved for reuse, according to Daniel Hashim of Rice University. He worked with an international collaboration including researchers at The Pennsylvania State University, University of Illinois at Urbana-Champaign, Rensselaer Polytechnic Institute, Oak Ridge National Laboratory, and Arizona State University.

Matter & the Universe

Over the past decades, physics, astronomy, and cosmology have seen amazing advances driven by improvements in modeling, simulation, and data analysis algorithms.

The enormous and complex data sets generated by modern sky surveys and international physics experiments like the Large Hadron Collider are sifted, analyzed, stored, and fed into simulations that have become virtual experimental laboratories themselves.

In the safe confines of a simulation, it's possible to explode enormous supernovae, slow down and observe the dance of subatomic particles, or even to time travel, rewinding and replaying the evolution of our universe from the Big Bang through today and into the distant future.

Simulations also offer the only way for humans to perceive the stuff that makes up most of the universe: dark matter.

Thanks to a high-speed data pipeline, image analysis software and machine learning algorithms, the Palomar Transient Factory identified two supernova firsts in one year: an early stage supernova and a "progenitor" star system that turned into a supernova

Both of these finds led astronomers to train their telescopes on supernovae as they happened, rather than catching them at their fullest brightness, if at all.

"Large-scale simulations, vetted with observations from the world's most powerful telescopes, will advance new understanding of how the universe formed and evolved."

– The President's Council of Advisors on Science and Technology, *Designing a Digital Future: Federally Funded Research and Development in Networking and Information Technology*, 2010

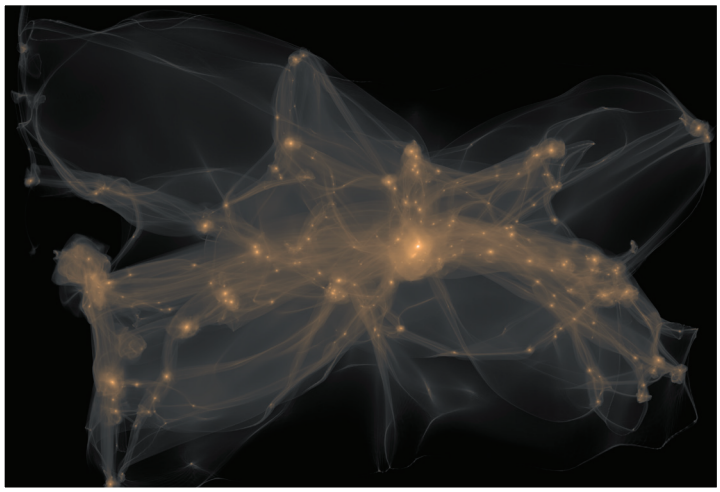

A new visualization technique offers a stunningly detailed look at the distribution of dark matter in the universe.

Dark matter makes up about 85 percent of our universe, but because it doesn't interact with light, we can't see it. Computer simulations and visualizations help humans to envision this important, but invisible stuff. Using a new technique developed at Stanford University's Kavli Institute for Particle Astrophysics and Cosmology (KIPAC) and the SLAC National Accelerator Laboratory, Ralf Kaehler, Oliver Hahn, and Tom Abel created this model to more accurately represent how dark matter forms in diaphanous sheets and concentrated clumps.

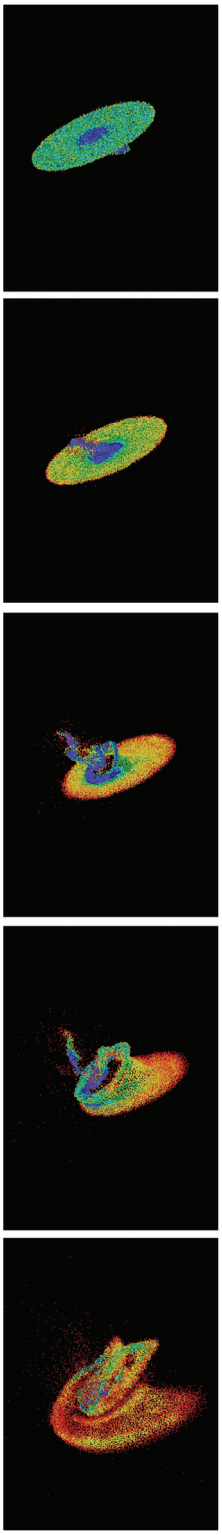

The Milky Way may have an invisible galactic neighbor lurking on the opposite side of the galaxy from Earth.

Sukanya Chakrabarti developed a mathematical method to uncover these "dark" satellites as part of her doctoral research at the University of California, Berkeley. Using disruptions in the cold hydrogen gas that extends far outside the visible galactic disk, her calculations revealed known satellites of the neighboring Whirlpool galaxy. Applied to our own Milky Way, they predict a dwarf satellite galaxy that would be obscured from our view by the dust and gas of the galactic plane. Her calculations could also be used to help astronomers detect the presence of galaxies made mostly of dark matter, which doesn't reflect light and can't be directly observed. Lawrence Berkeley National Laboratory's Prabhat translated her scientific data into an animated, 3D visualization that shows the effect of a satellite galaxy moving through the hydrogen disk of its larger host.

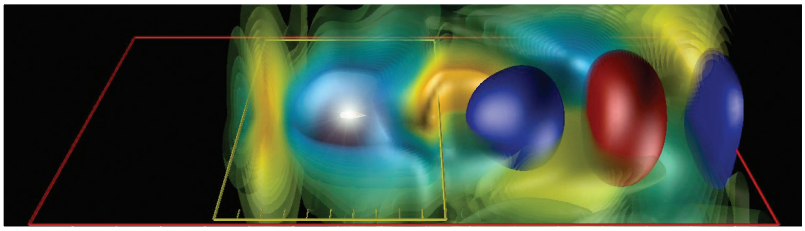

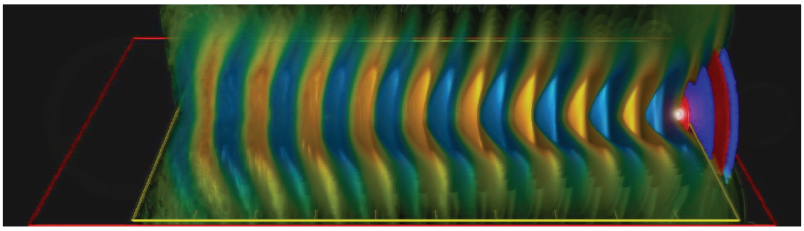

A breakthrough in computer modeling will hasten the design of smaller, more powerful particle accelerators.

Traditional accelerators, like the Large Hadron Collider, use magnets to boost particles to near-lightspeed along miles-long underground tunnels. Laser wakefield accelerators use powerful laser pulses to create plasma waves on which electrons could theoretically surf to LHC-class speeds, but in about 100 meters or less. Computer simulations play a critical role in their design, but the calculations take a long time. To speed the process, researchers at Lawrence Berkeley National Laboratory are creating simulations that fly into the oncoming waves at nearly the speed of light. This "moving frame" requires far fewer calculations than a static frame, speeding up calculations as much as 1 million fold. In the images above, the laser pulses are blue and red; the wakefields are pale blue and yellow. A stationary point of view (bottom) has many wavefronts per frame (yellow line). From a moving point of view (top), the frame contracts and the wave fronts are fewer and farther apart, resulting in far fewer (albeit more complex) calculations per frame.

Education & Outreach

In the 21st century, prosperity and innovation are inextricably linked with educating new generations in science, technology, engineering, and mathematics, or STEM.

Between 50 and 85 percent of U.S. gross domestic product growth over the last 50 years has been due to advancements in science and engineering, according to a 2010 National Research Council report on U.S. education, "Rising Above the Gathering Storm, Revisited: Rapidly Approaching Category 5." Yet, the U.S. ranks 15th in science literacy and 28th in mathematics literacy among top students.

CASC members are committed to reversing this trend, not only for the economic wellbeing of the nation, but because the most daunting problems facing the U.S.—energy, healthcare, climate change, and national security—will require the expertise of top STEM professionals.

Our members educate and inspire K-12 students and their teachers through workshops, seminars, summer camps, and online training and tools. CASC member institutions also reach out to undergraduates and graduates, offering mentoring, internships, and fellowships. Many of their efforts focus on promising, yet underrepresented pools of untapped American talent. Many centers also extend their reach to the general public, increasing STEM literacy among citizens and policy makers.

"STEM education will determine whether the U.S. will remain a leader among nations…."

– The President's Council of Advisors on Science and Technology, *Prepare and Inspire: K-12 Education in Science, Technology, Engineeing and Math for America's Future*, 2010

Girls tend to drift away from the sciences in the critical middle school years, and most never find their way back.

Whether they get Barbie's famous message that "math is hard," or simply lose interest for lack of encouragement, our nation loses sorely needed talent.

To motivate and engage girls in science, technology, engineering and math, Louisiana State University's Center for Computation & Technology (CCT) sponsors "Alice in Computation Land," a summer camp for girls in grades 6-8. Camp participants simulate a forest fire, model a population of frogs, and learn what it takes to create animation, games, and video. Most of all, they have fun doing it!

ALICE IN COMPUTATIONLAND

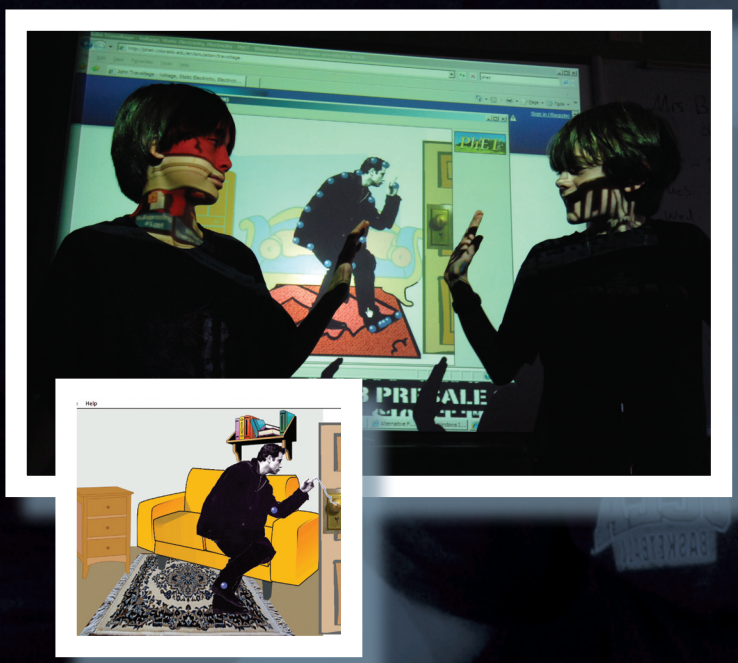

How do you capture the attention of students more interested in video games than in science and math? Try a video game.

The PhET Interactive Simulations project at the University of Colorado Boulder has built over 100 freely available, interactive simulations that teach math and science in an intuitive, game-like environment. Based on educational research, PhET simulations help students of all ages connect classroom teaching to real-world phenomena. To learn about static electricity, for example, scrub the foot of "John Travoltage" across a rug and experiment with build up, flow, and discharge. The project has served up over 70 million simulations since its founding in 2006. The Tech Museum in San Jose, California named PhET a laureate of The Tech Awards 2011 for applying technology to benefit humanity and spark global change.

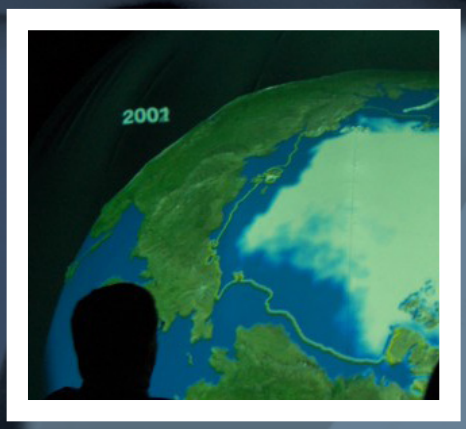

Adults need educating, too.

The Renaissance Computing Institute (RENCI) reaches out to decision-makers and the general public with models, simulations, and other technological tools and expertise that address everything from climate change to disaster preparedness and economic development. The RENCI at UNC Asheville Engagement Site offers a 16-foot visualization wall and an immersive GeoDome that showcase projects such as the Western North Carolina Vitality Index, a web-based tool covering over 160 metrics for 27 counties in the area; Forewarn, a web-based system to track unexpected forest change and potential threats; GroWNC, a listening and planning process to develop strategies for economic prosperity and sustainable development; and geospatial climate information viewers.

2013 CASC MEMBERSHIP

Arizona State University
Advanced Computing Center
Tempe, Arizona

Boston University
Center for Computational Science
Boston, Massachusetts

California Institute of Technology
Center for Advanced Computing Research
Pasadena, California

Carnegie-Mellon University & University of Pittsburgh
Pittsburgh Supercomputing Center
Pittsburgh, Pennsylvania

Case Western Reserve University
Core Facility in Advanced Research Computing
Cleveland, Ohio

City University of New York
High Performance Computing Center
Staten Island, New York

Clemson University
Computing and Information Technology (CCIT)
Clemson, South Carolina

Columbia University
Core Research Computing Facility
New York, New York

Cornell University
Center for Advanced Computing
Ithaca, New York

Florida State University
Department of Scientific Computing
Tallahassee, Florida

George Washington University
Institute for Massively Parallel Applications
Washington, District of Columbia

Georgia Institute of Technology
Atlanta, Georgia

Indiana University
Bloomington, Indiana

Lawrence Berkeley National Laboratory
Berkeley, California

Louisiana State University
Center for Computation & Technology (CCT)
Baton Rouge, Louisiana

Michigan State University
High Performance Computing Center
East Lansing, Michigan

Mississippi State University
High Performance Computing Collaboratory (HPC^2)
Mississippi State, Mississippi

Mount Sinai School of Medicine
New York, New York

National Center for Atmospheric Research (NCAR)
Boulder, Colorado

New York University
New York, New York

North Dakota State University
Center for Computationally Assisted
Science & Technology
Fargo, North Dakota

Northeastern University
Advanced Scientific Computing Center (ASCC)
Boston, Massachusetts

Northwestern University
Evanston, Illinois

NYU Langone Medical Center
New York, New York

Oak Ridge National Laboratory (ORNL)
Center for Computational Sciences
Oak Ridge, Tennessee

Pacific Northwest National Laboratory (PNNL)
Richland, Washington

Princeton University
Princeton, New Jersey

Purdue University
West Lafayette, Indiana

Rensselaer Polytechnic Institute
Scientific Computation Research Center (SCOREC)
Troy, New York

Rice University
Ken Kennedy Institute
for Information Technology (K^{2}I)
Houston, Texas

Rutgers University
Discovery Informatics Institute (RDI^2)
Piscataway, New Jersey

Stanford University
Stanford, California

Texas A&M University
Institute for Scientific Computation
College Station, Texas

Texas Tech University
High Performance Computing Center
Lubbock, Texas

The Ohio State University
Ohio Supercomputer Center (OSC)
Columbus, Ohio

The Pennsylvania State University
University Park, Pennsylvania

The University of Texas at Austin
Texas Advanced Computing Center (TACC)
Austin, Texas

University at Buffalo
Center for Computational Research
Buffalo, New York

University of Alaska Fairbanks
Arctic Region Supercomputing Center (ARSC)
Fairbanks, Alaska

University of Arkansas
High Performance Computing Center
Fayetteville, Arkansas

University of California, Los Angeles
Institute for Digital Research and Education
Los Angeles, California

University of California, San Diego
San Diego Supercomputer Center (SDSC)
San Diego, California

University of Chicago & Argonne National Laboratory
Computation Institute
Chicago, Illinois

University of Colorado Boulder
Boulder, Colorado

University of Florida
Gainesville, Florida

University of Georgia
Advanced Computing Resource Center (GACRC)
Athens, Georgia

University of Hawaii
Honolulu, Hawaii

University of Houston
Texas Learning and Computation Center
Houston, Texas

University of Illinois at Urbana-Champaign
National Center for Supercomputing Applications (NCSA)
Champaign, Illinois

University of Iowa
Iowa City, Iowa

University of Kentucky
Center for Computational Sciences
Lexington, Kentucky

University of Louisville
Louisville, Kentucky

University of Maryland
College Park, Maryland

University of Miami
Miami, Florida

University of Michigan
Center for Advanced Computing
Ann Arbor, Michigan

University of Nebraska
Omaha, Nebraska

University of Nevada
National Supercomputing Center for Energy
and the Environment (NSCEE)
Las Vegas, Nevada

University of New Mexico
Center for Advanced Research Computing
Albuquerque, New Mexico

University of North Carolina at Chapel Hill
Renaissance Computing Institute (RENCI)
Chapel Hill, North Carolina

University of Notre Dame
Center for Research Computing
Notre Dame, Indiana

University of Oklahoma
Supercomputing Center
for Education and Research
Norman, Oklahoma

University of South Florida
Research Computing
Tampa, Florida

University of Southern California
Information Sciences Institute
Marina del Rey, California

University of Tennessee
National Institute for Computational Sciences (NICS)
Knoxville, Tennessee

University of Utah
Center for High Performance Computing
Salt Lake City, Utah

University of Virginia
Alliance for Computational Science and Engineering
Charlottesville, Virginia

University of Wisconsin-Madison
Madison, Wisconsin

University of Wisconsin-Milwaukee
Milwaukee, Wisconsin

University of Wyoming
Advanced Research Computing Center (ARCC)
Laramie, Wyoming

Virginia Tech University
Advanced Research Computing
Blacksburg, Virginia

Yale University
New Haven, Connecticut

CASC

Coalition for Academic Scientific Computation

WWW.CASC.ORG

SUE FRATKIN, WASHINGTON LIAISON • PHONE 202.265.5410 • FAX: 202.332.8538 • EMAIL sue@casc.org