November 2006

Science & Technology Review

National Nuclear
Security Administration's
Lawrence Livermore
National Laboratory

Power of Thunder

Also in this issue:

- **Extracting Objects from Overhead Images**
- **Oxygen Proves Key to Evolution**
- **Shock Creates Laserlike Light**

About the Cover

As the article on p. 4 describes, Lawrence Livermore's Thunder supercomputer, funded by the Laboratory's Multiprogrammatic and Institutional Computing Program, has been used over the past year for five "Grand Challenge" projects. The unclassified, mission-relevant projects are allocated large banks of computing cycles to aggressively tackle problems for breakthrough science. Brought on line in 2004, Thunder is a 23-trillion-floating-point-operations-per-second cluster of 4,096 processors and runs on an open-source software environment.

Cover design: Amy Henke

About the Review

Lawrence Livermore National Laboratory is operated by the University of California for the Department of Energy's National Nuclear Security Administration. At Livermore, we focus science and technology on ensuring our nation's security. We also apply that expertise to solve other important national problems in energy, bioscience, and the environment. *Science & Technology Review* is published 10 times a year to communicate, to a broad audience, the Laboratory's scientific and technological accomplishments in fulfilling its primary missions. The publication's goal is to help readers understand these accomplishments and appreciate their value to the individual citizen, the nation, and the world.

Please address any correspondence (including name and address changes) to *S&TR*, Mail Stop L-664, Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, California 94551, or telephone (925) 423-3432. Our e-mail address is str-mail@llnl.gov. *S&TR* is available on the World Wide Web at www.llnl.gov/str.

Prepared by LLNL under Contract
No. W-7405-Eng-48

Science & Technology Review

November 2006

Lawrence Livermore National Laboratory

Contents

Features

3 **Expanded Supercomputing Maximizes Scientific Discovery**
Commentary by Dona Crawford

4 **Thunder's Power Delivers Breakthrough Science**
Livermore's Thunder supercomputer allows researchers to model systems at scales never before possible.

12 **Extracting Key Content from Images**
A new system called the Image Content Engine is helping analysts find significant but hard-to-recognize details in overhead images.

Research Highlights

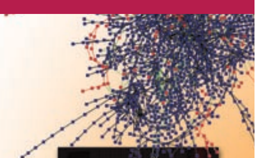

20 **Got Oxygen?**
Oxygen, especially oxygen metabolism, was key to evolution, and a Livermore project helps find out why.

22 **A Shocking New Form of Laserlike Light**
According to research at Livermore, smashing a crystal with a shock wave can result in coherent light.

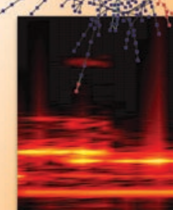

Departments

2 The Laboratory in the News

24 Patents and Awards

25 Abstracts

S&TR Staff

Scientific Editor
Harry B. Radousky

Managing Editor
Ray Marazzi

Publication Editor
Pamela MacGregor

Writers
Arnie Heller, Ann Parker, Gabriele Rennie, Maurina Sherman, and Katie Walter

Art Director
Kitty Madison

Designers
Amy Henke and Kitty Madison

Illustrator
Mark McDaniel

Proofreader
Carolin Middleton

S&TR, a Director's Office publication, is produced by the Technical Information Department under the direction of the Office of Policy, Planning, and Special Studies.

S&TR is available on the Web at www.llnl.gov/str.

Printed in the United States of America

Available from
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Road
Springfield, Virginia 22161

UCRL-TR-52000-06-11
Distribution Category UC-99
November 2006

Simulation improves understanding of turbulence

Using Livermore's BlueGene/L supercomputer, Laboratory scientists have discovered a turbulence Reynolds number effect that could accelerate flames in type Ia supernovae at rates exceeding those predicted by standard numerical simulations. The effect derives from the turbulence energy budget (the conversion of potential energy to kinetic energy) associated with unstable Rayleigh–Taylor flame fronts. Spontaneous mixing of fluids at unstable interfaces occurs in a variety of atmospheric, oceanic, geophysical, and astrophysical flows. Rayleigh–Taylor instability is a process that occurs when a light fluid pushes on a heavy fluid and the fluids seek to reduce their combined potential energy. This process plays a key role in all known types of fusion.

The simulation carried out by Andrew Cook and Bill Cabot of Livermore's Defense and Nuclear Technologies Directorate is a computation of evolving heavy fluid on top of a lighter fluid, subject to a gravitational field. This Rayleigh–Taylor simulation achieved a Reynolds number of 32,000 and is the first to reach past the mixing transition while resolving all hydrodynamic scales of motion. The findings, which appeared in the August 2006 issue of *Nature Physics*, have important implications for the study of turbulent combustion in supernovae, specifically, and for the understanding of turbulent convection in general.

Contact: Andrew Cook (925) 423-2856 (cook33@llnl.gov) or Bill Cabot (925) 423-9272 (cabot1@llnl.gov).

Research links warming sea surfaces to human influence

Scientists from the Laboratory's Program for Climate Model Diagnosis and Intercomparison, along with collaborators from 10 other research centers, have shown that rising sea surface temperatures (SSTs) in the hurricane breeding grounds of the Atlantic and Pacific oceans are likely to be directly linked to human activities. The team, whose findings appeared in the September 19, 2006, issue of the *Proceedings of the National Academy of Sciences*, used 22 climate models to estimate the magnitude of SST changes arising from natural processes and external forcing (human influence) from 1906 to 2005.

The researchers found an 84 percent likelihood that external forcing explains at least 67 percent of observed SST increases in the Atlantic and Pacific cyclogenic regions during the period studied. Previous research has associated the warming of SSTs with an increase in hurricane intensity. The research team used nearly all the world's climate models to estimate how the climate of an "undisturbed Earth" may have evolved. Comparing this undisturbed scenario with current global conditions, researchers posit that natural processes alone cannot explain the observed SST increases in the hurricane formation regions. They also note that although rising SSTs are not the sole cause of increased hurricane intensity, these increases are likely to be one of the most important influences on hurricane strength.

Contact: Benjamin Santer (925) 422-2486 (santer1@llnl.gov).

Expanded Supercomputing Maximizes Scientific Discovery

HIGH-PERFORMANCE computing has been a part of the Laboratory's genetic code since it was founded 54 years ago by physicists Ernest O. Lawrence and Edward Teller. Although computer simulation was in its infancy at the time, our Laboratory leaders recognized its potential to accelerate scientific discovery and engineering underpinning Livermore's national security missions. Our ability to use computers for simulating the physical world has evolved to the point where today three-dimensional simulation is an integrating component of the scientific method of observation, theory, and experiment. Livermore continues to develop increasingly more powerful, scientific computing tools and techniques that enable scientific discovery and the fulfillment of Laboratory missions.

Strategic leadership combined with a willingness to explore new directions in scientific computing over the years have established the Laboratory as one of the world's premier institutions for computer and simulation science. By leveraging the capabilities of the Advanced Simulation and Computing (ASC) Program, an important component of stockpile stewardship, our Multiprogrammatic and Institutional Computing (M&IC) Program has extended leading-edge simulation on scalable platforms to unclassified science in the national interest. M&IC works to balance investments in an effort to provide cost-effective capability platforms and meet the ever-growing demand for high-performance computing capacity from the broad cross section of Laboratory researchers.

Livermore's Thunder supercomputer is a shining example of our success with this high-performance computing strategy. Brought on line in 2004, Thunder is a 23-trillion-floating-point-operations-per-second (teraflops) cluster of 4,096 processors. When it debuted, Thunder was the world's second fastest supercomputer on the Top500 list according to the industry standard LINPACK benchmark. Two years later, Thunder still holds a respectable ranking of 14 in the highly competitive and rapidly developing supercomputing environment. Systems such as Thunder represent a shift from traditional parallel computers to Linux cluster technology, which offers increased power at a lower cost. M&IC's versatile systems are a product of the Laboratory's consistent investment, and the result is a powerful unclassified computing resource that is being used to push the limits of computing and its application to simulation science.

Over the last year, the M&IC Program has succeeded in using Thunder to run "Grand Challenge" projects—a select number of unclassified, mission-relevant projects that are allocated large banks of computing cycles to aggressively tackle problems for breakthrough science. These projects, as described in the article beginning on p. 4, range from research in climate change to the interface of water and vapor to nanofluidics to dislocation dynamics to protein folding. All of these projects involve collaborations with other leading research institutions such as the Scripps Institution of Oceanography, the National Center for Atmospheric Research, the IBM Almaden Research Center, the University of California campuses at San Diego, Berkeley, Davis, and Irvine, the University of Zurich, and the Polytechnic of Torino. Our capabilities and knowledge make us a desirable partner for collaborations that can accelerate scientific discovery to benefit both the Laboratory and the larger scientific community.

While Livermore's high-end classified systems BlueGene/L and ASC Purple attract the spotlight because of their respective number one and three world rankings on the LINPACK benchmark, systems such as Thunder are sometimes used to lay the scientific computing groundwork that is later transferred to the more powerful machines. Thunder is also paving the way for the next generation of supercomputers recently acquired as part of the Peloton procurement, which is bringing to the Laboratory an additional three clusters named Atlas, Zeus, and Rhea for a total of 77 teraflops. Atlas will take on M&IC's heavy lifting as the 44-teraflops capability machine for Grand Challenge projects starting in January.

The environment created to develop and sustain the ASC supercomputers for simulating nuclear weapons performance is also fostering the development of increasingly powerful unclassified systems for conducting complementary science. By ensuring that classified and unclassified computing work in complement, we've greatly expanded the high-performance computing resources available to Laboratory researchers in a broad range of disciplines. This symbiosis is another facet of the multidisciplinary team science that is a founding principle of Livermore. Our continued success is crucial not only to our national security missions and basic science but also to our ability to attract the top scientific talent needed to keep our Laboratory on the cutting edge.

■ Dona Crawford is associate director for Computation.

Thunder's Power

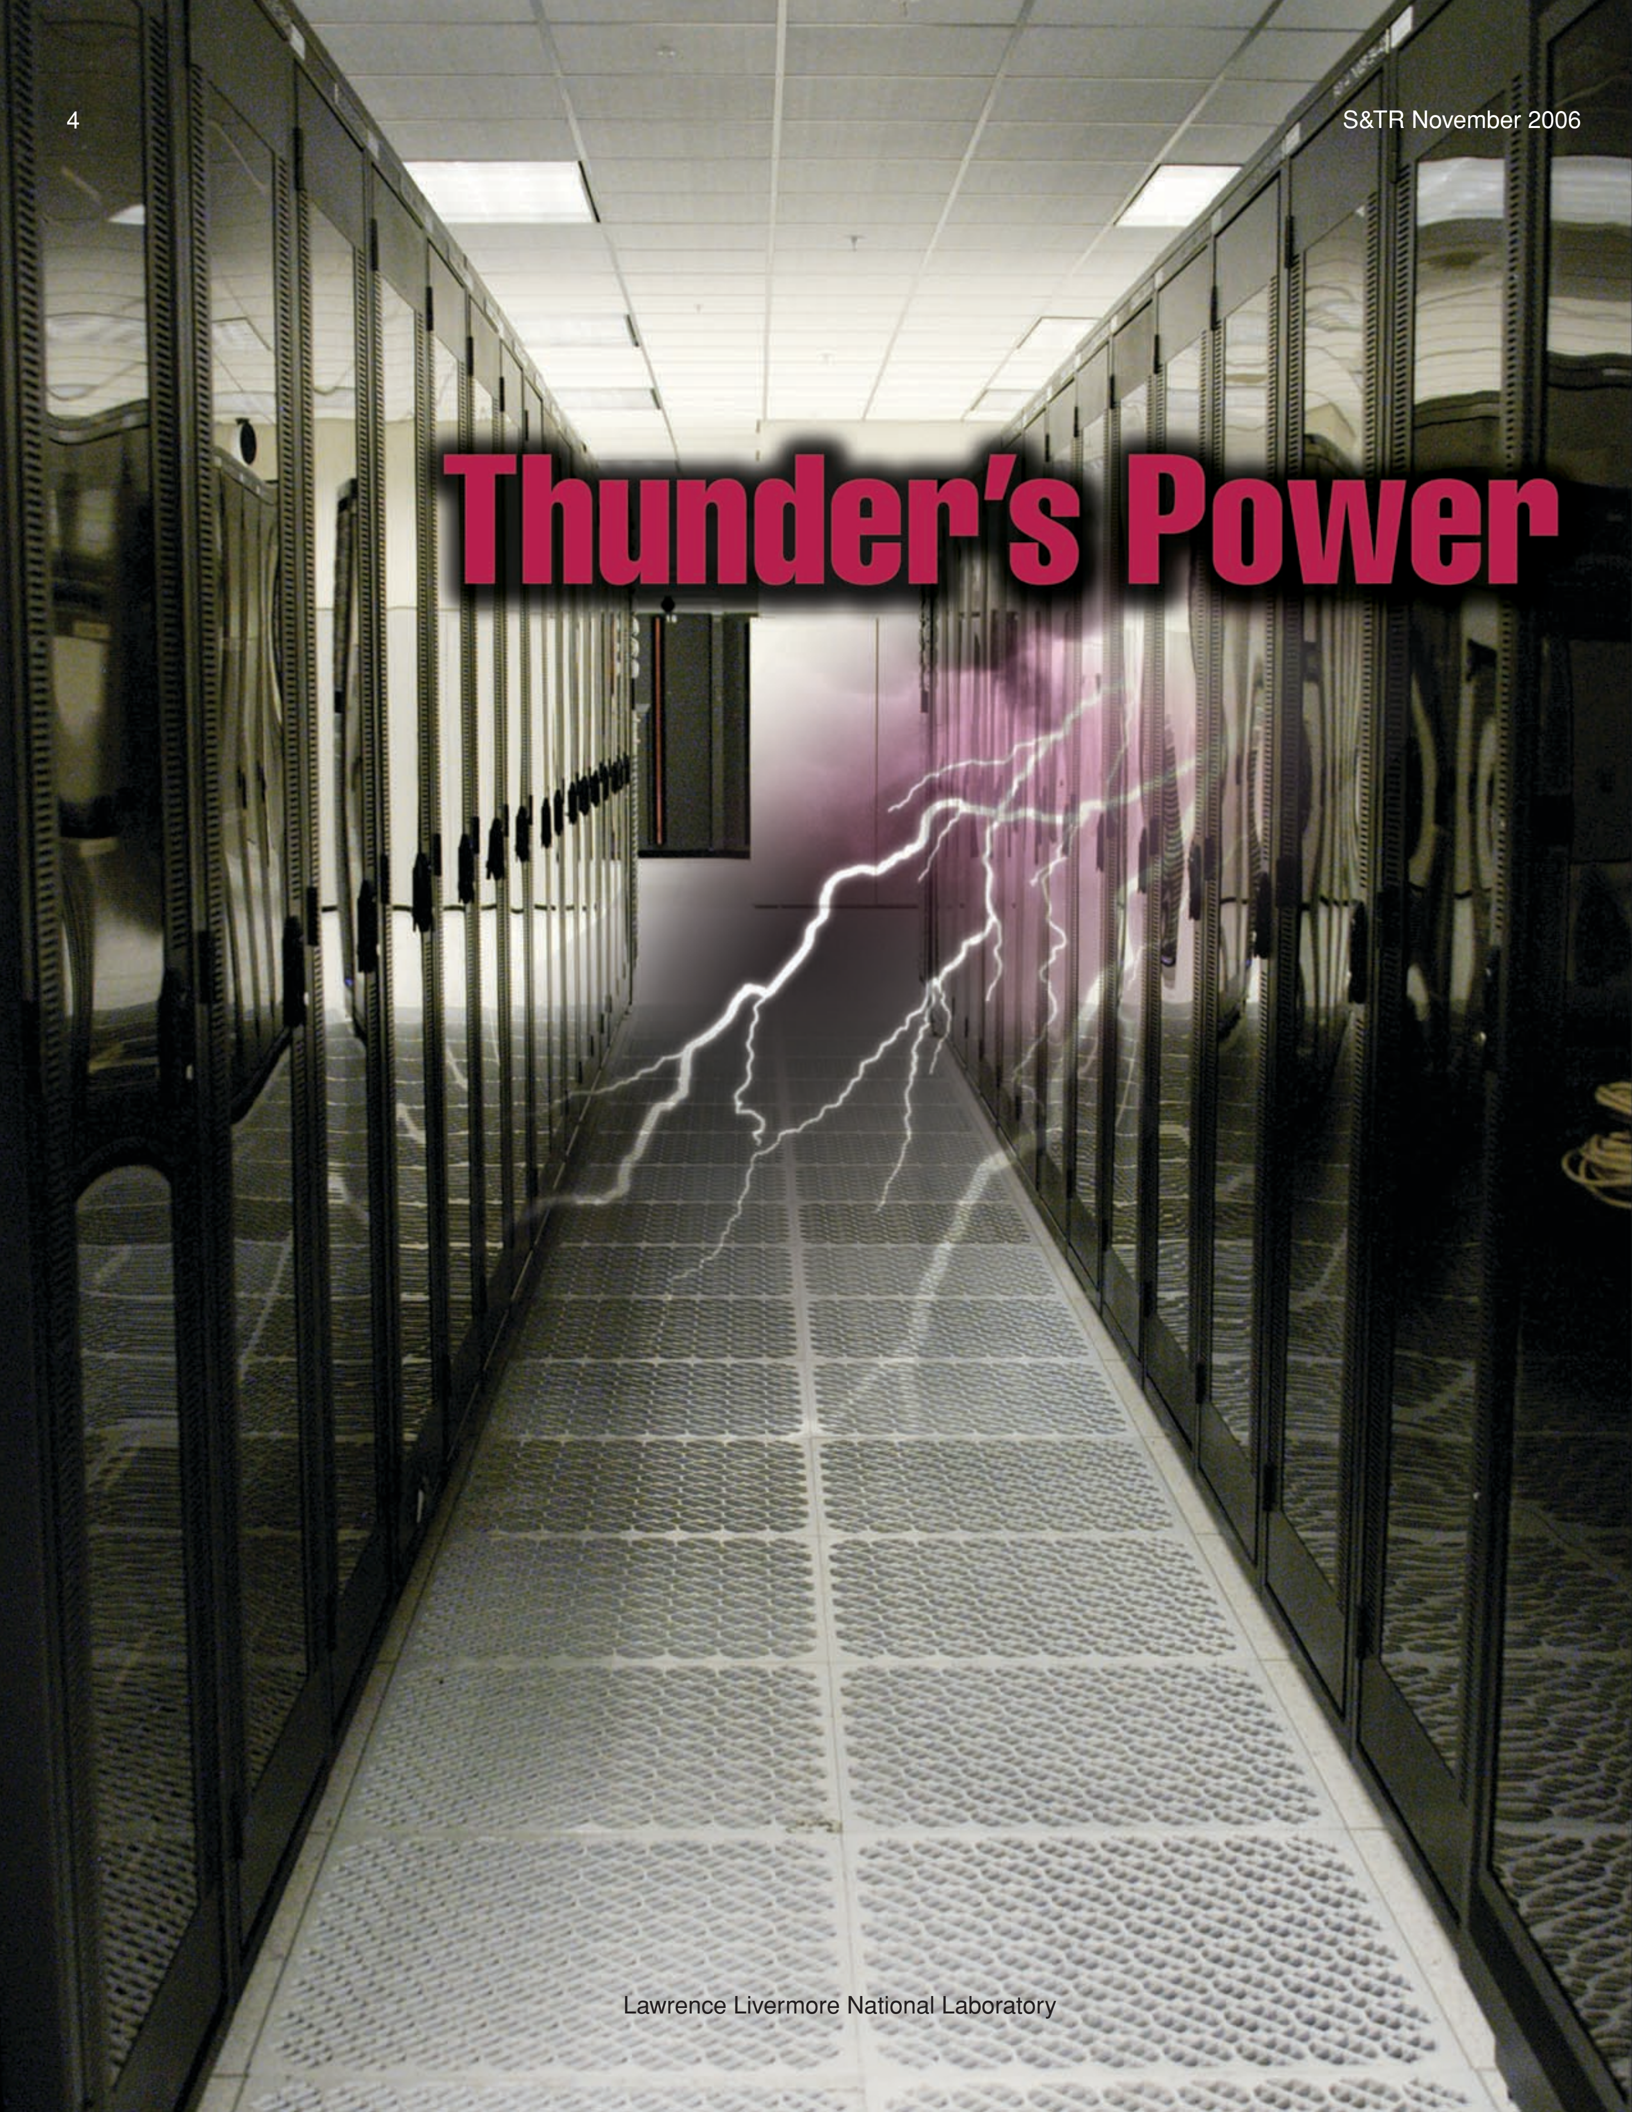

Delivers Breakthrough Science

Livermore's Thunder supercomputer allows researchers to model systems at scales never before possible.

WHETHER a computer is simulating the aging and performance of a nuclear weapon, the folding of a protein, or the probability of rainfall over a particular mountain range, the necessary calculations can be enormous. With each advancement in system hardware or software, Livermore researchers are ever ready to test the limits of computing power in search of answers to these and other complex problems. To assist them in their research, the Computation Directorate provides supercomputers that often rank high in performance. Today, the Advanced Simulation and Computing (ASC) Program's 360-trillion-floating-point-operations-per-second (teraflops) BlueGene/L supercomputer at Livermore holds the world's number one spot on the LINPACK benchmark, the industry standard used to measure computing speed.

As computing capability increases, Livermore's high-performance computers are assuming an increasingly central role in the Laboratory's research. In the 1990s, the Accelerated Strategic and Computing Initiative, now called the ASC Program, began using massively parallel computers, with the necessary computational speed and memory capacity to simulate complex multiphysics problems. In the last decade, the supercomputers' supporting role has grown to accurately and quantitatively connect disparate phenomena in basic and applied science research.

To exploit the machines' integrative capability and maximize potential scientific discoveries, physicist Michael McCoy of the Computation Directorate initiated the Multiprogrammatic and Institutional Computing (M&IC) Program in 1996 with institutional support from Jeff Wadsworth, then deputy director for science and technology, and director Bruce Tarter. The program provides Laboratory scientists and engineers access to ASC-class supercomputers for unclassified, mission-related research requiring the capability provided by big computers.

"M&IC provides the Laboratory's researchers a place to run unclassified calculations that are often as large as those used for the multiscale and multiphenomena calculations associated with the aging and performance of the nation's weapons stockpile. The simulations we produce differentiate the Laboratory from other institutions," says McCoy.

Power in Clusters

One advantage of parallel computers is their scalability. Traditionally, parallel computers relied on vendor-integrated systems and software and on proprietary interconnects, where the cost per teraflops improves very slowly. In 2000, McCoy proposed that M&IC switch to Linux cluster technology, which uses large groups of commodity microprocessors (computer systems that can be manufactured by multiple vendors), third-party interconnects, and open-source Linux operating system software. The first supercomputer assembled using this approach was the Multiprogrammatic Capability Resource (MCR), which includes 2,304 processors capable of 11 teraflops. In 2004, MCR was joined by Thunder, a 23-teraflops cluster of 4,096 processors. At its debut, Thunder was ranked the world's second fastest supercomputer.

Hal Graboske, then deputy director for science and technology, proposed allocating Thunder as a capability machine by awarding time on it to five Laboratory projects whose calculations required a significant percentage of the system. He dubbed the projects "Thunder Grand Challenges." These projects ranged from research in climate change to the interface of water and vapor to protein folding. Each Grand Challenge was strategically aligned with a Laboratory mission. Computer scientist Brian Carnes, who now leads M&IC, says, "The Thunder Grand Challenges have made important contributions to the scientific community because the machine's computational power allows researchers to accurately model systems at a scale never before possible."

(a)

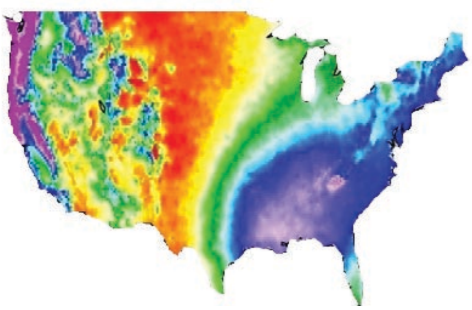

(b)

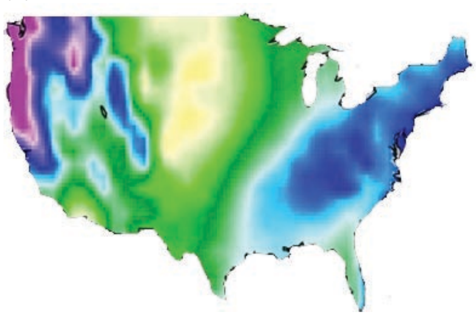

(a) Winter precipitation patterns from the Community Climate System Model agree well with (b) observed precipitation from 1971 to 2000, particularly in the western U.S.

Tracking Climate Change

Similar to the phenomena involved in stockpile stewardship, modeling climate change is a multiphenomena statistical science that requires multiple simulations. These runs take into account atmospheric and oceanic dynamics and physics, topography, vegetation, chemical constituents and aerosols, and the effects from clouds, precipitation, and ice. Climate change research has been focused mostly on global-scale models at resolutions no finer than 100 to 200 kilometers because of computer capability limitations. Doug Rotman, in Livermore's Carbon Management and Climate Change Program, is leading a Thunder Grand Challenge to model regional climate change at resolutions of about 10 kilometers. The calculations will allow, for the first time, evaluation of whether changes in certain regional phenomena, such as river flow, are due to anthropogenic (human-induced) factors.

The project team's strategy has been to perform multicentury, global simulations at 100-kilometer resolution, which is finer than that of any other century-long comprehensive climate run to date. Rotman says, "Comprehensive climate calculations must simulate centuries to millennia in order to properly represent the coupling between the atmosphere and oceans, particularly the deep ocean currents."

The Community Climate System Model achieves roughly 10 simulated years per computing day using 472 of Thunder's processors and includes a calculation that simulates a period of more than 1,100 years. Results of the global simulations serve as initial conditions and boundary data for 10-kilometer-scale atmospheric calculations that focus on areas the size of U.S. states. "No previous simulations could provide the resolution or length in time to assess internally driven regional climate variability, which is required to detect climate change at small scales," says Rotman. The calculations allow researchers to separate human-induced effects from natural variability at subcontinental scales.

The team collaborates with the Scripps Institution of Oceanography of the University of California (UC) at San Diego, the National Center for Atmospheric Research, and the University of Michigan. Researchers at Scripps will use the results of the regional calculations to drive a hydrology model to predict river flows. Comparison of these predictions with observational data will help ascertain the extent of regional climate change. The results will also have important implications for environmental issues such as air and water quality and the management of healthy ecosystems.

Water's Mysteries Uncovered

To run the necessary calculations for large simulations, Livermore researchers often rely on codes that use first principles to characterize material behavior. First-principles molecular dynamics simulations use the laws of quantum mechanics to describe the electrons in a system. The data are then used to compute the interactions between atoms without experimental input. Livermore has developed more than 50 classified and unclassified codes that use first principles.

Even a seemingly basic molecule, such as water, has many unresolved mysteries because researchers have not had the computational power and first-principles codes until now. Yet, understanding water is important for areas ranging from biological and physical sciences to atmospheric and planetary sciences. Thunder's power allows researchers to tailor calculations for individual studies by combining codes. Chemist Chris Mundy and his team have combined codes to resolve long-standing questions about water's behavior. (See *S&TR*, October 2005, pp. 12–18.) They were first to use Monte Carlo methods with first-interaction potentials to determine the vapor–liquid coexistence curve for water. "You can see water's vapor–liquid coexistence, and from the simulation, determine the thermodynamic properties," says Mundy. "This simulation has never before been done directly." Their results further understanding of chemistry at aqueous interfaces.

The team also has redefined the predicted phase diagram of water in the interior of large planets, confirming a theory that a superionic phase of water exists in the large planets. In this phase, oxygen atoms are fixed, or frozen, while hydrogen atoms move freely. Mundy says, "There has been a great debate about the existence of superionic water. Our calculations reveal that this exotic phase may exist."

Controversy also exists among experimentalists over the structure of water on Earth due to difficulty in interpreting results such as x-ray data of hydrogen-bond energies at different temperatures. Simulations performed on Thunder in collaboration with Richard Saykally's research group at UC Berkeley, which is well known for its experimental studies on water, are helping to resolve this controversy. Mundy's team also collaborates with researchers from the University of Minnesota, UC Irvine, and the University of Zurich.

One of the team's next goals is to use Thunder and molecular dynamics simulations to study the biochemical synthesis of uracil, an RNA nucleotide involved in enzyme production. The reaction that occurs in an enzyme is one of the most disputed mechanisms among scientists. The simulations will include a system of more than 30,000 atoms and will be the largest calculation of enzyme catalysis ever conducted. The study may help further understanding of how protein motion affects the efficiency of important enzymatic reactions.

Free from Energy Barriers

As proteins are synthesized, they fold into specific shapes that are related to their intended function. They fold on average in one-thousandth of a second and measure 2 to 3 nanometers long. Because a protein's structure is tied to its function, Livermore scientists, and many others worldwide, conduct studies of folding processes. Much of this work is influenced by experiments conducted for the Critical Assessment of Structure Prediction biennial event, which is organized by Livermore's Chemistry, Materials, and

(a)

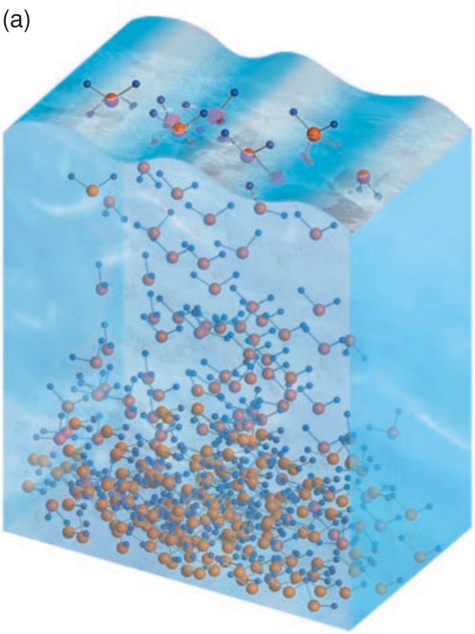

(b)

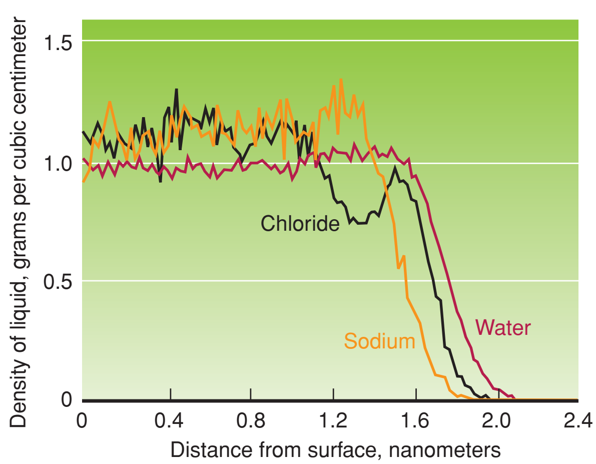

(a) The hydrogen atoms (blue) in water molecules have a tendency to orient themselves into the air at the surface. This orientation may increase reactivity with other types of molecules. (b) Simulations of sodium chloride and water molecules show greater interaction between molecules at the surface and a decrease in interaction for molecules in bulk liquid.

Life Sciences Directorate. (See *S&TR*, December 2004, pp. 12–17.)

Predicting the three-dimensional structure of a protein from its amino-acid sequence is extremely complex. However, determining the final structure is only half the problem. The other half is understanding how proteins fold, that is, the sequences of structural changes they undergo before reaching their final native structure.

Scientists estimate humans have between 50,000 and 100,000 different proteins. Many of the protein chains are composed of hundreds of amino acids. Because protein-folding processes are complex, researchers have mostly studied smaller proteins composed of chains of about 50 amino acids. One reason proteins are a challenge to study is that they exist in a watery environment. Researchers must therefore consider the individual atoms defined by the amino-acid sequences as well as those from the surrounding water molecules. Complicating this scenario is that each atom is dynamic and its state is in constant flux, producing fluctuating energy barriers. A protein must overcome these barriers as it moves into its folded state.

Simulating this ever-changing energy landscape has traditionally been limited to small systems over inadequate time scales. Theoretical physicist Farid Abraham is leading a Thunder Grand Challenge to simulate protein folding over a broad range of temperatures to reveal the thermodynamic pathways.

Abraham and collaborators Jed Pitera and William Swope at IBM Almaden Research Center simulated folding of the 1BBL protein, whose 40 amino acids are part of a larger protein involved in enzymatic activity. Researchers at other institutions have studied 1BBL, and their experimental results indicate that it folds in a downhill energy process. “Experimentalists can’t measure each wiggle of a protein as it folds,” says Abraham. “However, each of the atoms’ dynamics determines the folding process.”

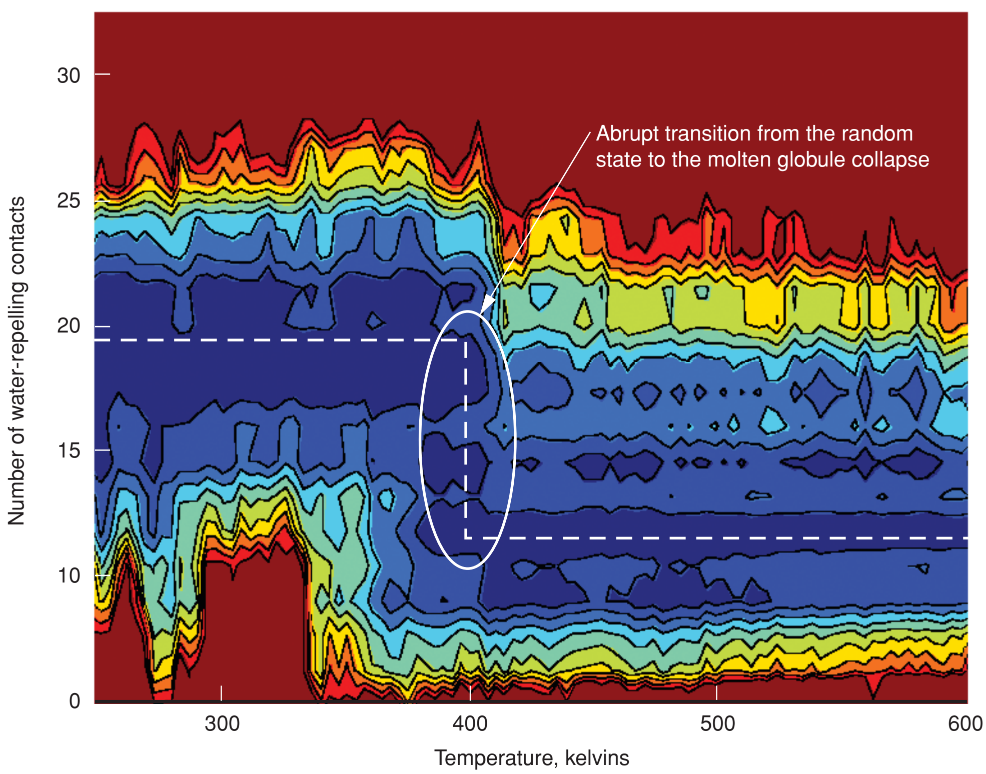

The Replica Exchange Molecular Dynamics technique simulates protein-folding thermodynamics over a broad range of temperatures. Results for the 1BBL protein show an abrupt transition from the unfolded state to a globular state at a temperature of 380 kelvins.

The team uses a new technique called Replica Exchange Molecular Dynamics (REMD) to simulate different copies (replicas) of the system at the same time but at different temperature values. Each replica evolves independently by molecular dynamics. High-temperature simulation segments facilitate the crossing of the energy barriers. Low-temperature segments explore in detail the conformations present in the minimum energy basins appropriate for the native state of the protein. The result of this swapping between temperatures is that high-temperature replicas help the low-temperature replicas to jump across the energy barriers of the system.

Simulations for the 1BBL protein comprised about 65,000 atoms, most of which were water molecules. The team created 256 replicas of this system and simulated each replica at a different temperature, ranging from 250 to 600 kelvins (body temperature is 300 kelvins). Each replica used four processors running concurrently for more than 2 million central-processing-unit hours.

“Researchers usually use supercomputers to simulate large systems,” says Abraham. “In this case, we’re using a supercomputer to simulate many small systems, while also taking advantage of its power to study a high-temperature state for better sampling of the low-temperature state. There has been an ongoing controversy among experimentalists over the last four years about whether the new downhill paradigm or the old two-state picture describes 1BBL folding. Our results confirm the experimental findings that the protein folds in a downhill energy process.” However, Abraham says that the story is more complicated. “The 1BBL protein does have an initial energy barrier it has to overcome in order to collapse into a globular package, but then it slithers into its native state through a downhill path, establishing the new folding paradigm."

The simulations are three times larger and four times longer than previous studies. However, Abraham says it's not just having access to the best supercomputers in the world that enables the Laboratory's discoveries. "Our results were possible because of M&IC's dedication week after week for a year. You can have big computers, but it's the individuals who are devoted to the science that make the difference."

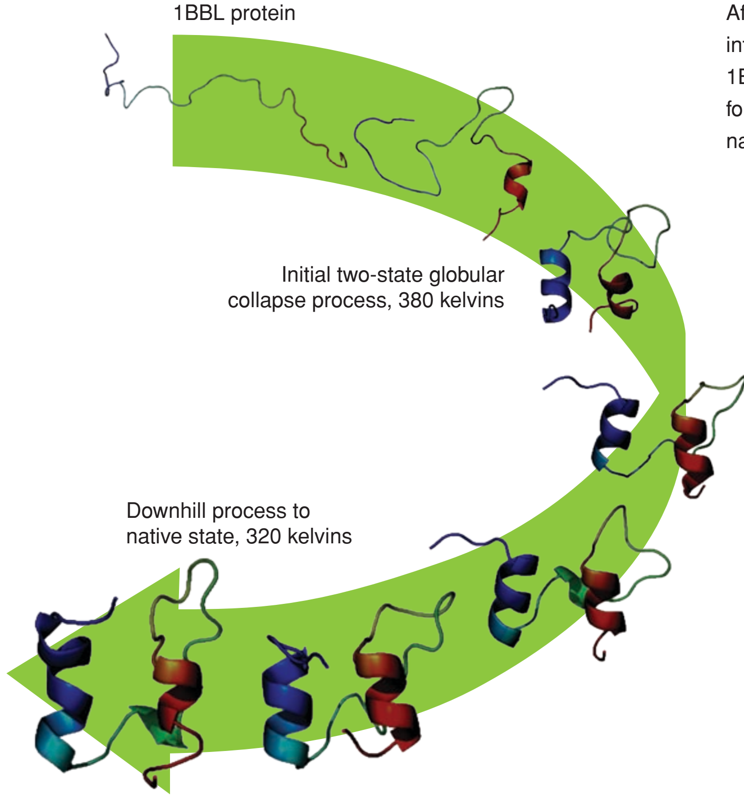

After an initial collapse into a globular state, the 1BBL protein continues to fold until it reaches its final native structure.

Stress–Strain Can Build Strength

Whether studying an organic material, such as protein, or an inorganic material, such as metal, scientists must understand the origin of microstructural features to predict the material's properties. In a metal, atoms stack in an orderly fashion, forming a crystalline lattice. However, some regions of the metal are less ordered. These disorderly areas have impurities or misaligned planes of atoms, called dislocations, that affect a metal's properties.

Plasticity occurs in a metal as dislocations move, or slip, along well-defined crystal planes and directions. To predict a metal's strength, researchers must understand its dislocation behavior. Dislocation dynamics—the interaction among dislocation lines—may be responsible for strain hardening, a property of metals in which a material's strength increases as deformation increases.

The speed at which dislocations occur within an area in a metal can cause it to fail. An important part of ASC's stockpile stewardship work is to examine metal strength under conditions of high strain rate and pressure. Materials scientist Vasily Bulatov leads a Grand Challenge team that is modeling in unprecedented detail the mechanics of dislocation motion by following the evolution of dislocation lines in response to applied stress. The team is using Thunder and the code Parallel Dislocation Simulator (ParaDiS) to develop an experimentally validated and computationally efficient model of crystal strength. The team has focused on molybdenum and tantalum because these metals are similar to some materials in the nuclear stockpile. (See *S&TR*, November 2005, pp. 4–11.)

Dislocation dynamics simulations involve tracking millions of dislocation lines over millions of time steps. Thunder allows the team to surpass traditional computational studies that model systems two to three orders of magnitude smaller in size and shorter in time scale. "Thunder is the muscle ParaDiS needs," says Bulatov. "We can investigate, for the first time, microstructural origins of stress–strain behavior."

The team's simulations are the first to show that during strain, colliding dislocations could partially merge into a junction bounded by two nodes. These multijunctions tie the dislocation lines into a tight knot that seems to confer strength in metals. The researchers also found that microstructures with many multijunctions can multiply dislocations at a faster rate than microstructures with fewer multijunctions. The greater the number of multijunctions formed, the greater the amount of stress required to deform a metal (strain) before it breaks, and the harder the resulting metal.

Guiding Nanoscale Designs

The ability to predict material performance at micro- and nanoscales (billionths of a meter) helps researchers design chemical and biological detection instruments for homeland security. To guide the design of such small-scale instruments, Livermore researchers use simulations to predict how a component's material properties may change as dimensions are reduced. Subtle changes in a material's electronic and structural properties can be crucial to a microdevice's performance. The simulations also predict how the device's surrounding environment may change.

For example, in addition to the diverging theories on water's structural properties, debate exists regarding the properties of confined water, that is, water trapped between layers of another substance. Understanding how the structural and dynamic properties of water change as it passes through confined regions is extraordinarily challenging for both experimentalists and theorists. Physicist Eric Schwegler, who leads Livermore's Quantum Simulations Group, is conducting a Grand Challenge to simulate the electronic and structural properties of confined water at the nanoscale as it comes into contact with materials such as silicon carbide, graphite sheets, and carbon nanotubes. The team is focusing on these materials because their high electrical conductivity and biocompatibility are useful characteristics for microfluidic devices designed to detect biothreat agents. (See *S&TR*, January/February 2006, pp. 11–17.)

To model the systems, the team uses a Livermore-developed, first-principles molecular dynamics code called Qbox. Designed for massively parallel systems, Qbox is used for conducting research on condensed matter subjected to extreme temperatures and pressures. The Qbox code calculates the electronic structure of a system with methods based on density functional theory, which defines a molecule's energy and electronic structure in terms of its electronic density.

The team requires about 560 of Thunder's processors and the Qbox code to simulate systems containing between 500 and 600 atoms for 20 to 25 picoseconds. Silicon and carbon have different electronic properties, one of which is their affinity or aversion to water. One of the surprises the team found is that surfaces with an aversion to water tend to have a more pronounced effect on the structural and dynamical properties of confined liquid water than do surfaces with an affinity to water. For example, on the surface of graphite, which has an aversion to water, complex hydrogen-bond ring structures form at the surface. In confinement, the probability of forming these structures increases. Water molecules confined in carbon nanotubes demonstrate the same hydrophobic tendency.

The results can be used to prepare nanoscale materials that are most compatible for an intended purpose. They may also

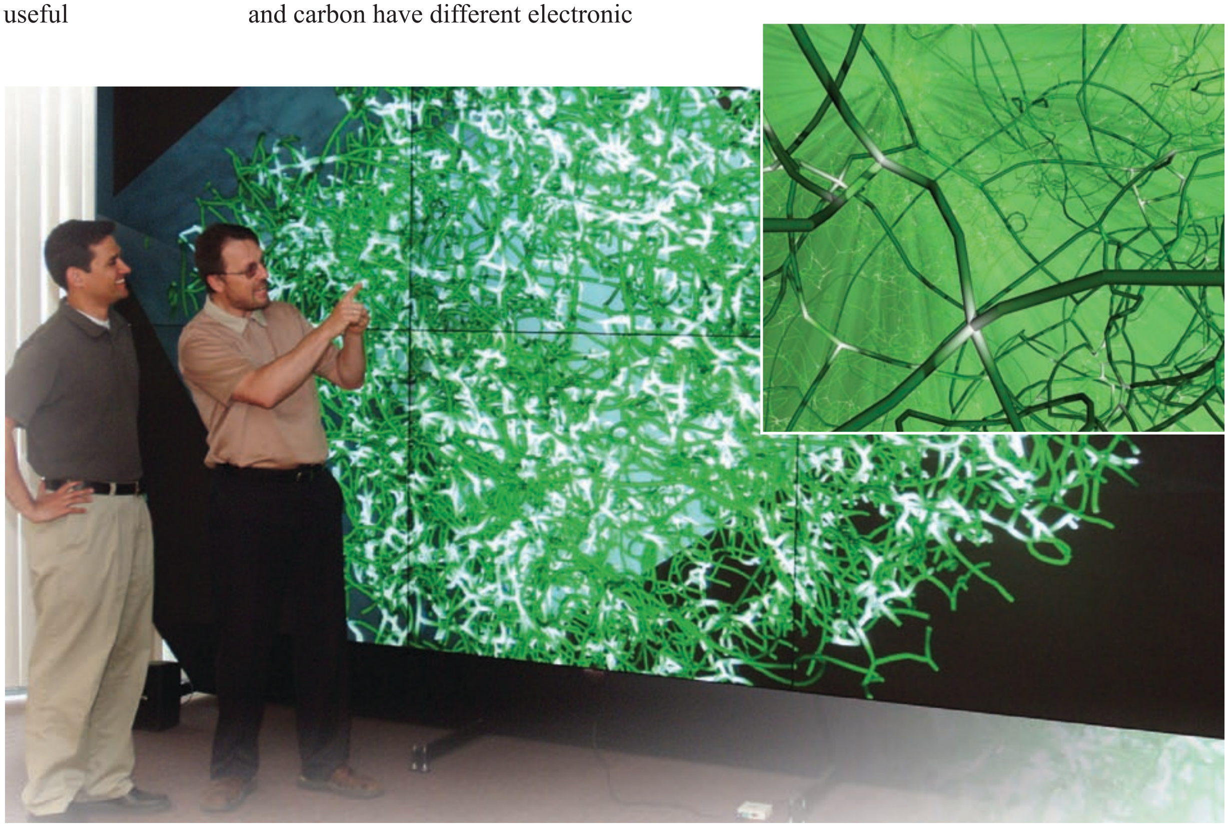

Tom Arsenlis (left) and Vasily Bulatov (right) use the Parallel Dislocation Simulator code to study how metals deform and fail. The white areas in the inset show the formation of multijunctions, which tie three or more dislocations into tight knots, conferring strength to metal.

help protein-folding studies because the confined state of water may affect folding processes in an environment in which molecules have an aversion to water. "Access to Thunder through M&IC has enabled us to perform a predictive and systematic study of nanoscale confinement with respect to length scales, dimensionality, and interface effects without having to compromise on the level of theory," says Schwegler.

Schwegler's team includes collaborators from UC Berkeley, UC Davis, Massachusetts Institute of Technology, and Polytechnic of Torino, Italy. In fact, all of the Grand Challenge projects involve researchers outside the Laboratory. Because large volumes of data are produced in these studies, M&IC is developing the Green Data Oasis collaboration, which will enable external collaborators to more efficiently share data with Livermore researchers. Carnes says, "The powerful unclassified systems at the Laboratory produce vast amounts of numerical results. A storage capability outside the Laboratory's firewall allows our researchers to share data efficiently so they can achieve programmatic goals. The platform also enables UC students and faculty access to Livermore science."

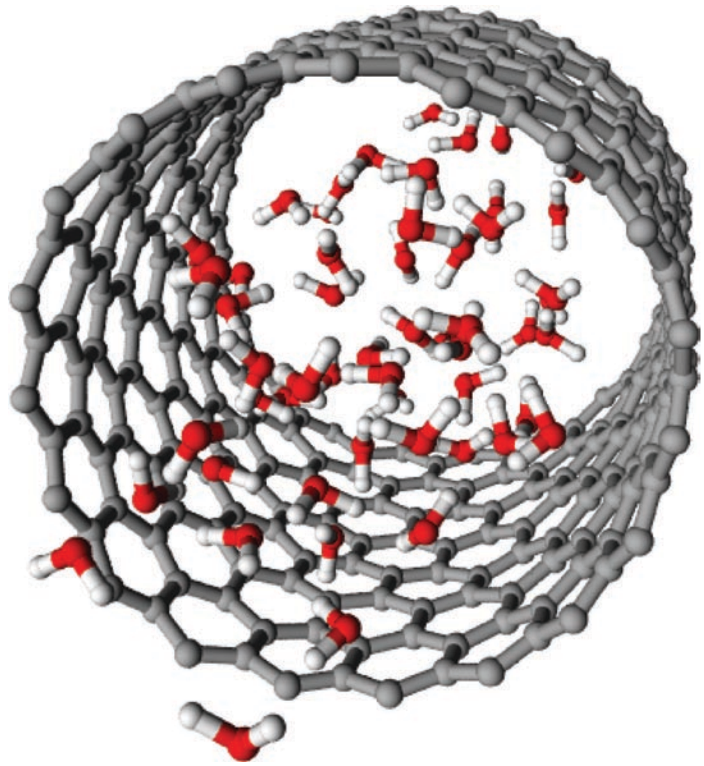

A simulation of water molecules (red and white) confined within a carbon nanotube demonstrates carbon's hydrophobic tendency.

Paving the Way for Atlas

When it isn't running simulations for one of the Grand Challenge projects, Thunder is being used to advance other scientific areas. For example, a team led by physicist Fred Streitz performed simulations for early studies on the nucleation and growth of grain structures in metals. The work was transferred to BlueGene/L, and in 2005, Streitz's team won the Gordon Bell Prize for reaching over 100 teraflops using a scientific application. (See *S&TR*, July/August 2006, pp. 17–19.) In another Livermore effort, physicist Evan Reed used Thunder to simulate coherent light from crystals. (See the article on p. 22.)

M&IC is currently gearing up for Atlas, a 44-teraflops machine that also will serve as the capability supercomputer for Grand Challenge projects beginning in January 2007. Because Atlas will roughly double Thunder's peak performance, Carnes expects Atlas simulations to make a large scientific impact on future Grand Challenge projects. "Consistently investing in institutional computing is an essential component of our science and technology investment strategy," says Carnes. "The addition of Atlas and other systems allows our researchers to push the limits of computing and its application in simulation science to ensure that Livermore remains a leader in science and technology."

—Gabriele Rennie

Key Words: 1BBL protein, Atlas supercomputer, climate change, confined water, dislocation dynamics, Green Data Oasis, Multiprogrammatic and Institutional Computing (M&IC) Program, Parallel Dislocation Simulator (ParaDiS) code, plasticity, protein folding, Qbox code, Replica Exchange Molecular Dynamics, Thunder supercomputer.

For further information contact Brian Carnes (925) 423-9181 (carnes1@llnl.gov).

Extracting Key Content

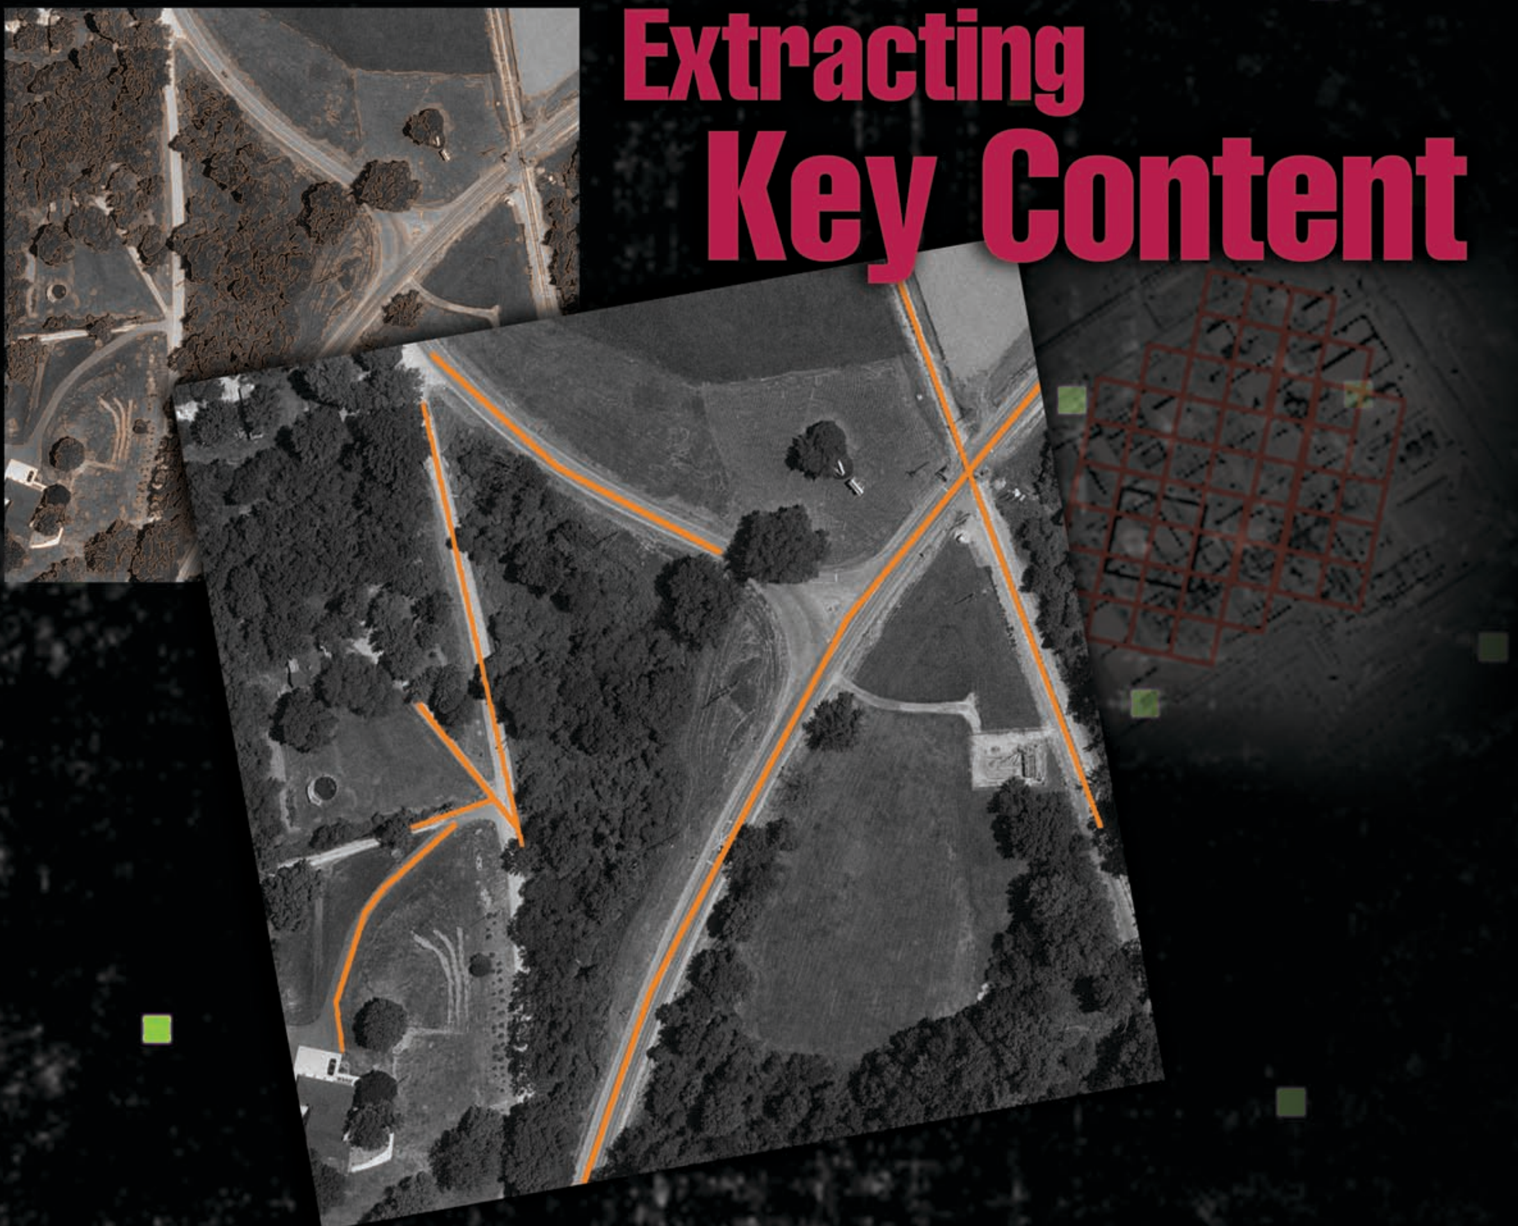

from Images

A new system called the Image Content Engine is helping analysts find significant but often obscure details in overhead images.

ADVANCEMENTS in imaging technologies, particularly in remote sensing systems, are producing vast amounts of data that can easily overwhelm human analysts. A team of Livermore engineers, computer scientists, and physicists has come to the aid of overburdened analysts who need to quickly analyze large volumes of overhead images. The team's new extraction system, called the Image Content Engine (ICE), allows analysts to search massive volumes of data in a timely manner by guiding them to areas in the images that likely contain the objects for which they are searching.

ICE was developed under a three-year Laboratory Directed Research and Development (LDRD) Strategic Initiative begun in 2003. It encompasses a new approach for the computer-aided extraction of specific content information from different kinds of images, but especially those images taken with overhead sensors.

Imagery analysts in the Laboratory's Nonproliferation, Homeland and International Security (NHI) Directorate are using ICE to identify objects such as specific types of buildings or vehicles. A variation of ICE is being evaluated for adoption by the National Ignition Facility's (NIF's) Optics Inspection Analysis Team. In addition, the system may be applicable to other fields, including nondestructive analysis, biological imaging, astronomy, and supercomputer simulations of nuclear weapon performance.

"We need new ways to handle the increasingly vast amounts of data generated by sensors," says engineer David Paglieroni, technical lead and co-principal investigator of ICE. Paglieroni, who leads the Engineering Directorate's Imagery Sciences Group, notes that information overload is a serious problem in the intelligence community because analysts face a flood of data from different sources.

Jim Brase is the other ICE co-principal investigator and head of the Optical Science and Technology Division of the Physics and Advanced Technologies Directorate, a group that develops advanced detectors for astronomical research and national security. The ICE team also includes engineers Barry Chen, Chuck Grant, Aseneth Lopez, Doug Poland, and George Weinert; computer scientists Jim Garlick and Marcus Miller; and postdoctoral researchers Siddharth Manay and Faranak Nekoogar.

Brase notes that the laboratory has been a leader in data mining, which involves finding items of interest in large amounts of data. For example, Livermore computer scientists are helping to tease out the most relevant features in three-dimensional visualizations of scientific simulations. (See *S&TR*, November 2004, pp. 12–19.) Livermore researchers have also developed several algorithms to extract items of interest in supercomputer visualizations of simulated weapons performance. The goal of these efforts is to enable analysts in all of the Laboratory's programs to focus on the most important details, which are often buried within massive amounts of visual information.

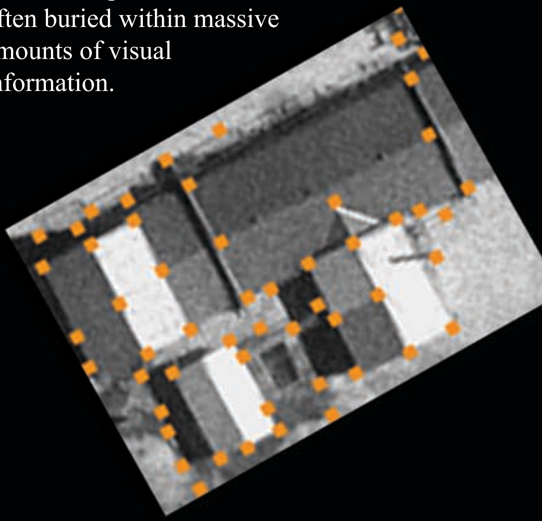

Extracting Desired Content

The ICE development team worked closely with imagery analysts in NHI's Z Division (International Assessments and Knowledge Discovery). August Droege, leader of the Precision Intelligence Group, explains that analysts pore over large amounts of imagery, either at light tables or at computer workstations, often looking for obscure objects. "Our analysts are very fast, but there can be thousands of images to sort through," says Droege.

The resolution and the amount of area covered in an image can vary widely, with many images covering enormous areas. For example, a commercial satellite image with 1-meter-per-pixel resolution (6,000 rows by 10,000 columns of pixels) covers an area of about 60 square kilometers. Wes Spain, leader of Z Division, says, "Efforts over the past 20 years aimed at analyzing this kind of data with computers have had limited success."

ICE can accommodate images acquired with different types of overhead sensors and at varying resolutions. The software is also able to account for the fact that images are taken at different times of the day, during different seasons, and under changing weather conditions. Images often contain distracting background clutter, and potential objects of interest can be in full or partial shadow, occluded by trees, or obscured by snow or clouds. "Finding specific objects, such as particular types of buildings or vehicles, in overhead images that cover hundreds of square kilometers, is a difficult task," says Paglieroni.

The ICE architecture can run on different computing platforms and operating systems, such as Windows or Linux operating systems, laptops or powerful clusters of computers, and isolated or networked processors. The ICE software contains a library of algorithms, each of which focuses on a specific task. The algorithms can be chained together in pipelines configured through a graphical user interface. Each pipeline is designed to perform a specific set of tasks for extracting specific image content.

One of the most useful ICE algorithms, known as gradient direction matching (GDM), was developed by Paglieroni. GDM uses a novel approach for implementing equations to rapidly "pull" objects of interest out of images that have large amounts of visual clutter. The algorithm compares pixel gradient directions (the direction of flow from dark to light) in an image being analyzed to pixel gradient directions that are perpendicular to unoccluded edges in submitted models. As a result, GDM is relatively insensitive to image brightness and contrast variations. For overhead images, ICE uses GDM for matching objects to a variety of models, for extracting vertices (where two or more lines meet) of specified polygons, and for computing certainties associated with

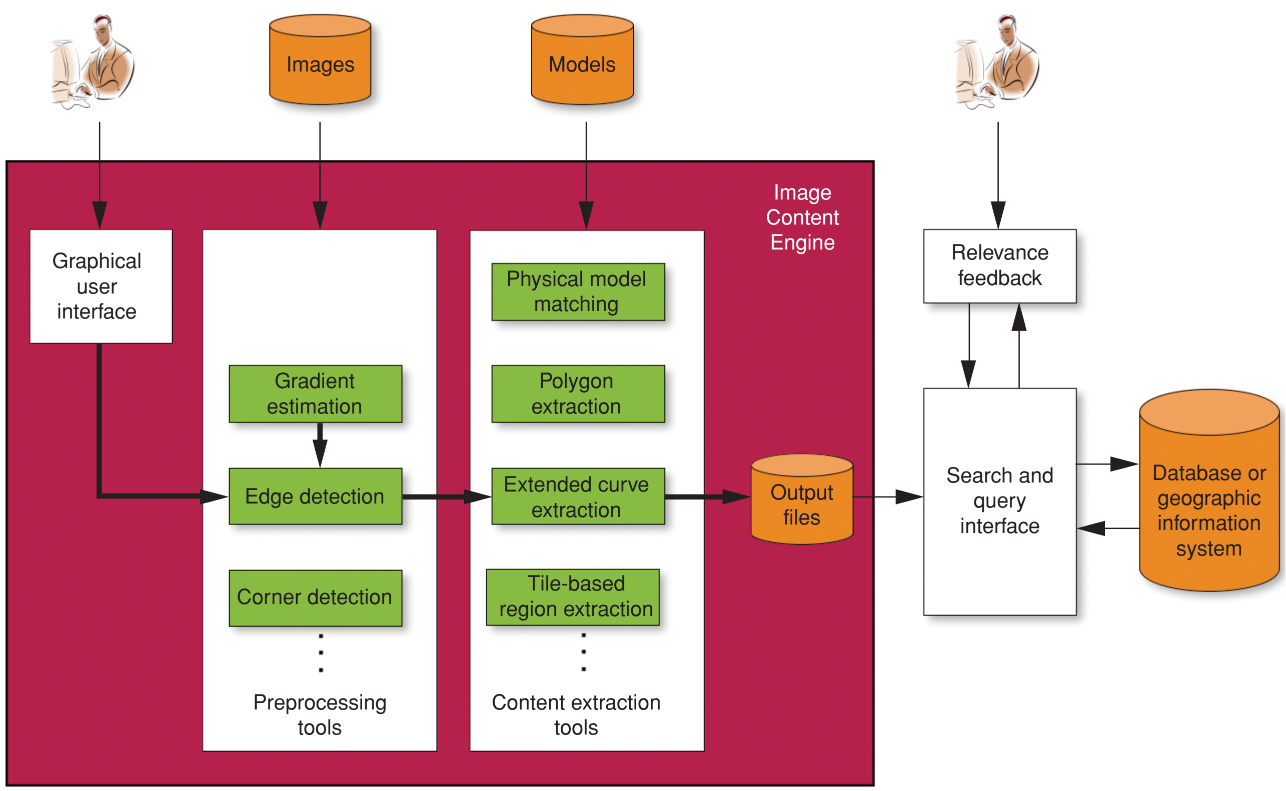

The Image Content Engine (ICE) software contains a library of algorithms, each of which focuses on a specific task. The algorithms can be chained together into pipelines configured through a graphical user interface. In this example, a user requests ICE to extract roads (extended curves) from images.

matches to submitted polygons. GDM has also been successfully used in NIF optics inspection to search for diffraction patterns that signal defects. "GDM has been a gold mine for us," says Paglieroni. "We consider ourselves very fortunate when one algorithm can be used to support more than one important processing task."

ICE also provides tools for extracting regions, extended curves, and polygons from images. The region extraction tool breaks images into small, adjacent square tiles containing one or more pixels. For each tile, the algorithm searches for spectral or textural characteristics and then groups tiles with similar features into regions. This tool is useful for separating distinct areas such as forests, bodies of water, plowed fields, and clusters of buildings from the image background.

The extended curve extraction tool is used to find lines of communication, such as roads, power lines, and canals, in overhead images. ICE uses a hierarchical approach in which mistakes made in processing at the pixel level are corrected at successively higher levels, such as the line-segment or curve levels. The ICE team developed a novel approach to consolidate collections of broken line segments that are nearly parallel into single, consolidated line segments. These segments are then assembled into consolidated curves. This approach, which is still under development, allows ICE to extract lines of communication from even highly cluttered scenes, a capability that other image-analysis tools rarely possess. (See the figure on p. 16.)

The polygon extraction tool, which is also under development, is useful for automatically extracting man-made objects, such as buildings and moving vehicles, from overhead images. The ability to perform such automated extractions facilitates computer-aided searches for objects in overhead images. The tool assigns each polygon vertex a pixel location, sharpness, and orientation. By analyzing groups of vertices, the algorithm

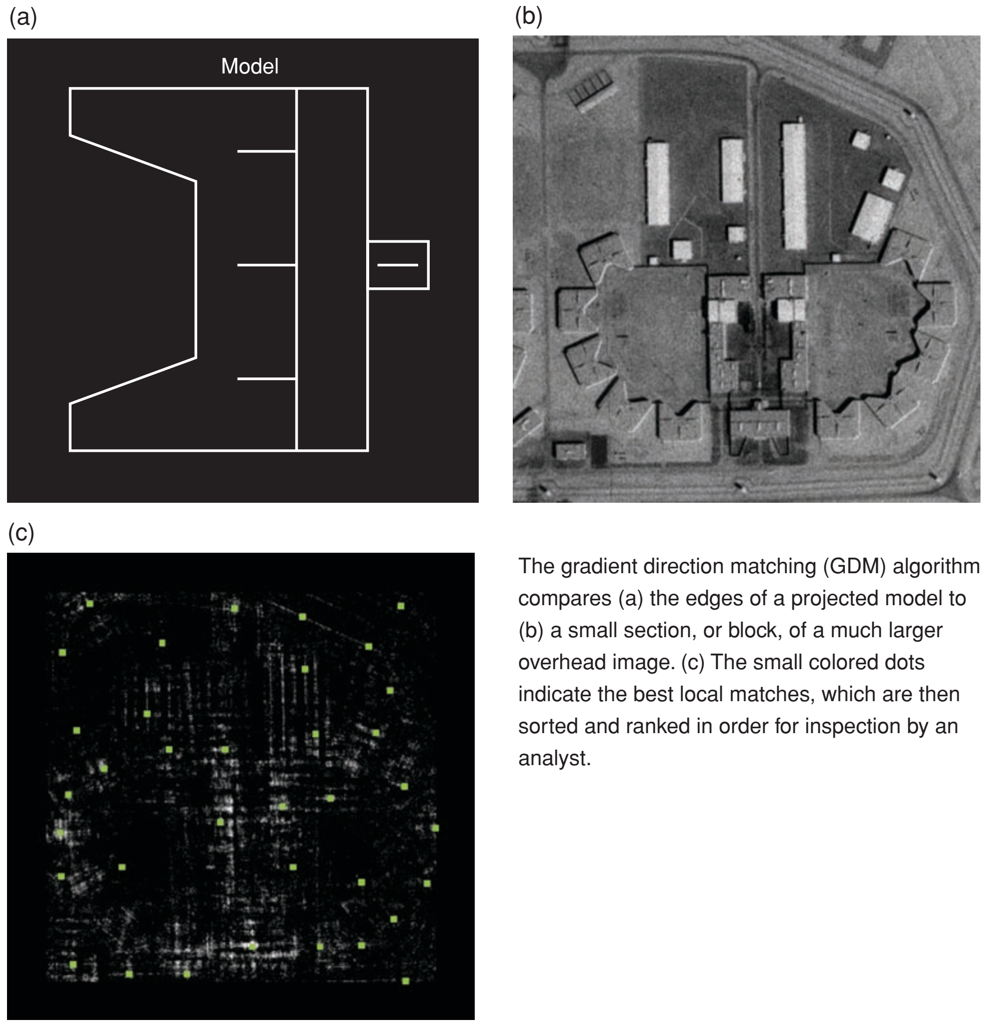

The gradient direction matching (GDM) algorithm compares (a) the edges of a projected model to (b) a small section, or block, of a much larger overhead image. (c) The small colored dots indicate the best local matches, which are then sorted and ranked in order for inspection by an analyst.

ICE extracts a region composed of tiles (green boxes) whose features are consistent with built-up areas.

quickly extracts polygons of prescribed geometry, independent of their position and orientation in the image. If the model specifications are relaxed, polygons of arbitrary size or with a particular ratio of width to height can be found. “We are encouraged that GDM appears to be outstanding at extracting vertices, which has traditionally been a difficult task,” says Paglieroni. (See the top figure on p. 17.)

The ICE team envisions that by combining diverse content extraction tools, analysts will be able to search for complicated patterns containing different pieces of content. As an analyst adds pieces of content to be extracted from an image, the false-alarm rate will drop. That is, fewer areas in the image will be mistakenly identified as a feature of interest.

(a)

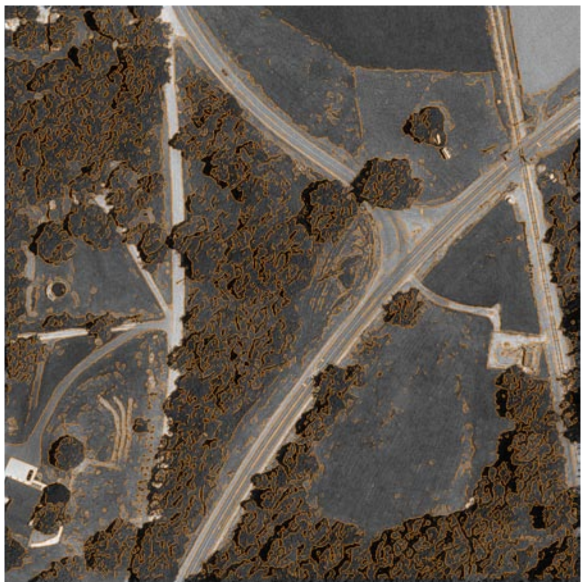

(b)

An example of extracting extended curves from an image shows (a) traditional edges, as might be found with other image-processing systems and (b) consolidated lines and curves.

Specifying and Matching Models

ICE provides an interface for specifying the sets of images to be processed and how they are to be processed. A user either specifies the model of an object to search for from a registry furnished by ICE or supplies the model to match. The user-supplied model can be a specified polygon or a physical model with two or three dimensions (in meters) based on a detailed drawing or, alternatively, with dimensions derived from a reference photograph.

ICE divides images into overlapping blocks, and the amount of overlap depends on the size of the object being searched. At each pixel location, the edges of the submitted model are projected onto blocks at typically 75 different orientations. The gradient direction perpendicular to the edge boundary of the projected model is computed for each pixel in each block using GDM. The best match over all orientations at every pixel is saved for subsequent ranking. “This processing stage is computationally expensive and may require parallel-processing compute clusters,” says Weinert. With adequate computational horsepower, ICE can process hundreds of images in a few hours.

In the search and query stage, ICE creates image thumbnails (small image blocks) of the closest matches to submitted models. These matches are sorted in order of decreasing similarity to the model and presented to an analyst for visual inspection and interpretation. Much like an Internet search engine, the program assigns each match a score, ranging from 0 to 1.00 (0 to 100 percent). Although parallel computing may be required for ICE to process large sets of images in a timely manner, ranked thumbnails can be generated quickly on a personal computer.

When an analyst clicks on a thumbnail, the surrounding context for that thumbnail is displayed in a larger window. "The object of interest may not be contained in the thumbnail with the highest score, but the probability is high that the object of interest will be present among relatively few top-scoring thumbnails," says Weinert. Because analysts can miss targets in thumbnails that are too dark or light or that have poor contrast, a user can enhance the brightness and contrast settings. For example, a user can brighten an object in shadow or darken an object in snow.

ICE can compute geographic coordinates (latitude and longitude) for matches extracted from images. The ICE team is developing the capability to link those matches to geographic information system data such as street and topographic maps. Also in the works is the ability to store ICE matches in commercial databases and semantic graphs (advanced techniques that capture complex relationships among detected objects of interest) along with nonimage data generated by other systems.

(a)

(b)

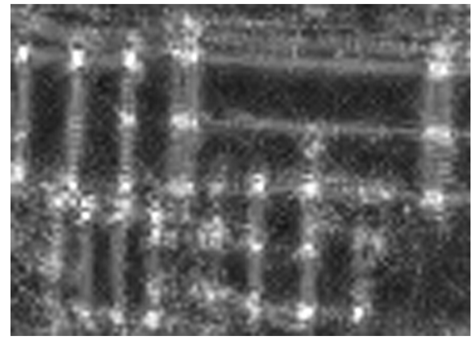

(c)

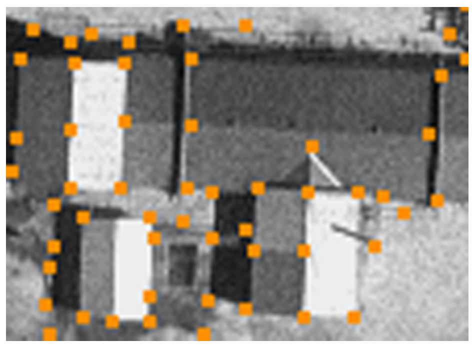

(a) An example of 90-degree polygon corner extraction using the GDM algorithm shows how the algorithm detects the flow of light to dark gradient magnitudes in the image and (b) matches the gradients to a model of 90-degree polygon corners. (c) Polygons in the imag e are then extracted from the candidate corners.

(a)

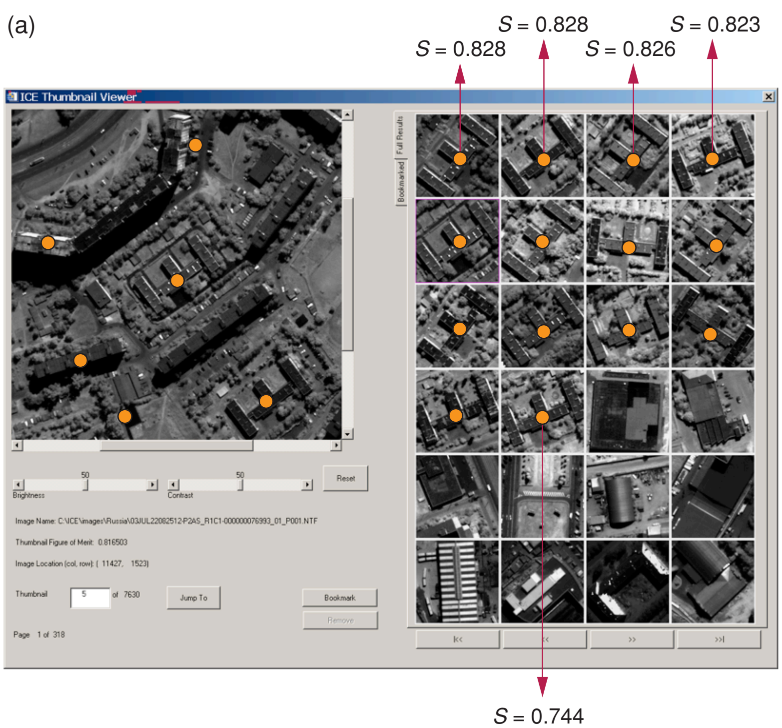

(b)

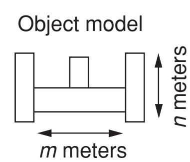

(c)

(a) In the search and query stage, ICE creates image thumbnails of the closest matches to (b) a submitted model. These matches are sorted in order of decreasing similarity to the model and presented to an analyst for visual inspection and interpretation. In this example, the upper left-hand thumbnail has a similarity (*S*) of 0.828 (the highest of any match). An analyst can click on any thumbnail and see it in a larger context. The thumbnails can also be adjusted for brightness and contrast. (c) Each match is captured as a vector and stored in an output file of matches.

ICE Proves Itself

For more than a year, ICE has demonstrated significant potential to increase productivity by focusing imagery analysts' attention on possible objects of interest. "ICE gets the eyes of analysts on the highest priority images," says Spain. "It serves almost like a triage function, telling us which images are the most deserving of our attention."

The program can also be a huge timesaver. "We can potentially do so much more with ICE," says Droege. She estimates that ICE could reduce more than one week of intensive analytical work into a couple of hours. "ICE significantly increases the probability of finding something of interest. It does the heavy lifting for us," she says.

Droege is especially appreciative of the program's flexibility. Analysts can ask the program to look for a specific structure that matches exact measurements or for one that merely suggests an object such as an airfield. "We may want to search just for a structure in a desolate area where one might not expect anything to be built," she says.

Paglieroni explains that ICE provides a computer-assisted approach to analysis. In contrast, the computer-automated approach attempts to replace human abilities for analyzing and interpreting images with computer programs. However, computers are not good at interpreting images. "The computer-assisted approach uses the strengths of computers to enhance the strengths of human analysts," he says.

Scrutinizing NIF Optics

The NIF Optics Inspection Analysis Team in conjunction with the ICE collaborators has adapted the GDM algorithm to strengthen NIF's computerized optics inspection system. The project, called Finding Rings of Damage in Optics (FRODO), is an Engineering Directorate technology-base project. Technology-base projects adapt new techniques to the specific needs of Livermore programs.

NIF currently has two bundles of eight laser beams each, with a third bundle undergoing installation. When completed, NIF will have thousands of square-shaped optics that guide, amplify, and focus light from 192 laser beamlines (24 bundles) onto a tiny target. Camera systems throughout the laser take pictures of individual optics as part of the NIF Optics Inspection Analysis System, which analyzes the images to automatically monitor the condition of optics. Up to 80 images are taken of the optics in each laser beamline every time a full inspection is requested. When all 192 beams are operational, some images will be required multiple times per day. The goal is to automatically detect and characterize changes over time, providing information to help managers decide whether to repair, replace, or continue using an optic.

Engineer Laura Kegelmeyer, assigned to NIF from the Engineering Directorate's Signal, Image Processing, and Control Group, is the principal investigator of the FRODO project. Kegelmeyer says that NIF optics must be frequently inspected to operate safely, ensure quality performance, and determine when optics need to be refurbished. "The slightest imperfection can affect the uniformity of the laser beam, and some imperfections can grow under repeated exposure to laser pulses," says Kegelmeyer. "We want to find any defect less than 500 micrometers in diameter so

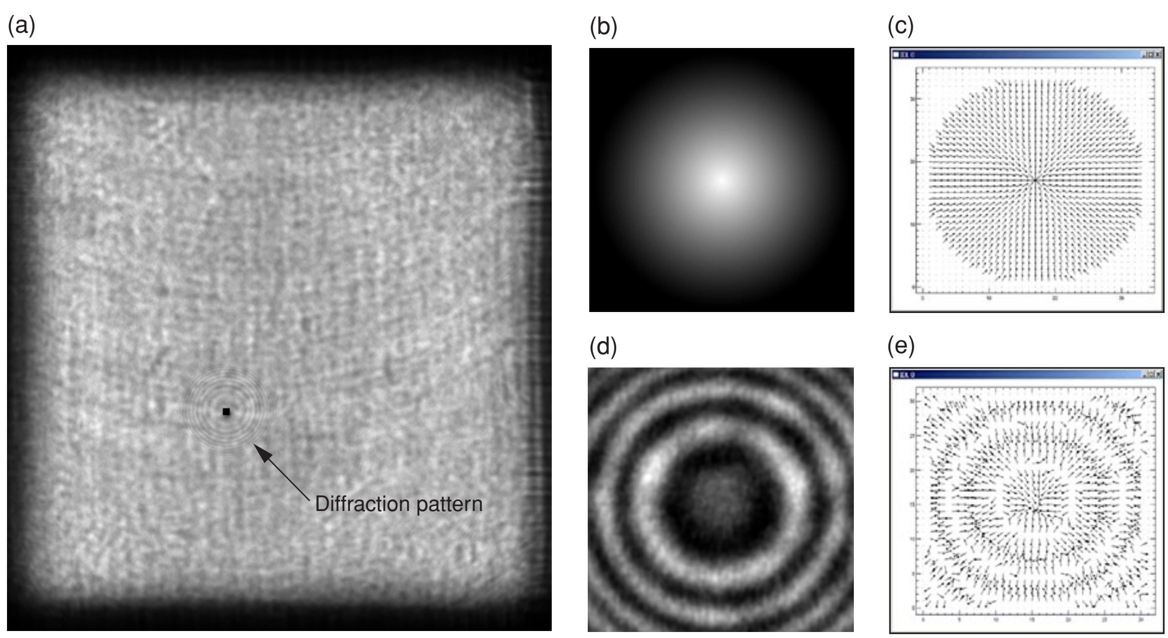

The GDM algorithm has performed extremely well in detecting diffraction patterns in test images of small flaws in laser optics. (a) In the test image, a faint diffraction pattern is detected. (b–c) To detect the diffraction rings, the operator uses a model that produces a field of light-to-dark gradient vectors flowing away from the center of a circle. (d–e) The GDM algorithm then identifies regions in the test image whose light gradient directions best match those of the luminance disk. After GDM identifies potential rings, ICE ranks the rings by how well they fit the model.

the optic can be refurbished before the defect grows too large."

Custom algorithms developed at Livermore currently detect and characterize defects on NIF optics using a "direct" approach to find irregularities in images of the optics. However, the location of some optics limits the detail that can be obtained in direct images.

FRODO is a complementary approach that searches for indirect evidence of defects in the form of diffraction ring patterns. These circular patterns of light appear on images of optics that are located downstream from the optics with imperfections. "The situation is analogous to dropping a pebble into a pond," says Paglieroni. "The larger the pebble, the stronger the ripples." The pebbles are defects that occur in upstream NIF optics but which manifest themselves as diffraction patterns closer to the camera. The further away from the defect, the larger are the concentric ring patterns.

The GDM algorithm has performed extremely well in analyzing images of optics and detecting diffraction patterns caused by distant flaws. NIF engineers are hopeful that this indirect approach will complement their direct inspection approach by helping detect flaws that are outside the camera focal range or for which image resolution is limited.

When attempting to detect diffraction rings, the FRODO software invokes the GDM algorithm by using a luminance disk, which is a model that produces a field of light-to-dark gradient vectors all flowing toward the center of a circle. This model is similar to those used by ICE when searching for specific objects in overhead images but has a round shape rather than a shape similar to a road or building.

The algorithm's final output is a set of estimates for ring locations on the image as well as the estimated size of the upstream defect and the distance to the optic. Kegelmeyer says, "Our goal is to demonstrate the robustness of the algorithm in real and simulated images and in images that have no defects, single defects, or multiple defects of different sizes. We have found that the algorithm can consistently and efficiently identify the associated diffraction patterns."

In certain instances, however, the algorithm encounters problems. Extremely small diffraction rings are difficult to detect as are rings cut off near the edge of an image when two rings overlap more than 50 percent. The FRODO team has optimized the algorithm parameters and the postprocessing steps to minimize the number of false alarms. "These false-alarm mitigation techniques look promising," says Kegelmeyer. Once this issue is resolved, the GDM algorithm can be incorporated into the suite of NIF optics inspection analysis codes.

Building on Success

Livermore imagery analysts' success with ICE has sparked interest by several federal agencies involved in national security. In the meantime, imagery analysts are pushing ICE to its limits by searching for increasingly smaller objects.

ICE has potential applicability to other areas of experimental science, including physics, biology, and environmental science, in which mining massive archives of complex measurement data is an essential research activity. One possibility under consideration involves using ICE for automatic extraction of data generated by the Large Synoptic Survey Telescope (LSST). (See *S&TR*, November 2005, pp. 12–18.) The telescope, scheduled for completion in 2012, will provide digital imaging of objects in deep space across the entire sky. LSST will create 24 gigabytes of data every 30 seconds, unprecedented in astronomical data gathering. Effectively managing the vast amount of LSST data is the most challenging aspect of the project. Computer scientists Miller and Garlick are involved in developing computational approaches for automatically mining the thousands of digital images that LSST will record daily.

A new LDRD Strategic Initiative, called Predictive Knowledge Systems (PKS), is building, in part, on the capabilities of ICE. PKS will pull together multiple sources of information, such as imagery, radio intercepts, and other sensor data, and correlate the data in space and time. Brase, who is principal investigator for PKS, says the project is aimed at nuclear nonproliferation and homeland security applications. In disciplines where the sheer amount of information can be overwhelming, ICE is providing a way to find the proverbial needle in a haystack. In this case, finding that needle may save lives.

— Arnie Heller

Key Words: gradient direction matching (GDM) algorithm, Image Content Engine (ICE), Large Synoptic Survey Telescope (LSST), National Ignition Facility (NIF), Predictive Knowledge Systems (PKS).

For further information contact David Paglieroni (925) 423-9295 (paglieroni1@llnl.gov).

Got Oxygen?

WITHOUT oxygen, life as we know it would not exist. And we have the lowly, much-maligned pond scum to thank.

For the first third of Earth's history—about 1.5 billion years—oxygen was locked up in molecules, and anaerobic (without oxygen) bacteria were the only life forms. Earth's atmosphere was composed mostly of carbon dioxide, sulfur dioxide, and nitrogen. Oxygen made up less than 1 percent of the atmosphere and was toxic to existing life forms.

About 3 billion years ago, for reasons no one knows, oxygenic photosynthesis slowly began to evolve in a form of pond scum, or algae, known as cyanobacteria. Cyanobacteria convert sunlight, carbon dioxide, and water into carbohydrates, producing free oxygen as a waste product. These three food groups were plentiful, and the cyanobacteria population thrived, increasing the atmosphere's oxygen content.

Earth's anaerobic microbes found the increasing oxygenation of the atmosphere challenging in the extreme. Many life forms disappeared. Others adapted by developing enzymes and processes to mediate oxygen's toxicity. The presence of oxygen offered organisms opportunities for respiration and the biosynthesis of entirely new classes of molecules.

As these adaptations proceeded, the geologic record shows that the first multicellular organisms began to appear, eventually producing the multiplicity of life forms we know today. Oxygen, especially oxygen metabolism, was key to evolution, but why?

A Livermore team sought to find at least a partial answer in a project that computationally examined metabolic "networks" in 70 modern aerobic and anaerobic genomes. This bioinformatics project was the first to specifically examine the effects of the presence or absence of oxygen on metabolic networks. The groundbreaking work was performed by Lawrence postdoctoral fellow Jason Raymond and researcher Daniel Segre, who is affiliated with both Livermore's Chemistry, Materials, and Life Sciences Directorate and Boston University.

Planting Seeds

The metabolic network expansion algorithm used by Raymond and Segre begins with a set of randomly chosen "seed" compounds that are allowed to react according to enzymatic reaction rules enumerated by the Kyoto Encyclopedia of Genes and Genomes (KEGG). The KEGG database is a collection of data from across all known genomes. These metabolic networks thus correspond not just to the reactions tenable within any one organism but to the metabolic potential of the entire biosphere, or at least that portion whose DNA has been fully characterized.

The current KEGG database encompasses 6,836 reactions extending across 70 genomes and involving 5,057 distinct compounds, which together produce a huge number of possible combinations. The

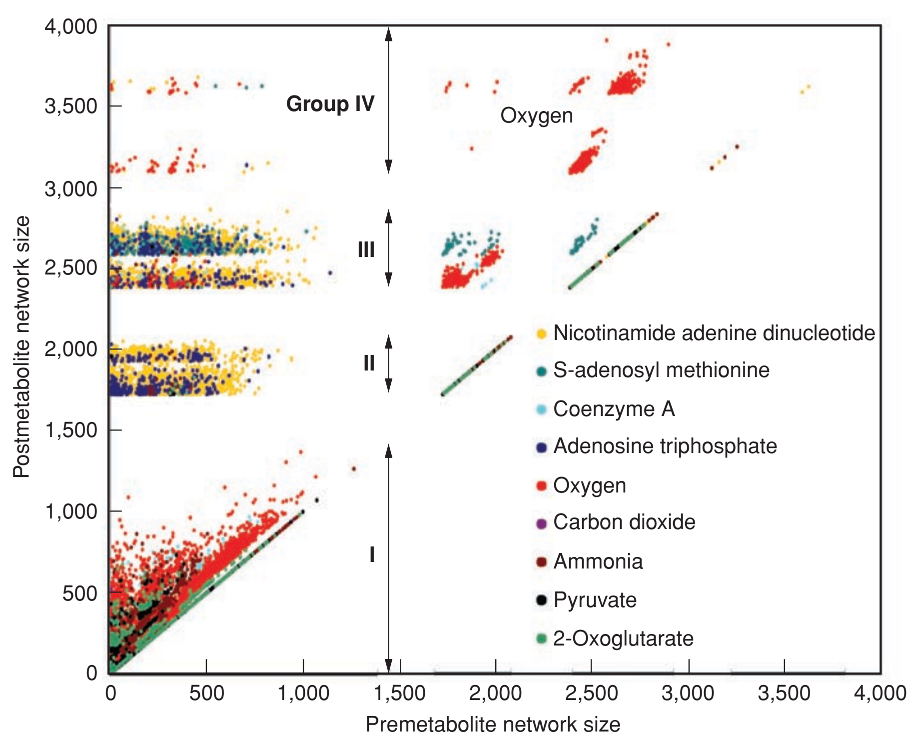

This graph shows the effects of various metabolites on the total number of reactions in metabolic networks. Only networks that include molecular oxygen as a metabolite appear in group IV, which contains the largest, most complex networks. Other transitions, from less to more complex, are determined by the availability of other metabolites.

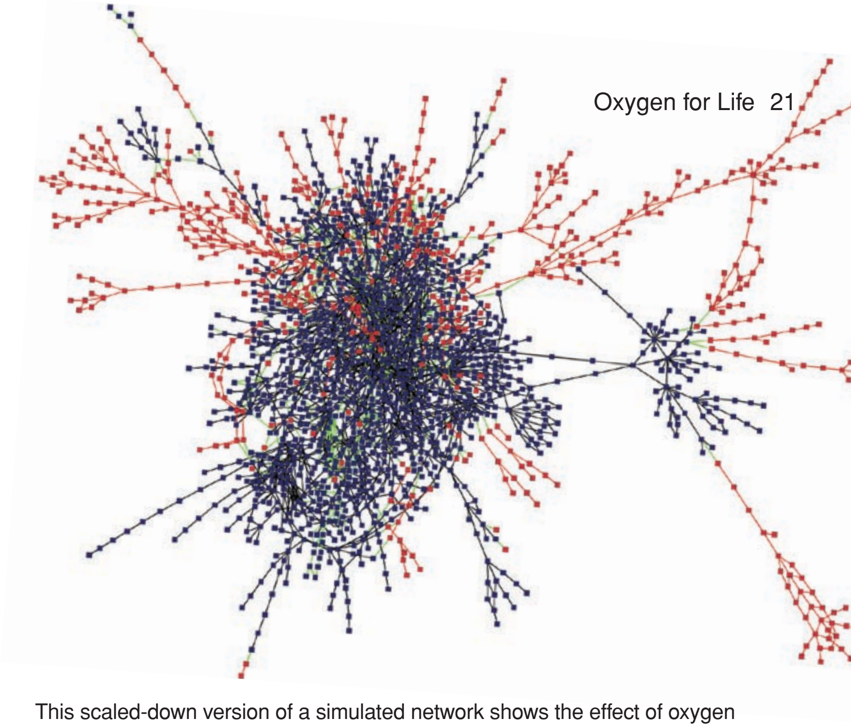

This scaled-down version of a simulated network shows the effect of oxygen on an initial seed set. Blue nodes (metabolites) and lines (reactions) represent the metabolic network before oxygen is added. Red metabolites and reactions are contingent on the availability of oxygen. Green lines represent the rewiring of pathways to allow the use of molecular oxygen.

team used Monte Carlo statistical random sampling to select a much smaller but still highly variable list of seed conditions for the simulation of about 100,000 reaction networks.

In each simulation, a metabolic reaction could occur only if all its reactants were present in the seed set. The products of those reactions then augmented the seed set, resulting in new types of reactions that could occur. This process was repeated until no new metabolic products were generated, and no other reactions were possible.

The effects of nine specific metabolites—all known by biochemists to be crucial to modern organisms—were examined. These included oxygen, carbon dioxide, and ammonia as well as nicotinic adenine dinucleotide and adenosine triphosphate, which are cofactors. These metabolites were first withheld from the metabolic networks and then added in order to examine the effects of their exclusion or inclusion. Although oxygen is only the third most often used compound in metabolism, it is by far the most effective at increasing the number of reactions that occur within large metabolic networks.

"We knew oxygen would matter," says Raymond, "but we also wanted to know how other compounds, such as coenzyme A, would increase complexity. We found that only with oxygen do the largest, most complex networks appear."

Four Simple Groups

When the researchers compared simulations of seed sets that included oxygen with those that did not, oxygen's effect on the overall metabolic system quickly became apparent. "We expected every type of combination of metabolites to be different," says Raymond. "However, we were surprised to find that all the combinations fell into just four distinct groups."

Despite the widely varying initial conditions of networks within the four groups, more than 95 percent of the reactions and metabolites within each group are identical. The networks in smaller groups are mostly nested within the four larger groups; for example, most reactions and metabolites in group II networks are a subset of those in group III and IV networks. This grouping is a consequence of the special architecture of metabolic networks, which are composed of highly connected "hub" metabolites that, when created, can be enzymatically converted to many other compounds.

All of the networks in group IV, which had the largest and most interconnected networks, include molecular oxygen. Networks from this group contain as many as 1,000 more reactions than the networks generated without oxygen. About half of these oxygen-dependent—or oxic—reactions do not explicitly use oxygen but belong to metabolic pathways that rely on it, much like the early stages in an assembly line are crucial to later steps. These pathway-level effects would not have been found in studies of individual reactions.

The image above right is an example of the effect of oxygen on a group III network, a network that is quite similar to that found in modern anaerobic organisms. This network includes 1,861 metabolites and 2,652 reactions that represent the many ways in which simple carbon compounds, starting with carbon dioxide, are transformed into essential molecules such as amino and nucleic acids, lipids, and cofactors. The blue nodes (metabolites) and lines (reactions) represent the group III metabolic network before the addition of oxygen. The red nodes and lines are oxic network metabolites and reactions—the end transformation of the group III network into a much larger and more connected group IV network that is analogous to the metabolic networks of aerobic organisms.

Although the total number of reactions and metabolites increased by 150 percent in the oxic network, the density of the network—the average number of lines between nodes—increased only slightly. One example of this increase is highlighted by the green lines in the network, which correspond to reactions that are found only in an oxic network but use at least one anoxic metabolite. These reactions represent the replacement or "rewiring" of anoxic pathways to take advantage of oxygen. An example of rewiring is in the biosynthesis of thiamin and vitamin B12, which are cofactors required by many anaerobic and aerobic organisms but synthesized more efficiently in aerobes using oxygen.

This evolution of new reactions and pathways makes available an entirely new set of metabolites that in turn promote the evolution of new types of enzymatic reactions. When single-celled microbes ruled the biosphere over 2 billion years ago, this kind of adaptation may have provided the molecular underpinnings for the development of increasingly complex, multicellular—and oxygen-dependent—life.

—Katie Walter

Key Words: anaerobic organisms, bioinformatics, evolution, metabolic network expansion, oxygen, oxygenation.

For further information contact Jason Raymond (925) 423-6507 (raymond20@llnl.gov).

A Shocking New Form of Laserlike Light

A Livermore scientist has predicted a new method for producing coherent light, one of the few since the invention of lasers nearly half a century ago. This novel, laserlike light is made by whacking certain kinds of crystals with a sharp mechanical shock.

"In the past, we expected to generate mostly random photons and possibly some 'sparks' when shocking a crystal," says Livermore's Evan Reed, who came up with the mechanical-shock concept for generating coherent light. "We found one can also produce coherent light in the terahertz region—a frequency band below infrared but above that used by cell phones." Narrow bandwidth radiation in the terahertz range, like that generated in these shocked crystals, is an area of interest and active research at the Laboratory. Coherent radiation could potentially be used in a host of applications, such as a diagnostic for shock waves and for explosives detection. Other applications may also arise. When lasers were invented in 1958, they were considered novelties with limited practical use, yet laser light is now a part of everyday life.

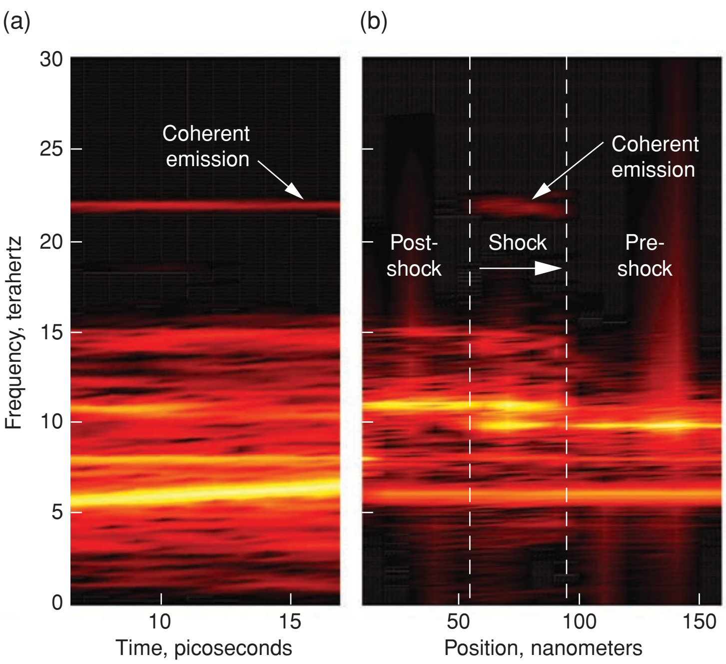

In a molecular dynamics simulation of shocked sodium chloride, coherent light is emitted at 22 terahertz. (a) The emission of the light is shown as a function of time while the shock is propagating. (b) The generated radiation is shown as a function of location within the shocked crystal, indicating the 22-terahertz coherent signal is generated at the shock front (between the dotted lines).

Coherently Speaking of Light

Anyone who has observed the red light from a laser pointer or a grocery store checkout scanner has seen coherent light in the visible part of the spectrum. The photons emitted in coherent light all have the same wavelength and frequency. (In incoherent light, such as that emitted from a light bulb, the wavelengths vary randomly.) Not only are the photons at the same wavelength in coherent light, the waves are synchronized, rising and falling in concert. This characteristic is what gives laser light its special properties and its high brightness.

In lasers, a pulse of light is focused on a lasing material, and the light's energy is absorbed by the electrons in the material. The electrons "hop up" to a specific, higher energy level, then fall back down to their normal energy state when stimulated by a photon, releasing a photon of the same frequency. Thus, a single photon stimulates others to be released simultaneously. This phenomenon is what gives lasers their name: light amplification by stimulated emission of radiation. The mechanism by which a shocked crystal releases coherent light is quite different.

The Crystallization of an Idea

Reed's research on coherent light stems from work he performed on shock waves propagating through photonic crystals while he was a graduate student at the Massachusetts Institute of Technology (MIT). Two years ago, Reed came to Livermore under a Lawrence Fellowship. (For more information about Lawrence Fellowships, see *S&TR*, November 2002, pp. 12–18.)

In the lattice of photonic crystal structures are "energy gaps" where certain bands of energy are not allowed, and thus no photons of corresponding frequencies are allowed to propagate through. "I thought it might be possible to get these crystals to emit light of specific frequencies, and the frequencies would be in the terahertz range because of the small lattice constants of ionic crystals," says Reed. "I began testing this theory to see if those possibilities might have some physical basis."

In his calculations, Reed used molecular dynamics and Maxwell-equation computer simulations, the LAMMPS code, and several other Livermore-developed codes to approach the problem. To ensure accurate simulation results, Reed needed to model tens of millions of atoms for tens of picoseconds (1 picosecond is a millionth of a millionth second). A simulation of this size required the capabilities of Livermore's 23-trillion-floating-point-operations-per-second Thunder supercomputer. Using Thunder, Reed and his colleagues, Richard Gee from Livermore's Chemistry, Materials, and Life Sciences Directorate and Marin Sojacic and J. D. Joannopoulos from MIT, modeled what would happen if dielectric crystal arrays were subjected to a sharp mechanical shock. For their simulations, Reed and his colleagues used crystals of sodium chloride (regular table salt) and then launched a virtual shock wave of the sort produced by laser light, a gas gun, or a projectile at the highly ordered arrays of atoms.

"We hit the crystal lattice arrays with a high-amplitude planar shock wave—a wave with a 'flat' front—and watched how the lattice responded," says Reed. The wave front excites the atoms in the first plane, causing them to move in a synchronized manner. The periodic motion of the atoms is then replicated in the next plane as the planar shock moves forward through the array. "Because the salt atoms carry a charge, the movement of these charged particles creates a temporally periodic electric current in the crystal," says Reed. "The charged particles move together, with the same amplitude and frequency, producing coherent radiation, or laserlike light."

This effect, Reed noted, occurs at the front of the shock wave. The frequency of the emitted light is determined by the shock wave speed and the crystal lattice structure. "Most of the Laboratory's research on effects of shock waves focuses on what happens long after the front has passed," he says. "A great deal of work is being done on turbulence, chemical reactions in explosives, and other effects of shock waves on a material. Calculating and experimentally measuring what happens at the front of the shock wave is an entirely different area of research."

Results of the simulations showed that, indeed, coherent light was produced in the 1- to 100-terahertz range. This part of the spectrum has not been examined in shock experiments, because most shock wave probes are either at much higher or lower frequencies. "Detectors in this range are currently much less sensitive than those for other frequencies," says Reed.

Stepping Up to Experimentation

"I didn't know in the beginning if anything would come of the theory," says Reed. "The Laboratory invested in this basic science question, and now it's come to fruition." Reed and his colleagues from the Laboratory and MIT are moving on to the experimental phase of the research. Reed recently received funding from Livermore's Laboratory Directed Research and Development Program to conduct shock-induced light experiments.

"We have gained confidence through our simulations that radiation will be present given certain experimental parameters," says Reed. "We're using well-established tools such as molecular dynamics simulations. We have just stretched their use to a different regime, calculating electromagnetic signatures of nonequilibrium phenomenon that have not been examined before." Reed,

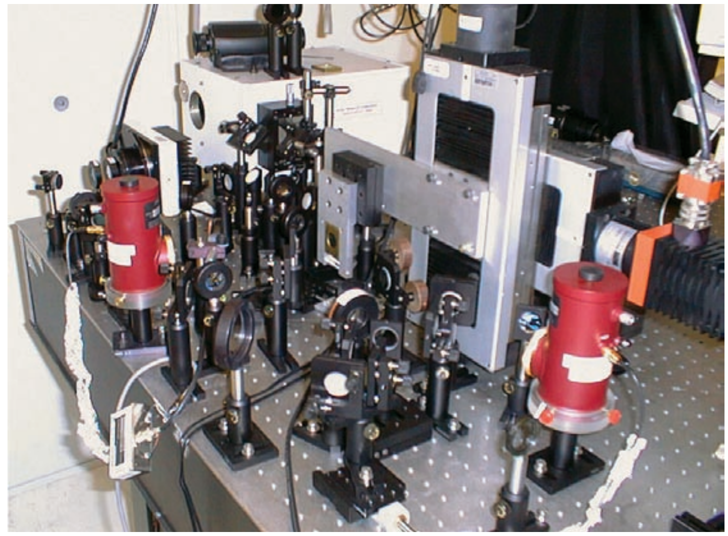

Researchers at Livermore and the Massachusetts Institute of Technology will use this experimental apparatus to attempt to produce coherent light from mechanically shocked crystals.

postdoctoral researcher Michael Armstrong, and researchers from Los Alamos National Laboratory are collaborating on experiments at the Los Alamos Center for Integrated Nanotechnologies to search for this signal in the terahertz range.

A number of possible applications beckon. Laserlike light in the terahertz range could be used as a diagnostic for understanding shock waves, including shock speed, and the crystallinity of a substance. Laserlike light in this frequency range could also propagate through walls and containers, enabling a unique imaging capability for security applications or aiding the remote sensing of explosives. In addition, signals could be used in spectroscopy to provide a window on a relatively unknown part of the spectrum.

Smash a crystal with a shock wave and probably the last thing one would expect to emerge is laserlike light. Yet, the physical world is full of surprises—and it's those surprising twists and turns that take scientists into unknown territory and often yield the most interesting results. Sometimes, it's a matter of keeping one's eyes open for the unexpected light.

—Ann Parker

Key Words: coherent light, crystal lattice arrays, ionic crystals, laser, molecular dynamics simulations, shock wave, terahertz frequency, Thunder supercomputer.

For further information contact Evan Reed (925) 424-4080 (reed23@llnl.gov).

Each month in this space, we report on the patents issued to and/or the awards received by Laboratory employees. Our goal is to showcase the distinguished scientific and technical achievements of our employees as well as to indicate the scale and scope of the work done at the Laboratory.

Patents

System and Method for Characterizing Voiced Excitations of Speech and Acoustic Signals, Removing Acoustic Noise from Speech, and Synthesizing Speech
Greg C. Burnett, John F. Holzrichter, Lawrence C. Ng
U.S. Patent 7,089,177 B2
August 8, 2006
This system can characterize human or animate voiced excitation functions and acoustic signals in order to remove unwanted acoustic noise that often occurs when a speaker uses a microphone in common environments. The system can also synthesize personalized or modified human or other animate speech on command from a controller. A low-power electromagnetic (EM) sensor detects the motions of windpipe tissues in the glottal region of the human speech system before, during, and after a user produces speech. A voiced excitation function is derived from these tissue motion measurements. The excitation function provides information to enhance noise removal from human speech and allows accurate transfer functions of speech to be obtained. Previously stored excitation and transfer functions can be used to synthesize personalized or modified human speech. The EM sensor and acoustic microphone systems can be configured to enhance noise cancellation and enable multiple articulator measurements.

Awards

Science Spectrum magazine has recognized three Laboratory scientists as being top performers in their respective fields. The award recipients were honored in September 2006 at the Minorities in Research Science Conference in Baltimore, Maryland.

Robert Shepard of the Administration and Human Resources Directorate was awarded the **Emerald Honors Educational Leadership Award**. Shepard is the founding executive director of the Science and Engineering Alliance, Inc. (SEA), a nonprofit education research consortium of four historically black colleges and universities. SEA, headquartered in Washington, DC, collaborates on research projects for government agencies, the private sector, and other universities to increase participation of faculty at underrepresented academic institutions in a broad range of scientific and technical disciplines. Shepard is a long-time advocate for recruiting and retaining minorities in the Livermore technical workforce.

Hope Ishii of the Physics and Advanced Technologies Directorate was honored with a **Trailblazer Award**, which recognizes outstanding minority and women professionals whose leadership and innovative thinking on the job and in the community extend throughout and beyond their industry. Ishii studies astromaterials and has developed a method for extracting comet dust from the National Aeronautics and Space Administration's Stardust mission that was returned to Earth in January 2006.

Dean Williams of the Computation Directorate was also selected to receive a **Trailblazer Award**. Williams leads a variety of software projects aimed at data analysis and visualization of climate models. He is also active in Livermore's recruiting efforts and participates in the summer student program to help ensure the Laboratory continues to hire qualified minorities into scientific fields.

Livermore's **Plant Engineering Department** has won two awards for a construction safety video promoting a safety-first culture at the Laboratory. "Contractors Construction Safety Orientation at LLNL" earned both an **Aegis Award** and a **DV Award**, which is sponsored by *Digital Video Magazine*.

The video, which communicates the importance of creating and sustaining a safe and productive work environment, is required viewing for all subcontractors performing construction work at Livermore. The video has been produced in both English and Spanish.

The Aegis Award recognizes excellence in video and film productions by such groups as news and media organizations, government agencies, universities, and Fortune 500 companies. The DV Award acknowledges outstanding creative and technical achievements in digital video production such as independent films, commercials, corporate communications, and industrial videos.

Abstracts

Thunder’s Power Delivers Breakthrough Science

Livermore’s Multiprogrammatic and Institutional Computing (M&IC) Program provides Laboratory researchers with powerful supercomputers for unclassified, mission-related research. The program’s 23-trillion-floating-point-operations-per-second Thunder supercomputer is being used for five Laboratory projects, called Grand Challenges, whose calculations require a significant percentage of the system. One Grand Challenge is modeling regional climate change at resolutions of about 10 kilometers. Another team is resolving long-standing questions about water’s behavior, including the coexistence of phases for liquid water and vapor and a superionic phase of water that is believed to exist in large planets. Another Grand Challenge team is helping experimentalists better understand the complex processes involved in protein folding by simulating folding over a broad range of temperatures. Researchers are also modeling in unprecedented detail the microstructural origins of stress–strain behavior in metals to predict material performance. In addition, simulations on Thunder are helping researchers design chemical and biological detection instruments for homeland security.

Contact:
Brian Carnes (925) 423-9181 (carnes1@llnl.gov).

Extracting Key Content from Images

A team of Livermore engineers, computer scientists, and physicists has developed the Image Content Engine (ICE), which allows analysts to search massive volumes of image data in a timely fashion. The system guides researchers to areas in the images that likely contain the objects for which they are searching. ICE encompasses a new approach for computer-aided extraction of specific content information. Imagery analysts in the Laboratory’s Nonproliferation, Homeland and International Security Directorate are using ICE to spot items such as specific types of buildings, vehicles, or roads. A variation of ICE is being evaluated for use by the National Ignition Facility’s Optics Inspection Analysis Team. The system may also have application to other fields, including nondestructive analysis, biological imaging, astronomy, and supercomputer simulations of nuclear weapon performance.

Contact:
David Paglieroni (925) 423-9295 (paglieroni1@llnl.gov).

U.S. Government Printing Office: 2006/670-074-10023

A Fast Diagnostic for Flu

A Livermore device can diagnose influenza in just two hours and distinguish between five respiratory viruses.

Also in December

- *A multiagency team led by Lawrence Livermore and Sandia national laboratories is developing restoration procedures to decontaminate a major airport after an attack with biological weapons.*

- *A new technique detects small changes in skeletal calcium that may signal bone disease.*

- *Laboratory scientists are designing a system that uses gamma rays to look deep inside the atom.*

Coming Next Month

University of California
Science & Technology Review
Lawrence Livermore National Laboratory
P.O. Box 808, L-664
Livermore, California 94551