Lawrence Livermore National Laboratory

# Science & Technology Review

July/August 2014

## ACCELERATING DRUG DISCOVERY

*Also in this issue:*

**Chaotic Solar System Birth**

**Advanced Computing Tools**

**Reducing the Hydrofracturing Footprint**

## About the Cover

Computational chemistry efforts, such as those described in the article beginning on p. 4, may help streamline the drug discovery process, thereby enabling researchers to bring new therapies to clinical trials and the marketplace more rapidly and with a higher rate of success. This work is also relevant to national security, because broad-spectrum antibiotics can often serve a dual role as countermeasures for bioterrorism threats. Livermore researchers are using supercomputer-based modeling and simulation to produce a robust predictive capability that will help support and guide the drug development process. On the cover is a snapshot from a simulation created by computational biochemist Tim Carpenter of a cytochrome P450, a protein that is crucial in breaking down drugs in the body.

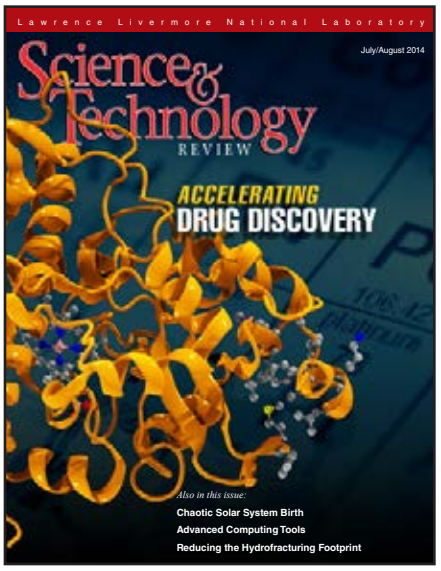


Cover design: Amy E. Henke

## About *S&TR*

At Lawrence Livermore National Laboratory, we focus on science and technology research to ensure 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 eight 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.

The Laboratory is operated by Lawrence Livermore National Security, LLC (LLNS), for the Department of Energy's National Nuclear Security Administration. LLNS is a partnership involving Bechtel National, University of California, Babcock & Wilcox, Washington Division of URS Corporation, and Battelle in affiliation with Texas A&M University. More information about LLNS is available online at www.llnsllc.com.

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 Web at str.llnl.gov.



National Nuclear Security Administration

Prepared by LLNL under contract
DE-AC52-07NA27344

Science & Technology REVIEW

July/August 2014

Lawrence Livermore National Laboratory

# Contents

**S&TR Staff**

**Scientific Editor**
Gregory C. Burton

**Managing Editor**
Ray Marazzi

**Publication Editor**
Pamela MacGregor

**Writers**
Rose Hansen, Arnie Heller, Caryn Meissner, and Katie Walter

**Art Director**
Amy E. Henke

**Proofreader**
Carolin Middleton

**Print Coordinator**
Charlie M. Arteago, Jr.

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

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

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-52000-14-7/8
Distribution Category UC-99
July/August 2014

## Feature



## Research Highlights

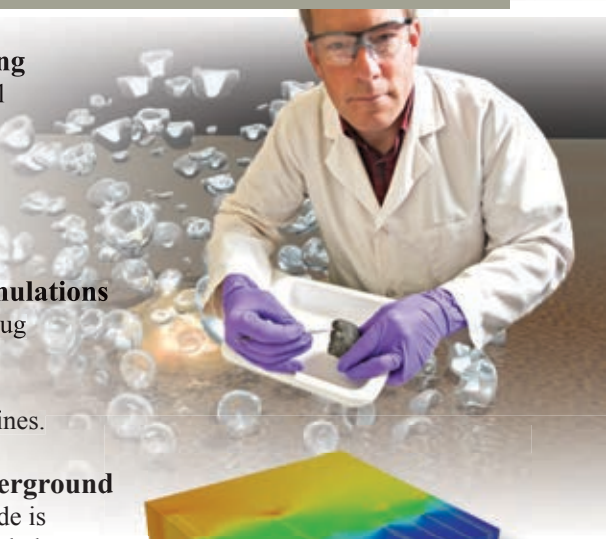



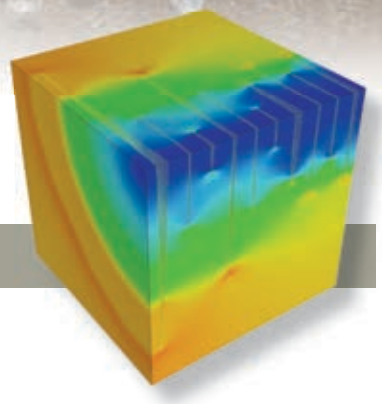



## Departments

## Light Shines on Novel Additive Manufacturing Approach

For nearly a century, electrophoretic deposition (EPD) has been used to coat materials by depositing particles of various substances onto the surfaces of items such as ceramics, metals, polymers, and even living cells. A limitation of EPD is that material can be deposited only across an entire surface and not in specific, predetermined locations. A team of researchers at Lawrence Livermore has overcome this limitation with a new technique called light-directed EPD, which uses photoconductive electrodes and direct-current electric fields to dynamically pattern surface material.

Light-directed EPD allows the buildup of material in targeted areas where the light comes in contact with the photoconductor's surface, thus enabling the creation of arbitrarily patterned, three-dimensional multimaterial composites over large areas with fine resolution. To create a proof-of-concept logo (shown center) using the new additive manufacturing process, the researchers first deposited a layer of tungsten nanoparticles on surface areas illuminated through a laser-cut aluminum mask. The mask was then changed, along with the solution of nanoparticles, to deposit alumina ceramic material. In the future, the masks will be replaced by a digitally projected mask for a completely automated deposition system.

Light-directed EPD has the potential to elevate traditional EPD from a single-layer, single-material coating process to a true additive manufacturing technique that allows for unique composites to be formed. For example, void areas can be precisely created in a part to control polymer material behaviors for energy absorption, or such areas can be formed within cellular material to create veins for manufactured organs. "This work represents a large step in advancing electrophoretic deposition as a method of fabricating complex, three-dimensional patterned composites," says Andrew Pascall, engineer and lead author of the team's paper published in the April 9, 2014, edition of *Advanced Materials*.

***Contact: Andrew Pascall (925) 423-1926 (pascall1@llnl.gov).***

## Big Data Supercomputer for Collaborative Research

Catalyst, Lawrence Livermore's latest supercomputer, is available to industry and academia collaborators for testing emergent big data technologies, architectures, and applications. A resource for the National Nuclear Security Administration's (NNSA's) Advanced Simulation and Computing (ASC) Program, Catalyst represents a major departure from classic simulation-based computing architectures common at Department of Energy laboratories. Its architecture is modified from a Cray® CS300™ high-performance computing cluster. Developed by a partnership with Lawrence Livermore, Cray, and Intel, the machine boasts nearly a terabyte ($10^{12}$ bytes) of volatile and nonvolatile (NVRAM, or flash) memory per compute node, creating more than 300 terabytes of system memory. The 150-teraflops (trillion floating-point operations per second) cluster runs the NNSA-funded Tri-Lab Open Source Software, or TOSS, which provides a common user environment across NNSA tri-lab clusters (Los Alamos, Sandia, and Lawrence Livermore national laboratories).

Catalyst's novel architecture opens new opportunities to combine floating-point focused capability with data analysis in one system, providing insights into the technologies the ASC Program will require over the next decade to meet mission needs in high-performance simulation and big data computing. The broad range of big data problems Catalyst is well suited to address includes bioinformatics, business analytics, graph networks, machine learning, and natural language processing.

"Our purpose is to use Catalyst as a test bed to develop optimization strategies for data-intensive computing," says Fred Streitz, director of Livermore's High Performance Computing Innovation Center, which manages private sector access to the system. "We believe that advancing big data technology is a key to accelerating the innovation that underpins our economic vitality and global competitiveness."

***Contact: Fred Streitz (925) 423-3236 (streitz1@llnl.gov).***

## Americans Use More Energy in 2013

Each year, the Laboratory releases energy flowcharts that illustrate the source and amount of energy consumed by the nation. According to the most recent energy flowcharts (flowcharts.llnl.gov), use of renewable, fossil, and nuclear energy increased in 2013.

Wind energy continued to grow steadily, increasing 18 percent from 1.36 quadrillion British thermal units, or quads, in 2012 to 1.6 quads in 2013. New wind farms continue to come online with bigger, more efficient turbines. Also, the transportation sector used more renewable energy, specifically biomass that is converted to ethanol. Natural gas prices rose slightly in 2013, reversing some of the recent shift from coal to gas in the electricity production sector. Overall, natural gas consumption increased by 0.6 quads. Losses in the electricity sector were more than offset by increased gas use in the residential, commercial, and industrial sectors. "2013 was a cold winter," says A. J. Simon, group leader for Energy at Lawrence Livermore. "We expect to see continued high gas consumption in 2014 due to another tough winter on the East Coast." Petroleum use also increased last year. With oil prices remaining relatively constant, this increase was likely due to the nation's modest economic expansion. Finally, nuclear energy use was greater in 2013. Simon says, "It's likely fewer reactors were down last year for refueling than in previous years."

The majority of energy consumption in 2013 was for electricity generation (38.1 quads), followed by the transportation, industrial, residential, and commercial sectors. Energy use in the residential, commercial, and industrial sectors all increased slightly. Overall, Americans used 2.3 quads more in 2013 than the previous year.

***Contact: A. J. Simon (925) 422-9862 (simon19@llnl.gov).***

Commentary by Glenn A. Fox

# Combating Biological Threats through Drug Development

CONCERNS about underground nuclear testing and the effects of ionizing radiation on humans prompted the Atomic Energy Commission to establish a biological research program at Lawrence Livermore in 1963. Since then, the program's mission has evolved to address changing national needs, while making important scientific breakthroughs in the process. For instance, the early work on radiation led to the discovery of flow sorting and chromosome painting, which enabled researchers to study DNA damage and the creation of chromosome-specific clone libraries in new ways. The Human Genome Project, the largest biological research project ever undertaken, followed these key discoveries.

In the years since 9/11, Livermore biologists have concentrated their efforts on detecting, assessing, and combating biological threats, both natural and deliberately developed. Solving national problems in biosecurity and human health demands highly collaborative interdisciplinary teams working at the frontiers of basic and applied sciences. The Biosciences and Biotechnology Division is the nexus for biology research at the Laboratory. Our biologists also work closely with chemists, engineers, physicists, and computer scientists from across the Laboratory on endeavors such as delivering technologies that rapidly detect a pathogen once it is released, processing samples with possible bioterrorist agents, cleaning contaminated facilities, and treating people exposed to pathogens, to name just a few.

The Laboratory has a distinguished track record in developing and deploying biosecurity and biomedical capabilities. Many of these efforts are enabled by the tools and expertise available through our biology-related capability centers. For instance, Livermore's Select Agent Center, home to the only Department of Energy Biosafety Level 3 Facility, allows researchers to conduct experiments on a wide range of microorganisms for developing much-needed biodetection capabilities. The Center for Accelerator Mass Spectrometry, a particularly versatile facility, is used to perform a Livermore-developed technique called microdosing, which allows biomedical scientists to safely test in humans whether the active ingredients of a drug are absorbed and how long each ingredient persists in the body. Researchers at the Livermore Microarray Center can analyze and compare the concurrent activity of thousands of genes at a time, advancing research in such areas as cell response to radiation exposure and the genetic makeup of plague-causing bacteria. Finally, the Forensic Science Center is home to nationally recognized experts who often work at the intersection between biology and chemistry in support of various counterterrorism efforts.

Of growing importance to our biological program are Livermore's world-class computational resources. Only recently have we begun to see computer modeling and simulation approach the level of experiment and capture some of the complexity of biology. We are just scratching the surface now, but future generations of high-performance computing (HPC) systems will serve as an equal and crucial partner to experiment. Accomplishing accurate biological simulations requires significant computational horsepower and advanced codes that have been fine-tuned to run on these machines. Our biologists face a formidable computational challenge in analyzing the reams of data produced by modern biology experiments—an effort well-aligned with Livermore's growing emphasis on big data analysis.

The tantalizing possibilities that HPC can offer biomedical research are exemplified by the research efforts discussed in the article beginning on p. 4. Here, we describe how Livermore researchers are expediting the development of drugs and medical countermeasures for new pathogens by addressing key scientific barriers in the drug discovery and development process. The current process is not sufficiently agile to respond to unexpected biological threats in a timely fashion. Bringing a drug to market can take 10 years or more, while new threats can arise suddenly. The 1918 influenza pandemic, for instance, may have killed as many as 25 million people in its first 25 weeks. Fortunately, the work our researchers are conducting suggests that portions of this process can be significantly accelerated with the help of HPC-based modeling and computational expertise. More broadly, performing biology research in silico will become increasingly important in the coming years, as Livermore biologists pursue challenging research agendas in the areas of human health, agriculture, and biosecurity.

■ Glenn A. Fox is the associate director for Physical and Life Sciences.

# A New Model for Pharmaceutical

*Supercomputing-based modeling may help validate and accelerate drug research.*

THE emergence in recent years of multidrug-resistant bacteria, which Livermore physicist Monte LaBute describes as "nature's own bioterrorists," underscores both the challenge of antibiotics development and the need for novel antibiotics. However, the pharmaceutical industry considers antibiotics in particular a poor investment because they are prescribed for only a short period, while drugs that treat a chronic condition such as high blood pressure may be prescribed for the remainder of a patient's life. Drug companies often devote 10 to 15 years and more than a billion dollars to bring a new drug to market. This cost and time commitment can act as a powerful disincentive to developing new drug therapies.

Unless antibiotic development is revitalized, a growing portion of infections, particularly in hospitals, may become impossible to treat. In addition, bioterrorists could potentially introduce new bacterial threats, for instance, by releasing multidrug-resistant microbes in a large city. In the event of such an attack or of a sudden but natural emergence of a new bacterial strain, the current drug development process and the limited industry interest in antibiotic development would make a timely response difficult.

Pharmaceutical companies need a faster and more accurate way to identify promising drug compounds and evaluate the efficacy and safety of new drugs.

# Research

Such a capability could reduce costs and risks, thus allowing companies to bring antibiotics and other drugs more quickly to market. The solution, according to Livermore researchers, may lie in supercomputer-based modeling and simulation.

While pharmaceutical companies do engage in computer-aided drug design, most only have access to personal computers or midsize clusters with a few hundred cores, which are not powerful enough to run comprehensive simulations in a reasonable time frame. The companies may also lack the expertise to use these tools effectively. Livermore computational biologist Felice Lightstone and her team in the Physical and Life Sciences Directorate are pioneering high-performance computing (HPC) modeling techniques that they hope will accelerate the development of medical countermeasures such as antibiotics.

In 2011, Lightstone's team collaborated with Trius Therapeutics, Inc., to develop what may be the first new class of antibiotics in 30 years. In the process, the team demonstrated how computer-based screening can minimize costly laboratory experiments and help shorten the chemistry phase of drug development from three or four years to just six months. Lightstone notes, "Our utility is in the speed and the number of calculations we can do. Sometimes the need is for

Computational chemistry efforts such as those by Felice Lightstone's team at Lawrence Livermore may help streamline the drug discovery process, thereby enabling researchers to bring new therapies to clinical trials and the marketplace more rapidly and with a higher rate of success.

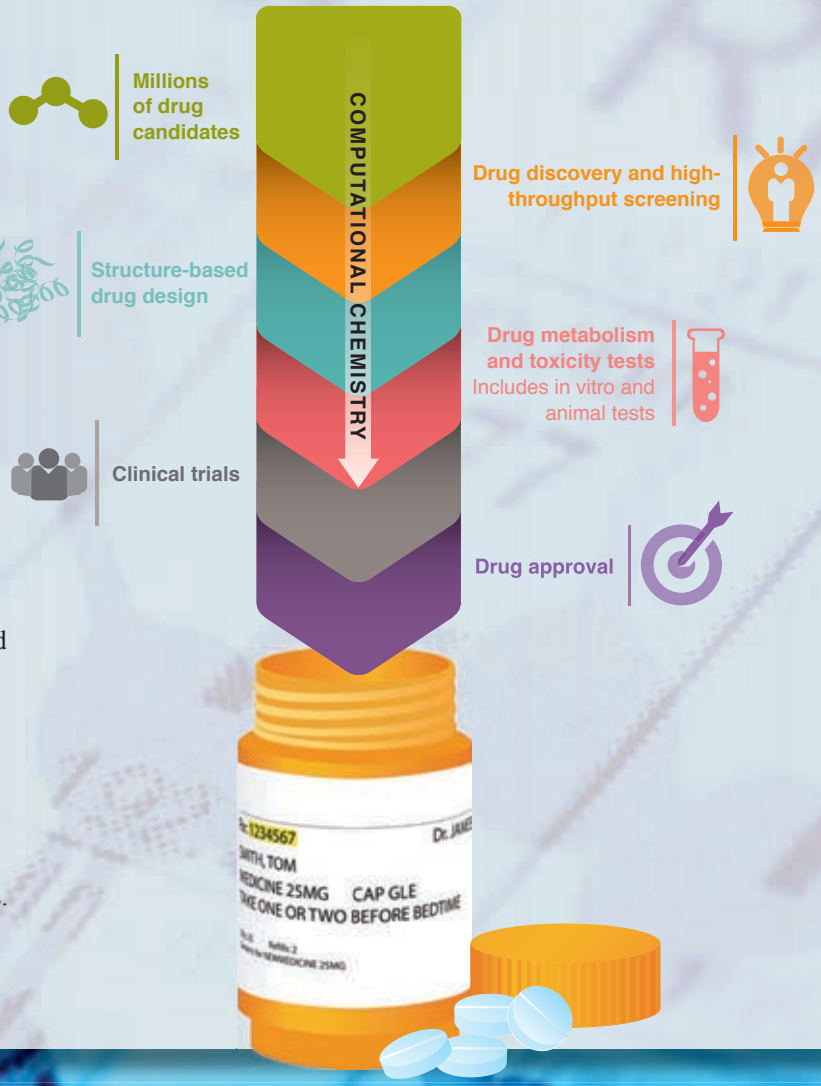

many small calculations and other times for a massively large one. We have the computing resources and expertise at Livermore to do both."

Through funding from Livermore's Laboratory Directed Research and Development Program, Lightstone's team is now in its third and final year of a project that shows promise in helping prioritize therapeutic candidates and mitigate the risk of failure in clinical trials. Computational biologist Sergio Wong explains, "Following our work with Trius involving the chemistry phase, or drug design phase, we realized that improving the success rate of drugs in clinical trials would be even more beneficial. Years of clinical trials and potential failures still lie ahead after the six-month design phase. Ninety-five percent of medicines studied in humans during clinical trials fail to be both safe and effective. Reducing that failure rate to even 70 percent would result in a huge savings in both cost and time for drug companies." Lightstone's project includes several thrust areas aimed at quickly zeroing in on which drug candidates are likely to be effective. Another focus is to enumerate the side effects, known as adverse drug reactions (ADRs), of a particular drug.

## Virtual Screening, In Parallel

Most drug compounds consist of molecules that work by binding with a protein receptor and either activating or inhibiting its behavior. Virtual screening is a computational technique used to identify which of the 30 million or more candidate drugs contained in publicly available chemical and pharmaceutical databases are most likely to bind to a targeted receptor. This technique can help to efficiently reduce the list of drug candidates to a manageable number for synthesis and testing.

Virtual screening includes two main steps: molecular docking (see *S&TR*, April 2001, pp. 4–11; June 2002, pp. 4–10) and rescoring based on docking results. Postdoctoral researcher Xiaohua Zhang explains, "Molecular docking predicts which molecules are most likely to interact favorably with a particular protein receptor. Then a more accurate and computationally demanding molecular mechanics method reevaluates the top-ranking combinations."

Most molecular docking programs run on personal computers or small workstations and scale poorly to high-performance systems, thus limiting the number of molecules that can be screened within a reasonable time frame. To expand the use of this potentially powerful research tool, Zhang and colleagues have optimized a popular docking program for parallel computing. The resulting application—called VinaLC, where LC stands for Livermore Computing—implements a hybrid programming scheme to coordinate work and keep available processors busy. It also collects results from thousands of processes without overwhelming interprocess communication channels.

In tests on several Livermore supercomputers, including the 20-petaflop (quadrillion floating-point operations) Sequoia system, VinaLC demonstrated excellent scaling results. It easily outperformed the few other parallel docking programs that can run on HPC systems without sacrificing accuracy.

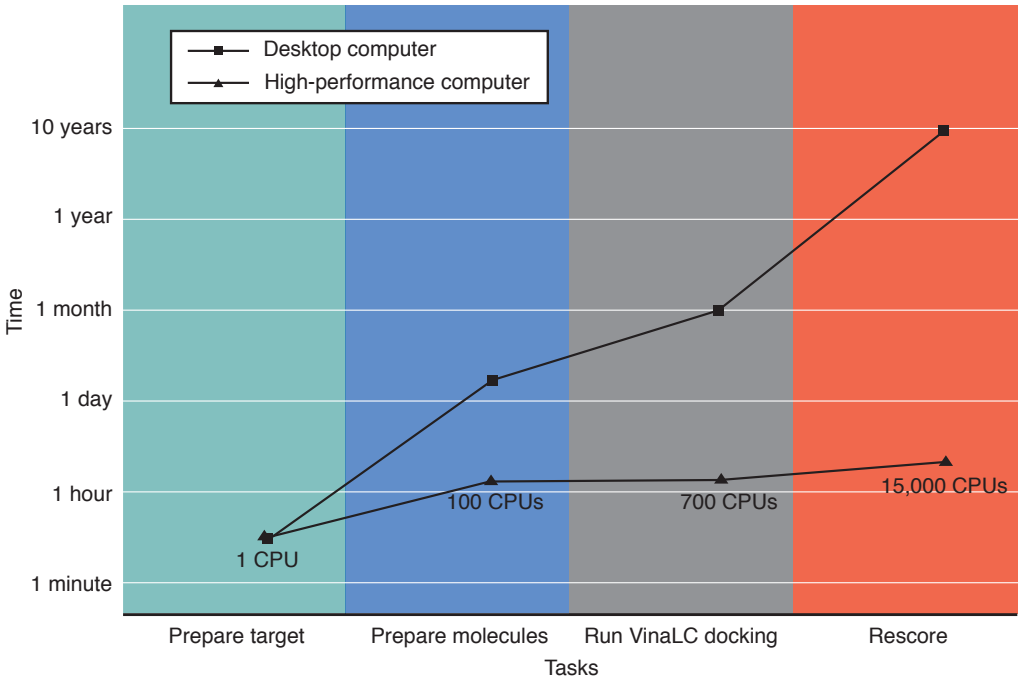


High-performance computing (HPC) systems can drastically reduce the current virtual screening time frame and increase the feasibility of more accurately screening extremely large compound databases. In this example, researchers used Livermore's parallel docking program VinaLC to test about 40,000 molecules against a single drug target. They then used a more computationally intensive method to rescore (reevaluate) the top 20 docking poses. These high-fidelity simulations calculated the free energy in the system. (CPU is central processing unit.)

The code completed 1 million flexible compound molecular docking calculations in just 1.4 hours using 15,000 central processing units (CPUs) on Livermore's Sierra supercomputer.

Lightstone's team has also integrated VinaLC into its in-house, high-volume virtual screening system. This system automates and expedites molecule preparation, binding-site selection, docking calculations, and rescoring of the top pairings based on docking results. All applications in the system have been optimized for HPC. With Sierra, the researchers rescored a total of 700,000 compounds against 38 protein targets, in what may be the largest rescoring calculation to date. With such a large number of calculations, the team was able to statistically determine the optimal number of docking poses—the orientation of the compound with respect to the target—that should be kept for rescoring. The team found that 5 to 10 poses per compound provided a good compromise between accuracy and computational expense.

These efforts demonstrated that running calculations on a supercomputer dramatically reduces the time it takes to perform virtual screening, enabling the comprehensive evaluation of larger pools of candidate drugs. The team is now working to enhance the virtual screening system's speed, accuracy, and ease of use.

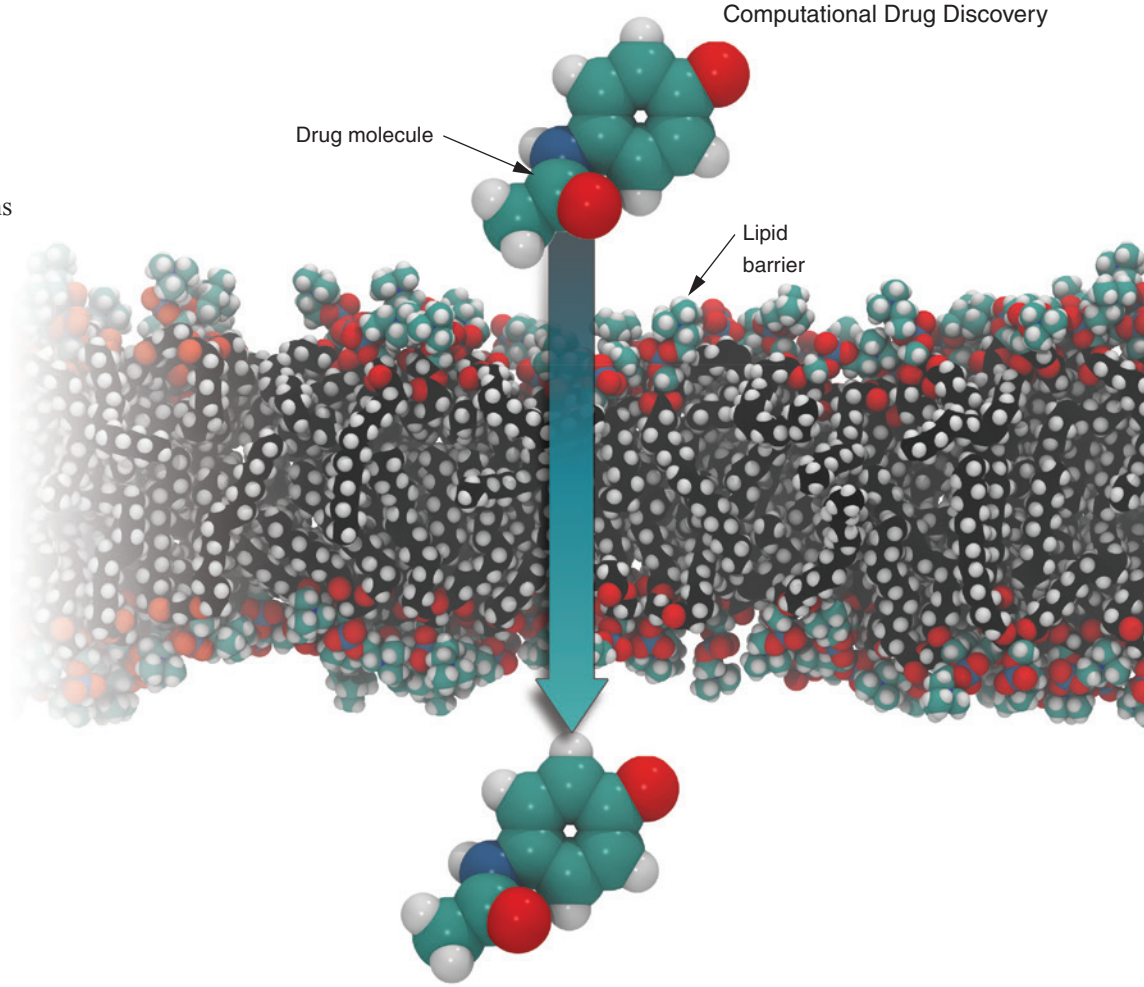


This simulation shows a drug compound crossing the lipid barrier between the bloodstream and the brain. Predicting such behavior could help drug researchers better gauge drug efficacy and catch potentially serious side effects at an early stage of drug development.

## Selective But Predictable Barrier

The easiest way for drug therapy to travel from the bloodstream to tissues is through gaps between the cells in blood vessel walls. Blood vessels in the brain, however, have especially tight cell junctions, creating a highly selective barrier. Predicting whether a given compound will successfully travel from the bloodstream through these cells to reach the brain is a crucial part of drug design. Some drugs, such as those targeting Alzheimer's disease or depression, need to enter the brain to be effective, but drugs targeting other parts of the body must be kept out, lest they cause unexpected and potentially serious side effects.

Testing whether a drug will cross the blood–brain barrier typically occurs five to six years into a drug development project, during in vitro or animal studies. Livermore researchers have been evaluating whether computer modeling can predict this behavior at a much earlier stage. In their modeling protocol, the barrier is represented by two leaflets of closely packed phospholipid molecules, which are surrounded by water. This thin membrane has a hydrophobic center and hydrophilic outer edges.

The researchers investigated the model's predictive capability by simulating a chemically diverse set of 12 compounds, all of which have been well studied experimentally. The technique they used, called umbrella sampling, involves subdividing each drug's interaction with the phospholipid bilayer into 100 separate molecular dynamics simulations. Each 45-nanosecond simulation features a single drug compound at a slightly different position within the system. Together, the simulations follow a compound's progression through the bilayer and out the other side. The model applies a force to the drug to keep it close to the center of the simulation, although it is free to move laterally and rotate to its optimal orientation. Based on a measure of how much force is needed to keep the compound in position, the scientists can

estimate the relative energy values across the system for that specific compound and thus how likely the compound is to diffuse through the barrier.

Most of the drugs required more force to keep them in the middle of the bilayer than in the water. Computational biochemist Tim Carpenter, who is leading the simulation effort, says, "In general, we found that if the compound had a positive free energy in the middle of the bilayer, it would not cross the barrier. However, if the compound had a negative value, it would cross the bilayer and enter the brain." Results correlated well with experimental data and also compared favorably with existing computational techniques, most of which use empirical rather than first-principles methods.

Having proven the feasibility of their approach, the researchers now hope to develop a more realistic model, incorporating a greater variety of lipids. The current model simulates only passive diffusion of molecules. Future models may also incorporate active uptake, because some drugs mimic compounds the brain needs to fool certain proteins into pumping the compounds across the blood–brain barrier.

## Searching for Unexpected Targets

"A drug introduced into the body interacts with many things in addition to those intended," says Wong. "It is the sum of these effects that gives us the clinical results we see." Interactions may include off-target drug bindings, when the drug binds with more protein types than just the target protein, or more complex disruptions to a biological pathway, either of which could produce unanticipated ADRs. Enabling drug researchers to identify and evaluate the severity of these reactions early in the drug development process, when redesign is less costly and time consuming, is the goal of an ongoing collaboration between Lightstone's team and LaBute. Together, they are exploring the feasibility of predicting ADRs with statistical modeling.

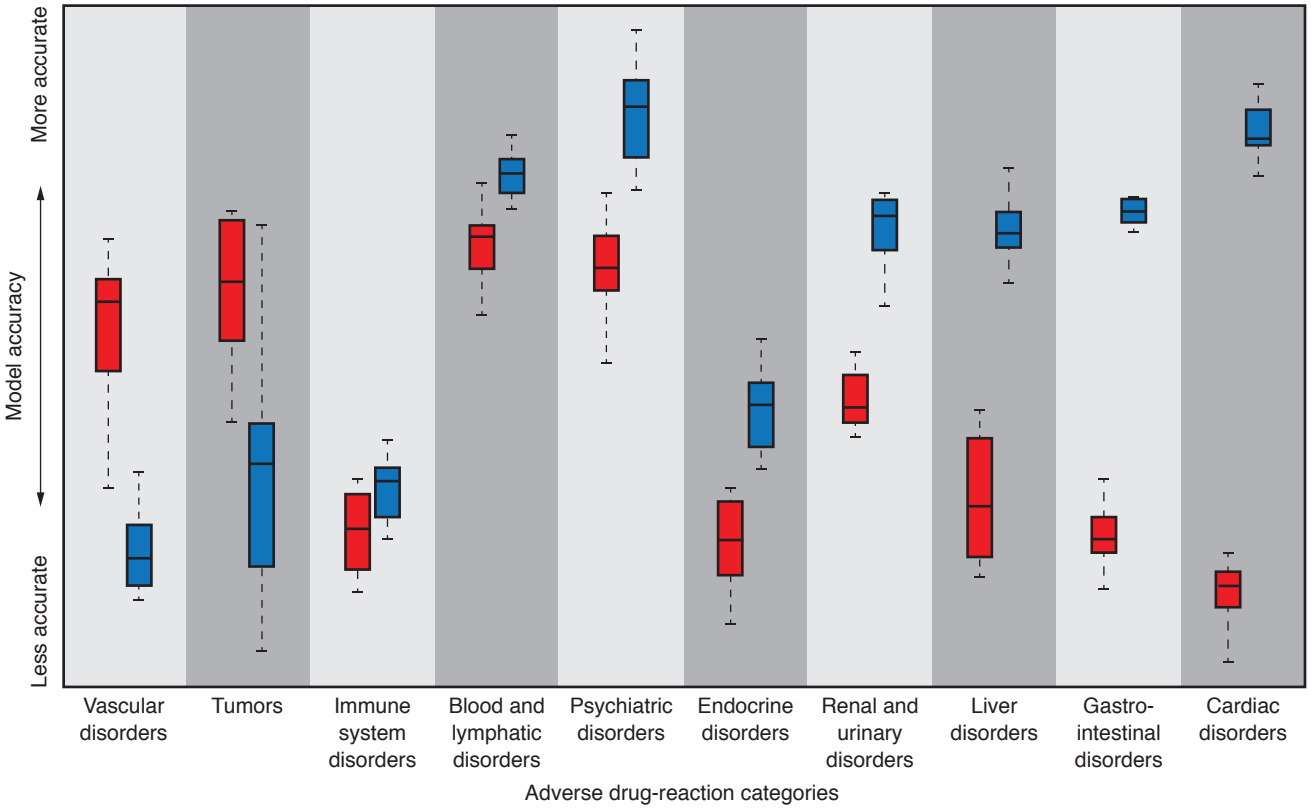


A Livermore-developed virtual adverse-drug-reaction (ADR) screening program integrates in-house molecular docking data with publicly available data sets. Models tested with VinaLC docking calculation data (red) and with experimental binding data (blue) are roughly comparable at predicting ADRs. This graph displays model quality averaged across 10 rounds of cross validation for 10 categories of serious or lethal ADRs, for a total of about 500 drugs. The VinaLC docking calculations evaluated about 500 protein receptors per drug.

The team's primary approach has been to use VinaLC to calculate binding possibilities between known drugs and all available human protein structures. Then using data on known ADRs for these drugs, the researchers are attempting to statistically associate reactions with binding events identified through molecular docking. Molecular docking should allow the team to spot more off-target effects than in vitro testing. "Pharmaceutical companies screen between a dozen and 50 receptors for ADRs in vitro, but that's a minuscule fraction of human proteins," notes LaBute.

For comparison purposes, the team is also analyzing correlations found between existing publicly available data sets, including those for drug structures, known ADRs, biological pathways, and proteins that cause side effects. Unfortunately, the available data have limitations. Public information on known ADRs consists mostly of on-target reactions and thus is not always a complete listing for a given drug. Also, Livermore researchers are focusing on serious or lethal ADRs. Those reactions have a clearer cause and effect and are more likely to cause a drug to fail. However, serious ADRs are underrepresented in the data because documented drug failures are often proprietary information—not available to the public.

LaBute has prepared and tested his models through cross validation, an iterative technique often used when data are sparse or incomplete. Modeling results so far are encouraging. The on-target data set was better at predicting certain types of ADRs, particularly more complex reactions such as endocrine and gastrointestinal disorders. The off-target molecular docking data were better at predicting other ADR categories, such as vascular disorders and cancers. In time, the off-target predictive approach may prove a good complement to existing ADR screening methods. The team is also investigating how to enhance predictions by using machine learning to identify new patterns in existing data sets or by mining the scientific literature for relevant information to expand the data pool.

## Mapping Free Energy

When breaking down a drug, the body may produce one or more by-products, called metabolites. Some of these metabolites may be toxic, even when the drug itself is not. Postdoctoral researcher Yue Yang is leading an effort to use HPC for predicting at the early stage of drug discovery which metabolites will be produced and in what relative quantities. Yang's team has extensively modeled acetaminophen, a popular and well-studied pain reliever.

The cytochromes P450 (CYPs), a family of proteins, play a crucial role in breaking down drugs in the body. During acetaminophen metabolism, several CYPs produce mainly N-acetyl-p-benzoquinone imine (NAPQI), a toxic by-product associated with liver damage, while other CYPs produce mainly nontoxic metabolites. Livermore researchers have hypothesized that the reason different P450 members produce different metabolites must have to do with how and where the CYPs bind with acetaminophen during drug metabolism.

To explore this idea, Yang's team performed molecular docking calculations to find the top-ranked binding poses for five CYP protein types. With Sierra, the researchers then ran large-scale molecular dynamics simulations and two-dimensional umbrella sampling to determine the leading binding arrangements. Calculating binding

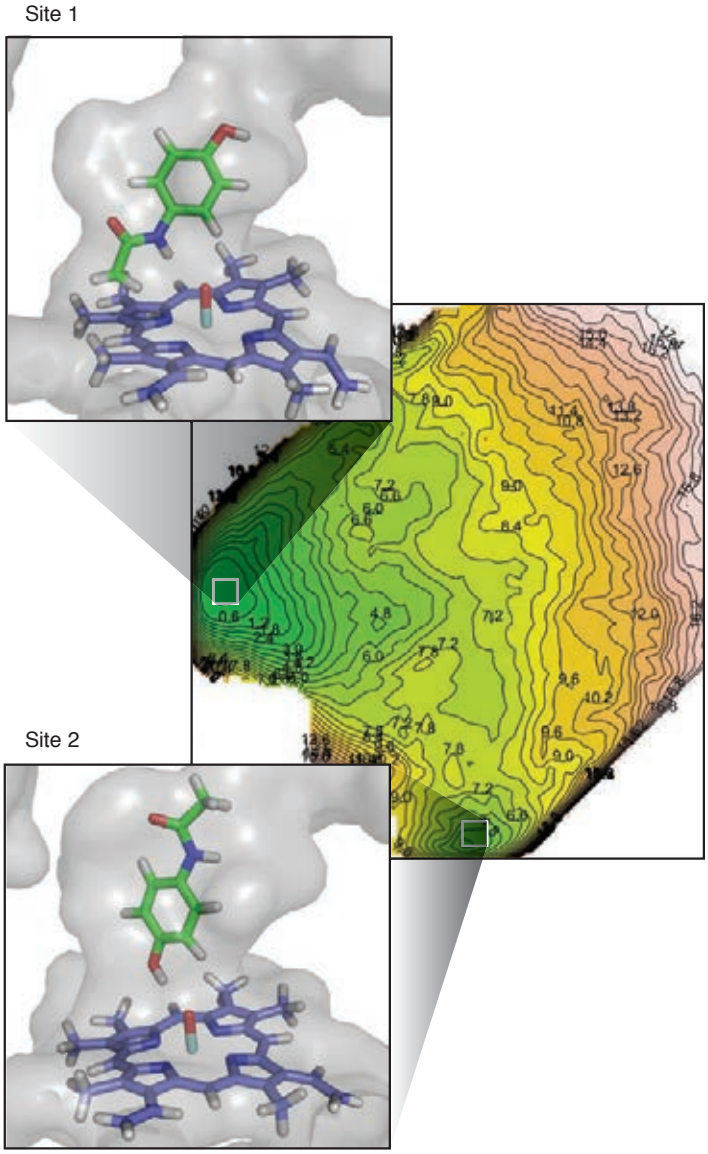


Livermore researchers calculated the most favorable binding arrangements and locations for several key proteins during acetaminophen metabolism to help determine which reactions produce which by-products. Free-energy mapping (right) and binding calculations (insets) show that the protein CYP2E1 prefers site 1, a configuration and binding location leading to N-acetyl-p-benzoquinone imine (NAPQI), a by-product of acetaminophen that can cause liver damage. Site 2, another possibility, would produce a nontoxic by-product; however, it has a higher free-energy value, so a reaction is much less likely.

probability across large binding sites such as those on the CYPs is computationally intensive. Each CYP required about 280 10-nanosecond simulations. The whole project took 800,000 CPU hours to complete.

With the simulation data, the team could analyze the binding free-energy landscape for each acetaminophen–CYP reactive complex and create a map of the spatial positions of the interacting molecules in the system along with the corresponding energy levels. A system usually seeks to achieve a minimum of free energy, so the lower the binding free energy, the more probable the binding mode. Using the maps, the researchers identified several of the most stable binding sites on each CYP and the metabolite each binding reaction would produce.

CYP2E1 was of particular interest. Although this protein is not expressed in great quantities in the human body, it is responsible for a large portion of acetaminophen metabolism and is the primary producer of NAPQI. Interestingly, CYP2E1's top-scoring pose from the docking calculations is one that should lead to a nontoxic metabolite, but the molecular dynamics simulations and the free-energy profile indicated that a different site and configuration were more energetically favorable. This alternate mode generates NAPQI consistent with experimental results.

The study successfully reproduced the metabolites reported experimentally for each CYP, suggesting that the technique will be able to identify possible unwanted reactions in new drug entities. "This study shows that the relative binding free energy for the drug and different binding modes play an important role in determining the distribution of toxic and nontoxic metabolites," says Yang. "Previously, researchers looked only at the chemical reaction barrier to determine the most probable metabolites. Now, we can more rationally predict toxic metabolites in silico, before drug testing, which is the most expensive step in the drug discovery and development process."

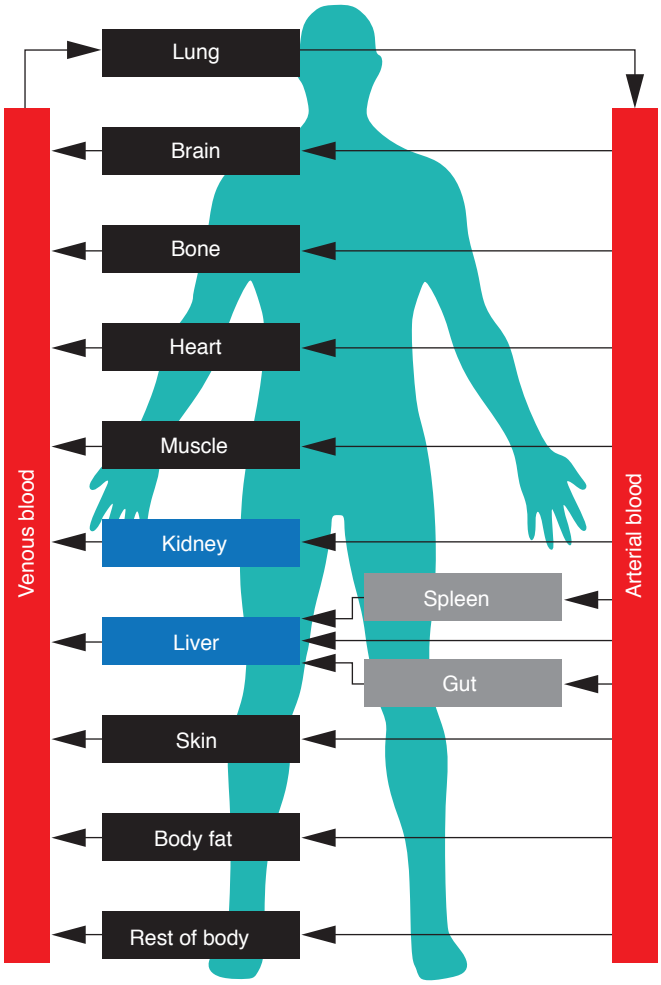


In Livermore's 14-compartment physics-based kinetic model of the human body, various tissues are connected by the circulatory system, and key processes in the human response to drugs are included. This model will help researchers assess toxicity risks early in the drug development process.

## Existing Data, New Conclusions

Acetaminophen also served to validate a study that examined ADRs and drug metabolism on a larger scale. Lightstone and her team are developing a first-principles kinetics model to better understand how drugs behave and change in the body, based on the drug's structure. This physics-based model will aid them in assessing efficacy and toxicity early in a drug's development. The model represents the body as a series of compartments (organs) linked by blood. It can be used to track the flow of a drug into and out of the compartments and determine how the drug interacts with specific proteins in each organ.

More broadly, the model enables scientists to see how processes such as absorption, desorption, distribution, metabolism, and elimination change the concentration and chemistry of a drug in the body over time. The researchers used the model to simulate how acetaminophen is metabolized in the body and whether certain behaviors—fasting, binge drinking, and chronic overuse—increase the likelihood of liver damage resulting from acetaminophen consumption.

How much damage a drug causes depends on its concentration and the length of time any toxic by-products of the drug stay in the system. A normal dose of acetaminophen, for instance, results in a small amount of NAPQI that can quickly be detoxified by an antioxidant called glutathione in the liver. In fact, the simulation showed that even as much as five times the recommended acetaminophen dose should be safe for healthy and well-fed individuals. However, both fasting and high levels of alcohol consumption can prevent glutathione from regenerating at a normal rate. For people whose glutathione supply is depleted, the simulation indicated that even small overdoses could cause NAPQI to accumulate in the liver, bind to various liver proteins, and cause damage.

Livermore's kinetic model was validated by comparing predictions with clinical outcomes. The model accurately replicated some existing clinical results and provided model validation, but it also shed light on metabolism of this common drug. Ali Navid, who led the study, notes that computational tools enable drug research that might be difficult or even unethical to conduct on patients. For instance, a researcher would never ask a patient to binge drink or overdose on acetaminophen simply for study purposes, but understanding the physiological and pharmacological changes associated with these fairly common activities is crucial for toxicity prevention and treatment. Furthermore, without modeling insights, doctors may struggle to weigh the effects of patients' daily habits and to pinpoint why, for instance, a mild acetaminophen overdose causes liver damage in a particular patient.

Notes Navid, "That's the beauty of developing computational tools for pharmaceutical research—we can use information gathered by physicians and clinicians over the past 25 years and draw new conclusions from it. For example, we found that taking acetaminophen after an extended period of fasting results in more damaging side effects than taking it while consuming a bottle of vodka. These results aren't necessarily intuitive." Such findings may help clinicians settle a long-standing debate on the role of alcohol consumption in acetaminophen-induced liver damage. Next, Navid will use the same framework to simulate how drug–drug interactions such as ciprofloxacin and caffeine affect the body.

## Multiscale and Multidisciplinary

The kinetic model, the cornerstone for Lightstone's project, will allow her team to integrate many different levels and types of simulations. Brain-barrier permeability estimates, metabolism site predictions, and mappings of off-target effects, for instance, will eventually feed into the kinetic model to gain a more dynamic understanding of how a drug compound will interact with the human body.

Enhancing the model is an enduring goal for the team. "We are a long way from simply uploading a person's genome into a computer program that tailors a therapy, such as was depicted in *Star Trek*," says Navid. "However, we may be as close as 25 years from using the structure of a drug to predict a patient-specific outcome. We have made large strides in determining the blueprint of human beings and other organisms. We still need more knowledge of specific dynamic characteristics such as how environmental conditions or personal history affect biological interactions. These interactions form the foundations of our models and are crucial for predicting the outcome of therapeutic treatments."

As this project nears conclusion, other ideas that would marry biomedical research with HPC are gaining momentum at the Laboratory. One is an HPC for biology incubator, modeled on the successful hpc4energy incubator. (See *S&TR*, June 2013, pp. 4–12.) A biology incubator would partner computational biologists and biomedical companies with HPC resources to solve a problem or advance a company's research. Also, in a follow-on project to the successful Cardioid heart modeling (see *S&TR*, September 2012, pp. 22–25), Livermore is helping launch a national initiative to model the human body at multiple scales, which would greatly advance digital drug and biological countermeasure development.

"Throughout our project, we've drawn on such areas as machine learning, chemistry, biology, and physics," says LaBute. "This multiscale, multidisciplinary project was well suited for a national laboratory."

*—Rose Hansen*

**Key Words:** adverse drug reaction (ADR), blood–brain barrier, computational chemistry, cytochrome P450 (CYP), drug development, high-performance computing (HPC), kinetic model, molecular docking, molecular dynamics simulation, N-acetyl-p-benzoquinone imine (NAPQI), Sequoia, Sierra, supercomputer, umbrella sampling, virtual screening.

***For further information contact Felice Lightstone (925) 423-8657 (lightstone1@llnl.gov).***

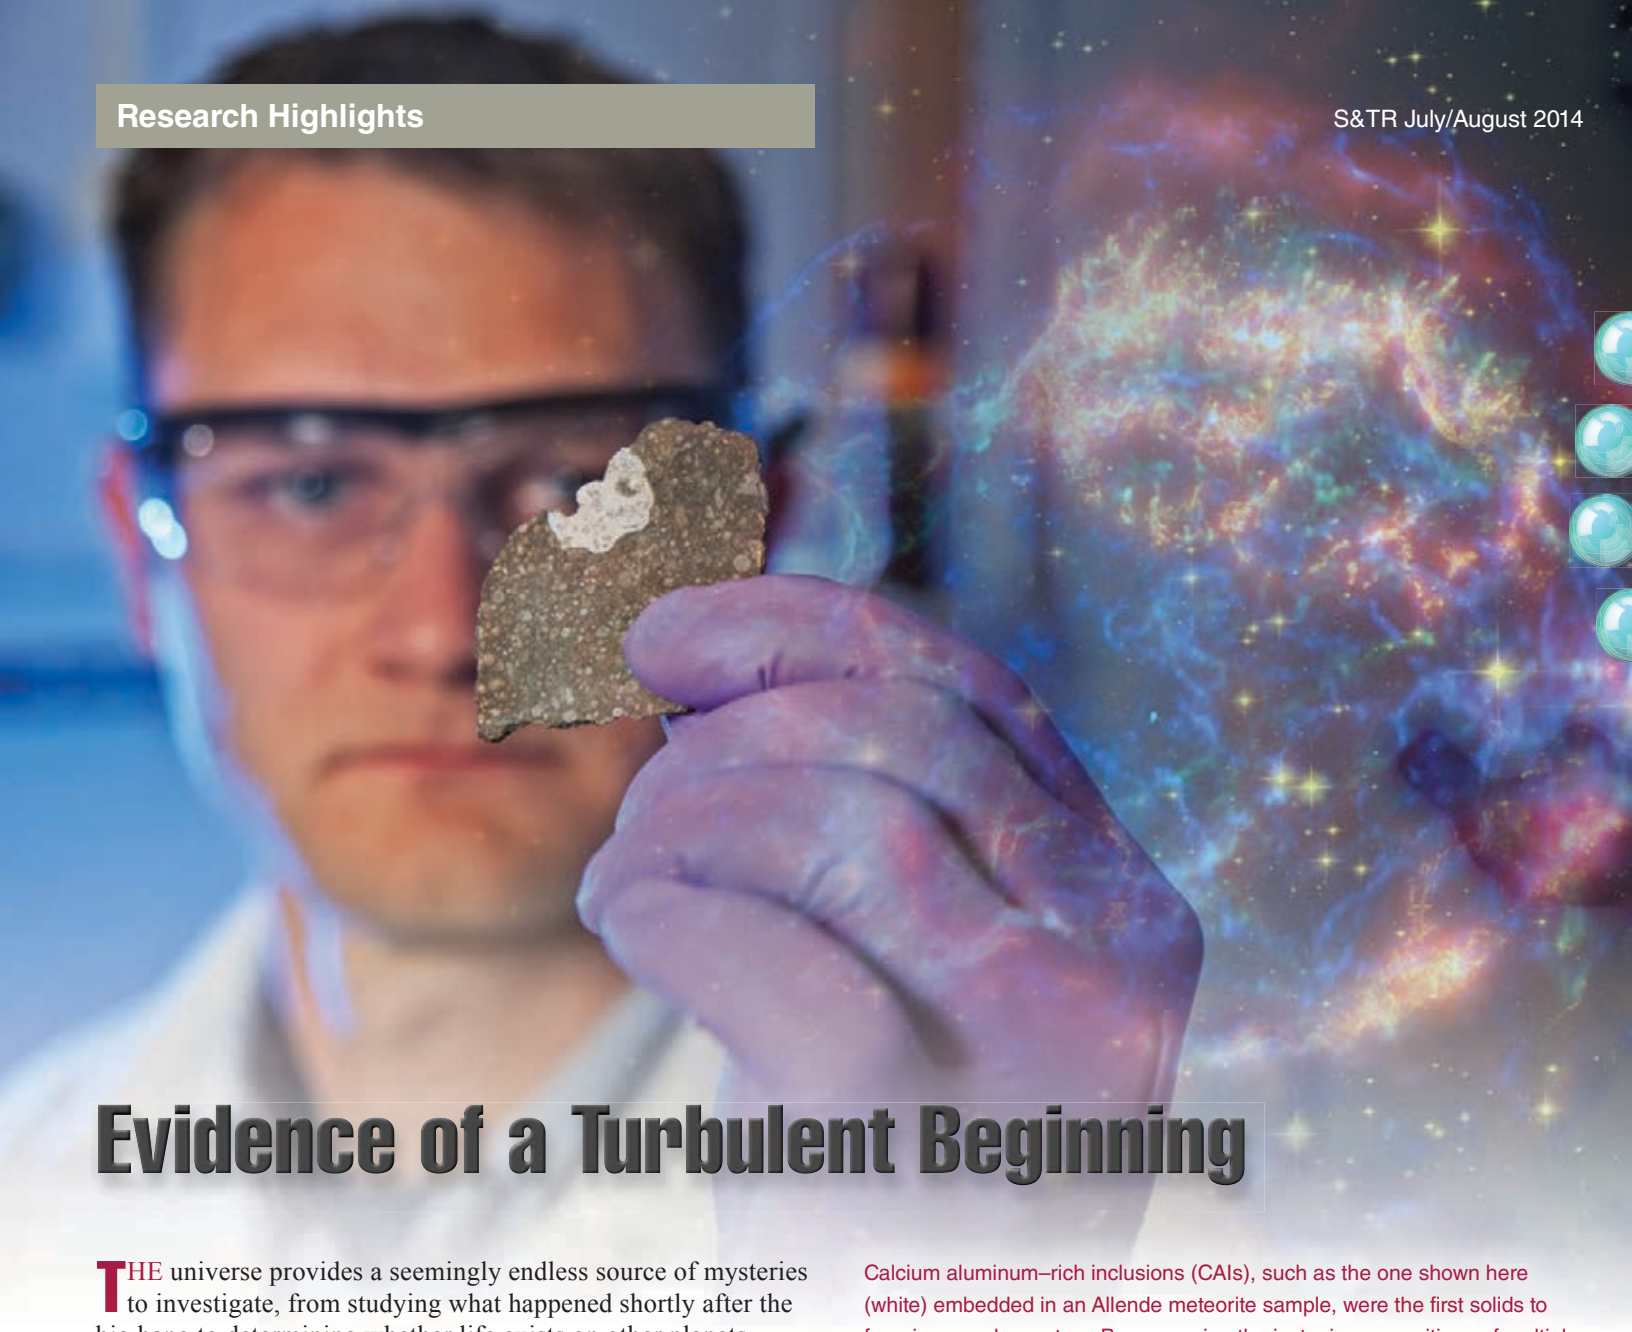

# Evidence of a Turbulent Beginning

Calcium aluminum–rich inclusions (CAIs), such as the one shown here (white) embedded in an Allende meteorite sample, were the first solids to form in our solar system. By measuring the isotopic compositions of multiple elements in individual CAI samples, Livermore researchers Greg Brennecka (shown) and Lars Borg have discovered systematic differences between CAIs and terrestrial materials. The data indicate that a nearby supernova (similar to the one top right) injected material into the solar system sometime between the formation of the first solids and the formation of the planets. (Photo by George A. Kitrinos; Cassiopeia A supernova image courtesy of NASA's Chandra X-Ray Observatory.)

THE universe provides a seemingly endless source of mysteries to investigate, from studying what happened shortly after the big bang to determining whether life exists on other planets. Understanding our solar system's formation and evolution is one of the foremost areas of research among planetary scientists, including astrophysicists and geochemists, who seek to answer how, when, and why the solar system began. Of particular interest is deciphering the conditions of the system's nascent environment, an actively debated topic in scientific circles. New research from a Livermore-led team suggests our solar system's birth was chaotic and included input from a nearby supernova explosion early in its evolution.

Researchers Lars Borg and Greg Brennecka, along with colleague Meenakshi Wadhwa at Arizona State University, used Livermore-developed techniques to evaluate isotopic signatures in calcium aluminum–rich inclusions (CAIs) from the Allende meteorite, which fell to Earth in northern Mexico in 1969. These submillimeter- to centimeter-sized inclusions within the meteorite matrix were the first solids to form in our solar system, predating the terrestrial planets (Mercury, Venus, Earth, and Mars) by more than 1 million years. Thus, the isotopic character of CAIs provides details about the composition of the system at its earliest stages and the astrophysical events that may have contributed later to the elemental makeup of planets. Previous studies of these materials

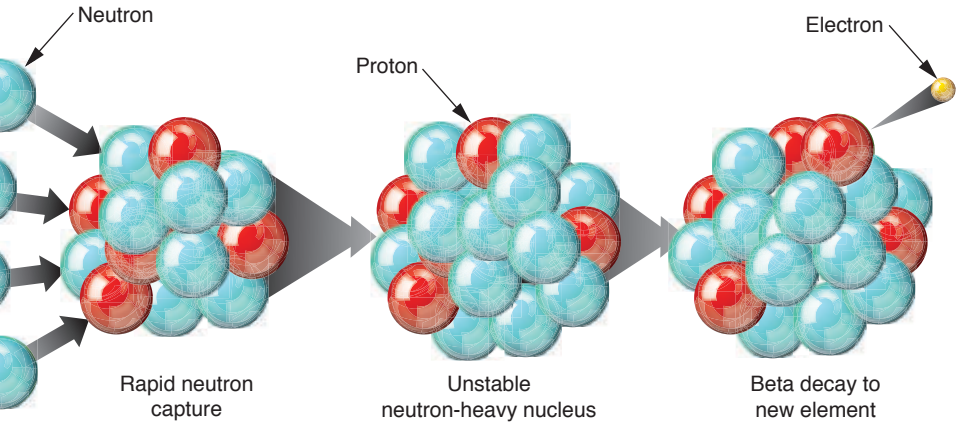


R-process nucleosynthesis produces heavy elements from lighter seed nuclei. In this process, rapid neutron captures (left) produce heavier isotopes of the seed element, until an unstable isotope with a beta-decay lifetime shorter than the neutron capture rate is formed (center). Beta decay of a neutron, which produces a new proton and an emitted electron, transforms the unstable neutron-heavy seed element isotope into a very heavy isotope of the new element (right). Eventually, when the neutron flux declines, these very heavy species will decay back to stability, producing the most neutron-rich stable isotope of the newly created element.

focused on the isotopic compositions of single elements and how they differ from terrestrial samples. Borg and Brennecka have taken a fundamentally different tack, examining the compositions of multiple elements in the same sample set.

Their approach provides clear evidence that the cloud of matter condensing to form the Sun and planets was showered with material produced by a nearby type II supernova explosion. "The data indicate elements with isotopic composition typical of supernova material must have been injected into the solar system sometime between the formation of the first solids and the formation of the planets," says Brennecka. The work, a spin-off of Brennecka's doctoral research studying uranium isotope compositions in CAIs, provides important new details on the formation of the solar system and answers some of cosmochemistry's long-standing questions about the isotopic abundances observed on Earth.

## A Glimpse into the Past

An isotope has the same atomic number as its "parent" element but a different mass number based on how many neutrons are contained in its nucleus. In stellar environments, nucleosynthetic processes, such as proton- and neutron-capture reactions, and radioactive decay contribute to the creation of heavy elements and their isotopes throughout the universe. Elements heavier than nickel are produced by p-process (proton capture), s-process (slow neutron capture), or r-process (rapid neutron capture) reactions. Brennecka says, "The combination of these reactions is ultimately responsible for the isotope abundances present in all solar system materials."

CAIs are thought to have formed in the innermost part of the solar system's protoplanetary disk. They are isotopically anomalous to materials that formed later in the system's development, but the reasons for this characteristic are not well understood. To study the discrepancy, Borg and Brennecka dissolved whole CAIs from the Allende meteorite in acid and isolated individual elements by passing the samples through multiple ion chromatography columns to separate targeted elements from the sample matrix. Borg says, "We studied elements such as strontium, molybdenum, barium, neodymium, and samarium. These elements have multiple stable isotopes, and each one has variable inputs from p-, s-, and r-process nucleosynthesis." Using mass spectrometry instrumentation at Livermore and Arizona State University, the team determined that the isotopic compositions of these elements differed slightly from the isotopic compositions of the terrestrial planets.

"By examining multiple heavy elements in the same sample set, we uncovered a systematic difference in their isotopic abundances," says Brennecka. "These isotopic 'fingerprints' indicate that CAIs originated in a homogeneous reservoir different from that of solar system materials that later formed the planets." The systematic difference suggests that after the formation of CAIs, the early solar system must have been injected with a small amount of material containing heavy elements created by r-process nucleosynthesis. This process occurs when the nuclei of lighter elements capture large numbers of neutrons, then beta decay to stability, increasing the atomic number of the nucleus. The most likely astrophysical process producing the conditions needed for initiating and sustaining r-process nucleosynthesis is a core-collapse (or type II) supernova.

## Timing Is Everything

Type II supernovae are the end product of short-lived, massive stars that quickly exhaust their available fuel, gravitationally collapse, and violently expel material to interstellar space. These events are categorized by the isotopic masses they produce, with certain mass ranges occurring at a higher frequency than others. As a result, materials created in close proximity to these events have distinct isotopic signatures. "Surprisingly, what we found in our

analyses were excesses in r-process isotopes below mass 140 and deficits of those isotopes above mass 140," says Borg. "These data provide physical evidence for the multiple, differing sources of r-process nuclides."

"The basic idea is that something relatively depleted in r-process isotopes below mass 140 and relatively enriched in r-process isotopes above mass 140 would be required to make up the differences between the CAI reservoir and the terrestrial planet reservoir," says Brennecka. "These systematic differences between the first solids and everything that formed later give us a window for when the supernova would have injected material into the protoplanetary disk. They also tell us that the neighborhood our solar system 'grew up' in was rough. Possibly, many violent events were happening as our solar system was forming."

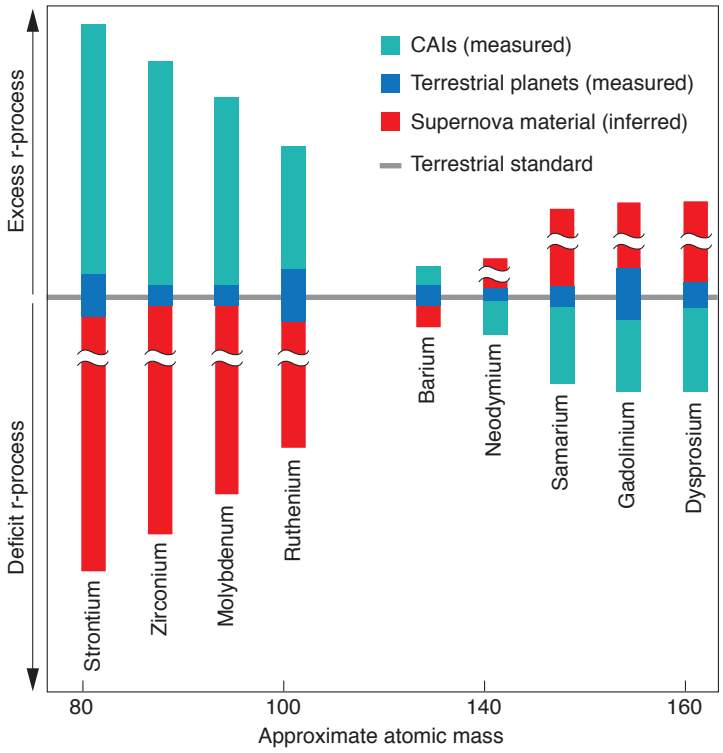


This graph shows the measured isotopic character of CAIs from the Allende meteorite (turquoise), the inferred isotopic character of supernova material (red), and the measured isotopic character of the terrestrial planets in our solar system (blue). For the "average" composition of the terrestrial planets to develop, supernova material must have been mixed in with those elements over time and space. A type II supernova is the most likely astrophysical process to produce the right conditions for r-process nucleosynthesis. Although the amount of supernova material mass is still unknown (thus the uncertainty in the magnitude of the anomaly), the isotopic character can be inferred.

The work performed by Borg, Brennecka, and Wadhwa is important for several reasons, not the least of which is that it is really interesting science. "Through our analyses, we know what the isotopic fingerprint looks like for this supernova," says Brennecka. Borg adds that the system the researchers developed for studying multiple isotopes in a single sample is particularly noteworthy. "It's not a trivial task to isolate each of the elements from an individual sample," he says. Furthermore, no other laboratory has studied more than a few elements in the same CAI sample set. To date, the Livermore team has examined more than 10 elements in the same sample set, spanning almost the entire mass range of the periodic table. This work has additional applications for other Department of Energy research areas, from understanding thermal neutron capture to isotope fingerprinting of nuclear materials, both of which are important to the Laboratory's national security missions.

Many questions remain regarding our solar system's inception and evolution. However, Borg and Brennecka offer a new method by which to probe for answers to these unknown questions. The next big challenge for the team is to study other CAIs to assess whether these materials also demonstrate evidence for multiple sources of r-process nuclides or whether the anomalies are unique to the Allende meteorite. "We have raised the bar for the geochemistry and cosmochemistry fields going forward," says Borg. "We are proving that single-element analysis gives us only part of the picture. By analyzing multiple elements in the same sample set, we are gaining a better understanding of what was happening during our solar system's formation."

*—Caryn Meissner*

**Key Words:** Allende meteorite, calcium aluminum–rich inclusion (CAI), cosmochemistry, geochemistry, ion chromatography, isotope, isotopic signature, nucleosynthesis, planetary science, rapid neutron capture (r-process), solar system, type II supernova.

***For further information contact Lars Borg (925) 424-5722 (borg5@llnl.gov) or Greg Brennecka (925) 423-2304 (brennecka2@llnl.gov).***

# Supercomputing Tools Speed Simulations

In late 2013, an international team of scientists was in the midst of simulating a collapsing cloud of 15,000 bubbles using Lawrence Livermore's Sequoia supercomputer, when the calculations suddenly stopped. In quick response, the team turned to a Livermore software tool called STAT (Stack Trace Analysis Tool) to locate which of more than 6 million computing threads (calculations) was causing the problem. Within a few minutes, STAT traced the hangup to a particular microprocessor core (computing engine). The team went on to complete the pioneering simulation and win the 2013 Golden Bell Prize for outstanding achievement in high-performance computing (HPC). (See *S&TR*, January/February 2014, p. 20.)

STAT is the product of a small team of computer scientists comprising the Development Environment Group (DEG) in Livermore's Computation Directorate. In addition to STAT, the group has developed tools designed to boost performance and productivity such as AutomaDeD and SPINDLE. Another set of tools is being developed under the PRUNER project to help with the reproducibility of large simulations. The group works closely with the Laboratory's Center for Applied Scientific Computing, which supports the demanding computing requirements of Livermore scientists.

The group's tools are designed to work on massively parallel machines such as Sequoia, one of the world's most powerful supercomputers. Sequoia has 1,572,864 processing units, or cores,

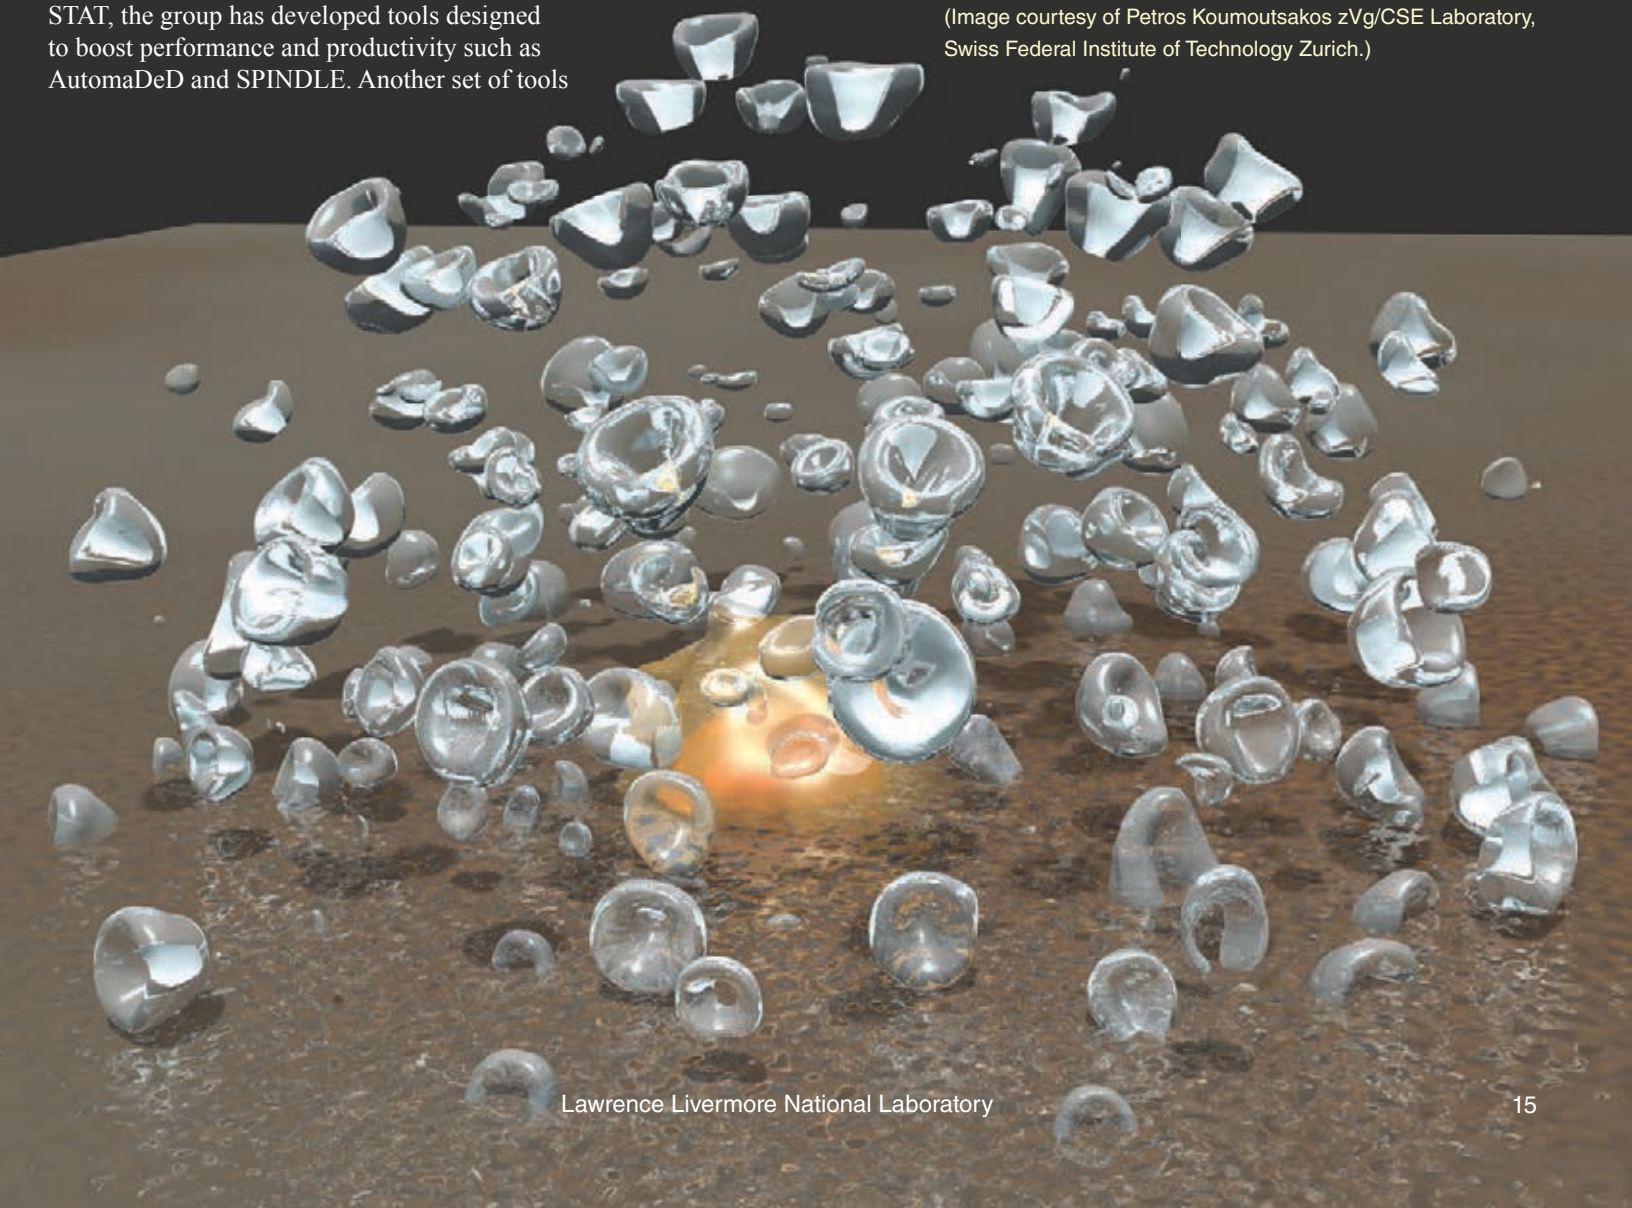

Lawrence Livermore scientists and collaborators set a new supercomputing record in fluid dynamics by resolving unique phenomena associated with a cloud of collapsing bubbles. The work earned the team the 2013 Gordon Bell Prize. See the simulation at http://www.youtube.com/watch?v=zfG9soIC6_Y. (Image courtesy of Petros Koumoutsakos zVg/CSE Laboratory, Swiss Federal Institute of Technology Zurich.)

and a peak performance speed of 20 petaflops (quadrillion floating-point operations per second). This processing power is compressed into 96 racks, each the size of a large refrigerator. (See *S&TR,* July/August 2013, pp. 4–13.)

Massively parallel supercomputers break a problem into tiny parts that are solved simultaneously. For many applications, such as simulating complex physical phenomena, this computer architecture has replaced vastly slower serial processing, in which tasks are performed sequentially by a single processing element. The Laboratory has been a leader in using parallel supercomputers since their inception. As a result, advanced simulations at Livermore have become as important to scientific exploration as theory and experiment.

According to computer scientist and DEG member Dong Ahn, developing tools for supercomputers requires expertise that resides in just a few research centers worldwide. This expertise includes skills in programming and debugging massively parallel supercomputers as well as anticipating the tools needed for next-generation machines, on which applications are expected to routinely use millions of processors. As parallel supercomputers become more powerful, Livermore computer scientists develop new methods to maximize the potential of such machines. Says Ahn, "Livermore is an applied laboratory, and our research must have practical value to users."

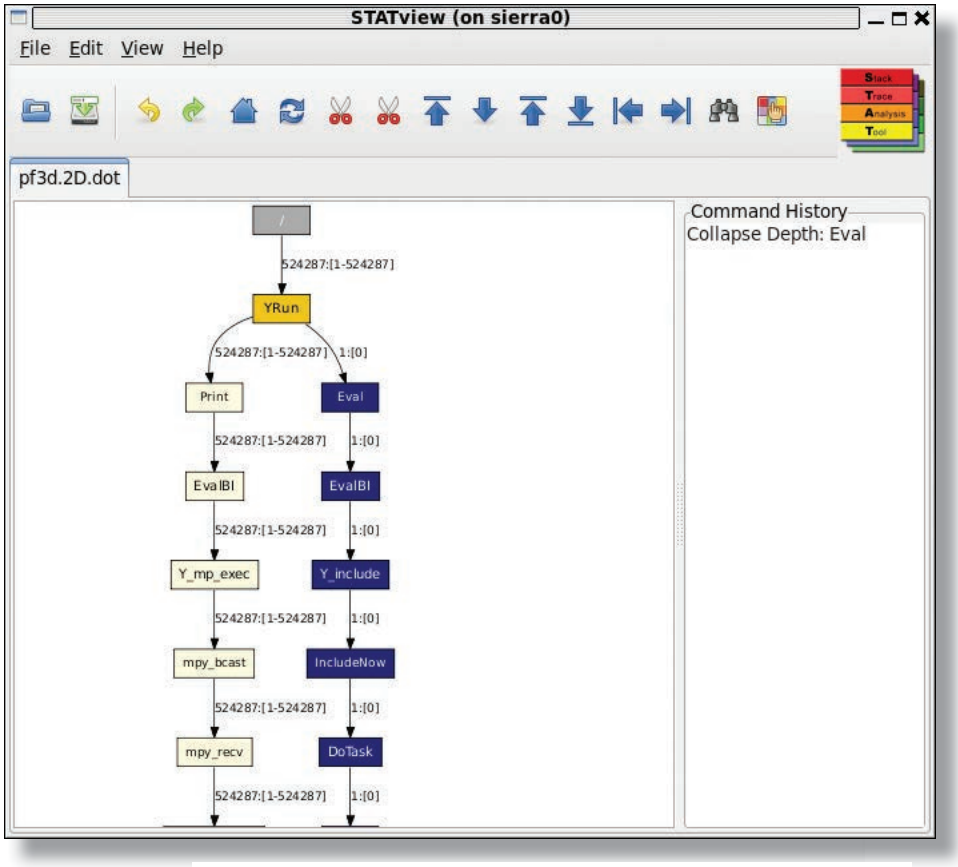


This computer screen shot shows how the STAT tool analyzed 542,288 calculations. STAT formed a graphical "tree" representing a snapshot of an application that was hung up. The tree revealed the problem was in a single calculation (on the right, in blue) and resulted from a programming error.

## Debugging Six Million Calculations

Many of DEG's tools focus on finding bugs. In the world of parallel computing, debugging has become a difficult and complex task. A massively parallel application is a big search space in which errors can reside. Sequoia has nearly 1.6 million cores, each running four threads of execution (calculations). Often bugs emerge only at large scales, overwhelming users with the complexity of the task to isolate the problem. "If something breaks, we need to know quickly what went wrong in one or more of 6 million threads," says HPC systems engineer Adam Bertsch. "We also need to know why it went wrong."

"Many traditional debugging techniques are precluded by the sheer amount of resources that must be examined," says computer scientist Scott Futral, DEG leader and along with Bertsch, member of the Gordon Bell Prize–winning simulation team. Because Livermore applications are continually refined, developers can spend 25 percent or more of their time debugging and optimizing codes, a practice that has become increasingly costly.

In response to the need for an advanced debugging tool aimed at machines such as Sequoia, DEG members collaborated with researchers from the University of Wisconsin and the University of New Mexico to design STAT. This highly scalable tool is capable of identifying errors in computer codes running on machines with more than 1 million processor cores. In 2011, STAT won an R&D 100 Award as one of the year's top innovations. (See *S&TR*, October/November 2011, pp. 14–15.)

STAT is used throughout the Department of Energy's supercomputer community. It is most effective for diagnosing calculations that are "hung up," although the tool has also proved useful for isolating other problems. STAT indicates

where in the code all of the processes are at a given point in time, giving the user insight into where the bug may lie. With a strong graphical user interface, STAT produces two-dimensional (2D) and 3D graphs in the form of treelike structures. The 2D tree represents a single snapshot of the entire application, while the 3D tree presents a series of snapshots from the application captured over time.

STAT played a significant role in validating Sequoia as racks of nodes were added over several months. "As we added racks, we had the opportunity to prove STAT's scalability," says Ahn. The tool helped both early users and system integrators of Sequoia to quickly isolate errors, including issues that manifested only at extremely large scales. In one case, STAT rapidly diagnosed a deadlock in a simulation using over 500,000 cores, allowing the user, who had tried unsuccessfully for weeks to solve the problem using a traditional method, to complete his project on schedule.

For the simulation project that earned the Gordon Bell Prize, STAT helped researchers achieve an ultrahigh-resolution simulation of cloud cavitation collapse. That work set a simulation record in fluid dynamics with 14.4 petaflops of sustained performance. When the calculation suddenly stopped, recalls Bertsch, STAT quickly scanned all 6 million calculations and isolated a problem in one of the processor cores. The team replaced the processor that contained the identified core, and the application was able to proceed. The resulting simulation represented a 150-fold improvement over previous simulations for this type of application and a 20-fold reduction in time to complete the task.

Although it has proven itself many times, STAT is considered a "lightweight" tool that may not always locate a bug if the problem is something other than a hung calculation. For this reason, the group has extended STAT's debugging features with the DysectAPI tool. Still in early testing, DysectAPI is designed to enable users to "program their intuition" so as to construct various higher level debug queries. The tool represents a new approach to debugging a computer program that runs on more than 100,000 processors. The method screens out unnecessary information to allow the user to rapidly zero in on the cause of a crash, fault, or other bug. Ahn says that one could use STAT to first perform a "triage" to locate the general area of the problem and then apply the DysectAPI tool to pinpoint the problem.

## Increasing Computational Efficiency

DEG experts have also developed AutomaDeD, a tool that uses artificial intelligence to automate the debugging process for massive simulations. AutomaDeD has two major functions: identifying abnormal computational tasks and regions of code, and finding the least-progressed task. The first function is accomplished by detecting outliers and the second by ordering processes according to their relative progress. This work involves developing and rapidly detecting problems when system performance deviates statistically from the model. AutomaDeD creates probabilistic behavioral models of how simulations should

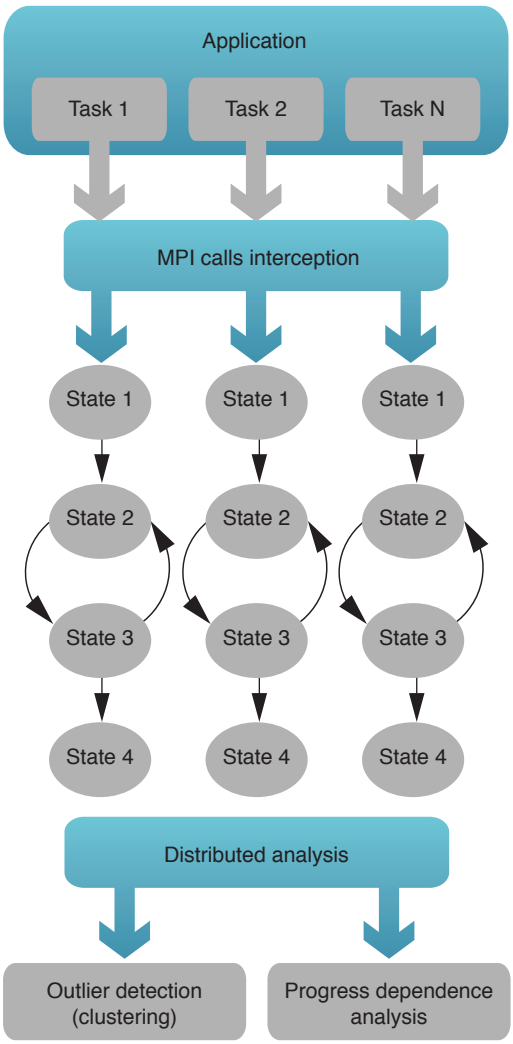


The AutomaDeD tool uses artificial intelligence to go beyond STAT and automate debugging when hundreds of thousands of tasks are being performed simultaneously. Message passing interface (MPI) is used by Livermore's Sequoia and other massively parallel supercomputers to pass information among the machine's hundreds of thousands of computing engines or cores. AutomaDeD identifies abnormal computational tasks and malfunctioning regions of code by detecting outliers (bottom left) as well as the least-progressed tasks (bottom right).

work. When a failure occurs, these models are analyzed to find the origin of the failure.

SPINDLE, another tool from DEG, addresses problems that can occur when millions of cores simultaneously open an application consisting of thousands of shared libraries. The tool builds a cache server to quickly send data from the libraries to the compute nodes. Ahn explains that many applications retrieve libraries of code and data that are shared by every processor, which can greatly slow down processing. SPINDLE's novel approach to loading coordinates simultaneous file system operations so that the file system does not become a bottleneck. This tool is an example of middleware infrastructure, which sits "on top" of system software. SPINDLE has proven to be highly scalable. In one test, system performance at 64 nodes without SPINDLE was similar to system performance at 1,280 nodes with SPINDLE—a 20-fold improvement.

DEG scientists, in collaboration with the University of Utah, are also studying the reproducibility of large simulations under a project called PRUNER. This work is focused on obtaining a fundamental understanding of simulation failures that occur only occasionally, or seemingly without a pattern, and then developing tools to detect and remedy them. Ahn says it may seem counterintuitive, but when a large supercomputer duplicates the same long string of calculations, it can occasionally give slightly different results, or failures can occur such as a crash. This so-called nondeterminism is often the bane of parallel software development, and it can be costly to fix. Many sources of nondeterminism exist such as the sheer scale of computing, a programmer's assumptions, and the order in which calculations are performed. Under the PRUNER project, tools are being developed to detect, control, and eliminate sources of nondeterminism. "These tools would be helpful in validating programs," says Futral.

## Anticipating CORAL

The group is already anticipating the next generation of massively parallel supercomputers, scheduled to appear in 2017. A collaboration of Oak Ridge, Argonne, and Lawrence Livermore (CORAL) national laboratories will deliver these machines. Livermore's system will join Sequoia in serving the National Nuclear Security Administration's Advanced Simulation and Computing Program in support of nuclear stockpile stewardship. The next-generation system will perform up to 200 peak petaflops, about 10 times faster than Sequoia's 20 petaflops.

CORAL represents an important step toward the long-awaited exascale (extreme scale) systems. Ahn says that although supercomputer simulations are used in virtually every research area at Lawrence Livermore, many scientific challenges require computing at the exascale ($10^{18}$ flops). These exascale systems, which are likely to debut at Livermore and other Department of Energy national laboratories early in the next decade, will deploy millions of processing elements or cores. Because of their size, simulations run on exascale machines will present challenges in diagnosing both software and hardware faults, problems to which traditional methods and tools are unsuited.

Ahn emphasizes the role played by academic partners, including the University of Wisconsin and the University of New Mexico for STAT; the Technical University of Denmark for DysectAPI; Purdue University for AutomaDeD; the Jülich Supercomputing Centre in Germany for SPINDLE; and the University of Utah for PRUNER. In the same collaborative spirit, all supercomputing tools developed at Livermore are open source, meaning anyone can use them and are invited to improve them.

With an eye on the fast-changing supercomputer future, Livermore computer scientists are preparing for new generations of giant machines. In particular, the onset of extreme computing may require equally extreme software tools, but Ahn and his colleagues are confident those tools will be in hand.

*—Arnie Heller*

**Key Words:** AutomaDeD, debugging, DysectAPI, exascale, Gordon Bell Prize, high-performance computing (HPC), PRUNER, Sequoia, SPINDLE, STAT (Stack Trace Analysis Tool).

***For further information contact Dong Ahn (925) 422-1939 (ahn1@llnl.gov).***

# A Closer Look at the Rocky Underground

OVER the last several years, the U.S. has begun tapping its enormous domestic reserves of shale gas and oil using hydraulic fracturing techniques. In fact, just a few months ago, the U.S. surpassed Russia as the largest supplier of natural gas in the world. Hydraulic fracturing, also known as hydrofracturing or fracking, promises to provide the nation with an inexpensive and abundant source of energy that is significantly cleaner to burn than coal, the most widely used energy source today.

GEOS, an advanced supercomputer code developed under a Laboratory Directed Research and Development project, offers hydrofracturing planners a new tool to predict the behavior of Earth's subsurface. The GEOS code is just one piece of a larger Laboratory effort to increase hydraulic fracturing efficiency and decrease its costs and environmental impacts. A team led by geologist Rick Ryerson has spent three years developing GEOS to guide natural gas extraction from shale formations and thereby maximize the yield of underground shale gas reservoirs. This project also plays an important role in a Laboratory mission to improve the nation's energy security. Successful projects that increase affordable domestic energy sources help make the U.S. less dependent on imported oil.

## Problem, Solution

In conventional hydrocarbon deposits, oil and gas fill a spongelike trap in rock, making extraction relatively easy. In contrast, shale oil and gas reside in extremely tight rock that must be broken up so the fluids trapped within can flow. A hydrofracturing well travels vertically to the depth of a shale gas deposit and then proceeds horizontally for many kilometers. A combination of water, chemicals, and sand is pumped through the well bore and into the rock at engineered intervals. Under very high pressure, the mixture forces open small fractures in the rock, with the sand acting as a collection of tiny wedges—known as proppants—that keep the fractures open.

Oil and gas production in shale and other low-permeability rocks requires the hydraulic stimulation of vertically oriented, radial fractures that emanate from a horizontal well bore. Natural gas flows from the fractures in the shale into the pipe and to the surface. (inset) Proppants, typically sand or ceramic spheres, hold the fractures open. (Rendering by Kwei-Yu Chu; not to scale.)

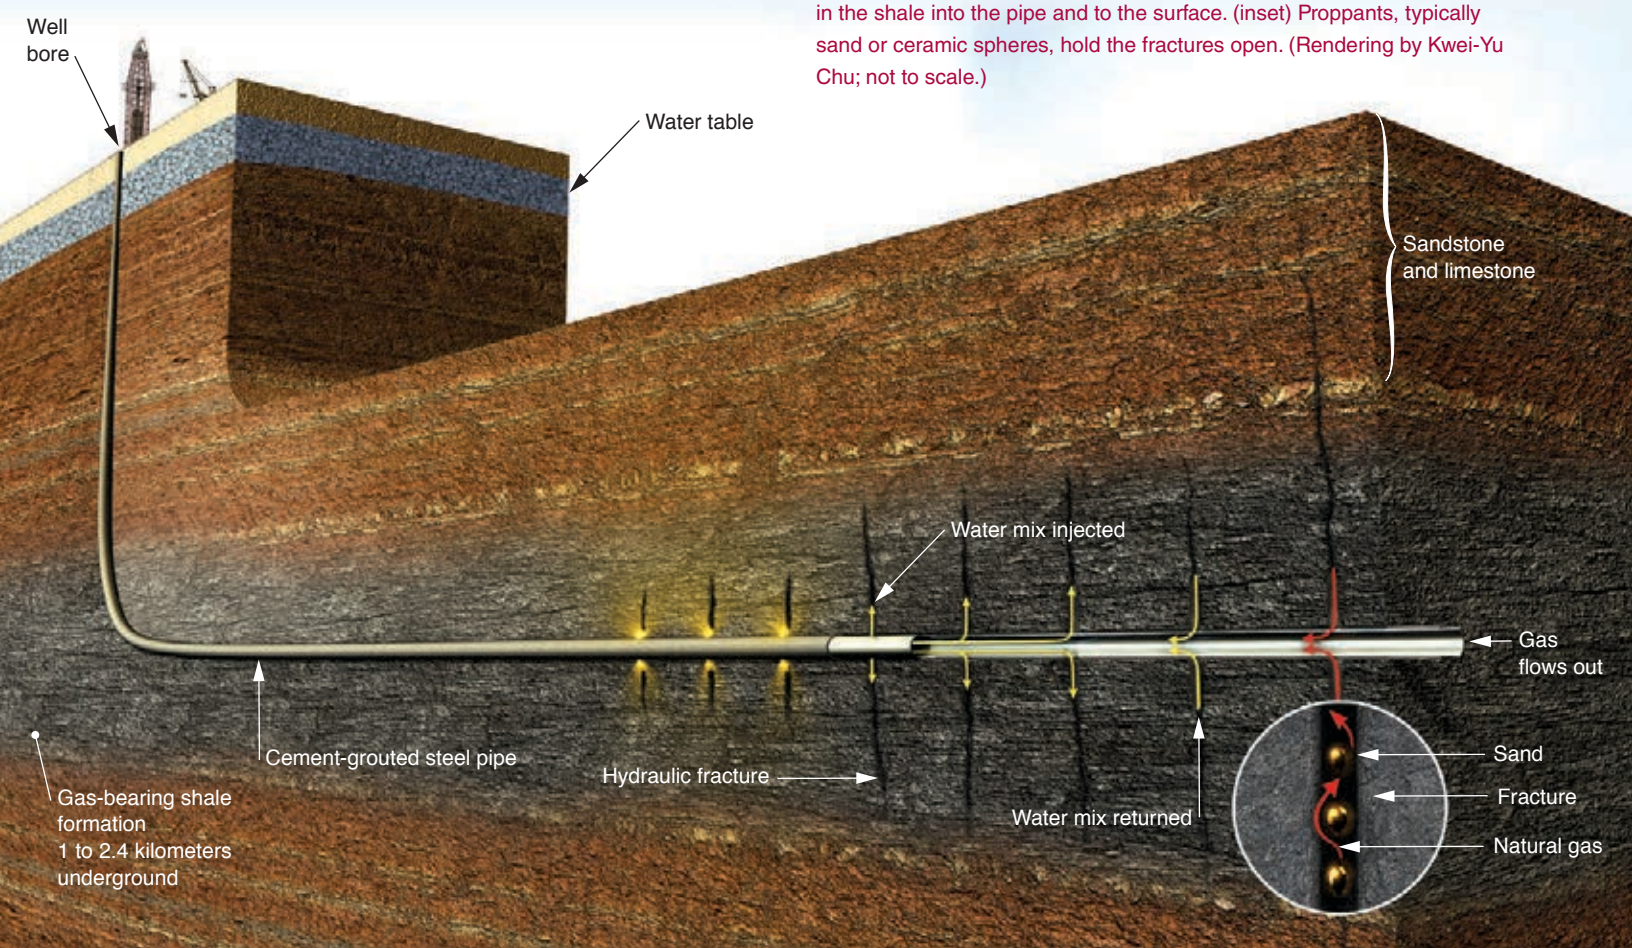

This "stimulation" process succeeds when the pressure of the injected materials exceeds the minimum horizontal stress of the rock. Hydrofracturing greatly increases both the rate and total amount of oil or gas that can be extracted. The method was developed in the 1940s, but it was rarely used until American geologists became aware of the huge volumes of gas-saturated sandstones with permeability too low for economical recovery. Modern fracturing matured in the 1980s and 1990s through several joint industry–Department of Energy projects. Use of hydrofracturing grew rapidly in the late 1990s when energy companies learned how best to stimulate the extraction of shale gas from the Barnett Shale in north Texas. Today, underground deposits previously considered difficult or even impossible to access are being extracted because their permeability has been enhanced through hydraulic fracture stimulation.

One challenge, however, is that well operators generally have limited information about the fractures they are creating. They cannot predict which fractures will develop or how a fracture will perform. Developing geologic simulation tools for oil and gas extraction is extremely difficult, given the complex physical processes at work in the reservoir and the range of temporal and spatial scales involved—from seconds to years and micrometers to kilometers. Uncertainty about a reservoir's properties or the stress conditions within the rock at a distance from the fractures adds to the complexity. For conventional oil and gas extraction, industry typically depends on simulations based on extensive field experience. This method works extremely well when the operation is mature with an accompanying large database. However, an operation in a new locale presents significant challenges, with engineers having to fall back on trial and error.

"Livermore experts are addressing these challenges," says Ryerson. "With our decades of experience in underground nuclear testing, the Laboratory is well acquainted with the best means of analyzing the aftereffects of high-pressure underground activity. We have a great deal of knowledge about how rocks break." GEOS code authors Scott Johnson, Randolph Settgast, Pengcheng Fu, Stuart Walsh, and Joshua White have combined Laboratory expertise in geomechanics, seismology, and high-performance computing to examine fluid-induced rock fracture propagation in the shale gas setting. Their aim was first to understand the dominant mechanisms and sensitivities involved and then to provide practical, real-time guidance for optimizing extraction operations. The team's three-dimensional GEOS code can simulate subsurface fracture networks in a variety of geologic settings.

The oil and gas industry has taken note, and several companies are partnering with the Laboratory to use GEOS, which is a more expansive and higher fidelity code than current commercial codes. To date, Baker Hughes, ION Geophysical, Shell Exploration & Production, and ExxonMobil have partnered with Lawrence Livermore.

## GEOS at Work

Johnson likens a successful geomechanics simulation code to a pair of video cameras that have different lenses and can speed up or slow down the frame rate as the situation changes. One camera is kept on time-lapse mode to survey long-term overall behavior. A high-speed camera zooms in on an area of interest when something starts to happen, and its feed is then synchronized with that of the time-lapse camera. Together, they provide a more complete view of a scene, eliminating the need to observe the entire scene at high spatial and temporal resolution. Simulations can be performed in a similar way by splitting operations over different temporal and spatial scales and by applying different physical models at each scale. This type of adaptive scaling is an essential feature of the GEOS computational framework.

During a GEOS simulation, an implicit solver is used at a coarse scale the majority of time, allowing for fast computation of an entire reservoir's evolution. When a fracturing event is imminent, the implicit solver is transitioned to an "explicit solver" that captures the fast temporal evolution of the event and then synchronizes with the implicit solver to continue its simulation. The two "cameras" work together to provide the best view of activity in the reservoir.

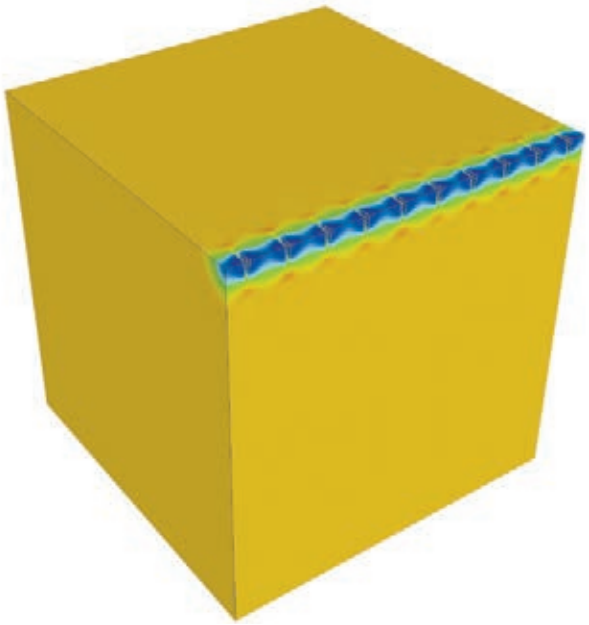

(left to right) Two snapshots of a GEOS simulation illustrate the growth of 10 simultaneously pumped, hydraulically driven fractures in shale.

Coupling between physics models occurs in a similar manner. Most of the time, the hydraulic behavior of rock with preexisting joints and fractures is governed by slow flow through the rock's natural porosity. When the reservoir is being stimulated under pressure, the temporal scales are dramatically reduced, and faster fluid-driven fracture propagation replaces the slower behavior. The microseismic events that happen when a fracture occurs are passed off to a wave propagation code for predicting ground motion. Similarly, an asymmetric permeability tensor, which defines fluid flow direction, can be passed off to fluid and thermal transport codes for longer-term production estimates.

The only way to "see" where fractures are occurring in shale is to locate the microseismic events that accompany each fracture. Each such "earthquake" is typically comparable to the energy of a milk jug falling off a table. Every microseismic event—whether from a new fracture or caused by changes to an existing fracture—must be tracked. But rather than incorporate the real geometry of each microcrack, continuum damage mechanics allows for GEOS to predict the statistical average response of cracked solids. In this way, GEOS can translate fine-scale cracking effects to larger scales.

GEOS also applies uncertainty quantification to predictions of reservoir stimulation. Uncertainties include operational parameters as well as a range of uncertainties about the reservoir's geologic model. Livermore has been a pioneer in the development of solvers for handling uncertainty in a system. The work began with quantifying uncertainties in nuclear weapon systems and has since spread to global climate research and other disciplines. The GEOS team has linked the code with the Laboratory's PSUADE code, a suite of uncertainty quantification modules that addresses high-dimensional sampling, parameter screening, global sensitivity analyses, optimization, and other important factors. In summary, says geologist Roger Aines, "GEOS accounts for all the physics we know, in three dimensions."

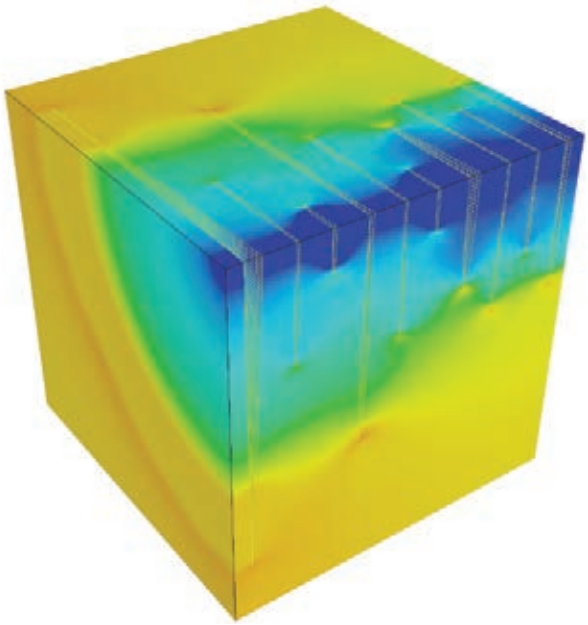

## The Bigger Picture

Since the shale gas boom began, the price of natural gas has dropped, which is in sharp contrast to the soaring price of gasoline. If heavy trucks—which now run largely on imported diesel fuel—could be made to run on compressed natural gas, the savings would be enormous. Yet, despite the benefits that may accrue from replacing coal and gasoline, including more than a 50-percent reduction in carbon dioxide emissions, shale gas extraction faces environmental opposition. Hydrofracturing uses large quantities of water that must be trucked to a well field. Disposing of used water and chemicals by pumping them into the subsurface has caused noticeable seismic activity.

As part of the Laboratory's larger efforts to increase fracturing efficiency and decrease environmental impacts, researchers are also looking at ways to replace sand proppants. Sand is heavy and requires heavy "goo" to transport it. If a light proppant could be developed, clean water alone could be used to transport proppants to the subsurface, without the added chemicals. The clean water could be returned to the surface, eliminating the need to pump used goo underground for disposal.

The Laboratory's additive manufacturing team is developing a lightweight proppant microcapsule with a shell made of the same material as the heat-resistant spatulas used in kitchens. This creative team has also taken an experimental sample of actual rock, imaged it, and printed an exact duplicate in clear plastic so they can watch in real time how new proppants behave. Livermore's multidisciplinary team of scientists and engineers is thus working to increase the efficiency and reduce the environmental footprint of the nation's domestic oil and gas production from hydraulic fracturing techniques.

*—Katie Walter*

**Key Words:** fracking, GEOS code, hydraulic fracturing, hydrofracturing, oil and gas extraction, proppants, shale gas, solver.

***For further information contact Rick Ryerson (925) 422-6170 (ryerson1@llnl.gov).***

*In this section, we list recent patents issued to and 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

### Distributed Road Assessment System

**N. Reginald Beer, David W. Paglieroni**
U.S. Patent 8,681,036 B2
March 25, 2014
This distributed road assessment system detects damage on or below the surface of a paved structure or pavement. The system includes road assessment pods and a server. Each pod has a ground-penetrating radar antenna array and a detection system and is mounted on a vehicle. As the vehicle travels down a road, the system detects damage as return signals and transmits information about each occurrence to the road assessment server, which maintains a database of information on the detected occurrences of road damage. When the server receives information on newly detected occurrences of road damage for a portion of a road, it determines which new occurrences correspond to which previously detected occurrences of road damage.

### Device for Detection and Identification of Carbon- and Nitrogen-Containing Materials

**Alexander Ivanovich Karev, Valery Georgievich Raevsky, Leonid Zavenovich Dzhilavyan, Valery Dmitrievich Laptev, Nikolay Ivanovich Pakhomov, Vasily Ivanovich Shvedunov, Vladimir Ivanovich Rykalin, Louis Joseph Brothers, Larry K. Wilhide**
U.S. Patent 8,681,939 B2
March 25, 2014
This device detects and identifies carbon- and nitrogen-containing materials through the process of photonuclear detection. The device may comprise a racetrack microtron, a breaking target, and a water-filled Cherenkov radiation counter.

### Method and System for Edge Cladding of Laser Gain Media

**Andrew James Bayramian, John Allyn Caird, Kathleen Irene Schaffers**
U.S. Patent 8,682,125 B2
March 25, 2014
A gain medium to amplify light at a gain wavelength with reduced transverse amplified spontaneous emission includes an input surface with an opposing output surface. A central region in the gain medium includes a gain material and extends between the input surface and the output surface along a longitudinal optical axis of the medium. An edge-cladding region surrounds the central region and extends between the input surface and the output surface along the longitudinal optical axis of the medium. The edge-cladding region includes the gain material and a dopant that absorbs light at the gain wavelength.

### Shape Memory System with Integrated Actuation Using Embedded Particles

**Patrick R. Buckley, Duncan J. Maitland**
U.S. Patent 8,685,000 B2
April 1, 2014
This shape-memory system consists of a shape-memory material body with embedded magnetic pieces for integrated actuation. The system apparatus can be actuated to perform an activity on a subject by first using a device to position a shape-memory material body in a subject. This shape-memory material body is formed into a specific primary shape, reformed into a secondary stable shape, and controllably actuated with the magnetic pieces to recover the specific primary shape and perform the activity on the subject.

### Mechanically Robust, Electrically Conductive Ultralow-Density Carbon Nanotube–Based Aerogels

**Marcus A. Worsley, Theodore F. Baumann, Joe H. Satcher, Jr.**
U.S. Patent 8,685,287 B2
April 1, 2014
In this method of making a mechanically robust, electrically conductive, ultralow-density carbon nanotube–based aerogel, nanotubes are first dispersed in an aqueous media or other media to form a suspension. Reactants and a catalyst are then added to the suspension, creating a reaction mixture. The mixture is cured to form a wet gel, and the wet gel is dried to produce a dry gel. The dry gel is pyrolyzed to produce the carbon nanotube–based aerogel. The aerogel is made of a porous carbon material that is 5 to 95 percent by weight carbon nanotubes and 5 to 95 percent carbon binder.

### Multi-Pass Amplifier Architecture for High Power Laser Systems

**Kenneth R. Manes, Mary L. Spaeth, Alvin C. Erlandson**
U.S. Patent 8,687,270 B2
April 1, 2014
This main amplifier system includes a first reflector that receives input light through a first aperture and directs the input light along an optical path. The input light is characterized by a first polarization. The system also includes a first polarizer that reflects light characterized by the first polarization state. The system further includes a first and second set of amplifier modules. Each set of modules includes an entrance window, a quarter wave plate, numerous amplifier slablets arrayed substantially parallel to each other, and an exit window. A set of mirrors reflects the light exiting the first set of amplifier modules onto the second set of amplifier modules. A second polarizer reflects light characterized by a second polarization state.

## Awards

The **Precision Strike Association** awarded the 18th annual **William J. Perry Award** to the **BLU-129/B Team**, which includes researchers from Lawrence Livermore, the U.S. Air Force, and Aerojet Rocketdyne, Inc. The award honors the immediate and long-term impact of BLU-129/B on combat operations. As part of this joint effort, the Laboratory leveraged its long-term investments in computational codes, engineering expertise, and computing and manufacturing infrastructure to develop the munition in record time. The weapon represents a new class of innovative munitions, integrating disruptive technologies that significantly reduce collateral damage. The ability to couple sophisticated guidance systems with weapons that have a more accurate lethal footprint has been profound. The award is named in honor of former Secretary of Defense William J. Perry (1994–1997) and recognizes exceptional contributions to precision strike systems in the private or public sector by an individual or team.

**Mark Zelinka** of the Laboratory's Program for Climate Model Diagnosis and Intercomparison earned funding from **NASA** under its **New Investigator Program in Earth Science**. The NASA program, aimed at researchers who have earned their Ph.D. within the last five years, will fund Zelinka up to $100,000 per year over three years. With the funding, Zelinka will use detailed data from NASA satellites along with climate models to examine the vertical structure and optical properties of clouds. His goal is to understand how and why clouds change as the planet warms and what may be the implications of these changes for Earth's climate. The NASA award is designed to support outstanding scientific research and career development of scientists and engineers at the early stage of their professional careers. The Earth Science Division places particular emphasis on an investigator's ability to promote and increase the use of space-based remote sensing through the proposed research.

Two Laboratory researchers won the prestigious **Ernest Orlando Lawrence Award** for their contributions to the **Department of Energy**'s (DOE's) missions in science, energy, and national security. Livermore seismologist **Stephen Myers** was recognized for his work advancing national security and nonproliferation by developing seismic monitoring technologies to locate nuclear explosions. Former Livermore scientist **Siegfried Glenzer**, an 18-year Laboratory veteran who joined SLAC National Accelerator Laboratory at Stanford University last year, was recognized for his work advancing fusion and plasma sciences at Livermore's National Ignition Facility.

The E. O. Lawrence Award honors midcareer scientists and engineers for exceptional contributions in research and development supporting DOE's National Nuclear Security Administration (NNSA). Named for the physicist who cofounded Lawrence Livermore, it comes with a citation signed by the Secretary of Energy, a gold medal bearing the likeness of E. O. Lawrence, and $20,000.

Livermore's **Leslie Positeri** and **Brian Molyneaux** received the **Learning Champion of the Year Award** at **Skillsoft**'s 2014 conference. This award honors individuals who show remarkable skills and innovation in delivering a successful learning program. Positeri is a senior training and development specialist in the Strategic Human Resources Management Directorate, and Molyneaux is the eLearning information technology lead from the Business Directorate. Together, they manage the Laboratory's U-Learn Program.

Laboratory scientists **Jennifer Pett-Ridge** of the Chemical Sciences Division and **Todd Gamblin** of the Center for Applied Scientific Computing were selected by **DOE**'s **Office of Science Early Career Research Program** to receive up to $2.5 million over five years for their proposed projects.

Pett-Ridge will focus on understanding how changes in climate (rainfall and temperature regimes) in the tropics may affect the capacities of soil microbial communities that drive decomposition, nutrient availability, and carbon stabilization. The work will involve both field-scale ecosystem manipulations and controlled laboratory experiments. Her research will take advantage of Livermore resources such as the nanometer-scale secondary-ion mass spectrometer and the Center for Accelerator Mass Spectrometry as well as collaborations with DOE's Joint Genome Institute and Environmental Molecular Sciences Laboratory user facilities.

Gamblin will accelerate the adaptation of scientific simulation codes to increasingly powerful supercomputers, a process that currently can take up to six months for advanced applications. Increasingly complex machine architectures and applications are making this process even slower. Gamblin's research is particularly important as the high-performance computing (HPC) community prepares to ramp up computing speeds from petascale (quadrillions

of floating-point operations per second) to exascale systems that will be as much as 1,000 times faster. HPC experts believe the first exascale systems will come online in the 2020 time frame.

Three teams of Laboratory researchers and two individuals were honored with **NNSA Defense Programs Awards of Excellence** for work performed in 2012.

The **Primary Design Code Team** developed and implemented a new capability that has enabled unprecedented high fidelity in simulations of a wide range of programmatic applications. Examples include nuclear weapons disablement techniques, high-explosives initiation, safety-related experiments, effects of three-dimensional geometric features on weapons performance, and inertial-confinement fusion. The new capability provides a significantly improved representation of the underlying physics for many fine-scale phenomena.

The **L1 Milestone Team** used advanced science-based models and novel uncertainty quantification (UQ) methods to assess the quality of early-phase hydrodynamic models and to make a prediction, with quantified uncertainties, of a hydrodynamic experiment at Los Alamos National Laboratory. The team developed and exercised a functional UQ methodology that incorporated both focused and integral physics experiments. This adaptive methodology allowed the team to complete UQ studies in a fraction of the time required by past methods. This achievement enabled, for the first time, a systematic analysis of all relevant focused physics experiments related to early-phase hydrodynamics and how they inform science-based models. The effort clearly described the weaknesses and strengths of the various models and where improvements are needed.

The **Department of Defense/DOE TATB Technical Working Group**—composed of many members including Lawrence Livermore researchers—played a critical role in finalizing a domestic production capability for the high-temperature explosive TATB (triaminotrinitrobenzene) at BAE Systems's Holston Army Ammunition Plant in Tennessee and successfully qualified laboratory- and pilot-scale TATB production through the readiness campaign. Several years of effective leadership led to the restoration of a full-scale production capability for a material that is critical to the safety of both the conventional and nuclear weapons stockpiles. The determination and expertise of the group's diverse participants enabled this complex endeavor.

Through years of consistent effort, **Paul Nowak** has developed a modern three-dimensional computational physics capability that has significantly redefined the simulation approach to nuclear weapons design and assessment, as used in the annual assessment review and in developing options for future life-extension programs. This capability combines improved physics fidelity with a computational efficiency that allows routine use on massively parallel supercomputers.

Over the last three decades, **Wigbert Siekhaus** has developed a complete and comprehensive model of the aging behavior of a major component in the detonation system used in Lawrence Livermore and Los Alamos nuclear explosives. Documentation of this extensive work was completed in 2012. His results aided the decision-making process behind the forwarding of a Livermore weapon system for consideration as the warhead in the proposed long-range stand-off weapon. The corrosion of gold wires by lead–tin–indium solders has been a potential source of performance degradation in detonators since the early days of the Manhattan Project. It also is a source of concern for electronic components fielded in long-duration space missions supported by NASA.

A team of researchers led by **Andreas Kemp** received a **2014 Leadership Computing Challenge Award** from **DOE**'s **Office of Advanced Scientific Computing Research** (ASCR) to further pursue the study of short-pulse laser interactions with solid-density plasmas using supercomputer simulations. The research project entitled "Laser-Driven Relativistic Electron Beam Filamentation in Solids" was awarded 30 million central-processing-unit hours on the Titan Cray XK7 supercomputer at Oak Ridge National Laboratory, the world's second fastest supercomputer. The ASCR Leadership Computing Challenge Awards provide large computing time allocations on leadership-class supercomputer systems for projects of interest to DOE. The award will allow for large-scale simulations to be performed in support of Kemp's 2012 DOE Early Career Research Program Award. This computer modeling work applies high-energy-density laser physics (HEDLP) to such astrophysical problems as gamma-ray bursts and the origin of cosmic rays. Other applications range from fusion energy to the controlled amplification of desirable radiation or particle beams of interest for possible medical use. Computer modeling is essential to the design and interpretation of experiments in HEDLP.

## Abstract

### A New Model for Pharmaceutical Research

The pharmaceutical industry spends increasing amounts of time, money, and effort to develop new drugs. Researchers at Lawrence Livermore assert powerful computational tools can help the industry improve the drug-development process. A team of computational biologists and biochemists is endeavoring to demonstrate how high-performance computing can accelerate the early stages of drug development and reduce the risks associated with subsequent clinical trials. Their work is also relevant to national security, because broad-spectrum antibiotics often serve a dual role as countermeasures for bioterrorism threats. The researchers are using supercomputer-based modeling and simulation to screen drug–protein pairings, identify unintended binding events, predict whether a drug will cross the blood–brain barrier, and study drug breakdown within the human body. Ultimately, many levels and types of simulations will be linked to produce a robust predictive capability that will help support and guide the drug development process.

***Contact: Felice Lightstone (925) 423-8657 (lightstone1@llnl.gov).***

U.S. Government Printing Office: 2014/770-075-10004

# Sharper Focus for Adaptive Optics

Photographer Marshall Perrin, Space Telescope Science Institute

Advanced adaptive optics technologies are opening new frontiers for imaging exoplanets and x rays.

***Also in September***

*• A software tool for mapping computer networks offers cybersecurity situational awareness.*

*• A collaboration between Livermore and Texas A&M University is combining high-performance computing with advanced materials science to improve aneurysm treatment.*

*• Livermore weapons and additive manufacturing experts are developing customized foams for stockpile stewardship applications.*

Coming Next Issue

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

PRSRT STD
U.S. POSTAGE
PAID
San Bernardino, CA
PERMIT NO. 3330

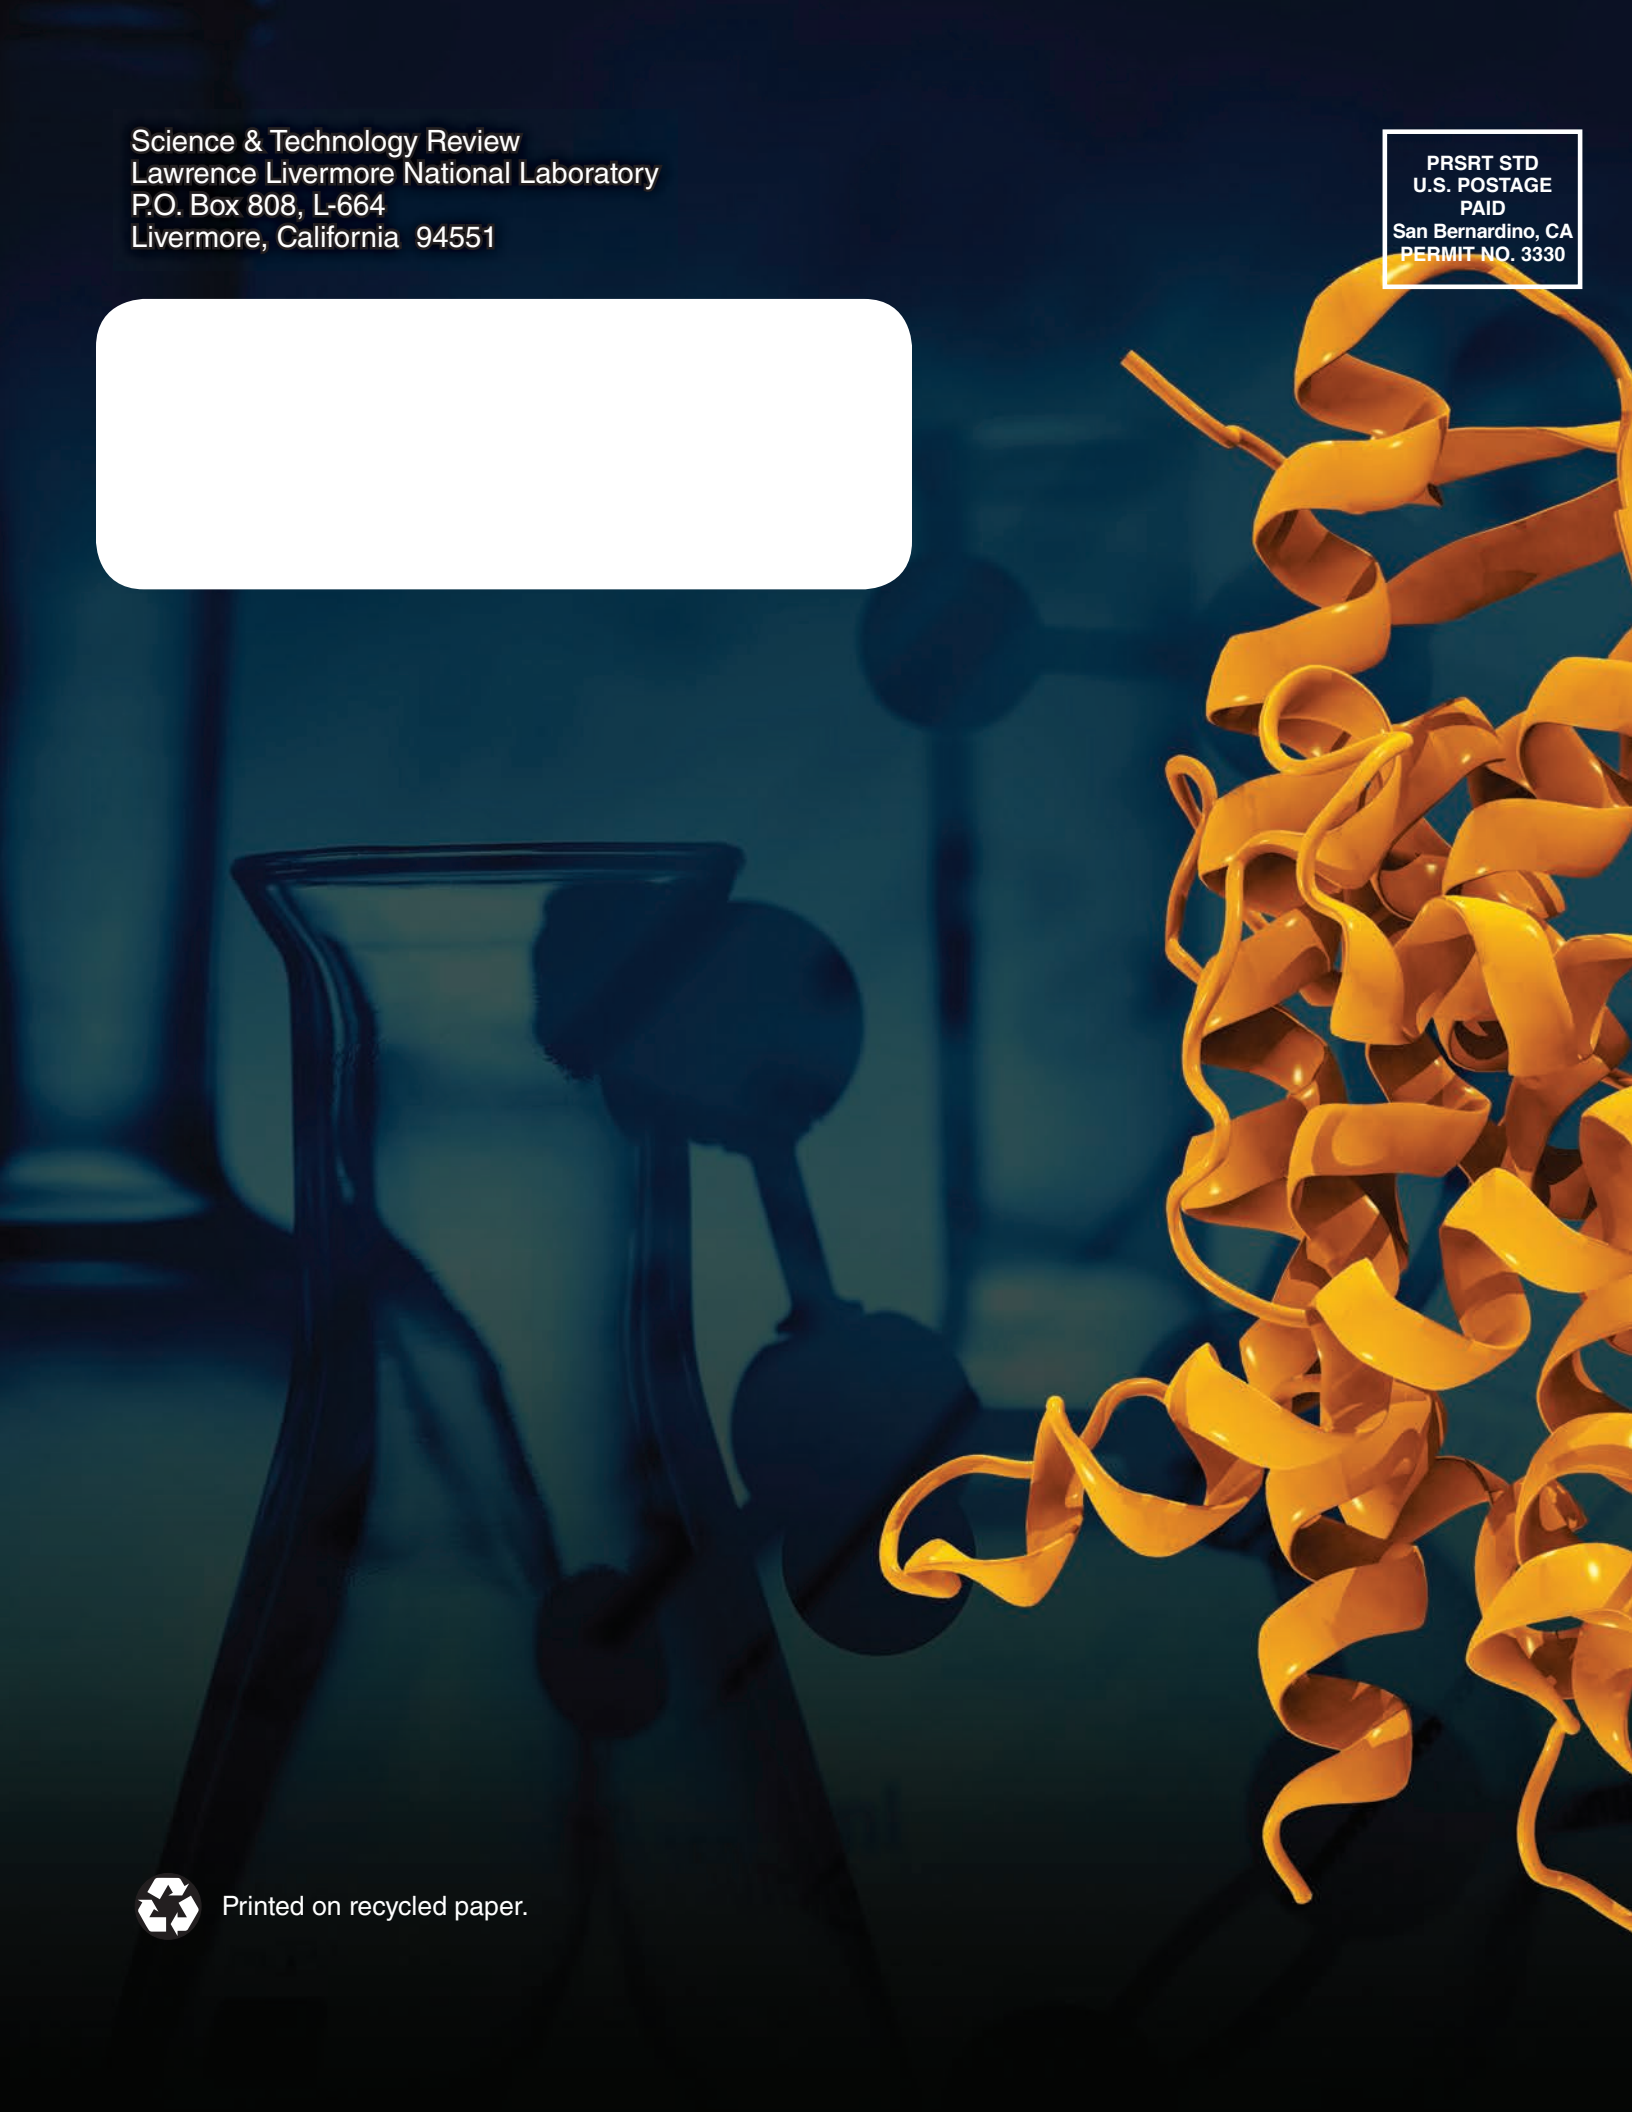

Printed on recycled paper.