January/February 2013

Science & Technology
REVIEW

Solving the Data
Science Challenge

Also in this issue:

Graphite Transition Surprises

A New Tool for Desalination

Diffractive Optics for Space

About the Cover

Lawrence Livermore scientists are developing novel algorithms and computational methods to address the data overload challenge that plagues fields as diverse as bioinformatics, cybersurveillance, and climate science. As the article beginning on p. 4 describes, their efforts range from devising methods for predicting the evolution of viruses, to creating tools for tackling streaming data in the cybersecurity arena, to fashioning accessible intuitive portals and analytics for climate scientists worldwide.

Cover design: George A. Kitrinos

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.

Contents

Feature

Research Highlights

Departments

Changing the Dynamics of Bulk Materials

Lawrence Livermore researchers have developed a method to fabricate graphene-based bulk materials from polymer-derived carbon foams. Carbon atoms are selectively removed from a network composed of both unstructured carbon and graphite nanoplatelets. "The new technique is inexpensive, scalable, and yields mechanically robust, centimeter-size monolithic samples composed almost entirely of interconnected networks of single-layer graphene nanoplatelets," says Ted Baumann, who developed the synthetic approach.

The graphene bulk materials have an ultrahigh surface area and may thus be used for energy storage systems such as supercapacitors, where energy is stored by polarization of the graphene electrode–electrolyte interface. Graphene bulk material could also be used as an electrically conductive network to support the active material in battery applications. An emerging field in which the bulk material could be applied is capacitive deionization, a desalination method in which ions are removed from electrolytes to create a clean water source. (See the article on p. 16.)

The advantage of using bulk materials rather than composite materials is their superior stability, which allows for longer lifetimes, higher conductivity (less loss during charging and discharging), and the ability to tune the pore structure. "This concept is potentially game-changing in the area of materials science," says Juergen Biener, lead Livermore author of the cover article in the September 25, 2012, issue of *Advanced Materials*. Other institutions contributing to this effort include the Karlsruhe Institute of Technology, Technische Universität Darmstadt, and Technische Universität Hamburg–Harburg.
Contact: Ted Baumann (925) 423-5543 (baumann2@llnl.gov).

Cold Cases Heat Up

Livermore scientist Bruce Buchholz and a team of international collaborators have found a multidisciplinary approach to identifying the remains of missing persons. Using "bomb pulse" radiocarbon analysis developed at the Laboratory (see *S&TR*, April/May 2010, pp. 15–17), combined with recently developed anthropological analytic and forensic DNA techniques, the researchers were able to identify the remains of a missing child four decades after discovery of the body.

In 1968, a child's cranium was recovered from the banks of a northern Canadian river. Initial analyses using the technology at the time concluded that the cranium came from the body of a seven- to nine-year-old child, but no identity could be determined. At the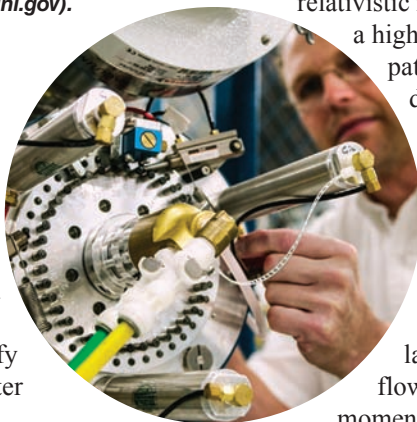 Laboratory's Center for Accelerator Mass Spectrometry (shown center), researchers recently conducted radiocarbon analysis of enamel from two of the child's teeth and determined a more precise birth date—within one to two years. Forensic DNA analysis, conducted at Simon Fraser University in Canada, indicated the child was a male,

and the obtained mitochondrial profile matched a living maternal relative to the presumed missing child. The multidisciplinary analyses indicated an age at death of approximately four and one-half years and resulted in a legal identification 41 years after the discovery of the remains. The effort highlights the potential of combining radiocarbon analysis with anthropological and mitochondrial DNA analyses to produce confident personal identifications in forensic cold cases dating to within the last 60 years.

"Thousands of John and Jane Doe cold cases exist in the United States," says Buchholz, who conducted the radiocarbon analysis in the case. "We could provide birth and death dates for many of these cases." Other institutions participating in the research are the Karolinska Institutet in Sweden and the British Columbia Institute of Technology. The research appeared in the September 2012 issue of the *Journal of Forensic Sciences.*
Contact: Bruce Buchholz (925) 422-1739 (buchholz2@llnl.gov).

New Regimes Explored in Fast Ignition

In an experiment led by the Laboratory's Yuan Ping using the Titan laser at the Jupiter Laser Facility, researchers performed the first time-resolved measurements on the hole-boring process in the relativistic regime with a novel diagnostic called Specular FROG (frequency-resolved optical gating). The data show an unexpected slowing down of hole boring even when the laser intensity is still increasing over time.

When a laser beam is reflected from a mirror, a finite amount of momentum—carried by light—is transferred to the mirror so that the mirror is slightly pushed. At high laser intensities into the relativistic regime, the light pressure is intense enough to push a high-density plasma and create a channel along the light path. The channel helps to guide the laser beam and deliver energetic electrons to ignite the core of fuel. Previous models on hole boring only take into account the ion momentum. Electron momentum has been ignored because electrons have much less mass. In the relativistic regime, however, electrons become so energetic that their momentum is no longer negligible. According to Livermore's Andreas Kemp, these electrons quickly leave the laser–plasma interaction region. A return current flows to compensate the charge but not the loss of momentum. As a result, the hole boring becomes slower than previous predictions.

The results provide a new understanding of relativistic laser–plasma interaction, in particular how the energy and momentum are partitioned among different groups of particles. The research team includes scientists from Lawrence Livermore, General Atomics, University of California at San Diego, Ohio State University, and Princeton University. The team's research appeared in the October 5, 2012, edition of *Physical Review Letters*.
Contact: Yuan Ping (925) 422-7052 (ping2@llnl.gov).

Big Data in Science Is a Big Challenge

WHEN the *Harvard Business Review* takes on the topic of data science and announces that data scientist is "The Sexiest Job of the 21st Century," one knows that "big data" has become the latest big news. However, data science has been at the core of the Laboratory's mission from the very beginning, when data was measured in kilobytes, as compared to the petabytes that are processed and analyzed today.

We live in a world awash in data, a world of complex, interconnected systems and networks, from the vast systems that control and run various aspects of the country's infrastructure, to systems embedded in manufacturing processes, to the smartphones and tablets on which we depend in our personal lives. Here at Lawrence Livermore, data science involves taking massive amounts of raw data in a variety of forms—for example, ocean-surface temperatures, spectra gathered from telescopes scanning the night sky, and DNA sequences—and analyzing it to extract information about relationships, patterns, and connections, and then presenting that information in forms that can be used to make decisions.

Today, the challenges of big data sweep across all our mission areas—from biosecurity and counterterrorism to nonproliferation and weapons systems. Data science also plays a major role in science-focused areas such as climate, energy, astrophysics, and high-energy physics. Any mission that is involved with collecting and analyzing enormous quantities of data will turn to data science as part of the analysis process. The difference between the Laboratory and a commercial entity is that our focus is on ensuring national security—the driver behind our efforts.

All our missions are affected by the need to push the limits and applications of data science. To be relevant in the current century, we must excel in this field. To that end, the Laboratory has established a new core competency initiative in data science, headed by James Brase, deputy associate director of Data Science in the Computation Directorate. As the article beginning on p. 4 describes, the initiative will focus on pattern discovery, predictive models and simulation, and the underlying data-intensive computing and management that enables us to handle high data volumes and rates.

The pattern discovery area involves machine learning—exploring methods and algorithms that will allow data-science tools to evolve to detect ever-changing patterns and activities. One such effort is Ana Paula de Oliveira Sales's project to develop a method for analyzing streaming data as it evolves.

The data-intensive computing thrust will explore the computer architecture required to meet data-science challenges. For example, a Laboratory Directed Research and Development project is exploring new computer architectures that use flash memory. (See *S&TR*, January/February 2012, pp. 23–25.)

Data science has huge challenges ahead, but vast rewards beckon. At Livermore, we are engaged in forward-leaning areas of investigation, pushing the boundaries in all directions, whether the subject is biosecurity, such as typified by the effort to predict the viability of mutated viruses, or climate science research, such as the Earth System Grid Federation, an international collaboration whose portals include Livermore's Program for Climate Model Diagnosis and Intercomparison.

To succeed in these areas and elsewhere, addressing the data tsunami is vitally important. The Laboratory will contribute with the excellence of effort that it brings to bear on all its missions, providing the country's decision makers with the information they need to do their jobs.

■ Dona Crawford is associate director for Computation.

Dealing with Data Overload

Livermore researchers develop methods for keeping up with the tsunami of data.

WE may think we have problems managing our ever-increasing stream of electronic personal "data," whether that information comes in the form of e-mails, social network updates, or phone texts. However, the challenges we face are miniscule compared with those faced by scientists who must parse the growing flood of scientific data critical to their work. From sequences to simulations to sensors, modern scientific inquiry is awash in electronic information. At Lawrence Livermore, computer scientists supporting various projects and programs are developing methods and algorithms that provide new ways to tame, control, and understand these large amounts of data.

"Four key steps are involved in solving data-science problems," explains Dean Williams of Livermore's Global Security Principal Directorate. "One must organize the data—arrange the numbers and bits into meaningful concepts. One must prioritize—choose the most useful data when time and resources are limited. One must analyze—find meaning in images, graphs, data streams, and so on. Finally, it's important to make it easy for researchers to use the data; that is, we must create the methods and systems that help users query, retrieve, and visualize."

Three "V's" can sum up the type of data and the challenges involved: the variety, the velocity, and the volume. For example, in biological mission areas, the variety is high, but the velocity is low, with the volume that needs to be manipulated ranging from gigabytes to terabytes. By contrast, in the cybersecurity arena, variety and velocity are high, and the volume, which is continually changing as data streams past, can be very large. For climate research, the variety and velocity of data are also high, with an enormous volume accumulating in databases worldwide (from petabytes to exabytes).

The Laboratory, with its broad expertise in analysis, experience in storage technologies, and strong institutional computing culture, is addressing the data-science challenges of its programs and projects. These efforts range from devising methods for predicting the evolution of viruses, to creating tools for tackling streaming data in the cybersecurity arena, to fashioning accessible intuitive portals and analytics for climate scientists worldwide.

Getting Ahead of Viral Evolution

Scientists have been sequencing genomes for decades, with full genetic sequences now completed for more than 1,000 organisms as well as viruses—subcellular organisms with genomes consisting of RNA or DNA. Viruses are of particular interest because of their worldwide impact on health and welfare, the difficulty of combating them in nature, and the potential for their use as biological threat agents. (See *S&TR*, September 2012, pp. 6–13.) RNA viruses have exceptionally high mutation rates, enabling them to form mixed-variant virus populations, often referred to as "quasi-species." The high genetic variability within quasi-species helps these viruses adapt to different environments and hosts. Understanding such genetic diversity, especially in pathogenic viruses, is critical to developing accurate diagnostics and therapeutics. Although the information in existing DNA sequence databases is incomplete, the variety and volume of new data can be overwhelming.

in the Scientific Realm
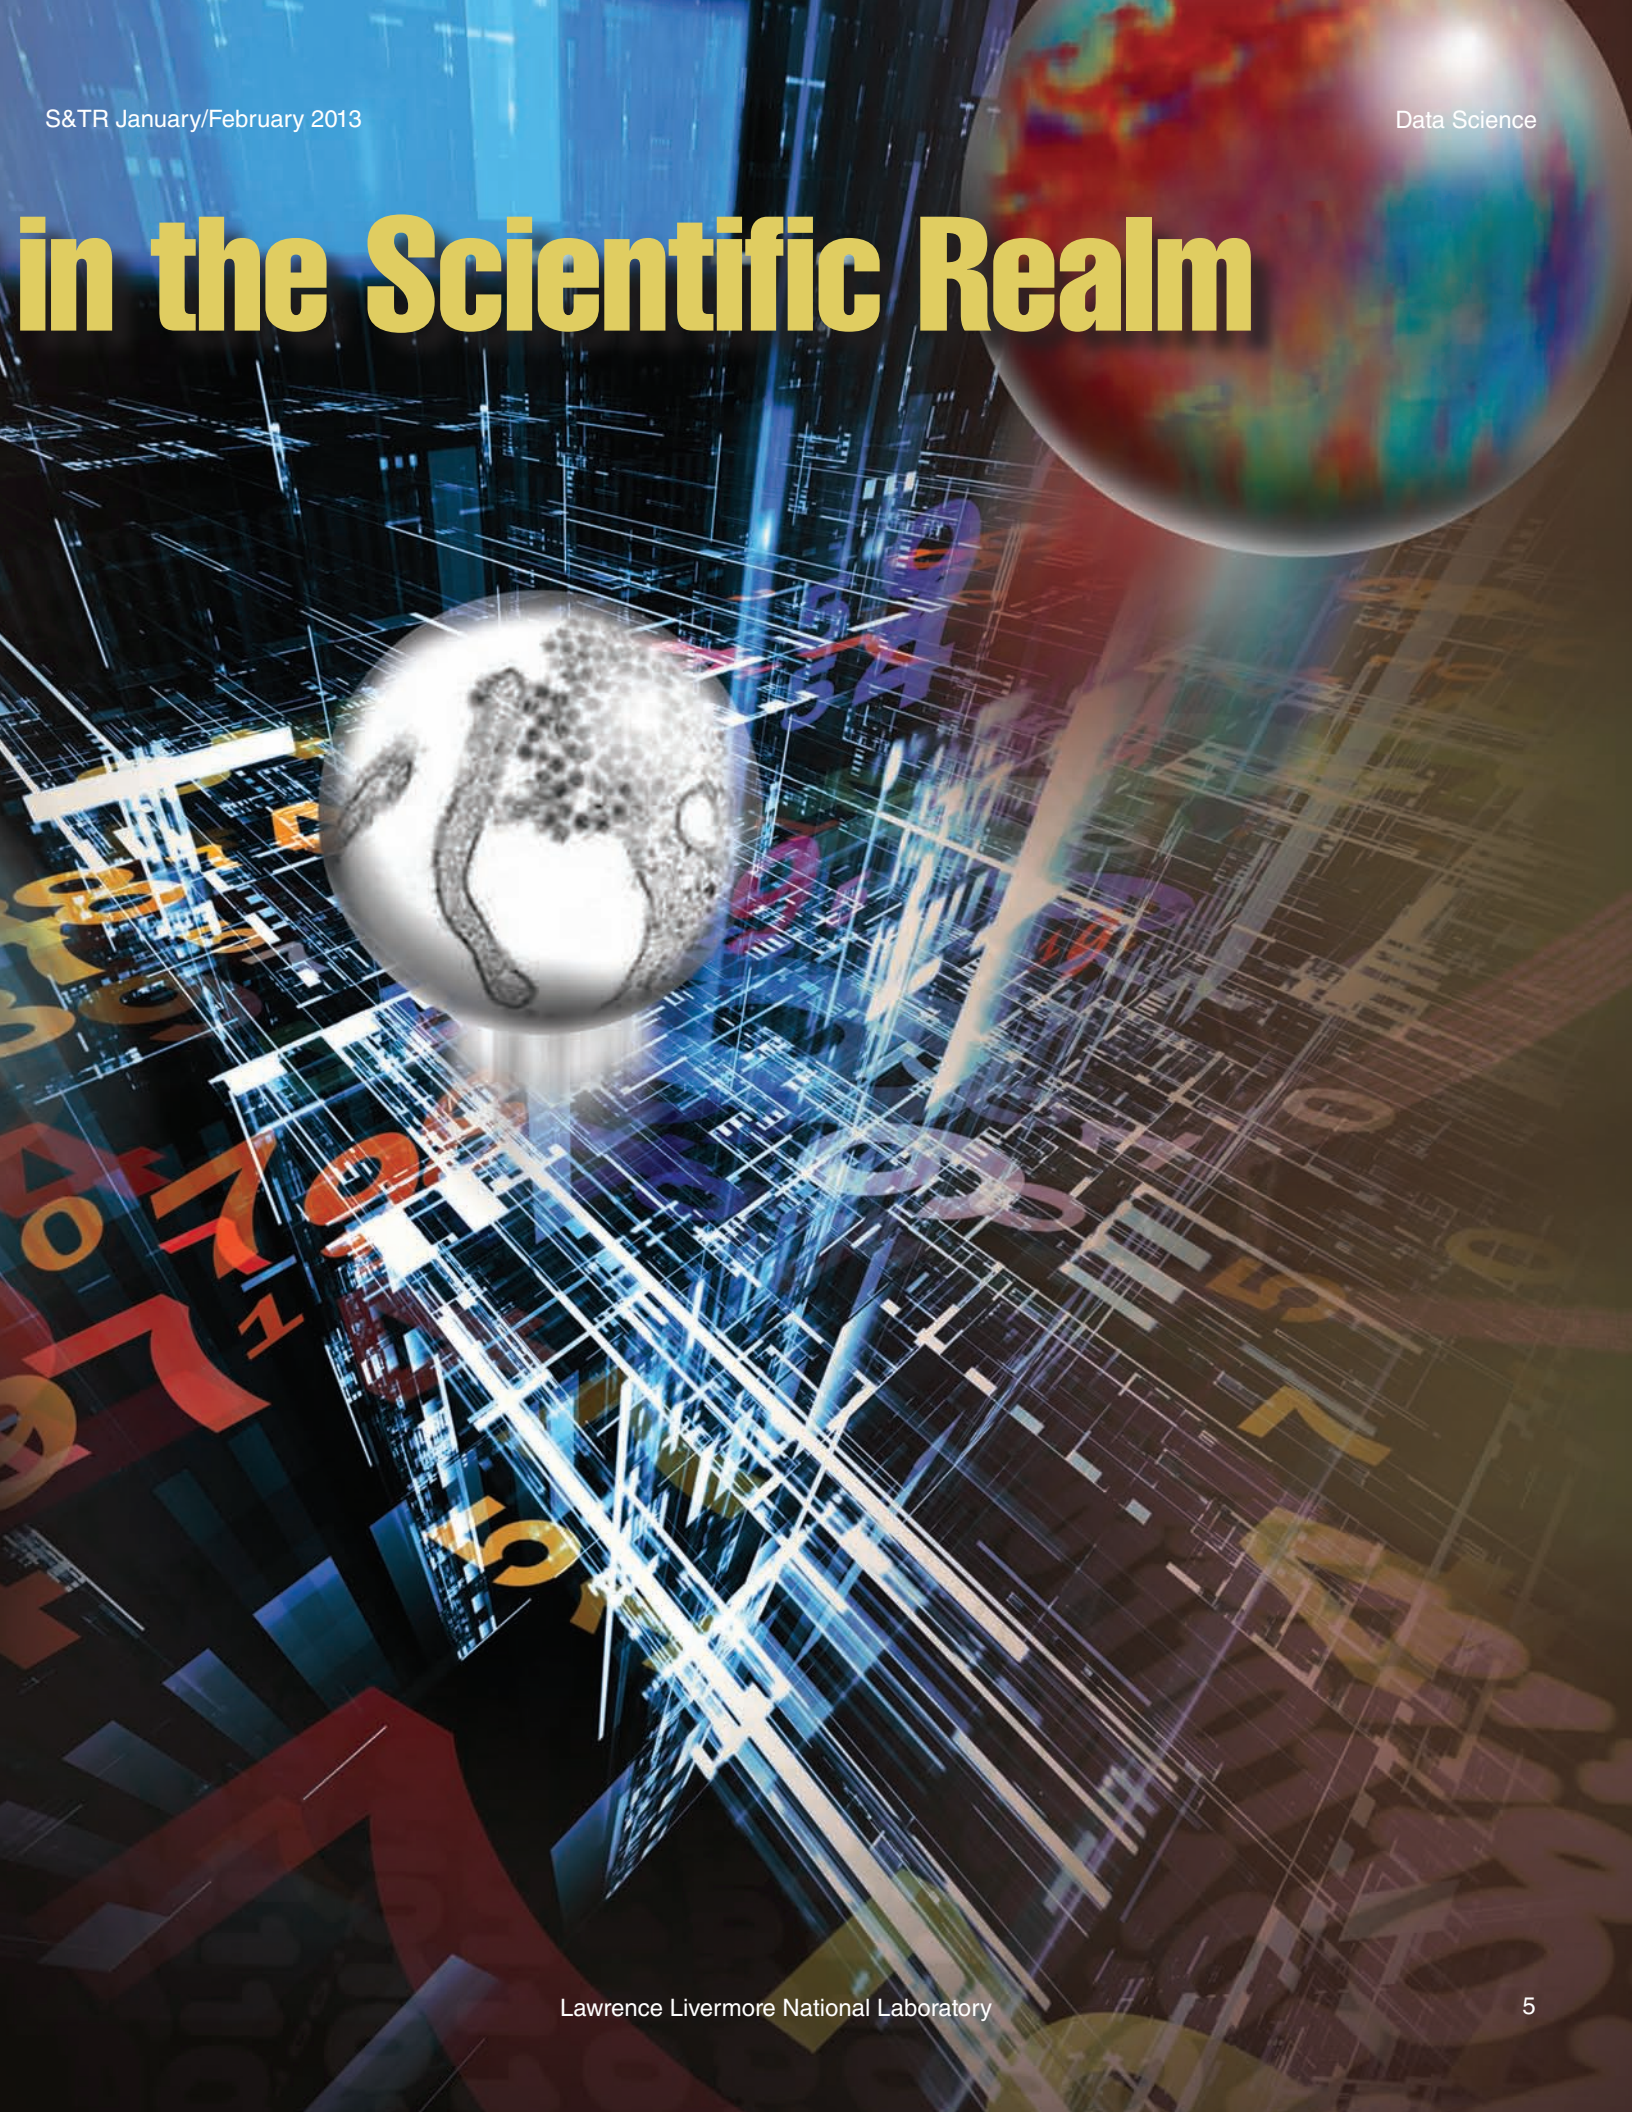
To address this issue, computer scientists Adam Zemla and Tanya Kostova Vassilevska are developing a computational system called GeneSV and a stochastic simulation model to help predict viral evolution that could lead to the emergence of new strains or quasi-species clouds. Zemla's GeneSV allows characterization of possible sequence variations within a viral genome and makes predictions about the viability of a potential viral mutation. Vassilevska's model uses GeneSV results to further simulate viral evolution. Zemla and Vassilevska collaborated with researchers from the University of Texas Medical Branch (UTMB) at Galveston, who conducted experiments to test GeneSV's predictions. "This collaboration with experimentalists was like a dream project for a bioinformatician such as myself," says Zemla. "We designed and developed new algorithms, used these algorithms to create predictions and hypotheses, and worked with Galveston experimentalists to test and validate the predictions."

Funded by the Defense Advanced Research Projects Agency (DARPA), Livermore scientist Pejman Naraghi-Arani led a team to create these computational modeling tools that, when used with a novel microfluidics platform to grow viruses under many conditions, could

evaluate and predict aspects of viral evolution. For this project, the team focused on a fast-evolving RNA flavivirus, the Dengue virus type 2 (DENV-2). Flaviviruses—a genus that includes the Dengue, West Nile, and yellow-fever viruses—mutate quickly and are of particular interest to biologists, health care specialists, and biodefense researchers.

"The RNA viruses have mutation rates between 10^{-3} and 10^{-5} per base per generation. In a single replication event of a virus, such as DENV-2 in a host cell, the majority of the produced genomes will have at least one variation," says Zemla. "Each host cell produces hundreds or thousands of progeny viruses, which can infect other host cells. This exponential growth of the number of multivariants helps the virus spread in the organism."

Public genomic databases used by researchers contain tens of thousands of genomic sequences as well as three-dimensional protein models of different viruses. However, this information encompasses only a part of the genetic diversity of viral species. "Furthermore," says Zemla, "the databases tend to be biased toward the dominant viral genome. It's unlikely that the full array of viable viral genotypes of a given species will ever

be represented in these databases, because of sampling biases and the fact that these quasi-species clouds are constantly evolving." Because in vivo research and experiments cannot keep up with the millions of viral mutations and determine which genomic shifts lead to a viable virus, in silico predictive systems such as GeneSV are invaluable tools as they become more fully developed.

GeneSV uses information from existing genomic sequences, related protein sequences, and constructed protein-structure homology models to classify new variants and assess their viability. The system starts with a mutated gene sequence and searches the databases for similar sequences in closely related organisms. For example, if the input is a sequence from an unknown variant of DENV-2, GeneSV first looks in databases at sequences of other Dengue serotypes. (A serotype is a group of closely related microorganisms distinguished by a characteristic set of antigens.) It also examines sequences from more distantly related viruses, such as West Nile, and compares those sequences to the unknown mutant.

GeneSV identifies positions in the genomic sequences where the compared virus types—unknown and known—are different or similar. Then GeneSV

Using the Livermore-developed GeneSV computational system, Laboratory scientists analyzed never-before-seen designer mutants of an RNA polymerase gene (gene NS5) from the infectious Dengue virus type 2 (DENV-2) clone. Experimental tests conducted at the University of Texas Medical Branch at Galveston showed that the algorithm correctly predicted viability of evaluated mutants more than 80 percent of the time. In this transmission electron micrograph of Dengue virus virions (a virion is a single virus particle), the virions show up as a dark cluster of dots in the center. (inset) A capsid model simulates one virion.

(bottom left) A fragment of the DENV-2 genomic sequence shows highlighted mutation position candidates for GeneSV system analysis and experimental tests. (center) Two components of the GeneSV system are StralSV sequence–structure analysis and the AS2TS structure modeling algorithms. (right) The final structural AS2TS model of the NS5 polymerase gene shows target mutation positions.

Sequence–structure analysis


```
ACCAACTCAGTGCTGCCCCCCCCAATATC
GACGACCAACCGCTGTTCTGGCCCGCCA
GTCAAGTATGCGCTACGCCATGATAAAT
CAGACGGGCTTCTTCTTCGCCACTGATG
ACTTGAATAGGGCAGCCGAGGATCTCGG
TTGCAGTACGCAAGAATTCGTCGGCATT
AGAGGCCAGGCGCAAGGTTGCCGTTGCG
ACCACGGCAGCATAAAGTGAGGTGGAAC
TTAGCATCATTTATCAGCCTGGTCGCTT
ACGGTTATTCCGTGGTGCTAGCATCTAA
ACAGTTTGAATCTAAAGTTTTACTCTAT
GTCGAATTGATTTGTGT......
```

AS2TS structure modeling

Structural model: NS5 polymerase

converts the nucleic-acid sequences to protein sequences and continues similarity searches for closely related organisms on the protein level, using the same methodology as before. Finally, GeneSV builds protein structural models and conducts similarity searches against protein-structure databases.

For each level of evaluation—genomic sequence, protein sequence, and protein structure—GeneSV creates predictions and identifies characteristics that can be assigned to particular positions within the unknown, mutated sequence. One such characteristic is the location of the region in question on the protein structure, for instance, whether that region is buried inside or exposed on the surface. Another characteristic is sequence regions conserved within closely related organisms. "For example," says Zemla, "to invade a host cell, the virus binds to a specific receptor on the cell. The regions and residues responsible for binding need to be conserved in some ways, but at the

same time, they must be able to mutate to avoid detection by the immune system."

The algorithms in GeneSV were designed to allow detection of such characteristics, with each algorithm focusing on a different aspect. For analysis on the protein-structure level, the system leverages the StralSV (structure alignment–based sequence variability) and AS2TS (amino-acid sequence to tertiary structure) algorithms developed under previously funded Laboratory Directed Research and Development (LDRD) projects to identify regions of sequence or structure variability. This diverse combination of sequence- and structure-based approaches results in a detailed analysis of the proteins of interest within a given genome.

For example, GeneSV was used to estimate the frequencies of mutations observed in codon (nucleotide triplet) positions in different genes within the set of all currently available DENV-2 genomic sequences. Results showed that the most mutable regions are in the envelope protein within segments characterized as helical, exposed, and predicted as antigenic determinants.

For the DARPA project, the team used GeneSV to evaluate two kinds of possible mutations for the NS5 polymerase gene from the DENV-2 virus: single-point mutations and compensatory mutations. (Compensatory mutations involve a process in which two or more residue positions, sometimes closely located in three-dimensional space, mutate simultaneously to preserve protein function.) GeneSV identified a set of proposed 32 single-mutation types, 31 of which had never been observed in any publicly available DENV-2 sequence. The UTMB scientists engineered each of the mutated viruses and experimentally confirmed that in 26 of the 32 mutations, GeneSV correctly predicted the viability of the mutants generated. GeneSV also generated five possible compensatory mutations in which double amino-acid substitutions occurred in two different parts of the genomic structure. Four of the five GeneSV predictions were experimentally proven to be correct.

The success of Phase I of the DARPA project has Zemla and colleagues hopeful for the future. With information such as that generated by GeneSV, medical researchers would have an advantage against possible attacks from all kinds of viruses. "Our approach can be applied to genomic sequences from any organism," says Zemla. "The current version of GeneSV has shown potential to help characterize possible variations in genomic sequences and sequence annotation efforts."

Finding Relevant Data in the Flood

The world is riddled with viruses of all kinds and not just of the biological variety. In the cybersecurity and surveillance arenas, threats akin to viruses abound, but they may go undetected within vast amounts of data. Livermore's Ana Paula de Oliveira Sales is leading an LDRD project focused on improving the ability to analyze streams of data that arrive at very high rates.

Her work could not only help fighters of cyber-crime but also benefit biosurveillance and real-time energy-distribution efforts. The challenges include the variety of data and the velocity at which it streams.

"In many domains of science, our ability to collect data continues to outpace our ability to analyze data," Sales explains. The deficit is particularly apparent in areas that involve streaming data. "For example, consider the common electronic communication that surrounds us every day, such as Twitter, texts, and Internet searches," says Sales. "For cyber-surveillance experts, monitoring the constant, data-rich activity is difficult. Storing all the information is impractical. Our project focuses on narrowing this gap between data collection and analysis rates by analyzing data 'on the fly' as the information streams past."

Sales notes that some of the key statistical problems underlying many of

today's national security applications share a common need for continuously deployable, self-adapting learning systems. The goal of Sales's research is to create innovative computational learning algorithms that enable using sophisticated predictive-modeling techniques on modern streaming-data sources. This robust platform must quickly and accurately classify behavior and detect anomalies in large data streams, while effectively interacting with an ever-changing stream of data. "We need alternative predictive models for accomplishing standard learning tasks at a fraction of the computational cost," says Sales.

Sales and her colleagues designed and tested a highly customizable system to classify behavior and detect anomalies in large data streams. The code is implemented in high-performance C++,

and the approach is flexible and applicable to a vast array of domains, such as real-time threat detection, video surveillance, energy distribution, and biosurveillance.

The system combines Bayesian techniques of "particle filtering" and time-evolving composite mixture models. "Particle filters are a kind of stochastic 'survival of the fittest' algorithm, evolving over time and adapting to their environment—in this case, incoming data," says Sales. The system begins with a data set for developing initial models, or particles. For example, if the system were used for spam detection, each particle would represent a possible model of spam. A particle would look at every new e-mail and answer the question: How probable is it that this e-mail is spam? "The question would be answered by considering variables such as sender, e-mail subject line, and
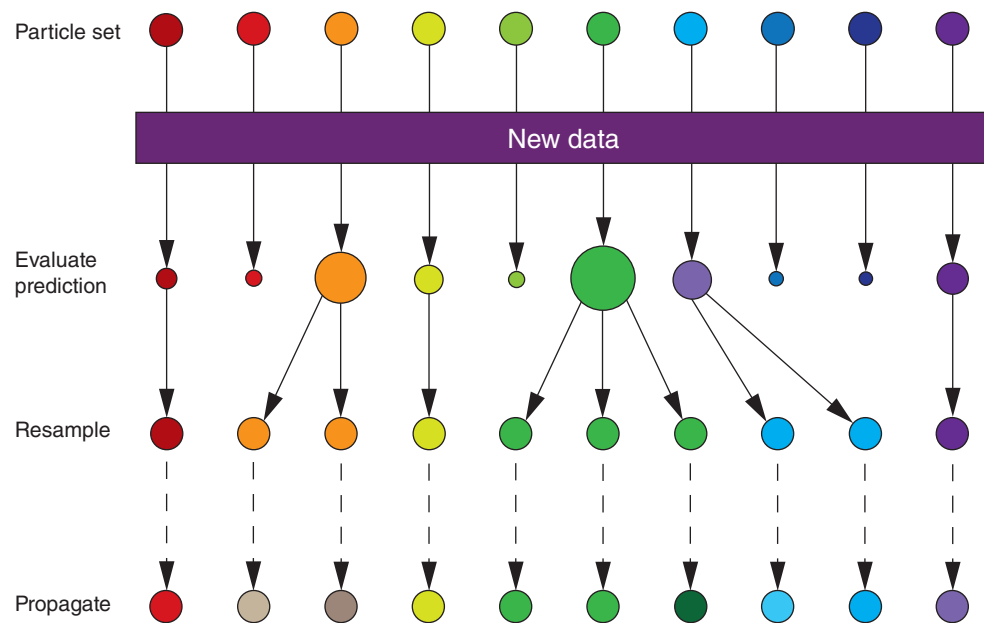
Particle filters are key to analyzing streaming data for anomalies. A particle set (a set of possible models) makes predictions about a new piece of data. The predictions are evaluated, and those models that are "more correct" are given more weight, while those that prove "less correct" receive less weight. In this way, the system evolves and changes character along with new, incoming data.

e-mail body," says Sales. "The beauty of using the composite mixture model is that it enables us to deal with a variety of data, numerical or categorical, and combinations thereof. Our model was intentionally built to be flexible and modular so it can be easily deployed to different domains."

When new data arrive to be analyzed, such as an e-mail message, the algorithm makes a prediction. In this case, the algorithm might determine that the e-mail message has a certain probability of being spam. The initial particles are resampled using weights that are proportional to how well each particle fits each new piece of information. As such, particles that are "more correct" receive more weight than those that are more incorrect, and they have a greater chance of surviving the resampling process. But even those that have less weight may continue to exist because of the stochastic nature of the system. "This feature is important," says Sales, "because it allows for greater diversity among the models. By having a variety of particles, we can be more confident about how well our ensemble represents the data."

The resampled particles then update their probability densities using relevant information from the data point (in this example, the e-mail message). Over time, the ensemble of particles becomes a more accurate representation of the data. This evolution allows the model to more effectively counter the adversary, while the adversary—in this example, the spammer—is also evolving to bypass the filter system.

To date, the Livermore-developed algorithm has primarily been used to analyze data sets that evolve gradually. A current research direction for the project, now in the second year of its three-year LDRD funding, involves teaching the algorithm to accommodate sudden-switch situations. Examples of

such situations include an abrupt change in power distribution as a result of a local blackout or the attack of a new computer virus. In these cases, the system needs to be able to "switch on a dime" when vital information changes suddenly and without warning.

Going forward, the project's primary focus is increasing computational performance, so that the system can be deployed for data streams with higher and higher rates of data arrival. One technique Sales and colleagues are pursuing is to make the approach more parallelizable by creating specialized small filters, or "ensembles of ensembles." Replacing a single, large particle filter with many smaller, parallel filters should significantly improve computational speed at little or no cost to prediction performance.

Sales expects that the most gains in computational speed will come from using adaptive sampling theory. "Updating the probability densities with each new data point is the most computationally expensive aspect of the model," says Sales. "We'll use adaptive sampling to intelligently choose data points with high information value for updating the model." The Laboratory has significant expertise in adaptive design of computer experiments in cases where data are scarce and data points sparse. "Using this technique, we could restrict expensive model updates to those observations that improve our understanding of the system dynamics," Sales explains.

Coding for Collaboration

Laboratory scientists are also devising methods that make it easier to deal meaningfully with enormous quantities of stored data. While the physical storage of exabytes (10^{18} bytes) of data is achievable today, the challenge is to make the data meaningful and widely used. Nowhere is this issue more apparent than in the

realm of climate research, where the amount of information at researchers' virtual fingertips continues to grow at an enormous rate. The types of information resources include observational data from National Aeronautics and Space Administration satellites and instruments; in situ, ground-based data such as albedo measurements and surface and air temperatures; simulations data; and reanalysis data (a mix of observational and model data). In climate research, all three of the basic challenges are present: volume, velocity, and variety of data.

The gathering and sharing of climate data is a key effort of the Coupled Model Intercomparison Project (CMIP), which is the worldwide standard experimental protocol for studying the output of coupled atmosphere–ocean general circulation models. Established in 1995 by the World Climate Research Programme's Working Group on Coupled Modelling (WGCM), CMIP provides a community-based infrastructure in support of climate model diagnosis, validation, intercomparison, documentation, and data access. Through CMIP, scientists are able to analyze general circulation models in a systematic fashion, a process that serves to facilitate model improvement. Virtually the entire international climate modeling community has participated in this project since its inception.

Williams, who leads the Earth System Grid Federation (ESGF) project at Livermore and abroad, remembers that in 2003, to gather the data used for the third CMIP (CMIP3), large-size "bricks" that could hold a single terabyte of data were shipped around the globe. Climate researchers loaded their data to send back to Livermore, where the entire 35 terabytes of data were then stored in a single centralized location. "We also created a Web portal so the user community could access the database," says Williams.

Fast forward to 2006, when WGCM agreed to promote a new set of coordinated climate model experiments. These experiments comprise the fifth phase of CMIP. CMIP5 is providing a multimodel context to assess the mechanisms responsible for model differences associated with the carbon cycle and with clouds, examine the ability of models to predict climate on decadal timescales, and more generally, determine why similarly forced models produce a range of responses.

For CMIP5, the total data are estimated to reach 3.1 petabytes (10^{15} bytes)—two orders of magnitude more than the total of CMIP3. To increase the accessibility and usefulness of the mountains of CMIP5 climate data, the international ESGF was formed (www.esgf.org). The federation grew out of the larger Global

Organization for Earth System Science Portals community. Collaborating partners in the federation include Livermore and other national laboratories as well as a host of other organizations such as the German Climate Computing Centre and the University of Tokyo Center for Climate System Research.

The ESGF portals are gateways to scientific data collections hosted at sites around the globe. Gateways are Web portals that allow the user to register and potentially access the entire ESGF network of data and services. Currently more than 24 portals are in use, including Livermore's Program for Climate Model Diagnosis and Intercomparison.

Instead of consolidating all the data from numerous locations, as was done for CMIP3, ESGF uses distributed storage in

conjunction with software that harvests and integrates the data. "Far too much data exists to ship it about physically, so we took a decentralized, cloud-type approach," says Williams. Modeling research centers download the ESGF software stack from the Livermore servers. Centers then use this software to publish their data to the federation for harvesting. Users access all data as if they are on one centralized archive system.

To date, ESGF has made 60 CMIP5 simulation model runs (more than 1.8 petabytes) from 25 climate research centers, which are available to users worldwide. ESGF provides access to 18 highly visible national and international climate data products, with more on the way. As a result, ESGF offers a promising option for

The Earth System Grid Federation (ESGF) peer-to-peer architecture is based on the principles of modularity and equality. Each system node (blue, orange, pink, chartreuse) can offer different services depending on how it is configured. Some nodes can be scaled differently, for example, to provide increased computational resources. All nodes interact as equals, so no single points of failure arise.
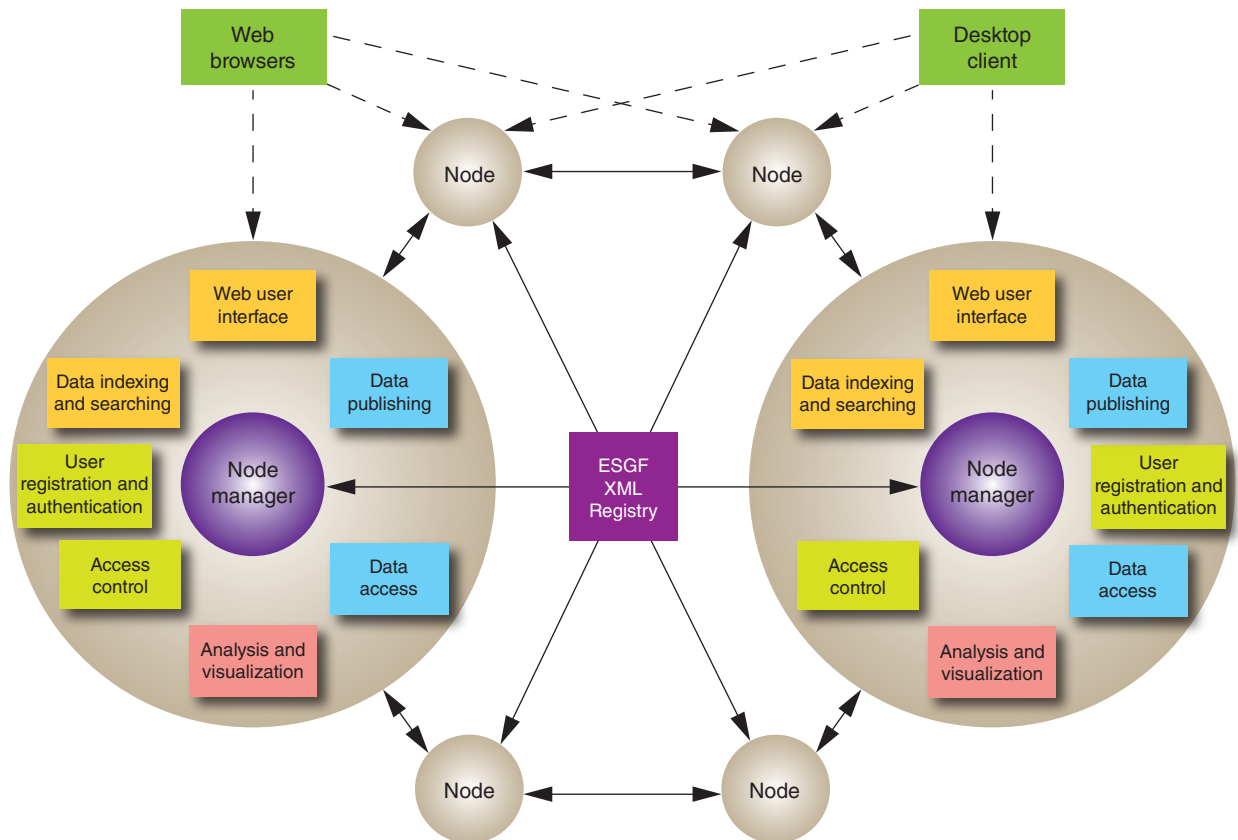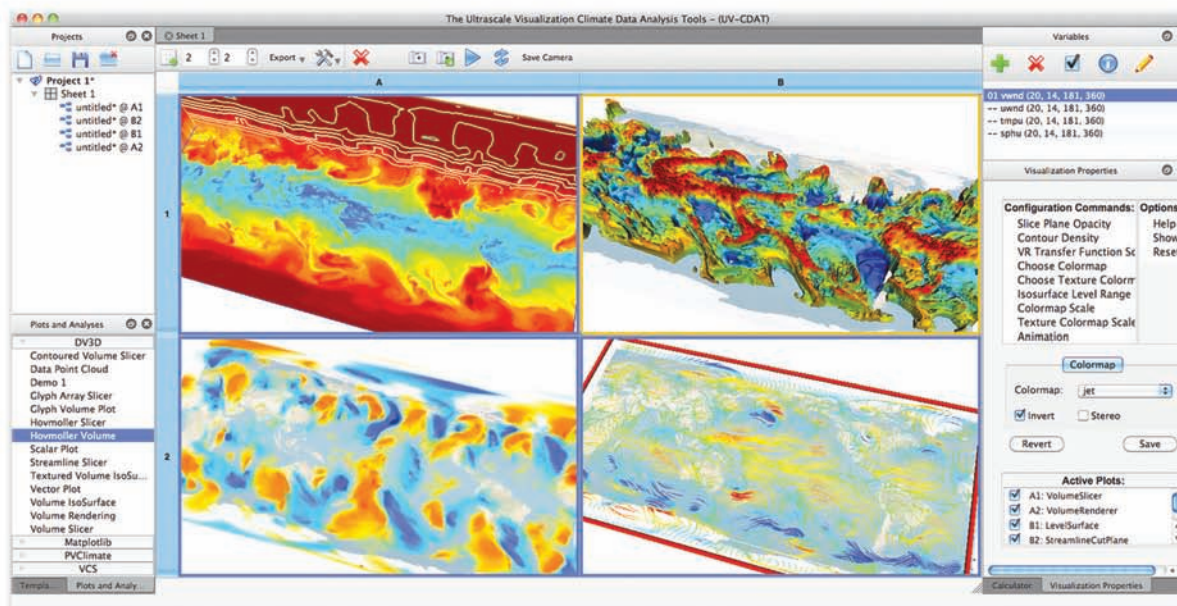
The Ultrascale Visualization Climate Data Analysis Tools display visualization results from different geographical data sets around the world via the ESGF distributed archive. (Shown here are atmospheric simulations.) Users can select from a variety of analytical and visualization tools such as CDAT, R, DV3D, VisIt, ParaView, and Matplotlib. The framework is flexible enough that other tools can be added at any time. Workflows can also be stored for later use and sharing with others.

building a collaborative knowledge system in the climate community.

The system makes all of this data-handling transparent to the user, while still allowing for local ownership. "If a modeling center improves a computer model and produces new output, the system handles archiving and notifications to those people who use the data," explains Williams.

The ESGF peer-to-peer architecture is based on the concept of a dynamic system of nodes that interact on an equal basis and offer a broad range of user and data services, depending on how each node is set up. This extensible and scalable system supports geospatial and temporal searches and includes a dashboard that shows system metrics, a user interface for notifications, and a rich set of analysis tools to help manipulate the data. For example, the Ultrascale Visualization Climate Data Analysis Tools have workflow scripts that automate scientific analysis and visualization, making it easy for users to re-run analyses and to work together,

which encourages collaboration and openness in scientific enquiry.

Williams envisions systems that will make it even easier for scientists to collaborate in the future. "I'd like to see all of the visualization tools put on the back end, with the results easily available on a laptop," he says. "We'd like to pull in more elements, such as Twitter and Whiteboard, and have people be able to use their tablets, smartphones, and whatever else may be coming in the years ahead." Williams is also looking toward the creation of an ESGF "Lite" version that would work similarly to today's social networks. By 2020, ESGF will embrace an estimated hundreds of exabytes; thus, tools for the future will be welcome.

Looking toward a Data-Rich Future

In nearly every arena of scientific inquiry imaginable, the amount of data needed to be organized, prioritized, analyzed, and utilized will continue to increase in velocity, variety, and volume. The data sources in many cases are huge and growing dynamically. Whether it's

forecasting ecological tipping points, predicting and mitigating energy grid instability, or identifying new computer viruses, the tools must evolve to meet and match the needs of researchers. "Prioritization will become ever-more important as data volumes and velocities grow," says Williams. "Organization will also gain importance as data sets become increasingly complex." The present challenges of analysis and utilization will remain as well. Livermore computer scientists and researchers are already addressing many of the challenges and have their sights set on the data-rich future that is just around the corner.

—Ann Parker

Graphite, a Quick-Change Artist

UNDER the glare of the brightest x-ray beam in the world, graphite crystals "melted" much faster than expected. The process took just 40 femtoseconds—40-quadrillionths of a second. In 1 second, light can travel almost to the moon. In contrast, in 1 femtosecond light can travel through a single sheet of plastic wrap.

An international team led by Livermore physicist Stefan Hau-Riege conducted the experiments at the Linac Coherent Light Source (LCLS) at SLAC National Accelerator Laboratory in Menlo Park, California. (See the box on p. 13.) Team members included scientists from SLAC and, in Germany, the University of Duisburg-Essen, Max Planck Advanced Study Group at the Center for Free-Electron Laser Science, Max Planck Institute for Medical Research, and Max Planck Institute for Nuclear Physics.

LCLS's tightly focused beam of "hard" x rays is 10^9 times brighter than previous-generation x-ray sources and 10^{21} times brighter than medical x rays. The laser's pulse lengths are very short, and its wavelength is about the size of an atom. These features allow researchers to capture images of objects that are ultrasmall, such as the DNA helix or a water molecule, and that move ultrafast, such as single atoms and possibly electrons.

The laser's fast-pulse, strobelike beam can capture stop-action shots of moving molecules, revealing details never seen before. Then, almost instantaneously after LCLS images a sample, the sample is blasted to smithereens by the x-ray beam.

Why Study Graphite?

Graphite, the softest polymorph of carbon, was chosen for this study in part because its structure is well known. Imaging experiments are important to better understand biological molecules, all of which are based on carbon. Proteins are of particular interest because of their importance in living cells. Scientists generally use x-ray crystallography to reveal the macromolecular structure of proteins. However, success relies on growing crystals of sufficient size, which is often not possible. LCLS offers an excellent alternative for examining many proteins as well as other types of living molecules.

The graphite experiments allowed Livermore researchers to also study materials under extreme conditions (very high temperature, pressure, and density), an effort that relates to the Laboratory's role in maintaining the safety, security, and reliability of the nation's nuclear weapons stockpile. A key task for stewardship of the stockpile is to fully understand and be able to predict precisely how materials behave when exposed to extreme conditions. Drawing on decades of experience with weapons, researchers have developed computer codes to model material behavior. The challenge, however, says Livermore physicist Frank Graziani, is to test the codes to ensure their accuracy.

Says Graziani, "LCLS gives us the opportunity to study material properties that are accessible to both simulation and experiment." Experimental results from LCLS complement those at the National Ignition Facility, which has a laser system that can generate even higher pressure, temperature, and density regimes.
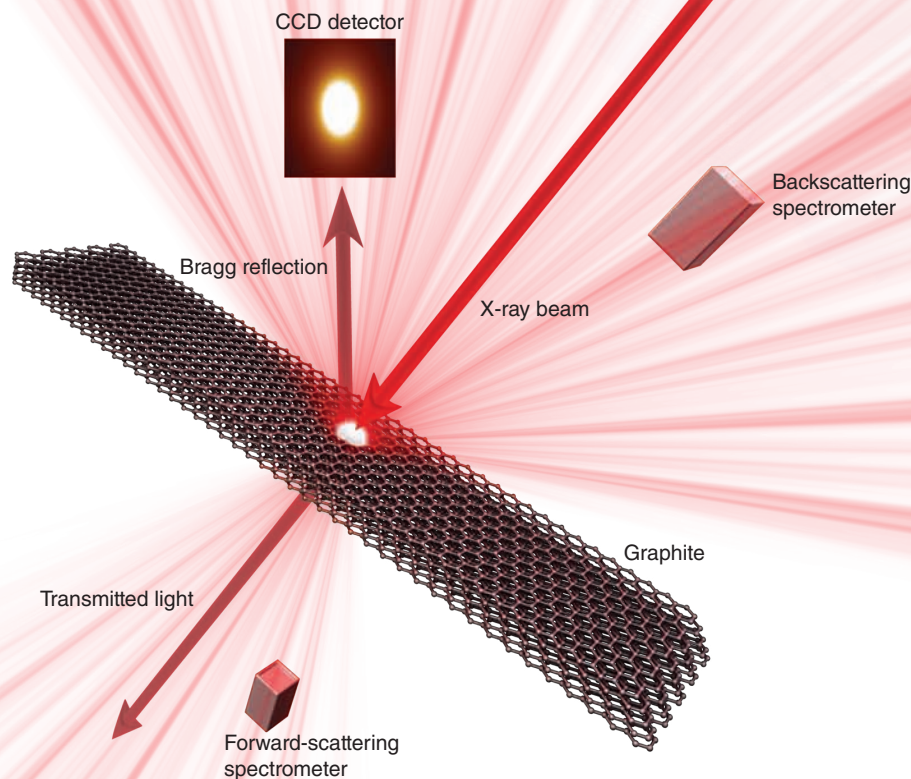
At the Linac Coherent Light Source

The Linac Coherent Light Source (LCLS) at SLAC National Accelerator Laboratory in Menlo Park, California, uses a portion of SLAC's existing 3.2-kilometer-long linear accelerator (linac) to create the beam of electrons that generate x-ray laser light. LCLS is the brightest x-ray source in the world and the first such machine to probe matter with hard x rays. Most lasers, including Livermore's powerful National Ignition Facility, excite electrons that are bound in gas, liquid, or solid-state materials. The LCLS's lasing medium is a beam of "free" electrons moving through a vacuum chamber at almost the speed of light.

The first LCLS experiments began in October 2009, and a team from Lawrence Livermore helped commission the laser. (See *S&TR*, January/February 2011, pp. 4–11.) Early experiments achieved an in-depth understanding of how the beam interacts with matter as well as its operating characteristics, including the exposure required to damage materials. LCLS is designed to capture unprecedented details of molecules in motion that are important for research in chemistry, materials science, high-energy-density physics, and medicine. The LCLS project is a collaboration of SLAC; Lawrence Livermore, Argonne, Brookhaven, and Los Alamos national laboratories; and the University of California at Los Angeles.

Livermore experts designed and fabricated the optics that transport the x-ray beam to chambers in two experimental halls. These mirrors help control the size and direction of the x-ray beam. Livermore also fabricated diagnostic detectors that help LCLS operators fine-tune the rest of the machine to ensure it produces laser light most efficiently. Most of these optics and diagnostics are in use today to support LCLS operation for various user experiments. Livermore's work on LCLS is funded in part by the Laboratory Directed Research and Development Program.

CCD detector

The setup for the graphite experiments simultaneously uses a charge-coupled-device (CCD) detector, which collects Bragg-reflected light, and crystal x-ray spectrometers, which collect light from inelastically scattered photons, to observe extremely fast changes in graphite. (Rendering by Kwei-Yu Chu.)
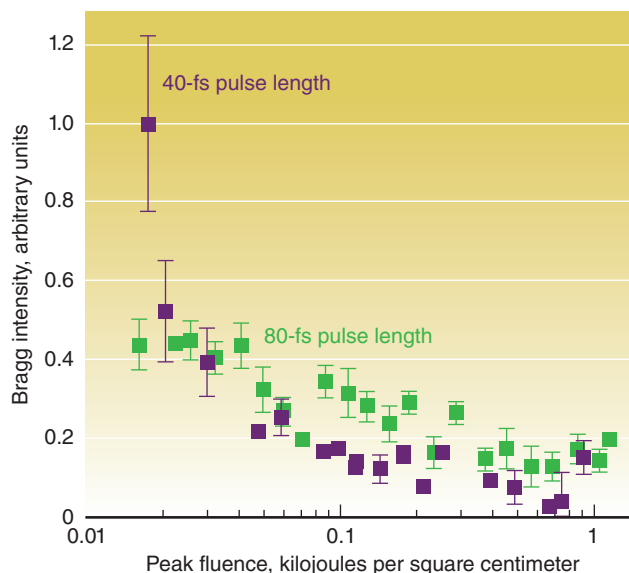
Bragg reflection

Backscattering spectrometer

X-ray beam

Graphite

Transmitted light

Forward-scattering spectrometer

Graphite Surprises

In the graphite experiments, Hau-Riege's team measured the breakdown of the solid crystal using two characterization techniques. In one method, a charge-coupled-device detector directly images the "Bragg signal" produced by x rays elastically (without change of energy) reflecting off the graphite at an angle precisely determined by the crystalline structure. A second method uses x-ray spectrometers to measure the energy spectrum of scattered photons as the crystal breaks down.

The experiments were performed using 2-kiloelectronvolt x rays in pulses that were 40, 60, and 80 femtoseconds in duration with varying amounts of pulse energy up to 2.8 millijoules. Using the two types of diagnostics and differing pulses, the team was able to extract the time dependence of the interaction process. At low x-ray intensities, the crystal structure stayed intact, and the Bragg signal was very strong. Also, diffuse scattering was weak. With increasing intensities, the Bragg signal degraded as a result of ionization and the motion of the atoms. The diffuse scattered signal increased correspondingly. The combined measurements allowed the team to determine the speed at which the crystal melts through changes in ion temperature.

Two surprises came out of the experiments. First, the Bragg signal dropped off at unexpectedly low levels of peak fluence.

40-fs pulse length

80-fs pulse length

Bragg intensity, arbitrary units

Peak fluence, kilojoules per square centimeter

The Bragg intensity drops off at just 0.02 peak fluence—a surprisingly low level. The measured pulse-averaged Bragg intensity as a function of the peak fluence is shown for two different pulse lengths—40 and 80 femtoseconds (fs).
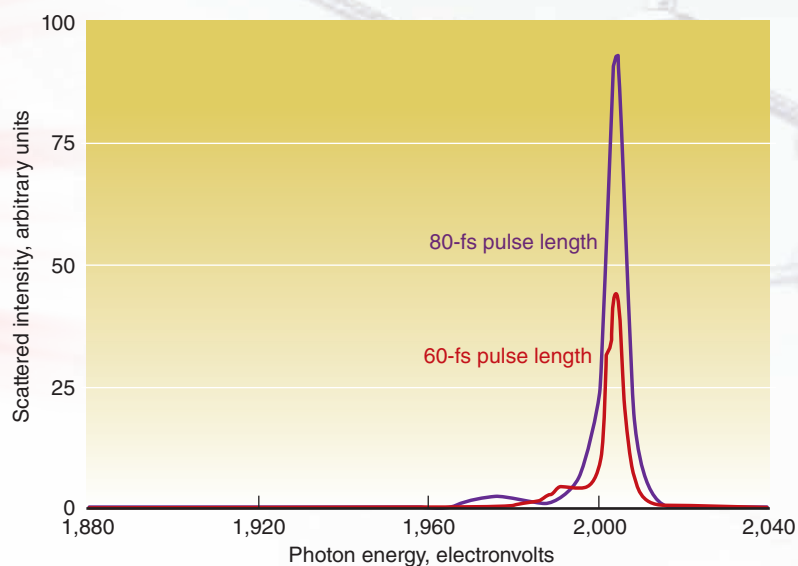
The scattering intensity increased sharply as a function of the pulse length (60 and 80 fs), indicating that rapid heating and transition from crystal graphite to warm, dense matter occurs much faster than theory predicted.

For the 40-femtosecond pulses, the team observed a prominently large drop in the Bragg reflectivity (by a factor of 2.5) as the fluence increased above 0.02 kilojoules per square centimeter. A similar effect was not seen for the 80-femtosecond pulses because of a lack of very-low-fluence data. The cause of this drop in reflectivity may stem from phenomena involving the electrons and not solely from the motion of the atoms.

The diffuse scattering spectra gathered in the experiments also revealed that the rate at which the graphite ions heated up was much faster than predicted by simulation models. The elastic-scattering intensity increased with pulse energy and pulse length. The spectral intensity sharply peaked at the energy of the incoming light (2 kiloelectronvolts) because most of the scattering was elastic (a phenomenon called Rayleigh scattering). These data were used to determine the integrated Rayleigh intensity (normalized to the incoming pulse energy) as a function of peak-pulse intensity and pulse duration, and from that, ion temperature was determined as a function of time.

The solid graphite crystal became warm, dense matter— too cold to be a plasma but no longer solid—much sooner than expected. Results from both characterization techniques showed that the solid crystal of graphite transitioned to warm, dense matter in 40 femtoseconds instead of more than

100 femtoseconds as predicted by computer simulations. In simulations using molecular-dynamics models, graphite's phase changes were significantly underestimated as were the timing and extent of the material's transformation.

These results may also have introduced a wrinkle in plans to study biological molecules if LCLS disturbs the sample before it can be imaged. However, Hau-Riege is quick to note that "these experiments are just the beginning." In fact, his team's experiments at LCLS were some of the first to be performed. Their efforts continue with experiments using silicon.

Hau-Riege gives high praise to LCLS as an experimental tool. "The facility's hard x-ray beam is the first available to us for experimental use," says Hau-Riege. "We would not have had these results without LCLS because we previously had no way to detect such phenomena."

—*Katie Walter*

Key Words: graphite, Linac Coherent Light Source (LCLS), phase change, hard x-ray laser.

For further information contact Stefan Hau-Riege (925) 422-5892 (hauriege1@llnl.gov).
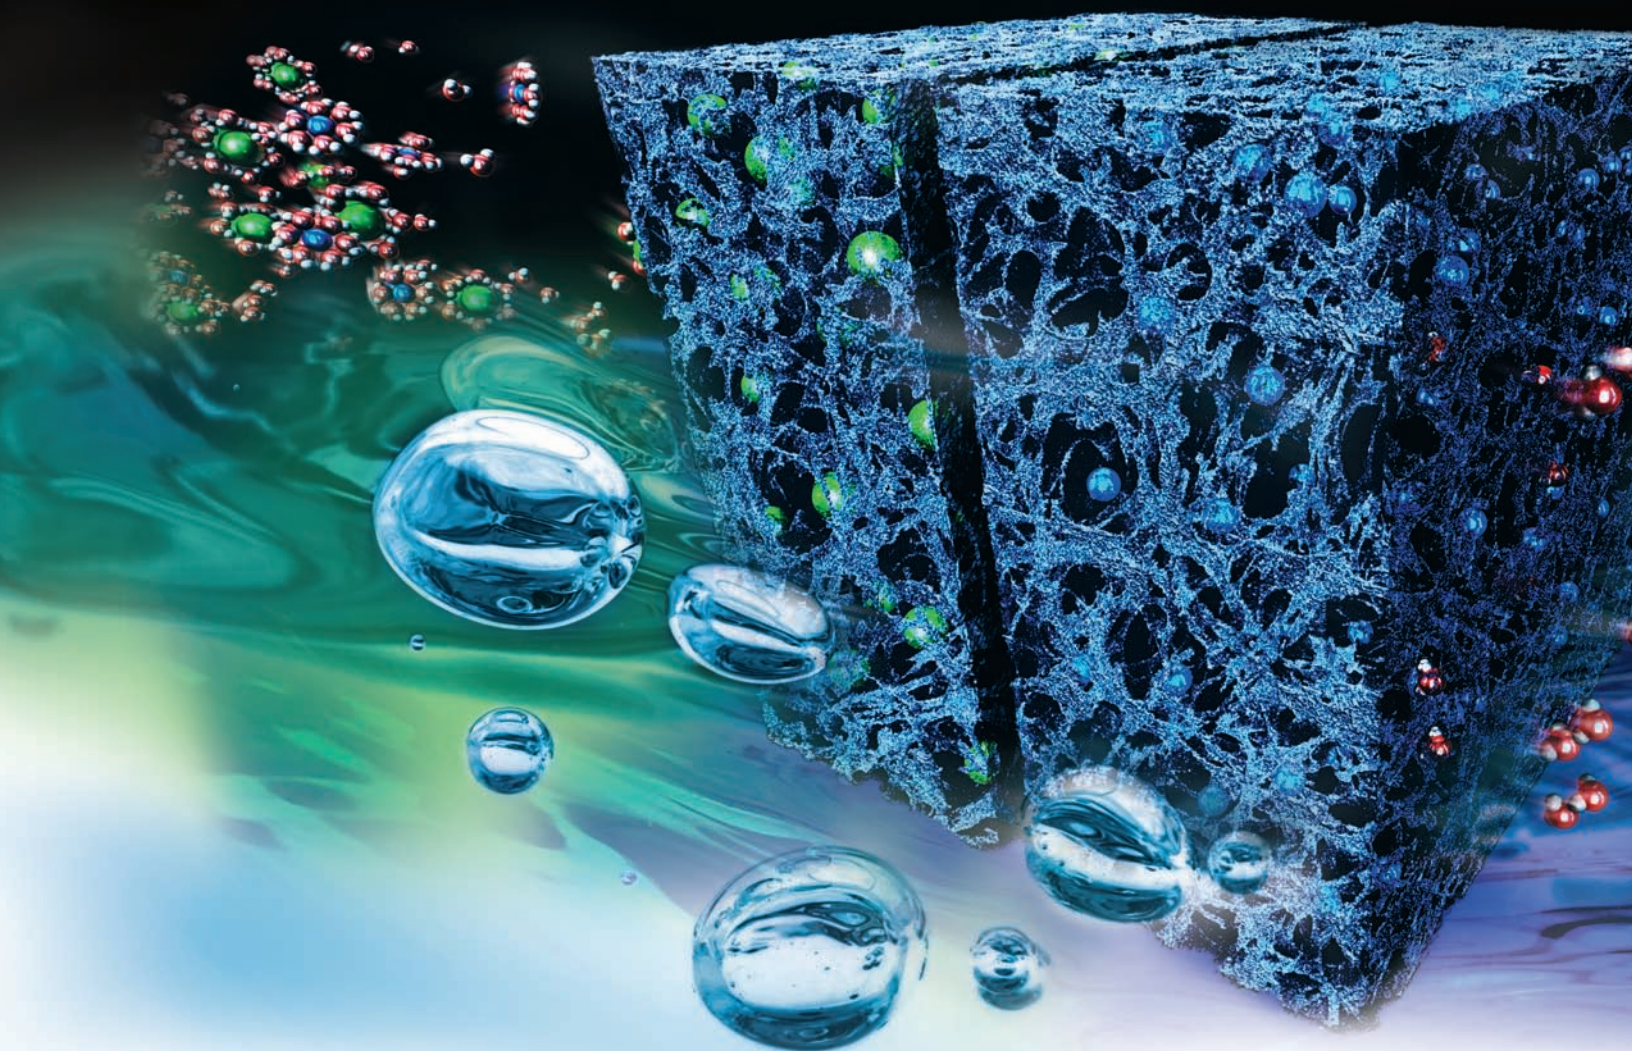
A Better Method for Desalinating Saltwater

I N "The Rime of the Ancient Mariner," Samuel Taylor Coleridge writes,

Water, water, every where,

nor any drop to drink

capturing perfectly the dilemma we face when looking to the ocean for drinking water. Oceans cover 70 percent of Earth's surface and make up 97.5 percent of the world's water, yet this water is of no use to the more than one billion people who don't have access to clean water.

As the population continues to increase, the demand for freshwater for use in industry and households is increasing. At the same time, freshwater supplies are decreasing as glaciers melt, rivers become polluted, and freshwater is simply wasted. Water scarcity is a national security issue, and new sources of freshwater are essential to ensuring peace and prosperity on a worldwide scale.

One solution to the problem is desalination. Many areas that do not have access to potable water do have access to salty water such as seawater, brackish surface water, or groundwater. If a cost-effective way to extract the freshwater could be found, many cities, states, and countries could solve their water problems.

Desalination is hardly a new idea, but current methods such as reverse osmosis (RO) and flash distillation require extensive infrastructure, are energy intensive, and are therefore expensive. RO requires energy to produce the pressure needed to push the water through the membranes. Flash distillation requires energy to heat the water for distillation processes.

Large Pores Let Water Flow Through

Physicist Michael Stadermann was originally investigating how the pore structure of carbon materials affects the performance of capacitors for energy storage, when chemist Ted Baumann showed him a new carbon aerogel with some interesting properties.

Flow-through electrode capacitive desalination (FTE-CD) uses a new hierarchical porous carbon material to create a device geometry in which the feed stream passes directly through the electrodes, resulting in significant improvements to the salt- removal rate. In this rendering, saltwater enters the carbon aerogel electrodes (blocks of porous material) on the left, leaving behind sodium (green dots) and chloride (blue dots) ions. The clean water exits on the right. (Rendering by Kwei-Yu Chu.)
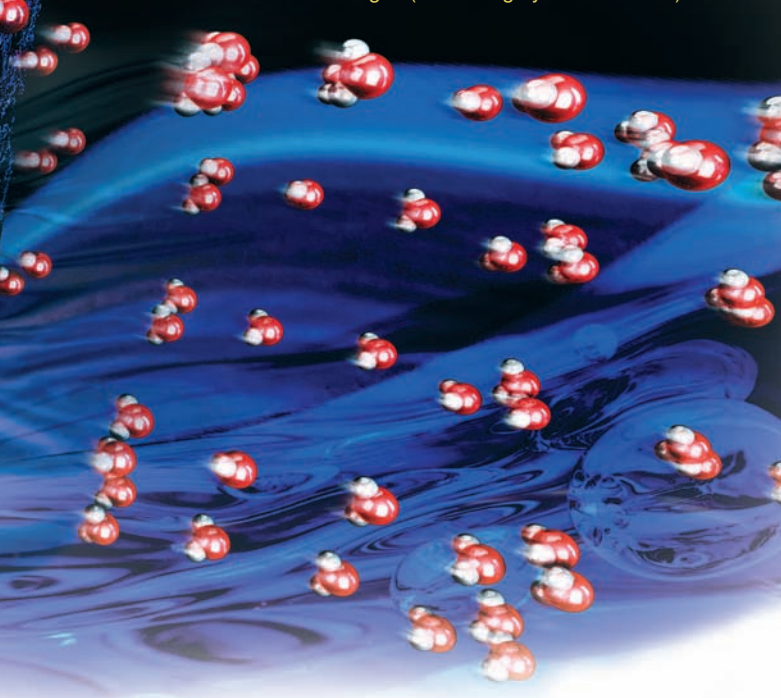
Aerogels, a special class of ultralow-density materials, have a complicated, cross-linked internal structure that gives them the highest internal surface area per gram of any known material.

The new carbon aerogel has the electrical properties of traditional aerogels but has much larger pores with diameters of micrometers versus nanometers. It also is mechanically robust and can be machined and fabricated into different sizes and shapes.

A key feature of this aerogel is that the walls of the big pores are themselves porous, with pores just nanometers in size. Hierarchical porosity allows water to easily flow through the large pores at low pressure, while the small pores give the material an enormous surface area. One gram of the aerogel contains up to 3,000 square meters of surface area, an area equal to more than 10 tennis courts.

This new aerogel did not quite meet the specifications for energy storage that Stadermann required for his original project, but it does look promising for capacitive water desalination. In capacitive desalination, saltwater flows through a capacitor with two electrically charged electrodes, and the positively and negatively charged sodium and chloride ions that make up the salt are captured on the surface of these electrodes. The aerogel is an ideal material for the electrodes because it is porous enough to let the water through and has a vast surface that can hold the ions.

Stadermann and Baumann, along with Lawrence Scholar Matthew Suss and his advisor at Stanford University, Juan Santiago, began exploring how to optimally use the material in capacitive desalination. Capacitive desalination using carbon aerogels was pioneered at Livermore in the 1990s. The first devices used aerogels with small pores, and the water flowed between the electrodes. The method required long desalination cycles and could only treat

0.5 micrometers
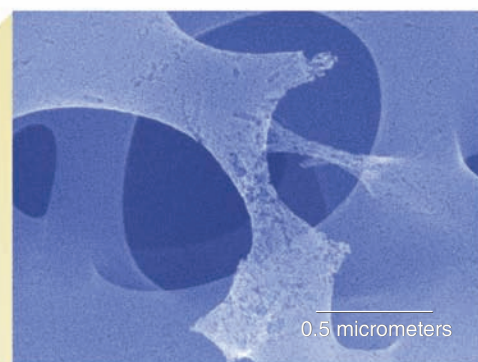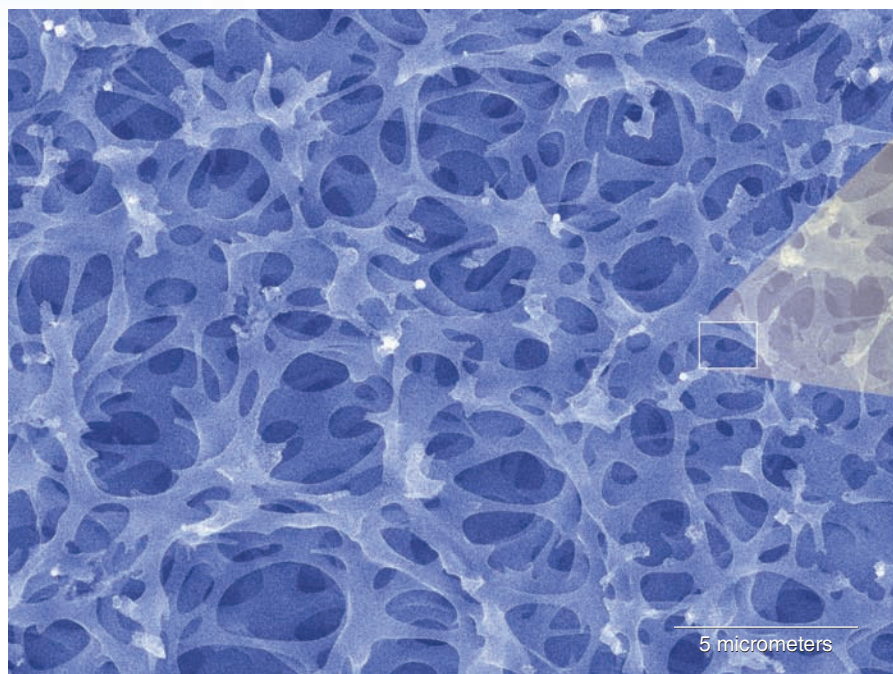
This micrograph of carbon aerogel shows that the pores in the FTE-CD module are large enough to let water through. (inset) The walls of the big pores are themselves porous. These smaller pores give the material its enormous surface area.

5 micrometers

moderately brackish water in a single charge. In another method, initially proposed in the 1970s, water flowed directly through the electrodes. That method failed because the electrode materials were unreliable.

Because the new carbon aerogel has such large pores, the researchers wondered if they could use monolithic aerogel materials with flow-through electrodes. Their idea was to build a capacitor with two hierarchical carbon-aerogel electrodes. The water would flow through the entire capacitor instead of through a small channel between the electrodes. Thus, the capacitor could potentially desalinate more water in a shorter period.

Developing a Prototype

Suss developed a computer model to explore the feasibility of the idea. When the model showed that significantly faster and more efficient desalination could be achieved by flowing water through the electrodes, he built a prototype flow-through electrode capacitive desalination (FTE-CD) module. (See the figure at right.) When a small potential (about 1 volt) is applied to the capacitor, the entire internal surface of the aerogel electrostatically attracts salt ions: sodium ions to the negative electrode and chloride ions to the positive electrode. As the pump pushes in more saltwater, the freshwater is flushed out, and more and more salt is captured in every electrode. Because the water now flows through the entirety of each electrode, all the desalinated water that was formerly trapped in the aerogel can be collected.

When the module's surface becomes saturated with salt, it can be cleaned by simply removing the voltage. The charge can be used again by transferring it to another module, thereby recovering the energy, and the ions can be released into the remaining water and rinsed out.

A single module can serve as one stage of desalination. As the water becomes less salty, it can be moved into a new module. Ocean water can be desalinated in eight

This prototype FTE-CD module, shown next to a glass of water for scale, provides one stage of desalination and can be coupled with more modules, depending on the level of desalination desired. The size and shape of the module could be made compatible with those in reverse osmosis plants and could conceivably be used in retooling such a plant.
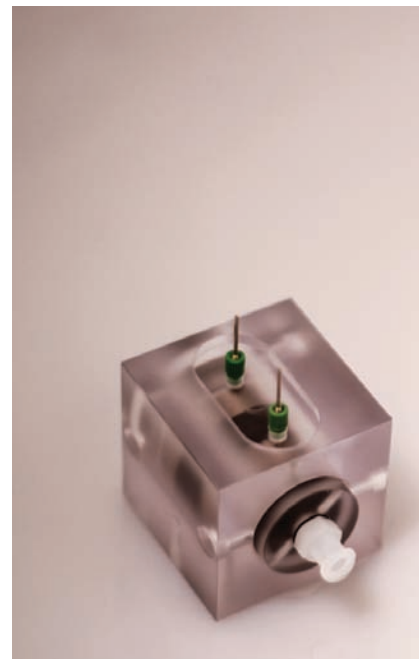
passes, and brackish water such as that from river mouths, saline aquifers, or wastewater can be desalinated in two or fewer passes. Experiments using FTE-CD modules showed that the device removes three times as much salt per charge and desalinates water 10 to 20 times faster than other capacitive devices.

Looking back to the start of the project, Stadermann says, "Desalination is not the kind of work the Lab typically does, but we're good at it because we have the infrastructure and a history in this field. With drought throughout the Midwest, our research is in the national interest and within the scope of a national laboratory."

Scalable Desalination

Desalination plants require so much infrastructure and energy that they are only cost effective at a large scale, but a great need

exists for smaller operations. FTE-CD modules are quite scalable and can be used individually or in stages. They require very little energy to charge the capacitors and very little pressure to pump water through the modules. The energy is simply stored in the capacitor, so it can be easily recovered or transferred to other modules. Furthermore, because desalination is done in stages, the salinity of the resulting water can be precisely controlled.

The team is continuing work on improvements to the FTE-CD modules. Solar power could potentially be used as the energy source. Also, researchers could develop a gravity-fed module to force water through the capacitors.

A wide range of possible applications exist for the FTE-CD modules including for industries that use partially desalinated water for their processes. The military might provide soldiers

with small desalination units for their packs, or slightly larger desalination systems could be carried in military transport vehicles. Other possible applications include using the modules to filter water in disaster areas or to soften household water supplies.

Some industries require large amounts of water and are looking at ways to switch from freshwater to brackish water or ocean water. For example, FTE-CD could be used for hydraulic fracturing, or "fracking" as it is commonly called, where large amounts of water and other fluids are forced through shale or other rock formations to release petroleum or natural gas. Salinity of water used for fracking can rise to more than three times the level of seawater, but the FTE-CD modules could keep the salinity at a consistent and desirable concentration.

FTE-CD modules could be built to be compatible with membrane modules in RO systems, without the infrastructure needed to support RO. As the cost for FTE-CD decreases, RO modules could potentially be replaced with FTE-CD modules. Cost-effective, modular, and scalable desalination would solve one of the world's most pressing problems, and FTE-CD is one tool that will bring us closer to the goal of freshwater for all.

—*Karen Rath*
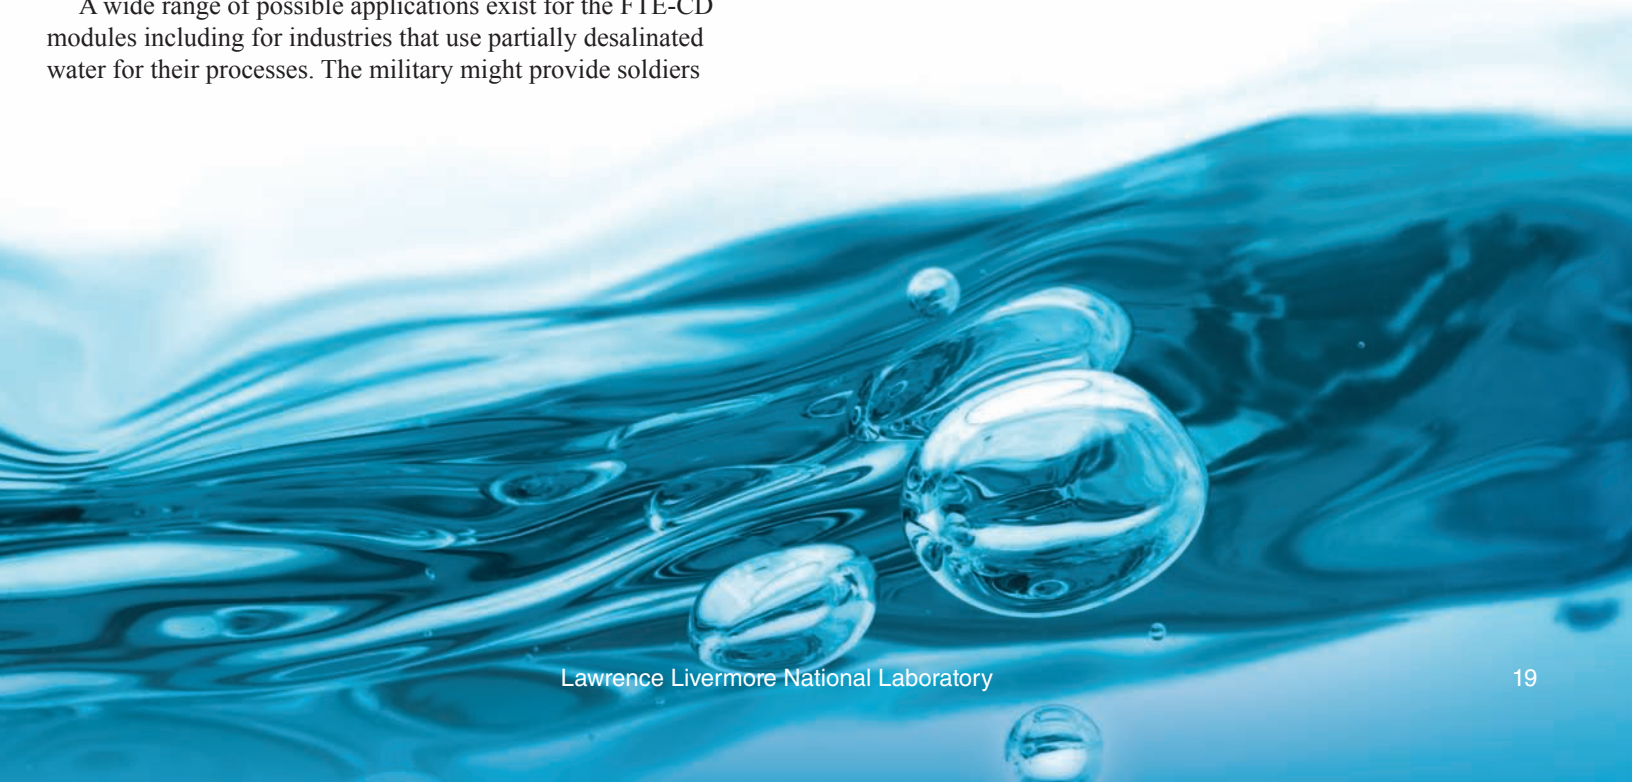
Lawrence Livermore, Ball Aerospace and Technologies Corporation, and NeXolve Corporation are preparing diffractive membrane optics for the Membrane Optic Imager Real-Time Exploitation (MOIRE) project.

Developing Lightweight Optics for Space

A general rule for telescopes is that bigger is better. Improving our ability to view dim and distant objects, or see finer details on nearer objects, requires a telescope with better resolution and more light-gathering power, which generally means a bigger aperture. With no atmospheric distortion to affect the resolution of large-aperture space-based telescopes, the sky is the limit. Scientists are only restricted by the practicalities of device fabrication and transportation. Because the materials must be transported to space at great cost and effort, strict limits are imposed on the size and weight of the telescope optics, structures, and other components.

When the National Aeronautics and Space Administration's next-generation James Webb Space Telescope is launched in 2018, it will have a segmented aperture with a diameter of 6.5 meters—seven times the light-collecting area of the Hubble Space Telescope. As early as 1996, though, Lawrence Livermore physicist Rod Hyde was working on a design that would take advantage of the Laboratory's expertise in optics design and fabrication to put a much larger telescope—with a 20-meter-diameter aperture—into space. (See *S&TR*, March 2003, pp. 12–18.) To make this possible, Hyde and his colleagues seemingly took a step backward in telescope technology.

High-performance telescopes such as Hubble and the future Webb typically are reflective telescopes, meaning their primary optical element is a curved mirror. In contrast, a transmissive telescope consists of a long tube with a lens (rather than a mirror) at one end to collect and focus light and an eyepiece at the other through which to look. Optics for transmissive telescopes are significantly less sensitive to surface imperfections than mirrors, but standard transmissive lenses are far too heavy for transporting to space. Pairing transmissive telescope technology with diffractive optics—a notable manufacturing strength for the Laboratory—would allow researchers to lighten the launch load.

Transmissive diffractive optics, also called Fresnel lenses, are patterned on one surface, with features often too small to be visible to the naked eye. These tiny, patterned features are tailored to bend light of particular wavelengths. (See *S&TR*, September 1995, pp. 24–33.) The height of these surface features is on the order of the wavelength of light. Because light is focused by the surface features and not by refraction of bulk glass, diffractive lenses can be made far thinner and lighter than standard lenses. The disadvantage of diffractive optics–based telescopes is their extremely limited bandwidth. A diffractive light collector focuses the various wavelengths of light at different points in space, severely limiting the light signal at the focal plane. Fortunately, incorporating a second, inverse diffractive lens with a pattern carefully matched to the light collector corrects these focusing aberrations and provides enough photons at the focal plane for imaging.

Fitting a 20-meter-diameter sheet of glass, no matter how thin, into a rocket was infeasible, so Laboratory researchers designed a segmented lens that could be neatly folded for transport and then opened on arrival in orbit. With funding from the Laboratory Directed Research and Development Program and the Defense Advanced Research Projects Agency (DARPA), they built a 5-meter primary lens called Eyeglass to test how well the components folded, unfolded, and assembled. Livermore's diffractive optics group, led by optics engineer Jerry Britten, fabricated the precision meter-scale diffractive optics segments that the design required. The prototype, completed in 2002, generated interest. However, perceived risk and budget constraints prevented a large-scale follow-on, and the project was dormant for nearly a decade.

A New, More Flexible Eyeglass

In 2010, DARPA held a design competition to demonstrate technology for a video imaging system in geosynchronous Earth orbit. The DARPA project, Membrane Optic Imager Real-Time Exploitation (MOIRE), called for a large-aperture transmissive diffractive optical space telescope that could image an area greater than 100 square kilometers with a video update rate of at least one frame per second. Ball Aerospace and Technologies Corporation, together with the Laboratory and NeXolve Corporation, was selected by DARPA to prepare a design for MOIRE that would demonstrate the manufacturability of a light-collecting lens, the structures to hold the optics tight and flat, and the additional optical elements needed to turn a diffraction-based optic into a wide-bandwidth imaging device. Lawrence Livermore's area of responsibility is diffractive optics fabrication.

In many ways, MOIRE is a direct successor to Eyeglass. Perhaps the most significant differences between the two efforts are in the optical material and manufacturing techniques. Glass
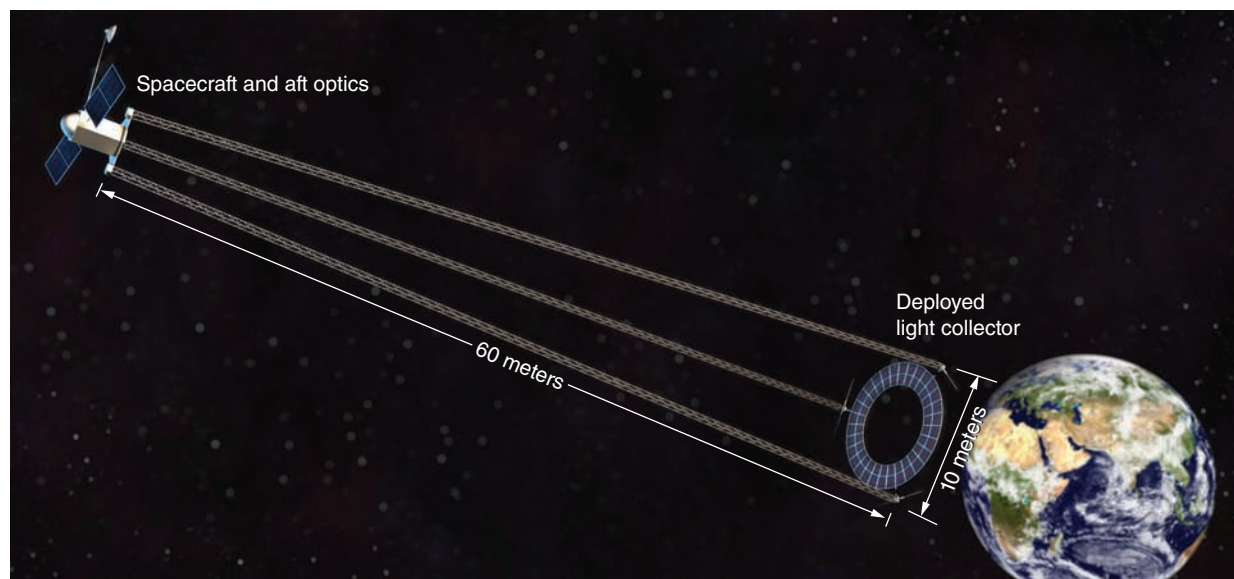
Upon arrival in orbit, the MOIRE spacecraft would deploy solar arrays, unfold the telescope tube, and then unfurl a micrometer-thin, diffractive optic membrane to form a massive segmented light collection lens (not to scale). (Rendering by Ryan Chen.)

Spacecraft and aft optics

60 meters

Deployed light collector

10 meters

offers excellent optical properties, but concerns about the launch survivability of large, thin glass plates, along with a desire to minimize weight, drove researchers to select a lightweight, packable, flexible membrane material for MOIRE optics. The NeXolve-created polyimide membrane is virtually weightless and a mere 20 micrometers thick, with a near-zero coefficient of thermal expansion. Significant temperature fluctuations in space would cause materials with high coefficients of thermal expansion to stretch significantly and thereby distort images.

A transmissive telescope typically requires vastly greater distances for focusing than reflective telescopes. Eyeglass designers initially intended to put the light-collecting lens into orbit several kilometers from the eyepiece and other electronic components. However, launching two spacecrafts and keeping them precisely aligned with one another would have introduced additional complexity and risk.

A change to the manufacturing process allowed the MOIRE team to shrink the length of the telescope. Says Britten, "With Eyeglass, we used a wet-etching process and were limited to approximately 10-micrometer-wide optical surface features. With an ion mill, we can now create high-fidelity, submicrometer features on meter-scale optics." Using equipment and techniques developed to create large diffraction gratings for the National Ignition Facility and its Advanced Radiographic Capability (see *S&TR*, December 2011, pp. 12–15), the diffractive optics group was able to shorten the required focal length for a 10-meter-diameter telescope to about 60 meters. Furthermore, the MOIRE team's design calls for a single-launch payload, significantly reducing both risk and transport costs.

Delivering Precision Optics

The word *moiré* refers to an interference pattern, such as one used during the fabrication of diffractive optics. It is an appropriate name for a project where diffractive optics are central to success. The process for creating each of MOIRE's diffractive optical segments is complex and takes several weeks because of the sheer size of the optics.

Pattern transfer begins with fabricating a master pattern in chrome on glass. A membrane surface is then covered uniformly with a photosensitive coating, and the master pattern

Lightweight, flexible, packable optics and other telescope components are essential for the launch to geosynchronous orbit. The MOIRE primary lens at the top of the rocket is designed to fold similar to an umbrella for launch and transport.
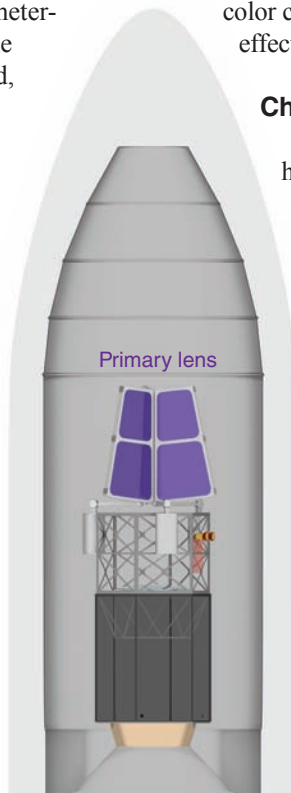
Primary lens

is placed similar to a stencil over the membrane. When the photosensitive coating is exposed, a pattern is created in the coating. Finally, the pattern is transferred permanently into the membrane using an ion-beam etcher, and the coating is stripped off. The result is a membrane surface that is patterned with precisely spaced grooves of uniform depth.

Working with Ball and NeXolve, Britten's group has manufactured optics for two phases of the multiyear MOIRE project. Membranes were supplied by NeXolve and printed and etched at the Laboratory. For the first phase of delivery, completed in mid-2011, the diffractive optics experts created a sample 80-centimeter-diameter off-axis circular diffractive optical element with 4-micrometer-wide features that will serve as one segment in the giant light-collecting lens. Britten's group and the team at Ball tested the accuracy and precision of the master pattern and pattern transfer with promising results. Diffraction efficiency was measured at better than 30 percent, approaching the approximate 35-percent theoretical maximum efficiency for the design. This performance established confidence in the MOIRE team's ability to fabricate precision large-scale membrane optics. For the Laboratory's contribution to Phase 2, to be completed in early 2013, team members have significantly scaled up optics production. They will deliver six 80-centimeter-diameter trapezoidal optics that comprise one-eighth of the ring-shaped primary lens, several spare optics, and the smaller optical components at the "back end" of the telescope—diffractive color correctors that turn a transmissive telescope into an effective imaging device.

Challenges and Enhancements

While membrane optics are lightweight and flexible, handling and etching a pattern in the surface of a material analogous to household plastic wrap has posed significant challenges throughout the project. The material may undergo almost no thermal expansion, but it does have a large moisture expansion coefficient, meaning that it expands and contracts substantially as humidity levels change. This behavior can distort the diffraction pattern or shift the built-in alignment fiducials that match up the lens segments.

Fortunately, an effort is under way to quantify and anticipate membrane behavior. The group has noted, for instance, that the photoresist-coating process consistently causes the material to stretch and sag. Understanding this effect will allow the diffractive optics group to compensate for it in the master pattern or correct for it through adaptive optics—in effect, to create a corrective lens for the telescope. The researchers are also experimenting with optimal methods of mounting the membrane and maintaining
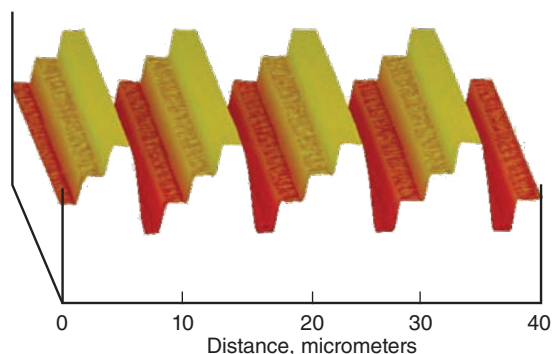
(left) Three-dimensional atomic force microscopy (AFM) is used to image the surface of a four-level glass master optic. The vertical scale (690 nanometers high) and sidewall slope is exaggerated by software and the AFM tip width, respectively. (right) With a four-level etched membrane pattern, light is diffracted in one direction, producing a single offset image of optics engineer Jerry Britten.

tension to improve smoothness and radial uniformity throughout the manufacturing process.

The diffractive optics group, in conjunction with NeXolve, has also engaged in a more ambitious research effort. The researchers want to improve imaging quality by changing the diffraction feature shape of the primary light collector and thereby increase light transmission and reduce background noise. This risk-reduction demonstration, performed in addition to the Phase 1 and 2 optics delivery, has entailed switching from a two-level pattern, which splits most of the light into two paths, to a four-level stepped pattern, which concentrates the light into the imaging path while simultaneously reducing stray light, thus increasing the signal-to-noise ratio of the telescope. Patterning multiple levels of features requires two rounds of masking and etching with submicrometer accuracy, a process deemed virtually impossible to perform on an unstable membrane surface. Instead, the partners decided to fabricate a membrane with the complex pattern built into it from the start.

Preliminary testing with NeXolve confirmed the feasibility of the idea, and the equipment used to write the chrome mask pattern was upgraded so it could generate the more complex features. The Laboratory's diffractive optics group then proceeded to fabricate a glass master optic patterned with the four-level stepped structure. Using the master, NeXolve created replica membranes with the negative of this pattern formed in the surface during manufacturing. The first round of membranes has demonstrated impressive performance and confirmed that the Livermore and NeXolve team can produce sophisticated multilevel patterns at the requisite scale. Diffraction efficiency has risen from 35 to 55 percent, and background noise from stray light has been cut from 30 percent to less than 1 percent.

MOIRE Takes Flight

Phase 2 of MOIRE culminates in April 2013 with an integrated demonstration of a partial 5-meter-diameter ground-based telescope. If this demonstration achieves its goals, launch and deployment of a 10-meter-diameter telescope with more efficient diffractive optics in geosynchronous orbit could follow. That telescope would weigh an astonishing seven times less than a reflective telescope of the same size. While more work must be done before MOIRE reaches the skies, Lawrence Livermore's optics fabrication expertise has helped to successfully demonstrate that a large diffractive membrane optic can be used as a building block for developing a lightweight, segmented-aperture telescope. Laboratory investment in optics research and manufacturing infrastructure made as part of National Ignition Facility construction has confirmed its value yet again with the MOIRE effort.

If MOIRE is successful, other lightweight transmissive space telescopes are likely to follow, with scientific, defense, or even communications missions. Although MOIRE is a telescope that will look down at Earth rather than at distant stars, the principles are the same for both. The eventual goal for MOIRE, like Eyeglass before it, is a space telescope with a 20-meter-diameter lens. With a space telescope of that size, astronomers could observe weather patterns on Saturn with 100-kilometer resolution or planetary nebula 600 light years away with 1-astronomical-unit resolution. For a project that not long ago looked to be relegated to the annals of astronomical engineering history, the future is looking bright.

—*Rose Hansen*

For further information contact Jerry Britten (925) 423-7653 (britten1@llnl.gov).

Patents and Awards

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.

S&TR January/February 2013

Patents

Nanoporous Carbon Actuator and Methods of Use Thereof

Juergen Biener, Theodore F. Baumann, Lihua Shao, Joerg Weissmueller
U.S. Patent 8,231,770 B2
July 31, 2012
An electrochemically drivable actuator includes a nanoporous carbon-aerogel composition that exhibits charge-induced reversible strain when it is wetted by an electrolyte and a voltage is applied. When the nanoporous carbon aerogel is wetted by an electrolyte, the actuator has a mechanism that applies a charge to induce reversible strain in the aerogel. A method is described for electrochemically actuating an object, which includes causing charge-induced reversible strain in a nanoporous carbon-aerogel composition wetted with an electrolyte to actuate the object by the strain.

Multiplex Detection of Respiratory Pathogens

Mary McBride, Thomas Slezak, James M. Birch
U.S. Patent 8,232,058 B2
July 31, 2012
These kits can be used to detect respiratory pathogens such as influenza A (including subtyping capability for H1, H3, H5, and H7 subtypes), influenza B, parainfluenza (type 2), respiratory syncytial virus, and adenovirus. Genomic sequence information from the respiratory pathogens was analyzed to identify signature sequences, such as polynucleotide sequences, and determine whether a pathogen is present in a sample. Primer and probe sets are optimized for use in a polymerase-chain-reactor–based, multiplexed Luminex assay to successfully identify the presence or absence of pathogens in a sample.

Method and System for Modulation of Gain Suppression in High Average Power Laser Systems

Andrew James Bayramian
U.S. Patent 8,233,511 B2
July 31, 2012
A high-average-power laser system with modulated gain suppression has input and output apertures associated with a first laser beam extraction path. The system also includes a pinhole-creation laser with its optical output directed along a pinhole creation path. An absorbing material is positioned along both paths. The system has a mechanism that translates the absorbing material in a direction crossing the first laser beam extraction path and a controller that modulates the second laser beam.

Magnetic Separation of Devitrified Particles from Corrosion-Resistant Iron-Based Amorphous Metal Powders

Phillip D. Hailey, Sumner D. Day, Joseph C. Farmer, Nancy Yang, Thomas M. Devine, Jr., Larry Kaufman
U.S. Patent 8,245,661 B2
August 21, 2012
This system for coating a surface requires a source of iron-based amorphous metal that includes devitrified ferrite. The system sprays the metal on a surface, separating at least a portion of the devitrified ferrite from the spray before the spray reaches the surface.

Awards

The Laboratory's **Sequoia supercomputer** was selected for a **2012 Breakthrough Award** by *Popular Mechanics*. Sequoia, an IBM BlueGene/Q system, holds the No. 2 ranking on the industry-standard Top500 list of the world's fastest high-performance computing (HPC) systems. The annual Breakthrough Awards recognize the 10 most "world changing" innovations in fields ranging from computing and engineering to medicine, space exploration, and automotive design. The awards are given in two categories: innovators, whose inventions will make the world smarter, safer, and more efficient in the years to come; and products, which are setting benchmarks in design and engineering.

Clocking 16.3 petaflops (quadrillion floating-point operations per second), Sequoia has also been ranked No. 1 on the Green500 list as the most energy-efficient HPC system and No. 1 on the Graph 500 list thanks to its ability to solve big data problems. Sequoia is running unclassified applications as a part of the testing required before it transitions to classified stockpile stewardship computing for the Advanced Simulation and Computing Program in early 2013. Science being explored by Lawrence Livermore researchers includes high-energy-density plasmas and the electronic structure of heavy metals. Sequoia also has demonstrated its amazing scalability with a three-dimensional simulation of the human heart's electrophysiology. (See *S&TR*, September 2012, pp. 22–25.)

Isom Harrison, director of the Laboratory's Library since 1991, has received the **2012 Lifetime Service Award** from the **National Organization for the Professional Advancement of Black Chemists and Chemical Engineers** (NOBCChE) for his work as the organization's Western region chairperson. The award is bestowed by the NOBCChE Board of Directors on individuals whose contributions have been exemplary and served to enhance the day-to-day functioning of NOBCChE consistently over a long time. Harrison has a master's degree in organic chemistry and has been active in NOBCChE for the past 20 years, serving at the regional level with a focus on community outreach and education. "It is very gratifying to be recognized for the work you do, especially when the recognition is coming from your peers, most of whom are just as deserving," says Harrison.

Abstract

Dealing with Data Overload in the Scientific Realm

The world is awash in data, nowhere more so than in the scientific realm. Lawrence Livermore scientists are developing innovative schemes for dealing with the data overload that plagues fields as diverse as bioinformatics, cybersurveillance, and climate science. In the bioinformatics area, Laboratory computer scientists have developed GeneSV, which helps researchers assess whether a genetic sequence of a mutated virus will lead to a viable virus or not. In another effort, a robust platform that combines particle filtering and time-evolving computer models helps scientists perform real-time analysis of data streams, such as those encountered by spam filters, where data must be sampled "on the fly" and cannot be stored for later analysis. Finally, the Earth System Grid Federation uses distributed storage and software to federate petabytes of data from climate scientists around the world, harvesting and integrating data in a manner transparent to the user.

Contact: Dean Williams (925) 423-0145 (williams13@llnl.gov).

Engineered Munitions

Laboratory engineers combine advanced simulations, improved material compositions, and new fabrication techniques to develop munitions that minimize collateral damage.

Also in March

• Livermore researchers are making important progress toward achieving one of science's grand challenges—igniting fusion fuel in the laboratory.

• New mechanisms for regulating bone growth may lead to improved methods for strengthening or healing bones.

• Chaotic plasmas give birth to orderly electromagnetic fields.

Coming Next Issue

www.llnl.gov