Lawrence Livermore National Laboratory

Science & Technology Review

March 2017

PREPARING TO SCALE SIERRA

Also in this issue:

A Material to Defeat Weapons

Advanced Laser Alignment

Studying Failure in Granular Materials

About the Cover

Lawrence Livermore's next Advanced Simulation and Computing (ASC) Program supercomputer, Sierra, represents a break with the design of many past ASC systems. To perform efficiently on Sierra, Livermore's complex science applications must be significantly modified in many cases. As the article beginning on p. 4 describes, the Laboratory established a Center of Excellence to help coordinate and accelerate such preparatory efforts. Through the center, selected applications—those deemed critical for Laboratory missions and representative of the topical breadth of research at the institution—are receiving focused attention by application developers, hardware and systems software experts from vendor partners, and in-house experts in code optimization for next-generation systems. The cover image is an artist's depiction of a mountain embedded with "early-access" Sierra.

Cover design: Amy 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) 423-3893. Our e-mail address is str-mail@llnl.gov. *S&TR* is available on the Web at str.llnl.gov.

© 2017. Lawrence Livermore National Security, LLC. All rights reserved. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. To request permission to use any material contained in this document, please submit your request in writing to Public Affairs Office, Lawrence Livermore National Laboratory, Mail Stop L-3, P.O. Box 808, Livermore, California 94551, or to our e-mail address str-mail@llnl.gov.

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

Prepared by LLNL under contract
DE-AC52-07NA27344

Science & Technology REVIEW

March 2017

Lawrence
Livermore
National
Laboratory

Contents

Feature

3 **Platforms, Codes, and Facilities Form a Three-Pronged Supercomputing Strategy**
Commentary by Michel McCoy

4 **A Center of Excellence Prepares for Sierra**
Application developers are partnering with supercomputing experts and hardware vendors to ready a stable of mission-critical applications for the next generation of computing architectures.

Research Highlights

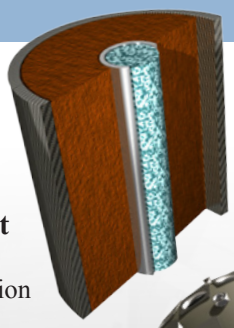

12 **Agent Defeat Efforts Strike Gold**
A highly effective payload material can destroy chemical and biological weapons facilities.

16 **Mighty ATLAS Supports Precise Alignment**
A new, extremely accurate laser tracker speeds alignment of critical diagnostics for National Ignition Facility experiments.

20 **Forcing Failure in Granular Materials**
Three-dimensional measurements of granular media under extreme stress reveal new information about material properties and energy dissipation.

Departments

2 The Laboratory in the News

24 Patents and Awards

25 Abstract

S&TR Staff

Scientific Editor
Ramona Vogt

Managing Editor
Ken Chinn

Publication Editor
Paul Kotta

Writers
Holly Auten, Rose Hansen, and Arnie Heller

Art Director
Mary Gines

Designers
Amy Henke and Santos Sosa Leon

Proofreader
Caryn Meissner

***S&TR* Online**
Rose Hansen, Lanie L. Helms, and Denise Kellom

Print Coordinator
Diana Horne

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-52000-17-3
Distribution Category UC-99
March 2017

Possible Bacterial Path of Antarctic Mercury

Livermore scientists Michael Thelen and Adam Zemla performed protein sequence analysis and structural modeling to predict key proteins involved in mercury methylation as part of research with the U.S. Geological Survey and the University of Melbourne to determine how highly toxic methyl mercury—the product of methylation—enters the Antarctic Ocean and accumulates in the marine food web there. The team's findings were published in the August 1, 2016, online edition of *Nature Microbiology* and highlighted in the journal's October 1, 2016, "News and Views" section.

Mercury from volcanoes and human activity such as burning fossil fuels circulates in the atmosphere, then deposits onto sea ice. Scientists are concerned that the fatty tissues of nearby fish could accumulate enough methyl mercury to pose a public health concern.

Thelen explains, "We examined the sequence data from DNA samples collected in sea ice and other Southern Ocean environments. A candidate sequence for a mercury methylation enzyme was found in a bacterial strain of the genus *Nitrospina*. Adam then used computational tools developed at the Laboratory to predict the structure and show how the active site of the enzyme would react with mercury." The researchers' analysis of samples collected in the second Sea Ice Physics and Ecosystem Experiment also found seawater under seasonal ice to contain significantly higher mercury concentrations than open-ocean water. Thelen adds, "Further research could physically confirm the proteins responsible for mercury methylation in *Nitrospina* and perhaps could help us find a way to inhibit or to reverse the reaction to reduce methyl mercury poisoning."

Contact: Michael Thelen (925) 422-6547 (thelen1@llnl.gov).

Game-Changing Forensics Technique Uses Hair Proteins

In an important breakthrough for forensic science, researchers have developed the first-ever biological identification method that unlocks the genetic information encoded in proteins of human hair. The groundbreaking technique could provide a second science-based, statistically validated way, in addition to DNA profiling, to identify people and link individuals to forensic evidence. The work by Lawrence Livermore, seven universities, and a Utah-based startup was announced in a paper published September 7, 2016, in the online journal *PLOS ONE*.

"We are in a very similar place with protein-based identification to where DNA profiling was during the early days of its development," says chemist Brad Hart, director of Livermore's Forensic Science Center and co-author of the paper. Proteins are chemically far more robust than nuclear DNA, which degrades quickly—a significant weakness in DNA-based forensics. In contrast, the team identified viable protein markers in human hair from skeletal remains in London cemeteries dating back 150 to 250 years. Hair samples from these and other subjects revealed each person to have unique protein markers that provide a forensically useful "fingerprint."

Proteins, including those in hair, are coded by the DNA that is unique to each person. The markers used by the scientists are protein variants resulting from DNA mutations known as single amino acid polymorphisms. The concept of basing identification on this facet of proteins was discovered by then Utah Valley University biochemist Glendon Parker (shown in the photo), who later came to Lawrence Livermore, in 2013. The researchers are seeking to establish a set of 90 to 100 protein markers that would be sufficient to distinguish an individual among the world's population using a single hair. Hart states, "This method will be a game-changer for forensics."

Contact: Brad Hart (925) 423-7374 (hart14@llnl.gov).

Explanation for Iron Springing Back to Shape

A team of Lawrence Livermore physicists has performed a series of calculations shedding light on an unexpected way that iron transforms under dynamic compression. The team describes its first-principle calculations in a paper published in the August 19, 2016, issue of *Physical Review Letters*.

More than a decade ago, pioneering shock experiments revealed that single-crystal iron transforms under compression but then largely springs back to its original, intact lattice—a startling behavior that had defied understanding. "This reversible transformation is reminiscent of what happens in a shape-memory alloy," explains Livermore materials scientist Michael Surh, lead author of the paper. Calculations by Surh and his colleagues used a carefully parameterized model of iron magnetic fluctuations that included the effect of the iron nuclei being moved as the material is squeezed. Previous work had neglected the role of magnetism.

In addition, a separate Livermore team's shock-compression experiments showed that this transformation can take much longer than previously predicted. Surh explains, "It turns out that the different behavior—fast versus slow, reversible versus irreversible—seen in iron depends on its changing magnetism." The researchers describe the role played by "magnetic frustration"—the competition between the tendencies of electron spins to be either aligned or misaligned with those of their neighbors. The team states that future simulations must also include magnetic effects to arrive at a fundamental understanding of the intriguing phenomenon.

Contact: Michael Surh (925) 422-2087 (surh1@llnl.gov).

Commentary by Michel McCoy

Platforms, Codes, and Facilities Form a Three-Pronged Supercomputing Strategy

THE feature article, "A Center of Excellence Prepares for Sierra," beginning on p. 4, details the initial key steps taken in executing Lawrence Livermore's high-performance computing (HPC) strategy for the next decade, including code preparation for the soon-to-arrive Sierra supercomputing system and its institutional companion. Also described is the paradigm shift underway in system architecture, as we move from a supercomputing world in which floating-point operations (flops) are expensive to achieve, to one in which bandwidth and memory are the most precious commodities. The shift may seem a small matter but has driven substantial revisions to our codes and, in some cases, compelled us to rethink code designs entirely.

Given the Laboratory's leadership role in the world of supercomputing and the mission-critical work we perform with our HPC capabilities, we necessarily take a broad view that goes beyond siting a single system and preparing codes to run efficiently on it. In fact, if we are to sustain our competence in computing and continue to lead in supercomputing, we must strategically focus on three interrelated areas over the coming decade: (1) platforms and the direction that vendors' business models are taking; (2) codes that effectively use these platforms, possibly in novel ways; and (3) facilities that can effectively operate very large computers.

As for platforms, the Department of Energy's National Nuclear Security Administration (NNSA) and Office of Science are both investing in developing exascale technology, including systems software, hardware components, codes, and engineering technology. The Exascale Computing Program's goal currently is to build the first exascale system by 2021, while NNSA's Advanced Simulation and Computing (ASC) Program aims to site production-ready exascale platforms soon thereafter. By 2023, an ASC exascale system should be in production at Lawrence Livermore. Many people at Livermore are already involved in exascale work, in both research and development and leadership capacities.

At the same time, hardware vendors are moving strongly in the direction of artificial intelligence, deep learning, and other brain-emulating cognitive processes. One could even say these will be front and center in vendors' business plans, with simulation being secondary. We will therefore need to invest cleverly so that such capabilities can be harnessed to support our needs. Sierra will feature elements likely to play a role in deep-learning computing, including use of graphic processing units and substantial amounts of nonvolatile memory. In this sense, we can view Sierra as an on-ramp, or training ground, for what is coming soon.

As for the second strategic element, the ASC Program is investing heavily in next-generation codes, in particular developing a next-generation nuclear weapons performance code. The Center of Excellence (COE) for Sierra and its institutional counterpart are playing an important role in code preparation for Sierra and beyond. The creation of an institutional COE for non-ASC codes was driven by the realization that the Laboratory cannot be healthy as an institution unless all its scientists have access to a similarly advanced computational environment. What must come soon is aggressive investment in cognitive simulation, that is, simulation assisted by cognitive processes. One such scenario is a cognitive process monitoring a simulation and redirecting its execution when it senses trouble or opportunity. Some intriguing early work already conducted at Livermore indicates that such techniques could potentially speed up some computational studies by an order of magnitude or more.

The third element of Livermore's supercomputing strategy is facilities. Computers are becoming more power-hungry and storage-heavy. The Laboratory is currently exploring a variety of strategies to create an efficient, sustainable facility that could support multiple exascale platforms simultaneously. The most ambitious and leading candidate is a facility—achievable by 2020—that would provide Livermore computers up to 80 megawatts and sufficient cooling capacity. Last year we built a new facility—Building 654—that is an engineer's dream of efficiency (if not an architect's vision of beauty). This new facility will serve our needs for many years and can be expanded in modular fashion for additional space.

In short, supercomputing at Livermore remains a vital and rapidly evolving core component of overall institutional strategy. We continue to take a long view. Livermore already has a vision of what supercomputing could look like in 10 years. Now it is time to execute that vision.

■ Michel McCoy is director of the Weapon Simulation and Computing Program.

Application developers and their industry partners are working to achieve both performance and cross-platform portability as they ready science applications for the arrival of Livermore's next flagship supercomputer.

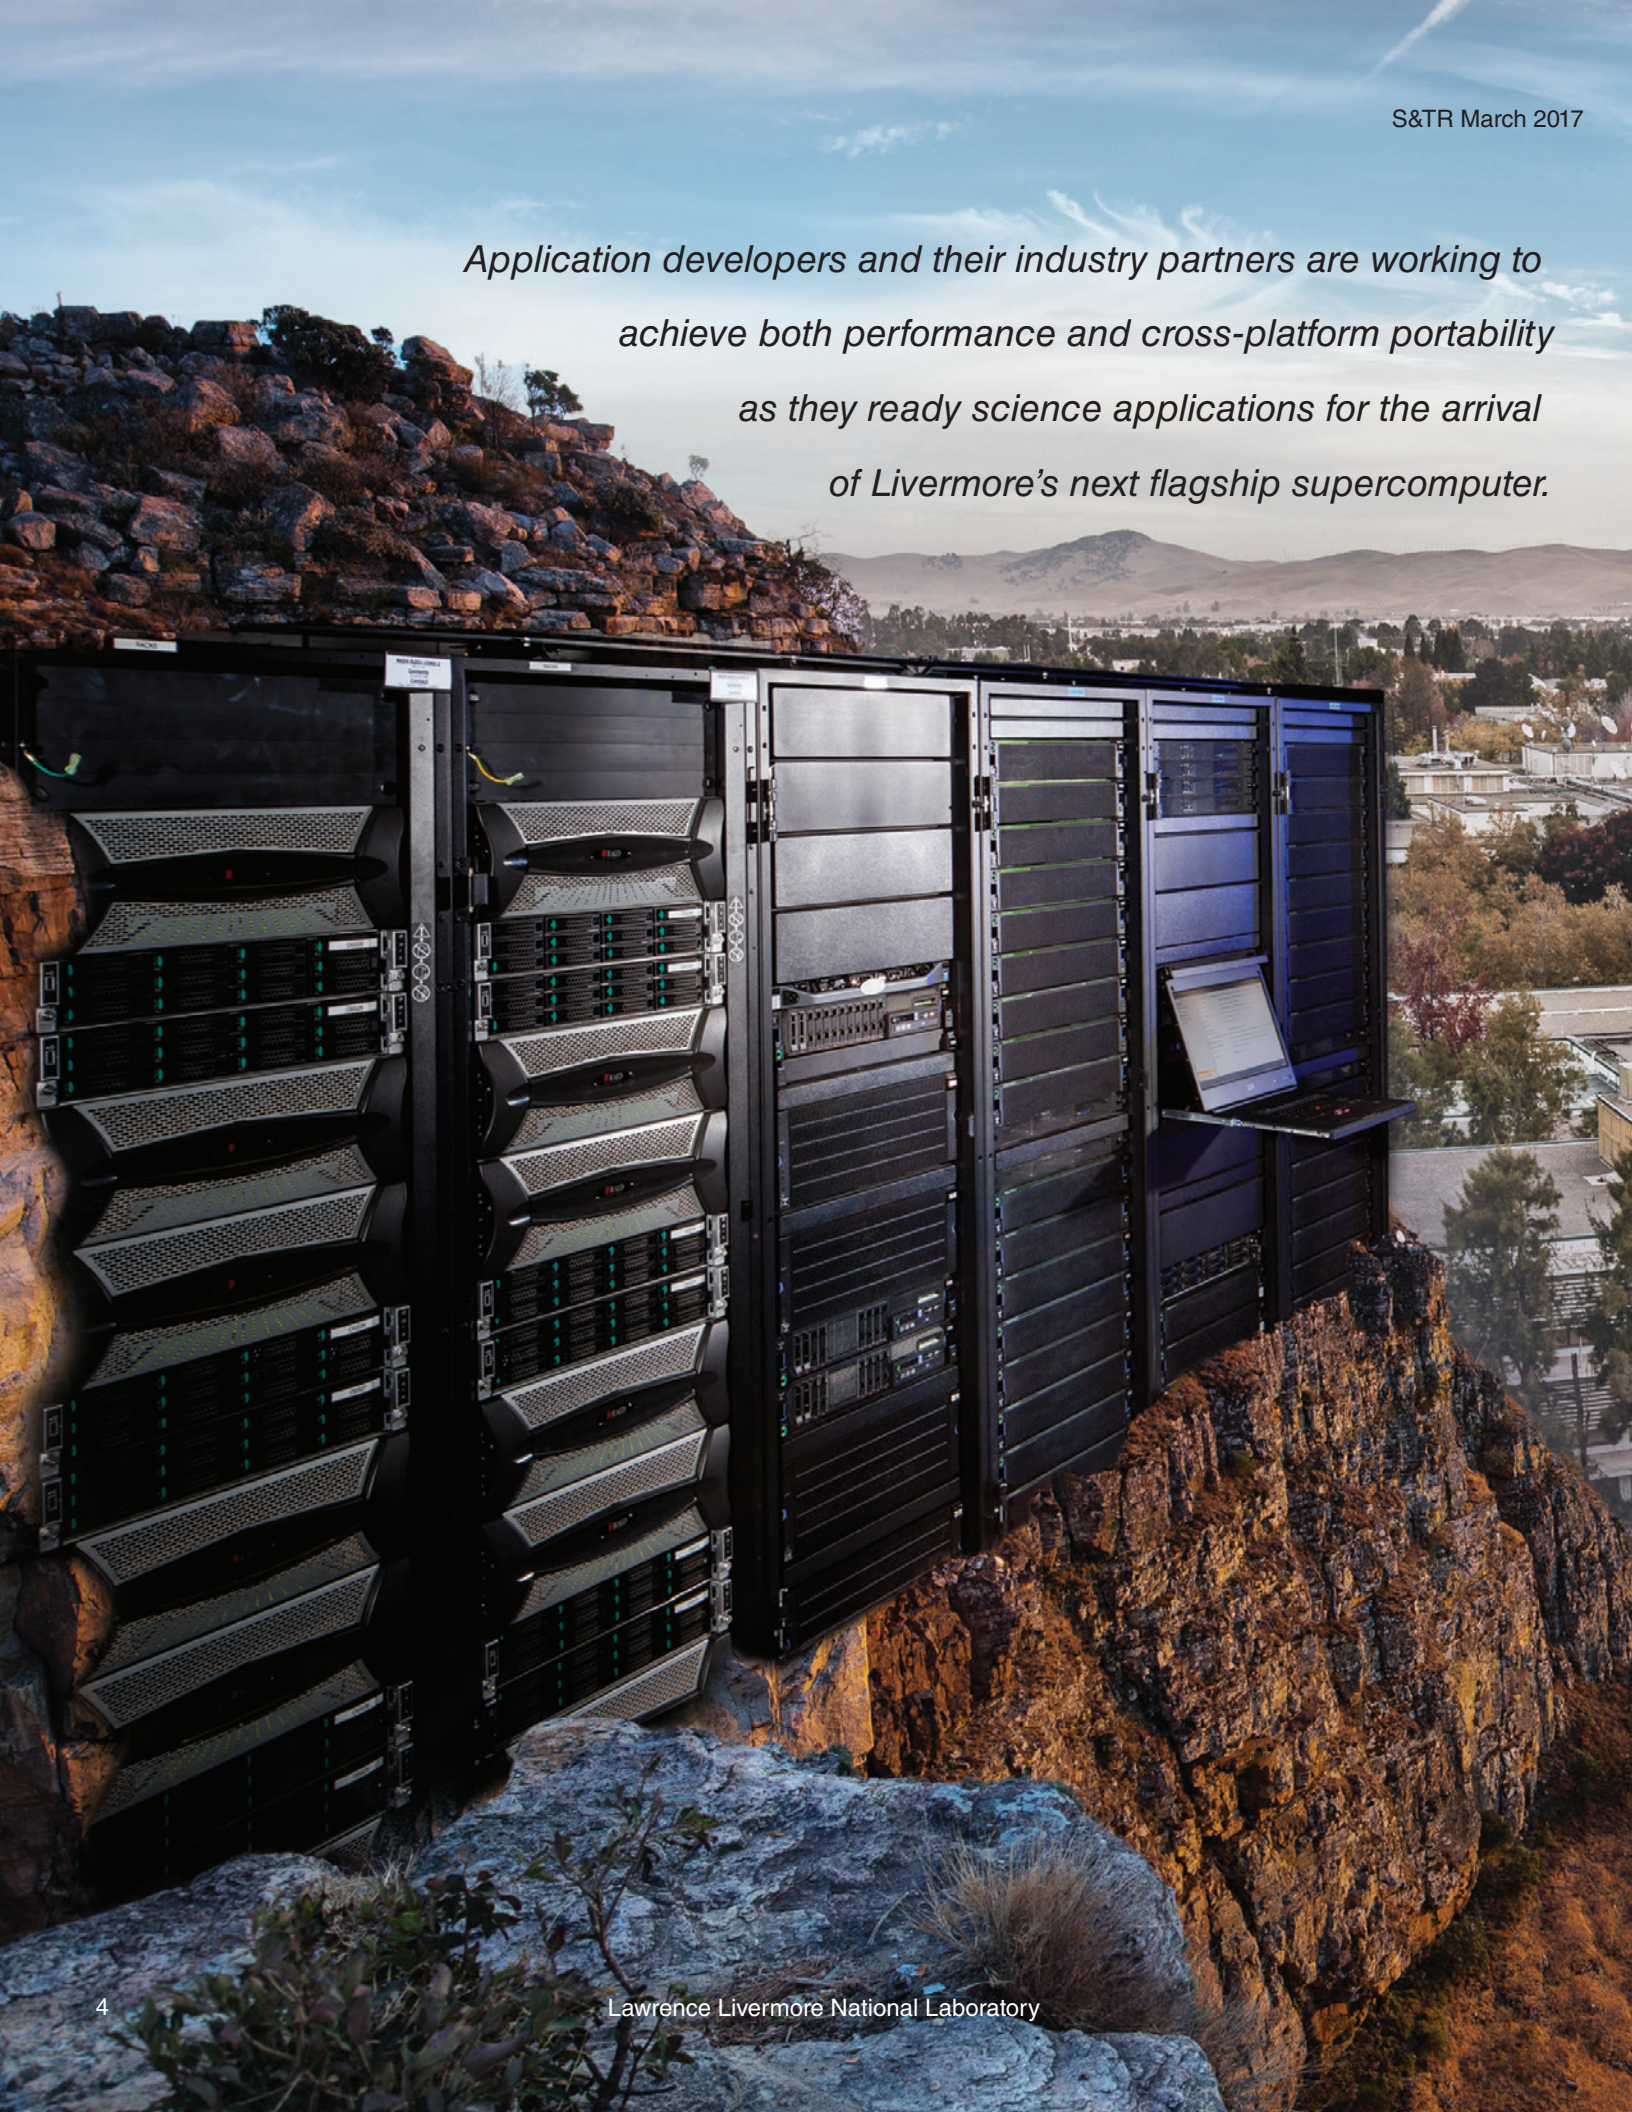

A Center of Excellence Prepares for SIERRA

IN November 2014, then Secretary of Energy Ernest Moniz announced a partnership involving IBM, NVIDIA, and Mellanox to design and deliver high-performance computing (HPC) systems for Lawrence Livermore and Oak Ridge national laboratories. (See *S&TR*, March 2015, pp. 11–15.) The Livermore system, Sierra, will be the latest in a series of leading-edge Advanced Simulation and Computing (ASC) Program supercomputers, whose predecessors include Sequoia and BlueGene/L, Purple, White, and Blue Pacific. As such, Sierra will be expected to help solve the most demanding computational challenges faced by the National Nuclear Security Administration's (NNSA's) ASC Program in furthering its stockpile stewardship mission.

At peak speeds of up to 150 petaflops (a petaflop is 10^{15} floating-point operations per second), Sierra is projected to provide at least four to six times the performance of Sequoia, Livermore's current flagship supercomputer. Consuming "only" 11 megawatts, Sierra will also be about five times more power efficient. The system will achieve this combination of power and efficienctly through a heterogeneous architecture that pairs two types of processors—IBM Power9 central processing units (CPUs) and NVIDIA Volta graphics processing units (GPUs)—so that programs can shift from one processor type to the other based on computational requirements. (CPUs are the traditional number-crunchers of computing. GPUs, originally developed for graphics-intensive applications such as computer games, are now being incorporated into supercomputers to improve speed and reduce energy usage.) Powerful hybrid computing units known as nodes will each contain multiple CPUs and GPUs connected by an NVLink network that transfers data between components at high speeds. In addition to CPU and GPU memory, the complex node architecture incorporates a generous amount of nonvolatile random-access memory, providing memory capacity for many operations historically relegated to far slower disk-based storage.

The Department of Energy has contracted with an IBM-led partnership to develop and deliver advanced supercomputing systems to Lawrence Livermore and Oak Ridge national laboratories beginning this year. These powerful systems are designed to maximize speed and minimize energy consumption to provide cost-effective modeling, simulation, and big data analytics. The primary mission for the Livermore machine, Sierra, will be to run computationally demanding calculations to assess the state of the nation's nuclear stockpile.

These hardware features—heterogeneous processing elements, fast networking, and use of different memory types—anticipate trends in HPC that are expected to continue in subsequent generations of systems. However, these innovations also represent a seismic architectural shift that poses a significant challenge for both scientific application developers and the researchers whose work depends on those applications running smoothly. "The introduction of GPUs as accelerators into the production ASC environment at Livermore, starting with the delivery of Sierra, will be disruptive to our applications," admits Rob Neely. "However, Livermore chose the GPU accelerator path only after concluding, first, that performance-portable solutions would be available in that timeframe and, second, that the use of GPU accelerators would likely be prominent in future systems."

To run efficiently on Sequoia, applications must be modified to achieve a level of task division and coordination well beyond what previous systems demanded. Building on the experience gained through Sequoia, and integrating Sierra's ability to effectively use GPUs, Livermore HPC experts are hoping that application developers—and their applications—will be better positioned to adapt to whatever hardware features future generations of supercomputers have to offer.

A Platform for Engagement

IBM will begin delivery of Sierra in late 2017, and the machine will assume its

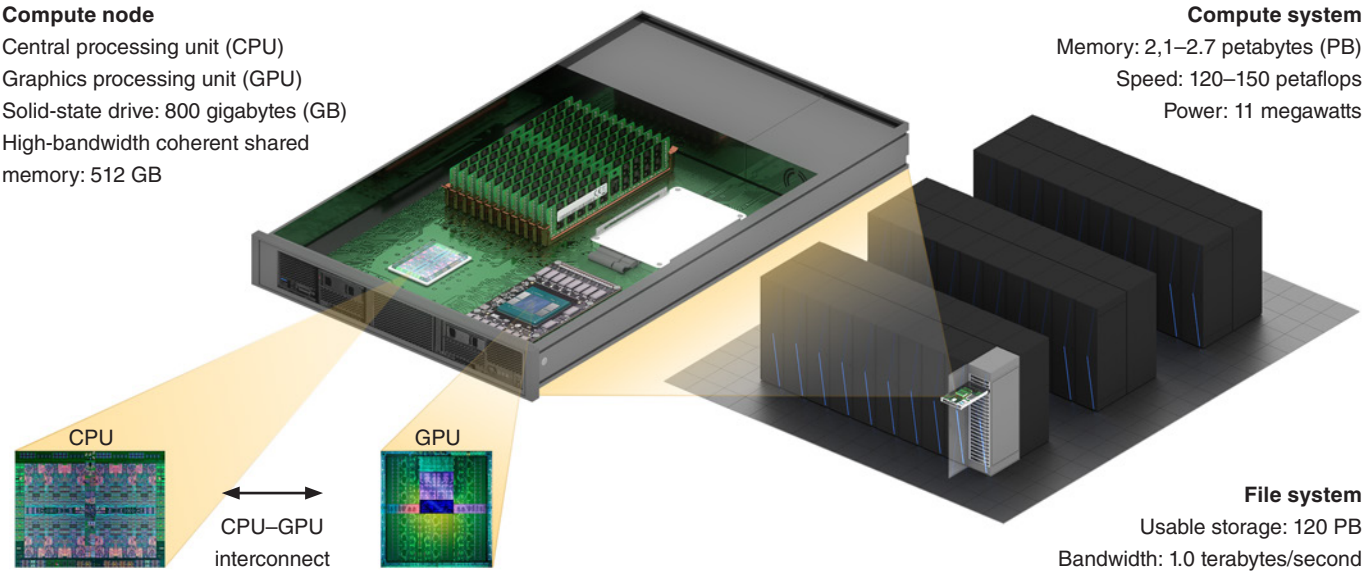

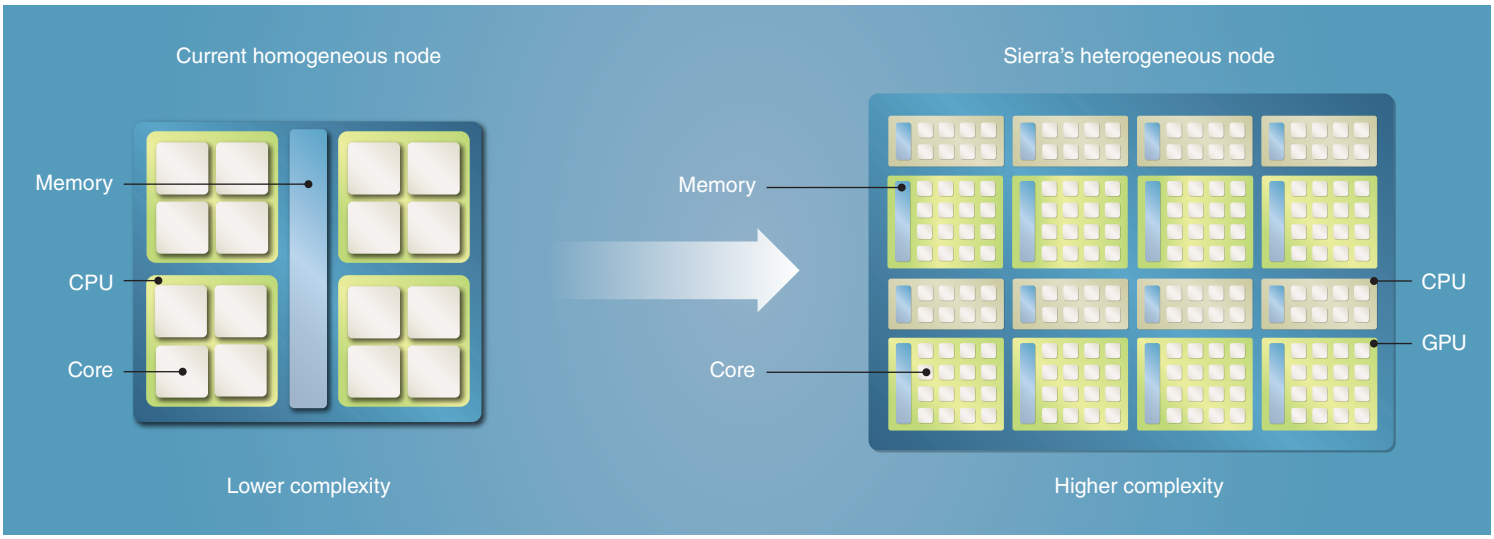

Compared to today's relatively simpler nodes, cutting-edge Sierra will feature nodes combining several types of processing units, such as central processing units (CPUs) and graphics-processing units (GPUs). This advancement offers greater parallelism—completing tasks in parallel rather than serially—for faster results and energy savings. Preparations are underway to enable Livermore's highly sophisticated computing codes to run efficiently on Sierra and take full advantage of its leaps in performance.

full ASC role by late 2018. Recognizing the complexity of the task before them, Livermore physicists, computer scientists, and their industry partners began preparing applications for Sierra shortly after the contract was awarded to the IBM partnership. The vehicle for these preparations is a nonrecurring engineering (NRE) contract, a companion to the master contract for building Sierra that is structured to accelerate the new system's development and enhance its utility. "Livermore contracts for new machines always devote a separate sum to enhance development and make the system even better," explains Neely. "Because we buy machines so far in advance, we can work with the vendors to enhance some features and capabilities." For example, the Sierra NRE calls for investigations into GPU reliability and advanced networking capabilities. The contract also calls for the creation of a Center of Excellence (COE)—a first for a Livermore NRE—to foster more intensive and sustained engagement than usual among domain scientists, application developers, and vendor hardware and software experts.

The COE provides Livermore application teams with direct access to vendor expertise and troubleshooting as codes are optimized for the new architecture. (See the box on p. 9.) Such engagement will help ensure that scientists can start running their applications as soon as possible on Sierra—so that the transition sparks discovery rather than being a hindrance. In turn, the interactions give IBM and NVIDIA deeper insight into how real-world scientific applications run on their new hardware. Neely observes that the new COE also dovetails nicely with Livermore's philosophy of codesign, that is, working closely with vendors to create first-of-their-kind computing systems, a practice stretching back to the Laboratory's founding in 1952.

The COE features two major thrusts—one for ASC applications and another for non-ASC applications. The ASC component, which is funded by the multilaboratory ASC Program and led by Neely, mirrors what other national laboratories awaiting delivery of new HPC architectures over the coming years are doing. The non-ASC component, however, was pioneered by Lawrence Livermore, specifically the Multiprogrammatic and Institutional Computing (M&IC) Program, which helps make powerful HPC resources available to all areas of the Laboratory. One way the M&IC Program achieves this goal is by investing in smaller-scale versions of new ASC supercomputers, such as Vulcan, a quarter-sized version of Sequoia. Sierra will also have a scaled-down counterpart slated for delivery in 2018. To prepare codes to run on this new system in the same manner as the COE is doing with ASC codes, institutional funding was allocated in 2015 to establish a new component of the COE, named the Institutional COE, or iCOE.

To qualify for COE or iCOE support, applications were evaluated according to mission needs, structural diversity, and, for iCOE codes, topical breadth. Bert Still, iCOE's lead, says, "We chose to fund preparation of a strategically important subset of codes that touches on every discipline at the Laboratory." Applications selected by iCOE include those that help scientists understand earthquakes, refine laser experiments, or test machine learning methods. For instance, the machine learning

application chosen (see *S&TR*, June 2016, pp. 16–19) uses data analytics techniques considered particularly likely to benefit from some of Sierra's architectural features because of their memory-intensive nature. Applications under COE purview make up only a small fraction of the several hundred Livermore-developed codes presently in use. The others will be readied for Sierra later, in a phased fashion, so that lessons learned by the COE teams will likely save significant time and effort. Still says, "At the end of the effort, we will have some codes ready to run, but just as importantly, we will have captured and spread knowledge about how to prepare our codes." Although some applications do not necessarily need to be run at the highest available computational level to fulfill mission needs, their developers will still need to familiarize themselves with GPU technology, because current trends suggest this approach will soon trickle down to the workhorse Linux cluster machines relied on by the majority of users at the Laboratory.

Performance and Portability

Nearly every major domain of physics in which the massive ASC codes are used has received some COE attention, and nine of ten iCOE applications are already in various stages of preparation. A range of onsite and offsite activities has been organized to bolster these efforts, such as training classes, talks, workshops focused on hardware and application topics, and hackathons. At multiday hackathons held at the IBM T. J. Watson Research Center in New York, Livermore and IBM personnel have focused on applications, leveraging experience gained with early versions of Sierra compilers. (Compilers are specialized pieces of software that convert programmer-created source code into the machine-friendly binary code that a computer understands.) With all the relevant experts gathered together in one place, hackathon groups were able to quickly identify and resolve a number of compatibility issues between compilers and applications.

In addition to such special events, a group of 14 IBM and NVIDIA personnel collaborate with Livermore application teams through the COE on an ongoing basis, with several of them at the Livermore campus full or part time. Neely notes that the COE's founders considered it crucial for vendors to be equal partners in projects, rather than simply consultants, and that the vendors' work align with and advance their own research in addition to Livermore's. For instance, Livermore and IBM computer scientists plan to jointly author journal articles and are currently organizing a special journal issue on application preparedness to share what they have learned with the broader HPC and application development communities.

Application teams are also sharing ideas and experiences with their counterparts at the other four major Department of Energy HPC facilities. All are set to receive new leadership-class HPC systems in the next few years and are making preparations through their own COEs. Last April, technical experts from the facilities' COEs and their vendors attended a joint workshop on application preparation. Participants demonstrated a strong shared interest in developing portable programming strategies—methods that ensure their applications continue to run well not just on multiple advanced architectures but also on the more standard cluster technologies. Neely, who chaired the

Livermore's Center of Excellence (COE), jointly funded by the Laboratory and the multilaboratory Advanced Simulation and Computing (ASC) Program, aims to accelerate application preparation for Sierra through close collaboration among experts in hardware and software, as shown in this workflow illustration.

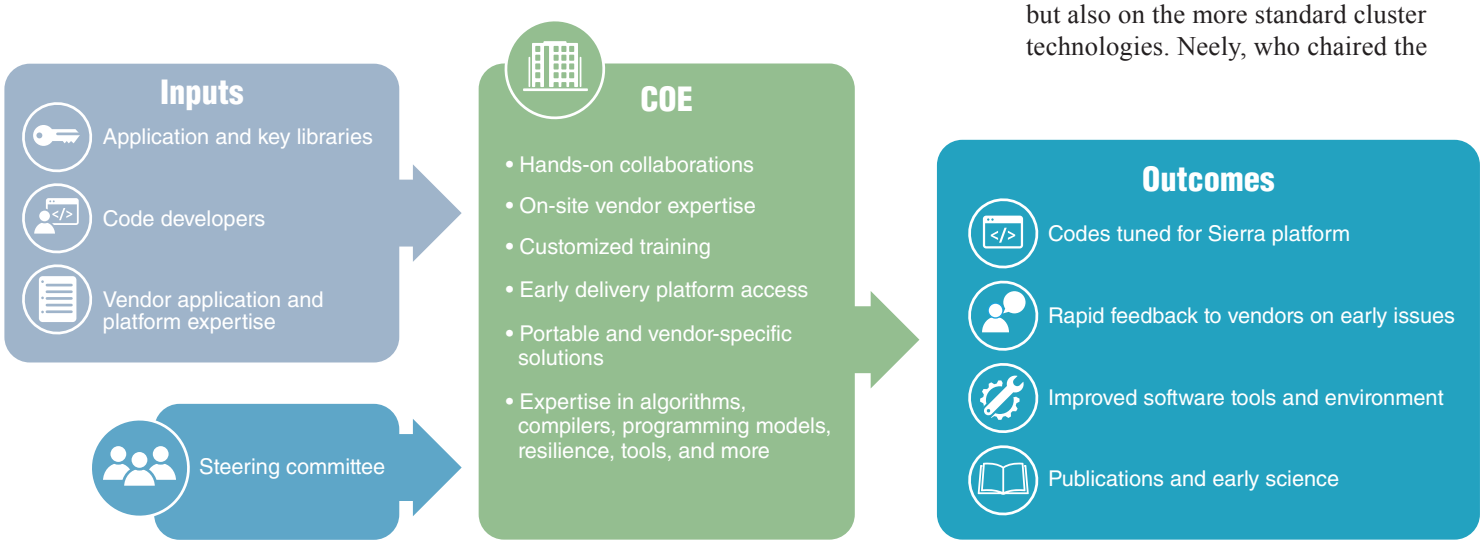

meeting, says, "Most of the teams in attendance need to run and support users on multiple systems, across either NNSA laboratories or the Leadership Computing Facilities of the Office of Science. This joint workshop was designed to examine the excellent work being done in these COEs and bring the discussion up a level, to learn how we, as a community, can achieve the best of both worlds—performance and portability." The consensus at the meeting was that the COE is indeed the right vehicle for exploring solutions to the daunting challenges they all face. Neely adds, "It may be too early to declare victory, but the COEs are working well."

A Smashing Development

Preparing an application for a new HPC system is an iterative process requiring a thorough understanding of the application and the new architecture. For instance, Sierra will feature more memory and a more complex memory hierarchy in the nodes than does Sequoia. Making full use of Sierra's features requires developers to carefully manage where data are stored across the memory hierarchy and to develop algorithms with a far greater degree of parallelism—the organization of tasks such that they can be performed in parallel rather than serially. Hitting performance targets on Sierra calls for 10 to 100 times more parallelism in applications than is present today. Developers must also identify the tasks best handled by the GPUs—which are optimized to efficiently perform the

A "Hit Squad" for Troubleshooting

In preparing Livermore's scientific applications to run on Sierra, everyone involved contributes something unique. Application developers understand the needs of their users and the performance characteristics of their applications. Vendors have in-depth understanding of the complex architecture of the coming machine, the systems software that will run on it, and the tradeoffs made during design and development. In-house experts on Livermore's Advanced Architecture and Portability Specialists (AAPS) team provide yet another essential contribution—robust expertise in developing and optimizing applications for next-generation architectures.

"The AAPS team pulls together exceptional talent in several key areas," says AAPS lead Erik Draeger. "We have computational scientists with years of hands-on experience in achieving extreme scalability and performance in scientific applications. Together, we have multiple Gordon Bell prizes. Our computer scientists have deep technical knowledge in the latest programming languages and emerging programming models. This breadth gives the team a clear sense of what is possible in theory and where things can get tricky in practice."

AAPS team members act as a troubleshooting "hit squad," according to Institutional Center of Excellence (iCOE) lead Bert Still. At any given time, team members are working with multiple application teams to provide hands-on support and advice on code preparation tailored to the unique characteristics of the code involved. Naturally, says Draeger, the biggest challenges arise with codes poorly suited for future architectures. Nevertheless, rewriting an application from scratch is rarely an option because of its size (sometimes millions of lines of code) and the effort already expended to create and enhance that code (sometimes years or even decades). In such instances, the AAPS team uses tools such as mini-apps (compact, self-contained proxies for real applications) to help determine precisely how much code needs to be rewritten and how to do so as efficiently as possible. (See *S&TR*, October 2013, pp. 12–13.)

The team's mandate also includes documenting and transferring knowledge about common challenges and successful solutions, particularly regarding strategies for portability—the ability of an application to run efficiently on different architectures with minimal modification. The biggest portability success thus far with Sierra may be developing the copy-hiding application interface (CHAI), an effort led by developer Peter Robinson, with AAPS team support. Draeger explains, "CHAI is an abstraction that allows programmers to easily write code that minimizes data movement and runs well on a variety of architectures. CHAI solves some of the key problems of writing code for both performance and portability."

The AAPS team is currently funded as part of Livermore's COE for Sierra, but Draeger hopes the team will endure well beyond the new system's arrival. He says, "High-performance computing is not likely to get any easier in the coming years, and we need to work together and share our experiences as much as possible if we're going to continue to be leaders in the field. AAPS is one important mechanism to achieve that goal."

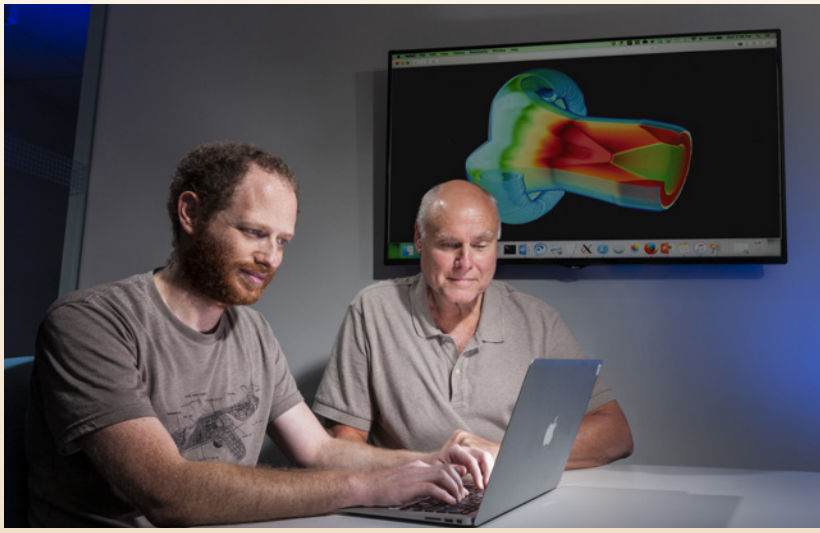

Ian Karlin (left), a member of Livermore's Advanced Architecture and Portability Specialists team, works with Tony DeGroot, part of the ParaDyn application team. (Photograph by Randy Wong.)

same operation across a large batch of data—and those to be delegated to the more general-purpose CPUs. At the same time, programmers will try to balance Sierra-specific optimizations with those designed to improve performance on a range of HPC architectures.

The iCOE team preparing the application ParaDyn readily acknowledges its challenges. ParaDyn is a parallel version of the versatile DYNA3D finite-element analysis software developed at Livermore beginning in the 1970s to model the deformation and failure of solid structures under impact. (See *S&TR*, May 1998, pp. 12–19.) The application performs well on today's HPC architectures but is proving difficult to prepare for Sierra. "We were fortunate in the past because our coding style worked well on Cray and today's HPC machines, including Sequoia," says Tony DeGroot. "Previously, we didn't have to make as many changes as other teams did to convert from machine to machine, but now we have to work harder to convert than some other teams do."

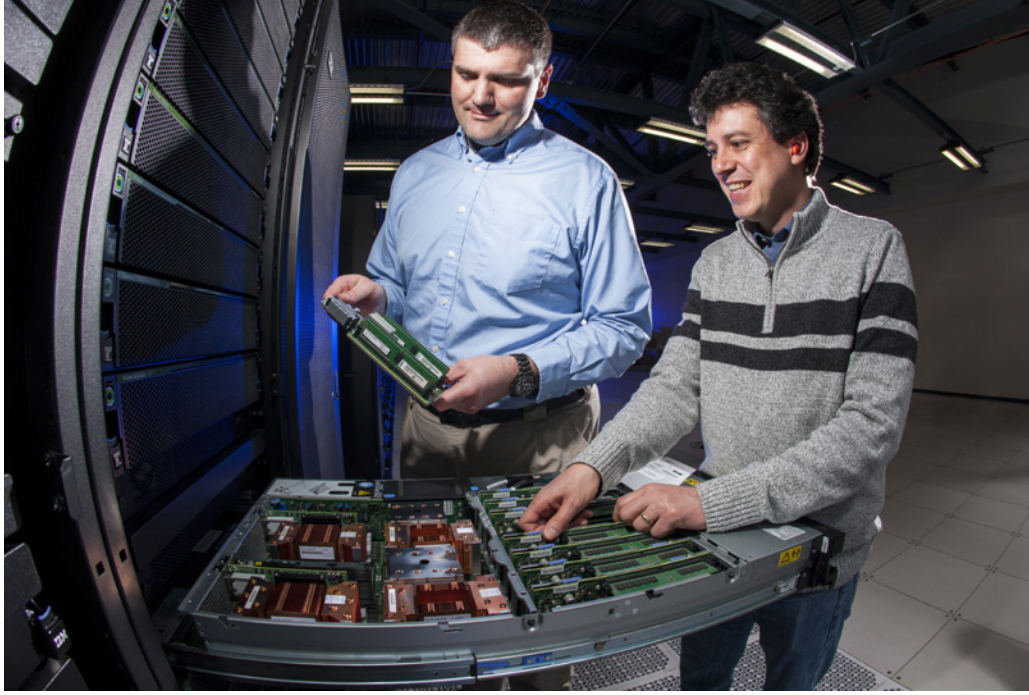

Adam Bertsch (left) and Pythagoras Watson examine components of the Sierra early-access machine, which is being used to help prepare codes to run on the full version of Sierra as soon as it comes online. (Photograph by Randy Wong.)

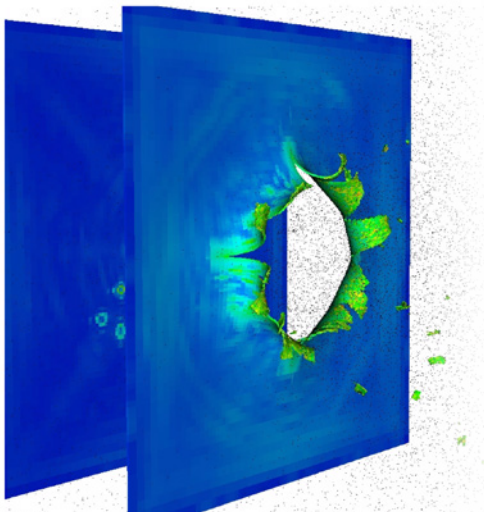

ParaDyn simulates the deformation and failure of solid structures under impact, such as the hypervelocity impact through two plates shown here. The ParaDyn team is presently exploring how to maintain the strong performance that the application achieves on CPUs while optimizing some aspects for Sierra's GPUs.

As with most teams, the ParaDyn crew began by assessing how the data move in the application and how work is divided up among its algorithms—the self-contained units that perform specific calculations or data-processing tasks. Edward Zywicz explains, "Unlike most physics codes, which have tens of millions of similar or identical chunks of a problem that can be processed in the same way, we have a large number of dissimilar groups. GPUs work best with large numbers of similar items. So the nature of our problems dictates some performance challenges." The ParaDyn team found the primary bottleneck to be data movement between GPUs and their high-speed memory. Consequently, the team is reorganizing and consolidating data and will need the help of the ROSE preprocessor (see *S&TR*, October/November 2009, pp. 12–13) to reduce the need for data movement. This strategy saves time and energy, the most precious resources in computing.

Pinpointing an application's performance-limiting features and testing the solutions are far more difficult when the hardware and systems software involved are still under development, as with Sierra. Hybrid CPU–GPU clusters at Livermore are a good stand-in but require far less parallelism to achieve good performance than does the more architecturally complex Sierra. Another crucial difference is that the hybrid clusters require the explicit management of data movement between CPUs and GPUs, which Sierra will not. DeGroot says, "These hybrid clusters help us identify bottlenecks and learn about our code and its performance, but they do not represent exactly how the code will perform on Sierra." To overcome this shortcoming, Livermore in late 2016 acquired two small-scale "early-access" versions of Sierra—one for ASC computing and another for other work. (These are separate from the permanent counterpart to Sierra being built by the M&IC Program.) These realistic stand-ins feature CPUs, GPUs, and networking components only one generation behind those of Sierra.

Heartening Results

Cardioid, a sophisticated application that simulates the electrophysiology of the human heart, is also being prepared for Sierra through iCOE. Developed to run on Sequoia by Laboratory scientists and colleagues at the IBM T. J. Watson Research Center, the powerful application

can, say, model the heart's response to a drug or simulate a particular health condition. (See *S&TR*, September 2012, pp. 22–25.) David Richards, one of Cardioid's developers, is leading the readiness effort. He explains, "Cardioid was highly optimized for Sequoia. We took advantage of features in Sequoia's CPU, memory system, and architecture to maximize performance. We are now trying to implement a more portable and maintainable design." What Cardioid may lose in performance will be more than made up for in improved features for researchers. Richards says, "Sierra will be six times more powerful in terms of raw flops than Sequoia. Even if we trade a little performance for flexibility, we will still be able to run more simulations, and run them faster, on Sierra." Potential collaborators have already expressed great interest in running immense, highly computationally demanding ensembles of Cardioid simulations, indicating the pent-up demand for a machine such as Sierra that can run more simulations simultaneously.

Cardioid has two major components developed specifically for Sequoia—a mechanical solver and a high-performance electrophysiology code. With help from iCOE, the Cardioid team intends to improve the portability and maintainability of the mechanics solver using components created and maintained at the Laboratory, such as the open-source MFEM finite element library. The team is also re-architecting the electrophysiology code for greater ease of use. Currently, the computational biologists and other researchers who use the code but lack advanced programming skills struggle to make complex modifications to the model. However, the new system will enable users to input their models and data in a familiar fashion. The system will then automatically translate that input into the format of OpenMP—a popular programming model that enhances the portability of parallel applications—and then produce a version of the code optimized to run on Sierra.

Code preparation is already netting promising results. When Cardioid team member Robert Blake optimized a representative chunk of the code for GPUs using OpenMP at a hackathon, he achieved 60 percent of peak speed for 70 percent of the code—a strong result. The Cardioid team also identified important issues with OpenMP that they are working to resolve in collaboration with the model's standards committee.

Looking Ahead

Thanks to the close and coordinated efforts of COE teams, vendors, users, and developers, a mission-relevant selection of Livermore's science applications will be ready when Sierra and its M&IC counterpart come online. All involved are fully confident that their preparatory efforts will enable scientists to make the most of the machines right from the start. Still says, "To further scientific discovery, we want to maximize use of the machines and minimize downtime for the five or so years of their operational lifespan."

Given the rapid evolution and obsolescence of computing hardware, the Laboratory is already planning for the generation of HPC systems to follow Sierra. Expected sometime around 2022, these systems will likely constitute another major leap in performance, up to exascale levels (at least 10^{18} flops), or 10 times the speed of Sierra. "We have a two-pronged approach for exascale," notes Neely, who is also involved in exascale planning. "We will attempt to carry forward as long as possible with existing codes. At the same time, we are exploring what we would do if untethered from existing designs to take advantage of the new architecture." (See *S&TR*, September 2016, pp. 4–11.) By increasing application portability,

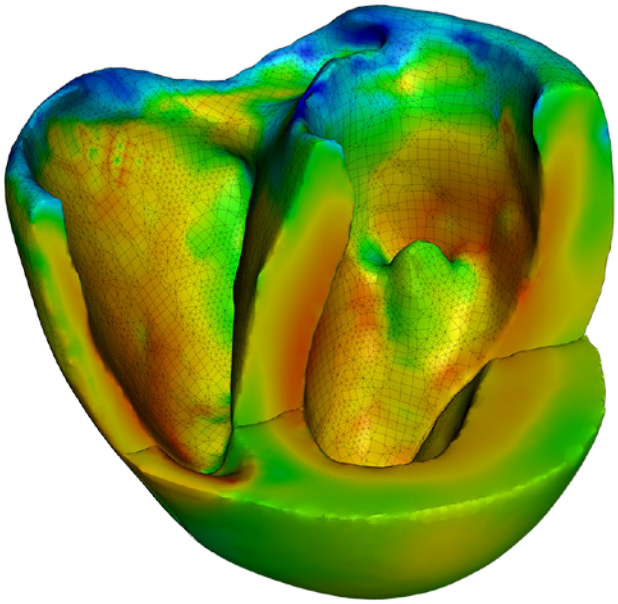

Developed by Lawrence Livermore and IBM to run on Sierra's predecessor, Sequoia, the highly scalable and complex Cardioid code replicates the heart's electrical system—the current that causes the heart to beat and pump blood through the body. Sierra is expected to enable Cardioid users to run larger ensembles of high-resolution simulations than was ever possible on Sequoia.

Livermore's twin COEs aim to extend the lifetime of many of Lawrence Livermore's immense, complex, and mission-relevant applications into the exascale era ahead.

—*Rose Hansen*

Key Words: Advanced Architecture and Portability Specialists (AAPS), Advanced Simulation and Computing (ASC) Program, Cardioid, Center of Excellence (COE), central processing unit (CPU), copy-hiding application interface (CHAI), graphics processing unit (GPU), high-performance computing (HPC), IBM, institutional Center of Excellence (iCOE), Linux cluster, Multiprogrammatic and Institutional Computing (M&IC) Program, node, nonvolatile random-access memory, NVIDIA, OpenMP, ParaDyn, petaflop, ROSE, Sequoia, Sierra, T. J. Watson Research Center.

For further information contact Rob Neely (925) 423-4243 (neely4@llnl.gov).

AGENT DEFEAT Efforts Strike Gold

THE U.S. military must be able to deal with a broad range of contingencies, including some that are extremely challenging technologically. One such challenge is to destroy a stockpile of chemical or biological agents before they can be deployed. In this scenario, a specialized "prompt agent defeat" (PAD) weapon is required. The ideal PAD weapon contains payload materials that are thermally persistent—burning long and hot enough to rapidly and thoroughly eliminate the agents—and that produce minimal blast overpressure, which could liberate a chemical or biological agent before it is destroyed and thereby cause collateral damage. Such weapons usually contain a high explosive to initiate and disperse the PAD payload.

Both CrashPAD, developed in 2004 to destroy agents stored in lightweight structures, and Shredder, a more penetrating version developed in 2007, use a payload of white phosphorus. However, this pyrophoric material spontaneously combusts on contact with air and is therefore dangerous to handle and store. When the effort to create a replacement for CrashPAD began, Lawrence Livermore was tasked with helping develop a payload material with less risk than white phosphorus but without sacrificing effectiveness.

Lawrence Livermore has contributed to such agent defeat research and development efforts—including Shredder and CrashPAD—in various capacities since 1999. The Laboratory's tasks have included concept development, design, materials

development, safety analysis, testing, and characterization. For its latest charge, the Livermore agent defeat team, in partnership with the Naval Postgraduate School (NPS), has developed a thermite-based payload that is far more effective than white phosphorus and much easier to work with. The breakthrough recently earned Livermore team lead John Molitoris and his colleagues their third Laboratory commendation for these efforts, a 2016 Gold Award from the Weapons and Complex Integration (WCI) Principal Directorate for exceptional innovation and excellence.

Formula for Success

Prior reactive materials experience and initial feasibility testing enabled the Livermore team to quickly determine that thermites, a type of pyrotechnic, were the most feasible payload replacement, potentially offering a level of thermal persistence and combustion comparable or superior to that of white phosphorus. Livermore's High Explosives Applications Facility (HEAF) provided the ideal environment for solving the project's two biggest scientific questions—the thermite formulation optimal for agent defeat and the best way to initiate the thermitic reaction. Molitoris observes, "We have a great facility and diagnostics to get the work done, but the true beauty of HEAF is that all of us—experimentalists, physicists, formulators, chemists, modelers, and technicians—are working under the same roof. Our scientific interactions are increased, and we get much more accomplished. For this type of work, there really isn't any other place like HEAF."

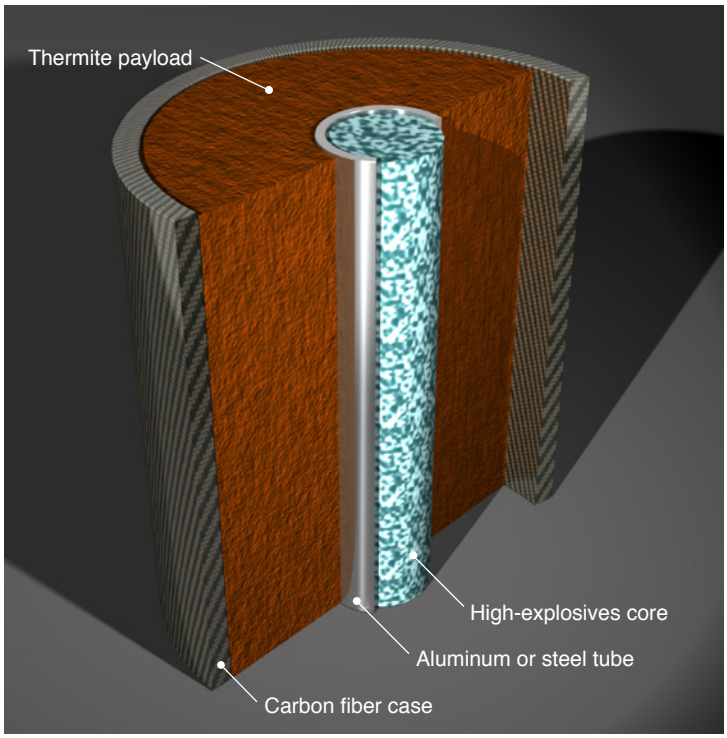

A cross-section drawing shows the components of the prototype agent defeat weapon that was subjected to performance tests in the firing tank at Livermore's High Explosives Applications Facility. (Rendering by Sabrina Fletcher.)

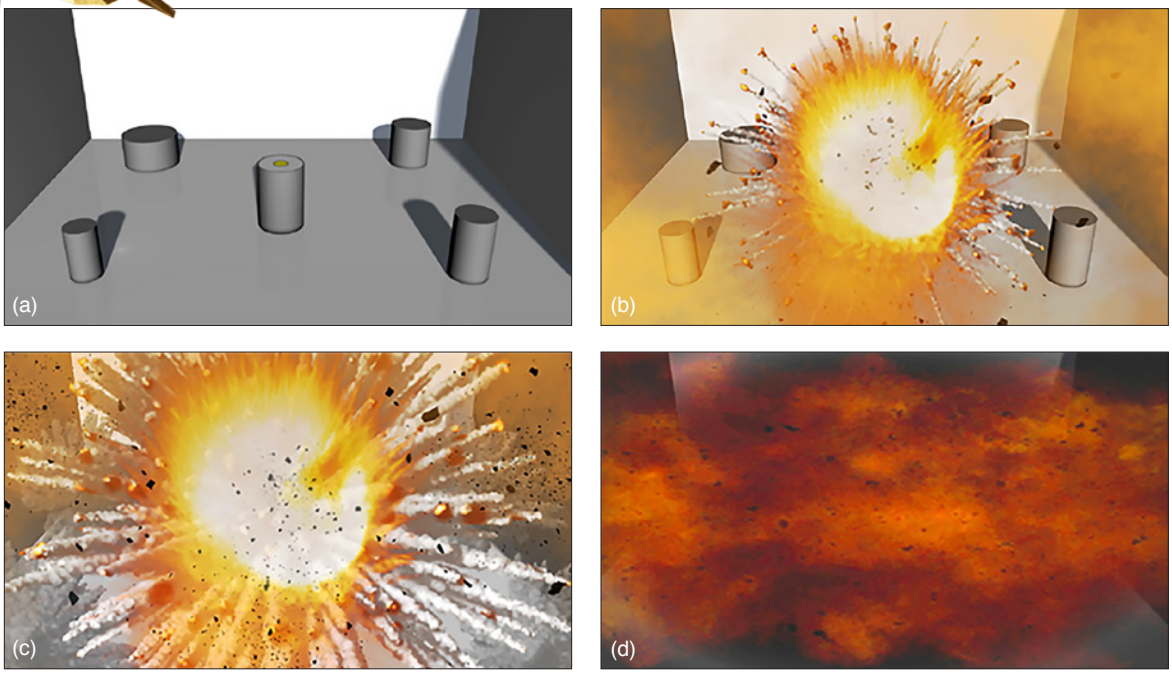

Livermore researchers are helping develop a weapon that can penetrate a storage facility and destroy chemical or biological agents inside before they can be deployed. Illustrations depict (a) a thermite-based payload (in the center of the space shown) delivered to a storage facility containing canisters of agents; (b) initiation of the weapon by high explosives; (c) creation of a thermitic fireball; and (d) establishment of a persistent thermitic cloud, which rapidly and efficiently destroys the agents. (Renderings by Sabrina Fletcher.)

The team has changed its approach over the years in response to the myriad challenges it has addressed. “Initially, we had to build specialized diagnostics to do the research, but now the effort is more about formulations and performance testing for specific applications,” notes Molitoris. Materials scientists Alex Gash and Octavio Cervantes have contributed heavily to recent work, as has formulator Mark Hoffman, whose expertise is in polymers and plastic-bonded explosives. Once Molitoris has formed the right team for a given application, the team determines what material properties will be required and formulates an initial set of viable solutions. Critical performance testing then feeds back into development of the formulation until they reach the best possible solution.

“Every time we make a material, we have to conduct performance tests to see how it behaves,” explains Cervantes. “In this case, we were able to gather the needed information with relatively simple tests.” The team began testing candidate formulas by igniting samples in horizontally positioned, clear plastic burn tubes. When initiated, a thermite’s metal fuel and metal oxide components undergo a self-contained and self-sustained chemical reaction, producing a burst of flame. The team used a high-speed camera to document how long and how well each material burned, but they soon determined that burn tubes were too chaotic an environment to make the required measurements. Instead, they developed a new vertical platform in which the test material is sandwiched between a cylindrical metal base and an igniter and burns from the top down. This test arrangement, which took nine iterations to perfect, proved more satisfactory for studying flame propagation and temperature gradients.

The more promising formulations were further evaluated in HEAF’s spherical firing tank, which is equipped with pressure and temperature measurement instruments and high-resolution optical and x-ray imaging capabilities. These experiments employed a cylindrical carbon fiber case—which x rays can easily penetrate—packed with a donut-shaped thermite payload around a core of high explosives. The high explosives and thermites were separated by a metal liner. During performance tests, the tank’s multichannel flash x-ray system (see *S&TR*, December 2004, pp. 22–25, and March 2006, pp. 11–13) captured a sequence of radiographs to document the dynamic response of the PAD payload, from detonation to full thermite initiation and dispersal.

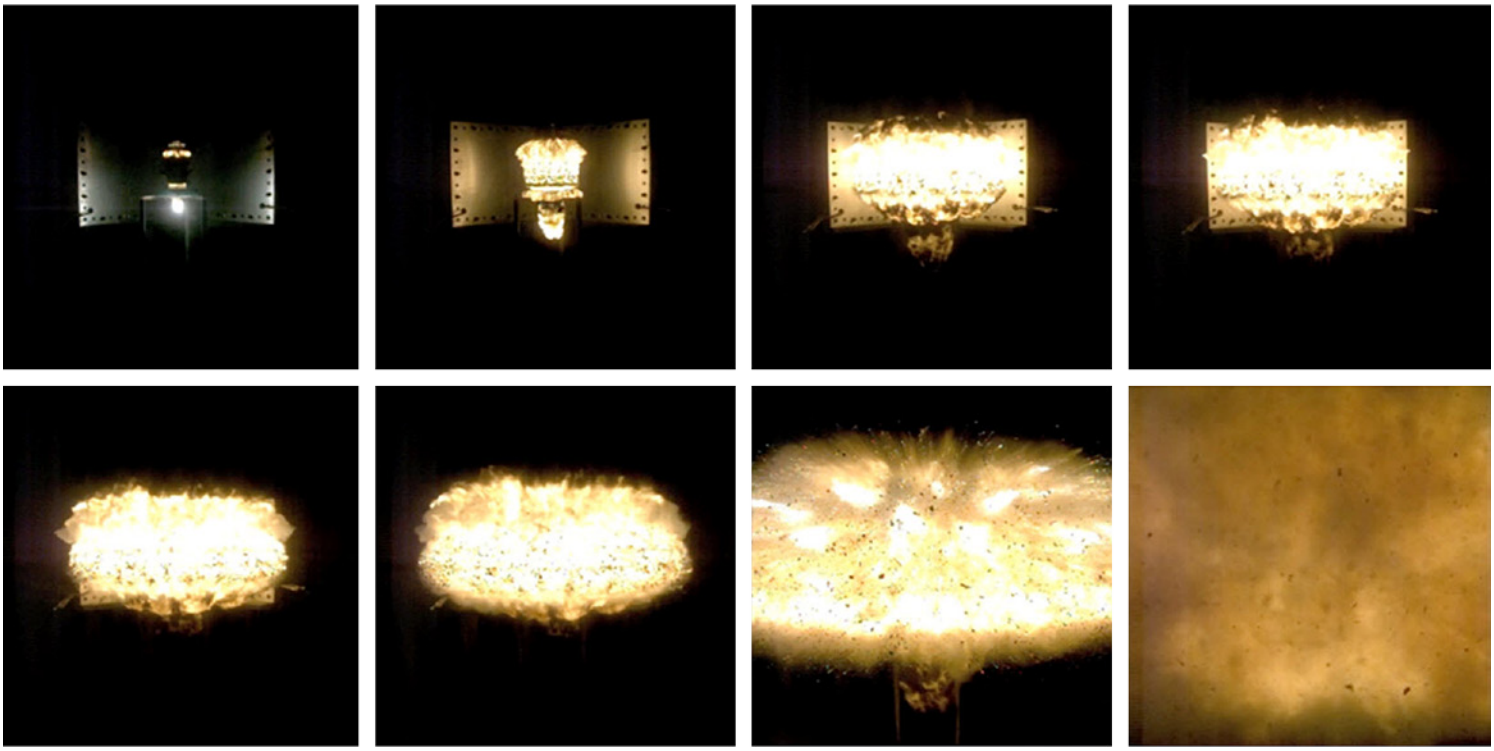

High-resolution images taken during firing-tank tests capture the evolution of a thermitic fireball. These images can be used to assess factors such as how efficiently the payload was initiated and dispersed.

Staff and students at the Naval Postgraduate School have characterized burn residue to help the Livermore team better understand the thermite reaction mechanism. A magnified sample is shown.

A First for Thermites

With the firing-tank test results, the researchers settled on a formula and initiation method that was both efficient and effective, producing a fireball that is hotter and more persistent, with less overpressure, than is possible with white phosphorus. Subsequent lethality tests performed with surrogate agents confirmed that Livermore's thermite payload was better at capturing and destroying its target than CrashPAD or other possible payloads. "Thermites were even more effective than we imagined," Molitoris says. The sponsors were very pleased with the results, but they soon challenged Livermore to take the research a step further—to create a plasticized (liquid) thermite that offered similar or identical performance to the newly proven granular (solid) formula, for even greater safety and ease of handling. If successful, the new payload would be used in a new incarnation of Shredder, dubbed the Agent Defeat Penetrator.

In addition to an accelerated timeline of only 15 months, the project entailed a high degree of risk because of the failure of all previous attempts by others to plasticize a thermite. Nevertheless, the Livermore team accepted the challenge. Even with the HEAF team's collective experience with formulating energetic materials, significant effort was needed to find a binder that, in just the right proportions, would produce a substance with the desired viscosity—like toothpaste or soft wax—and with a shelf life measured in years, all without hampering performance. The Livermore team evaluated possible formulas for performance, while NPS professors Joe Hooper and Young Kwon and their students shed light on the thermitic reaction by characterizing residue from stack burn and explosive tests and exploring the influence of the carbon fiber casing on performance test results.

During their experiments, the researchers noticed that plasticized thermites burned less intensely in stack tests than did regular thermites but performed exceptionally well in the HEAF firing-tank experiments. "We saw better and more realistic results under confined conditions—inside a container—than in the open conditions of the stack burn tests," notes Cervantes. "Flash x rays at HEAF have allowed us to understand scientifically how confinement makes the system work better." The team was ultimately able to deliver a plasticized thermite—the first of its kind—suitable for PAD weapons. The material met or exceeded all requirements in Livermore's performance tests and also exceeded expectations in the lethality tests. This successful plasticization of thermites for agent defeat is what earned the team its 2016 WCI Gold Award.

As the Agent Defeat Penetrator project transitions to a new stage—payload formula optimization and larger scale testing—Livermore researchers will continue to lend their considerable expertise in energetic and advanced materials development and testing. In collaboration with NPS, the team will further enhance the formula's properties by fine-tuning ingredient proportions, particle size, and material mixture. They will also lead scaled-up tests at the Laboratory's Site 300 facilities. Molitoris says, "Scaling up can be one of the most difficult tasks because of the need to preserve not only the defeat mechanism itself but all the other properties of the thermite fill, including performance, insensitivity, gas production, amenability to cast curing, and shelf life." If their two-decade track record in agent defeat research and development is any indication, Molitoris and his HEAF colleagues will surely rise to the challenge once again.

—Rose Hansen

Key Words: Agent Defeat Penetrator, CrashPAD, Gold Award, High Explosives Applications Facility (HEAF), Naval Postgraduate School (NPS), prompt agent defeat (PAD), pyrotechnic, Shredder, thermite.

For further information contact John Molitoris (925) 423-3496 (molitoris1@llnl.gov).

Mighty ATLAS Supports Precise Alignment

ONE of the most demanding aspects of preparing for an experiment at the National Ignition Facility (NIF) is reliably and precisely aligning the myriad components involved—up to 192 laser beams; the target containing the material to be studied and at which the beams are directed; and a variety of optical, x-ray, and nuclear diagnostics. For instance, even a slight misalignment of a diagnostic with the target could result in crucial data not being captured accurately. A particularly exacting and potentially time-consuming aspect of alignment is fine-tuning the diagnostic instrument manipulators (DIMs) located inside the 10-meter-diameter target chamber. Proper alignment—to a demanding tolerance of 0.1 to 0.5 millimeters—ensures that desired data are acquired by the incredibly sensitive diagnostic instruments inside the DIMs.

Traditionally, DIM alignment has been handled by two telescope- and camera-based alignment systems, known respectively as the opposed port alignment system (OPAS) and the chamber interior viewing system (CIVS). However, using these systems takes alignment operators up to an hour before every shot and requires manual interpretation of the alignment. To improve accuracy and repeatability, reduce maintenance, and provide greater flexibility in alignment configuration, NIF engineers have developed the advanced tracking laser alignment system (ATLAS).

Currently undergoing final commissioning, ATLAS "represents a new approach for NIF diagnostic alignment," declares Dan Kalantar, senior scientist for the NIF target area. ATLAS uses a commercial laser tracker located inside the NIF target bay—a concrete silo 30 meters in diameter that encloses the target chamber—to locate and align more precisely the DIMs and their "snouts," the ends of the diagnostic instruments positioned closest to the target. The laser tracker is mounted on one of the target chamber's glass ports and looks into the target chamber from the target bay.

Checking Pitch and Yaw

Although many NIF diagnostics are in fixed positions inside the target chamber, a critical subset of diagnostics can be moved in and out, allowing the configuration to be customized to each experiment. These adjustable diagnostics ride on DIMs and require a more precise level of alignment than the fixed diagnostics. For many years, NIF had three DIMs—two positioned at the target chamber "equator" and one situated at the chamber's "north pole," looking down at the target.

Previously, each equatorial DIM was aligned by a dedicated OPAS telescope located at a port on the opposite side of the target chamber from the DIM. This arrangement permitted a direct

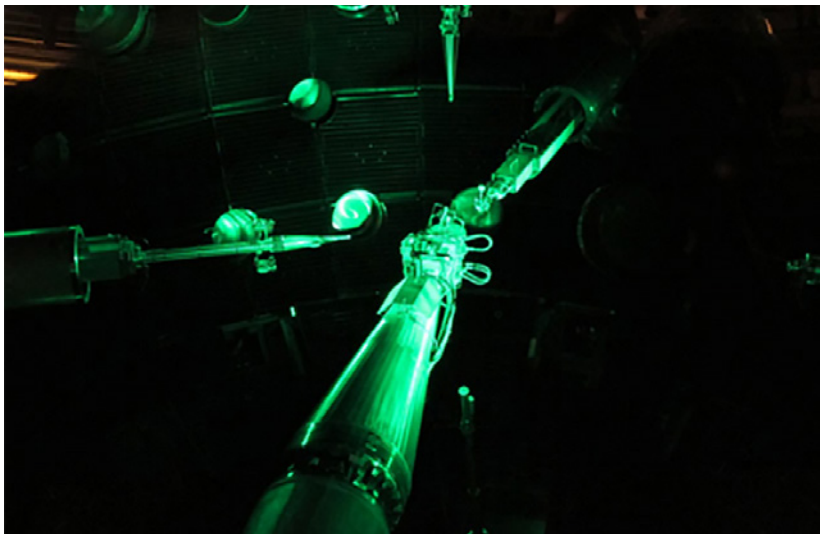

Target and diagnostic positioners inside the National Ignition Facility's (NIF's) target chamber are crucial to a successful experiment. An equatorial diagnostic instrument manipulator (DIM) enters from the left, and a polar DIM descends from the top of the chamber. NIF engineers have developed the advanced tracking laser alignment system (ATLAS) to streamline DIM alignment and improve accuracy and repeatability.

ATLAS has significantly improved alignment of the polar DIM (red) descending vertically from the top of the target chamber, as seen in this artist's depiction. (inset) ATLAS, mounted on one of the target chamber's glass ports, looks into the target chamber. (Illustration by Max Rodriguez.)

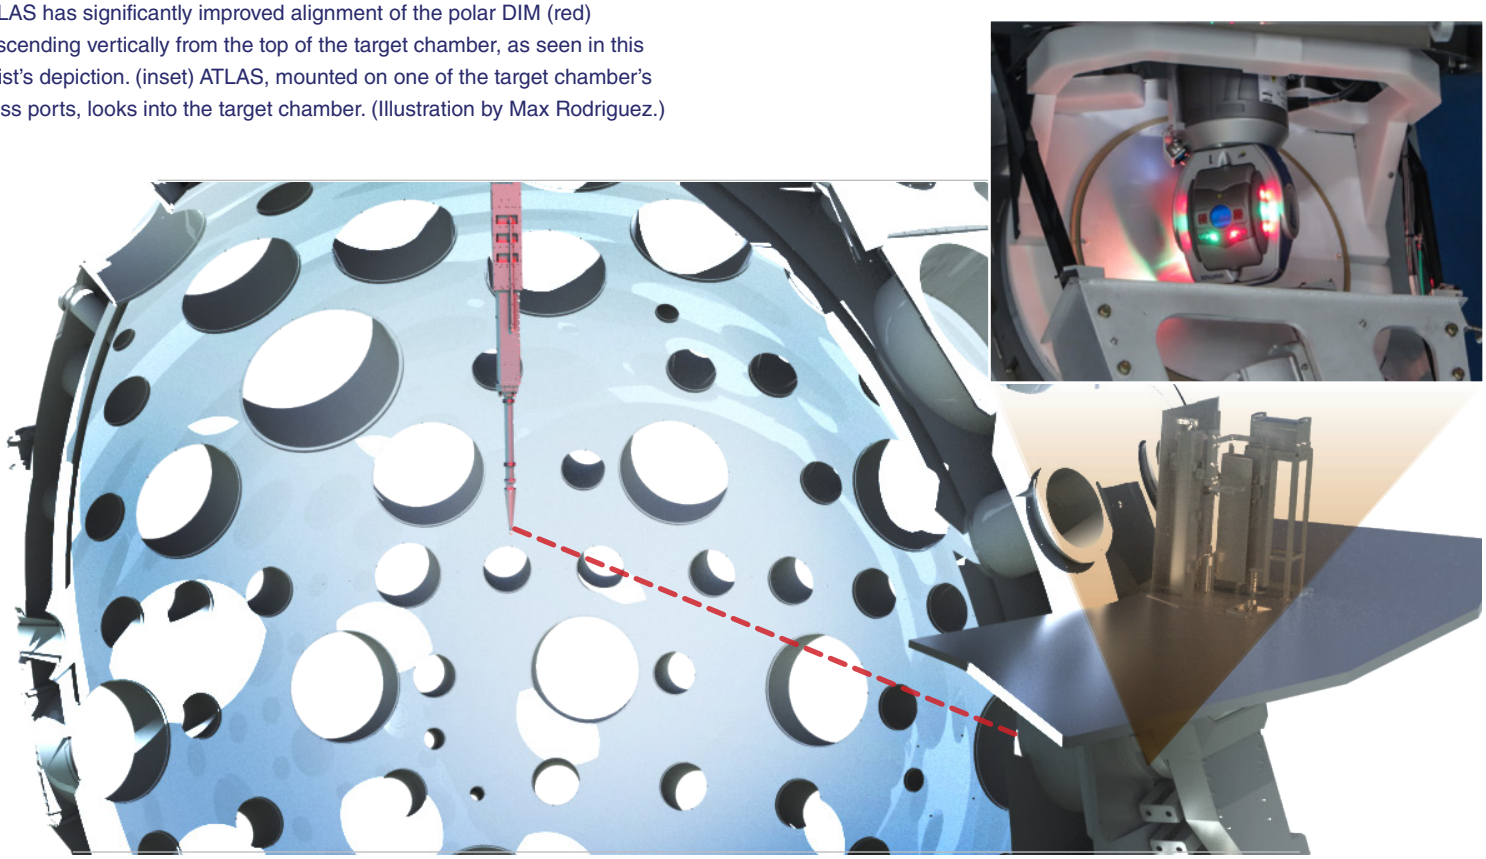

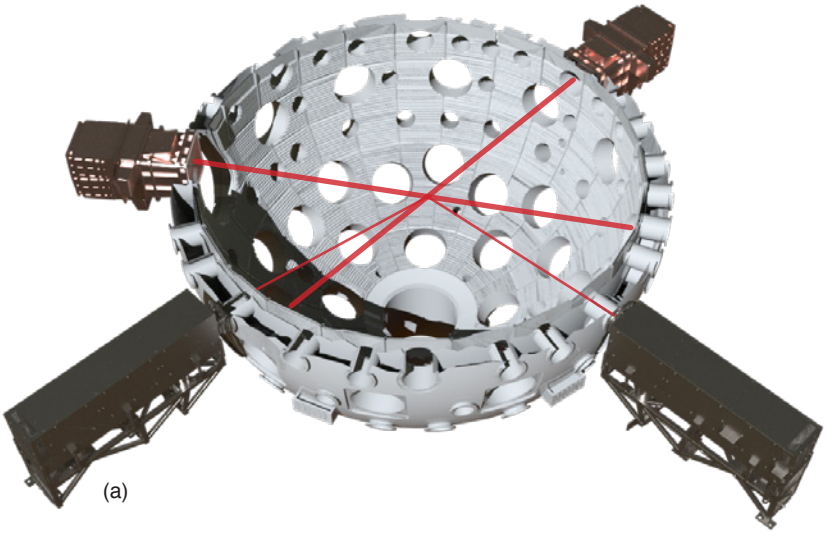

Compared to (a) the previous alignment system, the wide field of view afforded by (b) the ATLAS laser tracker permits faster and more accurate alignment of DIMs than was possible before. (Renderings by Max Rodriguez.)

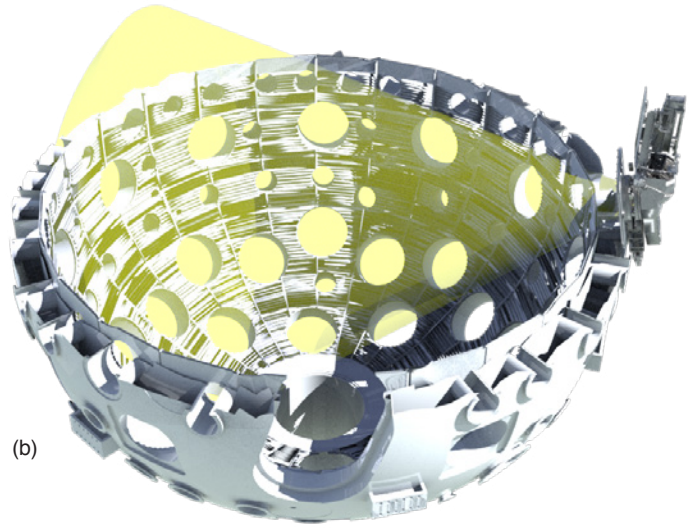

line of sight from the port to the diagnostic snouts and therefore facilitated alignment, including adjusting pitch (up and down) and yaw (left and right) angles. A particular challenge was aligning the polar DIM. An OPAS telescope could not be located opposite from the telescope because at the bottom of the chamber, instead of a port, is a two-person elevator for workers accessing the chamber. Consequently, NIF engineers used the CIVS video cameras, which provided a three-dimensional view of the polar DIM snout. However, these cameras were a temporary solution never engineered to be a precision alignment instrument. Hence ATLAS.

In parallel with ATLAS development, NIF engineers are adding two new manipulators in the target chamber, each capable of positioning both targets and diagnostics. The first new manipulator was installed in the NIF target bay in June 2016, and a second is scheduled for installation in 2017. The manipulators have their own diagnostics, which cannot be aligned with the OPAS or CIVS systems but instead rely on ATLAS.

ATLAS is also optimized to work with a new target alignment sensor, which, along with aligning the target and laser beams, ensures that the target is positioned in the exact center of the target chamber, so it can be seen by the surrounding diagnostics. The sensor is equipped with retroreflectors. These corner-cube mirrors act as reference points, telling ATLAS where the target will be. For every shot, ATLAS first measures the sensor to locate the target, after which, during alignment, the diagnostics are pointed at the target.

Improving on Human Judgment

Diagnostic alignments using the previous camera-based systems were performed by skilled alignment operators who relied on visual features and overlayed graphical images to assist with measurements. These operators positioned DIMs with high accuracy despite variations in visual features caused by lighting conditions, past shot debris, and radiation-induced artifacts on the video camera sensors. Although effective, manual alignment consumed an operator for up to one hour for each shot. Alignment based on ATLAS relies solely on objective measurement data from a laser tracker, not on an operator's skill or instincts.

Another key benefit of ATLAS is improved shot-cycle flexibility, reducing shot-cycle time and increasing overall efficiency. With the large field of view that ATLAS has from its perch above many of the equatorial manipulators, the system can measure DIMs even after the target has been aligned, whereas OPAS made it impossible to insert a target before the equatorial DIMs were aligned, because the target would block the line of sight through the target chamber center. In addition, ATLAS—including its polar DIM and manipulator diagnostics—takes up only a single port, in contrast to the multiple ports taken up by the previous telescopes and cameras.

During operation, the ATLAS laser tracker sends collimated low-power, eye-safe laser light across the target chamber to the DIM snouts. The tracker's servo system steers the laser beam to center on the snout-mounted retroreflectors, which bounce back more than 90 percent of the laser light. The reflected light re-enters the tracker, which analyzes the distance and angle to the retroreflector, providing a three-dimensional point location.

ATLAS measures a group of retroreflectors attached to the diagnostic snout to compute the pitch, yaw, and aim point of the snout. From the resulting measured location, the position of the DIM is adjusted so that the diagnostic will capture the desired view of the target. Finally, before the NIF laser beams fire, a shutter closes over the fused silica glass port window to shield ATLAS from debris. If an experiment will release large amounts of neutrons, which would freely pass through the metal shutter, the instrument is removed entirely from the target bay.

Part of an Effort to Increase Shot Rates

NIF facility manager Doug Larson explains that ATLAS development goes back to 2011. Effort was refocused in 2014, when Congress mandated that the NIF shot rate be increased and operations be optimized over three years—without additional funding. A 120-day study identified 20 areas for improvement to NIF equipment and procedures that could reduce the time and effort for fielding experiments. To meet these goals, over 80 individual projects were launched. ATLAS was one of the largest of these projects. Larson says, "ATLAS represents the final commitment of the 120-day study." (Another improvement was the new type of manipulator that can handle both targets and diagnostics.) In all, these and other improvements have roughly doubled the shot rate over a 2-year period, from 191 shots in fiscal year (FY) 2014 to 356 shots in FY 2015 and 417 shots in FY 2016.

ATLAS is expected to be fully commissioned by the summer of 2017 but is already being routinely used today to align the polar DIM. Lead ATLAS engineer Nick Shingleton says the new system is far exceeding the alignment tolerance of less than 500 micrometers, or half a millimeter. Shingleton explains that it is not surprising that ATLAS does a better job with polar DIM alignment, because the system was designed from the start as a precision alignment solution.

Integrating ATLAS into Operations

Shingleton notes that an important part of ATLAS development was writing new software to incorporate the alignment system into NIF's integrated computer control system, which governs every aspect of every shot. Several software engineering challenges had to be overcome, such as building a layer of front-end controls and seamlessly integrating the data captured by ATLAS into the overall system. As a result, all ATLAS alignment activities can now take place in the NIF control room. There, at special consoles, alignment operators oversee the collection and processing of ATLAS measurements and DIM alignment.

Although laser trackers are common in research and industry, ATLAS is unusual for propagating laser light through a glass window and into the target chamber vacuum, which causes the light beam to refract. Consequently, the raw measurements from the laser tracker must be compensated to account for this refraction and determine the true point. Compensation is performed automatically by an algorithm embedded in the control system, in a process transparent to users conducting the experiment.

Control system software engineer James Hoffman monitors the operations of ATLAS, seen in the top right of the photograph.

Once fully commissioned, ATLAS will help ensure that NIF operations run smoothly and with ever-increasing efficiency, and that critical data are reliably acquired by some of the most sensitive and advanced diagnostic instruments in science. Resting on the solid shoulders of ATLAS, experiments furthering national security, fundamental high-energy-density research, and energy security will all greatly benefit.

—Arnie Heller

Key Words: advanced tracking laser alignment system (ATLAS), diagnostic instrument manipulator (DIM), opposed port alignment system (OPAS), chamber interior viewing system (CIVS), National Ignition Facility (NIF), target bay, target chamber.

For further information contact Doug Larson (925) 422-1524 (larson8@llnl.gov).

Forcing Failure in Granular Materials

SUBJECTING small material samples to extreme pressure is a mainstay of research at Lawrence Livermore. Such high-energy experiments help solve national security challenges by advancing scientific knowledge of material properties. Experiments on granular media, which are prevalent throughout industry and nature, enable characterizations relevant to the Laboratory's defense, nonproliferation, and foundational science missions. For example, to interpret seismic signals from illicit underground weapons tests and to develop methods to counter these threats, scientists must first understand how sand and soil behave under external pressure. Doing so requires an intimate study of grains and their interactions with each other. As Livermore computational geophysicist Eric Herbold explains, "The response of a granular material as a whole depends on how each grain interacts with its neighbor."

Although the physics community has long studied granular materials under external compaction, a Laboratory collaboration has helped produce the first-ever three-dimensional (3D) measurements of interparticle forces in a granular material during external loading and unloading. Livermore postdoctoral researcher Ryan Hurley teamed up with Stephen Hall of Lund University in Sweden to image quartz grains subjected to external pressure. Their experiment used 3D x-ray diffraction (3DXRD) and x-ray computed tomography (XRCT) at the European Synchrotron Radiation Facility (ESRF) in Grenoble, France. Using a novel analysis of these measurements, Hurley led an effort to track the evolution of interparticle forces and the relationship between an individual grain's behavior and external pressure. Building on this success, Hurley and Herbold secured support from Livermore's Laboratory Directed Research and Development (LDRD) Program for new experiments that will focus on observing and detecting cracking, deformation, and other forms of material failure caused by external pressure. "Improving the current understanding of failure in granular materials will help

us scale up models of localized strength and failure to real-world scenarios," says Hurley.

Breaking New Ground

The Livermore team's experiments follow a specific protocol. A force-loading device grasps the top and bottom of a hollow cylinder that is roughly as wide as a nickel is thick and is packed full of tiny spherical grains. Where each grain meets another is known as a grain boundary. Multiple adjacent boundaries make up a contact network, through which the interaction of grains affects the overall material's physical properties. When the collection of grains is compressed, forces move through the contact network in paths called force chains and can create conditions for failure. This experimental setup enables scientists to examine force chains and contact networks around a fracture as well as the effects of these phenomena on material properties such as strength and stability.

In an initial experiment, Hurley and Hall packed a hollow quartz-glass cylinder with 77 grains of quartz, each grain measuring between 220 and 310 micrometers in diameter. Under the guidance of ESRF's Jonathan Wright, the team used the facility's diagnostic technology to record what happened inside the cylinder as the piston-driven device slowly applied vertical pressure to the sample. At regular intervals over the course of 70 minutes, the team initially held the pressure constant while illuminating the sample with a 1.5-millimeter-wide beam of x-ray radiation and slowly rotating the sample 180 degrees for 3DXRD and radiography measurements. These measurements were performed 24 times at different external pressures. After the experiment, XRCT algorithms converted radiography images into 3D renderings, and 3DXRD elucidated stress and strain within each grain. Special numerical algorithms then determined forces at each grain boundary from these data.

The collection of grains in the tightly packed cylinder mimics naturally occurring materials such as compacted soil. Studies have shown interparticle forces to move unevenly through these systems—a condition known as heterogeneity. However, the 3D quartz experiment revealed an interesting evolution—as the pressure increased, patterns in the force chains showed a homogeneous trend as they assembled and strengthened. Buttressed by the cylinder's walls, grains on the perimeter joined vertically connected force

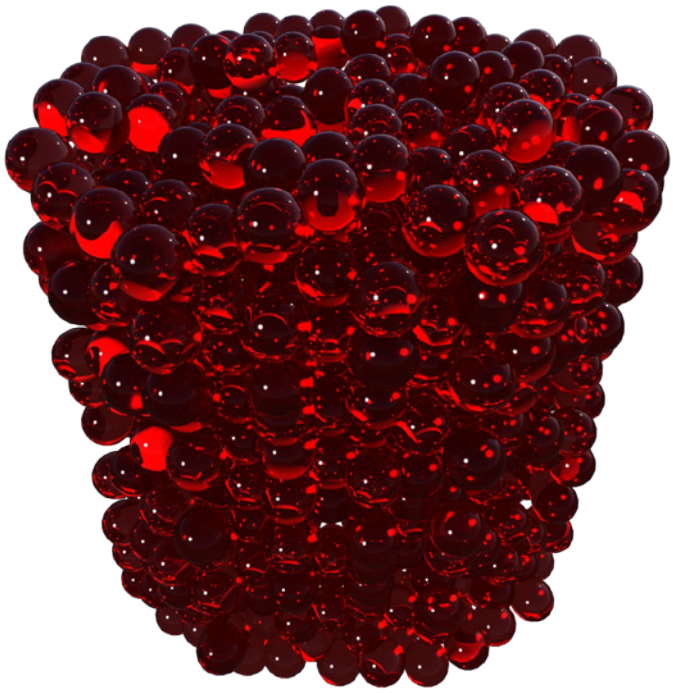

A rendering depicts an aluminum cylinder holding micrometer-scale grains for compression experiments. The entire cluster of grains measures less than 1.5 millimeters in any dimension. This preliminary experiment used a target containing 1,100 ruby grains. (Renderings by Adam Connell.)

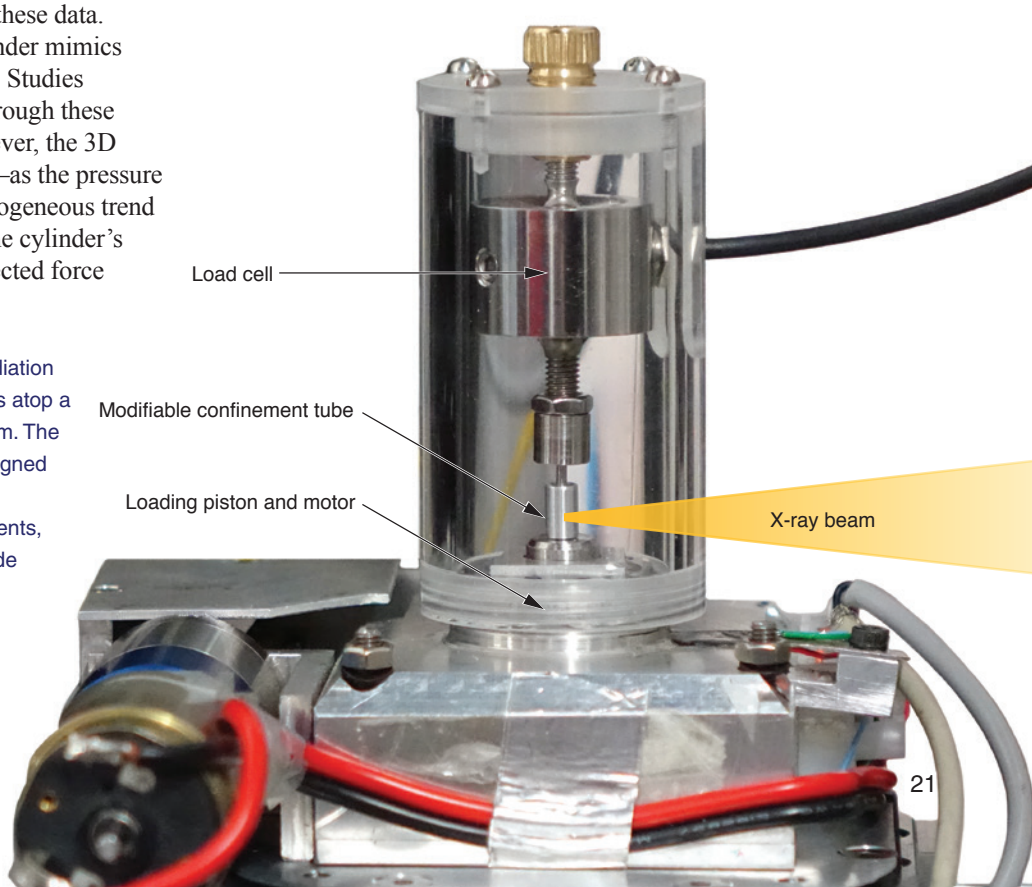

In this photograph, taken at the European Synchrotron Radiation Facility, a loading device used to test granular materials sits atop a mounting plate aligned with a 1.5-millimeter-wide x-ray beam. The device, fabricated by a Lund University collaborator, is designed to hold a cylinder target in place through a load cell while a piston applies pressure from below. For upcoming experiments, Livermore scientists Ryan Hurley and Eric Herbold will guide customization of the confinement tube and development of additional acoustic and ultrasonic technology.

chains, whose distribution shifted and became more uniform as the pressure continued toward a peak near 66 newtons, or about 15 pounds of force. Then, as the pressure ramped down and interparticle forces transferred to different grains, the system became more heterogeneous. This progression from homogeneity to heterogeneity suggests that, depending on external pressure, a different portion of grains may actually furnish a material's stiffness. Hurley notes that this heterogeneity also plays a role in strength and failure. He says, "A fractured grain may not be able to continue supporting the load, so the load shifts elsewhere."

When Failure Means Success

Before kicking off the LDRD project, Hurley again worked with Hall to conduct another experiment at the ESRF. The objective was to repeat the same process used for quartz with another material and more grains, compiling a preliminary data set to feed into subsequent experiments and analyses. This experiment used a similar loading protocol but a different type of sample—an aluminum cylinder full of 1,100 ruby crystals. "The friction and stiffness properties of quartz and ruby are on the order of what we find in natural sands and powders," explains Hurley. In addition, ruby grains are easier to find in small shapes, so more of them fit into the same cylinder volume. In this experiment, 3DXRD and XRCT measurements revealed that more grains cracked and fragmented as the load increased, providing valuable data about where failures occur and how force chains rearrange around those failures.

Beginning this spring, the Livermore team will embark on two new experiments with the continued engagement of Hall. The first will again test a ruby sample but apply compression on three axes. A cylinder made of a deformable membrane will allow the team to control the stress on all sides of the material. This so-called triaxial experiment is expected to yield data about strain patterns and the formation of localized deformation zones called shear bands. Shear localization is a key failure mechanism of granular materials, yet scientists do not fully understand what triggers shear bands or what effects they have on material properties after formation. For the second experiment, the team will test nonspherical grains. Spherical grains, such as those often found in additive manufacturing and pharmaceutical powders, pack in regular and predictable ways, but irregularly shaped grains should produce more complex force chains, strain evolution, and failure mechanics. This irregularity is found in many natural soils, polycrystalline solids, and amorphous metals.

Supplemental measurements may reveal aspects of failure processes elusive to other diagnostics. In particular, the Livermore team has the opportunity to use acoustic receivers to detect vibrations produced when grains fracture or a load changes its path. In addition, an ultrasonic transducer can send high-frequency waves through a granular material to show how waves travel along force chains. According to Herbold, measuring sound speeds across the acoustic transmission spectrum can help scientists interpret seismic activity moving through the ground—not only activity during an earthquake but also from an illicit nuclear test. Hurley notes additional applications in structural monitoring, soil testing, and material design. "Perhaps we'll see behavior we haven't seen before," he remarks.

The final phase of the LDRD project will involve calibrating Livermore's multiscale GEODYN-L code with accumulated data. Run on massively parallel computing systems, GEODYN-L lets researchers model wave propagation through fractured rock. Applications range from monitoring near-field explosives to direct metal laser sintering for additive manufacturing to calculating the impact effects of asteroid deflection. Once augmented with the

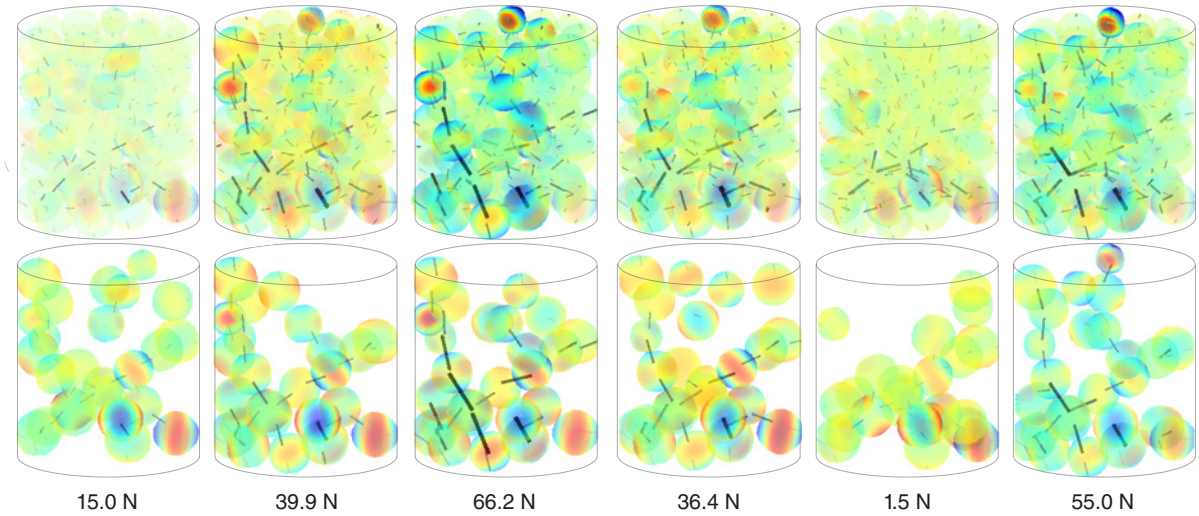

(top) Force inference results are shown for 6 of 24 load steps in an experiment, in which the peak was approximately 66 newtons (N). Strain among quartz grains is depicted along a color spectrum, with highest compression levels in blue and highest tension levels in red. (bottom) Prominent force chains along particle contact networks in the cylinder are shown, with isolated grains indicating where force chains diverge.

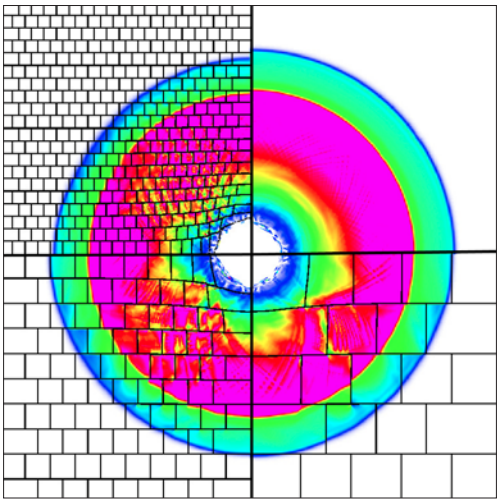

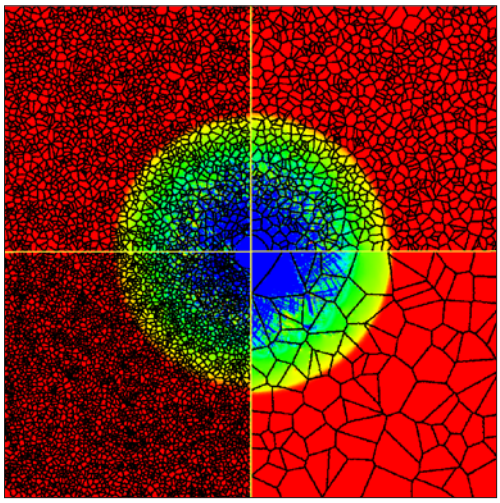

As a massively parallel physics code run on supercomputing systems, GEODYN-L has the capability to simulate nonlinear wave propagation through solid rock. In the two simulations shown, the changing intensity of a wave propagating from the center of a material sample is depicted as a colored spectrum within a matrix of solid particles (black-outlined shapes). Each simulation has quadrants of different particle sizes (or, in the left simulation, three sizes and a continuum in the upper right), demonstrating how particle size can be varied according to the material simulated. Data from the Livermore team's upcoming experiments will help improve the accuracy of frictional and tensile variables in the code.

quartz and ruby data, GEODYN-L will be able to predict further results with the team's new technologies. Herbold describes a complementary relationship between simulations and experiments. "Our models can include all of the nonlinear interactions between grains, but we don't know how the combined effects result in material response," he says. The team will perform sensitivity studies to analyze those responses and refine the code. In this way, modeling helps close gaps in experimental data, strengthening GEODYN-L's role as one of the Laboratory's valuable predictive capabilities.

Making Sense of the Ordinary

Lawrence Livermore's investigations of grains and their behavior go back many years, including such efforts as simulating powders under dynamic loading (see *S&TR*, June 2010, pp. 11–13) and modeling the microstructural effects of radiation on grain-boundary networks (see *S&TR*, March 2016, pp. 4–12). Multiple microscopy techniques have been used to characterize grains and grain boundaries in metals through a process called crystal orientation mapping (see *S&TR*, December 2014, pp. 16–19). Meanwhile, another LDRD project is measuring granular materials' responses to shock waves to improve the Laboratory's computational capability in shock physics.

Despite these advancements, making sense of the collected data presents a hurdle. Published studies use statistics and probability to quantify force chains and define force and grain boundary distributions. Beyond these, no standard metrics exist for quantifying the link between force chains and macroscopic behavior in granular materials. "The biggest challenge with these materials is that no first-principles models exist. We have to come up with the methods of quantification on our own," notes Hurley. "For instance, it's not clear what aspect of a contact network will be most important for predicting a particular material property, so we're exploring new ways of analyzing the data." Hall introduced the Livermore team to the Gini coefficient, which is normally used in economics to measure income inequality. In physics, the coefficient's representation of inequality within a frequency distribution helps describe force heterogeneity in various locations throughout a material sample. By calculating trends in vertically oriented forces in relation to exterior force from the piston, the Gini coefficient enables the team to generate numerical measurements of force homogeneity and heterogeneity.

Knowing how and when granular media fail supports the Laboratory's overall national security mission. From a geoscience perspective, grain behavior under stress has implications for load-bearing sand and soil infrastructure, such as in landslide and earthquake mitigation or penetration for hydraulic fracturing and recovering buried munitions. Herbold adds, "Predicting the mechanical response of a granular material can help our military design better runways in the field or identify suitable locations for underground structures." The research's relevance even extends to pharmaceutical pill packing and the storage of agricultural products. Observes Hurley, "Granular materials are everywhere."

—Holly Auten

Key Words: European Synchrotron Radiation Facility (ESRF), force chain, GEODYN-L code, Gini coefficient, granular material, Laboratory Directed Research and Development (LDRD) Program, three-dimensional x-ray diffraction (3DXRD), x-ray computed tomography (XRCT).

For more information contact Ryan Hurley (925) 422-1069 (hurley10@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. For the full text of a patent, enter the seven-digit number in the search box at the U.S. Patent and Trademark Office's website (http://www.uspto.gov).

Patents

Isotope-specific Arbitrary Material Flow Meter
Christopher P. J. Barty, John C. Post, Edwin Jones
U.S. Patent 9,476,747 B2
October 25, 2016

Depositing Bulk or Micro-scale Electrodes
Kedar G. Shah, Satinderpall S. Pannu, Vanessa Tolosa, Angela C. Tooker, Heeral J. Sheth, Sarah H. Felix, Terri L. Delima
U.S. Patent 9,485,873 B2
November 1, 2016

Incorporating an Optical Waveguide into a Neural Interface
Vanessa Tolosa, Terri L. Delima, Sarah H. Felix, Satinderpall S. Pannu, Kedar G. Shah, Heeral Sheth, Angela C. Tooker
U.S. Patent 9,486,641 B2
November 8, 2016

Li-ion Battery Thermal Runaway Suppression System Using Microchannel Coolers and Refrigerant Injections
Todd M. Bandhauer, Joseph C. Farmer
U.S. Patent 9,490,507 B2
November 8, 2016

Awards

The **Department of Defense** awarded a team from **Lawrence Livermore** and the **University of California at Merced** a **three-year, $768,803 grant** to study breast cancer metastasis. The work will focus on the metastatic role of extracellular shed vesicles (ESVs), packets of genetic information and other molecules long known to be involved in communication among cells and the performing of essential bodily functions. Recent studies suggest that tumors emit ESVs to prepare new sites in the body where cells separated from the tumor can land and grow into new tumors. Using accelerator mass spectrometry, the researchers hope to identify ESVs that act as biomarkers to better discriminate between malignant and benign tumors, thereby identifying dangerous tumors quickly and preventing overtreatment. The project could also lead to investigations of preventing metastasis by intercepting ESVs or stopping them from being released altogether.

The **Nuclear and Plasma Sciences Society** of the **Institute of Electrical and Electronics Engineers** selected Livermore researcher **Wayne Meier** for its **2016 Fusion Technology Award**. The award recognizes Meier's 40-year career of research and leadership advancing the science, technology, and integrated assessment of fusion power plants. During his career he contributed to studies on fusion chamber design, nuclear analysis, and systems integration and analysis, and authored or co-authored more than 150 technical papers. Meier retired in September 2016 as the deputy of Livermore's Fusion Energy Sciences Program and was notified of the award two weeks after retiring. The Fusion Technology Award recognizes outstanding contributions to research and development in the field of fusion technology and has been presented since 1989.

The **Federal Laboratory Consortium**'s (FLC's) **Far West Region competition** honored two Livermore teams for efforts related to technology transfer. FLC is a nationwide network of federal laboratories and provides a forum to develop strategies and opportunities for linking laboratory mission technologies with the marketplace.

An **Outstanding Technology Development Award** recognized the Wireless Battery Sensing and Failure Eliminator (WiBSAFE) technology, which Livermore researchers developed in collaboration with three companies. WiBSAFE can wirelessly monitor high-capacity power cells and is especially useful for cars, planes, spacecraft, and other environments where avoiding battery-related problems is crucial. Livermore personnel who worked on WiBSAFE include **John Chang**, **James Zumstein**, **Marianne Ammendolia**, **Jack Kotovsky**, **Todd Bandhauer** (who has since left the Laboratory), **Noel Peterson**, and **Joe Farmer**.

An **Outstanding Partnership Award** was presented to Livermore and two companies that licensed the Laboratory's geothermal silica extraction technology. A process for removing marketable colloidal and precipitated silica from geothermal fluids while coproducing potable water, the technology makes geothermal energy more competitive by offsetting production costs by about one cent per kilowatt-hour. The partnership team includes **Annemarie Meike** of Livermore's Industrial Partnerships Office.

Four Livermore scientists have been selected as **2016 fellows of the American Physical Society** (APS). Physicist **Adam Bernstein** was cited in the nuclear physics category for "pioneering work at the intersection of nuclear science and nuclear nonproliferation, including the development of antineutrino-based methods for monitoring the production of fissile material and large-volume detectors for rapid screening of cargo for the presence of fissile material." In plasma physics, physicist **Hui Chen** was recognized for "pioneering experimental research on relativistic positron generation using ultra-intense, short-pulse lasers." **Omar Hurricane**, chief scientist of the Inertial Confinement Fusion Program, was cited in the plasma physics category for "visionary leadership in experiments at the National Ignition Facility laser and innovative work in understanding instabilities in high-energy-density and inertial confinement fusion plasmas leading to the first laboratory demonstration of an alpha-heating-dominated, thermonuclear plasma producing a fusion energy exceeding its total stored energy." **James Trebes**, head of the Physics Division, was acknowledged by the APS Forum on Physics and Society for "contributions in laser physics and the application of physics to other disciplines, for leadership in multiple national security areas, and for contributions to education in the sciences and engineering."

Abstract

A Center of Excellence Prepares for Sierra

The architecture for Lawrence Livermore's next Advanced Simulation and Computing (ASC) Program supercomputer, Sierra, represents a break with the design of many past ASC systems but anticipates future computing trends. For Livermore's complex science applications to perform efficiently on the new system, which will arrive in late 2017, application developers must, in many cases, significantly modify their existing applications. A Center of Excellence (COE) was established at the Laboratory in late 2014 to help coordinate and accelerate preparatory efforts, and to document problems and solutions unearthed. Through the COE, selected applications—those deemed critical for Laboratory missions and representative of the topical breadth of research at the institution—receive focused attention by application developers, hardware and systems software experts from vendor partners, and in-house experts in code optimization for next-generation systems. These project teams must balance modifications intended to prepare and optimize the applications for Sierra with those that will help the applications run well on a range of computing architectures, as many of these applications will run on multiple systems at multiple laboratories. By making the applications more "portable," developers also help to prepare the codes for the following generation of supercomputers—exascale systems—and their still-to-be-determined architectures.

Contact: Rob Neely (925) 423-4243 (neely4@llnl.gov).

U.S. Government Printing Office: 2017/770-075-10025

Investing in the Nation's Future

The Laboratory Directed Research and Development (LDRD) Program has been driving scientific discovery for 25 years.

Also in April/May

- *Three-dimensional radioisotope batteries developed at Livermore provide a long-lasting, high-energy, reliable power source for remote applications.*
- *A team investigates surface-bound contamination and damage to high-power laser optics at the National Ignition Facility.*
- *Projects in the LDRD Laboratory-Wide category allow Livermore scientists to pursue innovative research paths.*

www.llnl.gov

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.