Lawrence Livermore National Laboratory

Science & Technology REVIEW

February 2021

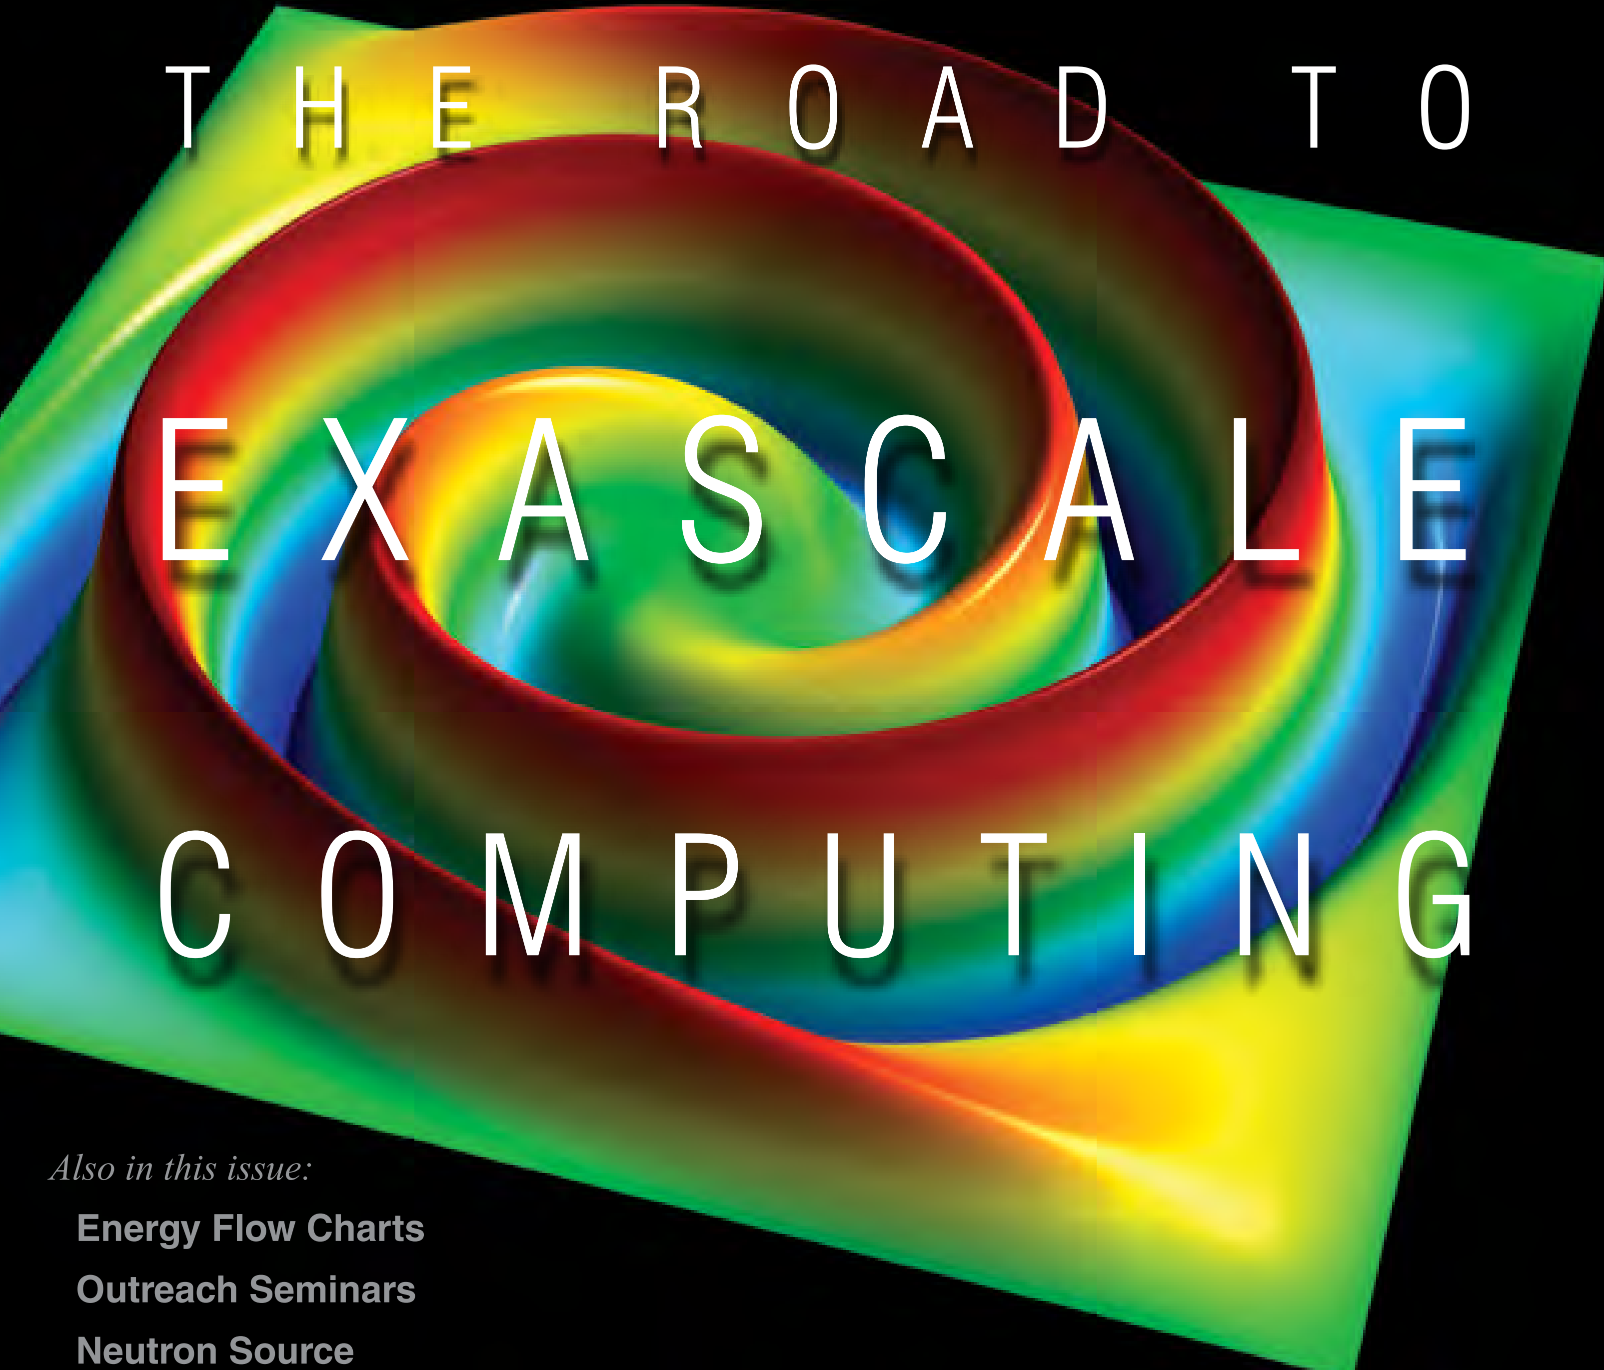

THE ROAD TO EXASCALE COMPUTING

Also in this issue:

Energy Flow Charts

Outreach Seminars

Neutron Source

About the Cover

The Department of Energy's Exascale Computing Project (ECP) and Lawrence Livermore's RADIUSS (Rapid Application Development via an Institutional Universal Software Stack) initiative benefit from strategically developed software tools. The front cover shows a simulation of advection under twisting rotation that uses high-order finite elements from Livermore's Modular Finite Element Methods (MFEM) software library and GLVis visualization tool. On the back cover, the logo (also created with GLVis) for the MFEM project illustrates the curved mesh and sub-element resolution used in high-order simulations. MFEM and GLVis are key components of the ECP's co-design Center for Efficient Exascale Discretizations (CEED) and RADIUSS. MFEM is also part of ECP's Extreme-Scale Scientific Software Development Kit (xSDK).

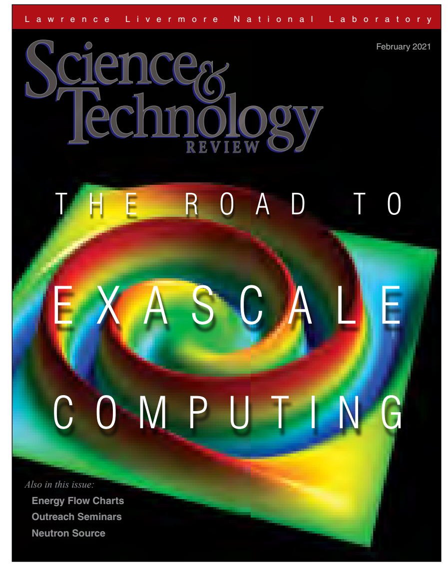

Cover design: Mark Gartland

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 managed by Lawrence Livermore National Security, LLC (LLNS), for the National Nuclear Security Administration (NNSA), a semi-autonomous agency within the U.S. Department of Energy (DOE). LLNS is a limited liability company managed by Bechtel National, Inc.; the University of California; BWXT Government Group, Inc.; and AECOM Energy & Construction, Inc. Battelle Memorial Institute also participates in LLNS as a teaming subcontractor. Cutting-edge science is enhanced through the expertise of the University of California and its 10 campuses and LLNS' affiliation with the Texas A&M University system. 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.

NNSA National Nuclear Security Administration

LAWRENCE LIVERMORE NATIONAL SECURITY LLC

Prepared by LLNL under contract DE-AC52-07NA27344

S&TR Staff

Scientific Editors
Dawn A. Shaughnessy
Holly Carlton

Managing Editor
Ken Chinn

Publication Editor
Suzanne Storar

Writers
Holly Auten, Ann Parker, J.C. Ross, Genevieve Sexton, and Suzanne Storar

Art Director
Mark Gartland

Proofreader
Deanna Willis

***S&TR* Online**
Lauren Casonhua, Rose Hansen, and Pam Davis Williams

Print Coordinator
Chris Brown

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 online 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-52000-21-02
Distribution Category UC-99
February 2021

Science & Technology Review

February 2021

Lawrence Livermore National Laboratory

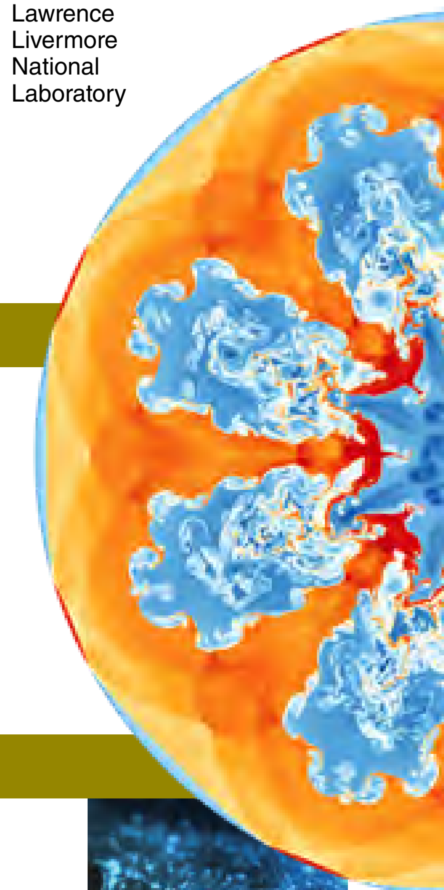

Contents

Feature

Research Highlights

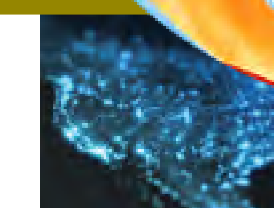

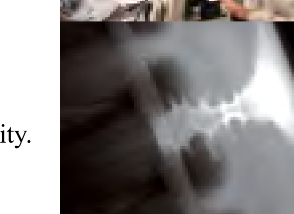

Departments

Extracting Rare-Earth Elements from E-Waste

A new process, based on a naturally occurring protein, could extract and purify rare-earth elements (REEs) from low-grade sources, outperforming human-made chelators and offering a new avenue toward a more diversified and sustainable REEs sector. Designed by researchers from Lawrence Livermore National Laboratory, in collaboration with Pennsylvania State University and Idaho National Laboratory, the process uses a protein, lanmodulin, that enables a one-step extraction and purification of REEs from complex metal mixtures, including electronic waste and coal byproducts.

"Lanmodulin has several unique and exciting properties," says Livermore researcher Gauthier Deblonde, lead author of the paper that appears in the July 20, 2020, issue of *Inorganic Chemistry*. "We were amazed to discover that a natural protein can be so efficient for metal extraction. This protein is the most REEs-selective macromolecule characterized to date and is able to tolerate industrially relevant conditions such as low pH, high temperature, and molar amounts of competing ions."

Current chemical processes to extract and purify REEs are complex and harmful to the environment. Extracting or recycling REEs from sources like electronic waste and coal byproducts, using natural products like lanmodulin, could be game-changing. REEs are essential for American competitiveness in the clean energy industry because they are used in many devices important to a high-tech economy and national security, including computer components, high-power magnets, wind turbines, mobile phones, solar panels, superconductors, hybrid/electric vehicle batteries, LCD screens, night vision goggles, and tunable microwave resonators.

Contact: Gauthier Deblonde, (925) 423-2068 (deblonde1@llnl.gov).

Exploring Earth's Chemical Origins

Massive compressive shearing forces generated by the tidal pull of Jupiter-like planets on their ice-covered moons may form a natural reactor that drives simple amino acids to polymerize into larger compounds. As reported in the July 27, 2020, cover article of *Chemical Science*, these extreme mechanical forces strongly enhance molecule condensation reactions, opening new possibilities for chemical origins of life on Earth and other rocky planets. "Compressive shearing forces are known to accelerate physical and chemical transformations in solid materials," says Livermore chemist Brad Steele, the study's lead author. "However, little is known about how these processes occur, especially for simple prebiotic molecules like amino acids, which can link."

The Livermore team focused on glycine, the simplest protein-forming amino acid and a known constituent of astrophysical icy bodies. To probe chemistry under such unusual conditions, the researchers developed a virtual rotational diamond anvil cell (RDAC) to enable rapid computational simulations of mechanically driven chemistry, or mechanochemistry, a relatively new field. RDACs add a shearing component to diamond anvil cell experiments compressing a sample between diamonds to access extremely high pressures (see illustration).

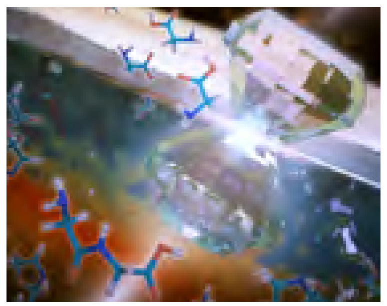
(Rendering by Veronica Chen)

The researchers determined that above a certain pressure, every shearing simulation predicted the formation of large polymeric molecules from the polypeptide glycylglycine to cyclic molecules and others with chiral centers. "Our study revealed a surprisingly complex chemistry coming from such a simple molecule," said Livermore scientist Will Kuo, one of the study's authors. The work points to compressive shearing forces as a potential driver for new and unusual chemistries in organic materials.

Contact: Brad Steele, (925)422-9902 (steele26@llnl.gov).

Simulating Bay Area Quakes at High Resolution

A Lawrence Livermore team has published new supercomputer simulations of a magnitude 7.0 earthquake on the Hayward Fault. Their work, reaching band frequencies of up to 10 Hz, presents the highest-ever resolution ground motion simulations from such an event on this scale. The study, published in the August 11, 2020, edition of the *Bulletin of the Seismological Society of America*, used the SW4 code developed at the Laboratory. Seismic waves as short as 50 meters were resolved across a region covering the San Francisco Bay Area—from Napa to San Jose and from the Sacramento-San Joaquin Delta to the Pacific Ocean. Previous simulations lacked the performance and memory to model such high-frequency motions on such a large domain. Calculations were made on the IBM Sierra supercomputer featuring NVIDIA graphics processing units.

Additional analysis accounted for the effects of soft soils that cover urbanized areas of the Bay Area. Seismologist Arthur Rodgers, lead author of the study, says, "Soft soils deform nonlinearly when subjected to strong shaking. We used a recently developed empirical model to correct ground motions for the effects of soft soils not included in the Sierra calculations. These improved the realism of the simulated shaking intensities."

High-frequency shaking is critical for evaluating seismic hazards and damage risk to buildings, homes, transportation, and utilities. Supercomputer simulations allow scientists to estimate the time-varying, three-dimensional pattern of shaking for an earthquake of interest. As computing power increases, such simulations will become easier and more accessible to earthquake scientists and engineers.

Contact: Arthur Rodgers, (925) 423-5018 (rodgers7@llnl.gov).

To Exascale and Beyond

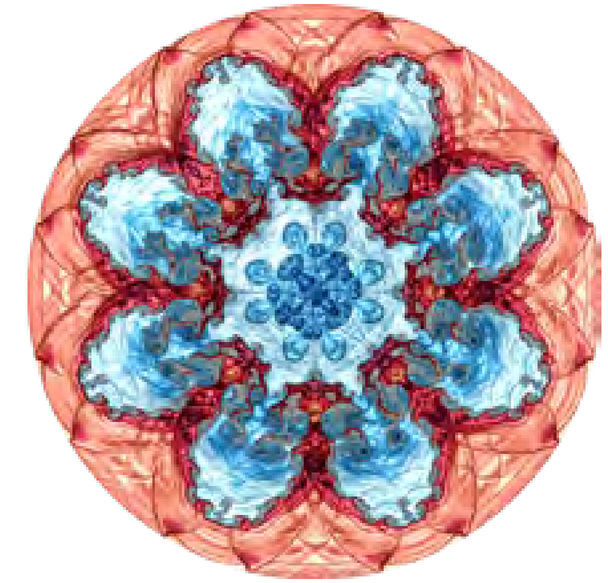

This multi-material inertial confinement fusion–like implosion was generated with Livermore's BLAST code, which relies on high-order mathematical software to produce simulations.

WORKING in the high-performance computing (HPC) field can feel like riding in a rocket. It takes you on an exhilarating, white-knuckle ride into uncharted space. You simultaneously appreciate the technology currently propelling you forward and wonder how far and fast the rocket can go. Pushing technological and scientific boundaries is how we explore new possibilities.

The next phase of computing technology is known as the exascale era, when computers will be able to process an exaflop—a quintillion (10^{18}) calculations per second. With a vast increase over existing systems in compute power, storage, and memory, exascale computing will allow us to peer more deeply into the inner workings of physical systems. We will uncover scientific surprises and learn more than ever before. Such insights will help advance Lawrence Livermore's missions.

This potential drives our participation in the Department of Energy's Exascale Computing Project (ECP), in which Livermore staff hold key leadership positions and participate in innumerable projects. As the feature article beginning on p. 4 explains, the ECP is forging a path toward a comprehensive and reliable exascale computing ecosystem that will benefit research in national security, foundational science, healthcare, energy, and other areas.

Hardware commands most of the attention in discussions about HPC and exascale capabilities. Indeed, supercomputers are tangible, visible assets that enable researchers to study scientific processes in detail. Livermore has a track record of standing up uniquely capable systems, and we are fortunate to house and manage some of the world's most powerful machines. For example, Sierra and Lassen are our largest heterogeneous systems, which use both central processing units (CPUs) and graphics processing units (GPUs) to perform calculations more quickly and efficiently than their predecessors. Leveraging GPUs is one step toward a more economical architecture that addresses the challenges in transistor technology associated with the ending of Moore's Law and Dennard scaling. Debuting in just a few years, our first exascale system, El Capitan, will consume considerably less electrical power (around 30 megawatts) with GPUs than it would without (at least 90 megawatts).

As HPC systems become more complicated, so too does the programming of applications that run on them. We constantly make tradeoffs between hardware decisions and software design, and in some ways, software is the more difficult part of achieving exascale computing. The lifespan of software must be considered well beyond a single project's timeline or any one machine. We saw in the transition from earlier systems to Sequoia's CPU-based architecture, and again to Sierra's hybrid GPU–CPU design that porting codes from one type of machine to another is a monumental task. We need to design software today that will still be relevant on HPC systems in 10 or even 20 years from now—even though we have limited clarity about what those systems might look like.

If we are able to design and build software that is robust against upcoming changes to computer architecture, we will have greatly reduced our mortgage for the future. Our very large software portfolio will only be able to run dependably on new machines if we embrace a development model of reusing components while providing a middle layer that insulates scientific applications from the details of the underlying hardware. The feature article describes many of our innovative software libraries and tools, including those developed under the RADIUSS (Rapid Application Development via an Institutional Universal Software Stack) initiative. The ECP and RADIUSS share this sustainability strategy, focusing on modular components and practical implementation to plan for an uncertain future.

The exascale threshold is often touted as the finish line in HPC, but it is merely a mileage marker in an ongoing competition with the rest of the world. Even as we prepare for El Capitan, we must continue innovating for the computing landscape beyond GPUs and exaflops. Creative changes in architecture will have major implications for how we build software. The next phase of computing technology is always on the horizon. No one knows yet what it will look like, but we at Lawrence Livermore aim to find out.

■ Bruce Hendrickson is the associate director for Computing.

THE EXASCALE SOFTWARE PORTFOLIO

The exascale era will introduce the world to supercomputers, like the Laboratory's forthcoming El Capitan, that perform calculations faster than ever before. As high-performance computing hardware becomes more complex, so too does software design and development in support of these advanced machines. Livermore's forward-looking software projects inspire innovation both inside and outside the Laboratory.

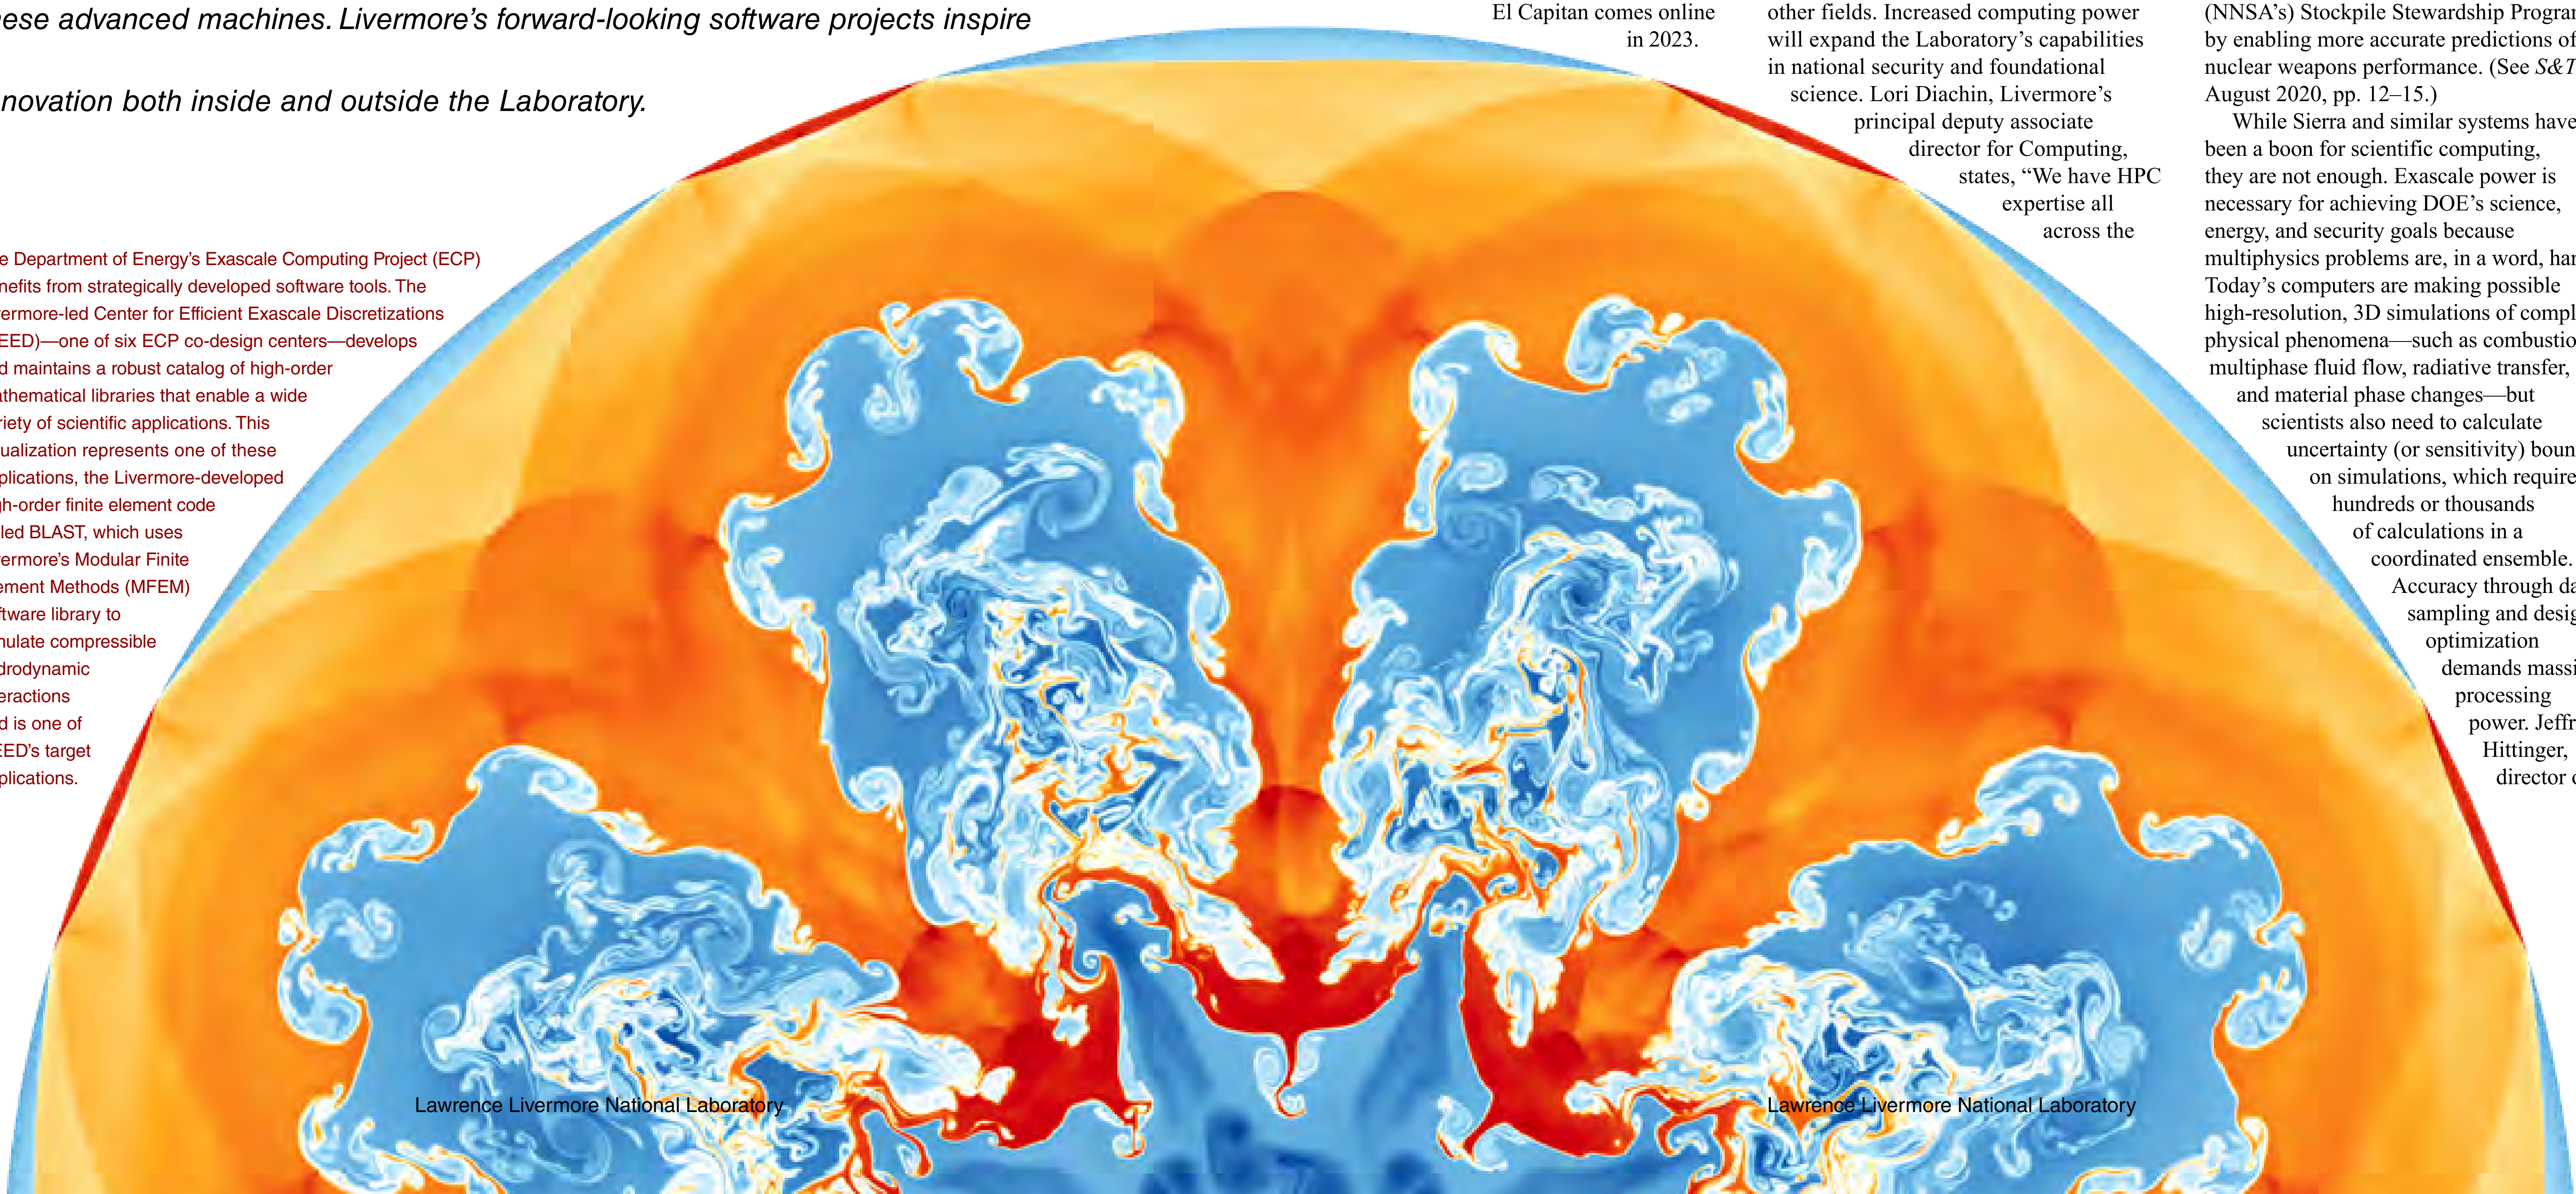

The Department of Energy's Exascale Computing Project (ECP) benefits from strategically developed software tools. The Livermore-led Center for Efficient Exascale Discretizations (CEED)—one of six ECP co-design centers—develops and maintains a robust catalog of high-order mathematical libraries that enable a wide variety of scientific applications. This visualization represents one of these applications, the Livermore-developed high-order finite element code called BLAST, which uses Livermore's Modular Finite Element Methods (MFEM) software library to simulate compressible hydrodynamic interactions and is one of CEED's target applications.

EXASCALE supercomputers will process information a thousand times faster than the systems that introduced the possibilities of predictive simulation a decade ago. Lawrence Livermore will be among the world's first high-performance computing (HPC) centers to deploy an exascale-class system, capable of 10^{18} floating-point operations per second (flops), when El Capitan comes online in 2023.

The Laboratory wields a large portion of the Department of Energy's (DOE's) HPC resources. These computing investments—hardware technology, software infrastructure, and scientific applications—have aided discoveries in nuclear and high-energy-density physics, materials science, climate change, energy efficiency, biological processes, and many other fields. Increased computing power will expand the Laboratory's capabilities in national security and foundational science. Lori Diachin, Livermore's principal deputy associate director for Computing, states, "We have HPC expertise all across the Laboratory. Contemplating the kinds of science we'll be able to do on exascale machines is exciting."

The heart of this effort is the predictive capability that comes from modeling, simulation, and visualization. For example, Sierra, one of the world's most powerful supercomputers, supports the National Nuclear Security Administration's (NNSA's) Stockpile Stewardship Program by enabling more accurate predictions of nuclear weapons performance. (See *S&TR*, August 2020, pp. 12–15.)

While Sierra and similar systems have been a boon for scientific computing, they are not enough. Exascale power is necessary for achieving DOE's science, energy, and security goals because multiphysics problems are, in a word, hard. Today's computers are making possible high-resolution, 3D simulations of complex physical phenomena—such as combustion, multiphase fluid flow, radiative transfer, and material phase changes—but scientists also need to calculate uncertainty (or sensitivity) bounds on simulations, which requires hundreds or thousands of calculations in a coordinated ensemble. Accuracy through data sampling and design optimization demands massive processing power. Jeffrey Hittinger, director of

Livermore's Center for Applied Scientific Computing, explains, "Efficient exploration of the solution space is a huge hill to climb."

As the exascale era begins, two major initiatives leverage and expand Livermore's HPC capabilities. The spotlight in this feature is software. The Exascale Computing Project (ECP) brings together many national laboratories to address many of the challenges inherent in their scientific and national security missions. At Livermore, the RADIUSS project—Rapid Application Development via an Institutional Universal Software Stack—aims to benefit scientific applications through a robust software infrastructure.

The Exascale Threshold

The exascale threshold—one thousand times faster than petascale—is incredibly difficult to reach. Although supercomputers are becoming more powerful, hardware manufacturers are nevertheless approaching limitations of processor speed and chip size. For decades, machines were built with a large but limited number of processors and memory modules, and each new version offered more capability for the same price, energy cost, and footprint. Now, physical constraints have resulted in increasing costs and energy consumption for small gains in computing performance. This situation has inspired new designs with the potential to improve the performance of scientific applications—but not without innovation in software.

The introduction of graphics processing units (GPUs) into HPC systems has opened the door to new computational possibilities. For certain types of calculations and applications, GPUs consume less energy and take up less space than central processing units (CPUs). Parallel processing capability thus increases, and a computer's workload can be balanced accordingly. Machine learning algorithms work well on GPUs, running faster at lower floating-point precision. Sierra's generation of computers is known as heterogeneous, or hybrid, because their architectures combine GPUs and CPUs. The DOE's first three exascale systems—El Capitan, Argonne National Laboratory's Aurora, and Oak Ridge National Laboratory's Frontier—will also take advantage of GPUs.

Diachin, who also serves as the ECP's deputy director, explains, "CPUs and GPUs are now physically closer together, but latencies exist due to the location of data in memory." During a calculation, data has to move near GPUs to take advantage of them, which means algorithms must be designed and executed with this in mind. Furthermore, HPC hardware varies by vendor, requiring customized software as well as interoperability solutions when switching platforms.

The HPC community quickly realized that monolithic code bases were no longer sustainable in this context, and software programming needed a new paradigm. Diachin states, "Transitioning to advanced architectures means investing in reusable software solutions that bridge the complexity between applications and diverse hardware. The problem involves a lot of exploration and hard work." Hittinger adds, "Heterogeneous computing architectures are complicated, especially at extreme scales. Multiphysics codes require significant computational resources and an HPC ecosystem that includes software libraries and tools designed to work on these machines. We are simplifying computer programming and improving software quality so scientists can focus on science."

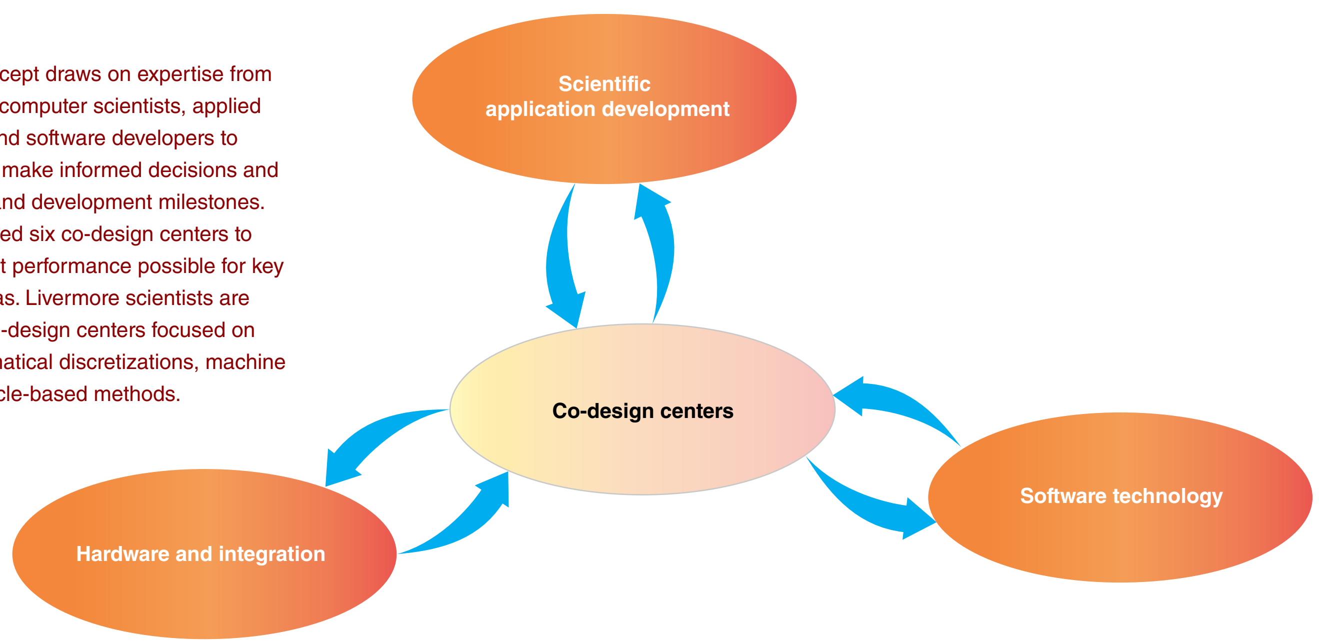

The co-design concept draws on expertise from domain scientists, computer scientists, applied mathematicians, and software developers to help organizations make informed decisions and achieve research and development milestones. The ECP established six co-design centers to achieve the highest performance possible for key computational areas. Livermore scientists are involved in ECP co-design centers focused on high-order mathematical discretizations, machine learning, and particle-based methods.

A Sum Greater than the Parts

The ECP launched in 2016 as the U.S. government was already looking toward procuring exascale-capable computers. The 7-year, $1.8 billion effort is funded by the DOE's Office of Science and NNSA and includes most national laboratories and approximately 1,000 researchers. According to Diachin, Livermore was a natural fit for the project. "Our laboratory has a strong reputation for fielding world-class HPC systems," she says. "We bring decades of computing experience to bear on the most challenging scientific problems." (See *S&TR*, September 2016, pp. 4–11.) Like Diachin, several Livermore researchers hold ECP leadership roles.

All ECP research and development activities revolve around the delivery of a sustainable exascale computing ecosystem that supports mission-critical applications. By enabling higher fidelity solutions to scientific problems, the ECP aims to advance scientific discovery, strengthen national security, and improve industry competitiveness. Diachin states, "Collectively, we are creating tools and capabilities that the individual players would not otherwise be able to create."

The project is organized into three focus areas: application development, software technologies, and hardware and integration. Scientific application development is the ECP's top priority, with two dozen teams working to demonstrate simulation capabilities at a large scale. For example, a Livermore-led team is refining multiple codes that simulate physics processes relevant to stockpile stewardship. Software technology teams concentrate on the underlying software infrastructure that helps applications run accurately, quickly, and reliably. Meanwhile, hardware and integration teams work with industry vendors and HPC facilities on power-efficient and affordable HPC designs, testbed support, and ECP software deployment logistics.

The software development effort includes evaluating the tools and features that will reduce exascale development costs. For instance, a flexible exascale ecosystem will need to integrate applications composed of independently developed parts, each with its own programming language and parallel programming model. Similarly, data management and visualization tools are essential to collecting, analyzing, moving, and storing simulation data. "We will accomplish exascale computing by working together on common solutions," notes Livermore computer scientist Rob Neely, deputy program director for Weapon Simulation and Computing, who oversaw ECP software technology developed within NNSA's Advanced Simulation and Computing (ASC) portfolio and at NNSA laboratories.

Coordinated Development

One of the ECP's most critical tasks is effective collaboration to avoid redundancy and ensure interoperability of the ecosystem's components. This coordinated development coalesces under the concept of *co-design*, which draws on expertise from domain scientists, computer scientists, applied mathematicians, and software developers. "Co-design centers centralize the effort on commonly recurring algorithms by working with multiple application teams to provide highly optimized solutions," states Tzanio Kolev, who leads the ECP's co-design Center for Efficient Exascale Discretizations (CEED).

Livermore researchers lead or participate in three of the ECP's six

Livermore's Lassen (above) and Sierra supercomputers exploit graphics processing units (GPUs) for increased computing power. With a peak performance topping 20 petaflops, the Lassen system is a smaller, unclassified version of Sierra. (Photo by Garry McLeod.)

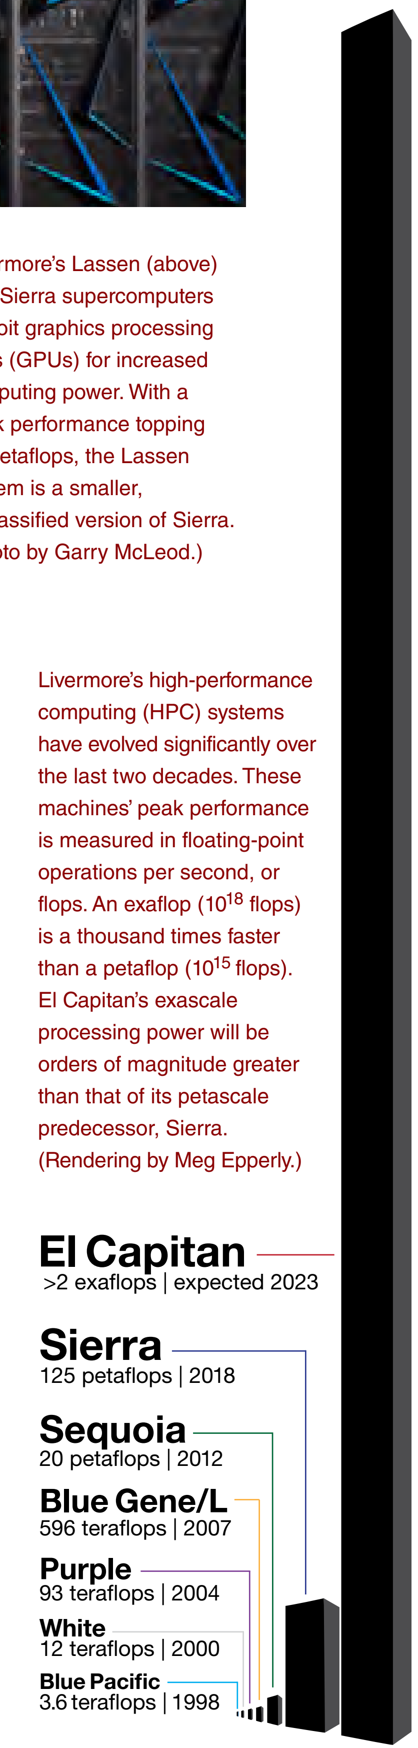

Livermore's high-performance computing (HPC) systems have evolved significantly over the last two decades. These machines' peak performance is measured in floating-point operations per second, or flops. An exaflop (10^{18} flops) is a thousand times faster than a petaflop (10^{15} flops). El Capitan's exascale processing power will be orders of magnitude greater than that of its petascale predecessor, Sierra. (Rendering by Meg Epperly.)

co-design centers. These interdisciplinary centers are organized around key scientific computing subjects, and their work helps achieve the ECP's research and development milestones. For example, CEED's purview is improving computational accuracy and efficiency of simulations via finely calibrated mathematical discretization libraries. (See the box below.) The ExaLearn co-design center is developing a scalable machine learning and artificial intelligence software framework for use in scientific applications and at experimental facilities. The Co-design Center for Particle Applications specializes in particle-based simulations of molecular dynamics and other particle interactions.

Mathematical Foundations

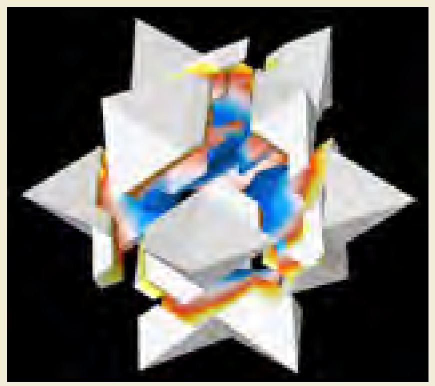

Versatile, high-order math libraries give scientific applications a boost in performance and accuracy when run on HPC systems. The MFEM software library and GLVis visualization tool produced this image of a heat diffusion simulation on a 3D unstructured tetrahedral mesh and its parallel decomposition.

Large-scale, complex scientific applications cannot run on new computing architectures without a foundation of rigorous and efficient mathematical solutions, such as high-order discretization methods. Computational mathematician Tzanio Kolev explains, "Scientific applications in any field of study must incorporate robust mathematical calculations that are accurate and predictive." For example, finite element numerical methods naturally describe scientific phenomena relevant to Laboratory missions, such as compressible fluid flow, heat transfer, design optimization, and additive manufacturing. These methods provide efficient solvers for partial differential equations, which define many real-world processes in a rigorous mathematical form. The process of discretization transforms continuous mathematical functions into discrete components, making those functions understandable to a computer. Cumulatively, these sophisticated techniques exploit a supercomputer's data parallelism and memory access, improving performance by orders of magnitude over traditional low-order methods.

In the exascale era, researchers must modify such algorithms so scientific codes can make the most of graphics processing units (GPUs). Accordingly, math libraries play a significant role in the Exascale Computing Project's (ECP's) software portfolio. High-order solution algorithms are well suited for GPU-based architectures but difficult to execute. Kolev states, "Controlling the arithmetic intensity and ensuring the accuracy of these methods is a mathematically big challenge. We are working on different ways to accomplish this, such as by developing novel matrix-free algorithms and solvers." One of the ECP's co-design centers, the Center for Efficient Exascale Discretizations (CEED), is dedicated to making high-order methods as practical and efficient as possible so scientists do not need to reinvent or optimize these parts of their code.

Led by Kolev, CEED combines experts from Lawrence Livermore and Argonne national laboratories and five universities. The Center's goals include developing a comprehensive software suite of libraries, solvers, application programming interfaces, and programming models; improving the tools that transition low-order applications to using high-order methods; and defining community standards such as format specifications for high-order data and operators. Since CEED's inception, the team has published nearly 50 scientific papers and given more than a dozen presentations on these technologies.

Among the Center's projects are "mini-apps" that capture key physics properties and are used to benchmark scientific applications' performance. In true co-design spirit, ECP teams and vendors use mini-apps such as Laghos (Lagrangian High-Order Solver) to model compressible gas dynamics and fluid flow. Laghos solves ordinary differential equation systems through novel use of mass and force matrices, resulting in less data storage and fewer memory transfers.

Another important tool in CEED's software suite is the Modular Finite Element Methods (MFEM) library, which provides building blocks for developing finite element algorithms. Researchers use MFEM to run simulations on a wide variety of machines—from personal laptops to the largest GPU-powered supercomputers. MFEM's development under CEED directly benefits next-generation codes for the National Nuclear Security Administration's Advanced Simulation and Computing program.

With its version 4.0 release in 2019, MFEM leapt from a solution optimized for central processing units (CPUs) to one that also supports GPUs. The upgrades were tested on Lassen, one of Livermore's GPU-based supercomputers. A more recent incremental release offers optimized support for the specific type of GPUs that will be used in the DOE's first exascale systems. Kolev describes this effort as the most difficult task the CEED team has accomplished to date. "Imagine you replaced a car's internal combustion engine with an electric engine. From the driver's perspective, the car still runs, but everything is different under the hood," he says. GPUs are not as "smart" as CPUs, so MFEM's algorithms were revised to express tasks simply and independently. The library was also refactored to accommodate memory movement between GPUs and CPUs and to support many different hardware platforms while also remaining compatible with machines that use only CPUs. The result is a highly flexible math library that will be adaptable to exascale demands.

Software Sustainability

In a software development context, *sustainability* means ensuring that software remains relevant and useful, works correctly, and is regularly updated to meet users' evolving goals. "In the ECP, software sustainability includes support for more applications and larger scale machines. Software tools must be interoperable in this context," explains Ulrike Meier Yang, Mathematical Algorithms and Computing group leader at Livermore and head of the ECP's Extreme-Scale Scientific Software Development Kit (xSDK) effort.

Focusing on numerical libraries, the xSDK team works toward the seamless integration of software packages needed by ECP applications. According to Yang, math libraries are usually independently developed with their own software strategies on different platforms and built with different compilers, which can lead to a variety of issues when they are used in combination. She notes, "Our team achieves consistency across libraries through activities such as normalizing build processes and avoiding namespace conflicts. We vet new libraries for inclusion in the Kit and resolve any incompatibilities." The xSDK provides both the turnkey aggregation of a range of mathematical software as well as a set of community policies and documentation that encourage standardization.

So far, the xSDK contains 23 math libraries including several developed by Livermore scientists, such as the HYPRE library of high-performance preconditioners and solvers, the SUNDIALS collection of nonlinear and differential equation solvers, and the Modular Finite Element Methods (MFEM) library. Most scientific applications use various subsets of the full catalog, which are tested on different operating systems at ECP partner sites. "Interoperability means we serve a range of user scenarios," states Yang. "The xSDK software suite and its accompanying policies are valuable because they are versatile, reliable, and cost-effective."

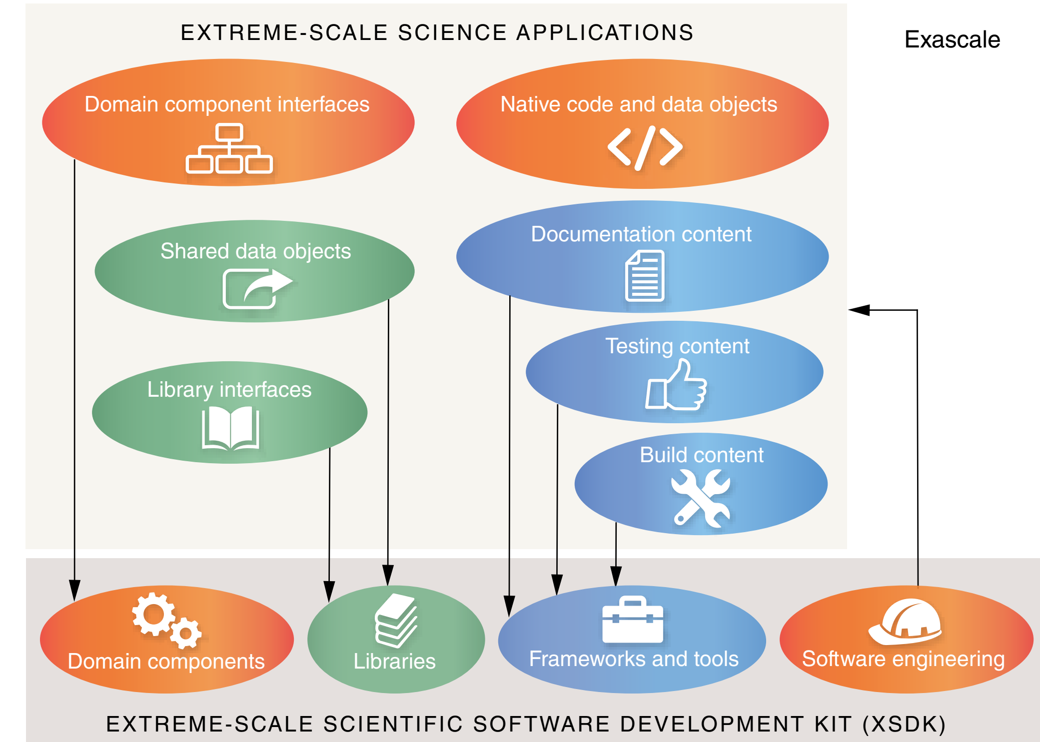

The ECP's Extreme-Scale Scientific Software Development Kit (xSDK) provides a standardized aggregation of math libraries for use by scientific applications on exascale-capable machines. The xSDK team's methodical approach to software sustainability helps the ECP deliver dependable software to scientific applications.

Institutional Improvements

Livermore's RADIUSS project benefits from the ECP's productivity and insights. Hittinger notes, "The ECP has invested a lot in software sustainability and maintenance. We don't want to lose that momentum at the Laboratory." Neely, who in addition to his ECP role serves as the RADIUSS project lead, adds, "RADIUSS is inspired by the ECP's efforts to make software highly dependable for users."

An additional motivator is NNSA's ASC program, which for years has funded applications that simulate and predict the performance and safety of nuclear weapons—as well as the software to execute these codes on supercomputers. According to Neely, this investment naturally dovetails with RADIUSS objectives. He says, "The ASC program has always recognized the need for software sustainability to ensure the most reliable simulations. RADIUSS leverages the tools developed via ASC, allowing us to build complicated software out of simpler, modular components."

The project aims to strengthen versatile HPC software and broaden its usage at Livermore and across the scientific application community. Neely explains, "Scientific computing is at an intersection of evolving architectures, increasingly complex simulations, and the need to change software approaches accordingly. Ultimately, we are advocating for adoption of Livermore's scalable, stable open-source software in the broader community. RADIUSS encompasses a production-quality set of tools for every scientific application developer to use, including users outside of the programs that fund development of these software products."

RADIUSS builds on expertise from computer scientists and software developers all over the Laboratory to encourage common development standards and provides another venue for developers to encourage adoption of their products. The team tackles software sustainability within the entire HPC ecosystem—computing infrastructure, code installation, application integration, data management and visualization, testing guidelines, documentation, and more. The project's software tools individually provide solutions for specific use cases and collectively offer a flexible "shopping list" for HPC users. (Many of these tools are also part of the ECP's software portfolio.) Neely emphasizes, "We are adopting, creating, and promoting best practices and a standard way of developing and releasing open-source software at the Laboratory."

For instance, performance and workflow optimization are key areas that RADIUSS promotes. Running a scientific application on a supercomputer is not as simple as clicking a button. Codes contain scripts that perform calculations, pull in or generate data, or execute other tasks. How these actions are completed depends on their interaction with interconnected computing nodes, and coordination of individual jobs that run during large-scale simulations is tricky. ASC software tools like Flux allow users to schedule and manage HPC resources, while the Caliper library lets users customize performance measurements for their applications.

The RADIUSS portfolio also includes portability and memory management tools that tackle hardware-related challenges. Application codes written for CPU-based systems need to work on newer GPU-based systems. Moreover, heterogeneous architectures vary, so codes also need to adapt to any available supercomputer. This challenge was anticipated in Livermore's weapons program a decade ago, motivating the development of tools like RAJA Portability Suite, which provides abstractions of calculation loops

Intelligent Memory Allocation

The scenario is familiar to high-performance computing centers that run large scientific codes on heterogeneous machines—limited memory resources require strategic memory management, but multiphysics codes and the hardware they run on can vary widely. Researchers must consider where data will be stored as the simulation is processed, and how best to move data to and from available compute nodes for optimal performance. Central processing units (CPUs) store more data, but graphics processing units (GPUs) are faster.

Lawrence Livermore's RADIUSS project—Rapid Application Development via an Institutional Universal Software Stack—addresses such execution challenges with a suite of advanced software tools that ultimately improve a code's performance. For instance, Umpire is a memory management solution developed by researchers David Beckingsale, Marty McFadden, Kristi Belcher, and Rich Hornung. Like an umpire makes decisions on a baseball diamond, Umpire determines how to allocate data among a supercomputer's complex memory resources and accommodates a range of device specifics and programming models. Principal investigator Beckingsale notes, "Instead of forcing users to commit to one technology-specific implementation, Umpire creates a memory resource for everything it detects on each system. Users do not have to know anything about the hardware or decide how to manage memory while their codes run."

Umpire works through an application programming interface (API) that abstracts and unifies memory allocation. The API allows for multiple complementary tasks such as querying compute nodes for availability, adjusting memory-pooling methods to speed up allocations, transferring simulation data between GPUs or between GPUs and CPUs, and tracking which data is stored where. "For large-scale physics applications, not enough memory exists on GPUs alone. Data must move around dynamically," explains Beckingsale.

Many of the Laboratory's production codes—from stockpile stewardship to seismic monitoring—rely on Umpire. It leverages other RADIUSS tools and can work alone or in tandem with Livermore's RAJA portability software, which helps move codes from one type of computing architecture to another. According to Beckingsale, the work to adapt codes to the current generation of GPU-based machines will pay off when the next generation arrives. He states, "We will take this same approach to memory management on exascale computers."

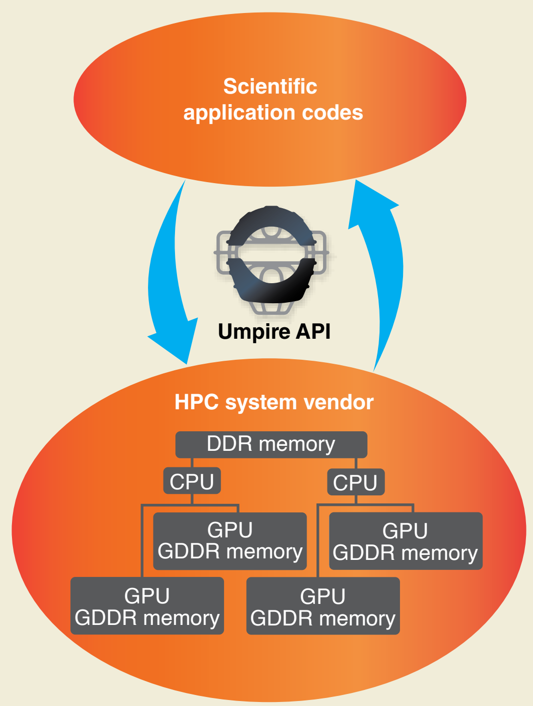

Like an umpire makes decisions on a baseball diamond, Livermore's Umpire software determines how to allocate a supercomputer's complex memory resources such as double data rate (DDR) and graphics double data rate (GDDR) integrated circuits. Multiphysics codes and the hardware they run on can vary widely, so Umpire accommodates a range of device specifics and programming models, ultimately improving a code's performance.

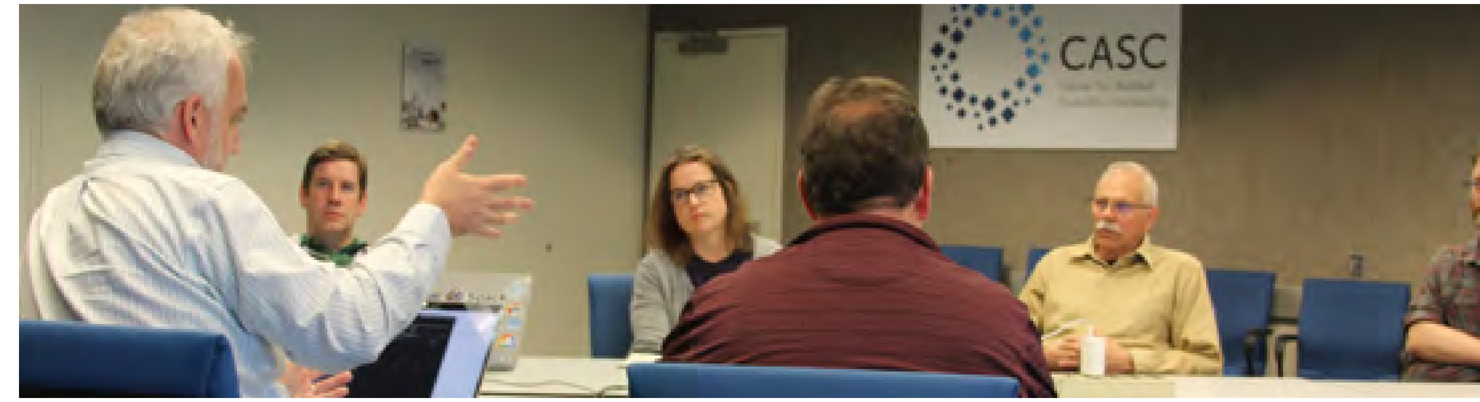

Rob Neely (far left) leads Livermore's RADIUSS team—Rapid Application Development via an Institutional Universal Software Stack—in tackling software sustainability within the entire HPC ecosystem. (Photo by Meg Epperly.)

to target machine-specific programming models and constructs, as well as memory allocation and movement decisions. (See *S&TR*, August 2020, pp. 12–15.) "Most people want to take advantage of the GPU revolution," says Neely. "RADIUSS is reducing overheads for application teams, providing a pathway to next-generation architectures, and building a knowledge repository of local expertise." (See the box on p. 10.)

Open for Collaboration

Though much of Livermore's scientific application portfolio is necessarily classified, a culture of openness in unclassified software development has taken root at the Laboratory. Open-source software (OSS)—the practice of releasing licensed code and inviting outside feedback and contributions—is valuable to many projects and essential when external collaborators are involved, as in the ECP. (See *S&TR*, January/February 2018, pp. 4–11.) For example, the xSDK team consists of developers at five laboratories and six universities, so their software efforts must be accessible to all participants. Hittinger emphasizes the importance of open-source development in projects large and small, stating, "More community input makes for better software. External contributors help us identify bugs, evolve features, and attain broader usage."

Although RADIUSS primarily benefits the Laboratory, Neely points out that input from the open-source community is constructive. He says, "All of the products under the RADIUSS banner are open source. We hope to share our work and experiences with, and learn from, other national laboratories and HPC centers." RADIUSS team member David Beckingsale adds, "Our software benefits greatly from engagement with the HPC community, vendors, and university collaborators. Public-facing development helps give users confidence in our projects and see that they are actively maintained."

OSS can quickly build momentum among users and developers, as recent successes with the Spack package manager and the Scalable Checkpoint/Restart framework have shown. Both of these Livermore-led OSS projects won 2019 R&D 100 Awards for innovation. (See *S&TR*, July 2020, pp. 8–11.) Both are also part of the ECP's software portfolio and promoted by RADIUSS. In fact, all software tools and libraries mentioned by name in this article are open source.

Beyond Exascale

As the exascale era dawns, Livermore researchers and software developers take a holistic view of the supercomputing landscape, where versatility and scalability are crucial to high performance—regardless of machine. Diachin states, "Interest in deploying the ECP's software is global, so our software must be compatible with many different computing architectures and be performance portable." Hittinger adds, "The best software stack insulates scientific codes against future architecture changes."

Exascale-capable systems like El Capitan will come online in the next few years. The first exaflop calculation will be run, and global supercomputer rankings will shift accordingly. However, Hittinger points out, "An exaflop is a milestone, not the finish line." A more meaningful moment will come when scientists can solve a problem with exascale computing capability that they could not have solved previously. He continues, "Computing constantly evolves. We won't simply stop at the next breakthrough."

—*Holly Auten*

Key Words: Advanced Simulation and Computing (ASC) program, application programming interface (API), Center for Efficient Exascale Discretizations (CEED), central processing unit (CPU), co-design, Department of Energy (DOE), El Capitan, exascale, Exascale Computing Project (ECP), Extreme-Scale Scientific Software Development Kit (xSDK), floating-point operations per second (flops), graphics processing unit (GPU), hardware, high-performance computing (HPC), Modular Finite Element Methods (MFEM), National Nuclear Security Administration (NNSA), open-source software (OSS), Rapid Application Development via an Institutional Universal Software Stack (RADIUSS), software sustainability, RAJA Portability Suite, Umpire.

For more information contact Lori Diachin (925) 422-7130 (diachin2@llnl.gov) or Rob Neely (925) 423-4243 (neely4@llnl.gov).

S&TR February 2021

CHARTING THE NATION'S **ENERGY USE**

ENERGY powers the nation, flowing from resources such as oil, the Sun, and natural gas to our homes, industries, businesses, and infrastructure in a complex, interconnected web. The multifaceted nature of U.S. energy supply and demand can yield complicated answers to simple questions: Which resources provide the most power? What sectors still rely on coal? Is geothermal energy on the rise? How much energy do we consume as a nation, and how efficiently is it used?

In the mid-1970s, Lawrence Livermore National Laboratory began developing a way to visualize answers to questions about energy use, creating easy-to-understand energy flow charts—one-page diagrams of U.S. resources and their consumption over a one-year period. Each flow chart draws upon immense data sets compiled by the Department of Energy's Energy Information Administration (EIA), presenting the information in an easily understandable form.

Nearly 50 years of diagrams chronicle the nation's energy use and how its dependency on different resources has changed as new technologies emerged and national priorities changed. The energy flow charts, released each spring and available in a searchable format at flowcharts.llnl.gov/commodities/energy, help scientists, analysts, and the nation's decision makers visualize the complex interrelationships involved in powering the nation. By bringing clarity to these important issues, this ongoing project provides information critical to national energy and resource security.

Energy Use Visualized

Lawrence Livermore's A.J. Simon, the project's principal investigator, notes that part of the Laboratory's national security mission is to explore ways to help achieve U.S. energy and resource security in an evolving energy and climate landscape. Part of that work requires knowing the flow of energy in the here and now, and how that compares to the past. Mapping the energy flow through the national economy relates to that mission. "The project feeds directly into the Laboratory's energy security and climate resilience mission to secure and expand the supply and delivery of affordable, clean energy," says Simon. "This mission requires us to understand and communicate the ins and outs of energy production, distribution, and utilization."

Project analyst Hannah Goldstein creates the energy flow charts (see example on p. 15), which take the form of Sankey diagrams. Sankey diagrams are used for data visualization in process engineering, energy management, manufacturing, and other areas. The energy flow charts, reading left to right, present material, energy, and cost flows from resource to final disposition. Goldstein explains, "The size of each box and width of each line provide a general correlation to the amounts of energy involved. Thus, the bigger the box or the thicker the line, the larger the energy flow."

Raw energy resources fall into one of three categories: renewable (solar, hydro, wind, geothermal, and biomass), nuclear, and fossil (natural gas, coal, and petroleum). The energy amounts for the resources are represented in quads—a quadrillion (one thousand trillion) British Thermal Units (BTUs). For comparison, 20 million passenger cars consume about a single quad of energy in a year.

In the center of the diagram, the flow lines split and recombine, showing how portions of energy from different resources move through the energy system to fulfill the needs of various sectors of the nation's economy and infrastructure. All of these resources contribute some fraction of their total energy to the production of electricity— from nuclear energy, which contributes 100 percent of all 8.46 quads to electricity generation annually, to petroleum, which is primarily used for transportation but still contributes 0.19 quads to niche electricity generation applications.

Whether from generated electricity or directly from the raw resources, energy flows continue to boxes representing economic end use sectors: residential (houses and apartment buildings), business (including office buildings, retail, restaurants, stores, and hospitals), industrial (including manufacturing, agriculture, and construction), and transportation (cars, trucks, planes, trains, ships, etc.). To the far right, a set of boxes indicates the total energy expended by these sectors, from keeping cars on the road to heating and cooling homes.

The final box on the right side of the energy flow diagram represents rejected energy, energy released back into the environment. As Simon notes, the first law of thermodynamics states that energy cannot be created or destroyed. “The second law says that whenever we transform energy from one form to another, we lose some energy along the way. This lost or ‘rejected’ energy is often released as heat. For example, when driving a car, approximately 20% of the energy from the fuel reaches the wheels. The rest goes back into the environment in the form of heat released through the radiator or the tailpipe.”

Fifty Years of Change

Early charts were visually complex, including details such as imported and exported energy resources. “The results sometimes had the unfortunate characteristic of looking like a plate of tangled spaghetti,” says Simon. Today’s charts are streamlined, with a focus on the data and presentation. “We zero in on what the audience wants to know and how best to portray that information,” says Goldstein. “For example, we look at whether some of the resources can be bundled together or if more details are required. We also consider how best to lay out the lines of flow to make the chart easier to read.”

Livermore’s energy flow charts show, over time, how the country’s energy use has changed. “For instance,” says Simon, “when the oil crisis hit in the 1970s, the country began shifting away from its heavy reliance on oil for generating electricity to using more coal and natural gas, which were domestically abundant fuels at the time. In the late 1980s, nuclear power became a significant source of electricity. Starting in the mid-2000s, the charts show solar and wind becoming meaningful contributors. Before then, the lines representing these renewable resources were barely a pixel in width. In the earliest charts, they aren’t even listed.”

Just as utilization of resources has changed over time, so have the production and appearance of the energy flow charts themselves. Originally, a researcher reviewed sparsely published estimates of energy use, and a graphic designer drew the chart by hand, using rulers and ink. In the decades following the energy crisis of the 1970s, the EIA and its predecessor organizations developed more sophisticated, comprehensive, and accessible data products describing energy resources, conversion, and use. Graphics programs became available in the 1990s, and the project team built their own Java-based desktop program in the 2010s.

With the introduction of commercially available Sankey diagram software, production and updating became even easier. Goldstein notes, however, producing the charts requires more effort than pouring data into the Sankey software. “Creating a final chart is an art,” she says.

Many of Lawrence Livermore’s early energy flow charts were visually complex, showing myriad details such as this one for 1976, which includes information on imported and exported energy. Charts created in the ’70s and ’80s were hand-drawn, whereas later versions were created with various software tools.

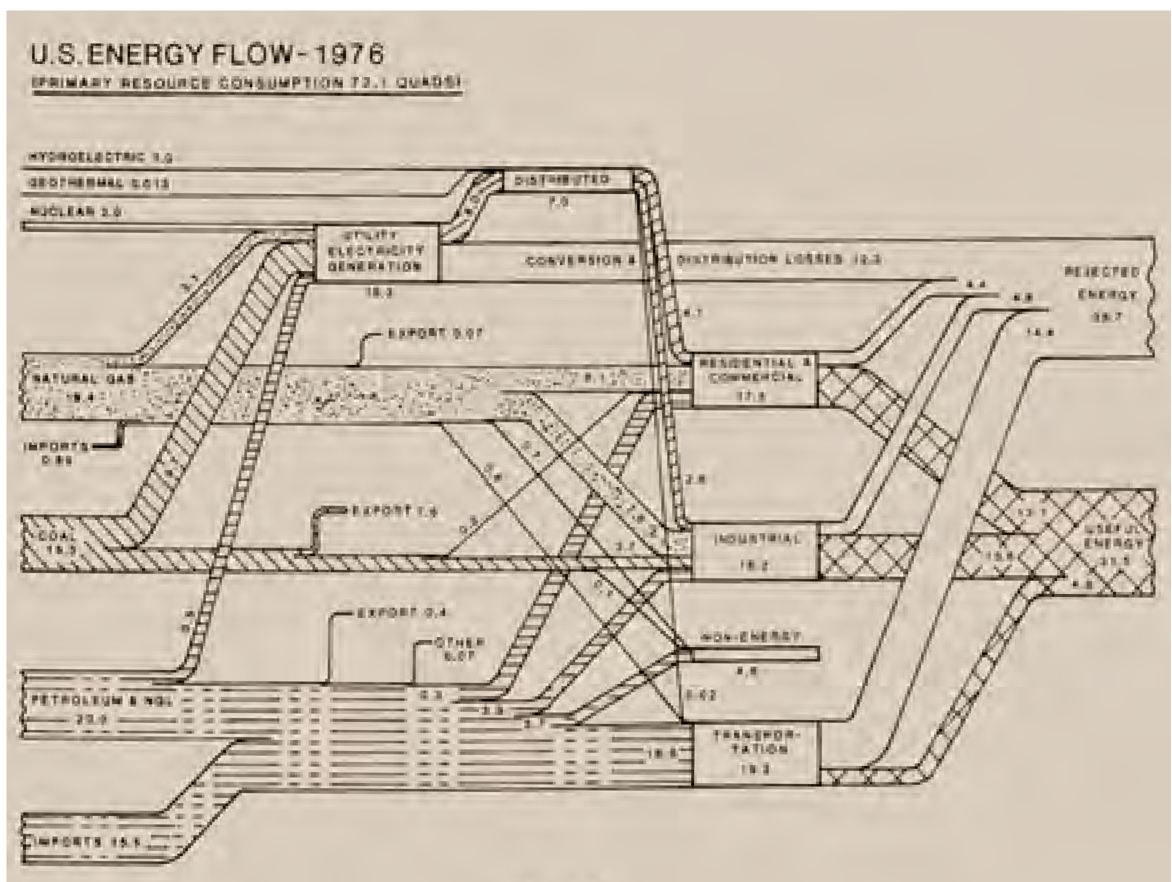

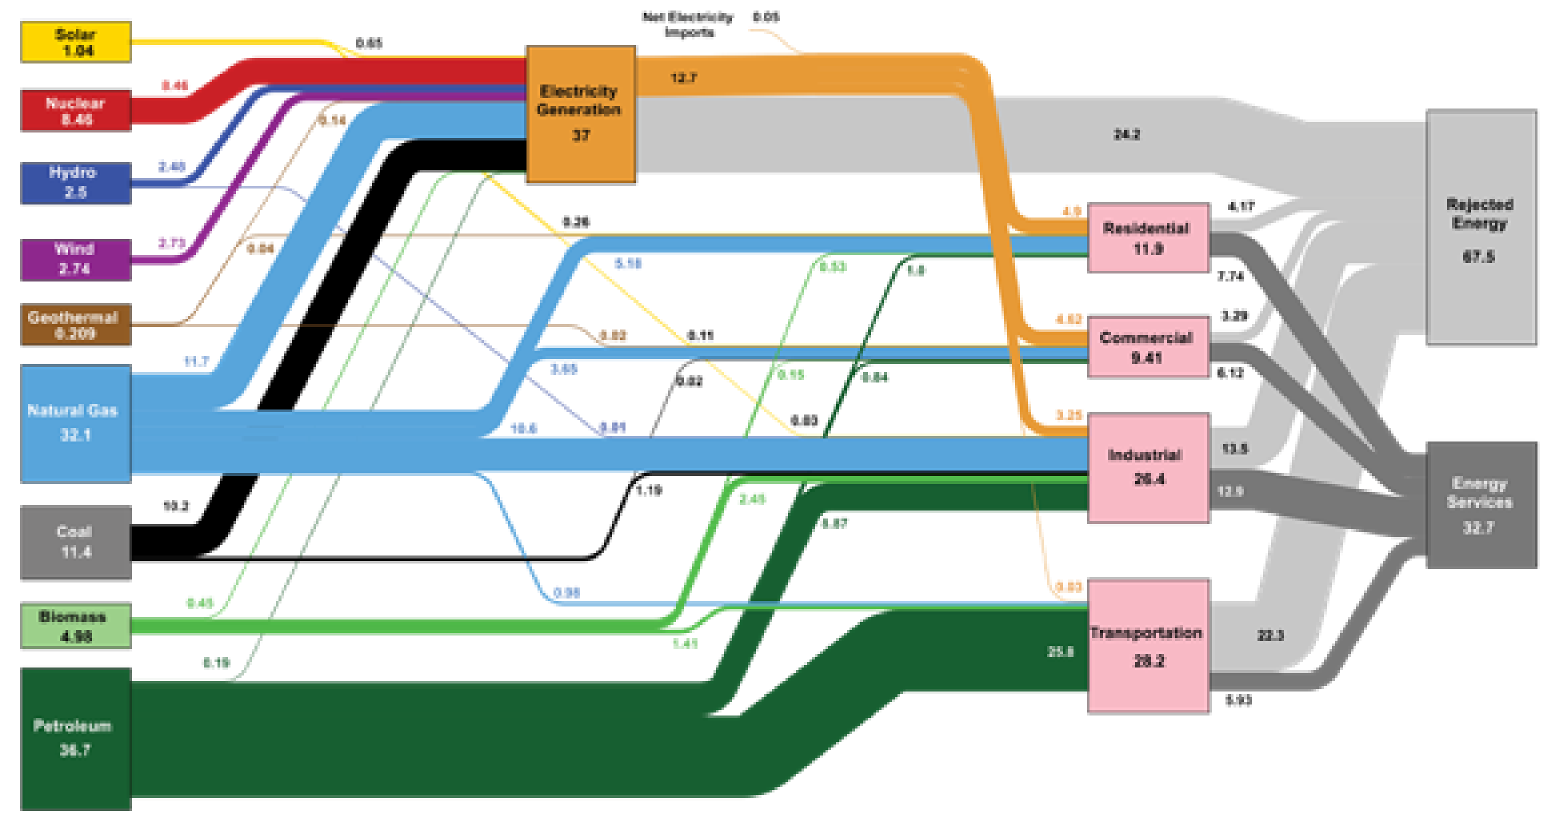

Energy flow charts produced by Lawrence Livermore researchers present a wealth of information about (left) where the nation’s energy comes from, (center) how much of each resource is used in various sectors, and (right) how efficiently that energy produces end-use services.

Future Flow

As the project has garnered more attention and interest over the years, the researchers have tackled requests for custom charts such as the 2010 introduction of a flow chart focused on energy use in the U.S. transportation and residential sectors. Such analyses provide insights that can be used, for instance, in identifying underused resources or opportunities for better technology. In 2010, the Laboratory worked with DOE’s Office of Energy Policy and Systems Analysis and the National Energy Technology Laboratory to produce an atlas of hybrid energy/water Sankey diagrams depicting energy use and water flow in each state.

Simon and Goldstein continue to respond to organizations such as the Department of Defense and others as they seek to better understand different energy scenarios. “We don’t build models of energy futures, but we do help other researchers visualize them,” says Simon.

Some requests for diagrams go beyond simple energy use, seeking charts that present the complex relationships among energy, water, and carbon. Goldstein notes, “We recently released carbon dioxide emission flow charts up to the year 2018, nationwide, and up to the year 2017 statewide. All these charts are now on the website as well.” Goldstein has also drawn on data from the International Energy Agency for international energy charts.

The usefulness and visibility of the Laboratory’s effort continues to expand. The production of the charts and their level of detail may have changed, but the effort has real staying power. “The Lab has been producing these for almost half a century,” says Simon. “Clearly, they are valued. We keep seeing new users show up and previous requesters return. They understand that these diagrams are great snapshots that help explain our nation’s complex energy systems in single, very visual images.”

—*Ann Parker*

Key Words: electricity generation, energy flow charts, energy resources, energy use renewables, fossil fuel, Sankey diagrams.

For further information contact A. J. Simon (925) 422-9862 (simon19@llnl.gov).

CELEBRATING SCIENCE AND ENGINEERING OUTREACH

IN the summer of 2010, an event hosted by the Lawrence Livermore Science Education program left a lasting impression on Nan Ho, a biology professor at Las Positas College (LPC). Richard Farnsworth, then the science education program manager, had organized the presentation "Engineer vs. Scientist" with physicist Chris Eberle and engineer Rick Cross, both from Lawrence Livermore. "Listening to these two people exchange ideas, after years of working together, I was fascinated by how they approached the same, complex problem in different ways but engaged with each other as team members," says Ho. "That richness of perspectives seared into my memory. I also realized that I would not have been able to learn how much exciting work happened at the Lab if I had not been 'past the gate.'"

As a Department of Energy Academies Creating Teacher Scientists (DOE ACTS) program fellow, Ho had the opportunity to attend presentations and observe Livermore's research firsthand. DOE ACTS, designed by the department's Office of Science, aimed to create outstanding science and math instructors armed with scientific information and research experience to serve as positive change agents and leaders in their teaching communities.

Ho asked Farnsworth to explore cultivating a connection between Livermore scientists and LPC students and faculty. At the time, the only way for those without a Livermore badge to learn about the Laboratory's work was by attending Science on Saturday lectures. Beth Vitalis, a Livermore biomedical scientist, adjunct biology professor, and colleague of Ho's at LPC, introduced the joint seminar concept to her colleague

Lawrence Livermore National Laboratory and Las Positas College (LPC) have jointly hosted a science and engineering seminar series on the LPC campus for over 10 years, cultivating connections between the Laboratory and LPC students. (Photo courtesy of Doug Jorgensen, *The Independent*.)

Kris Kulp, then a group leader in the Biosciences and Biotechnology Division and now the division leader. Ho also recruited Neal Ely, LPC dean of Math, Science, Engineering, and Public Safety, to join Vitalis, Kulp, and Farnsworth to launch the LLNL/LPC Science and Engineering Seminar Series.

A Perfect Idea

The first presentation, in October 2010, featured Kulp and physicist Kuang Jen Wu in "Biology as a Team Sport: How Physicists, Chemists, Statisticians, and Biologists Work Together." Kulp and Wu discussed using time-of-flight secondary ion mass spectrometry to image biological samples. Kulp recalls, "The speaker series was a perfect idea, and the seminar was great fun. We really emphasized team science. Kuang Jen and I each gave a piece of the presentation, so that the students could see how the two different disciplines worked together to tackle the same problem. The auditorium was standing room only, and the students were enthusiastic and engaged."

Over the next 10 years, approximately 70 scientists and engineers presented their work to the LPC community. As the speaker series evolved, so too did the topics. "Some topics were

Following his 2012 talk, Laboratory Director Parney Albright (above) toured LPC's new science facilities with students attending the seminar series.

The seminar series, attended by thousands of students over the last decade, supports and inspires students attending LPC to build science and research skills. (Photo courtesy of LPC.)

really technical, but we found speakers who we knew would appeal to college students and the local community," says Ho. From micro- and nanotechnology and bioengineering to experimental plasma physics, seismology, and public safety, scientific divisions from the Laboratory have presented to a packed auditorium at LPC. A highlight of the series came in fall 2012 when then Laboratory Director Parney Albright presented "Running a National Lab and What It Takes to Succeed at One." Albright provided an overview of the Laboratory missions and his personal

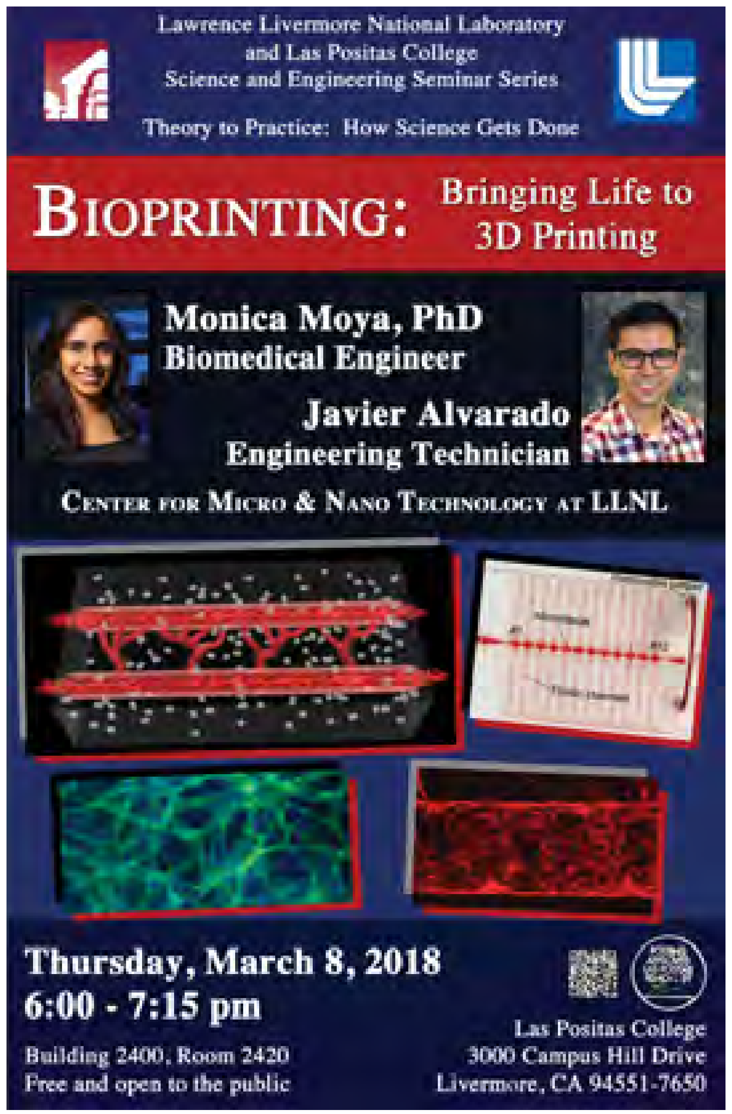

The March 2018 event "Bioprinting: Bringing Life to 3D Printing" featured the work of LLNL's Monica Moya and Javier Alvarado, who had been an LPC student the year before.

take on how to succeed in a science career, emphasizing the need for expertise, breadth of experience, and a commitment to excellence.

In addition to lectures, the series offers an exclusive opportunity for select science and engineering students to meet in small groups with the Livermore scientists and show them around the LPC science facilities. "At first, the students can feel somewhat nervous about talking to the Livermore scientists," says Ho. "But someone always asks, 'How did you become a scientist or engineer? Do you have any advice?' And the common response is that everyone has a unique professional journey loaded with unexpected twists and turns. Having scientists speak to the LPC students gives them a better sense of what a career in science or engineering involves, and that it is attainable. These experiences remove some of the mystery. The students aren't intimidated; their horizons expand. And the connection between the students and the speakers is really powerful." The seminar series has facilitated many professional networking relationships over the years that have yielded internships and even full-time employment positions at the Laboratory.

A Terrifying Leap

In spring 2017, Javier Alvarado, just completing an internship in Livermore's Materials Engineering Division, decided to take classes at LPC to refresh his math skills and explore biology and computer-aided design. Seeing posters for an upcoming LLNL/LPC lecture, "The Science behind 3D Printing," part of the Theory to Practice: How Science Gets Done series, Alvarado spotted an opportunity he knew he could not miss. "I was entering the field of 3D bioprinting and wanted to know more about 3D printing aspects—the materials, the challenges," says Alvarado. "I thought, 'The speakers could be people I work with in the future.'"

Then, in March 2018, Alvarado, now a full-time employee, was invited by his mentor, Monica Moya, to co-present with her at the LLNL/LPC seminar to discuss his work in 3D cell culture, biomaterials, and applying 3D bioprinting techniques to model the neurovascular environment *in vitro*. The prospect of speaking about his research in front of 300 people initially terrified Alvarado, but three words into the talk, he loved the experience.

During Alvarado's portion of the presentation, "Bioprinting: Bringing Life to 3D Printing," he demonstrated how diversity in thought, knowledge, and process can come together to answer big questions. Alvarado described in simple, accessible terms how he applies the scientific method and the engineering design process together to tackle problems, saying, "I have an idea. Let me try it and keep troubleshooting until I figure it out."

"After the presentation, and after talking to the students, I felt great seeing how my work can have a positive impact," says Alvarado. "If underserved students can get a better sense of

what it takes to be an engineer, maybe that visibility will attract them to a career in science and engineering they might not have considered." Alvarado's family sat in the crowd that day, including his grandfather who was suffering from dementia. "Having recently lost my grandfather to COVID-19, the experience of presenting at LPC has even more significance to me and my current path," says Alvarado, who is pursuing his master's in Bioengineering at UC Riverside with support from Livermore's Education Assistance Program while working full time at the Lab. He went on to present a second time in Lawrence Livermore's Science on Saturday lecture series for middle and high school students in 2019.

A Vibrant Celebration

Over the years, many people have worked behind the scenes on the collaboration. "The seminar series provides a great opportunity for LPC students to see the variety of science and engineering research efforts going on inside the Lab, and to learn about internships and future career opportunities," says Joanna Albala, Livermore's Science Education program manager. "Sadly, Dick Farnsworth passed away last March, but he would be so proud to know that his work continues." Albala credits LPC team members including Adeliza Flores, Barbara Zingg, Caryl Shill, Jean O'Neil-Opipari, Mike Ansell, Robin Rehagen, and other faculty and staff who have supported the series and encouraged students to attend. "From running the video cameras to designing fliers and accompanying the Livermore guests on tours of the LPC campus labs, the series is an authentic team effort, energized and sustained by Laboratory members as well as faculty, students, and staff," says Albala.

To kick off the seminar series' tenth year in October 2019, Jessica Osuna, an operations scientist at the Livermore's National Atmospheric Release Advisory Center spoke about interdisciplinary science in support of public safety. The following month, Carolyn Koester, associate director of the Forensic Science Center, presented "Responding to Chemical Warfare," and Kulp was honored for her 10 years of service on the seminar planning team with a plaque presented by Ho. "The enduring nature of the seminar series is a testament to its importance," says Kulp. "Engaging and inspiring the next generation of scientists and engineers is vital to the Lab's pursuit of scientific excellence. Interacting with talented students is a great way to start! We hope to sustain the seminar series well into the future and see more success stories like Javier Alvarado's." LPC interim president Roanna Bennie also honored the partnership with Albala accepting the award on behalf of the Laboratory. The wording on the plaque captured the spirit and mission of the two partners: "Creating a Vibrant Celebration of Science and Engineering for the Community."

The seminar series continues to engage audiences despite moving to a virtual format in response to COVID-19 restrictions.

Lawrence Livermore and LPC continue to share the Laboratory's exceptional science and technology resources with the community. Pictured in the front row, from left, are Joanna Albala (Livermore), Kris Kulp (Livermore), former president Roanna Bennie (LPC), Carolyn Koester (Livermore), Jean O'Neil-Opipari (LPC), Barbara Zingg (LPC), and dean Nan Ho (LPC), joined by additional LPC faculty and LPC Fire Technology students at the 10-year anniversary event. (Photo courtesy of Jean O'Neil-Opipari, LPC.)

Ho, now dean of Science, Technology, Engineering, and Mathematics at LPC, still possesses the same passion for educating and inspiring the next generation of scientists and engineers while showcasing the opportunities and capabilities Lawrence Livermore offers. "Thousands of students have attended these seminars over the last decade," says Ho. "The fact that the series is entirely volunteer-based and has lasted for over 10 years is remarkable. How often do you get that kind of commitment and connection? I would love to see it celebrate another 10 years."

—Genevieve Sexton

Key Words: 3D bioprinting, Department of Energy Academies Creating Teacher Scientists Program (DOE ACTS), education, engineering, Las Positas College (LPC), Lawrence Livermore Science Education program, materials engineering, outreach, team science.

For further information contact Joanna Albala (925) 422-6803 (albala1@llnl.gov).

To learn about upcoming seminars, go to laspositascollege.edu/llnl. To see past seminars, go to laspositascollege.edu/llnl/past.php

DENSE PLASMA FOCUS BACK IN THE SPOTLIGHT

Chris Cooper, Alex Povilus (left) and Gustavo Bartolo (right) make adjustments to MJOLNIR, LLNL's Dense Plasma Focus, to increase neutron yield, creating a radiological imaging source for the Laboratory. (Photo by Garry McLeod.)

USING neutrons to "look through" fast-moving experiments requires an intensely bright neutron source originating from a very small spot. One tool that suits this scientific role is the dense plasma focus (DPF). The DPF uses a strong magnetic field to compress a plasma into a discharge called a "pinch" that creates a short, intense pulse of x rays, neutrons, and directed ion beams.

Introduced in the 1960s, the relatively compact size and uncomplicated design of the devices made them attractive as the foundation for fusion power plants, although that application was eventually discounted after decades of research. More recently, the DPF has been considered as a neutron source for radiography and other national security applications. An important complement to x-ray imaging, fast neutron imaging enables researchers to examine the lighter elements in objects of study. Unlike current commercial and industrial neutron imaging sources, the DPF generates an incredibly bright but very short flash (less than 100 nanoseconds) that is well-suited for taking still-frame pictures of highly dynamic processes.

The inability to model the DPF plasma with sufficient fidelity, however, has limited researchers' detailed understanding of the physics needed to develop methods that produce a consistent, reliable, and intense neutron pulse. "Decades ago, the DPF was essentially a black box operated by rules of thumb," says Livermore scientist Alex Povilus.

After a multiyear Livermore research program, the DPF's time to shine may have finally arrived. A team led by researcher Andréa Schmidt has been using hundreds of computer simulations to understand—physically—what occurs in a DPF discharge. With this knowledge, Schmidt's team has developed a deeper understanding of the effect of different DPF design aspects—pulsed-power capacitor bank characteristics, anode length and shape, operating voltage, device pressure—on its neutron pulse. Now, a prototype DPF design and ongoing modifications are informed by modeling predictions. The ability to guide experimental design using simulations and then challenge the simulations with experimental results has enabled the team to make remarkable progress in DPF design. The team's ultimate goal is to create a new flash neutron imaging capability for the nuclear weapons enterprise.

Thor's Hammer

The prototype DPF, named MJOLNIR for MegaJOuLe Neutron Imaging Radiography, yields experimental data to bolster new simulations, further refining the design toward demonstrating the desired intense neutron pulse. MJOLNIR (pronounced "myole-ner), known to Norse mythology and Marvel comics fans as the name of Thor's hammer, was proposed in a naming contest by Povilus, the lead experimentalist.

MJOLNIR is the latest and largest iteration in a series of smaller projects developing the DPF. Initial efforts, seen as high-risk but high-reward, were funded by Livermore's Laboratory Directed Research and Development program. Early simulations informing an effective, predictable DPF (see *S&TR* July/August 2013, pp. 22–25) were performed using LSP, the computationally intensive particle-in-cell (PIC) code. Faster-running fluid codes could not adequately model ion distributions in a DPF's plasma discharge and, as a result, could not correctly predict properties of the resulting neutron pulse, limiting their utility.

Schmidt's first PIC simulations, although primitive compared to the simulations run today, matched the neutron yield for experimental DPF devices. Her team scaled to bigger devices and modeled other teams' experiments. "Skeptics thought our models were slow, therefore, not useful," says Schmidt. "As we learned more, performance improved and generated confidence in our approach. A sponsor stepped up to fund a project to build a larger machine for our experiments."

In this image of visible light emission from the pinching plasma discharge, a dense plasma focus (DPF) produces a plasma sheath that pinches, over the course of tens of nanoseconds, eventually becoming unstable and producing an energetic ion beam. The ion beam's bombardment of a dense plasma target produces neutrons.

Schmidt's team later developed a hybrid methodology using both fluid and PIC methods in Chicago code, the successor to LSP, making it possible to run simulations with larger electrodes and sufficiently accurate physics models in a reasonable computational time. The Livermore team ran models in parallel with construction of the current prototype, continually improving the design and proposing novel electrode shapes for future experiments.

The length and shape of the anode and cathode, as well as other components of the DPF device, were designed to reliably produce the most neutrons possible based on hundreds of kinetic plasma simulations. (Photo by Garry McLeod.)

The MJOLNIR design-build team: (from left to right) Rick Anaya, Andréa Schmidt, Steve Chapman, Justin Angus, Matt McMahon, Ed Koh, Tony Link, Don Max, Clément Goyon, Harrison Flores-Alimboyoguen, James Mitrani, Yuri Podpaly, Dave Van Lue, Drew Higginson, Ihor Holod, Chris Cooper, Steve Hawkins, Alex Povilus, Michael Anderson. Not pictured: Tony DaCosta, John Reed, Kurt Walters. (Photo by Ian Fabre.)

High Neutron Yields

After an 18-month building period, MJOLNIR fired its first shots in August 2018, achieving neutron yield on the seventh shot. "Generating neutron yield that quickly is a remarkable accomplishment," says Michael Zika, deputy director of Transformational Weapon Science in the Weapon Physics and Design Program within Livermore's Weapons and Complex Integration directorate.

The neutron yield in MJOLNIR's initial shots reached 2.5×10^{10}. Optimizing pulsed power and anode/cathode design have led to an order-of-magnitude increase in neutron yield, with recent shots achieving up to 4×10^{11} neutrons per pulse at 2.2 MA peak current. After operating at 1 MJ of stored energy, experiments are under way to apply 2 MJ, the highest energy level ever applied to a DPF.

"Past DPF projects tried and failed to achieve high neutron yields at the plasma currents we're aiming for," says Schmidt. "Without our predictive modeling capabilities, past researchers could observe their source get brighter or dimmer when they made modifications, yet they had a limited understanding of what was happening in the plasma."

MJOLNIR's neutron source is located in a 25-foot diameter pit, one level above capacitor towers providing a total of 2 MJ and one level below ground level, where the project's diagnostic digitization equipment is shielded from the electromagnetic pulse and ionizing radiation generated each time the device fires. "We're diagnostics-rich for greater certainty in gathering experimental data from firing," says Povilus. "We're also able to reconfigure measurements and deploy new diagnostics rapidly to help better understand features we see in the pinch physics."

Once the MJOLNIR prototype demonstrates that a DPF can meet requirements for a flash neutron radiographic source, the team will create a MJOLNIR-like machine for use at facilities such as Livermore's Site 300. If successful at Site 300, other x-radiographic facilities in the National Nuclear Security Administration complex, such as the Dual Axis Radiographic Hydrodynamic Test (DARHT) facility at Los Alamos National Laboratory and Scorpius, an x-ray source being built for the Nevada National Security Site, may be interested in complementing their capabilities with neutron radiography.

The Right Shot at the Right Time

MJOLNIR can provide a less expensive platform for developing diagnostics as well as a compact source that can go to the experiment, rather than the other way around. "Not everyone is willing to maintain an accelerator facility for a flash radiation system," says Povilus. "MJOLNIR is a compact alternative that could be fielded within existing facilities to provide additional radiograph capability. If you need a quick pulse, this is an economical solution for many radiography applications."

Capacitor towers providing 2 megajoules of stored energy fill the floor below the chamber where the plasma and neutrons are produced. (Photo by Garry McLeod.)

Looking toward other potential applications, flash neutron radiography lends well to viewing fuel flow through engines that cannot be imaged effectively with traditional x-ray radiography. Miniature versions of a similar DPF source have been considered to replace radiological sources currently used in petroleum and geothermal well logging.

The DPF, discounted decades ago for fusion power plants, is enjoying a resurgence of interest as short-pulse neutron sources for national security applications. The Livermore team, powered by detailed simulations, has advanced understanding of the DPF plasma discharge beyond what earlier researchers could have imagined. The interplay between simulations and experiment has enabled development of a machine to support new applications.

Livermore scientist and MJOLNIR team member Steve Chapman finds satisfaction in improving a technology that held so much unrealized promise in the past. "We're coming back for a long-forgotten tool on the shelf," says Chapman. "The DPF wasn't dead, only dreaming."

—Suzanne Storar

Key Words: capacitor bank, dense plasma focus (DPF), flash neutron imaging, MJOLNIR, neutron source, particle-in-cell (PIC) code, plasma discharge, predictive modeling, pulsed power, radiography, z-pinch.

For further information contact Andréa Schmidt (925) 423-9636 (schmidt36@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 eight-digit number in the search box at the U.S. Patent and Trademark Office's website (uspto.gov).

Patents

Nanostructured Layer for Graded Index Freeform Optics
Eyal Feigenbaum
10,612,145 B2
April 7, 2020

Heat Treatment to Anneal Residual Stresses During Additive Manufacturing
James A. DeMuth, Andrew Bayramian, Bassem S. El-dasher, Joseph C. Farmer, Kevin J. Kramer, Alexander Rubenchik
10,618,111 B2
April 14, 2020

Compact Absorptivity Measurement System for Additive Manufacturing
Alexander M. Rubenchik, Manyalibo Joseph Matthews, Johannes Trapp
10,646,960 B2
May 12, 2020

Mechanical Reticulation of Polymeric-Based Closed Cell Foams
Jennifer N. Rodriguez, Duncan J. Maitland, Thomas S. Wilson
10,647,037 B2
May 12, 2020

System and Method for Computed Axial Lithography (CAL) for 3D Additive Manufacturing
Brett Kelly, Robert Panas, Maxim Shusteff, Christopher Spadaccini, Hayden Taylor, Indrasen Bhattacharya
10,647,061 B2
May 12, 2020

System and Method for Light Assisted Friction Stir Processing and Welding of Metallic and Non-Metallic Materials
Joseph C. Farmer, Alexander M. Rubenchik, Raymond J. Beach, Robert J. Deri, Edward I. Moses, Bassem S. El-Dasher, Sarath K. Menon, Terry McNelley
10,647,064 B2
May 12, 2020

Concentric Semi-Circular Split Profiling for Computed Tomographic Imaging of Electronic Beams
John W. Elmer, Alan T. Teruya
10,649,102 B2
May 12, 2020

Systems and Methods for Additive Manufacturing to Encapsulate Transformative Colloidal Suspensions
Julie A. Jackson, Eric Duoss, Alexandra Golobic, Mark Christian Messner, Christopher Spadaccini, Kenneth J. Loh
10,661,549 B2
May 26, 2020

Awards

Livermore scientist **Félicie Albert** has been elected a **Kavli Fellow** of the **U.S. National Academy of Sciences**, the Laboratory's seventh Kavli Fellow since the program started in 1989. As a new Kavli Fellow, Albert was invited to present a poster about next-generation x-ray light sources during the academy's annual Kavli Frontiers of Science U.S. symposium in July 2020. The symposium is sponsored by the Kavli Foundation, founded by the late Fred Kavli to advance science and promote the understanding of scientific research.

Four Livermore scientists—**Richard Berger**, **Laurent Divol**, **Max Fenstermacher**, and **Art Nelson**—have been selected as **2020 fellows** of the **American Physics Society** (APS). The new fellows represent physics expertise including laser plasma physics, magnetic fusion plasmas, theoretical and computational understanding of plasma interactions, and soft x-ray and free electron laser platforms. APS fellowships are awarded after extensive review and rely on nomination and recommendation by professional peers. The APS is a nonprofit organization with over 55,000 members dedicated to the advancement of physics knowledge.

Livermore atmospheric scientist **Ben Santer** has been honored with the **American Geological Union's 2020 Bert Bolin Award**. The award is presented annually for groundbreaking research or leadership in global environmental change through cross-disciplinary, interdisciplinary, and trans-disciplinary research in the past 10 years. Santer's recent work has attempted to identify anthropogenic fingerprints in a number of different climate variables, such as tropopause height, atmospheric water vapor, the temperature of the stratosphere and troposphere, ocean heat content, and ocean surface temperatures in hurricane formation regions.

The **Association for Women in Mathematics** has named Livermore computational scientist **Carol Woodward** as a **2021 fellow**, recognizing her commitment to supporting and advancing women in the mathematical sciences. Woodward's research focuses on nonlinear solvers and time integration methods and software, and she is part of the Department of Energy's (DOE's) FASTMath SciDAC Institute to improve numerical software for use in DOE applications. She is also working with the DOE Exascale Computing Project in developing time integration software for use in large-scale combustion, phase field, and cosmology applications.

Abstract

The Exascale Software Portfolio

As a leader in high-performance computing (HPC), Lawrence Livermore wields a large portion of the Department of Energy's (DOE's) HPC resources to advance national security and foundational science. The Sierra supercomputer supports the National Nuclear Security Administration's (NNSA's) Stockpile Stewardship Program by enabling more accurate, more predictive simulations. This generation of computers is known as heterogeneous, or hybrid, because their architectures combine graphics processing units and central processing units to achieve peak performance well above 100 petaflops. (A petaflop is 10^{15} floating-point operations per second.) The next generation's processing capability—at least an exaflop (10^{18} flops)—will be many times greater. HPC software must adjust to these new hardware standards. As the exascale era begins, two major initiatives leverage and expand Livermore's HPC capabilities, with a spotlight in this issue on software. The Exascale Computing Project (ECP), a joint effort between the DOE Office of Science and NNSA, brings together several national laboratories to address many hardware, software, and application challenges inherent in the organizations' scientific and national security missions. Within the Laboratory, the RADIUSS project—Rapid Application Development via an Institutional Universal Software Stack—aims to benefit scientific applications through a robust software infrastructure.

Contact: Lori Diachin (925) 422-7130 (diachin2@llnl.gov) or Rob Neely (925) 423-4243 (neely4@llnl.gov).

U.S. Government Printing Office: 2021/770-075-20002

Coming Next Issue

Fusion Supports the Stockpile

The safety and reliability of our nation's nuclear deterrent depends on innovation along the path to fusion energy.

Also in the next issue:

- *The driving range of hydrogen-powered vehicles gets a boost with the help of quantum mechanics.*
- *A simple, yet powerful, health assessment tool in development could enable blood tests in extreme environments, such as Earth's orbit.*
- *Livermore scientists track the mysterious migration of plutonium in radioactive fallout.*

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

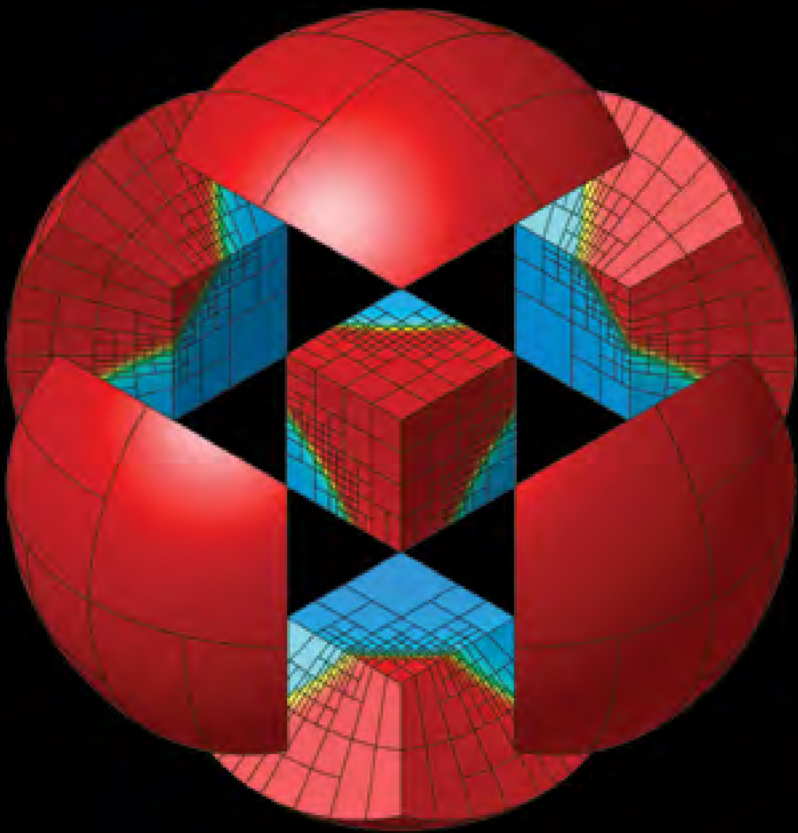

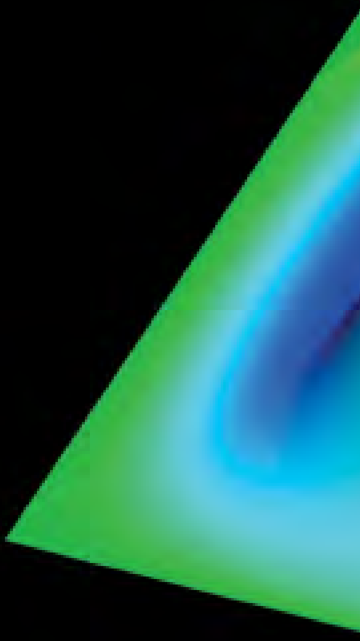

Printed on recycled paper.