COMPUTATION

2014 ANNUAL REPORT

PRODUCTION TEAM

LLNL Associate Director for Computation
Dona L. Crawford

Deputy Associate Directors
James Brase, Trish Damkroger, John Grosh, and Michel McCoy

Scientific Editors
John Westlund and Ming Jiang

Art Director
Amy Henke

Production Editor
Deanna Willis

Writers
Andrea Baron, Rose Hansen, Caryn Meissner, Linda Null, Michelle Rubin, and Deanna Willis

Proofreader
Rose Hansen

Photographer
Lee Baker

3D Designer
Ryan Chen

Print Production
Charlie Arteago, Jr., and Monarch Print Copy and Design Solutions

LLNL-TR-668095

Prepared by LLNL under Contract DE-AC52-07NA27344.

CONTENTS

MESSAGE FROM THE ASSOCIATE DIRECTOR

For more than six decades, Lawrence Livermore National Laboratory (LLNL) has pioneered the computational capabilities required to advance science and address the nation's toughest and most pressing national security challenges. This effort, dating to the Laboratory's earliest years, has been punctuated by the acquisition of leading computer systems and their application to an ever-broadening spectrum of scientific and technological problems. In 2014, this tradition continued as the Computation Directorate provided the computing architectures, system software, productivity tools, algorithmic innovations, and application codes necessary to fulfill mission-critical national security objectives for the U.S. Departments of Energy (DOE), Homeland Security, and Defense, as well as other federal and state agencies.

As Livermore scientists continue to push the boundaries of what is scientifically possible, they require computational simulations that are more precise, that cover longer periods of time, and that visualize higher fidelity, more complex systems. This requires that Computation advance the state of the art in both high performance computing (HPC) and the supporting computer science.

Larger and more complex simulations require world-class platforms. To sustain the nation's nuclear weapons deterrent without the need for full-scale underground nuclear tests, we are preparing for another major advance in HPC that will make our physics and engineering simulation models more predictive. The Collaboration of Oak Ridge, Argonne, and Livermore (CORAL) is making this advance possible. As a result of CORAL, three DOE systems—one for each laboratory—are being acquired. LLNL's system acquisition contract for Sierra was signed in November 2014 with an expected delivery date late in CY17. Sierra is a next-generation supercomputer and Computation's next capability platform for the National Nuclear Security Administration and the Advanced Simulation and Computing Program. At the moment, Sierra only exists as a proposal within contract language. A significant amount of research, development, design, and testing must still take place to move Sierra from proposal to reality. LLNL and Oak Ridge National Laboratory have partnered with the hardware vendors IBM, NVIDIA, and Mellanox to form several working groups responsible for co-designing the new architecture. Co-design is an essential part of developing next-generation architectures well suited to key DOE applications. By incorporating the expertise of the vendors' hardware architects, the system software developers, and the DOE laboratory experts—including domain scientists, computer scientists, and applied mathematicians—Sierra will be well equipped to handle the most demanding computational problems.

In addition to the design and development of the platform, other preparations are underway to ensure that Sierra can hit the computer room floor running. An important aspect of a successful launch is having applications that are capable of utilizing such an immense system. We are adapting our existing large applications using incremental improvements such as fine-grained threading, use of accelerators, and scaling to millions of nodes using message processing interface—with the 20-petaflop Sequoia Blue Gene/Q machine providing a virtual laboratory for these explorations.

As these codes evolve, our teams also must ensure that they remain stable, fast, and accurate. For more than a decade, Livermore has been developing the Automated Testing System, which runs approximately 4,000 tests nightly across Livermore Computing's (LC's) Linux and IBM Blue Gene systems. These tests generate diagnostic information used to ensure that Livermore's applications are ready for future systems and challenges.

As large-scale systems continue to explode in parallelism, simulation codes must find additional ways to exploit the increasingly complex design. This can be especially difficult for simulations that must replicate phenomena that evolve, change, and propagate over time. Time is inherently sequential, but thanks to novel work by researchers at Livermore, Memorial University, and Belgium's Katholieke Universiteit Leuven, existing large-scale codes are becoming capable of parallelizing in time. By using a new multilevel algorithm called XBraid, some applications can now solve problems up to 10 times faster. As its name suggests, XBraid allows simulations to "braid" together multiple timelines, eliminating the need to be solved sequentially. These braided solutions are solved much more coarsely, and the results are fed back into the algorithm until they converge on a solution that matches the expected results from traditional sequential algorithms, within defined tolerances.

Taking advantage of LC's valuable resources in an orderly manner is no small task. With a sizable number of complex applications, Livermore requires robust means to manage the competing demands for a platform's time, network capacity, power consumption, and software licenses as resources. To address this issue, we have begun the development of FLUX, a new scheduling solution that allows for improved management of LC resources. Computing jobs that must increase or decrease in size or that need to migrate to new resources due to failing components are paradigms that are beyond most schedulers; all but a few fail to operate smoothly with the number of systems LC maintains. FLUX allows LC to effectively handle these requirements and provide users a level of access to systems that few other computing centers in the world can offer.

Computation encourages LLNL scientists to leverage our HPC expertise to advance a broad range of science critical to the nation, from elucidating the structure of proteins to optimally operating a smart electric grid to understanding the factors behind population migration in times of extreme disruption. Allocating time on some of our most powerful supercomputers is an important investment the Laboratory makes in our science and technology. One of the ways Computation supports LLNL projects and programs is through the annual Computing Grand Challenge program. Now in its ninth year, the Grand Challenge program awarded more than 15 million central processing unit-hours per week to projects that address compelling, large-scale problems, push the envelope of capability computing, and advance science. This year marked the first time that institutional demand for cycles on the 5-petaflop Vulcan system exceeded the available allocation, an encouraging indication of the growing interest in advancing science by pushing the state of the art in HPC. Only seven countries in the world possess more computing resources in total than LLNL makes available for unclassified computing.

Traditionally, Livermore has focused on delivering flops at the tera-, peta-, and future exa- scales. However, new paradigms of computing are emerging that are not limited by the number of calculations that can be done but instead by the amount of data that can be moved through the system. So-called "big data" problems are becoming more common and are increasingly critical to LLNL's national security missions. To accelerate solutions to these challenges, LC is providing Livermore scientists with platforms on which to perform big data research as well as developing methods for bringing traditional big data tools to our existing HPC systems. These tools allow users to make use of Hadoop, a common big data framework, while taking advantage of LC's parallel file systems and large number of nodes. Work has also commenced that will push the performance of Hadoop even further. For example, 800 gigabytes of flash memory has been installed on each of the Catalyst system's 304 compute nodes. Memory and storage in the system allow Catalyst, a first-of-a-kind supercomputer, to appear as a 300-terabyte memory system. In this configuration, our scientists reduced the TeraSort runtime, a big data benchmark, to a little more than 230 seconds—much shorter than traditional Hadoop. Also this year, Catalyst was made available to industry collaborators through Livermore's High Performance Computing Innovation Center to further test big data technologies, architectures, and applications.

Computation also provides expertise to other mission-critical aspects of LLNL, such as the major software development for the National Ignition Facility (NIF). Among the team's many achievements this year was the addition of the Advanced Radiographic Capability (ARC). ARC allows NIF to produce short bursts of intense x-rays and image NIF targets during shots. The resulting images provide a "movie" with tens of trillionths of a second resolution. Incorporating this functionality required 18 months of development, followed by the deployment of software that touched nearly every aspect of the facilities control systems. This work deepens the paths of inquiry for NIF by providing additional data.

All of these endeavors rely on people, and LLNL has some of the best in the world. We are deeply committed to preparing the next generation of experts to meet future challenges. Computation welcomed its third-largest scholar class ever, hosting 131 students and faculty—76 graduate students, 49 undergraduates, 1 high school student, and 5 faculty members—from 89 universities and 8 countries. Specific specialties within the scholar program have been emerging over the past several years—the Cyber Defenders and the Co-Design Summer School. Now in its fifth year, the Cyber Defenders matched each of the 21 computer science and engineering students with an LLNL mentor and assigned them a real-world technical project wherein they apply technologies, develop solutions to computer-security-related problems of national interest, and explore new technologies that can be applied to computer security. The 2014 inaugural Co-Design Summer School brought nine graduate students from all over the world to LLNL to consider some of the most challenging problems at the boundary between designing application codes and developing high performance computers.

The foundation of our success in deploying advanced computing capabilities, creative technologies, and innovative software solutions lies with the people of Computation. I am grateful to our sponsors for their continued support and, most of all, to the Computation staff who, year after year, demonstrate an extraordinary combination of technical expertise and dedication to the nation. The breadth of skill and creativity they represent are what sets Computation apart as a world- renowned organization.

Dona Crawford
Associate Director, Computation

AN AWARD-WINNING ORGANIZATION

The stories in this annual report present a cross-section of Computation's accomplishments in research, high performance computing (HPC), software applications, and information technology and security. In addition to the projects highlighted in the report, several Computation personnel and projects received prestigious external recognition in 2014. Selected accomplishments are featured in this section.

Sequoia Tops Graph 500 List of Big Data Supercomputers

Sequoia ranked No. 1 on the Graph 500 ranking in November, after completing the largest problem scale ever attempted—scale 41—with a performance of 23.751 teraTEPS (trillions of traversed edges per second). The team employed a technique developed by IBM.

The biannual Graph 500 list looks at graph-based data problems, a foundation of most analytics work, and the ability of systems to process and solve complex problems. Sequoia's performance reflects its ability to push the boundaries of data-intensive computing, an accomplishment that is critical to Lawrence Livermore National Laboratory's (LLNL's) national security missions. In addition to achieving the top Graph 500 ranking, LLNL computer scientists also demonstrated scalable Graph 500 performance on small clusters and even a single node. To achieve these results, Livermore computational researchers combined innovative research in graph algorithms and data-intensive runtime systems.

Funding Innovative Cancer Research

Livermore computational scientist Amanda Randles received a Director's Early Independence Award from the National Institutes of Health (NIH) to pursue research that will help predict likely sites for cancer to metastasize, a method that combines personalized massively parallel computational models and experimental approaches.

The NIH Common Fund award provides funding to encourage exceptional young scientists to pursue high-risk, high-reward independent research in biomedical and behavioral science. Randles will receive about $2.5 million over five years. The goal of the project is to develop a method to simulate flow of realistic levels of cells through the circulatory system, thereby gaining insight into mechanisms that underlie disease progression and localization. "Building a detailed, realistic model of human blood flow is a formidable mathematical and computational challenge requiring large-scale fluid models as well as explicit models of suspended bodies like red blood cells," says Randles. "This will require high-resolution modeling of cells in the blood stream, and necessitate significant computational advances."

Early Career Award for Scaling Simulation Codes

Computer scientist Todd Gamblin was selected by the U.S. Department of Energy's (DOE's) Office of Science Early Career Research Program to receive up to $2.5 million in funding over five years for a project to accelerate the adaptation of scientific simulation codes to increasingly powerful supercomputers, a process that currently can take up to six months for complex applications. Increasingly complex machine architectures and applications are making this process even slower.

Figure a. Robin Goldstone, Dona Crawford, and Maya Gokhale with the certificate for No. 1 on the Graph 500 at SC14.

Figure b. Amanda Randles

Figure c. Todd Gamblin

Figure d. Dean Williams leads the award-winning, multi-institutional Ultrascale Visualization Climate Data Analysis Tools team.

Figure e. Leading the effort to conserve energy and reduce costs by consolidating data centers at LLNL are (from left) Marriann Silveira, Anna Maria Bailey, and Andy Ashbaugh.

Figure f. Abhinav Bhatele

Gamblin's research is particularly important as the HPC community prepares to ramp up computing speeds from petascale (quadrillions of floating point operation per second) to exascale (as much as 1,000 times faster). The dynamically changing behavior of modern simulation codes makes existing techniques for modeling their performance difficult to apply. Under a project entitled "Statistical Methods for Exascale Performance Modeling," Gamblin proposes to develop statistical models of applications that can represent adaptive, data-dependent code behavior in a manner that can be scaled up for more powerful computing systems. In addition, the project will develop techniques to reduce the complexity of application models so that application developers understand them.

Climate Analysis Tool Wins National Tech Transfer Award

The multi-institutional team responsible for developing, refining, and distributing the Ultrascale Visualization Climate Data Analysis Tools (UV-CDAT) was chosen as the 2015 recipient of the Federal Laboratory Consortium for Technology Transfer's (FLC's) Interagency Partnership Award. UV-CDAT is a powerful tool set that aids climate researchers in solving complex data analysis and visualization challenges by integrating more than 70 disparate scientific software packages and libraries.

This award, one of the FLC's highest honors, recognizes the efforts of laboratory employees from at least two different government agencies who have collaboratively accomplished outstanding work in the process of transferring a technology. Institutions recognized by the FLC for the award include Lawrence Livermore, Lawrence Berkeley, Los Alamos, and Oak Ridge national laboratories; Goddard Space Flight Center; and the National Oceanic and Atmospheric Administration's Earth System Research Laboratory.

UV-CDAT also won a regional FLC award in 2014 for Outstanding Partnership. The team is led by Dean Williams.

Consolidating Data Centers Effort Earns Award from DOE

The Laboratory's effort to conserve energy and reduce costs by consolidating data centers at LLNL received a Sustainability Award from DOE. The program, which began in 2011 and uses LLNL's HPC Strategic Facility Plan as a guide, has so far closed 26 data centers, representing 26,000 square feet of space and resulting in an annual savings of $305,000 in energy bills and $43,000 in maintenance costs. Leading the effort are Anna Maria Bailey and Marriann Silveira of Computation and Andy Ashbaugh of Operations and Business Services.

"Our institutional HPC needs will grow in the future and it is important for the Laboratory to have a Data Center Sustainability Plan to ensure we manage our computing resources in an efficient and cost effective way," says Bailey, LLNL's HPC facilities manager.

Young Achiever in Scalable Computing

Computer scientist Abhinav Bhatele was recognized for his HPC research by the IEEE Technical Committee on Scalable Computing. The Young Achievers in Scalable Computing Award acknowledges individuals within five years of receiving their Ph.D. degree who have made outstanding, influential, and potentially long-lasting contributions in the field of scalable computing.

Bhatele's work focuses on how researchers can use present and future HPC architectures with the highest possible efficiency. A key part of this work involves optimizing communication on different supercomputers, which is one of the primary bottlenecks for scalability. Bhatele has helped show that network topology—the path the data takes between nodes or other components—is a key factor in application performance. The topology-aware algorithms for task mapping and load balancing that he has developed have demonstrated significant performance improvements for applications such as OpenAtom, pF3D, and Qbox. This research could ultimately benefit machine architects designing new networks, computer scientists developing new algorithms, and even customers trying to decide which machine to buy based on their parallel workloads.

Recognition from Top News Service

The partnership that produced the first-of-a-kind Catalyst supercomputer was selected for an HPCWire "Best HPC Collaboration Between Government and Industry" award by readers and editors of the publication.

Computation Deputy Associate Director Trish Damkroger was named one of HPCWire's "People to Watch" in 2014. The 16 people selected by HPCWire this year are among the most talented and outstanding individuals in the HPC community. Recipients are selected from a pool of potential candidates in academia, government and industrial end-user and vendor communities. Damkroger was also the General Chair of the HPC industry's international Supercomputing Conference, SC14, which took place in New Orleans in November.

CORAL CONTRACT AWARDED AND NONRECURRING ENGINEERING BEGINS

Preparatory work commences to bring the next advanced technology system, Sierra, to Livermore

"We need diversity of technologies and vendors as well as systems that will provide value to the DOE laboratories."

On November 14, 2014, Secretary of Energy Ernest Moniz announced that a partnership involving IBM, NVIDIA, and Mellanox was chosen to design and develop systems for Lawrence Livermore (LLNL) and Oak Ridge (ORNL) national laboratories. The LLNL system, Sierra, will be the next advanced technology system sited at LLNL in the Advanced Simulation and Computing (ASC) program's system line that has included Blue Pacific, White, Purple, BlueGene/L, and Sequoia. As the next advanced technology system, Sierra will be expected to address the most demanding computing problems that the ASC Program and its stockpile stewardship mission face. To achieve this goal, the system must provide the largest capability available to ASC applications and incorporate novel technologies that foreshadow the future directions of the Department of Energy's (DOE's) large-scale systems on the path to exascale computing.

The partnership's design for Sierra uses IBM Power architecture processors connected by NVLink to NVIDIA Volta graphics processing units (GPUs). NVLink is an interconnect bus that provides higher performance than the traditional Peripheral Component Interconnect Express for attaching hardware devices in a computer, allowing coherent direct access to GPU and memory. The machine will be connected with a Mellanox InfiniBand network using a fat-tree topology—a versatile network design that can be tailored to work efficiently with the bandwidth available. Sierra is expected to be at least seven times more powerful than LLNL's current advanced technology system, Sequoia.

Sierra is part of the CORAL procurement, a first-of-its-kind collaboration between ORNL, Argonne, and LLNL that culminated in three pre-exascale high performance computing (HPC) systems to be delivered in the 2017 timeframe. The CORAL competitive bid process began in January 2014 with the release of a single request for proposals. In March, the proposals were evaluated by a collaborative team of more than 100 experts from the three CORAL laboratories, and two different options were chosen for the three laboratories. CORAL was established by DOE to leverage supercomputing investments, to streamline procurement processes, and to reduce the costs to develop supercomputers.

"Our collaborative goal was to choose two systems that, as a set, offer the best overall value to DOE. We need diversity of technologies and vendors, as well as systems that will provide value to the DOE laboratories," says Bronis de Supinski, chief technology officer for Livermore Computing (LLNL's supercomputing center). "Diversity helps to offset risk and ensure that future systems will continue to meet our evolving needs."

The Argonne and ORNL systems will help meet the future mission needs of the Advanced Scientific Computing Research program within the DOE's Office of Science, while Sierra will serve the mission needs of the ASC Program within the National Nuclear Security Administration. The ORNL system, called Summit, will have the same architecture as Sierra, which demonstrates the synergies between the missions of the two parts of DOE.

Now that the contracts have been awarded to the IBM partnership, Nonrecurring Engineering (NRE) work to maximize the impact and utility of the resulting LLNL and ORNL systems has begun. NRE includes nonrecurring expenses paid to the vendors for design and engineering

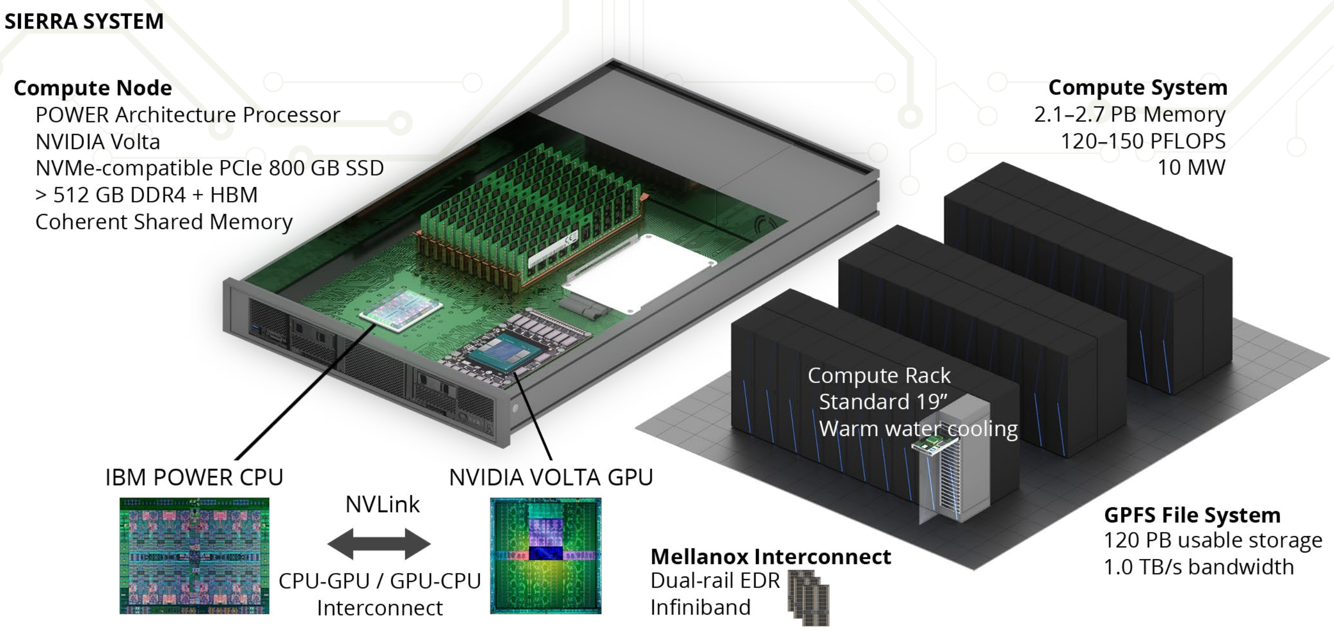

The U.S. Department of Energy awarded IBM contracts valued at $325 million to develop and deliver the world’s most advanced “data centric” supercomputing systems at Lawrence Livermore and Oak Ridge national laboratories to advance innovation and discovery in science, engineering, and national defense. These OpenPOWER-based systems minimize data motion and energy consumption, providing cost-effective modeling, simulation, and big data analytics.

BRONIS R. DE SUPINSKI
BRONIS@LLNL.GOV

milestones specific to Sierra. This separate contract precedes the “build” contract to provide accelerated or modified development to enhance usability or effectiveness of the final system. The NRE contract provides significant benefit by creating a Center of Excellence (CoE) that will foster interaction between laboratory domain scientists and vendor experts as actual applications are ported and optimized for the new architecture. The NRE contract will also support exploration of motherboard design and cooling concepts; GPU reliability, file system performance, and open-source compiler infrastructure; and advanced systems diagnostics and scheduling along with advanced networking capabilities.

Several working groups that bring together the three laboratories and the IBM partnership have been formed to ensure the future Sierra and Summit systems meet DOE requirements. These working groups are now hubs of activity, addressing the programming environment, node design, and various other topics that will ensure the usability and performance of the final systems. The CoE functions effectively as one of the working groups. Working group discussions often turn into a co-design process to meet goals. Co-design draws on the combined expertise of vendor experts, including hardware architects and system software developers, and laboratory experts, such as domain scientists, computer scientists, and applied mathematicians—working together to make informed decisions about hardware and software components. Activities have also begun for the build contract, the first milestones of which are nearing completion.

A small, early-access system scheduled for delivery in 2016 will have an earlier generation of the IBM Power processor architecture, NVIDIA Pascal GPUs, and a version of NVLink. This early access system will support interactions on several critical topics, such as development of an effective compiler infrastructure. “It will be a complete precursor system, so we can explore the capabilities and begin to deploy some early software systems on the machine,” says de Supinski.

PREPARING CODES FOR A TECHNOLOGY TRANSITION

Exploring viable paths to optimize application performance today and on next-generation platforms

The computational codes and algorithms that run on the capability-class computers at Lawrence Livermore National Laboratory (LLNL) span a wide range of physical scales and are useful not only for gaining scientific insight but also as testbeds for exploring new approaches for evolving challenges, including massive (million-way) concurrency, an increased need for fault and power management, and performance bottlenecks. To ensure that Livermore's major applications and codes capture an increasing amount of tomorrow's high performance computing systems' peak performance, significant research must be done today to study the impact of and plan mitigation strategies for future computing challenges.

"The tradeoffs involved are not just memory versus FLOPs. Both matter, but what we really need is everything in balance."

LLNL computer scientist Ian Karlin, who is conducting research in this area, explains, "If we have a complete understanding of how our applications perform today and if we identify and address the pain points and performance bottlenecks, we can extrapolate those improvements into performance gains that are portable across systems and can be sustained in future supercomputing architectures."

Unfortunately, there is no one-size-fits-all solution to improving code performance; each code package behaves differently on different computing platforms and must be considered individually. As computing architectures evolve, the job of determining exactly what factor is limiting a code's performance gets more challenging. "Our original assumption was that most of our codes are performance limited by memory bandwidth, but we have found multiple examples where that assumption does not hold true," says Karlin.

After surveying the contents of the high-energy-density physics code HYDRA, Karlin and computational physicist Steve Langer observed that none of HYDRA's code packages were reaching the bandwidth limit on LLNL's current roster of supercomputers. Nor were the packages limited by floating-point operations (FLOPs). Instead, the culprit for a few of the code packages was integer instructions, which perform tasks such as data movement or value testing (If X = Y, then Z). For example, HYDRA's diffusion package was executing integer instructions at more than 50% of peak on the IBM Blue Gene/Q Sequoia supercomputer, and two packages were executing at more than 75% of integer peak.

Codes experiencing performance limitations tend to have irregular data access patterns, the Livermore researchers found. For instance, unstructured mesh codes and algorithms often require that a processor perform more integer operations to find and retrieve data than in other types of codes. Indeed, comparisons showed that some of the code packages were hitting an integer operation limit well before they hit a bandwidth limit.

Other packages were staying well below the bandwidth, FLOP, and integer instruction limits, yet they were experiencing performance on Sequoia that was relatively worse than expected compared to how they were performing on non-specialized commodity Linux clusters. For instance, only considering bandwidth and FLOPs, the HYDRA code would be expected to perform 1.5 to 3 times worse per node on Sequoia than on a Linux cluster, but the difference was actually closer to 5 times worse. Karlin and Langer concluded that something other than bandwidth and FLOPs is limiting

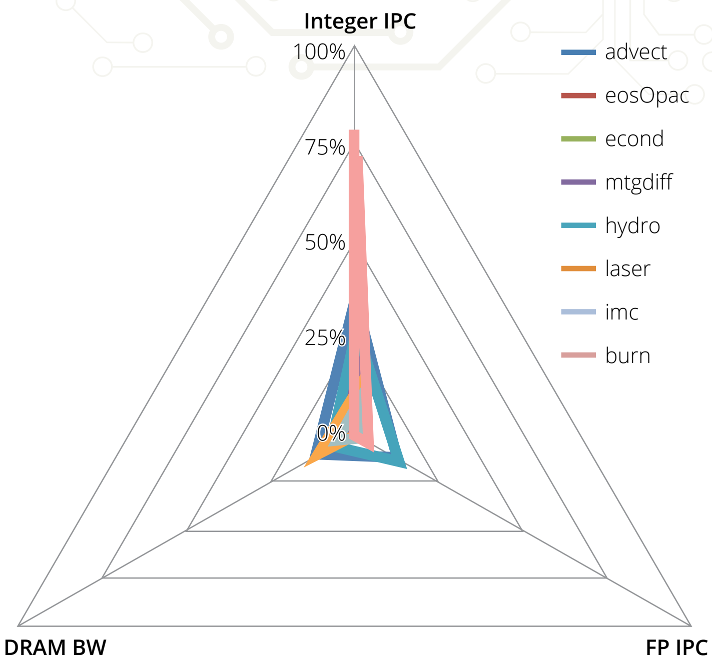

This three-dimensional chart shows the performance of HYDRA's code packages on the Sequoia supercomputer. Some packages are reaching significant percentages of integer instruction per cycle (IPC) peak; however, none are close to dynamic random-access memory bandwidth (DRAM BW) and floating point (FP) IPC peak.

performance. They suspected that many codes are also being hampered by memory latency—the time delay between when data is requested by the processor and when it begins to arrive. Latency is largely driven by the distance between the processors and memory and thus is difficult to improve. However, a physics code can be written or modified to minimize the number of times the same piece of data needs to be retrieved from memory, which could reduce the impact of a latency problem. "Because FLOPs are getting cheaper relative to data motion, we need to look more closely at algorithms that perform more floating-point operations relative to data motion," says Karlin.

To this end, Karlin leads the computer science team working on SHOCX, a Laboratory Directed Research and Development project that explores the suitability of higher-order finite element methods for future computing platforms. The computer science team works with mathematicians and physicists to tune performance and incorporate algorithmic method changes into Livermore's BLAST code. Their collaboration has resulted in transitioning to partial finite element assembly algorithms, as well as new iterative linear solvers that reduce data motion and should scale better on next-generation platforms. Overall, these changes, along with computer science optimizations, reduced data motion by more than 10 times and sped up the code by more than 4 times. According to Karlin, some of the algorithmic changes the team implemented have created additional room for improvement, which they will continue to explore.

Karlin and his colleagues' findings are already informing the development of Livermore's next advanced technology system, Sierra, and co-design efforts in general. It will take close coordination between hardware and software vendors and application scientists to create systems that balance both newly uncovered and traditionally considered performance factors. Karlin observes, "We need to take a more nuanced view. The tradeoffs involved are not just memory versus FLOPs. Both matter, but what we really need is everything in balance. We are capable of taking these other balance points into account when we build new machines."

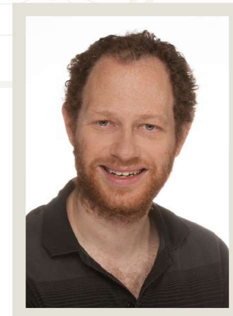

IAN KARLIN
KARLIN1@LLNL.GOV

FLUX: A FRAMEWORK FOR RESOURCE MANAGEMENT

Developing a toolset for solving data center bottlenecks

"Users will benefit from pluggable schedulers with deeper knowledge of network, I/O, and power interconnections, and the ability to dynamically shape running work."

Large computer centers that house dozens of large-scale systems with unique capabilities must have a method for efficiently scheduling their resources for use. In the case of Livermore Computing (LC), those resources include extremely large Linux clusters, such as the 46,656-core, 970-teraflop Zin, as well as myriad smaller support systems for generating, visualizing, analyzing, and storing data that is critical to fulfilling Lawrence Livermore National Laboratory's national security missions. LC developers have a long history of developing state-of-the-art software—including SLURM and its predecessors—that allows users to run and manage their simulation codes across multiple clusters. However, current resource and job management approaches cannot keep up with the challenges of increasing system scales and interplays, such as those that occur between compute clusters and file systems.

Flux is a next-generation resource and job management framework, currently in development, that will expand the scheduler's view beyond the single dimension of "nodes." Instead of simply developing a replacement for SLURM and Moab, the team behind Flux is creating a framework that will enable new resource types, schedulers, and framework services to be deployed as data centers continue to evolve.

A resource manager tracks and monitors the hardware deployed in the data center, then arbitrates access as customers submit work they would like to run. The job scheduling algorithms must not only determine when and where resources will be available that meet the user-specified requirements, they must also implement an allocation policy. Job placement in both space and time is critical to achieving efficient execution and getting the most work done for the time, power, and money spent. Flux addresses this issue by making smarter placement decisions and by offering greater flexibility and more opportunity for adaptation than current resource management software. These solutions will help scientific researchers and computing users more effectively harness the power of LC capabilities. For example, with a holistic view of the data center's input/output (I/O) bandwidth capability and utilization, Flux will be able to avoid the "perfect storm" of I/O operations that can occur when a naïve scheduler places I/O-intensive work without regard to I/O availability.

In Flux, each job is a complete instance of the framework, meaning the individual task can support parallel tools, monitoring, and even launch sub-jobs that are, like fractals, smaller images of the parent job. Because each job is a full Flux instance, users can customize Flux for use within their jobs. For example, a user desiring to launch many small, high-throughput jobs could submit a large, long-running parent job, and inside it load a specialized scheduler that is streamlined for high throughput. Panning outward in scale, schedulers operating at a larger granularity can move resources between child jobs as bottlenecks occur and employ pluggable schedulers for resource types that do not exist today.

"We are providing more capable resource management through hierarchical, multi-level management and scheduling schemes," says Becky Springmeyer, LC deputy division leader and Flux project leader. "Users will benefit from pluggable schedulers with deeper knowledge of

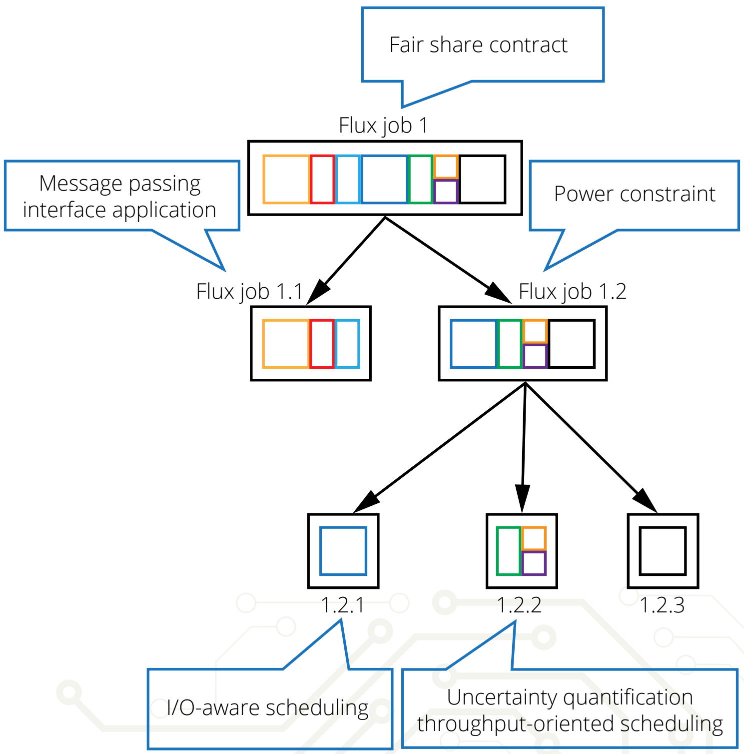

Flux, a resource and job management framework currently in development at Livermore, offers a hierarchical, multilevel approach that will allow users to dynamically schedule computing resources with methods tailored to specific workloads.

network, I/O, and power interconnections, and the ability to dynamically shape running work. One of the challenges we face in designing Flux is making sure its framework is general and extensible enough to support resources and use cases that are only now emerging in research. Our team includes researchers in power-aware and I/O-aware scheduling." Flux is being designed with input from system developers, computer science researchers, and end users, as well as external organizations that operate large computer centers. In addition, Livermore's co-design efforts with code-development tools, such as STAT, Spindle, and TotalView, will provide a highly scalable and composable code-development environment for LC users. Users will be able to pick and choose the programming tools they need and seamlessly use them together under Flux's framework. For example, users of the Kull code could scalably launch the application with Spindle and debug it using TotalView or STAT if necessary.

Flux is open-source software that will be available to high performance computing centers around the world. Already, developers outside of LLNL have shown interest in Livermore's prototyping efforts and have contributed ideas and feedback via the Flux collaboration space on GitHub. Flux developers are working with the University of Delaware to develop the I/O-aware scheduling component of Flux, and the team plans to expand its research collaborations with other academic institutions for elements such as elastic resource and job management.

Flux testing on Livermore's next generation of commodity technology systems will commence when the first system arrives in 2016.

BECKY SPRINGMEYER
SPRINGMEYER1@LLNL.GOV

IMPROVED PERFORMANCE DATA VISUALIZATION FOR EXTREME-SCALE SYSTEMS

Livermore research blends performance analysis with information visualization

"We are bringing together the visualization and performance analysis communities to develop high performance computing tools that enable our users to more easily and accurately decipher raw performance data."

Software debugging techniques, performance analysis, and optimization methods are becoming increasingly important high performance computing (HPC) tools for next-generation architectures. Such innovations are necessary to diagnose and remedy problems that slow a system's operation and reduce its utilization. Martin Schulz, a computer scientist at Livermore's Center for Applied Scientific Computing, is an HPC expert whose knowledge of performance visualization and analysis is leading to advances in both research areas.

Exascale machines will deploy millions of processing elements or cores and are expected to run a quintillion floating-point operations per second. Extreme-scale computing will be beneficial for simulating complicated physical, chemical, and biological processes that cannot be realized on today's most powerful machines. However, the increasingly complex applications and architectures required for exascale systems will also make performance analysis more difficult. One of the key challenges in assessing performance is transforming raw data into a user-friendly, intuitive format. Visualization of performance data can play a significant role in this task.

Together with colleagues at Livermore, Schulz is leading Performance Analysis and Visualization at Exascale (PAVE) efforts for providing higher fidelity visualizations of performance data across hardware, software, and communication domains. "We have developed a new paradigm of projecting and visualizing performance data obtained from one domain onto other domains for faster, more intuitive analysis of applications," says Schulz. Called the HAC (hardware, applications, communication) model, this framework provides a more comprehensive approach to data acquisition and analysis. "HAC accounts for the interplay between on-node performance, domain decomposition, and an application's intrinsic communication pattern," he says. "By taking data from each of these domains and projecting it to, visualizing it on, and correlating it with the other domains, we obtain valuable information regarding how parallel application codes behave."

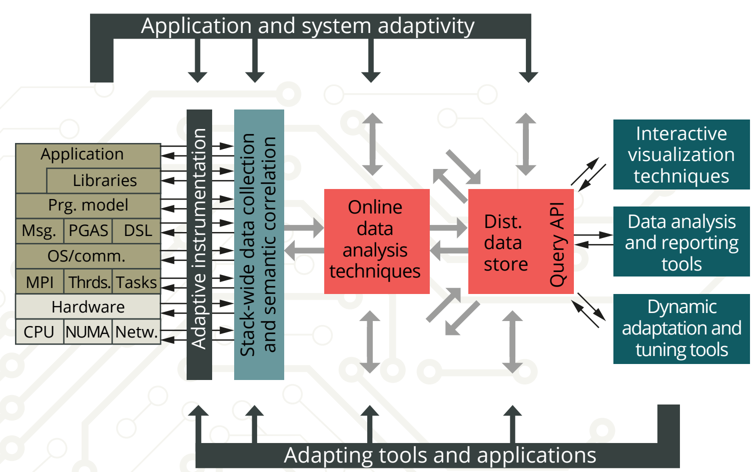

The software architecture of PIPER illustrates how performance data will be collected and analyzed at different levels of the exascale software stack. The output of the performance analysis is either displayed intuitively to users using visualization tools, or it is fed back to the exascale runtime for online optimization.

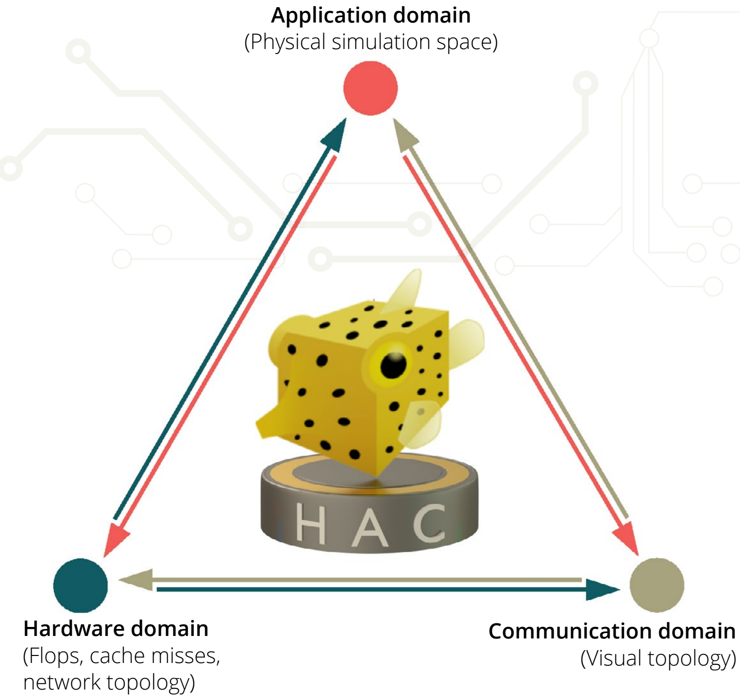

The Performance Analysis and Visualization at Exascale project provides a new paradigm for projecting and visualizing performance data from the hardware, applications, and communication (HAC) domains of a system. BoxFish is a visualization tool that represents this approach.

Seeing the benefits of exascale computing and encouraging its development, the Department of Energy's Office of Advanced Scientific Computing Research within the Office of Science started the X-Stack Program. This program promotes basic research that advances HPC capabilities in the area of programming models and tools and addresses key software challenges to building an exascale software stack. Led by Schulz, the Performance Insight for Programmers and Exascale Runtimes (PIPER) project is a multi-institutional effort that also includes Pacific Northwest National Laboratory, Rice University, and the Universities of Maryland, Utah, and Wisconsin. The project goal is to develop new techniques for measuring, analyzing, attributing, and presenting performance data on an exascale software stack. Schulz says, "Performance tools are essential to enable users to optimize application and system code and must include intuitive summaries and visualizations to help identify performance problems." PIPER will collect and analyze data at different levels of the exascale software stack. The output of the performance analysis is either displayed to users using visualization tools, such as Boxfish (based on the HAC model), or it is fed back to the exascale runtime stack for online optimization.

Additional efforts underway address techniques to optimize memory access for improved system performance. "Currently, memory systems—which provide a computer system with the means to read and write stored information—are falling behind processor development, " says Schulz. "This discrepancy has implications for a system's overall performance and power efficiency." Schulz is involved in creating MemAxes, a tool that visualizes and aims at improving communication between cores on a node and the associated memory system(s). Using sampling techniques that aggregate fine-grained data over time slices, users can extract performance information from data sources. Improved visualizations of this data help identify communication problems and speed resolution of the issue.

Ravel, another tool under development, will help to visualize execution traces on large-scale parallel programs. Execution traces record a history of process events and interprocess messages in a parallel application. Trace visualization allows users to browse this event history and search for insights into the observed performance behavior. In contrast to existing tools used for this purpose, Ravel first extracts a logical event timeline. Individual events can then be plotted in a display that enables users to depict sources, patterns, and the evolution of delays in an execution.

According to Schulz, the goal of this extensive work is to create an integrated framework for displaying and visualizing performance information that is critical to optimizing its ability to run code. To meet this objective, experts in the visualization and performance analysis communities must come together in a meeting of the minds. Toward this end, Schulz and colleagues have taken steps to integrate the two communities by encouraging activities and conferences in which experts in each field can collaborate with one another. "We are bringing together the visualization and performance analysis communities to develop HPC tools that enable our users to more easily and accurately decipher raw performance data, and with that capability, improve performance utilization," says Schulz. "Within the supercomputing paradigm, this joint work between performance analysis and information visualization is becoming its own academic subfield."

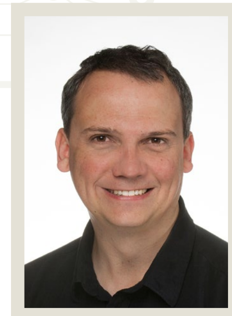

MARTIN SCHULZ
SCHULZM@LLNL.GOV

MACHINE LEARNING STRENGTHENS PERFORMANCE PREDICTIONS

Using machine learning to model and characterize the performance of adaptive applications

Unlike typical computer programs, in which only a single stream of instructions executes at a time, supercomputer programs require coordinating potentially millions of tasks simultaneously. The multiphysics simulations that run on Lawrence Livermore National Laboratory's (LLNL's) high performance machines divide models of the physical world into smaller chunks and distribute them across millions of processors. Each processor calculates the behavior of one of these small pieces, frequently sharing its results with the other processors through an ultra-low latency network. If the processors do not work efficiently, time in the simulated world slows to a crawl. LLNL scientists spend many painstaking hours ensuring that simulations can run as fast as possible, and the difficulty of this task is growing as supercomputers become more complex.

One particular challenge is writing instructions that will execute efficiently on specific hardware. However, a more vexing problem is that the performance of many applications is data-dependent—that is, the particular piece of the model being simulated on a given processor may affect its performance. Adaptation is key to many of the numerical algorithms used on high performance computers, without which many simulations would not be sufficiently accurate or feasible. The task of pinpointing the source of performance problems is complicated by the complexity of the hardware, thus the number of cases that must be optimized quickly becomes impossible to handle. Porting production codes to new platforms is already time consuming, and today's turnaround of four to six months may become years on next-generation machines.

"There is a desperate need for insightful modeling techniques that can simplify analyzing and tuning HPC applications."

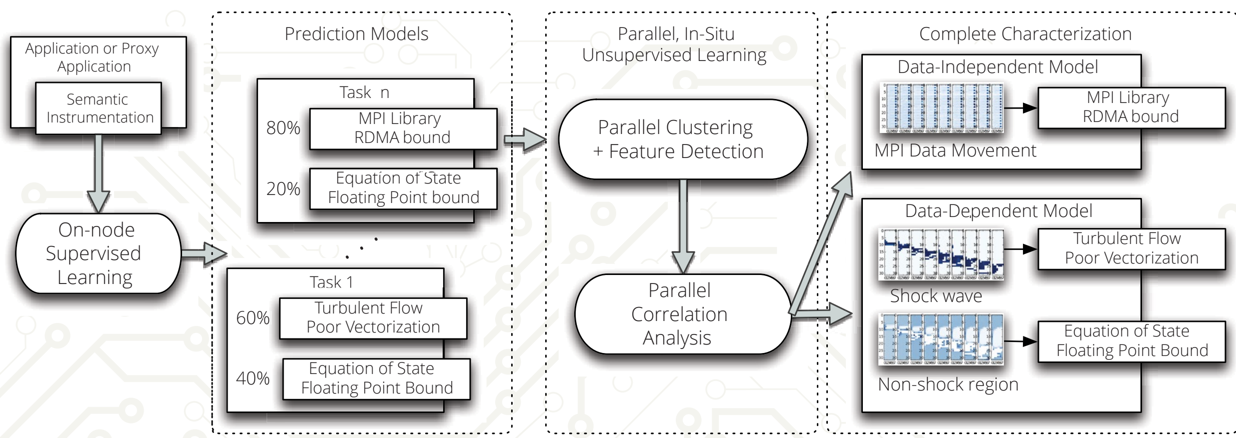

The Modeling for Exascale Applications and Data performance modeling architecture starts with a proxy application, and then builds models to predict its local performance. These models are analyzed in parallel to produce a data-dependent performance model, which can be used to understand problems in the code.

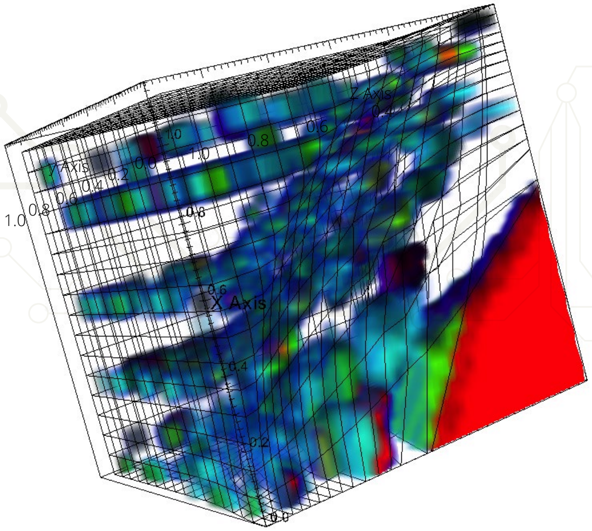

Performance data mapped onto the material mesh from the LULESH hydrodynamics application shows memory latency is highest on regions shared by multiple concurrent tasks.

Funded through the Department of Energy's Early Career Research program, computer scientist Todd Gamblin is leading a project designed to accelerate the adaptation of existing scientific simulation codes to perform on next-generation exascale systems. Gamblin's project leverages machine learning to develop algorithms that can predict the performance of application codes even when they adapt based on their inputs. Rather than requiring scientists or engineers to tune their codes manually with explicit instructions, codes tuned using machine learning could "learn" from past executions. "There is a desperate need for insightful modeling techniques that can simplify analyzing and tuning high performance computing (HPC) applications," says Gamblin.

Over the course of the five-year project, which began in 2014, Gamblin and his team will develop statistical models of HPC applications that can represent data-dependent and adaptive code behavior. Such models will help predict the performance of realistic applications and identify where performance problems originate in a system. In addition, the team will develop techniques to reduce the complexity of application models. "We are treating performance analysis as a data science problem," says Gamblin. "The ability to accurately measure, analyze, and visualize performance data will allow us to increase throughput of the entire computation center to do more complex science simulations that are important to the Laboratory's missions."

The statistical models under development will use both supervised and unsupervised machine-learning techniques to determine performance predictors and better characterize data at run times. Supervised learning methods, which learn from past examples of good performance, will be used to predict the performance of algorithms for single sets of inputs. Unsupervised models, on the other hand, look for structure in large data sets. "We will use unsupervised learning techniques such as correlation analysis and clustering to identify features that predict data-dependent performance of applications," says Gamblin. "Combined, these models will be able to predict the performance of full-system, multiscale, adaptive applications." Additional techniques will be developed for comparing the performance of different codes to one another.

To ensure the team's methodology—called Modeling for Exascale Applications and Data—will work for full production applications, it will be initially run using proxy apps (smaller, more portable applications that mimic larger codes of interest). As an example, the LULESH proxy app models a very small but computationally intensive part of the Laboratory's sophisticated ALE3D multiphysics code. While ALE3D uses many physics models that interface with a unified mesh and physics data model, LULESH uses a much simplified, smaller mesh. "Currently, there is no way to validate proxy apps to production codes and little or no work has been done on modeling the truly dynamic, data-dependent algorithms that dominate the workloads of integrated applications," says Gamblin. "By using our models to validate proxy applications against codes, we will ensure that proxy apps can be trusted as representative of production code performance."

Gamblin will receive up to $2.5 million in funding over the next five years for this project. He says, "The idea is for these models to provide simulation developers with insights that allow them to quickly optimize the performance of their code, ensuring that applications can take full advantage of the performance of future exascale machines." Work such as this is imperative for the Laboratory to continue performing the cutting-edge scientific simulations it needs to understand the complex scientific processes that underpin its mission objectives. Thanks to work like Gamblin's, Livermore will be ready to take scientific exploration to the next step when exascale systems come online.

TODD GAMBLIN
TGAMBLIN@LLNL.GOV

PLANNING HPC RESOURCES FOR THE INSTITUTION

Preparing Livermore for the next generation of HPC systems

"HPC institutional planning is not just about bringing in new hardware. We must also proactively prepare our HPC applications to make effective use of next-generation systems."

Livermore offers two main tracks of high performance computing (HPC) to staff researchers: the first, advanced technology (AT) systems, are at the leading edge of HPC technology and run the most demanding simulations that push the boundaries of computational science. These huge simulations leverage the opportunity to run in parallel on millions of processor cores and offer potentially game-changing science—as evidenced by the ground-breaking work into lithium-ion batteries that John Pask and his colleagues have done to revolutionize consumer electronics and to pave the way to the elimination of fossil fuels. The second track of HPC technology, commodity technology (CT) systems are, as the name suggests, built from commodity components that are more cost effective, easier to use, and satisfy a critical "bread and butter" HPC need at the Laboratory. While the AT systems are normally split among a few dozen users who time-share a large portion of the system at once, the CT systems may have hundreds of concurrent users who are each running multiple smaller parallel jobs to satisfy their daily research needs. The limited lifespan of both tracks of HPC resources demands ongoing institutional planning throughout the technology lifecycle.

The Lab's Multiprogrammatic and Institutional Computing (M&IC) Program brings these tailored, cost-effective HPC services to Livermore's programs and scientists. The M&IC strategy for procuring new HPC systems is to leverage the Advanced Simulation and Computing (ASC) procurement cycle for both AT and CT systems. Such procurements allow M&IC to support the Laboratory Directed Research and Development program, the Laboratory's discovery-focused Grand Challenge program, and the directorate strategic capability banks that keep the Laboratory's disciplines strong and help retain a top-tier scientific workforce. In particular, the M&IC Institutional Computing Executive Group (ICEG), the governing body of Livermore's institutional computing, plans for future HPC needs and

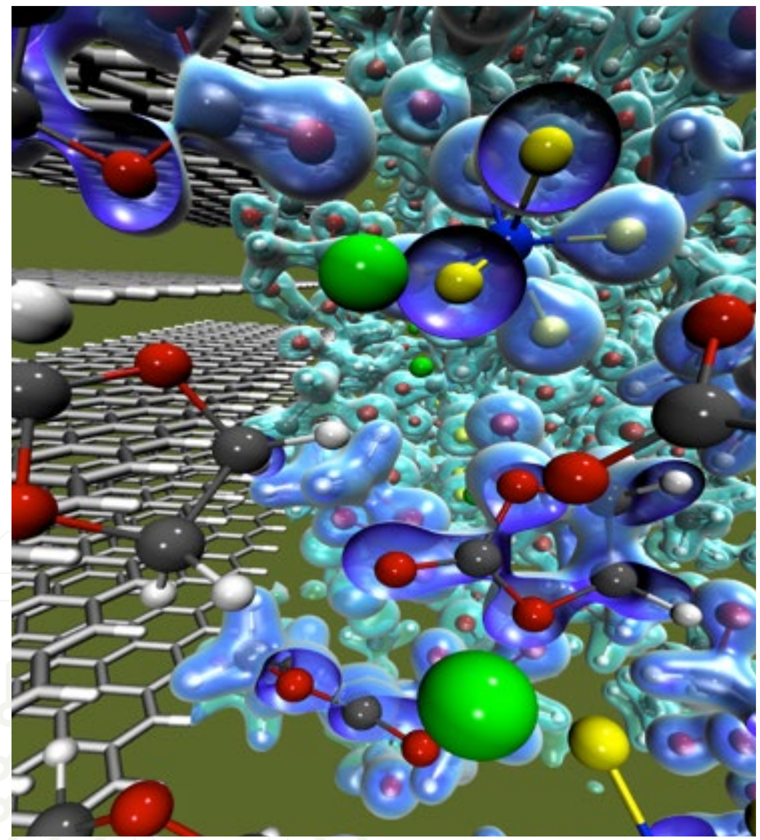

Improving lithium-ion batteries from first principles. This visualization shows the quantum molecular dynamics of the solid–electrolyte interphase layer. Visualization by Liam Krauss.

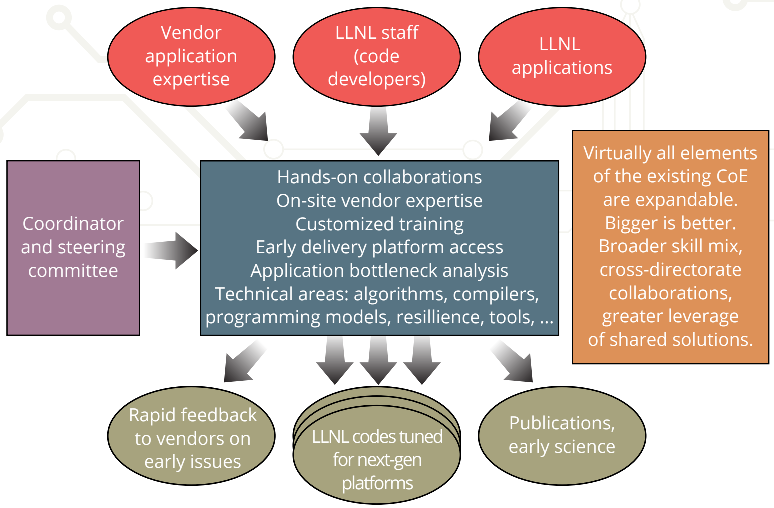

Next-generation application development will be accelerated through a "Center of Excellence" vendor partnership.

documents investment recommendations for consideration by the Lab's Senior Management Team. The ICEG plays a crucial role in recommending both appropriate architectures and the level of investment for future institutional HPC procurements.

Yet hardware planning is only half of the challenge. "HPC institutional planning is not just about bringing in new hardware," explains Brian Carnes, M&IC's program director. "We must also proactively prepare our HPC applications to make effective use of next-generation systems." HPC vendors design future technologies mainly to respond to the communication, data management, data analysis, and gaming needs of the public; however, given Livermore's advanced science and technology goals, the Laboratory's expectations and demands for HPC do not always align with the rapidly evolving mainstream architectures. To bridge this gap, Livermore code developers must partner with vendors in co-design efforts and adapt their applications to run on next-generation architectures, a difficult and time-consuming process that is eased by planning and preparing early in anticipation of the challenges.

The increased emphasis on application preparation stems from past experience. "When Vulcan arrived in 2013, we discovered that many of our institutional codes were not quite ready to take full advantage of the new advanced technology hardware," says Carnes. "That resulted in a preparedness latency we're hoping to avoid when the 2017 AT systems arrive." To help prevent this latency, the ICEG has asked the institution to participate in the ASC Center of Excellence (CoE) effort. The CoE effort brings in vendor experts to work with Livermore code developers to proactively prepare the Laboratory's HPC applications so that these applications will effectively utilize the new AT system as quickly as possible once the technology arrives. As part of this effort, the ICEG has been compiling an institutional strategic code portfolio to identify which code efforts should be targeted first for institutional funding.

Institutional HPC plays a significant role in Livermore's scientific and recruiting activities. "We want to keep our workforce engaged in leading-edge discovery for the benefit of the nation and the Lab," says Carnes. "By continuing to provide AT systems to meet institutional needs, disciplines can keep their workforce at the forefront of technology, and stay competitive for external computational opportunities."

Through strong and consistent investments, M&IC's computing resources have been used across the Laboratory to push the limits of computing and its application to simulation science and data analytics. All of Livermore's programs and disciplines have been bolstered by access to world-class HPC systems. Through these efforts, Livermore has become a preeminent laboratory in computational science, brought about ground-breaking scientific insights, and facilitated the recruitment and retention of Livermore's leading physical and computer scientists.

BRIAN CARNES
CARNES1@LLNL.GOV

ENHANCING DATA-INTENSIVE COMPUTING AT LIVERMORE

New platforms improve big data computing on Livermore's high performance computers

> *The work we have done illustrates the synergy between big data and HPC, and puts LC in a leadership position to meet the needs of both camps going forward.*

Increasingly large datasets and the growing diversity of data drive the need for more capable data-intensive computing platforms. At Livermore, this concern takes on additional significance since the Laboratory's work uses big data to pursue a safer, more secure world for tomorrow. The Laboratory's high performance computing (HPC) capabilities offer exceptional opportunities for data analysis and processing. However, the specialized nature of Livermore Computing's (LC's) hardware presents challenges to traditional approaches to data-intensive computing.

"LC's HPC systems have been tailored to run scientific simulations very well. Unfortunately, this is not the optimal architecture for many data-intensive computing applications," explains Robin Goldstone, a member of LC's Advanced Technologies Office. Goldstone and her team have been exploring solutions that can bring LC's expertise to bear on the Laboratory's growing demand for big data computing platforms. "We recognized that we needed to take a look at Hadoop, a solution that has been requested by numerous customers. We set out to see how we could tweak our traditional HPC systems to meet the needs of these big data customers."

The Hadoop ecosystem—which includes MapReduce, Hbase, and newer frameworks such as Spark and Storm—has gained widespread adoption in part due to its relatively modest computing hardware requirements. Clusters of inexpensive commodity servers with local hard drives can run Hadoop effectively since the software has been designed from the ground up to tolerate failure. In contrast, HPC applications typically do not tolerate failure, which causes HPC systems to demand more expensive hardware and complex recovery mechanisms that will achieve resilience. For these reasons, HPC systems are typically dismissed as being "overkill" for frameworks like Hadoop.

However, since LC already has these HPC systems deployed, the question becomes whether such systems can efficiently run Hadoop in place of an entirely separate set of commodity-class resources. To answer this question, LC purchased a small, generic Hadoop cluster in order to gain experience in deploying and managing such a system. This cluster, named

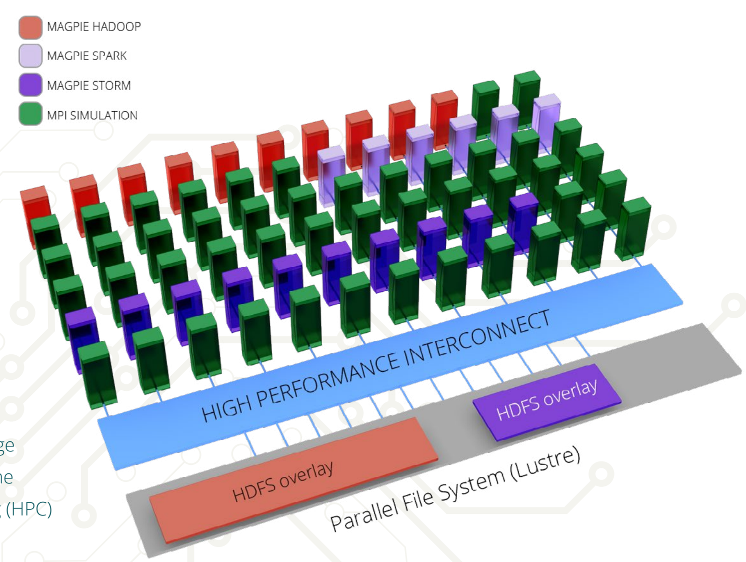

Magpie enables running Hadoop-style applications alongside traditional message passing interface (MPI) applications on the specialized high performance computing (HPC) systems at Livermore.

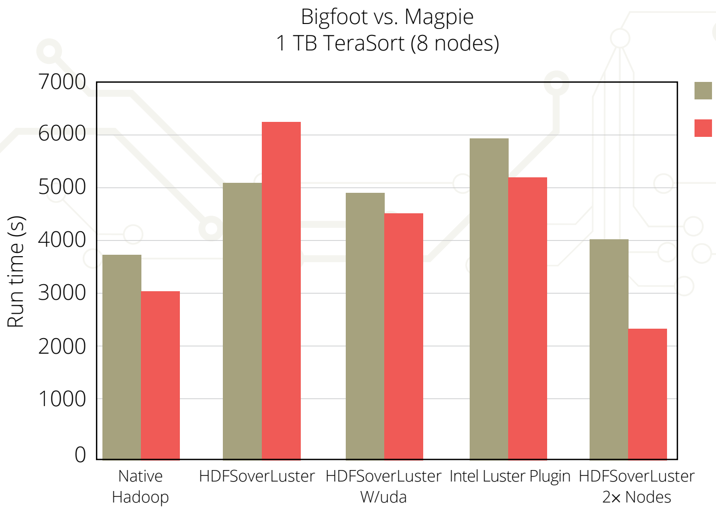

Results of the TeraSort benchmark tests that ran on Bigfoot and the HPC cluster. The various tests sought to determine which configuration achieved the most efficient approach to conducting big data analysis on an HPC system.

ROBIN GOLDSTONE
GOLDSTONE1@LLNL.GOV

Bigfoot, allowed Goldstone's team to evaluate the operational impact of supporting this platform while simultaneously providing a place to conduct trials between commodity and HPC systems.

The HPC-centric approach involved the development of a software package, named Magpie, which allows Hadoop and similar data analytics frameworks to run on LC's HPC systems. Magpie accomplishes this task by instantiating the framework within the context of a batch job—rather than on a persistent, dedicated cluster—and by reading and writing from the Lustre parallel HPC file system instead of local disk drives.

With both Bigfoot and Magpie in hand, Goldstone's team assessed whether Magpie could replace the need for dedicated Hadoop clusters at Livermore. Using the de facto Hadoop benchmark, TeraSort, the team ran a sort function on the Bigfoot cluster and then on an equivalent number of nodes on one of LC's HPC systems. The team additionally formulated several TeraSort configuration options and Magpie tunables to test the two systems. The results showed that the benchmark performed at best 50% slower on the HPC cluster than on Bigfoot when using an equivalent node count.

The team performed two additional experiments with more encouraging results. In the first test, the TeraSort benchmark ran on the HPC cluster using double the number of nodes as on Bigfoot. This time, the HPC cluster won the comparison, achieving a 33% reduction in runtime over Bigfoot. While this success might not appear to be a fair test, it demonstrates the "surge" capability that LC offers users—with thousands of cluster nodes already deployed, LC can quickly accommodate a customer's need to scale up their analysis. To do the same on a dedicated Hadoop cluster would require months of lead time to purchase and deploy additional hardware.

In the second experiment, the team employed Catalyst, a new LC HPC system equipped with non-volatile random access memory, more commonly known as "flash storage." Each Catalyst compute node contains 800 GB of high-performance Peripheral Component Interconnect Express-attached flash storage, which Magpie can use in place of Lustre for storing Hadoop intermediate data files. It is this intermediate input/output (I/O) that puts the most strain on the Lustre file system, so the Magpie developer conceived that a modest amount of fast local storage could significantly improve I/O performance. The team's testing validated this theory, achieving a 2× performance improvement when running a TeraSort across 295 of Catalyst's nodes that use the local flash storage.

"Reducing data motion is not just a big data issue," says Goldstone, elaborating on the value of these outcomes. "Our HPC simulation customers are also feeling the pain of moving data, and we see architectures like Catalyst and the future Sierra system as the path forward. The work we have done illustrates the synergy between big data and HPC, and puts LC in a leadership position to meet the needs of both camps going forward."

INTERWEAVING TIMELINES TO SAVE TIME

A new parallel-in-time method dramatically decreases solution time for various simulations

> *We're performing time-stepping simulations by combining multiple timelines of differing accuracies to get the solution significantly faster. The key point to the algorithm is that by "braiding" together the timelines in parallel, you don't have to solve any individual timeline sequentially.*

Computer simulations in numerous areas rely on the solution of time-dependent problems, from the evolution of a supernova to the propagation of seismic waves to the metabolism of pharmaceuticals in the human body. Generally, solutions to these problems are reached by addressing one time step after another, in order. Unfortunately, the sequential nature of traditional time stepping has left the programmers who maintain and enhance these codes and the scientists employing them with a shared dilemma: how to coax more efficiency from such codes during a pivotal transitional period in computer architecture.

For decades, these codes have benefited from a steady increase in computer chip speeds with each new generation of computer system. So long as the chips continued to get faster, the time it took to compute each individual time step was reduced. Spatial refinements made to a code to increase accuracy often must be balanced by temporal refinements, but chip speed increases enabled more time steps to be calculated without increasing the total compute time. Chip speeds have plateaued in recent years, though, meaning that any further increases in the number of time steps will simply increase the total compute time.

"Now, to take advantage of bigger machines and see speed-ups, all parts of these codes need to be able to take advantage of parallelism," explains Lawrence Livermore computational mathematician Jacob Schroder. This includes time, one of the final frontiers in algorithm parallelization. Schroder and a team of Livermore scientists, in collaboration with researchers at Memorial University and Belgium's Katholieke Universiteit Leuven, have developed a method for solving all of the time steps simultaneously, with the help of a new multilevel algorithm called XBraid and the massively parallel processing capabilities of today's high performance computing (HPC) systems. The approach has already been shown to dramatically decrease solution time for various simulations, some by as much as tenfold.

While they are not the first researchers to explore the idea of solving time steps out of order, their method offers advantages over most other methods. In particular, XBraid obviates the need for a complete code rewrite, which for a complex simulation code could be an enormous time investment. Schroder observes, "One of the real novelties of our approach is that it's largely nonintrusive. Our goal is for users to come to us with an existing sequential time-stepping application they've been working on for

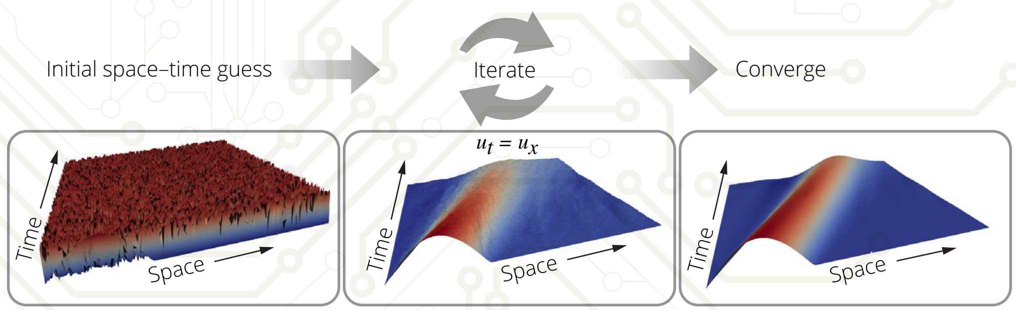

Visualization of the XBraid method for the advection equation, which describes a sine wave moving from left to right. A random initial guess across the entire space–time domain is iterated on until it converges to the solution. This is in stark contrast to sequential time stepping, which would serially propagate the sine wave from one time step to the next.

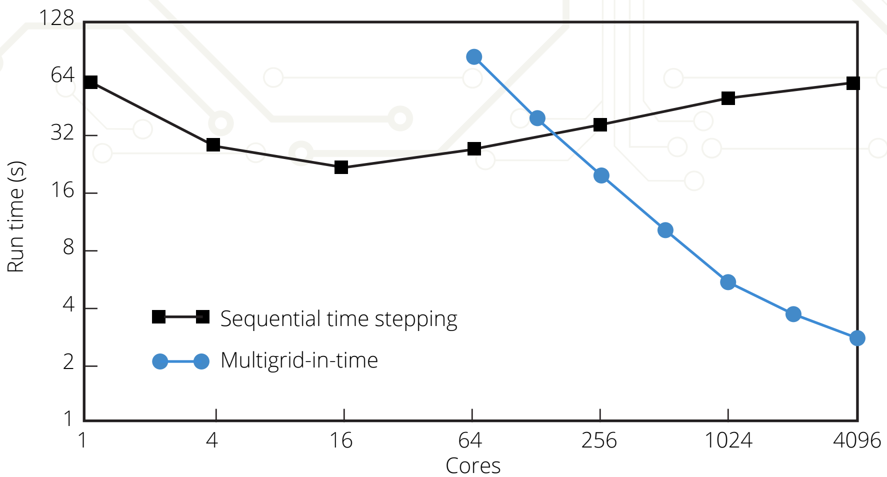

As shown in this heat equation example for a 129 × 129 × 16,385 space–time grid, there is a crossover point beyond which significant speedups are realized through parallel time stepping. The parallel-in-time approach uses more memory and requires more operations than sequential time stepping and thus is slower if only a few processors are available. However, due to its increased parallelism, as more processors become available, it is able to use them fully.

JACOB SCHRODER

SCHRODER2@LLNL.GOV

maybe 10 or 20 years. Then, all they have to do is wrap it with some code using a conceptually simple interface."

For this effort, the XBraid team applied their extensive experience in developing scalable multigrid spatial solvers to the time dimension. Multigrid approaches solve systems at various levels of granularity. "We're combining multiple timelines of differing accuracies to get the solution significantly faster. The key point to the algorithm is that by 'braiding' together the timelines, you don't have to solve any individual timeline sequentially," explains Schroder. The solver begins with a "guess" of the solution and then uses an algorithm to generate an improved guess. This procedure repeats until the iterations converge to the solution. By solving coarser (and less computationally expensive) versions of the same problem and feeding those solutions back into the finer scale version of the problem, the iterative process is accelerated. Since the computational cost is proportional to the number of equations, a large problem can be solved in the same time as a smaller one simply by increasing the number of processors working on the calculation. Importantly, the solutions from XBraid and sequential time stepping are identical, up to a user-defined tolerance.

The benefits of this parallel-in-time approach are twofold. By incorporating XBraid into parallel codes solving for large numbers of time steps, researchers eliminate a performance bottleneck and reach a solution more quickly. They can also make better use of today's and tomorrow's HPC systems. "Exascale computers will be even more massively parallel than Sequoia," says Schroder. "We need our algorithms to be able to make the best use of that level of parallelism."

In the year and a half since the project's genesis, the XBraid team has developed foundational theory for the methodology with support from the Department of Energy's Office of Science and developed proof of concept demonstrations and software using funding from Livermore's Laboratory Directed Research and Development program. Both linear and nonlinear parabolic problems, their first target area, have demonstrated excellent results, with XBraid solving the problems up to 10 times faster than before. XBraid also shows promise for fluid dynamics calculations. For instance, a project to develop a next-generation helicopter design code in collaboration with the Department of Defense, while still in early stages, has demonstrated an eightfold speedup with XBraid.

In 2015, the team will continue to expand and improve XBraid and the types of problems that it can help solve. This will include validating the parallel-in-time method on hyperbolic problems such as wave propagation, a challenging problem type for multigrid methods. They also hope to expand awareness and usage of the open-source XBraid software in the broader scientific community. "Parallel in time is a timely topic right now," Schroder says, "as people grapple with big changes in computer architecture."

REMAPPING ALGORITHM BOOSTS BLAST SIMULATIONS

New approach enables BLAST to simulate larger time steps in a way that preserves accuracy

> *"At higher orders, unexpected things can happen. A lot of questions people thought were settled suddenly became very interesting."*

Shock hydrodynamic simulations, used to understand the behavior of materials at very high pressures, provide essential information for Lawrence Livermore stockpile stewardship work and complement and support laser experiments performed at Livermore's National Ignition Facility (NIF). However, high compression rates and hydrodynamic instabilities make accurate modeling challenging. Computational mathematician Tzanio Kolev explains, "Implosions at NIF generate enormous pressure. This produces shock waves generating and interacting with each other, and dealing with these discontinuities computationally is difficult. Typically these are modeled by an Arbitrary Lagrangian–Eulerian (ALE) approach, where during the Lagrangian phase the mesh evolves with the simulation. One of the challenges with this method is that, if we're not careful, the mesh will intersect and tangle. Also, we have multiple materials to represent. Modeling the physical discontinuities at shock fronts and material interfaces is challenging mathematically."

Achieving higher-quality simulations of shock hydrodynamics problems requires the development of more advanced numerical algorithms. To that end, a team of computer scientists, mathematicians, and physicists from the Computation and Weapons and Complex Integration (WCI) organizations, led by Kolev and WCI's Rob Rieben, created a new high-order ALE framework as an alternative to the standard low-order ALE solution methods. The framework is based on a high-order finite element approach, which uses simple element equations over many small curved geometric subdomains to approximate a more complex equation over a large domain. "Higher-order elements have more vertices, or control points, positioned around them, allowing us to curve the element boundaries and the geometry inside," says Kolev. "This helps us more accurately follow the material flow."

The new numerical methods were implemented in the Livermore-developed BLAST research code. Over the past few years, the team has successfully demonstrated that the high-order ALE framework—and overarching BLAST code—can produce efficient, accurate, and robust simulations of a variety of challenging shock hydrodynamics problems, but they have also identified opportunities for improvement. In 2014, the research expanded to include performance analysis and improvement, led by Ian Karlin, and a Laboratory Directed Research and Development (LDRD) program strategic initiative (SI), led by WCI's Bert Still. Under the new SI, the team has incorporated a new high-order Discontinuous Galerkin (DG) remapping algorithm into the ALE framework. This enables BLAST to simulate larger time steps and to more accurately address mesh elements containing multiple materials.

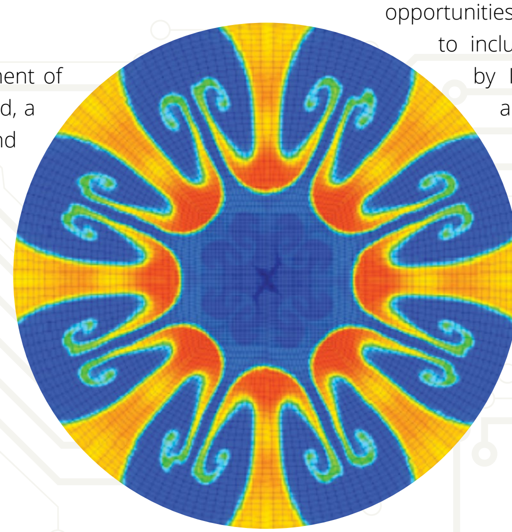

An inertial confinement fusion problem uses fourth-order finite elements to simulate the interactions of multiple materials.

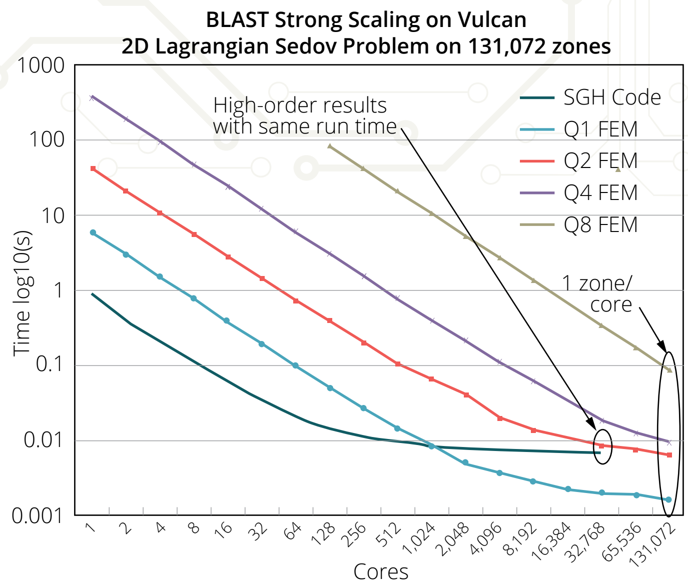

Because the communication overhead for high-order methods is not as significant as it is for low-order algorithms, it is possible to achieve excellent strong scaling results all the way down to the limit of a single computational zone, or element, per processor.

Sometimes mesh elements are unable to conform to a function as well as is desired, particularly for functions with steep gradients. For those portions of the calculation, the new DG algorithm "stops time" and institutes a remap phase, during which the function stays the same while the mesh evolves. Once it has been translated to a more appropriate mesh, the calculation continues from the point where it left off. Notes Kolev, "With high-order ALE, we can push the Lagrangian phase much farther than with low-order ALE codes, but finite elements thin out and time steps become small eventually. With our remap approach, we can run with much larger time steps in a way that preserves accuracy."

When the mesh changes, it can result in multiple materials within the same element, which can produce mathematical difficulties. The team has solved the problem by representing each material as a high-order function for purposes of remapping. "With our newest work, we are able to capture a mix of materials at a very detailed level," notes Kolev. Overall, the remapping algorithm has demonstrated excellent parallel scalability, geometric flexibility, and accuracy on a variety of model problems and single- and multiple-material ALE simulations. One of the most demanding calculations to date was a three-dimensional BLAST simulation involving three materials performed on 16,000 cores of the Vulcan supercomputer.

The team has now begun to apply high-order solution methods to other types of physics, beginning with radiation diffusion. "We're gradually extending it to more and more pieces of what a realistic multiphysics simulation would require," says Kolev.

Through their resolution of some long-standing numerical challenges in shock hydrodynamics simulation, the algorithms developed through this research support the Laboratory's national and energy security missions, but they also benefit research on parallel performance and next-generation computer architectures. Higher-order methods have greater FLOP/byte ratios, meaning that more time is spent on floating-point operations relative to memory transfer, an important characteristic of numeric algorithms for extreme-scale computing. "On top of their mathematical benefits, higher order lets us increase the arithmetic intensity," Kolev explains. "We can dial in how much we do with the data within each processor, each element, and each integration point. In fact, higher-order methods can often run in the same amount of time as lower order due to the increased computational efficiency." With LDRD funding, Computation and WCI researchers are characterizing and optimizing the performance of BLAST's high-order methods on different high performance computing systems.

High-order methods remain an area of opportunity for the researchers. Notes Kolev, "At higher orders, unexpected things can happen. A lot of questions people thought were settled suddenly became very interesting."

TZANIO KOLEV
KOLEV1@LLNL.GOV

CATCHING BUGS WITH THE AUTOMATED TESTING SYSTEM

Automated system provides streamlined testing and performance analysis

"Many of our programmatic codes at Livermore are too massive, dynamic, and complex to develop with manual testing. We simply had to automate the process."

Software developers use regression testing to help uncover new bugs in an application after enhancements, patches, or other changes have been made. It is an essential but repetitive activity, as it entails rerunning a representative set of test cases each time the code is updated to ensure that the code continues to run the same way and produce the same answers as before. For a large code that is updated many times a day by a sizeable team of computer scientists, such as some at Lawrence Livermore, regression testing would be impossible to keep up with if done in a traditional fashion.

Shawn Dawson, a computational scientist working with the Weapons and Complex Integration (WCI) Directorate, notes, "Many of our programmatic codes at Livermore are too massive, dynamic, and complex to develop with manual testing. We simply had to automate the process." Dawson is in charge of the maintenance and enhancement of one such solution—the Automated Testing System (ATS), a Livermore-developed, scalable, Python-based tool for automating the running of tests for an application. ATS currently operates on Livermore's Linux clusters and IBM Blue Gene machines and can be used to execute thousands of tests with a single command.

By helping projects to meet their milestones and maintain code integrity, ATS serves as a fundamental enabler for important projects in WCI and

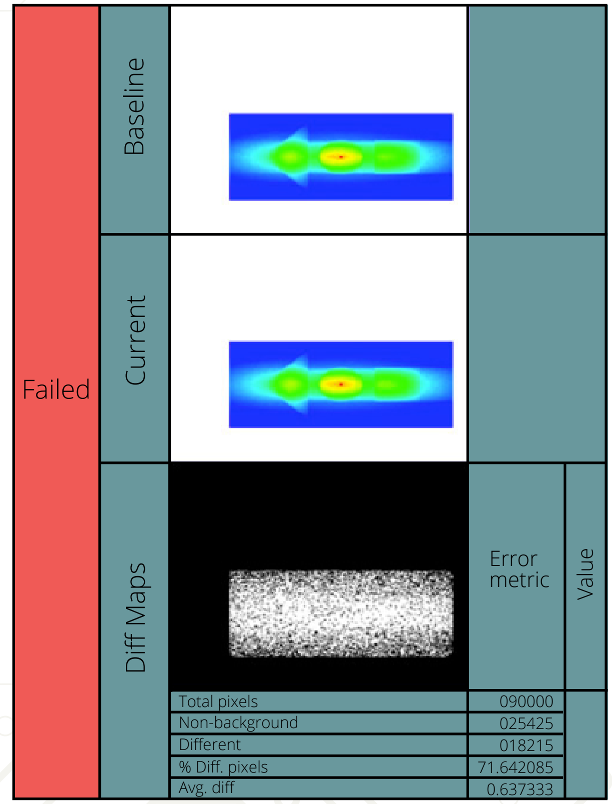

An extension to the Automated Testing System (ATS) enables the comparison of two images (top and middle) using the distributed, parallel visualization and graphical analysis tool, VisIt. Differences in the images are reported at the bottom. While the two images appear nearly identical, VisIt—through ATS—is able to detect subtle image differences and report these to the developers.

the Physical and Life Sciences Directorate. Tens of thousands of tests a day are run at Livermore with the help of ATS. Moreover, these project teams are able to make a more efficient use of resources by adopting the common testing framework, rather than each project developing their own ad hoc solutions.

It was dissatisfaction with commercial testing products that spurred the development of ATS a decade ago. Dawson, who is both a developer for and a customer of the tool, notes, "Commercial codes didn't suit the needs of the codes we develop. We run on unique systems, not simply desktop PCs or Macs. Also, those codes don't have the flexibility to extend testing or compare data." With ATS, Dawson and his colleagues have been able to introduce various extensions into the tool, such as—most recently—a capability for testing and comparing image files.

Not only do Livermore developers test the code before committing to each proposed change, they also run multiple larger nightly test suites on several computer platforms to test the correctness and performance of the code with different configurations and compilers. ATS's customer codes average about 4,000 tests each night, and each test produces multiple plots, resulting in a substantial amount of data. ATS organizes the output of these nightly runs into a text-based list and emails it to the project team for review the next morning. It also creates a graphical web-based report for the team to compare the curves side by side and quickly spot bugs introduced to their code. Scientists can easily post-process the testing results for deeper analysis.

Dawson's team strives to enhance the tool on an ongoing basis, for instance by incorporating user feedback. "Testing is a complex issue. We can't ever test 100% of the code, as the lines depend on the type of data input, but we are testing the large majority of our use cases. When a user finds an area we missed, we incorporate it into the test suite." Customer needs also drive the development of new features for the tool. In 2015, the ATS team will begin archiving metrics related to each test run, a much-requested capability. This requires connecting ATS to a Laboratory-developed database called Tracker. With this enhancement, users will be able to run queries and generate reports on the data.

While ATS was designed for regression testing, it also can be used to look at code performance over time and across platforms. "The tool is really useful for comparing one machine to another. I use it frequently to compare compiler releases," says Dawson. He anticipates employing ATS for acceptance testing for the new compilers and libraries during the integration of Livermore's next advanced technology system, Sierra. "ATS was designed to be easily extendable to new machines," he adds. "As we move to the next advanced technology system, ATS will be ported to the machine, and we will provide support for it."

ATS will also support code preparation for Sierra and other future machines. ATS user Patrick Brantley notes, "ATS is a very flexible framework that allows us to deliver a more reliable simulation code to our users. As we contemplate potentially significant changes to our simulation code to enable its use on advanced computing architectures, we will continue to rely on ATS to help us maintain the integrity of our code."

The author wishes to acknowledge the team of computer scientists working on codes in WCI who have developed and enhanced ATS.

SHAWN DAWSON
DAWSON6@LLNL.GOV

NIF DEPLOYS NEW ADVANCED RADIOGRAPHIC CAPABILITY

Thoughtful controls design and testing enables smooth integration for new laser capability

"One of the reasons ARC was so challenging was that all of the pieces of every control system software release were interrelated."

Less than two-millionths of a second elapse from the moment the initial laser burst is created to the completion of a typical high-energy-density science experiment at Livermore's National Ignition Facility (NIF). Obtaining precise information about the physical processes occurring in the target during this brief span has necessitated the development of a new generation of ultrafast, ultrahigh-resolution diagnostic devices such as the Advanced Radiographic Capability (ARC), which is currently being deployed on NIF. When complete, ARC will be used to produce short bursts of intense x-rays to image NIF targets as well as to enable new experiments in fusion ignition and high-energy-density stewardship science.

ARC is being integrated on an existing set of four NIF beam lines—known as a quad—to facilitate reuse of the existing NIF main amplification system. The modified NIF quad has been designed to rapidly and autonomously switch between ARC (short pulse) and NIF (long pulse) operations based on experimental requirements. In ARC mode, it will use a split-beam configuration, propagating two short-pulse beams for each NIF aperture. Staggering the arrival of the eight ARC beamlets onto special back-lighter targets will enable creation of an x-ray "movie" with tens of trillionths of a second resolution.

Incorporating ARC functionality into NIF has been a substantial undertaking, requiring major controls development and other system updates. A team of software engineers, led by Gordon Brunton, worked for 18 months to develop and deploy the system and software controls enhancements necessary for ARC operations. Integration had to be planned carefully to minimize disruption to ongoing NIF shot operations. Brunton notes, "One of the reasons ARC was so challenging was that all of the pieces of every control system software release were interrelated. Failure to complete one piece on time could have jeopardized the whole schedule. But due to the work of the whole team, we were able to meet every milestone."

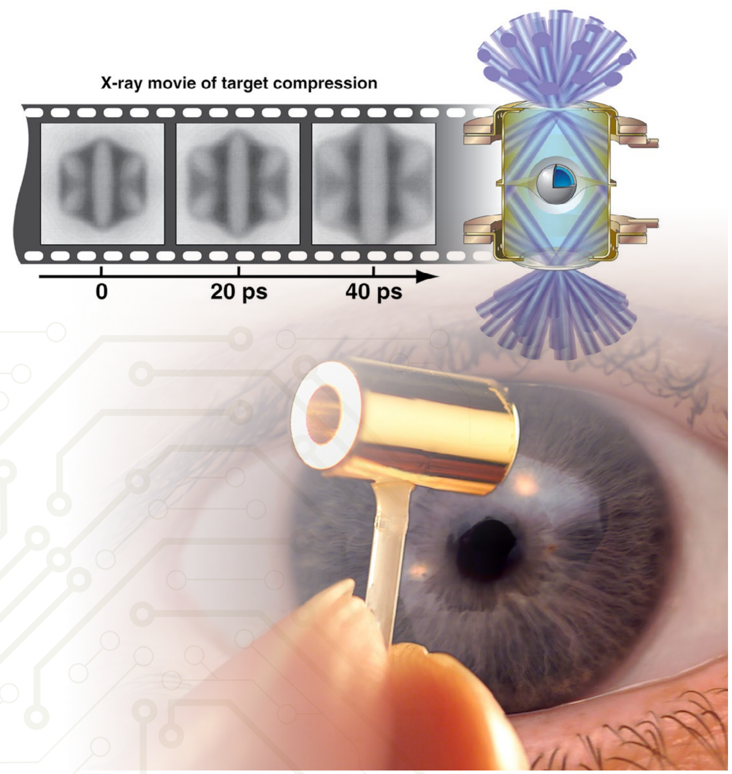

By staggering the arrival of the Advanced Radiographic Capability's beamlets onto back-lighter targets, x-ray images resembling a movie strip (upper image) will be acquired, capturing the dense core of fuel in a National Ignition Facility target (lower image) as it is being compressed.

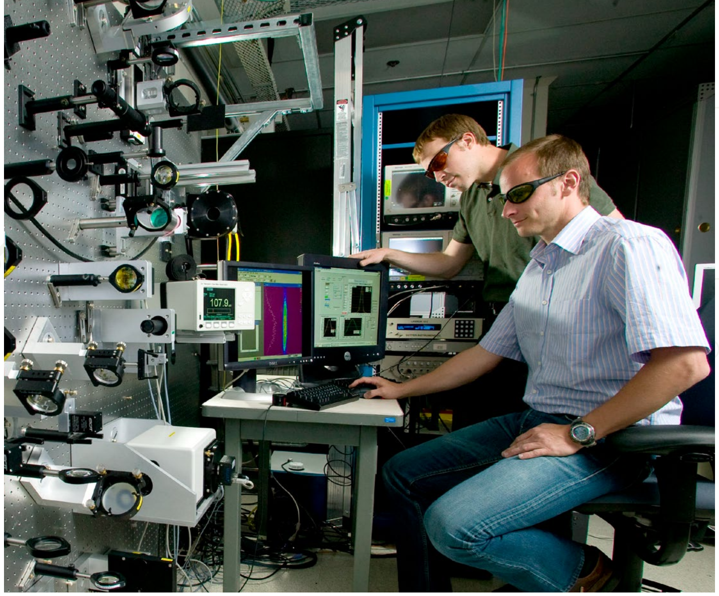

Ultra-short laser pulse diagnostics such as Frequency Resolved Optical Gating, shown here, necessitated substantial software development.

ARC increases the amount of complex laser equipment that must be monitored and coordinated for a successful NIF experiment. In fact, the number of control points for electronic, optical, and mechanical devices on the modified NIF quad grew by 70%. "It felt like building another new beam line from scratch," says Brunton. Throughout the system design phase and in close coordination with the system experts, the team sought opportunities for leveraging existing types of control hardware. This reuse strategy significantly reduced the deployment schedule, cost, and risk. "Writing software for a new control type takes months, whereas implementing control points for an existing type only takes days," Brunton adds. "Overall, the reuse saved us many years of effort."

Operationally qualifying the ARC system requires verification of many aspects of the short pulse. To allow these verifications to be performed and continuously monitored on each ARC shot, a comprehensive suite of short-pulse diagnostics have been integrated into the system, several of which were developed specifically for ARC. Due to the specialized equipment, the new and modified software deployed—representing about 15% of the new control points—primarily related to diagnostics.

All NIF experiments are performed with the support of the experiment automation system. This suite of software applications manages the full life cycle of a laser shot, from shot goal acquisition to shot data archiving. The software is based on a data-driven workflow engine within a state machine model. The model breaks a shot cycle into numerous operational phases, and each phase is populated with workflow nodes that perform well-defined, reusable, automated activities (for instance, pulse shaping). Shot goals are used to autonomously reconfigure the system based on experimental needs. It is this data-driven flexibility that has facilitated the integration of ARC into the workflow automation with minimal software framework modifications.

Due to the scale of the ARC system modification, Brunton's team knew that comprehensive off-line qualification would be critical for minimizing risk to NIF operations. For testing, every control point in NIF was given an emulated equivalent that closely mirrored the behavior of the real hardware device. During off-line integration and formal quality assurance, they used these capabilities extensively to complete qualification of most software modifications over multiple phases. This strategy greatly reduced the on-line NIF facility time required for qualification of the modifications. The final major release phase concluded with shot automation capability deployment in the summer of 2014.

An extensive series of commissioning shots are scheduled on ARC over the next several months, and the system is expected to begin experimental operations later in the year. Brunton's team, which will continue to support ARC controls through commissioning and operations, is pleased to see the ARC effort at last approaching completion. Brunton remarks, "Even though significant progress has been made on the fusion ignition challenge, ARC remains an important diagnostic capability. It will help us understand more about inertial confinement fusion experiments by providing improved data on important parameters such as compression, symmetry, and fuel mix. In addition, it opens up opportunities for further experimentation, allowing us to look deeper into dense, novel materials that current diagnostic methods can't see into."

GORDON BRUNTON
BRUNTON2@LLNL.GOV

LEVERAGING DATA-INTENSIVE COMPUTING FOR SLEUTHING SEISMIC SIGNALS

Developing core technologies to change the way geophysical monitoring systems work and data is processed

Since the 1960s, scientists at Lawrence Livermore National Laboratory have been developing methods to locate, identify, and distinguish nuclear explosions from earthquakes and other types of seismic activity. Today, high performance computers and new data-intensive computing techniques are improving our country's ability to monitor nuclear explosions and allowing researchers to better define seismically active faults.

"An important step for utilizing big data has been the use of distributed-data platforms, such as Hadoop and Storm."

In the last decade, the ability to monitor international seismic events has grown more computationally challenging. The data collected from seismology sensors deployed around the world has been increasing at exponential rates, and the advent of ultra-cheap sensors promises to maintain this trend. With ubiquitous sensors comes the need to store and process enormous data volumes.

"Livermore and other organizations tasked with tracking seismicity around the world now must monitor, analyze, and respond to data that is one million times more voluminous than what we were dealing with 10 years ago," explains Stan Ruppert, who leads a team of computer scientists tasked with developing data analytic capabilities to address the problem. "Today we are handling millions of signal sources, hundreds of thousands of recording sensors, and billions of derived parameters and measurements. The growing data volume is overwhelming the community's traditional processing techniques."

One of the most significant problems associated with combining data from different sensors is organizing the metadata that come with them. Often, the data collected are of unreliable quality, updates are unpredictable, key values are inconsistent, and data are provided in inconsistent formats. In addition, duplicate and conflicting metadata must be coalesced into a single logical entity. Currently, these issues are solved by manual data fixes or complex data manipulations buried in application code—methods that are unsustainable given the substantial increase in incoming sensor data and the accompanying exponential growth in metadata.

A team of computer scientists supports Livermore's Geophysical Monitoring Program, including (from left) Douglas Knapp, Douglas Dodge, Stan Ruppert, Jessie Gaylord, and Steven Magana-Zook.

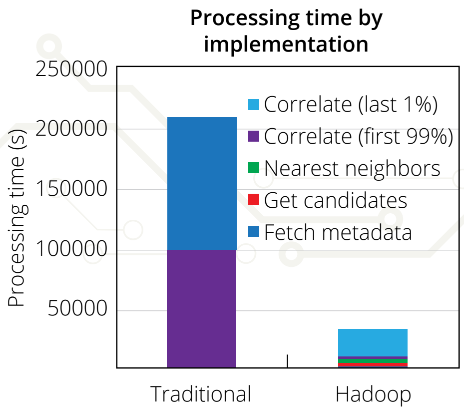

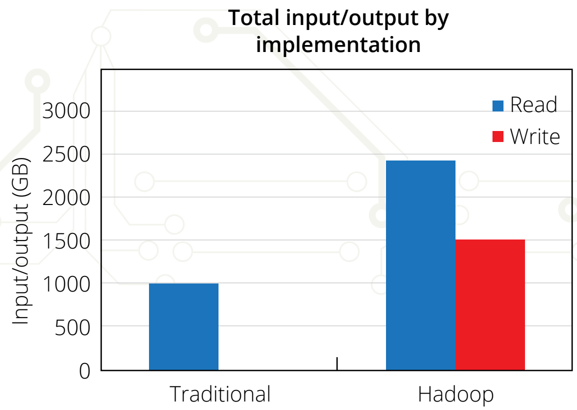

Using Livermore Computing's Bigfoot (Hadoop) cluster, scientists achieved a 19× performance increase and better scaling on a 1-terabyte dataset of seismic events.

STAN RUPPERT
RUPPERT1@LLNL.GOV

To address this issue, the Livermore team is designing a new metadata processing architecture that easily handles data provenance for incoming data in any format, isolates downstream applications from ingestion and transformation details, and preserves the complete provenance of the data. The proposed architecture leverages NoSQL data warehousing techniques developed for large-scale text processing and analysis. The new metadata processing design will provide a highly scalable and maintainable solution that also improves the quality and availability of critical sensor metadata.

Big data techniques are also being used to help understand new seismic events by comparing the waves generated to "templates" of past underground nuclear tests and other seismicity. A wide variety of nuclear proliferation/detonation detection capabilities rely on recognizing the specific signature of an event in a large volume of streaming real-time or archived data. If templates of the target signatures can be defined, then automated template matching (e.g., correlation, pattern recognition) offers a potential way to detect, locate, and identify the source of interest, even when the signal is buried in a noisy background.

An important step for utilizing big data has been the use of distributed-data platforms, such as Hadoop and Storm. A partnership between Livermore's Geophysical Monitoring Program (GMP) team and Livermore Computing (LC) has helped demonstrate the potential of data-intensive computing to solve earth science problems.

To better understand the distribution and behavior of correlated seismic events, the team cross-correlated a 50-terabyte global dataset consisting of more than 300 million seismograms. Running on a conventional distributed cluster, the dataset took 42 days to complete. In anticipation of processing much larger datasets, the team re-architected the system to run as a series of MapReduce jobs on a custom LC Hadoop cluster, dubbed "Bigfoot." The results were impressive—a 19× performance increase on a 1-terabyte test dataset. Given the results, the team expects to be able to rerun the entire 50-terabyte dataset in 2 days versus the original 42 days.

"We think Hadoop will lead to a viable signal correlation architecture for processing streaming seismic data in the next few years," says Doug Dodge, who co-authored a paper detailing the results on Bigfoot. "This will dramatically improve our ability to conduct research on massive seismic datasets."

Fundamental algorithmic transformations were required to achieve this performance increase. No single data-intensive tool (e.g., MapReduce) will solve every analytic workflow, so the LC and GMP teams are investigating other software techniques and hybrid hardware, such as Livermore's Catalyst system. One prototype is helping to ensure the quality, reliability, and accuracy of seismic data archived by Incorporated Research Institutions for Seismology (IRIS), the world's largest seismic data distributor. In partnership with IRIS, a representative sample of global seismic data (45 terabytes) derived from the IRIS repository is being analyzed with LLNL tools from the distributed computing ecosystem and with techniques from digital signal processing, statistics, and machine learning to facilitate the extraction of quality data by the world-wide seismology community.

These innovative seismic-monitoring technologies should ultimately extend to other data-rich domains also important to Livermore's national security missions, including applied electromagnetics, streaming signals, and the electric grid.

MANAGING APPLICATION PORTABILITY FOR NEXT-GENERATION PLATFORMS

RAJA simplifies parallel software portability

"The challenge is to transform codes once so they can run efficiently on every platform—by exploiting different parallelization strategies and programming models."

Over the past two decades, simulation code performance for the Advanced Simulation and Computing (ASC) Program has improved dramatically as central processing unit (CPU) clock rates increased and computer architectures remained fairly stable. Application developers have focused primarily on scaling to ever-larger node counts as computing platforms grew in size. Now, advanced technology (AT) system node architectures are becoming more complex and diverse as hardware vendors strive to deliver performance gains within constraints such as power usage. New technologies, such as multilevel memory systems and many-core and heterogeneous processor configurations, are challenging application developers to expose massive amounts of fine-grained on-node parallelism, which can take different forms depending on the underlying hardware.

The challenges are particularly acute for ASC multiphysics codes, which are essential tools for Livermore's nuclear stockpile stewardship mission. A typical large integrated physics code contains millions of lines of source code and tens of thousands of loops, in which a wide range of complex numerical operations are performed. Variations in hardware and parallel programming models make it increasingly difficult to achieve high performance without disruptive platform-specific changes to application software. The challenge is to transform codes once so they can run efficiently on every platform by exploiting different parallelization strategies and programming models.

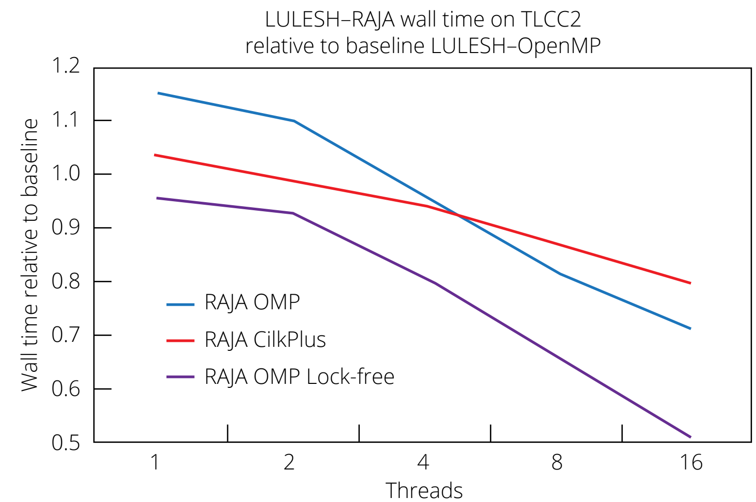

The potential of RAJA is shown in this runtime comparison between different RAJA variants of LULESH and the baseline OpenMP version of LULESH on a single Tri-Lab Linux Capacity Cluster (TLCC2) node. The blue curve shows OpenMP overheads at lower thread counts due to compiler issues when RAJA and OpenMP are combined. However, the purple curve shows that the overheads can be overcome and substantial performance gains achieved by invoking complex execution patterns via RAJA, such as lock-free task scheduling. Note: a single line source code change is required to switch between the RAJA variants.

To address this challenge, computer scientists Richard Hornung and Jeff Keasler are developing RAJA, a programming approach that insulates applications from platform-specific hardware and parallel programming model concerns. RAJA is designed to integrate with existing codes and provide a development model for new codes to be portable from inception. Basic insertion of RAJA enables a code to run on different platforms. Then, architecture-specific tunings can be pursued within the RAJA layer without substantial application code disruption.

The fundamental conceptual abstraction in RAJA is an inner loop, where the overwhelming majority of computational work in most physics codes occurs. The main idea promoted by RAJA is a separation of loop bodies from their iteration patterns. This allows encapsulation of various execution issues including, but not limited to, single instruction multiple data vector parallelism, multithreading and graphics processing unit (GPU) device dispatch, asynchronous task parallelism, fine-grained rollback for resilience, and data permutations to increase locality and/or reduce synchronization dependencies.

"We performed an initial assessment of RAJA to determine whether it can support the programming styles used in Livermore's ASC codes and simplify access to fine-grained parallelism," says Hornung. Hydrodynamics packages in ARES and KULL were used to evaluate basic RAJA usage. LULESH, a proxy for the Lagrange hydrodynamics algorithm in ALE3D, was used to demonstrate RAJA flexibility and more advanced concepts (see figure).

Hornung explains, "These codes use very different software constructs, mesh types, and methodologies for looping over mesh-based data. Parts of our evaluation used the RAJA reference implementation while others required specialization of RAJA concepts for specific code situations. Such customization is a fundamental design goal of RAJA."

Nearly all loops in the ARES Lagrangian hydro package were converted to RAJA. Nominal runs on the Blue Gene/Q platform (Livermore's current AT system, Sequoia) yielded a 50% speedup by introducing four-way OpenMP inner loop threading. KULL performance was studied on a loop-by-loop basis; some loops saw no benefit while others saw close to perfect 4× speedup. Such speedups are due to using hardware threads, which is not done typically in ASC's MPI-only codes since running multiple MPI tasks per core is generally not possible due to memory limitations. Thus, these gains are a pure performance win over the status quo. In contrast, similar loop threading on Livermore's Intel-based machines showed no performance benefit. This was expected due to relatively high performance of sequential execution on standard CPUs compared to fine-grained OpenMP threading. However, the study showed that RAJA allows the option to choose an appropriate execution for each platform and propagate that choice easily through a large source code base at compile time.

This work proved key aspects of RAJA. It is sufficiently flexible to support Livermore production codes; it can enhance readability and maintainability, as opposed to other programming models, such as directive-based, that can "clutter" code; and it can enable easy exploration of different programming models and data layouts (for example, switching between OpenMP and CUDA and reordering loops and data is simply a matter of recompiling a code).

Based on these investigations, several avenues of future work will be pursued, including: continued RAJA development and exploration of portability issues related to future AT platforms, such as Trinity (Intel many-integrated core) and Sierra (GPU); compiler features and optimization issues required for RAJA will be worked with vendors through Trinity and Sierra "centers of excellence" and other co-design activities; and several Livermore production codes are actively moving toward adopting RAJA.

The authors wish to acknowledge members of the ARES and KULL code teams who helped with the RAJA evaluations.

RICHARD HORNUNG
HORNUNG1@LLNL.GOV

JEFF KEASLER
KEASLER1@LLNL.GOV

TRAINING TOMORROW'S CYBERSECURITY SPECIALISTS

Cyber Defenders program emphasizes teamwork and innovation

"We want to give students the depth and hands-on training they might not get in a standard computer science or engineering program."

Lawrence Livermore National Laboratory's (LLNL's) Cyber Defenders program is a unique summer student experience focused on training the next generation of computer security professionals. Cybersecurity is one of the fastest growing technology sectors, the popularity of which is largely founded on an increasing need to combat illegal network hacking and data theft and to protect government and critical infrastructure resources—problems that affect government agencies, corporations, and individual citizens alike.

The Cyber Defenders program is led by computer scientists Celeste Matarazzo and Evi Dube. Several years ago, Matarazzo was attending a cybersecurity conference when she was struck with a realization: the field of cybersecurity needs new thinking—more diversity of thought—a concept that is very familiar to LLNL with its foundation of multidisciplinary science and technology.

"As a nation, we can't be secure without a diverse set of problem solvers to counter the cyber threat," says Matarazzo. No stranger to educational outreach, she realized that a fun yet practical introduction to cybersecurity might encourage more students to pursue careers in the discipline. In 2009, with Laboratory support, she established a Cyber Defenders summer internship program to provide hands-on training to potential future cybersecurity experts.

Six years later, Cyber Defenders is flourishing in its original intention to attract a varied and talented group of participants. The 2014 program included 3 academic faculty members and 27 student interns (12 undergraduates, 6 Master's degree students, 6 Ph.D. candidates, 1 law student, 1 high school teacher, and 1 high school student) from a candidate pool of more than 500 applicants. While Cyber Defenders is a technical internship, it also includes students from social science disciplines.

Cyber Defenders interns develop skills in areas such as intrusion detection and prevention, network monitoring and analysis algorithms, anomaly

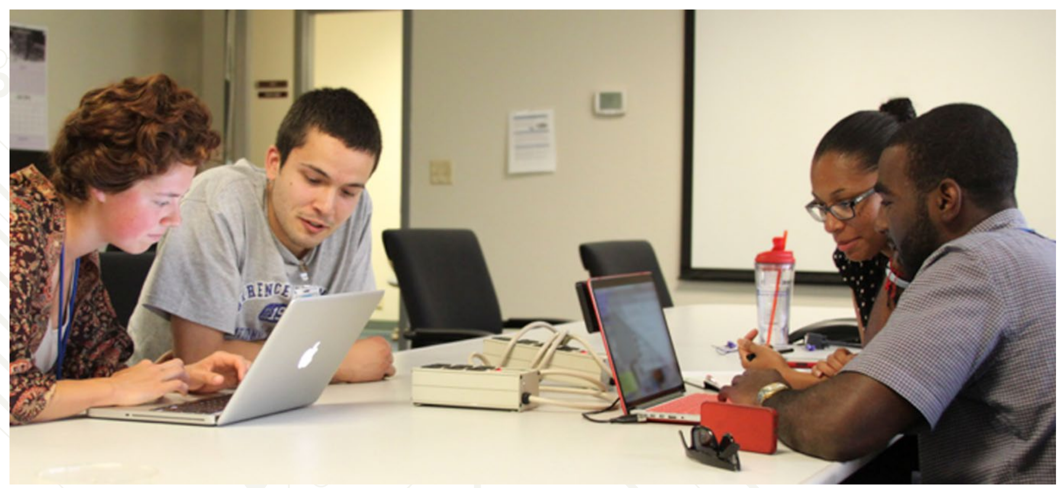

Four cyber defenders team up for a computer penetration testing challenge—one of the program's many competitions designed to encourage collaboration and provide engaging, real-world projects.

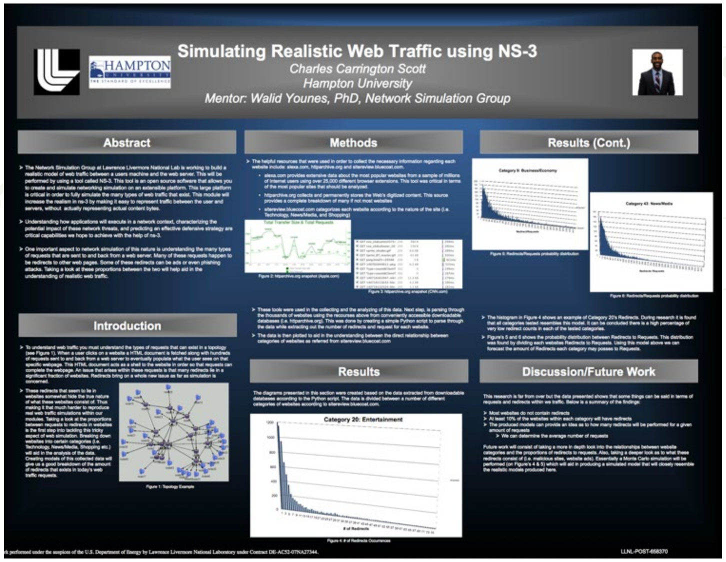

Charles Carrington Scott, a student from Hampton University, presented his work at LLNL's annual student poster symposium. Scott's research focused on NS-3, an open-source network simulator that is driving the Laboratory's realistic web traffic model.

detection, and machine learning. Working on real-world projects is paramount, Matarazzo says of the program's curriculum. "The overall topic of cybersecurity is exciting and is more than compliance and procedures," she says, "so to engage these bright, energetic cyber defenders, the activities need to capture their imagination and certainly not be boring."

Working closely with LLNL mentors, students complete a research project and share their results through a presentation and a poster session. They also attend lectures and seminars, participate in individual and group exercises, and explore new technologies that can be applied to computer security. "Our goals for the program are to excite people about the range of activities available in cybersecurity, to ensure that they leave the program with really strong skills, and to build a pipeline of skilled candidates for jobs at Livermore, at other national laboratories, and in government services," says Matarazzo. "We want to give students the depth and hands-on training they might not get in a standard computer science or engineering program."

A differentiator for the program is the inclusion of the full spectrum of cyber issues. It introduces participants to role-playing and debate exercises and offers a one-day workshop on the legislative, policy, and privacy topics of cyber security, taught by LLNL staff and visiting professors.

One of the more popular activities is a summer-long penetration testing competition in which participants use Raspberry Pis (credit-card size computers) to solve challenge problems in topics such as protocol processing and web security. Another highlight is a one-week cyber defense competition, called Tracer Fire, which is held in collaboration with colleagues from Sandia National Laboratories. Aided by sugar, caffeine, and camaraderie, teams race to complete a series of computer security challenges that test skills such as securing servers from attackers, analyzing malware, and designing robust and secure networks. Just as cybersecurity professionals must continually adapt to remain a step ahead of cyber attackers, the teams must adjust their strategies in response to simulated events and evolving information. What makes the event unique, Matarazzo notes, is its multi-institutional nature: in 2014, interns and technical staff from Sandia's New Mexico campus helped develop the scenarios and led the competition, while interns from LLNL and Sandia California teamed with visitors from Charleston High School to compete.

The Cyber Defenders program will also be key to a new Cybersecurity Workforce Pipeline Consortium. In January 2014, the Department of Energy announced a five-year grant to support an alliance of 13 Historically Black Colleges and Universities, Livermore and Sandia, and a K–12 school district, working together to meet the growing need for trained cybersecurity professionals. LLNL will participate in a blend of research collaborations, mentoring, teaching, curriculum development, and recruitment of students, and Cyber Defenders will provide internship opportunities.

Whether Cyber Defender alumni go on to careers at Livermore or other national labs, as some already have, or they choose other career paths, Matarazzo hopes that they will keep the connections they have made through the internship program. "These connections could prove invaluable as a career resource," she says.

CELESTE MATARAZZO
MATARAZZO1@LLNL.GOV

EVI DUBE
DUBE1@LLNL.GOV

NEW CAPABILITIES FOR INFORMATION TECHNOLOGY SERVICE MANAGEMENT

ServiceNow will offer Livermore staff a cloud-based solution to request services or support

LLNL will be the first NNSA site to use ServiceNow.

Livermore's information technology (IT) teams are responsible for providing IT services and solutions that enable employee productivity and, as a result, further the overall national security mission of Lawrence Livermore National Laboratory (LLNL). The teams are committed to delivering highly available, well-managed IT services that strike a balance between agility, modernity, and compliance.

For an enterprise as large as LLNL, it is essential to have an information technology service management (ITSM) product that structures and automates the flow of work and service delivery, offers excellence in execution, and ensures operational efficiencies and cost effectiveness. Livermore's current ITSM product, which has been in use since July 2009, was due for a significant upgrade. Given the level of effort required to upgrade the product, an evaluation team was tasked with assessing the current product upgrade option and comparing it to other market-leading products, including ServiceNow.

After determining a set of requirements from Livermore customers, end users, and developers, both products were evaluated and rated during vendor presentations, hands-on assessments, and feedback from customers who were already using the products. Based on the reviews, ServiceNow was selected for Livermore to adopt as a pilot product. The pilot objectives included: validating the assessment criteria, ensuring LLNL security requirements could be met with a cloud solution, and giving Livermore's IT team hands-on experience with the tool, on which they could later capitalize during product migration.

Although ServiceNow is currently used by 13 Department of Energy (DOE) sites, including DOE Headquarters, LLNL will be the first National Nuclear Security Administration site to implement it.

ServiceNow has a government cloud architecture that includes full two-site redundancy (two sites within the U.S.) and high availability. In addition, ServiceNow has an active community of users and extensive online documentation. ServiceNow's "Share" site allows developers to exchange custom applications and modules.

As a cloud-based tool, ServiceNow's maintenance tasks, such as server and software patching, are the responsibility of the vendor rather than Livermore's Data Center and ITSM tools team. The time saved in IT staff effort can then be allocated to other tasks, such as designing and developing new services and improving existing tools. Leveraging cloud solutions such as ServiceNow allows Livermore's IT organization to be more responsive and flexible in addressing current IT infrastructure needs.

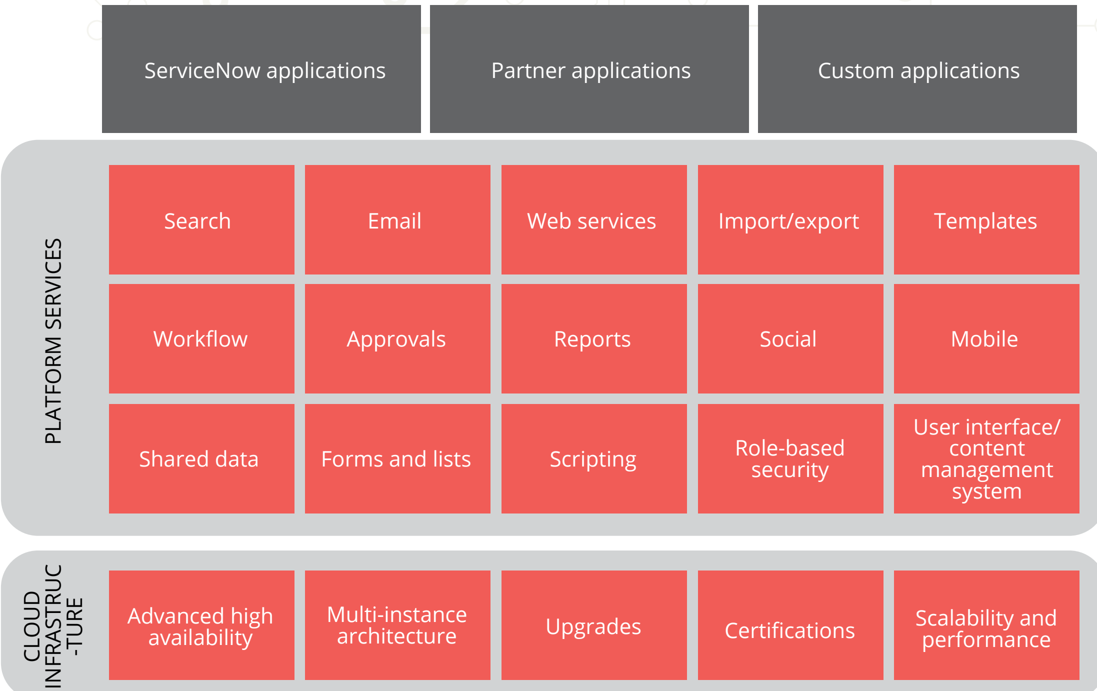

The ServiceNow Software-as-a-Service solution provides a tiered architecture designed with best-in-class cloud infrastructure and common ready-to-use platform services that become the building blocks of applications developed by ServiceNow, its partners, and its customers. With this model, Livermore can take full advantage of the applications that come out of the box, as well as those developed by partners and customers willing to share their applications. Equally important, the model provides the platform to develop custom applications that meet Livermore's unique requirements for service automation.

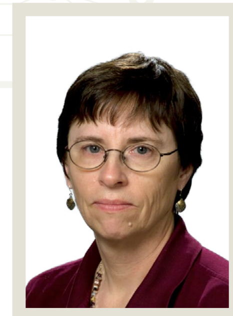

REA SIMPSON
SIMPSON6@LLNL.GOV

The foundation of LLNL's ServiceNow implementation will be a comprehensive service catalog that will enable users to easily select the products and services they need. Once submitted, users can track the status of their requests as they move through the fulfillment steps. With the service catalog providing a single point of presence, users get the benefit of a one-stop shop, leading to an improved customer experience.

The migration to ServiceNow will occur in phases over the next two years. FY15 efforts will include the implementation of the service catalog, migration of desktop support functions from the old product to ServiceNow, and automation of several request processes. The transition to ServiceNow will be completed in FY16.

APPENDICES

Academic Collaborators

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
Allen University	Abdollah Rabieh	Joint research	Malware intelligence harvesting for greater cyber defense	NNSA	Matt Myrick
Arizona State University	Stephen Johnston	Joint research	Peptide microarrays	DHS	Tom Slezak
Boston University	Jonathan Appavoo	Collaboration	Exascale operating systems and runtime	ASCR	Maya Gokhale
California Polytechnic State University, San Luis Obispo	Ignatios Vakalis	Joint research	Cybersecurity research; joint proposals	SMS	Celeste Matarazzo
Carnegie Mellon University	Christos Faloutsos	Joint research	Mining large, time-evolving data for cyber domains; joint proposals	—	Celeste Matarazzo
Carnegie Mellon University	Franz Franchetti	Joint research	Performance optimization of fast Fourier transform on Blue Gene/Q	ASC	Martin Schulz
Carnegie Mellon University	Franz Franchetti	Joint research	Optimization of numerical libraries for Blue Gene/Q	ASC	Martin Schulz
Carnegie Mellon University	Christos Faloutsos	Subcontract	Network summarization and fraud detection	LDRD	Brian Gallagher
Chalmers University of Technology	Sally McKee	Collaboration	Leveraging OpenAnalysis for alias analysis within ROSE	ASC	Dan Quinlan

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
Clark Atlanta University	Roy George	Joint research	Malware intelligence harvesting for greater cyber defense	NNSA	Matt Myrick
Clark Atlanta University	Peter Molnar	Joint research	Cybersecurity research; joint proposals	—	Celeste Matarazzo
Colorado State University	Stephen Guzik	Subcontract	A node-level programming model framework for exascale computing	LDRD	Chunhua Liao
Colorado State University	Michelle Strout and Sanjay Rajopadhye	Collaboration	Program analysis	ASCR	Dan Quinlan
Cornell University	Ken Birman	Joint research	Evaluation of scalable cloud computing technologies for use in Department of Energy systems and applications	ASCR	Greg Bronevetsky
Darmstadt University of Technology	Christian Bischof	Joint research	OpenMP performance tools	ASC	Martin Schulz
Dresden University of Technology	Wolfgang Nagel	Joint research	Improved analysis of message-passing interface traces and performance measurement infrastructures	ASC	Martin Schulz
Dresden University of Technology	Wolfgang Nagel and Andreas Knüpfer	Joint research	Semantic debugging of message-passing interface applications; trace-based performance analysis	ASC	Bronis de Supinski
Duke University	Joshua Vogelstein	Joint research	Brain network similarity	LDRD	Brian Gallagher
ETH Zürich	Alexander Grayver	Collaboration	Solvers for adaptive high-order geo-electromagnetic modeling	ASC/FRIC	Tzanio Kolev
ETH Zürich	Thorsten Hoefler	Joint research	Message-passing interface forum and advanced message-passing interface usage, performance and power modeling	ASC	Martin Schulz

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
Georgetown University	Heidi Wachs	Joint research	Cybersecurity research; joint proposals	—	Celeste Matarazzo
Georgia Institute of Technology	Raheem Beyah	Joint research	Cybersecurity research; joint proposals	—	Celeste Matarazzo
Georgia Institute of Technology	Polo Chau	Joint research	Dynamic networks	LDRD	Brian Gallagher
Georgia Institute of Technology	Richard Fujimoto	Subcontract	Research in reverse computation	LDRD	David Jefferson
Georgia Institute of Technology	Ling Liu	Collaboration	Graph algorithms	Programmatic; CAPS	David Buttler
Georgia Institute of Technology	George Riley	Collaboration	Research in null message synchronization in the context of the ns-3 network simulator	DOD Army Research	David Jefferson
Georgia Institute of Technology	Jarek Rossignac	Joint research	Compact streamable mesh formats	—	Peter Lindstrom
Georgia Institute of Technology	Richard Vuduc	Subcontract	Compiler support for reverse computation	ASCR	Dan Quinlan
Idaho State University	Yunrong Zhu	Collaboration	Multigrid solvers theory	ASC/FRIC	Tzanio Kolev
Imperial College	Paul Kelly and José Gabriel de Figueiredo Coutinho	Collaboration	Field-programmable gate arrays research	ASCR	Dan Quinlan
Indiana University	Jeremiah Willcock	Joint research	Binary analysis	ASCR	Dan Quinlan

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
Indiana University	Andrew Lumsdaine and Udayanga Wikramasinghe	Subcontract	Efficient message-passing interface runtimes; scalable data analysis	ASCR	Greg Bronevetsky
Johns Hopkins University	Allan Boyles	Collaboration	Seismoacoustic modeling for defense-related efforts	DOE	Shawn Larsen
Kansas State University	Bill Hsu	Collaboration	Dynamic time topic models	Programmatic; CAPS	David Buttler
Kyushu University	Koji Inoue	Joint research	Energy and power-aware high performance computing	ASC, ASCR	Martin Schulz
Louisiana State University	Lu Peng, Lide Duan, and Sui Chen	Subcontract	Characterizing the propagation of soft faults through numeric applications	ASCR	Greg Bronevetsky
Ludwig Maximilians University of Munich	Dieter Kranzlmüller	Joint research	Message-passing interface tool infrastructure and performance analysis, power- and energy-aware high performance computing	ASC, ASCR	Martin Schulz
Naval Medical Research Center	Vish Mokashi	Joint research	Microbial forensics	DTRA	Tom Slezak
Norfolk State University	Aftab Ahmad and Jonathan Graham	Joint research	Cybersecurity research; joint proposals	SMS	Celeste Matarazzo
North Carolina Agricultural and Technical State University	Gerry Dozier	Joint research	Cybersecurity research; joint proposals	—	Celeste Matarazzo
North Carolina Agricultural and Technical State University	Gerry Dozier	Joint research	Malware intelligence harvesting for greater cyber defense	NNSA	Matt Myrick
Ohio State University	Umit Catalyurek	Joint research	Task mapping of parallel applications using Chizu	LDRD	Abhinav Bhatele
Ohio State University	P. Sadayappan and Christophe Alias	Collaboration	Optimizing compiler program analysis	ASCR	Dan Quinlan

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
Pennsylvania State University	Jinchao Xu	Collaboration	Multigrid solvers theory	ASC/FRIC	Tzanio Kolev
Pennsylvania State University	Jinchao Xu and James Brannick	Subcontract	Multigrid methods for systems of partial differential equations	ASCR	Robert Falgout
Pennsylvania State University	Ludmil Zikatonov	Subcontract	Multilevel methods and upscaling techniques for elasticity problems	LDRD	Panayot Vassilevski
Polytechnic University of Puerto Rico	Alfredo Cruz	Joint research	Cybersecurity research; joint proposals	—	Celeste Matarazzo
Portland State University	Jay Gopalakrishnan	Collaboration	Discontinuous Petrov–Galerkin methods	ASC/FRIC	Tzanio Kolev
Purdue University	Saurabh Bagchi	Joint research	Anomaly detection and tracking in high performance computing	ASC	Martin Schulz
Purdue University	Saurabh Bagchi	Subcontract	Root cause analysis of faults in parallel systems	ASCR	Greg Bronevetsky
Purdue University	Zhiqiang Cai	Summer faculty	A posteriori error estimates for partial differential equations	ASC	Robert Falgout
Purdue University	Jennifer Neville	Joint research	Hypothesis tests for dynamic networks	LDRD	Brian Gallagher
Purdue University	Mithuna Thottethodi	Joint research	Optimized node mapping techniques	ASC	Martin Schulz
Purdue University	Mithuna Thottethodi	Joint research	Routing-aware task mapping	LDRD	Abhinav Bhatele

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
Queens University of Belfast	Dimitrios Nikolopoulos	Joint research	Power optimization for hybrid codes, epidemiology simulation	ASC	Martin Schulz
Rensselaer Polytechnic Institute	Mark Shephard and Onkar Sahni	Joint research	FASTMath: Frameworks, Algorithms, and Scalable Technologies for Mathematics	ASCR SciDAC	Lori Diachin
Rice University	John Mellor-Crummey	Joint research	Performance analysis, standardization for OpenMP	ASC, ASCR	Martin Schulz
Rice University	John Mellor-Crummey	Joint research	Sustained Performance, Energy, and Resilience Institute	ASCR SCiDAC	Bronis de Supinski
Rice University	John Mellor-Crummey, Keith Cooper, and Vivek Sarkar	Collaboration	Use of ROSE for compiler optimizations	ASCR	Dan Quinlan
Rice University	Vivek Sarkar, Jisheng Zhao, Vincent Cave, and Michael Burke	Joint research	Development of a static single assignment-based dataflow compiler framework for ROSE	ASCR	Greg Bronevetsky
Royal Institute of Technology, Sweden	Heinz-Otto Kreiss	Consultant	Adaptive methods for partial differential equations	ASCR Base	Anders Petersson
Rutgers University	Tina Eliassi-Rad	Subcontract	Cyber situational awareness through host and network analysis	LDRD	Celeste Matarazzo
Rutgers University	Tina Eliassi-Rad	Subcontract	Dynamic networks/network inference	LDRD	Brian Gallagher
Rutgers University	Manish Parashar	Joint research	ExACT Co-Design Center	DOE	Timo Bremer
RWTH Aachen University	Matthias Müller	Joint research	Message-passing interface correctness checking	ASC	Martin Schulz

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
RWTH Aachen University	Felix Wolf	Joint research	I/O and network interference on torus networks	ASC	Abhinav Bhatele
RWTH Aachen University	Felix Wolf	Joint research	Parallel performance analysis	ASC, ASCR	Martin Schulz
RWTH Aachen University	Felix Wolf and Matthias Müller	Joint research	OpenMP, message-passing interface performance analysis tools	ASC	Bronis de Supinski
Southern Methodist University	Dan Reynolds	Joint research/ subcontract	FASTMath: Frameworks, Algorithms, and Scalable Technologies for Mathematics	ASCR SciDAC	Lori Diachin
Southern Methodist University	Dan Reynolds	Subcontract	New time integration methods and support for multiscale solution methods in the LLNL SUNDIALS software suite	ASCR SciDAC	Carol Woodward
Southern Methodist University	Dan Reynolds	Subcontract	Time integration and solvers for materials simulations	NNSA	Carol Woodward
Stanford University	Juan Alonso and Gianluca Iaccarino	Subcontract	New modes of laser lethality	LDRD	Kambiz Salari
Stanford University	Sanjiva Lele	Subcontract	Development of a nonequilibrium wall model for the compressible flow solver CharLESX	DOE	Kambiz Salari
Stanford University	Olav Lindtjorn	Collaboration	Reverse-time seismic imaging for hydrocarbon exploration	CRADA	Shawn Larsen
Stanford University	Parvis Moin	ASC Predictive Science Academic Alliance Program Center	Center for Predictive Simulations of Multiphysics Flow Phenomena with Application to Integrated Hypersonic Systems	ASC	Dick Watson
Technical University of Denmark	Sven Karlsson	Joint research	Scalable debugging	ASC	Martin Schulz and Dong Anh

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
Technical University of Dortmund	Dmitri Kuzmin	Collaboration	Monotonicity for high-order remap	LDRD	Tzanio Kolev
Technical University of Munich	Arndt Bode	Joint research	Exascale computing	ASC	Martin Schulz
Technical University of Vienna	Markus Schordan	Collaboration	Compiler construction	ASCR	Dan Quinlan
Texas A&M University	Nancy Amato	Joint research	Load balance optimizations	ASC	Martin Schulz
Texas A&M University	Nancy Amato	Collaboration, Lawrence Scholar Program	Novel mechanisms to understand and improve load balance in message-passing interface applications	ASCR, ASC	Bronis de Supinski
Texas A&M University	Nancy Amato	Collaboration, Lawrence Scholar Program	Parallel graph algorithms	UCOP	Maya Gokhale
Texas A&M University	Yalchin Efendiev	Joint research	Scalable uncertainty quantification concepts for multiscale problems at extreme scale	ASCR	Panayot Vassilevski
Texas A&M University	Jean-Luc Guermond and Boian Popov	Collaboration	Finite elements for shock hydrodynamics	LDRD	Tzanio Kolev
Texas A&M University	Bjarne Stroustrup and Lawrence Rauchwerger	Joint research	Compiler construction and parallel optimizations	ASCR	Dan Quinlan
Tufts University	Scott MacLachlan	Joint research	Parallel multigrid in time	ASCR	Robert Falgout
UC Berkeley	Doug Dreger	Collaboration	Earthquake hazard	IGPP	Shawn Larsen

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
UC Davis	Matt Bishop and Sean Peisert	Joint research	Cybersecurity research, joint proposals, cyber defenders	—	Celeste Matarazzo
UC Davis	Soheil Ghiasi	Collaboration	Machine learning	Joint UC Davis/LLNL proposal	Maya Gokhale
UC Davis	Dipak Ghosal	Joint research	Interjob interference and performance variability on dragonfly networks	LDRD	Abhinav Bhatele
UC Davis	François Gygi	Collaboration	General Qbox development, new algorithms	—	Erik Draeger
UC Davis	Bernd Hamann	Joint research	Analysis and visualization of performance data	UC Fee	Timo Bremer
UC Davis	Bernd Hamann	Joint research	Performance analysis and visualization	ASC	Martin Schulz
UC Davis	Sean Peisert	Subcontract	Network resilience	LDRD	Celeste Matarazzo
UC Davis	Zhendong Su	Subcontract	ROSE support for rewrapping macro calls	ASCR	Dan Quinlan
UC Riverside	Michalis Faloutsos	Joint research	Cybersecurity research; joint proposals	—	Celeste Matarazzo
UC San Diego	Randy Bank	Subcontract	Solvers for large sparse systems of linear equations	ASCR	Robert Falgout
UC San Diego	Laura Carrington	Collaboration	Data-intensive architectures	UC Fee	Maya Gokhale

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
UC San Diego	Laura Carrington and Scott Baden	Joint research	Performance modeling	ASC	Martin Schulz
UC San Diego	Erik Gartzke and Jon Lindsey	Joint research	Cybersecurity research; joint proposals	—	Celeste Matarazzo
UC San Diego	Steve Swanson	Collaboration	Persistent memory emulator	—	Maya Gokhale
UC San Diego, Scripps Institution of Oceanography	Julie McClean	Collaboration	Ultra-high-resolution coupled climate simulations	BER	Art Mirin
UC Santa Cruz	Steve Kang	Collaboration	Persistent memory devices	UC Fee	Maya Gokhale
UC Santa Cruz	Carlos Maltzahn	Collaboration, Lawrence Scholar Program	Semantic file systems	LDRD	Maya Gokhale
United States Army Medical Research Unit, Kenya	John Waitumbi	Joint research	Pathogen diagnostics	—	Tom Slezak
University of Arizona	David Lowenthal	Joint research	Power-aware computing for message-passing interface programs; scalable performance models	ASCR, ASC	Bronis de Supinski
University of Arizona	David Lowenthal	Joint research	Power optimization and modeling	ASC, ASCR	Martin Schulz
University of British Columbia	Carl Olivier-Gooch	Subcontract	Aggressive mesh optimization	ASCR SciDAC	Lori Diachin
University of British Columbia	Carl Olivier-Gooch	Subcontract	FASTMath: Frameworks, Algorithms, and Scalable Technologies for Mathematics/Mesquite	ASCR SciDAC/ASCR Base	Lori Diachin

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
University of Cologne	Martin Lanser	Collaboration	Investigation of interpolation in algebraic multigrid for elasticity	DOE ASCR	Ulrike Yang
University of Colorado	Ken Jansen	Joint research	FASTMath: Frameworks, Algorithms, and Scalable Technologies for Mathematics	ASCR SciDAC	Lori Diachin
University of Colorado	Tom Manteuffel	Joint research	Solution methods for transport problems	ASC	Peter Brown
University of Colorado	Steve McCormick and Tom Manteuffel	Subcontract	Adaptive algebraic multigrid for graph mining problems	LDRD	Van Emden Henson
University of Colorado	Steve McCormick, Tom Manteuffel, John Ruge, and Marian Brezina	Subcontract	Error estimators for uncertainty quantification, adaptive mesh refinement, solvers for Stochastic partial differential equations, parallel adaptive algebraic multigrid/smoothed aggregation, and parallel solution of systems of partial differential equations	ASC	Robert Falgout
University of Delaware	John Cavazos	Subcontract	ROSE compiler project	ASCR	Dan Quinlan and Chunhua Liao
University of Delaware	Michela Taufer	Joint research	Massively scalable I/O-aware job scheduling	ASC	Dong Ahn
University of Illinois at Chicago	Tanya Berger-Wolf	Joint research	Network inference	LDRD	Brian Gallagher
University of Illinois at Urbana-Champaign	William Gropp	Joint research	Message-passing interface, hybrid programming models	ASCR, ASC	Bronis de Supinski
University of Illinois at Urbana-Champaign	William Gropp	Joint research	Optimization for algebraic multigrid	ASC	Martin Schulz
University of Illinois at Urbana-Champaign	Laxmikant Kale	Subcontract	Scalable topology-aware task embedding (STATE)	LDRD	Abhinav Bhatele

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
University of Illinois and IBM	William Gropp and Kirk Jordan	Collaboration	Modeling algebraic multigrid performance on multicore architectures	ASCR	Ulrike Yang
University of Karlsruhe	Wolfgang Karl	Joint research	Hardware transactional memory	ASC	Martin Schulz
University of Maryland	Jeffrey Hollingsworth	Joint research	Autotuning and tool infrastructures	ASCR	Martin Schulz
University of Maryland	Jeffrey Hollingsworth	Joint research	Sustained Performance, Energy, and Resilience Institute	ASCR SciDAC	Bronis de Supinski
University of Nevada, Reno	John Louie	Collaboration	Seismic modeling in the basin and range region	DOE	Shawn Larsen
University of New Mexico	Dorian Arnold	Joint research	Tool infrastructures	ASC	Greg Lee
University of North Carolina	Robert Fowler	Joint research	Sustained Performance, Energy, and Resilience Institute	ASCR SciDAC	Bronis de Supinski
University of North Carolina	Jan Prins	Joint research	OpenMP task scheduling	ASC	Martin Schulz
University of San Francisco	Jeff Buckwalter	Joint research	Performance modeling	ASC	Martin Schulz
University of Southern California	Robert Lucas and Jacqueline Chame	Joint research	Sustained Performance, Energy, and Resilience Institute	ASCR SciDAC	Bronis de Supinski
University of Tennessee	Jack Dongarra	Joint research	Empirical tuning	ASCR	Dan Quinlan

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
University of Tennessee	Jack Dongarra	Joint research	Sustained Performance, Energy, and Resilience Institute	ASCR SciDAC	Bronis de Supinski
University of Texas, Medical Branch	Yuriy Fofanov	Joint research	Genomic algorithms	DTRA	Tom Slezak
University of Texas, San Antonio	Shirley Moore	Joint research	Sustained Performance, Energy, and Resilience Institute	ASCR SciDAC	Bronis de Supinski
University of Turabo	Jeffrey Duffany	Joint research	Cybersecurity research: analysis and defense of large-scale smart meter networks; joint proposals	SMS	Celeste Matarazzo
University of Utah	Ganesh Gopalakrishnan	Subcontract	Compiler analysis of message-passing interface applications	ASCR	Greg Bronevetsky
University of Utah	Ganesh Gopalakrishnan	Collaboration	Message-passing interface optimizations	ASCR	Dan Quinlan
University of Utah	Ganesh Gopalakrishnan	Joint research	Identification and targeted elimination of non-determinism	ASC	Dong Ahn
University of Utah	Ganesh Gopalakrishnan and Mary Hall	Joint research	Semantic debugging of message-passing interface applications, and Sustained Performance, Energy, and Resilience Institute	ASCR, ASC	Bronis de Supinski
University of Utah	Mike Kirby	Joint research	Analysis of vector fields	DOE	Timo Bremer
University of Utah	Valerio Pascucci	Joint research	Performance analysis and visualization	ASC, ASCR	Martin Schulz
University of Utah	Valerio Pascucci	Subcontract	Performance analysis and visualization	DOE	Timo Bremer

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
University of Waterloo	Hans de Sterck	Subcontract	Numerical methods for large-scale data factorization	LDRD	Van Emden Henson
University of Wisconsin	Karu Sankaralingam	Joint research	Fault tolerant computing models for high performance computing	ASC	Martin Schulz
University of Wisconsin	Bart Miller and Ben Liblit	Joint research	Performance tools and tool infrastructures	ASCR, ASC	Martin Schulz
Utah State University	Renée Bryce and Steena Monteiro	Joint research	Statistical modeling of data-driven applications	ASCR, Lawrence Scholar Program	Greg Bronevetsky
Virginia Institute of Technology	Kirk Cameron	Joint research	Power-aware computing for hybrid systems	ASCR, ASC	Bronis de Supinski
Virginia Institute of Technology	Wu-chun Feng	Joint research	Hybrid computing programming models, power-aware computing	ASCR, ASC	Bronis de Supinski
Virginia Institute of Technology	Madhav Marathe	Joint research	Epidemiology simulation at scale	ASC	Martin Schulz
Virginia Institute of Technology	Madhav Marathe	Joint research	Mathematical and computational foundations of network sciences	SMS	Celeste Matarazzo
Voorhees College	Tim Kentopp	Joint research	Malware intelligence harvesting for greater cyber defense	NNSA	Matt Myrick

PUBLICATIONS

Journal Articles and Conference Proceedings

Abdel-Gawad, A., M. Thottethodi, and A. Bhatele, "RAHTM: Routing-Algorithm Aware Hierarchical Task Mapping," *International Conference for High Performance Computing, Networking, Storage and Analysis*, (IEEE Press), (LLNL-CONF-653568).

Addair, T. G., D. A. Dodge W. R. Walter, and S. D. Ruppert, "Large-Scale Seismic Signal Analysis with Hadoop," *Comput. Geosci.-UK* **66**, 145–154.

Adler, J. H., P. S. Vassilevski, "Error Analysis for Constrained First-Order System Least-Squares Finite-Element Methods," *SIAM J. Sci. Comput.* **36** (3), A1071–A1088.

Ahmed, M. F., et al., "Target Material Collection for High Energy Imaging Diagnostic," *SPIE Optics + Photonics 2014*, **9211**, 92110F (LLNL-CONF-657947).

Almendro, V., et al., "Inference of Tumor Evolution during Chemotherapy by Computational Modeling and In Situ Analysis of Genetic and Phenotypic Cellular Diversity," *Cell Rep.* **6** (3), 514–527.

Amendt, P., et al., "Low-Adiabat Rugby Hohlraum Experiments on the National Ignition Facility: Comparison with High-Flux Modeling and the Potential for Gas-Wall Interpenetration," *Phys. Plasmas* **21** (11).

Ayuso de Dios, B., A. T. Barker, and P. S. Vassilevski, "A Combined Preconditioning Strategy for Nonsymmetric Systems," *SIAM J. Sci. Comput.* **36**, A2533–A2556.

Bailey, P. E., et al., "Adaptive Configuration Selection for Power-Constrained Heterogeneous Systems," *43rd International Conference on Parallel Processing (ICPP-14)*, (LLNL-CONF-662222).

Barker, A. T., "A Minimal Communication Approach to Parallel Time Integration," Int. *J. Comput. Math.* **91**, 601–615.

Barker, A. T. and S. C. Brenner, "A Mixed Finite Element Method for the Stokes Equations Based on a Weakly Over-Penalized Symmetric Interior Penalty Approach," *J. Sci. Comput.* **58**, 290–307.

Be, N. A., et al., "Microbial Profiling of Combat Wound Infection through Detection Microarray and Next-Generation Sequencing," *J. Clin. Microbiol.* **52** (7), 2583–2594.

Bergmann, J. C., et al., "Discovery of Two Novel Beta-Glucosidases from an Amazon Soil Metagenomic Library," *FEMS Microbiol. Lett.* **351** (2), 147–155.

Bhatele, A., et al., "Optimizing the Performance of Parallel Applications on a 5D Torus via Task Mapping," *IEEE International Conference on High Performance Computing*, (IEEE Computer Society), (LLNL-CONF-655465).

Bhatia, H., V. Pascucci, and P. T. Bremer, "The Natural Helmholtz–Hodge Decomposition for Open-Boundary Flow Analysis," *IEEE T. Vis. Comput. Gr.* **20** (11), 1566–1578.

Bhatia, H., V. Pascucci, R. M. Kirby, and P. T. Bremer, "Extracting Features from Time-Dependent Vector Fields Using Internal Reference Frames," *Comput. Graph. Forum* **33** (3), 21–30.

Breslow, A. D., et al., "Enabling Fair Pricing on High Performance Computer Systems with Node Sharing," *Sci. Program.* **22** (2), 59–74.

Carver, J. C. and T. Epperly, "Software Engineering for Computational Science and Engineering: Introduction," *Comput. Sci. Eng.* **16** (3), 6–9.

Chen, M. J., et al., "An Efficient Bayesian Inversion of a Geothermal Prospect Using a Multivariate Adaptive Regression Spline Method," *Appl. Energy* **136**, 619–627.

Cinquini, L., et al., "The Earth System Grid Federation: An Open Infrastructure for Access to Distributed Geospatial Data," *Future Gener. Comp. Sys.* **36**, 400–417.

Clark, D. S., et al., "Survey of Pulse Shape Options for a Revised Plastic Ablator Ignition Design," *Phys. Plasmas* **21** (11).

Connors, J. M., J. W. Banks, J. A. Hittinger, and C. S. Woodward, "Quantification of Errors for Operator-Split Advection-Diffusion Calculations," *Comput. Method. Appl. M.* **272**, 181–197 (LLNL-JRNL-566435).

Crawford, D. L., "Life at the Leading Edge: Supercomputing @ Lawrence Livermore," *International Conference for High Performance Computing, Networking, Storage and Analysis*, (IEEE Press), (LLNL-PRES-664317).

Crawford, D. L., "Supercomputing," *IEA Combustion 2014*, (LLNL-PRES-655329).

Devault, A. M., et al., "Ancient Pathogen DNA in Archaeological Samples Detected with a Microbial Detection Array," *Sci. Rep.* **4**.

Dittrich, T. R., et al., "Design of a High-Foot High-Adiabat ICF Capsule for the National Ignition Facility," *Phys. Rev. Lett.* **112** (5).

Dobrev, V. A., T. V. Kolev, and R. N. Rieben, "High Order Curvilinear Finite Elements for Elastic–Plastic Lagrangian Dynamics," *J. Comput. Phys.* **257**, 1062–1080.

Dorf, M. A., et al., "Progress with the COGENT Edge Kinetic Code: Implementing the Fokker–Planck Collision Operator," *Contrib. Plasm. Phys.* **54** (4–6), 517–523.

Falgout, R. D. and J. B. Schroder, "Non-Galerkin Coarse Grids for Algebraic Multigrid," *SIAM J. Sci. Comput.* **36** (3), C309–C334 (LLNL-JRNL-645325).

Fattebert, J. L., M. E. Wickett, and P. E. A. Turchi, "Phase-Field Modeling of Coring During Solidification of Au-Ni Alloy Using Quaternions and CALPHAD Input," *Acta Mater.* **62**, 89–104 (LLNL-JRNL-636233).

Gan, Y. J., et al., "A Comprehensive Evaluation of Various Sensitivity Analysis Methods: A Case Study with a Hydrological Model," *Environ. Modeling Softw.* **51**, 269–285.

Gardner, S. N. and T. Slezak, "Simulate_PCR for Amplicon Prediction and Annotation from Multiplex, Degenerate Primers and Probes," *BMC Bioinformatics* **15**.

Geng, J., I. M. Navon, and X. Chen, "Nonparametric Calibration of the Local Volatility Surface for European Options Using a Second-Order Tikhonov Regularization," *Quant. Financ.* **14** (1), 73–85.

Giménez, A., et al., "Dissecting On-Node Memory Access Performance: A Semantic Approach," *International Conference for High Performance Computing, Networking, Storage and Analysis*, (IEEE Press), (LLNL-CONF-658626).

Gladden, J. M., et al., "Discovery and Characterization of Ionic Liquid-Tolerant Thermophilic Cellulases from a Switchgrass-Adapted Microbial Community," *Biotechnol. Biofuels* **7**.

Greenwald, M., et al., "20 Years of Research on the Alcator C-Mod Tokamak," *Phys. Plasmas* **21** (11).

Gyulassy, A., et al., "Stability of Dissipation Elements: A Case Study in Combustion," *Comput. Graph. Forum* **33** (3), 51–60.

Harvey, W., et al., "A Collaborative Visual Analytics Suite for Protein Folding Research," *J. Mol. Graph. Model.* **53**, 59–71.

Hilbrich, T., et al., "Memory Usage Optimizations for Online Event Analysis," *Exascale Applications and Software Conference (EACS 2014)*, (LLNL-CONF-652119).

Hurricane, O. A., "Fuel Gain Exceeding Unity in an Inertially Confined Fusion Implosion," *Nature* **506** (7488), 343.

Hurricane, O. A., et al., “The High-Foot Implosion Campaign on the National Ignition Facility,” *Phys. Plasmas* **21** (5).

Iftekharuddin, K. M., A. A. S. Awwal, S. S. Young, and G. M. Abdulla, “Optical and Hybrid Imaging and Processing for Big Data Problems,” *Opt. Eng.* **53** (1).

Isaacs, K. E., et al., “Combing the Communication Hairball: Visualizing Parallel Execution Traces Using Logical Time,” *IEEE T. Vis. Comput. Gr.* **20** (12), 2349–2358 (LLNL-JRNL-657418).

Isaacs, K. E., et al., “State of the Art of Performance Visualization,” *Eurographics Conference of Visualization*, (LLNL-CONF-652873).

Islam, T., S. Bagchi, and R. Eigenmann, “Reliable and Efficient Distributed Checkpointing System for Grid Environments,” *Journal Grid Computing* **12** (4), 593–613 (LLNL-JRNL-649440).

Jain, N., et al., “Maximizing Throughput on a Dragonfly Network,” *International Conference for High Performance Computing, Networking, Storage and Analysis*, (IEEE Press), (LLNL-CONF-653557).

Jiang, M., et al., “Multi-Threaded Streamline Tracing for Data-Intensive Architectures,” *IEEE Symposium on Large Data Analysis and Visualization*, (LLNL-CONF-645076).

Kamath, C., et al., “Density of Additively Manufactured, 316L SS Parts Using Laser Powder-Bed Fusion at Powers Up to 400 W,” *Int. J. Adv. Manuf. Tech.* **74** (1–4), 65–78.

Kegelmeyer, L. M., et al., “Optimizing Blocker Usage on NIF Using Image Analysis and Machine Learning,” *ICALEPCS2013: Proceedings of the 14th International Conference on Accelerator and Large Experimental Physics Control Systems*, (LLNL-CONF-644265).

Kim, H., C. Correa, and N. Max, “Automatic Registration of LiDAR and Optical Imagery Using Depth Map Stereo,” *IEEE ICCP*, 205–212 (LLNL-PROC-652419).

Kim, H., J. J. Thiagarajan, and P. T. Bremer, “Image Segmentation Using Consensus from Hierarchical Segmentation Ensembles,” *IEEE ICIP*, (LLNL-PROC-652239).

King, W.E., et al., “Observation of Keyhole-Mode Laser Melting in Laser Powder-Bed Fusion Additive Manufacturing,” *J. Mater. Process. Tech.* **214** (12), 2915–2925.

Laguna, I., et al., “Evaluating User-Level Fault Tolerance for MPI Applications,” *EuroMPI 2014*, (LLNL-CONF-656877).

Laney, D., et al., “Assessing the Effects of Data Compression in Simulations Using Physically Motivated Metrics,” *Sci. Program.* **22** (2), 141–155 (LLNL-CONF-642236).

Langer, S. H., A. Bhatele, and C. H. Still, “pF3D Simulations of Laser-Plasma Interactions in National Ignition Facility Experiments,” *Comput. Sci. Eng.* **16** (6), 42–50 (LLNL-JRNL-648736).

Levesque, G. A., et al., “Analyses on the Effect of Hot Spot Density on Material Consumption Rate,” *18th Joint International Conference of the APS Topical-Grp on Shock Compress of Condensed Matter/24th International Conference of the International Association for the Advancement of High Pressure Science and Technology*, (IOP Publishing Ltd), **500**, 152010.

Liao, Z. M., et al., “Damage Modeling and Statistical Analysis of Optics Damage Performance in MJ-Class Laser Systems,” *Opt. Express* **22** (23), 28845–28856.

Lindstrom, P., “Fixed-Rate Compressed Floating-Point Arrays,” *IEEE T. Vis. Comput. Gr.* **20** (12), 2674–2683 (LLNL-JRNL-652576).

Lindstrom, P. and D. Rajan, “Optimal Hierarchical Layouts for Cache-Oblivious Search Trees,” *IEEE International Conference on Data Engineering 2014*, (IEEE), 616–627 (LLNL-CONF-641294).

Liu, S., B. Wang, P. T. Bremer, and V. Pascucci, “Distortion-Guided Structure-Driven Interactive Exploration of High-Dimensional Data,” *Comput. Graph. Forum* **33** (3), 101–110.

Luffel, M., T. Gurung, P. Lindstrom, and J. Rossignac, “Grouper: A Compact, Streamable Triangle Mesh Data Structure,” *IEEE T. Vis. Comput. Gr.* **20** (1), 84–98 (LLNL-JRNL-503653).

MacKinnon, A. J., et al., “High-Density Carbon Ablator Experiments on the National Ignition Facility,” *Phys. Plasmas* **21** (5), (LLNL-JRNL-647814).

Marathe, A., et al., “Exploiting Redundancy for Cost-Effective, Time-Constrained Execution of HPC Applications on Amazon EC2,” 23rd IEEE International Symposium High Performance Distributed Computing (HPDC 2014), (LLNL-CONF-652718).

Execution of HPC Applications on Amazon EC2," *23rd IEEE International Symposium High Performance Distributed Computing (HPDC 2014)*, (LLNL-CONF-652718).

Masters, N., et al., "Evaluation of Observed Blast Loading Effects on NIF X-Ray Diagnostic Collimators," *20th Topical Conference on High-Temperature Plasma Diagnostics*, (Rev. Sci. Instrum.), **85** (11), 11D628 (LLNL-CONF-655173).

Maurer, E. P., et al., "An Enhanced Archive Facilitating Climate Impacts and Adaptation Analysis," *B. Am. Meteorol. Soc.* **95** (7), 1011.

McCandless, K. P., et al., "The Role of Data Driven Models in Optimizing the Operation of the National Ignition Facility," *ICALEPCS*, (JACoW), 1426 (LLNL-CONF-644538).

McCarthy, C. M., et al., "Visualizing the Five-Dimensional Torus Network of the IBM Blue Gene/Q," *1st Workshop on Visual Performance Analysis*, (LLNL-CONF-661000).

Mohan, A., T. Hacker, G. P. Rodgers, and T. Islam, "Batchsubmit: A High-Volume Batch Submission System for Earthquake Engineering Simulation," *Concurr. Comp.-Pract. E.* **26** (13), 2240–2252.

Mohror, K., A. Moody, G. Bronevetsky, and B. R. de Supinski, "Detailed Modeling and Evaluation of a Scalable Multilevel Checkpointing System," *IEEE T. Parall. Distrib. Sys.* **25** (9), 2255–2263 (LLNL-JRNL-564721).

Moody, J. D., et al., "Progress in Hohlraum Physics for the National Ignition Facility," *Phys. Plasmas* **21** (5).

Osei-Kuffuor, D. and J. L. Fattebert, "Accurate and Scalable O(N) Algorithm for First-Principles Molecular-Dynamics Computations on Large Parallel Computers," *Phys. Rev. Lett.* **112** (4) (LLNL-JRNL-641855).

Osei-Kuffuor, D. and J. L. Fattebert, "A Scalable O(N) Algorithm for Large-Scale Parallel First-Principles Molecular Dynamics Simulations," *SIAM J. Sci. Comput.* **36** (4), C353–C375 (LLNL-JRNL-649416).

Osei-Kuffuor, D., R. M. Maxwell, and C. S. Woodward, "Improved Numerical Solvers for Implicit Coupling of Subsurface and Overland Flow," *Adv. Water Res.* **74**, 185–195 (LLNL-JRNL-654773).

Ott, T., M. Bonitz, L. G. Stanton, and M. S. Murillo, "Coupling Strength in Coulomb and Yukawa One-Component Plasmas," *Phys. Plasmas* **21** (11).

Paradzik, M., et al., "Association of Kaposi's Sarcoma-Associated Herpesvirus (KSHV) with Bladder Cancer in Croatian Patients," *Tumor Biol.* **35** (1), 567–572.

Pavlu, V., M. Schordan, and A. Krall, "Computation of Alias Sets from Shape Graphs for Comparison of Shape Analysis Precision," *IET Softw.* **8** (3), 120–133.

Pearce, O., et al., "Load Balancing N-Body Simulations with Highly Non-Uniform Density," *ACM International Conference on Supercomputing (ICS)*, (ACM New York, NY, USA), 113–122 (LLNL-CONF-648577).

Pearce, R., M. Gokhale, and N. Amato, "Faster Parallel Traversal of Scale Free Graphs at Extreme Scale with Vertex Delegates," *International Conference for High Performance Computing, Networking, Storage, and Analysis*, (IEEE Press), (LLNL-CONF-658291).

Petersson, N. A. and B. Sjogreen, "Super-Grid Modeling of the Elastic Wave Equation in Semi-Bounded Domains," *Commun. Comput. Phys.* **16** (4), 913–955.

Pickering, B. P., et al., "Directive-Based GPU Programming for Computational Fluid Dynamics," *52nd Aerospace Sciences Meeting* (American Institute of Aeronautics and Astronautics).

Rosenstierne, M. W., et al., "The Microbial Detection Array for Detection of Emerging Viruses in Clinical Samples: A Useful Panmicrobial Diagnostic Tool," *PLoS ONE* **9** (6).

Santer, B. D., et al., "Volcanic Contribution to Decadal Changes in Tropospheric Temperature," *Nat. Geosci.* **7** (3), 185–189.

Sato, K., et al., "A User-Level InfiniBand-Based File System and Checkpoint Strategy for Burst Buffers," *IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid2014)*, (IEEE), 21–30 (LLNL-CONF-645876).

Sato, K., et al., "FMI: Fault Tolerant Messaging Interface for Fast and Transparent Recovery," *International Conference on Parallel and Distributed Processing Symposium 2014 (IPDPS2014)*, (IEEE Computer Society), 1225–1234 (LLNL-CONF-645209).

Schleife, A., et al., "Quantum Dynamics Simulation of Electrons in Materials on High-Performance Computers," *Comput. Sci. Eng.* **16** (5), 54–60.

Schone, R., et al., "Tools and Methods for Measuring and Tuning the Energy Efficiency of HPC Systems," *Sci. Program.* **22** (4), 273–283.

Scogland, T. R. W., W. Feng, B. Rountree, and B. R. de Supinski, "CoreTSAR: Core Task-Size Adapting Runtime," *IEEE T. Parall. Distr.*, **99** (1).

Scogland, T. R. W., B. Rountree, W. Feng, and B. R. de Supinski, "CoreTSAR: Adaptive Worksharing for Heterogeneous Systems," *The International Supercomputing Conference (ISC'14)*, (LLNL-CONF-662205).

Scogland, T. R., et al., "A Power-Measurement Methodology for Large-Scale, High-Performance Computing," *ACM/SPEC International Conference on Performance Engineering (ICPE)*.

Simakov, A. N., et al., "Optimized Beryllium Target Design for Indirectly Driven Inertial Confinement Fusion Experiments on the National Ignition Facility," *Phys. Plasmas* **21** (2).

Sjogreen, B. and N. A. Petersson, "Source Estimation by Full Wave Form Inversion," *J. Sci. Comput.* **59** (1), 247–276.

Sjogreen, B., H. C. Yee, and M. Vinokur, "On High Order Finite-Difference Metric Discretizations Satisfying GCL on Moving and Deforming Grids," *J. Comput. Phys.* **265**, 211–220.

Smalyuk, V. A., et al., "Measurements of an Ablator-Gas Atomic Mix in Indirectly Driven Implosions at the National Ignition Facility," *Phys. Rev. Lett.* **112** (2).

Smith, S., et al., "Online Transient Stability Analysis Using High Performance Computing," *2014 IEEE PES Innovative Smart Grid Technologies Conference (ISGT)*, (IEEE, Piscataway, NJ, USA), 1–5.

Thissen, J. B., et al., "Analysis of Sensitivity and Rapid Hybridization of a Multiplexed Microbial Detection Microarray," *J. Virol. Meth.* **201**, 73–78 (LLNL-JRNL-594212).

Town, R. P. J., et al., "Dynamic Symmetry of Indirectly Driven Inertial Confinement Fusion Capsules on the National Ignition Facility," *Phys. Plasmas* **21** (5).

Vassilevski, P. S. and U. Villa, "A Mixed Formulation for the Brinkman Problem," *SIAM J. Numer. Anal.* **52** (1), 258–281 (LLNL-JRNL-563632).

Vassilevski, P. S. and U. M. Yang, "Reducing Communication in Algebraic Multigrid Using Additive Variants," *Numer. Linear Algebr.* **21** (2), 275–296 (LLNL-JRNL-637872).

Vassilevski, P. S. and L. T. Zikatanov, "Commuting Projections on Graphs," *Numer. Linear Algebr.* **21** (3), 297–315.

Velsko, S. P., J. Osburn, and J. Allen, "Forensic Interpretation of Molecular Variation on Networks of Disease Transmission and Genetic Inheritance," *Electrophoresis* **35** (21–22), 3117–3124 (LLNL-JRNL-655754).

Wolfe, D. N., et al., "Perspective Piece: Current Strategic Thinking for the Development of a Trivalent Alphavirus Vaccine for Human Use," *Am. J. Trop. Med. Hyg.* **91** (3), 442–450.

Xie, J., X. Chen, and Y. Lai, "Decyling Numbers of Strong Product Graphs Involving Paths, Circuits, Stars, or Complete Graphs," *Int. J. Appl. Math. Stat* **52** (8), 83–90 (LLNL-JRNL-658728).

Xie, S. C., et al., "Interactions Between Cumulus Convection and Its Environment as Revealed by the MC3E Sounding Array," *J. Geophys. Res.-Atmos.* **119** (20), 11784–11808.

Yang, X. J., X. Chen, C. R. Carrigan, and A. L. Ramirez, "Uncertainty Quantification of CO_2 Saturation Estimated from Electrical Resistance Tomography Data at the Cranfield Site," *Int. J. Greenh. Gas Con.* **27**, 59–68.

Yeom, J., et al., "Overcoming the Scalability Challenges of Epidemic Simulations on Blue Waters," *IEEE International Parallel and Distributed Processing Symposium,* (IEEE Computer Society), (LLNL-CONF-648533).

Yi, S. A., et al., "Hydrodynamic Instabilities in Beryllium Targets for the National Ignition Facility," *Phys. Plasmas* **21** (9).

Yin, L., et al., "Stimulated Scattering in Laser Driven Fusion and High Energy Density Physics Experiments," *Phys. Plasmas* **21** (9), 092707.

Zhao, C. F., et al., "Quantifying Uncertainties of Cloud Microphysical Property Retrievals with a Perturbation Method," *J. Geophys. Res.-Atmos.* **119** (9), 5375–5385.

Industrial Collaborators

COMPANY	TOPIC	LLNL CONTACT(S)
Adaptive Computing Enterprises, Inc.	Moab workload manager	Don Lipari
Affymetrix	Microarray evaluation	Tom Slezak and Shea Gardner
AMD	DesignForward networking project	Edgar León
AMD	FastForward processor and memory project	Robin Goldstone and Kathryn Mohror
AMD	FastForward 2 node and memory projects	John May
AMD	Power and energy	Barry Rountree
AMD	Two-level memory, processing in memory	Maya Gokhale, Scott Lloyd, and Brian Van Essen
AOSense, Inc.	Gravity gradiometry in the detection of anomalous mass distribution in vehicles	Vijay Sonnad
Applied Micro and Red Hat	RHEL7 on ARM64	Trent D'Hooge
Argo Navis	Automated cache performance analysis and optimization in Open|SpeedShop	Kathryn Mohror and Barry Rountree
Arista Networks	Low-latency Ethernet networks	Matt Leininger
ARM	Processing in memory	Maya Gokhale and Scott Lloyd
ARM and Cray	FastForward 2 node project	Bronis de Supinski
Battelle	Terrorism risk assessments	Amy Waters and Lisa Belk
Broadcom and Cray	FastForward 2 node project	Bronis de Supinski

COMPANY	TOPIC	LLNL CONTACT
Catalyst	Exploring system software and applications algorithms	Matt Leininger
Cisco Systems, Dell, DataDirect Networks, Intel, NetApp, Mellanox Technologies, QLogic, Red Hat, Oracle, and Supermicro	Hyperion collaboration	Matt Leininger
Commissariat à l'Énergie Atomique	Resource management software	Don Lipari, Jim Garlick, Mark Grondona, and Dong Ahn
Cray	Scalable capacity clusters	Matt Leininger and Trent D'Hooge
Cray	Exploring the Chapel programming language using LULESH	Abhinav Bhatele
Cray	FastForward 2	Bronis de Supinski
Cray	DesignForward 1 networking project	Robin Goldstone
Cray	DesignForward 2	Bronis de Supinski
Cymer	Extreme ultraviolet simulation and analysis	Fred Streitz
Dell Computers	Scalable capacity clusters	Matt Leininger and Trent D'Hooge
Électricité de France	Aeroacoustics	Bill Henshaw
EOG Resources	Seismic processing	Shawn Larsen
ExxonMobil	Cooperative Research and Development Agreement: computational mathematics	John Grosh, Lori Diachin, Timo Bremer, and Martin Schulz
GAMS	Solvers	Barry Rountree

COMPANY	TOPIC	LLNL CONTACT
IBM	Advanced memory concepts FastForward 1 and 2 projects	Bronis de Supinski
IBM	CORAL interactions	Bronis de Supinski
IBM	Deep Computing Solutions	Fred Streitz and Doug East
IBM	Blue Gene/Q common development tools interface co-design	Dong Ahn
IBM	Evaluating the performance of algebraic multigrid on multicore architectures	Ulrike Yang
IBM	High performance storage system	Jerry Shoopman
IBM	Scalable systems, multiple areas	Bronis de Supinski
IBM	Tool interface for OpenMP	Martin Schulz and Ignacio Laguna
IBM	Flash storage systems	Maya Gokhale and Roger Pearce
IBM	DesignForward 2 memory research	Robin Goldstone
IBM	FastForward 2 advanced memory concepts	Maya Gokhale
IBM	DesignForward networking project	Kathryn Mohror
IBM Research	Improvements to CPLEX optimization software geared toward use cases in the electrical grid	Deepak Rajan
IBM Research	Operating systems	Maya Gokhale and Scott Lloyd
IBM Research and Optit SRL	Parallel decomposition schemes for solving large-scale electrical-grid stochastic optimization problems	Deepak Rajan

COMPANY	TOPIC	LLNL CONTACT
IBM Research and Knight Capital Group	Task scheduling with setup times	Deepak Rajan
IBM Research and North Carolina State University	Predictive performance anomaly prevention for virtualized cloud systems	Deepak Rajan
IBM Research, HP Labs, Knight Capital Group, and Bank of America	Scheduling heterogenous jobs in MapReduce environments	Deepak Rajan
IBM Research, Knight Capital Group, StonyBrook University, and Unscramble LLC	Mining large time-evolving graphs for proximity queries	Deepak Rajan
ICFO, Barcelona (Institute of Photonics)	Novel microarray reader	Tom Slezak
Incorporated Research Institutes of Seismology (IRIS)	LLNL hosts the Auxiliary Data Center for IRIS serving 200+ foreign and 200+ domestic collaborator institutions	Stan Ruppert
InfiniBand Trade Association	InfiniBand specifications body	Pam Hamilton
Intel	CRADA for system software research and development	Kim Cupps, Pam Hamilton, Chris Morrone, Martin Schulz, and Barry Rountree
Intel	DesignForward networking project	Matt Leininger
Intel	FastForward processor project	Matt Leininger
Intel	FastForward I/O project	Mark Gary
Intel	FastForward 2 node project	Matt Leininger
Intel	Many integrated core programming environment	Greg Lee
Intel	Power-limited high performance computing	Barry Rountree and Martin Schulz
Intel	Research and development for I/O systems	Mark Gary, Robin Goldstone, and Ned Bass
Intel	Simulations and Visualization of Interconnection Networks	Abhinav Bhatele

COMPANY	TOPIC	LLNL CONTACT
Intel and Cray	High performance architecture for data analytics	Matt Leininger, Robin Goldstone, and Trent D'Hooge
ION Geophysical Corporation	Oil exploration	Shawn Larsen
JTC1/SC22/WG21–The C++ Standards Committee– ISOCPP	C++ standards committee member	Mike Kumbera
John Deere	Design	Debbie May
John Deere	Sensors	Debbie May
Juelich Research Center/Juelich Supercomputing Center	Tools for performance analysis at scale	Bernd Mohr
Krell Institute/Argo Navis Technologies	Open|SpeedShop development and support and the component-based tool framework	Martin Schulz
Laboratory for Laser Energetics and Commissariat à l'Énergie Atomique	Miro and virtual beamline modeling and simulation codes	Kathleen McCandless
Life Technologies	Targeted microbial DNA amplification to enhance sequencing	Tom Slezak
Lightning Bolt Solutions, Ohio State University, and UC Berkeley	Improved models for electrical-grid optimization problems	Deepak Rajan
Mellanox	Long haul InfiniBand	Trent D'Hooge
Mellanox	Hadoop with InfiniBand remote direct memory access	Al Chu and Robin Goldstone
Mellanox	CORAL interactions	Bronis de Supinski
Micron	Processing in memory	Maya Gokhale and Scott Lloyd
MITRE Corporation	Subsurface modeling	Shawn Larsen

COMPANY	TOPIC	LLNL CONTACT
National Instruments	Object-oriented applications of Laboratory view on big physics data	Mike Flegel
NetApp	High performance I/O systems	Marc Stearman and Mark Gary
NSTec	Instrument calibration techniques	Steve Glenn
NVIDIA	FastForward processor project	Bronis de Supinski
NVIDIA	FastForward 2 node research	Robin Goldstone
NVIDIA	CORAL interactions	Bronis de Supinski
NVIDIA	DesignForward 1	Bronis de Supinski
OpenFabrics Alliance, Mellanox, and Intel	OpenFabrics enterprise distribution	Matt Leininger
OpenPower	Power CPU processors	Bronis de Supinski
OpenSFS	Lustre file system development and deployment	Terri Quinn and Chris Morrone
OpenZFS	ZFS file system development	Brian Behlendorf
OpenWorks	Valgrind memory tool and threading tool development	John Gyllenhaal
OSIsoft	Management and visualization of phasor measurement unit data	Ghaleb Abdulla
Pacific Gas and Electric	Department of Energy Office of Electricity project on supply chain cyber security	Dan Quinlan
ParaTools	Development and support of TAU performance analysis tool	Chris Chambreau
PTC	Windchill	Al Churby

COMPANY	TOPIC	LLNL CONTACT
Red Hat	Operating systems	Mark Grondona and Jim Foraker
Red Hat, Appro, Intel, and AMD	Hardware performance counters	Barry Rountree
Robert Bosch LLC	Computational fluid dynamics research using petascale systems	John Grosh
Rogue Wave Software	TotalView parallel debugger scalability and enhanced memory tools	Dong Ahn and Scott Futral
Rogue Wave Software	TotalView enhanced debugging for C++ applications	Matt Wolfe
Samplify	Data compression	Peter Lindstrom
San Diego Gas & Electric, Southern California Edison, Pacific Gas and Electric, and the California Public Utilities Commission	California Energy Systems for the 21st Century (CES-21)	John Grosh, Jamie Van Randwyk, and others
SchedMD	SLURM resource management software	Kim Cupps and Don Lipari
STFC Daresbury	National lab collaborations with industry	Fred Streitz
Tennessee Valley Authority and Applied Communication Sciences	Robust adaptive topology control for Advanced Research Projects Agency-Energy project	Deepak Rajan
TidalScale	Bioinformatics applications on virtual large-memory nodes	Alexander Ames

National Laboratory Collaborators

LABORATORY	TOPIC	LLNL CONTACT(S)
Argonne National Laboratory	CESAR Co-Design Center: tools and performance	Martin Schulz
Argonne National Laboratory	Exascale operating systems: power scheduling	Maya Gokhale, Brian Van Essen, Edgar León, Martin Schulz, and Barry Rountree
Argonne National Laboratory	Simulation technologies for multiphysics simulations	Carol Woodward
Argonne National Laboratory and Oak Ridge National Laboratory	CORAL procurement	Bronis de Supinski
Argonne National Laboratory, Lawrence Berkeley National Laboratory, Oak Ridge National Laboratory	SUPER: Institute for Sustained Performance, Energy, and Resilience	Bronis de Supinski and Daniel Quinlan
Argonne National Laboratory, Lawrence Berkeley National Laboratory, Los Alamos National Laboratory, Oak Ridge National Laboratory, and Sandia National Laboratories	SDAV: Scalable Data Management, Analysis, and Visualization	Eric Brugger
Argonne National Laboratory, Los Alamos National Laboratory, Lawrence Berkeley National Laboratory, Oak Ridge National Laboratory, Pacific Northwest National Laboratory, and Sandia National Laboratories	FastForward 2	Bronis de Supinski
Argonne National Laboratory, Lawrence Berkeley National Laboratory, Sandia National Laboratories; Rensselaer Polytechnic Institute	FASTMath: Frameworks, Algorithms and Scalable Technologies for Mathematics	Lori Diachin
Argonne National Laboratory, Brookhaven National Laboratory, Idaho National Laboratory, Lawrence Berkeley National Laboratory, Los Alamos National Laboratory, Pacific Northwest National Laboratory, Sandia National Laboratories	Energy Facility Contractors Group	Darrel Whitney and others
Atomic Weapons Establishment	Mesh generation co-development	Katie Lewis, Walt Nissen, Jack Middleton, and Cecilia Castillo
Commissariat à l'Énergie Atomique, Sandia National Laboratories, and Los Alamos National Laboratory	NNSA/CEA computer science collaborations	Kim Cupps, Matt Leininger, Trent D'Hooge, Chris Morrone, Rob Neely, Edgar León, Rob Falgout, Walt Nissen, Katie Lewis, Eric Brugger, Bert Still, and Ian Karlin
Los Alamos National Laboratory	ExMatEx Co-Design Center: tools and performance	Martin Schulz
Los Alamos National Laboratory	Monte Carlo N-Particle transport code	Lila Chase
Los Alamos National Laboratory and Sandia National Laboratories	Open|SpeedShop and component-based tool framework	Martin Schulz

LABORATORY	TOPIC	LLNL CONTACT(S)
Los Alamos National Laboratory and Sandia National Laboratories	Tri-laboratory common computing environment tools	Martin Schulz
Los Alamos National Laboratory and Sandia National Laboratories	TotalView enhanced debugging for C++ applications	—
Los Alamos National Laboratory and Sandia National Laboratories	Tri-Lab Operating System Software	—
Los Alamos National Laboratory, Sandia National Laboratories, and Oak Ridge National Laboratories	ExMatEx Co-design Center: materials in extreme environments	Milo Dorr
Los Alamos National Laboratory, Lawrence Berkeley National Laboratory, Oak Ridge National Laboratory, and Sandia National Laboratories	High Performance Storage System	Jerry Shoopman
Los Alamos National Laboratory, Sandia National Laboratories, Oak Ridge National Laboratory, and Pacific Northwest National Laboratory	DOE's Secret National Security Information (SNSI) Network: Architecture, Cyber Security, and Operations	Stan Ruppert, John Heck, and Laura Long
Los Alamos National Laboratory, Sandia National Laboratories, Oak Ridge National Laboratory, and Pacific Northwest National Laboratory	DOE's Secret Restricted Data (SRD) Information Network: Architecture, Cyber Security, and Operations	Paul Masi
Lawrence Berkeley National Laboratory	Extreme resilient discetizations	Jeff Hittinger
Lawrence Berkeley National Laboratory	High-order methods for kinetic simulation of plasmas	Milo Dorr and Jeff Hittinger
Lawrence Berkeley National Laboratory	Modeling throughput on dragonfly networks	Abhinav Bhatele
Lawrence Berkeley National Laboratory and Sandia National Laboratories	ExaCT Co-Design Center: Center for Exascale Simulation of Combustion in Turbulence	Rob Falgout and Ulrike Yang
Oak Ridge National Laboratory	Cybersecurity for energy delivery systems	Dan Quinlan
Oak Ridge National Laboratory	ExMatEx Co-Design Center: modeling	Martin Schulz
Oak Ridge National Laboratory	Solvers for implicit climate simulations	Carol Woodward and Aaron Lott

LABORATORY	TOPIC	LLNL CONTACT(S)
Oak Ridge National Laboratory and Lawrence Berkeley National Laboratory	SAMRAI performance capabilities	Brian Gunney
Pacific Northwest National Laboratory	CONRAD: Contested Operations Reporting and Defense network	Bill Orvis
Pacific Northwest National Laboratory	Cybersecurity for energy delivery systems	Dan Quinlan
Pacific Northwest National Laboratory	Performance analysis for the X-Stack	Martin Schulz
Pacific Northwest National Laboratory and Los Alamos National Laboratory	Terrorism risk assessments	Amy Waters and Lisa Belk
Pacific Northwest National Laboratory	Exascale operating systems/runtime	Maya Gokhale, Brian Van Essen, and Edgar León
Sandia National Laboratories	Cybersecurity for energy delivery systems	Jamie Van Randwyk
Sandia National Laboratories	ExMatEx Co-Design Center: Structural Simulation Toolkit	Martin Schulz
Sandia National Laboratories	Multiscale Climate SciDAC: CAM-SE	Aaron Lott
Sandia National Laboratories	Multiscale Climate SciDAC: Trilinos	Aaron Lott
Sandia National Laboratories	Exascale operating systems/runtime	Maya Gokhale and Scott Lloyd
Sandia National Laboratories, Commissariat à l'Énergie Atomique, Université Catholique de Louvain, UC Berkeley, Old Dominion University, Royal Military College of Canada, Icon Technology and Process Consulting Ltd., and CD-adaptco	International Meshing Roundtable Conference Committee	Katie Lewis