LAWRENCE LIVERMORE NATIONAL LABORATORY
COMPUTATION

2012 ANNUAL REPORT

The cover is a visual representation of text data depicting keywords and themes that are significant to the Computation Directorate as it fulfills its mission. As Lawrence Livermore National Laboratory engages a broader range of national security issues, Computation provides a correspondingly wider range of scientific expertise, especially in energy, intelligence, and cyber security.
Cover design by Daniel S. Moore

Production Team

LLNL ASSOCIATE DIRECTOR FOR COMPUTATION DONA L. CRAWFORD
DEPUTY ASSOCIATE DIRECTORS JAMES BRASE, PATRICIA DAMKROGER, JOHN GROSH, AND MICHEL MCCOY
SCIENTIFIC EDITORS MARC STEARMAN AND PUNITA SINHA
ART DIRECTOR DANIEL S. MOORE
PRODUCTION EDITOR DEANNA WILLIS
EDITORIAL ASSISTANCE CARYN MEISSNER
PROOFREADER PAMELA MACGREGOR
PHOTOGRAPHERS MICHAEL ANTHONY, LEE BAKER, AND NANCY RUTTER
PRINT PRODUCTION CHARLIE ARTEAGO, JR., AND MONARCH PRINT COPY AND DESIGN SOLUTIONS

LLNL-TR-636115

Contents

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344.

Message from the Associate Director

Delivered in early 2012, the deployment of the Sequoia supercomputer is emblematic of the depth and breadth of Lawrence Livermore National Laboratory's (LLNL's) expertise in creating and managing a world-leading high performance computing (HPC) ecosystem. Sequoia is a Blue Gene/Q system, developed in partnership with IBM, and was ranked as the world's fastest supercomputer on the June 2012 TOP500 list. At that same time, the 20-petaFLOP/s (quadrillion floating-point operations per second) machine was ranked as the most energy-efficient computer on the Green 500 list and No. 1 on the Graph500 list for its ability to solve "big data" problems. Early science work on Sequoia, in preparation for its classified stockpile stewardship work, achieved scientific computing breakthroughs, including the most detailed three-dimensional simulation of a beating human heart ever produced.

While we are well known for fielding state-of-the-art systems, the underlying disciplines that have allowed us to do this decade after decade receive less attention. These disciplines include, but are not limited to, expertise in HPC software development and its application, engineering, networking, and storage. Sequoia's success is built on nearly six decades of experience and previous HPC achievements from the UNIVAC in 1953 and the 1970s Cray systems to the more recent ASCI White, ASC Purple, and ASC Blue Gene/L/P/Q. In recognition of the technological innovation that led to Sequoia, *Popular Mechanics* honored IBM and LLNL with a 2012 Breakthrough Award.

Lawrence Livermore's ability to lead HPC development, as embodied in the Advanced Simulation and Computing (ASC) Program and the Computation Directorate, inspires the confidence of our leadership in Washington, D.C. In 2012, LLNL was designated the lead laboratory for the National Nuclear Security Administration (NNSA) and Department of Energy (DOE) Office of Science in managing the FastForward program to push HPC technologies toward the technically challenging goal of exascale (1,000 petaFLOP/s) computing. FastForward is a partnership between seven DOE national laboratories. A request for proposals was made in March 2012, and in June, research contracts totaling $62M were awarded to five leading HPC companies. In an era of fierce international HPC competition, the development of exascale computing capabilities becomes critical not only to our national security missions but to the nation's economic competitiveness in the global marketplace.

In support of economic competitiveness, Livermore stepped up its outreach in 2012, making LLNL supercomputing capabilities available to industry through efforts such as the "hpc4energy incubator," and the High Performance Computing Innovation Center (HPCIC) in the Livermore Valley Open Campus. hpc4energy was born out of a 2011 national energy summit in Washington, D.C. LLNL solicited and selected proposals from energy companies in early 2012. The companies selected were: GE Energy Consulting; Robert Bosch LLC; Potter Drilling, Inc.; ISO New England; United Technologies Research Center; and GE Global Research. These companies collaborated with LLNL scientists and used our HPC resources and predictive simulation expertise to find solutions to urgent energy-related problems. Company representatives and their LLNL collaborators reported on their progress at a November workshop, noting that HPC was a powerful tool for innovation. By all accounts the program was an unqualified success.

To ensure that our partners have access to the most advanced computers, we now offer them computing cycles on Vulcan, a new 5-petaFLOP/s Blue Gene/Q system that was installed in 2012 as a resource to industrial collaborations within HPCIC. Leveraging our capabilities—high performance computers, computer scientists, domain specialists, subject-matter experts—with industrial partners fosters the global HPC ecosystem and drives innovation to stimulate the U.S. economy, which is a crucial part of national security. As part of this effort, the Laboratory and IBM concluded an agreement in June 2012 forming Deep Computing Solutions within the HPCIC. The agreement broadens a 20-year partnership in HPC that has produced three No. 1 supercomputers on the TOP500 list. Under terms of the pact, computer and domain science experts from IBM Research and our Laboratory are working together with a broad range of American industry collaborators to devise HPC solutions that accelerate the development of new technologies, products, and services. Outreach to industry and academia will continue to grow, boosting industry and expanding the Laboratory's skill base.

While we are broadening our outreach to strengthen our computing expertise, our core mission of stockpile stewardship remains the same. Sequoia was the most significant capability NNSA added to stockpile stewardship in 2012. After a rigorous and challenging acceptance process and early, productive science runs, Sequoia's mammoth computational power will be used to assess physical weapons systems and provide a more accurate atomic-level understanding of the behavior of materials in the extreme conditions present in a nuclear weapon.

Although its primary use will be strengthening the foundations of predictive simulation through running very large suites of complex simulations called uncertainty quantification studies, the Sequoia platform proved equally suited to general-purpose science, as evidenced by the groundbreaking real-time heart simulation, Cardioid. Developed in collaboration with scientists at IBM Research, Cardioid was designed to run with high efficiency in the extreme strong-scaling limit. LLNL scientists were able to model a highly resolved whole heart, beating in near real time, representing a more than 1,200-time improvement in time-to-solution from the previous state of the art and performing to within 12% of real-time simulation. While heart modeling may not have the most obvious correlation to national security, the potential benefit to society, to human health, and to medical centers, pharmaceutical companies, and the medical device industry is indisputable. In addition, several techniques developed for the Cardioid code will find applicability in national security codes.

As part of Sequoia's shakeout, Los Alamos National Laboratory (LANL) researchers ran asteroid and turbulence simulations, and Sandia National Laboratories (SNL) scientists explored the properties of tantalum. Initial efforts by Livermore scientists also included a QBox first-principles molecular-dynamics code examination of the electronic structure of heavy metals, research of interest to stockpile stewardship. Qbox was developed at LLNL to perform large-scale simulations of materials directly from first principles, allowing scientists to predict the properties of complex systems without first having to carry out experiments. In addition, LLNL scientists investigated burn in doped plasmas, exploiting the full capability of Sequoia and the code developed for this purpose. These studies deepen scientists' understanding of the effect of dopants on burn, physics that is vital to capsule design for the National Ignition Facility, LLNL's large, laser-based inertial confinement fusion device.

As Sequoia was being assembled and tested, LLNL also managed and procured the second-generation Tri-Laboratory Linux Capacity Clusters, a joint effort between LLNL, LANL, and SNL to reduce the total cost of ownership of robust and scalable HPC clusters across the tri-laboratory complex. LLNL deployed several small- to large-scale commodity clusters that are being used to manage the nuclear weapons stockpile through simulation. These clusters have proven to be highly scalable, reliable, and cost-effective.

In addition to stockpile stewardship, LLNL continues to apply HPC systems to address other critical national challenges including climate change and the need for sustainable energy resources. Long-term planning of energy resource acquisition, environmental impact, and alternative energy models fall under the Laboratory's broader national security mission. Simulations run on LLNL's HPC systems are helping plan intelligent grid management and operations. Computation scientists are working with a global team to improve the scientific understanding of the uncertainty and the consequences of climate warming under various scenarios. Delivering solutions for these complex problems involve multidisciplinary approaches well within the core competencies of Computation.

Among many energy-related projects that demanded our time and expertise this year, and will for the next several years, is the California Energy Systems for the 21st Century initiative. The partnership, which was approved by the California Public Utilities Commissioners in December 2012, is a $150M, five-year agreement to improve the efficiency, security, and safety of California's utility systems. HPC and related domain expertise will be critical to the success of the partnership.

Computation at LLNL is driven by great scientific challenges—seemingly insoluble problems—that demand the best technology, domain expertise, and creativity we can muster. The nature of our work requires that we always look well into the future to the development of new computational tools and capabilities. Only by taking a long view can we begin to unravel longstanding scientific mysteries and find new opportunities for scientific discovery.

Dona Crawford

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, 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 2012. Several notable accomplishments are featured in this section.

A Global Visionary and Leader in HPC

Michel McCoy, whose pioneering work in HPC established LLNL as a world-renowned supercomputing center, was honored with the NNSA's Science and Technology Excellence Award.

For more than 16 years, McCoy has led the procurement and installation of numerous supercomputing systems ranked among the top 10 of the TOP500 list of the world's fastest supercomputers, notably ASCI White, Blue Gene/L, and Sequoia.

As director of LLNL's ASC Program, a deputy director for Computation, and head of the Integrated Computing and Communications Department, McCoy leads the Laboratory's effort to develop and deploy the HPC systems required for the three national weapons laboratories to fulfill their mission to ensure the safety, security, and reliability of the nation's nuclear deterrent without testing. In addition, McCoy has worked tirelessly to make HPC resources available to U.S. industry and the broad spectrum of Laboratory research programs in the belief that the more researchers use HPC, the faster computing technology and applications will develop for the benefit of the country.

The Science and Technology Excellence Award is the highest recognition for science and technology achievements in NNSA. McCoy is the first recipient of the award.

Michel McCoy received the National Nuclear Security Administration's Science and Technology Excellence Award for his "dedicated and relentless pursuit of excellence" from then NNSA Administrator Thomas D'Agostino.

Jeffrey Banks

Presidential Early Career Award

Computer scientist Jeffrey Banks received a Presidential Early Career Award for Scientists and Engineers (PECASE) for helping advance the basic mathematics needed to enhance computer simulations. PECASE is the highest honor bestowed by the U.S. government on science and engineering professionals in the early stages of their independent research careers. Banks was one of 96 early career scientists and engineers to be recognized in 2012.

Banks has dedicated his early science career to computational physics, scientific computation, and numerical analysis, especially pioneering contributions in numerical approximations to hyperbolic partial differential equations focusing on the development and analysis of nonlinear and high-resolution finite-volume and finite-difference methods. His research is about using computers to simulate problems from the physical sciences. Banks is particularly interested in wave problems such as those often found in fluid dynamics or electromagnetics.

The World's Most Powerful Supercomputer

In June, the Sequoia system landed the No. 1 spot on the industry-standard TOP500 list of the world's most powerful supercomputers, the Green500 list of the world's most energy-efficient HPC systems, and the Graph 500 list of systems able to solve big data problems. Sequoia dropped to No. 2 on the November 2012 TOP500.

Sequoia will be a critical tool in helping ensure the safety and reliability of the nation's aging nuclear deterrent. The system consists of 96 racks; 98,304 compute nodes; 1.6 million cores; and 1.6 petabytes of memory. In addition to being orders of magnitude more powerful than such predecessor systems as ASC Purple and Blue Gene/L, Sequoia is also 160 times more power efficient than Purple and 17 times more than Blue Gene/L. (See page 10.)

Sequoia added to its compendium of accolades with a 2012 Reader's Choice Award, which is selected by readers of *HPCwire*, the high performance computing news service. The supercomputer was also recognized by *Popular Mechanics* magazine as one of the 10 top "world-changing" innovations of the year with a 2012 Breakthrough Award.

Security engineer Matthew Myrick accepts a National Cybersecurity Innovation Award for the Master Block List from Debora Plunkett, National Security Agency Information Assurance Director.

Award-Winning Cybersecurity Tool

The development of a simple tool for quickly sharing information about cyber threats earned LLNL and collaborators recognition as one of the most important cyber security innovations of the year.

The Master Block List is a service and data aggregation tool developed by the DOE's Focused Advanced Persistent Threat Group, which is led by LLNL. The blocking service currently provides 10 DOE laboratories and plants with a real-time set of domain names that are known or suspected to be untrustworthy. Sites use the list to create filters and blocks to prevent cyber attacks. (See page 28.)

The task the Laboratory's cybersecurity team faces daily is daunting. In one 30-day snapshot, LLNL received 7.3 million e-mails of which 73% was spam or malicious—a statistic that ranges from 50 to 97% is typical for a national laboratory.

The National Cybersecurity Innovation Awards, established by the SANS Institute in 2011, recognize cyber security achievements that have resulted in tangible improvements in risk reduction on a large scale.

Computation's Institute for Scientific Computing Research hosted 144 participants in the 2012 summer program.

A Record-Setting Summer Program

The Computation Directorate had its second record-setting summer program in a row, hosting 144 students and faculty—96 graduate students, 33 undergraduates, 2 high school students, and 13 faculty—from approximately 90 universities and 18 countries. Students represent the future of computing at LLNL; their success will be our legacy.

Co-op Employer of the Year

LLNL received the Co-op Employer of the Year Award from the University of the Pacific (UOP) School of Engineering and Computer Science. This award is given to the employer that faculty, staff, and students determine has brought great contributions to their co-op program in the past and in the present.

LORI DIACHIN, DIACHIN2@LLNL.GOV

Computation's Core Competencies Reinforce LLNL's Superior Capabilities

In the past year, the Computation Directorate participated in an institutional effort to define a set of enduring core competencies that are critical to ensuring the relevance and vitality of Lawrence Livermore National Laboratory (LLNL). Computation's competencies were identified to be high performance computing (HPC) and computational mathematics, both of which support a wide variety of mission-critical applications at LLNL, programs in the Department of Energy's (DOE's) Office of Science and its Applied Energy Office, other federal agencies such as the Department of Defense, and industrial collaborations such as Exxon Mobil and the California Energy Systems for the 21st Century.

The HPC core competency includes the research, development, and deployment of the Laboratory's computer ecosystems, consisting of massive numbers of multicore processors connected via high-speed interconnects and running specialized system software. LLNL delivers a world-class HPC environment with constantly evolving hardware resources and a wealth of HPC research and application deployment expertise that enables application science to be performed on ever more complex computers. Our leadership is evidenced in part by the Sequoia system's number one ranking on the TOP500 list in June 2012, by being named a finalist for the Gordon Bell Prize at the 2012 Supercomputing Conference (SC12), and by fulfilling many leadership roles within DOE to help realize exascale systems. Going forward, we will focus on all levels of the software ecosystem for HPC systems, addressing exascale-relevant issues, such as resilience and power, and ensuring that LLNL computer facilities remain first-rate.

Our second core competency, computational mathematics, is the development and analysis of new models, algorithms, and software for predictive simulation of physical phenomena. LLNL has internationally renowned expertise in numerical partial differential equation (PDE) method development and analysis, high-order discretization techniques, linear and nonlinear solvers, verification, and uncertainty quantification. This expertise is captured in top-tier publications and widely used software. In the past year, an LLNL computational mathematics expert was recognized with a Presidential Early Career Award for Scientists and Engineers, and a Computation researcher was selected to co-lead efforts to help define DOE's exascale math programs. Our future computational mathematics work will address the continuing challenges of predictive simulation of complex physical phenomena by eliminating or reducing the data avalanche through mathematical analysis and ensuring that simulation can be used to confidently make decisions in a timely manner.

In addition to identifying existing core competencies, the Laboratory defined initiative areas aimed at establishing emerging competencies that are necessary to meet current and future mission needs. In Computation, we are leading the Lab-wide data science initiative. We are defining a strategy by identifying our key strengths, such as graph analytics, scientific visualization, image processing, machine learning, distributed data management, and persistent memory research, and by performing a gap analysis to target new areas of investment. We will integrate both aspects into a coherent program that addresses data challenges in cybersecurity, critical infrastructure resilience and protection, large-scale scientific simulations, and experimental user facilities.

Higher Order Methods Enable Efficient Seismic Modeling

Wave propagation phenomena are important in many DOE applications, such as nuclear explosion monitoring, seismic exploration, estimating ground-motion hazards and damage due to earthquakes, nondestructive testing, underground facilities detection, shale-gas applications, and acoustic noise propagation. In seismic exploration, computer scientists can improve the resolution in their seismic imaging by capturing shorter wavelengths in the simulations, which corresponds to higher frequencies. The estimation of earthquake-induced damage of large structures, such as nuclear power plants or dams, also require shorter waves to be resolved in ground-motion simulations.

During 2012, LLNL computer scientists began developing a fourth-order code for applied seismic wave propagation. Compared to the current second-order method, the new code is approximately eight times faster and uses eight times less memory. These results correspond to doubling the frequency resolution in simulations using the same computer hardware. The fourth-order method was also used for seismic source estimation based on solving a nonlinear optimization problem. One application of this work is to estimate the source parameters due to seismic events, such as hydrofracturing of rock, near shale-gas sites or geothermal energy plants.

Progress in 2012

In seismic applications, ground motion can be modeled by solving the elastic wave equation numerically. These simulations require significant computational resources. For example, simulating 100 seconds of motion due to an earthquake on the Hayward fault in the San Francisco Bay Area, including effects of realistic topography, heterogeneous materials, and viscoelastic dissipation, requires about 1 billion grid points and 10,000 time steps. On this grid, motions up to a frequency of 0.5 hertz can be captured using a second-order accurate code, such as WPP (developed with internal Laboratory Directed Research and Development [LDRD] funds at LLNL). The WPP simulation required approximately 4.5 hours of processor time on 2,048 cores of the Sierra machine at Livermore Computing.

A fundamental challenge in wave simulations is the rapid increase in computational cost that comes with increasing the frequency. The wavelength in the solution decreases when the frequency increases. If the frequency is doubled in a three-dimensional (3D) simulation, the total number of grid points must increase by a factor of 8. Furthermore, the number of time steps grows linearly with frequency. As a result, the computational cost increases by a factor of 16 each time the frequency is doubled. Distributing the simulation over many cores on a large parallel machine can only partially offset the high computational cost, because the calculation is inherently sequential in the time direction.

In a paper published in 1972, Kreiss and Oliger explained the theoretical benefits of higher order accurate methods for wave propagation. In this context, the order of accuracy measures the rate at which the error in the numerical solution decreases, as the grid is refined. To quantify how fine the computational grid needs to be, they introduced the concept of "number of grid points per wavelength." Kreiss and Oliger pointed out that the benefit of higher order accurate methods is not necessarily to make the numerical solution more accurate but to reduce the computational

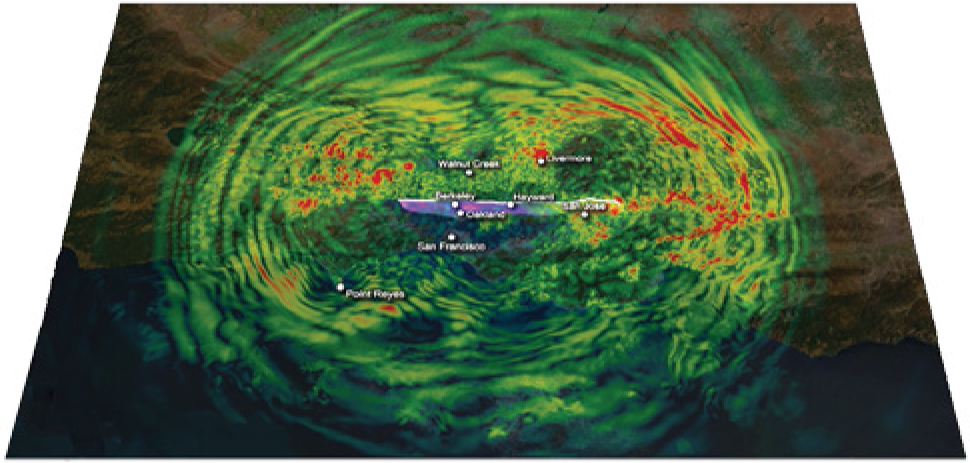

This simulation shows horizontal ground velocity caused by a hypothetical magnitude 7.0 earthquake on the Hayward fault in the San Francisco Bay Area. This earthquake simulation and others were shown every day at the California Academy of Sciences.

N. ANDERS PETERSSON, PETERSSON1@LLNL.GOV

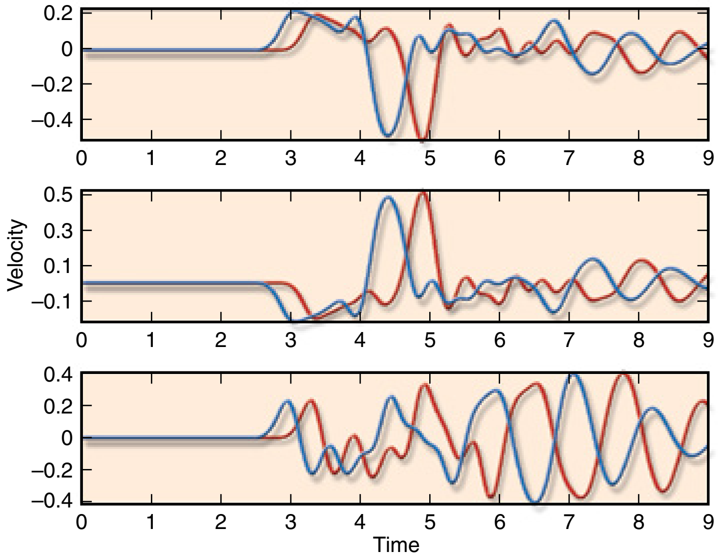

These graphs show how time-dependent ground velocity corresponds to the initial (red) and exact (blue) source parameters.

cost for obtaining a solution within an acceptable accuracy. Compared to a low-order method, a higher order method needs fewer grid points per wavelength for obtaining the same accuracy, and the error grows at a slower rate for long time simulations.

Even though higher order methods for simpler wave equations have been available for a long time, they were only recently generalized to work for seismic wave simulations. In the seismic application, it is of particular importance to impose the free surface condition on the realistic topography, allowing for heterogeneous elastic material properties and viscoelastic dissipation. Also, the full benefit of a higher order method is only realized if all parts of the modeling are performed to a matching high order, for example the time-integration, boundary conditions, seismic source models, local mesh refinement, and far-field closure of unbounded domains.

During 2012, we began implementing a fourth-order method for 3D seismic applications using finite-difference methods that satisfy the principle of summation-by-parts, which guarantees stability of the numerical solution. Preliminary comparisons indicate a significant gain in efficiency. For example, a small 3D problem that takes 600 seconds to solve using 128 cores with WPP, can now be solved in approximately 78 seconds with the new fourth-order method—almost eight times faster. The fourth-order method is faster than the second-order method because it only needs about half the number of grid points in each spatial direction to calculate a solution with comparable accuracy. Or, alternatively, if the fourth-order method uses the same grid, the frequency content can be doubled. For example, the fourth-order method would give clearer images of the subsurface in seismic exploration applications.

The improved efficiency of the fourth-order method makes gradient-based optimization methods tractable for practical 3D seismic applications. During 2012, we began developing numerical methods for estimating seismic sources from observed ground motions. In particular, we studied the inverse problem of determining the source parameters of a small seismic event to minimize the full waveform mismatch between observed and simulated ground motions. The mathematical properties of the fourth-order method can be used to derive an adjoint wave equation, which is used to calculate the gradient of the mismatch. The gradient of the mismatch constitutes the basic building block of the nonlinear optimization algorithm, which iteratively improves the source parameters. One application of this work is to estimate the source parameters due to hydrofracturing of rock near shale-gas sites or geothermal energy plants. Estimating the source mechanism provides important information about how the rock fractures, which is not available with current technology.

Many challenges remain. For example, one of the greatest problems facing seismology is the lack of detailed knowledge of subsurface wave propagation properties. In future work, we plan to improve the source estimation by attempting a simultaneous optimization of source and material properties.

Sequoia Integration Team Overcomes Challenges to Achieve Momentous Success

The challenges related to predictive simulation require that we build better science models and quantify the margins of uncertainty in our integrated calculations. Sequoia, a supercomputer with a peak performance of 20 quadrillion floating-point operations per second (petaFLOPS), is the first computer deployed in the DOE with the capability to address these challenges. Sequoia is also the first advanced technology platform that will tackle production environment simulations. Sequoia has 1.6 petabytes (PB) of memory and 1.6 million cores. This proliferation of cores is two orders of magnitude greater than any other LLNL production simulation HPC resource and represents a daunting code-porting task.

Sequoia racks began arriving at LLNL in January 2012, and all 96 racks were delivered by early April. The integration of the supercomputer was a long, challenging process. Several manufacturing defects were discovered during various phases of integration, which negatively impacted the machine's reliability. Lessons were also learned from installing and operating Sequoia's unique cooling infrastructure. Despite these challenges, Sequoia completed a 23-hour LINPACK benchmark run at 16.324 petaFLOPS to briefly take over the number one spot on the June 2012 TOP500 list of the world's most powerful supercomputers. More importantly, several scientists have achieved excellent results with their codes at scale on Sequoia.

Progress in 2012

The integration of any serial No. 1 advanced technology system is a daunting challenge. Despite Livermore's 60-year history deploying first-of-a-kind supercomputers, Sequoia proved to be more difficult than most. The integration schedule was exceptionally tight with five weeks between delivery of the final racks and the milestone to complete a full-system LINPACK run. During these five weeks, several problems introduced schedule delays. First, faulty adhesive on the electromagnetic interference gasket of the bulk power modules caused inch-long gaskets to occasionally fall off during assembly or repair. The problem, which also occurred in dozens of other Blue Gene/Q (BG/Q) racks around the world, was not detected until 48 racks were on the LLNL computer room floor. IBM made the difficult but correct decision to power off all the BG/Q racks and rework each bulk power module gasket. At Livermore, this process involved re-adhering 1,728 gaskets in the 48 racks, which took four days.

The next problem we encountered was false tripping of slow-leak detection by the water monitor. Each computer rack has a water monitor that measures the flow and pressure of the cooling liquid as it travels in and out of the rack. If pressure or flow differences are detected, the monitor assumes a leak exists and powers the rack off. In Sequoia's case, the racks were being turned off erroneously. Several full-system runs failed before we disabled slow-leak detection and relied on our under-floor sensor ribbon for leak

After a yearlong integration period, Sequoia's 96 racks, 1.6 million cores, 1.6 petabytes of memory, and 20 quadrillion floating-point operations per second are yielding impressive scientific results. *Photo by Ross Gaunt.*

KIM CUPPS, CUPPS2@LLNL.GOV

detection until the faulty algorithm in the monitor was resolved.

The third issue we encountered during the initial integration period was bent compute pins. Sequoia has 98,384 compute cards, and each connects to a system motherboard with pins that press into the motherboard's receptacles. During installation, we found that the compute pins were subject to bending if not inserted perfectly, resulting in inconsistent electrical connections. Once we identified the issue and determined the level of care and precision with which the pins must be handled, more than 6,000 cards were inspected and reseated.

The fourth issue resolved in phase I of the integration involved hot compute nodes in the racks. The compute cards that were being cooled by the last water to leave a drawer were hotter than what is considered acceptable. Our solution was to increase the flow of water to the racks, which decreased coolant temperature and resulted in better cooling for the cards in the hottest spots. To achieve this solution, we had to individually balance the water flow to each rack, take temperature and flow rates using both LLNL measuring equipment and the BG/Q water monitor, and then compare the diagnostic readings. Discrepancies, then had to be resolved, in both measured flow and temperature between our equipment and the BG/Q monitor. Eventually, we learned that the LLNL equipment was not as accurate as the monitor.

The final integration phase was scheduled to culminate with system acceptance by September 30, 2012, but this phase was not without its own set of challenges. We encountered more bent pins and various software issues, including the main service node running out of memory and unceremoniously crashing. However, the biggest issue was discovered near the end of August when the system began to experience an unusually high compute-node memory failure rate. IBM initially thought the problem was caused by a memory controller firmware issue, and they made the appropriate fixes. However, we eventually determined that the problem was in fact caused by microscopic cracks in the compute chip solder, which affixes the chip to the compute card. The cracks were detected by injecting a dye into the top of the chip and prying the chip off the board to see if the dye leaked into the substrate below. This "dye and pry" process uncovered a significant manufacturing defect. During a quality test in the manufacturing process, it was discovered that a faulty piece of testing equipment had applied unequal force to the compute card. The unequal force produced cracks in the solder. These cracks widened over time as the temperature of the compute card cycled during Sequoia integration. The memory controller resides in the corner where the solder was cracked resulting in a high memory failure rate. After identifying the problem, we replaced more than 25,000 compute cards, a process that was completed in mid-October. LLNL accepted Sequoia in December 2012.

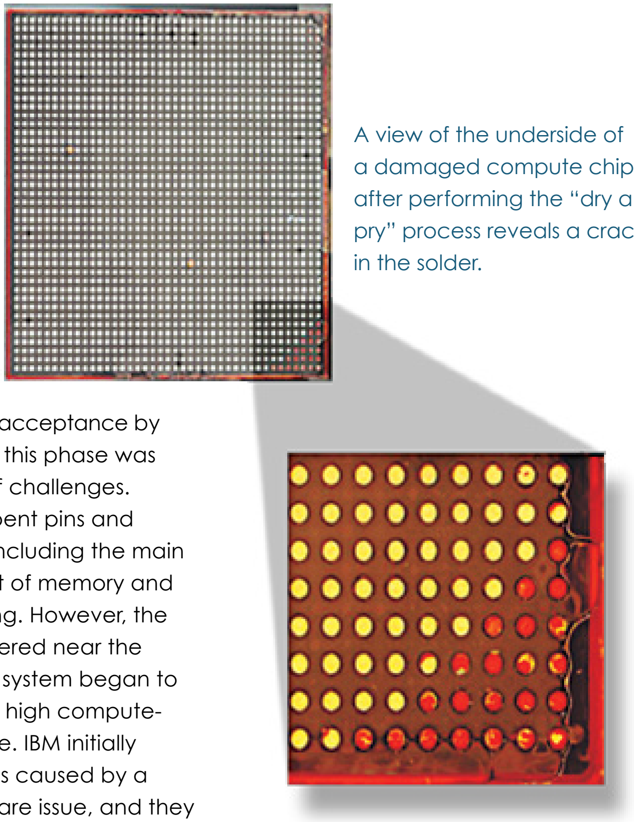

A view of the underside of a damaged compute chip after performing the "dry and pry" process reveals a crack in the solder.

Despite these very challenging conditions, two of five Gordon Bell Prize finalists ran their codes on Sequoia with impressive results. A team from Argonne National Laboratory achieved almost 70% efficiency with their code. An LLNL team, running the Cardioid code, demonstrated an accurate electrophysiological model of the human heart for hundreds of heartbeats, where previously only tens of heartbeats could be modeled (see page 16). High-fidelity models are critical in the development of drugs to efficaciously treat various heart conditions.

Given the innovation, complexity, and scale of Sequoia, it is not surprising that the integration was LLNL's most arduous in recent history. However, the challenges have not prevented early users from exploiting the massive capabilities provided by 1.6 million cores working harmoniously. Sequoia will move to the classified network in spring of 2013 and begin performing production environment simulations and uncertainty quantification runs.

COMPUTATION.LLNL.GOV/CASC/SUNDIALS

CAROL WOODWARD, WOODWARD6@LLNL.GOV

Nonlinear Solvers and Stiff Integrators Speed Up Simulation Codes

Nonlinear, time-dependent systems are ubiquitous in scientific simulations. When these systems exhibit fast modes decaying to stable manifolds, the systems are stiff, requiring implicit solution approaches for simulation codes to be efficient. Such is the case in applications like power grid simulation, dislocation dynamics, chemical combustion, and subsurface flow.

Implicit systems are particularly difficult to solve due to the need for fast and efficient nonlinear solvers for each time step solution. Computation researchers are developing solution approaches and software for nonlinear implicit systems. In the past year, new nonlinear solvers provided a 25% speedup on 512 processors for the LLNL dislocation dynamics code, ParaDiS. In addition, new thread parallelism in the LLNL-developed, open-source SUite of Nonlinear and DIfferential/ALgebraic equation Solvers (SUNDIALS) package provided speedup of a power grid simulation.

Progress in 2012

One challenge in solving nonlinear systems is that evaluating nonlinearities can be extremely expensive. In dislocation dynamics for materials science, calculating the forces between dislocations can dramatically dominate run times, and solvers generally require one evaluation per iteration. LLNL's ParaDiS code simulates dislocation dynamics using a nonstiff time integrator with a fixed-point nonlinear solver. We integrated a new accelerated fixed-point option into the nonlinear solver code, KINSOL, which is part of the SUNDIALS package. This Anderson accelerator mixes a series of prior residuals to form an update of the solution, allowing more progress within each iteration. Working with LLNL's ParaDiS team, we integrated this solver into ParaDiS and observed a 25% speedup on a 512-processor computation of the mechanical properties of single-crystal body-centered-cubic iron. Future work includes incorporating a stiff time-integration method from SUNDIALS into ParaDiS.

Simulating the electric power grid poses its own set of challenges. These nonlinear problems come from both power transmission (dynamic) and power distribution (steady). Transmission problems exhibit system discontinuities and require tracking of instabilities in the solution states. These characteristics constrain selection of adaptive time steps in numerical time integration. In addition, both systems lead to large, sparse linear problems within nonlinear solution updates, making parallelization of these solution processes challenging. We are investigating methods for selecting adaptive time steps that can accommodate discontinuous conditions as well as instabilities while still maintaining high-order accuracy where possible. New research is exploring how to effectively use thread parallel sparse direct solvers within the time-integration process.

Effectiveness of KINSOL's new accelerated nonlinear solver for the ParaDiS dislocation dynamics code is shown in serial. Total simulated time versus time step shows 66% increase with the accelerated method.

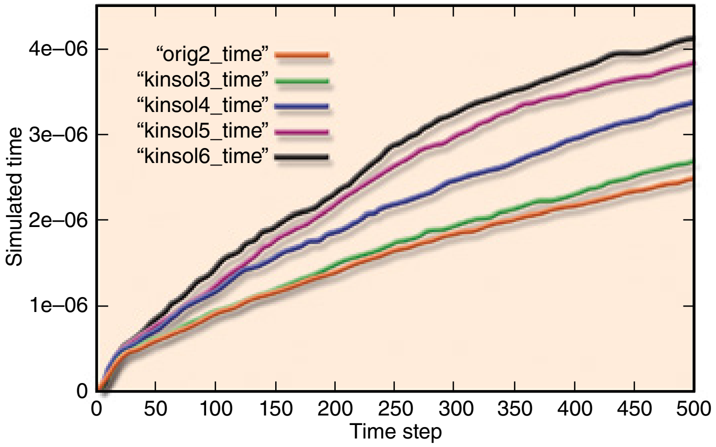

The transition to an increase in renewable energy devices will change the way power grids operate in the future. In particular, with more solar generators on homes and the inclusion of wind-energy generation, the separation of timescales and locations of power distribution and transmission is no longer clear. In a new LDRD project, Computation and Engineering researchers are investigating the coupling between distribution and transmission with the goal of developing an integrated modeling capability for simulations that require it. The SUNDIALS stiff time integrator and nonlinear solvers package will be critical to this work.

SUNDIALS, which is used in LLNL combustion chemistry, fusion, and subsurface flow applications, was downloaded more than 3,600 times in 2012 and more than 3,000 times in both 2011 and 2010. In one success story, the new thread-parallel capability in the SUNDIALS package helped LLNL and Iowa State University researchers achieve moderate speedups on a transmission power grid simulation.

MATTHEW LEININGER, LEININGER4@LLNL.GOV

TLCC2 Clusters Deliver Increased Scalability, Reliability, and Performance

Reducing the total cost of ownership of robust and scalable HPC clusters is a significant challenge that impacts many programs at LLNL. In 2007, LLNL experienced tremendous success when it led a first-of-a-kind effort to build a common capacity hardware environment, called the Tri-Lab Linux Capacity Clusters (TLCC1), at the three National Nuclear Security Administration laboratories—Lawrence Livermore, Los Alamos, and Sandia national laboratories. The TLCC1 experience proved that deploying a common hardware environment multiple times at all three sites greatly reduced (by an order of magnitude) the time and cost to deploy each HPC cluster.

Building on the success of TLCC1, the second generation of Tri-Lab Linux Capacity Clusters (TLCC2) was deployed in 2012. This year, LLNL and its laboratory and industry partners—Los Alamos, Sandia, Appro, Intel, QLogic, and Red Hat—sited 11 small- to large-scale commodity clusters that are critical to maintaining the aging U.S. nuclear weapons stockpile without underground testing. Since their deployment, the TLCC2 clusters have proven to be some of the most scalable, reliable, and cost-effective clusters that LLNL has ever brought into service.

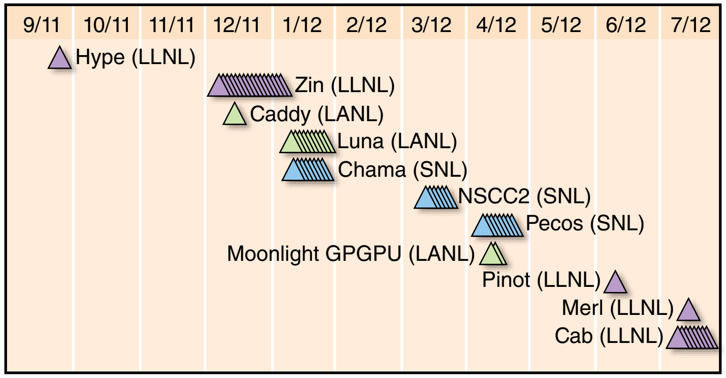

Eleven TLCC2 clusters were installed at Livermore, Los Alamos (LANL), and Sandia (SNL) in just 11 months. Each triangle represents a 162-node scalable unit building block of TLCC2 hardware.

Progress in 2012

The LLNL-led TLCC2 project is based on a unique approach for procuring, deploying, and operating commodity clusters. The TLCC2 systems are dedicated to the small- to medium-sized capacity and mid-range capability jobs required by LLNL's national security and institutional computing programs. Each cluster comprises one or more scalable units (SUs), and each SU is composed of 162 Appro nodes. The nodes use 2-socket 8-core Intel Sandy Bridge processors, contain 32–64 gigabytes of memory per node, and are connected together with a high-performance QLogic quad data rate InfiniBand interconnect.

Several TLCC2 platforms were sited at LLNL in 2012, including the 18-SU (2,916 nodes) Zin cluster, the 8-SU (1,296 nodes) Cab cluster, and a few smaller systems. A total of 70 TLCC2 SUs (11,340 nodes) forming 12 clusters were deployed across Livermore, Los Alamos, and Sandia. The TLCC2 hardware platforms are coupled to a common software environment, which has allowed the three laboratories to significantly reduce the total cost of ownership of these clusters relative to pre-TLCC procurement and deployment practices.

The TLCC2 platforms are highly reliable and represent the most scalable and high-performance commodity clusters LLNL has ever deployed. During the first year of operation, LLNL's TLCC2 platforms demonstrated hardware reliability more than 10 times greater than the TLCC1 platforms. This improvement is due to improved reliability of several system components, including memory, motherboards, power supplies, processors, and network cards. The platforms' overall performance and scalability for various LLNL scientific and engineering applications, including the pF3D code and Zrad3D benchmark, has set a new standard for commodity systems. pF3D and Zrad3D represent two of the most challenging multiphysics workloads run on TLCC systems. pF3D simulations on more than 4,000 CPU cores show that TLCC2 systems are running 3 times faster than TLCC1 systems. The Zrad3D benchmark, using more than 10,000 CPU cores, is running 2.5 times faster on TLCC2 systems compared to TLCC1. Zrad3D is also showing impressive scaling results to almost 45,000 CPU cores.

These successes in reliability, performance, and scaling are very beneficial to LLNL programs. The TLCC2 platforms allow users to run larger simulations with a higher job throughput than were previously possible on commodity systems. LLNL and its tri-lab partners will continue to leverage the TLCC experience. Planning for TLCC3 is currently under way.

PCMDI9.LLNL.GOV

DEAN N. WILLIAMS, WILLIAMS13@LLNL.GOV

Software Infrastructure Empowers the Study of Climate on a Global Scale

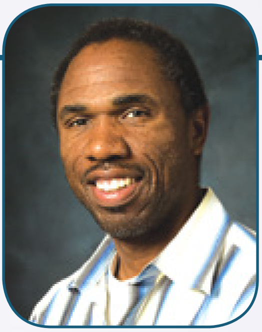

For more than a decade, LLNL has led the climate community's efforts to build an information and knowledge infrastructure that has revolutionized how large-scale climate model output is organized for worldwide dissemination and use. The Earth System Grid Federation (ESGF) is a multiagency international collaboration that provides interactive views of future climate changes based on projected natural and human factors. The newest iteration of ESGF offers an immense, computerized climate database that standardizes and organizes observational and simulation data from 21 countries, allowing scientists to compare models against actual observations. Today, 25,000 users (researchers and nonresearchers) from 2,700 sites on six continents are sharing data via ESGF and gaining an unprecedented understanding of climate change, extreme weather events, and environmental science. More than 2 PB of data have been downloaded to the climate community through ESGF, making it one of the most complex, successful big data systems in existence.

Progress in 2012

Laboratory computer scientists coordinated the most recent iteration of ESGF, which migrates and consolidates massive amounts of data from 25 international projects, including Climate Science for a Sustainable Energy Future, the North American Regional Climate Change Assessment Program, and the Coupled Model Intercomparison Project (CMIP). The data integration software objective is to help communities of data providers quickly merge or migrate information from their centralized data warehouses, operational stores, or other systems to an integrated, dynamic, federated, distributed infrastructure. LLNL sends providers a software stack that allows them to publish their data to an online portal. Similar to online shopping, researchers in the federation can query available data, add it to a "shopping cart," and then download it when they are ready. Replicas of the most popular data sets are archived at LLNL for downloading and backup.

The ESGF software stack includes four major components: (1) data node to publish and serve data through various protocols; (2) index node to harvest metadata and enable data discovery; (3) identity provider to register, authenticate, and authorize users; and (4) compute node for data reduction, analysis, and visualization. The platform unifies data, events, and services with secure federation portal access, common metadata, and application programming interfaces. Regardless of which secure portal is used, the enterprise presents a single collective view of the decentralized data warehouses. It supports a broad range of applications, including data migration, replication, versioning, master data management, and Web service integration.

ESGF relies on an active data hub—a unified data store relying on peer-to-peer architecture that allows nodes to interact on equal bases with flexibility in their configuration. Metadata shared among projects helps fully integrate the repository of data and components for usability and interoperability. ESGF also promotes standard conventions for data transformation, quality control, and data validation across processes and projects.

For the CMIP Phase-5 project, ESGF data exceeds 2 PB and includes simulations from 27 modeling centers in 21 countries. LLNL's Program for Climate Model Diagnosis and Intercomparison serves as a "program office" for CMIP, in which all the major coupled atmosphere–ocean general circulation models perform a common set of experiments. These models allow researchers to simulate global climate changes related to carbon dioxide increases in the atmosphere. The results of the simulations will be used in the fifth United Nations–sponsored Intergovernmental Panel on Climate Change assessment report, set for release in 2013.

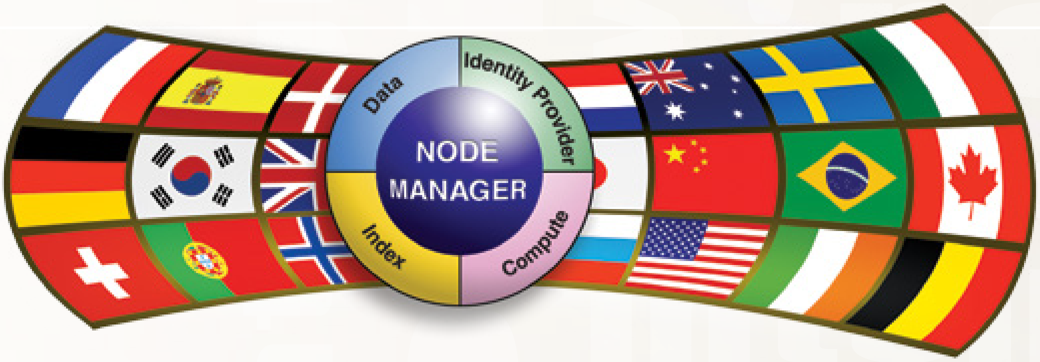

The Earth System Grid Federation seamlessly joins climate science data archives and users around the world. Researchers, policymakers, and others can easily access this data through a client, script, or Web-based interface.

PATRICIA DAMKROGER, DAMKROGER1@LLNL.GOV

Applications and Research

Applications and Research Portfolio Focuses on Programmatic Objectives and Innovation

The most effective way to describe the applications and research area of the Computation Directorate is to highlight the extraordinary capabilities of the people in the organization. The breadth of skill across the directorate allows Lawrence Livermore National Laboratory (LLNL) to be a global leader in diverse mission spaces. We hire computer scientists and mathematicians who become domain experts and specialists in high performance computing, control systems, data systems, equation of state, and thermal dynamics, among myriad fields. LLNL is unique in that we afford our employees the choice and opportunities to either remain world-renowned leaders in a specific discipline area or to broaden into domains experts, offering them continual challenges and a lifelong career path. We strive to maintain our vibrant community of experts by preserving, as a whole, the right mix of expertise and talent to meet current and future national security needs.

The goal of Computation's applications and research portfolio is to pursue pioneering methods, algorithms, and models in computer science and applied mathematics. The research arm of Computation is responsible for defining the path forward for Laboratory programs so that they remain at the cutting edge of innovation. Our research portfolio is tightly coupled to the drive for exascale and demands for big data solutions. This year, we shifted the directorate's internal investment portfolio to include an increased focus on data science. For example, many programs across the Laboratory are using Hadoop to solve their individual programmatic needs, but the MapReduce architecture is not the best answer for the entire institution. We need solutions for handling large quantities of real-time streaming data and massive distributed data sets to serve LLNL's mission needs for today and in the future. These issues are being strategically planned and thoroughly researched as we concurrently push toward exascale computing. Our researchers and scientists work closely with LLNL internal programs, academia, industry, and other national laboratories to provide solutions to compelling big data challenges.

Our application development staff are a critical element of our organization. They are ensconced within the programs so they can thoroughly understand the customers' problems, deliver solutions, and ensure their software products are usable, reliable, and maintainable. Our developers are well versed in modern languages and programming models. They also have a deep and fundamental understanding of the computer science domain and are often experts in their program's domain (e.g., biology, weapons, climate change, and energy technology). The opportunities for application developers at LLNL are vast and varied. Some are out in the field, deploying technologies around the world; others are part of around-the-clock operations or emergency campaigns. While their jobs vary, our application developers serve a common purpose to develop and deploy production-quality software systems and networking solutions that enhance the nation's security.

This section of the annual report highlights only a few select areas that the Computation Directorate impacts at LLNL, but it is a good representation of the diversity and breadth of our world-class employees.

return to Table of Contents

RETURN TO TOC

Scientists Create Simulation Capability for Modeling Cardiac Arrhythmias

Scientists at LLNL, working closely with colleagues at IBM, have developed the capability to study cardiac electrophysiological phenomena, such as arrhythmia in a human heart, discretized at near-cellular resolution with rapid time to solution. Previous scientific investigation had generally invoked simplified geometries or coarse-resolution hearts, with simulation duration limited to tens of heartbeats. This new capability is made possible by a highly scalable code called Cardioid, which mimics in exquisite detail the electrophysiology of the human heart, including activation of heart muscle cells and the cell-to-cell electrical coupling. Cardioid was developed to run with groundbreaking efficiency on Sequoia, a 20 quadrillion floating-point operations per second (petaFLOPS) IBM Blue Gene/Q system at LLNL. High-resolution simulation of the human heart can now be performed more than 1,200 times faster than published results in the field. The utility of this capability was demonstrated by simulating, for the first time, the formation of transmural (through the wall) re-entrant waves in a three-dimensional (3D) human heart. Such wave patterns are thought to underlie torsades de pointes, an arrhythmia that indicates a high risk of sudden cardiac death. This new simulation capability has the potential to impact a multitude of applications in medicine, pharmaceuticals, and implantable devices.

Progress in 2012

The code Cardioid simulates the electrophysiology of the heart by solving coupled reaction-diffusion equations on a finite-difference grid. The reaction term represents cellular biophysics, such as ionic channels and concentrations, and the diffusion term captures the coupling between cells, taking into account the anisotropic conduction due to the fiber structure of the heart. The reaction model is parameterized to simulate various types of cardiac cells, including M-cells, generally reported to be located near the mid-myocardium and believed to play a role in the formation of arrhythmias.

Anatomical models of human ventricles are reconstructed from cryosectional images from the Visible Human Project of the National Library of Medicine. The 2D images are stacked to form a 3D finite-element mesh, from which is derived a 3D Cartesian finite-difference grid for the solution of the reaction-diffusion equations. The inferred fiber geometry is used to calculate conductivity tensors for the diffusion computation. Simulated electrocardiograms (ECGs) are obtained by reconstructing the torso using methods similar to those used for reconstructing the heart, then computing surface potentials by solving a diffusion equation. ECGs can be directly compared to similar measures in patients with 12-lead configurations as are typically collected in clinical studies.

Cardioid has been highly optimized to make effective use of the Sequoia Blue Gene/Q architecture. The optimization requires exploiting multiple aspects of parallelism including nodes, cores, hardware threads, and single instruction, multiple data (SIMD) units. Computational work is divided across nodes by decomposing the heart into a collection of rectangular domains, with a one-to-one correspondence between domains and computational nodes. Work is further divided among the 16 cores on each node; typically 14 of the 16 cores are assigned to the reaction terms and the remaining 2 to the diffusion terms. Each core executes four hardware threads. This enables, for example, one thread to be loading data from memory while another performs arithmetic. Multiple execution threads also cover cache latencies and reduce pipeline stalls. Finally, data structures and algorithms are structured to utilize the quad SIMD double-precision floating-point unit capable of performing up to 8 floating-point operations per core, per cycle.

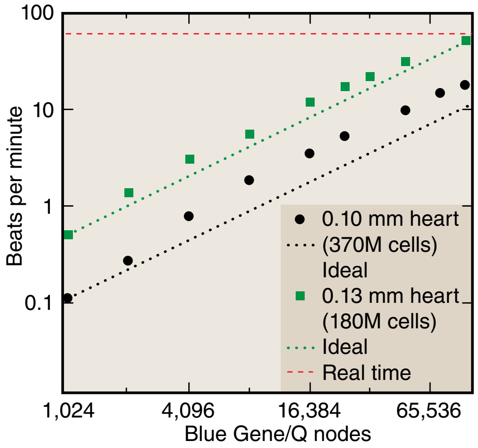

This graph shows strong scaling of Cardioid on Sequoia for two resolutions of a reconstructed human heart.

Cardioid operates in the strong scaling limit when run on the full Sequoia. Using a heart model with 370 million cells (0.1 millimeter resolution), each core is assigned fewer than 300 tissue cells. With only 20 variables evolved per grid point, the computing required for a time step can be accomplished in approximately 60 microseconds of wall time. Special care is needed to reduce overheads to fit within this time frame. For example, the halo exchange communication needed for the

ART MIRIN, MIRIN1@LLNL.GOV

diffusion operator is implemented using IBM's low-level systems programing interface, instead of the more conventional message passing interface (MPI), resulting in more than an order of magnitude reduction in communication latency. Instead of Open Multi-Processing (OpenMP) barriers for the threads, special atomic instructions provided in the Blue Gene/Q hardware are used, greatly reducing synchronization overhead.

Cardioid scales almost perfectly to Sequoia's 1.6 million processors and operates at close to 60% of peak theoretical performance when applied to the 0.05-millimeter heart. For a reconstructed human heart at 0.10-millimeter resolution (twice as fine in each direction as typically reported), we obtain a throughput rate exceeding 18 heartbeats per wall minute—more than 1,200 times faster than previously published results (when normalized to our fine grid and common timestep). For a human heart at 0.13-millimeter resolution, it takes 67 seconds of wall time to execute a 60-second simulation, which is approximately 12% slower than real time.

The scientific potential of Cardioid is illustrated by the first simulation of the formation of a transmural re-entrant wave pattern in a 3D human heart associated with the presence of a drug known to be arrhythmogenic. Following established experimental protocols, an extra stimulus is applied to a steady train of stimuli, while blocking the rapid potassium current in the reaction model to simulate the presence of the drug d-sotalol. The simulation captures the development of a re-entrant arrhythmia that forms around an M-cell island structure and continues to rotate in the presence of the simulated drug. Computed ECGs for cases with and without the drug show a delayed T-wave for the case with the drug, similar to the behavior seen in clinical records. Such behavior can prevent the heart from repolarizing, rendering subsequent heartbeats ineffective and possibly leading to sudden cardiac death. Although further investigation is needed to fine-tune the model and clarify the mechanisms associated with the development of this arrhythmia, the demonstrated ability to perform far more realistic simulations at high resolution over longer periods of time provides the potential to elucidate such detailed mechanisms that previously were out of reach.

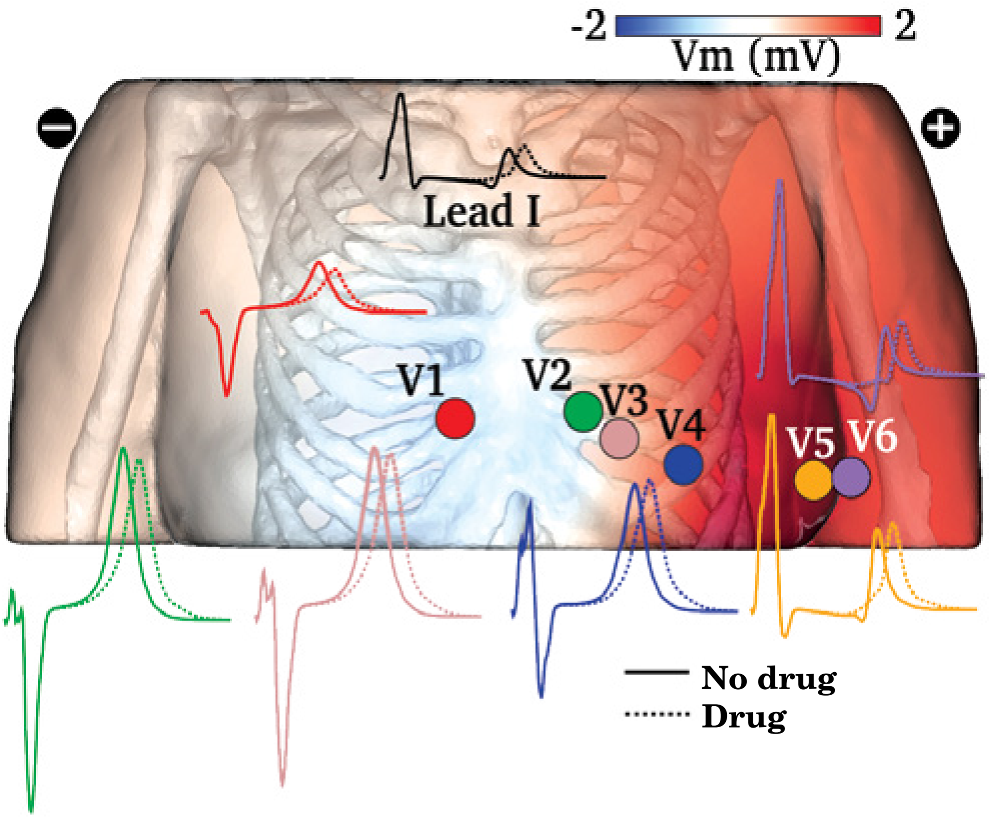

Electrocardiograms are simulated with and without a drug known to be arrhythmogenic. The case with the drug shows a delayed T-wave, the dashed-line wave on the right, which mimics the behavior observed in clinical records.

Exploiting Compiler Techniques to Address Exascale Challenges

ROSE is an open source compiler project at LLNL that is contributed to and used by research groups worldwide. Traditional compilers take source code written by humans (typically) and translate it to code that can be executed by machines. In contrast, ROSE is unique in that it is a source-to-source compiler addressing the custom requirements of application groups to build automated tools that operate directly on source code and output human-readable source code. A vendor compiler is then used on the ROSE-generated source code to create the machine code. This approach makes ROSE portable and more broadly targeted to custom or application-specific optimizations. Currently, ROSE supports the languages used in Department of Energy (DOE) applications, specifically C, C++, Fortran (77, 90/95, 2003, and 2008), OpenMP, Unified Parallel C, Compute Unified Device Architecture (CUDA), Open Computing Language (OpenCL), and Python. The LLNL team has also added Java language support and the capabilities to analyze binary executables.

One of the goals of the ROSE project is to support the research and development of programming models, domain-specific analyses, transformations, and optimizations for scientific applications running on current and future extreme-scale computers within DOE. Custom transformations (e.g., optimizations) can be automated on large-scale source code to focus on customized uses of hardware features or optimizations that are high level or specific to a particular application or parallel programming model. Since its inception more than 14 years ago, ROSE has continuously gained popularity, and research groups internal and external to DOE are still finding new ways to use it to meet their compiler needs.

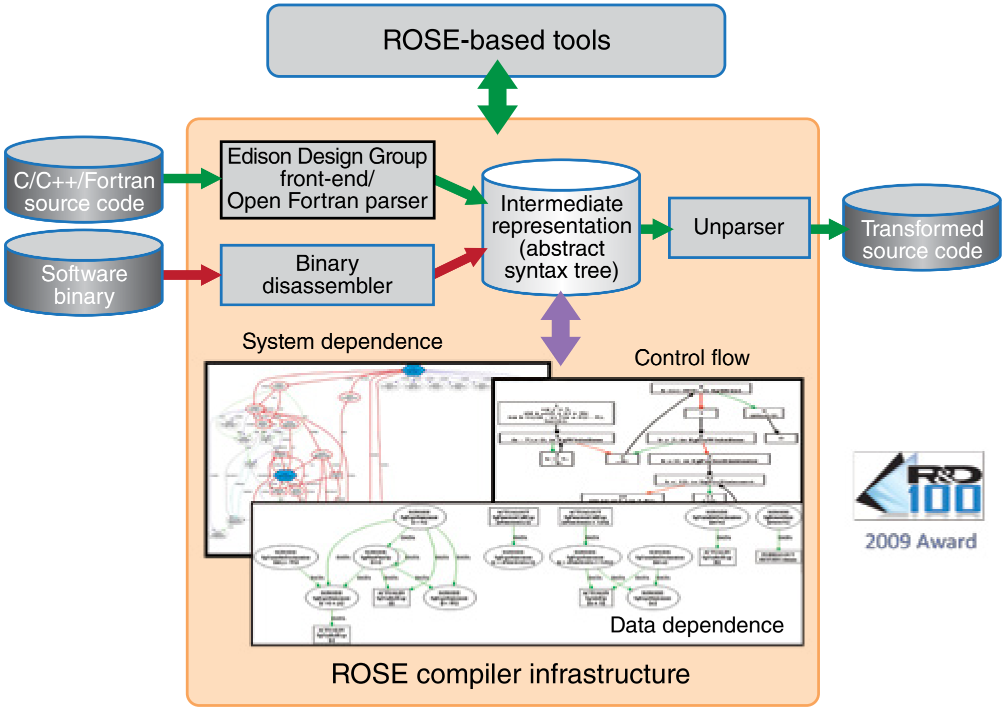

ROSE defines an open source infrastructure for building custom tools for arbitrary analysis and transformation of source code. Analyses include searching for parallelism and performance problems, and finding bugs. Transformations include instrumentation, optimizations, and generation of entirely new codes with parts specific to graphics processing units.

Progress in 2012

In 2012, we focused on 1) improving the support of large-scale C++ and Fortran 90 applications within DOE (e.g., C11/C++11, Fortran 2003, and Fortran 2008 support was added); 2) building a generic data-flow framework to support common compiler analysis such as constant propagation and liveness analysis; 3) exploring customized analysis and transformations to address exascale challenges including heterogeneity, power efficiency, and resilience; 4) adding CUDA and OpenCL language support to ROSE for graphics processing unit programming; 5) researching parallel programming models (e.g., MPI and OpenMP) and novel extensions for new architectures. As a result, our team published five research papers/posters on the topics of concurrency errors, OpenMP tasks, power optimizations, resilience, and MPI translation. We have also been working with collaborators to expand the use of ROSE to address a wider range of source code analysis and optimization problems. For example, we authored and co-authored 18 proposals that were submitted to DOE and the Department of Defense. As a result, LLNL now leads the successfully funded D-TEC (Domain Specific Language Support for Exascale) project with a focus on domain specific languages to simplify high performance computing (HPC) software development on exascale architectures. D-TEC is part of DOE's X-Stack program.

Many laboratories and universities are using ROSE to prepare for and analyze evolving computing requirements. For example, exascale architecture design teams at Argonne National Laboratory

DAN QUINLAN,
DQUINLAN@LLNL.GOV

CHUNHUA LIAO,
LIAO6@LLNL.GOV

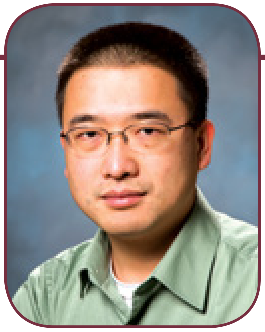

and LLNL used ROSE to characterize the behavior of their laboratory codes on their planned exascale systems. This approach to building new computing architectures, where both the hardware and the software are studied to understand how to build improved software and hardware recursively, is called co-design. Lawrence Berkeley National Laboratory, Sandia National Laboratories, University of Illinois at Urbana-Champaign, and UC San Diego used ROSE to support parameterized processor simulators within hardware/software co-design to define custom analysis and transformations. Their goal was to exploit novel exascale architecture designs (e.g., specialized instructions for scientific applications, active control of power usage by applications, and compiler optimizations). Active power management and design will be critical in DOE's exascale computer architectures. These architectures will need to be more power efficient (fewer watts per floating-point operations) to be feasible to build and maintain. LLNL's work with ROSE permits early experience with future design features to explore both the hardware features and how to write the software and/or design the compiler optimizations to use the hardware features; this is the essence of hardware/software co-design. Future systems architectures might also be heterogeneous.

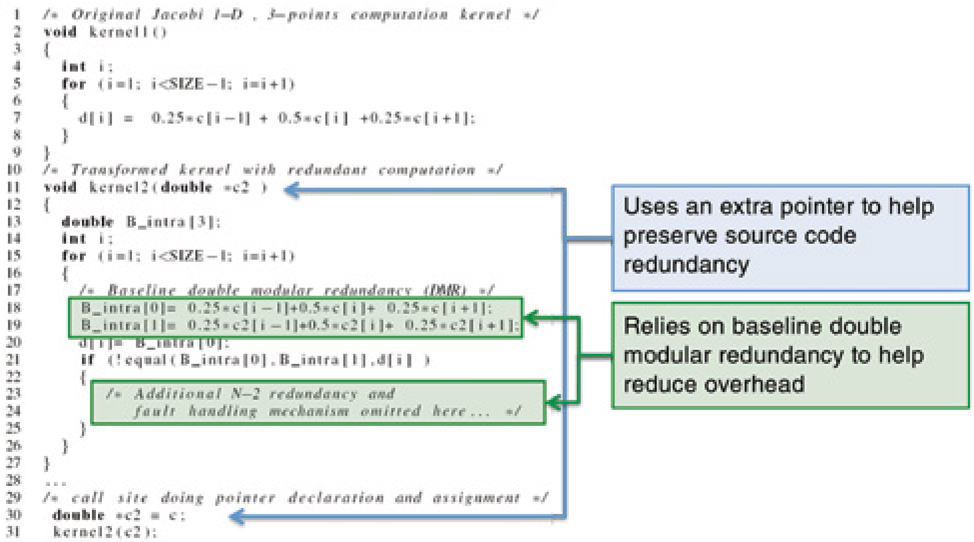

```
/* Original Jacobi 1-D , 3-points computation kernel */
void kernel1()
{
  int i;
  for (i=1; i<SIZE-1; i=i+1)
  {
    d[i] =  0.25*c[i-1] + 0.5*c[i] +0.25*c[i+1];
  }
}
/* Transformed kernel with redundant computation */
void kernel2(double *c2 )
{
  double B_intra[3];
  int i;
  for (i=1; i<SIZE-1; i=i+1)
  {
    /* Baseline double modular redundancy (DMR) */
    B_intra[0]= 0.25*c[i-1]+0.5*c[i]+ 0.25*c[i+1];
    B_intra[1]= 0.25*c2[i-1]+0.5*c2[i]+ 0.25*c2[i+1];
    d[i]= B_intra[0];
    if (!equal(B_intra[0],B_intra[1],d[i] )
    {
      /* Additional N-2 redundancy and
         fault handling mechanism omitted here... */
    }
  }
}
...
/* call site doing pointer declaration and assignment */
 double *c2 = c;
 kernel2(c2);
```

Uses an extra pointer to help preserve source code redundancy

Relies on baseline double modular redundancy to help reduce overhead

The ROSE team explored an idea of compiler-based triple modular redundancy, which can automatically harden critical code portions' execution. By carefully crafting the transformation (above), we significantly lowered the overhead of redundant computation while maintaining sufficient fault tolerance coverage. Our novel code duplication technique can also survive the aggressive compiler optimizations, which are designed to find and eliminate redundant computation.

A team at UC San Diego used ROSE to explore a directive-based programming model (MINT) to migrate and optimize stencil computations to use the CUDA programming model. Leveraging ROSE's source-to-source translation capability, they built a C-to-CUDA translator that achieved 78% to 83% of the performance obtained by aggressively hand-optimized CUDA. Similarly, the same group is using ROSE to explore MPI message optimizations. The MINT project is now released as part of ROSE and is an example of ROSE becoming an HPC community project. Another example of a project released via ROSE is the CodeThorn project, which is an LLNL collaboration with a colleague in Vienna, Austria.

The optimization of applications is critical within DOE. The ROSE project team, in conjunction with researchers from the DOE's Institute for Sustained Performance, Energy, and Resiliency, has dedicated significant effort to optimizing performance using a technique called auto-tuning. We have specifically focused on whole-program auto-tuning to support large-scale DOE applications. In auto-tuning, many different versions of parameterized optimizations are implemented and tested off-line, and the best version of the optimization is selected for the final compilation of an application. Auto-tuning permits an organized approach to the optimization of applications where internal compiler cost-modeling is intractable. This technique is especially useful when the performance is sensitive to minor perturbations in floating-point pipeline length, cache-line size, memory access stride, or any other case where the architecture is complex and not easily modeled within the compiler. Unfortunately, computer processor architectures are growing in complexity, and it is possible that auto-tuning will not become simpler in exascale architectures. Exascale architectures could be routinely optimized using this technology if researchers can make it more efficient for large-scale applications.

VAN EMDEN HENSON, HENSON5@LLNL.GOV

Novel Multilevel Approaches Enable Efficient Analysis of Large-Scale Graph Data

Graphs are highly useful tools whenever entities, or graph vertices, are connected by a relationship, or graph edge. For example, simple {vertices : edge} relationships could be {people : friendship}, {computers : network connection}, {intersections : road}. The graphs can be represented by matrices, and then spectral processing can be used to analyze the graphs. Spectral analysis has diverse applications from signal processing to medical imaging to the solution of partial differential equations. But the graphs can grow to extraordinary size ($O(10^{10})$, e.g., the Web, Facebook, or Twitter), making spectral processing problematic, mostly because the fundamental operation of matrix-vector multiplication is highly unscalable for extremely large power-law graphs. In 2012, the Eigensolvers project team delivered theory and tools for spectral analysis of such large-scale graphs, especially in three major applications areas: 1) partitioning, community-finding, and clustering, 2) distances and paths, 3) vertex importance and ranking.

Progress in 2012

Computing eigenvalues and eigenvectors of graph matrices is difficult and expensive even for small matrices, and while eigenpairs methods have been perfected for extremely large-scale matrices in applied physics problems, classical algorithms are unscalable and can fail altogether when applied to the power-law scale-free graphs that arise in network and data analysis. The Eigensolvers project was established in 2010 by the Laboratory Directed Research and Development (LDRD) Program to invent scalable methods for computing, using the eigenpairs for these graphs. An early discovery was that the matrix-vector product (or matvec), the fundamental computation of every linear algebra algorithm, was largely the culprit in creating the unscalability due to extremely poor interprocessor communications requirements imposed by the data. Special storage schemes for data have long been used to minimize communication; even with their use conventional matvecs scale as p^2 for these graph types, where p is the number of processors. We implemented a scheme that combined an unusual two-dimensional distribution of data to the processors with a novel ordering of the partial sums to produce a matvec that scales as $\sqrt{p}$ and achieves orders of magnitude speed-up over conventional matvecs for large graphs, greatly reducing the cost of the eigensolver for some problems.

A second approach taken in 2012 was to employ a vertex "disaggregation" technique we discovered in 2010 to speed up matvecs. This technique was tested on synthetic graphs with more than 200 million edges and showed good scalability on up to 25,600 processors of the Sequoia supercomputer. We combined disaggregation with a novel use of algebraic multigrid (AMG) as a preconditioner for a Locally Optimal Block Preconditioned Conjugate Gradient eigensolver. Conventional preconditioners are unscalable for these graphs, so we used an AMG preconditioner; these techniques have proven highly scalable on physics-based problems. However, even the use of a technique pioneered at LLNL known as "aggressive coarsening" is problematic, as it leads to dense coarse-level matrices for scale-free graphs, making the coarse-level matvecs prohibitively expensive. We devised a novel additive method, dubbed "+AMG," in which a correction from each coarse level is added rather than using the more conventional multiplicative correction. This approach avoids performing coarse-level matvecs with dense matrices. The +AMG method was tested on several real-world matrices, including Web data and two social network graphs, and it significantly outperformed both the Jacobi and BoomerAMG preconditioners on the test problems. Future work will focus on using these tools on clustering, identifying communities, and ranking problems.

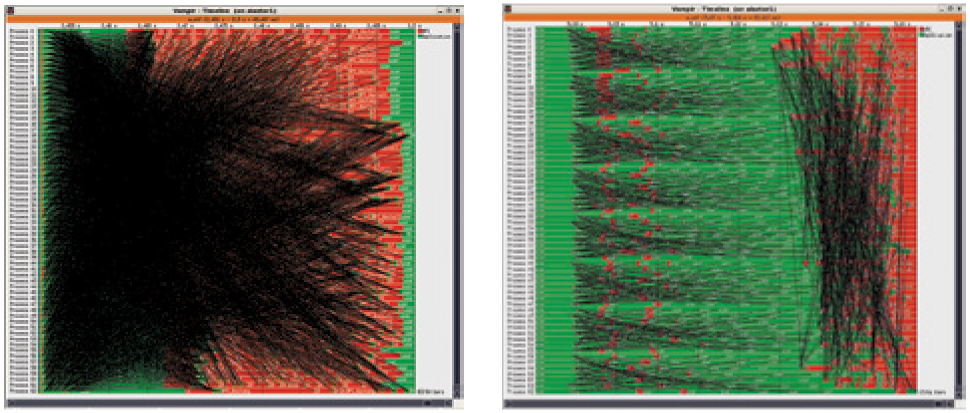

The Vampir™ diagrams show, for each processor, the time spent computing (green), waiting for communication (red), and communicating (black). Total time of the 1D (left) is compressed; it is many times that of the 2D (right).

PEER-TIMO BREMER, BREMER5@LLNL.GOV

Improving Airport Security through Advanced Computed Tomography

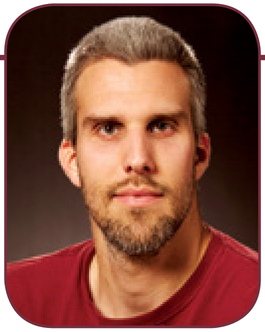

Computed tomography (CT) is a core technology in many applications crucial to national security, including stockpile stewardship activities and cargo inspection at airports. CT technologies can use x-rays, microwaves, or some other modality for imaging the inside of a sample or object. Our main challenge with CT is to find objects, materials, or other features of interest in potentially noisy and cluttered images corrupted by artifacts—rings, streaks, or other anomalies—and that have limited resolution. We are developing a new holistic processing pipeline to tightly integrate traditionally separated image-processing stages, including image segmentation and reconstruction. Our innovative techniques leverage Livermore's unique expertise and resources, enabling us to develop CT methods that reliably detect a broad range of features, such as physical threats in luggage, flaws in engineering components, or signs of aging in the nuclear stockpile.

Progress in 2012

LLNL is collaborating with partners in academia and industry to develop novel CT approaches for nondestructive testing and automated threat detection. Specifically, we are reworking the existing image-processing pipeline of acquisition, reconstruction, segmentation, and detection into an integrated and iterative process to produce more reliable and accurate imaging results.

Image information taken in the acquisition phase is typically ignored after reconstruction, but preliminary assessments have shown that this "discarded" data can provide important confidence information that is useful during segmentation. For example, the image below left is a simple test bag containing three metal rods. The rods are suspending a glass container that holds a smaller plastic bottle of water. The presence of metal produces major artifacts because it strongly attenuates the x-rays, and thus can shadow portions of the bag. The image below right illustrates this phenomenon. For each image pixel, the confidence is computed from the accumulated attenuation of all x rays contributing information to the pixel; the darker the shadow (or streak), the less confidence information is available. The image below left is reconstructed using LLNL's new processing technique; artifacts in the form of streaks can still be seen in both the water and the glass. These artifacts are difficult to compensate for in a stand-alone segmentation algorithm and can easily cause artificial merging or splitting of objects.

LLNL's technique uses confidence information to derive statistical measures of the likelihood that two neighboring pixels belong to the same object. The confidence information coupled with each individual measurement allows us to compensate for artifacts, such as streaks, by realizing that the corresponding differences in densities are not significant. As a result, segmentations will naturally integrate the uncertainty in the measurements and in the reconstruction by operating on the confidence values rather than the original image. This approach can potentially adapt to the high dynamic range of the data, making all scanned objects, including bags, appear more similar, and thus easier to process with a general segmentation algorithm. We are currently engaging industry partners, academia, and the Transportation Security Administration to integrate these concepts into a working system, evaluate their effectiveness in practice, and prepare a plan for quickly transitioning them into the field.

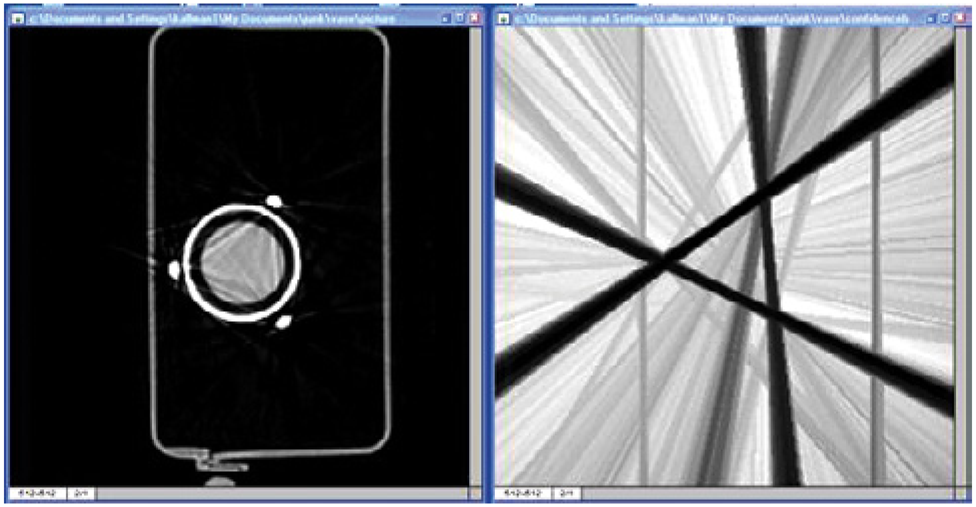

(Right) A simple test bag contains three metal rods, suspending a glass container that holds a smaller plastic bottle of water. (Far right) The metal produces major artifacts, thus shadowing other portions of the bag.

ROB NEELY, NEELY4@LLNL.GOV

Advanced Algorithm Research Develops Multiphysics for Exascale Architectures

Advanced Algorithm Technology for Exascale Multiphysics Simulation (AATEMPS) is an LDRD collaboration with the Weapons and Complex Integration Principal Directorate to develop the metrics, expertise, and insight required to transition LLNL's flagship multiphysics applications to the next generation of advanced computing architectures. The AATEMPS team consists of a broad cross section of software development researchers mirroring the multidisciplinary teams that make up the large Advanced Simulation and Computing (ASC) integrated code projects.

Within the next decade, the National Nuclear Security Administration (NNSA) plans to site at an ASC laboratory one or more advanced technology systems with a peak performance exceeding 1 exaFLOPS. This plan poses a significant challenge to system hardware designers because radical innovations will be required to achieve that speed within a 20-megawatt power budget—the limit at which operating costs of the machine make it feasible to deploy and operate. Equally (or perhaps more) challenging is the path ahead for application software developers, who must rethink how software is written to account for the features these new architectures will employ. As massively multicore processors and heterogeneous architectures emerge as likely candidates for the basis of exascale designs, the complexity foisted back onto the software developer significantly increases. In addition, the dominant performance driver is changing from computation to data motion. Obtaining high performance will require exploiting the many-fold increase in fine-grained concurrency, data locality effects of nonuniform memory access, and available SIMD and vector hardware. All of these characteristics will impact how the community thinks about algorithm design, and large multiphysics applications must find performance-portable solutions that provide a path to long-term productivity.

Progress in 2012

Arbitrary Lagrange–Eulerian (ALE) hydrodynamics and diffusion are two algorithms found in some of LLNL's multiphysics applications. The AATEMPS team chose to focus on these two packages because each possesses different algorithmic characteristics. ALE hydro includes irregular memory-access patterns and nearest-neighbor communication, while diffusion is characterized by more regular memory-access patterns and solution of a large sparse matrix involving both point-to-point and global communication. Also, both algorithms stress the memory bandwidth limits of emerging architectures. First, these two packages were incorporated into a research code called xALE to act as a surrogate multiphysics application. In addition, two mini-apps were developed for detailed studies: Mulard for radiation diffusion and LUAU3D for advection. The team also extensively used the LULESH proxy app, originally developed at LLNL for the Defense Advanced Research Projects Agency's Ubiquitous High Performance Computing Program, to study optimizations, data structure layouts, and programming models.

DOE has established a co-design process that encourages hardware and software designers to work together early in the design so that each side is fully informed of the trade-offs and compromises necessary to converge on a usable exascale system. AATEMPS is an important element of the LLNL ASC co-design strategy; the mini-apps developed for AATEMPS will subsist as representative examples of LLNL's software for use by the co-design community.

The AATEMPS project seeks to provide the ASC code teams facing exascale challenges with metrics, tools, and proven ideas that will set them on a path toward success. Although no one knows precisely what an exascale architecture will look like, we do know that the current way of doing business is insufficient and that we must radically rethink how we approach algorithm design and development.

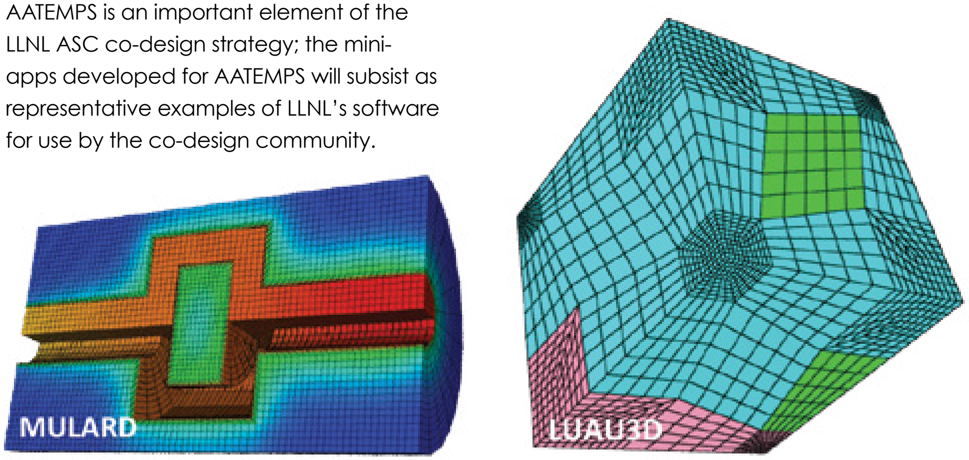

Multigroup **Advection**

New proxy applications are being developed within the Advanced Algorithm Technology for Exascale Multiphysics Simulation Project to help guide the transition of the Laboratory's multiphysics applications to the next generation of advanced computing architectures. The applications include a radiation diffusion mini-application called Mulard (left) and an unstructured mesh advection proxy called LUAU3D (right).

DOUGLAS EAST, EAST1@LLNL.GOV

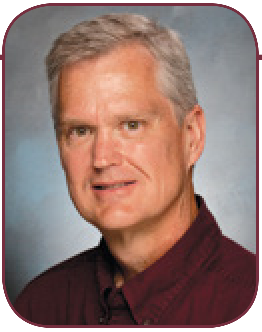

Vulcan: A New Resource for Industrial Collaboration

Vulcan is a prodigious unclassified computational resource that is the centerpiece of the High Performance Computing Innovation Center (HPCIC). Vulcan is an IBM Blue Gene/Q system with a peak performance of 5 petaFLOPS (one-quarter the size of Sequoia) and is tied with the fifth fastest supercomputer in the world on the November 2012 TOP500 list. Vulcan is shared between the tri-laboratory ASC Program and LLNL's Multiprogrammatic and Institutional Computing program, with 1.5 petaFLOPS available to industrial partners through the HPCIC in direct support of NNSA's mission to enhance U.S. economic competitiveness. Delivered in July 2012 and accepted in December 2012, Vulcan will be available for general laboratory and industrial use in May 2013.

The High Performance Computing Innovation Center makes available to industrial collaborators a prodigious resource in the 5-petaFLOPS Vulcan system as well as the unique breadth and depth of Livermore computational science and engineering expertise.

Progress in 2012

Vulcan is one of the largest computational resources available in the U.S. for industrial use through an innovative collaborative agreement that fully supports the development of proprietary technology. Industrial partners access Vulcan through the HPCIC, whose mission is to enhance American industrial competitiveness by broadening the adoption and application of HPC technology. Vulcan is not available solely as an on-demand computing resource but rather in support of collaborative development efforts involving computer scientists, domain scientists, and subject-matter experts from both LLNL and industrial partners.

Industrial use of the resource requires a nominal fee, which is based on a strict full-cost-recovery paradigm that was developed and approved for use in mid-2012. The fee establishes an allocation, which is a right to access a percentage of Vulcan for a period of time. The size of the allocation may be viewed, in effect, as a priority. The fee is a bundled rate that includes not only access to Vulcan but to the entire associated infrastructure available in the Livermore Computing HPC environment, including network access, file systems, archive storage, and help-line support. The FY13 rate for access to Vulcan is $3.55K per teraFLOP per year, which equates to just $0.0052 per CPU hour. This accounting mechanism is in sharp contrast to most typical usage schemes that charge directly for the number of cycles that are utilized by the project at a specified "dollars-per-CPU-hour" rate.

Another landmark development in 2012 was the establishment of the Deep Computing Solutions (DCS) collaboration with IBM as part of the HPCIC. Through DCS, industrial partners have available to them IBM expertise in the computational and domain sciences, particularly with respect to the optimization of applications for use on Vulcan. For an industrial partner, the availability of DCS significantly deepens the bench of available expertise that can be applied to their project and provides experience in the industrial hardening of research and development applications.

Since its establishment in 1952, LLNL has been at the forefront of supercomputing technologies and applications with a primary focus on the most demanding national security applications. Developing creative, actionable solutions for difficult problems is a Laboratory core competency. The HPCIC makes available to U.S. industry an incredibly capable resource in Vulcan and the unique breadth and depth of LLNL computational science and engineering expertise. Through the HPCIC, LLNL is poised to be a leader in helping industry take advantage of HPC technology to improve product design, quality, safety; shorten development lifecycles; reduce cost; and enable products of the imagination that will become tomorrow's innovations.

ALLAN CASEY, CASEY20@LLNL.GOV

NIF Archive Eases Access for the Scientific User Community

The National Ignition Facility (NIF) is the world's largest and most energetic laser, but more importantly it is a research center that allows scientists around the world to study inertial confinement fusion ignition and explore matter at extreme energy densities and pressures. Each experiment, or shot, potentially generates hundreds of gigabytes of data that then gets archived in a permanent repository. Given the vast amount of data contained in the NIF experimental results archive, it is imperative that scientists have the ability to interact with the archive visually, efficiently, and with tools that facilitate further experimental analysis. To this end, Computation's Shot Data Systems team improved the NIF archive by delivering interactive, Web-based dashboards, a virtual file system interface, a scripting language for programmatic data integration, and an experimental wiki to support collaboration.

Progress in 2012

The NIF experimental results archive is an Oracle database with software interfaces written on top of a low-level Java content management application programming interface (API). Using API to define classes in the archive ensures the objects will always be created unambiguously, that they will remain in the correct structures for the class of data, and that they will contain enough metadata to describe the object through its attributes. This structure provides the basis on which the archive's new features have been built.

New experiment dashboards, called the Archive Viewer, allow scientific users to custom assemble interactive browser widgets, such as portlets, that render a set of objects from the NIF archive. The dashboard can display graphs of laser power, a camera image, and tables of laser energy using the widgets that are capable of displaying graphs, images, and tables. Because each widget is customized with the type and attributes of the object it can display, no coding is required to display the data. A user only needs to define the layout of the widgets on the dashboard and then assign the object to be displayed to the relevant widget. This approach to visualization has reduced the development time required to visualize data in the archive from several weeks to a few hours.

Also added to the archive are two new features that make it easier for a user to directly access archival data. First, the archive has implemented a version of the Web Distributed Authoring and Versioning (WebDAV) protocol that allows users to browse the files in the archive as if they were navigating them in a directory structure, either in a Web browser or as a drive on the client computer. Second, the addition of a Groovy/Velocity-based scripting language allows users to programmatically package data from the archive for later download and analysis. The scripts can be run as needed or set up to execute automatically, following an event such as postshot analysis. The hierarchical nature of the archive enforced by API facilitates these forms of access, and both features greatly simplify the task of extracting data from the database.

The archive is now also integrated with an experimental wiki, a feature that reflects and enhances the collaborative nature of research. When scientists notice something of significance in the archive, they can post a comment directly to the wiki that can then be read by the rest of their team in real time. This eliminates the need to have both wiki and archive sessions running concurrently.

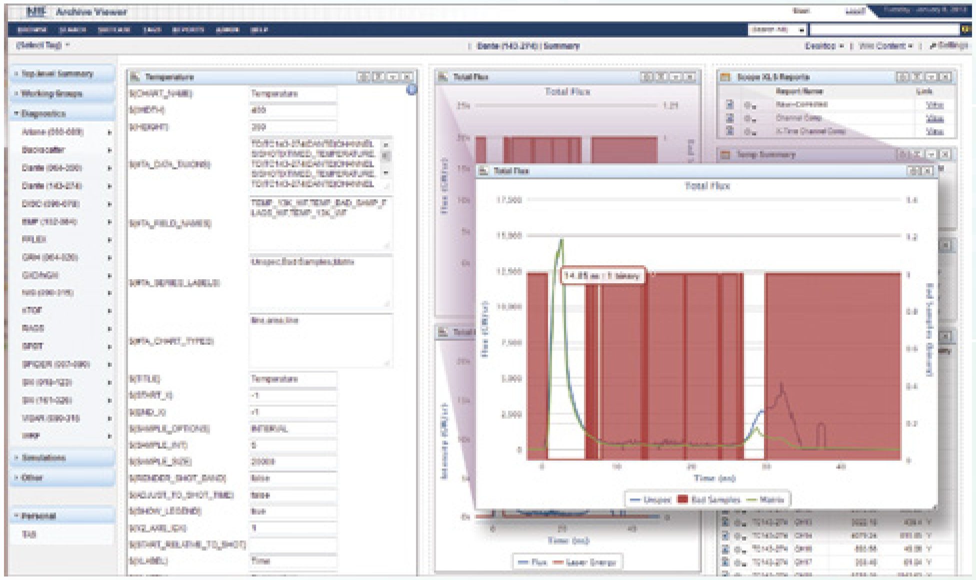

The National Ignition Facility's new Archive Viewer allows users to view data from experiments via customized portlets with data-driven configuration and interactivity.

DOUGLAS EAST, EAST1@LLNL.GOV

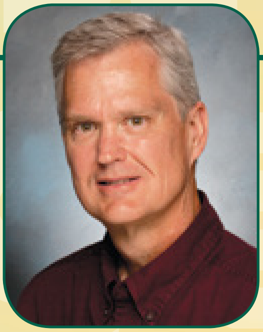

Institutional Services Pursue a Balance Between Security and Accessibility

For Lawrence Livermore National Laboratory (LLNL) to attract and retain the best workforce and realize productivity gains that maintain our position at the forefront of science, it is imperative that we continually improve our information technology (IT) services. Paramount to this evolution is enhancing user mobility and unifying the delivery of voice, video, and data services in a seamless and secure fashion.

Mobility frees employees from the confines of their offices and allows for more vibrant real-time interaction. We envision an environment where employees can fully utilize their laptops or tablet devices anywhere on site, at home, or on travel, and be able to easily and securely access necessary information just as if they are sitting in their offices. As such, software applications will not be limited to a desktop-sized screen but will be compatible with a range of mobile devices, including smartphones, tablets, and laptops. Work–life balance will be enhanced by employees integrating their workday and their personal-life with a minimum number of portable electronic devices to the extent they desire. Likewise, communication will be enhanced since employees will be able to contact others or be contacted no matter where they are by telephone, video, or other communication methods by a single identifier or username. As envisioned, these improvements will be easily utilized by the employee, be robust, and maintain the security and integrity of LLNL data.

We know employees can easily handle many IT tasks without the need to call on IT staff. For instance, at home we routinely reset forgotten passwords and install newly acquired software. Why not have the ability to do so in the workplace as well? In the necessary quest to secure our systems and networks, the ability of employees to solve issues themselves is often sacrificed. Unfortunately, these actions lead to delays for the resolution of otherwise simple issues because there is not enough IT staff to go around. We look forward to the continued expansion at LLNL of self-help (or Tier-0) support, enabling users to resolve their own problems without IT staff intervention and without compromising the security of the device or network.

LLNL has a long way to go in the fulfillment of this vision. We are lagging behind private industry in many ways but are also ahead of our contemporary laboratories in other ways. We face constrained budgets, and an anchor of existing policy, practice, and perspective sometimes further delays our progress. However, even in the face of these challenges, there has been notable progress over the past year. The articles in this section describe some of this progress, achieved by Computation staff under the aegis of the CIO Program.

Empowering Users with Tier 0 Support

LLNL offers different levels of computer support services depending on users' specific needs. Tier 1 support provides phone assistance for general user issues. Field services are available through Tier 2 support for those having hardware problems or other issues requiring hands-on attention. Tier 3 support provides centralized enterprise-wide services, including tools, encryption, patching, and software lifecycle management. In 2012, LLNL established Tier 0, which gives users various online tools for solving their computer problems without needing to engage LLNL's Institutional Service Desk (Tier 1) or Computer Support Units (Tier 2).

Tier 0 allows users to resolve their common computer issues more quickly. The new level of support also frees up Tier 1 and Tier 2 resources for addressing more complex problems. With Tier 0, users can check the backup status of their computers, remediate and unblock their systems, control when system management tasks are run, and look up answers to IT questions in an online archive. If users require more immediate assistance, they can receive face-to-face guidance from on-duty technical staff at our new service center.

Progress in 2012

LLNL's Tier 0 support incorporates several features for users to better manage their computer systems. BLAST (Backup Log Aggregation Server Tool) permits users to see the last time their computer systems were successfully backed up. By simply typing BLAST into a Web browser, users can view a dashboard with the backup status of all active computer systems registered in their names. Users can then contact local computer support to investigate any potential backup issues that could result in data loss if a system failure were to occur. BLAST also provides IT managers and system administrators with different dashboards and tools to monitor the overall status of backup operations in their areas.

The Laboratory's Blue Coat Blocking (BCB) institutional cybersecurity monitoring system prevents or "blocks" computers that are not compliant with LLNL cybersecurity rules from using the Internet or internal networks. In the past, users had to contact Tier 1 support to unblock their system. The new BCB Assistant—an applet installed locally on a user's computer—identifies security issues and provides a "fix now" option that can remediate most problems and remove the block. BCB Assistant is available on Windows, Macintosh, and Linux systems and is currently installed on thousands of computers across the Laboratory.

To maintain system security, LLNL performs regular maintenance on its computers, in part, through automated tasks that run periodically on individual systems. These tasks include multiple security and inventory scans as well as operating and application patches. Unfortunately, scans and patches require system resources and can run at inconvenient times for users. The Laboratory's novel Maintenance Manager permits Windows 7 users to control and schedule the various management tasks so they do not adversely affect system performance or user productivity.

Through Maintenance Manager, users can construct "maintenance plans" for their systems in which they list and order the required tasks and dictate when they run. System management tools are then constrained to operate according to the maintenance plans. Additionally, system administrators may develop a basic maintenance plan, or template, for systems under their purview. Users may or may not be free to edit the content of this plan, but they can select the time frame during

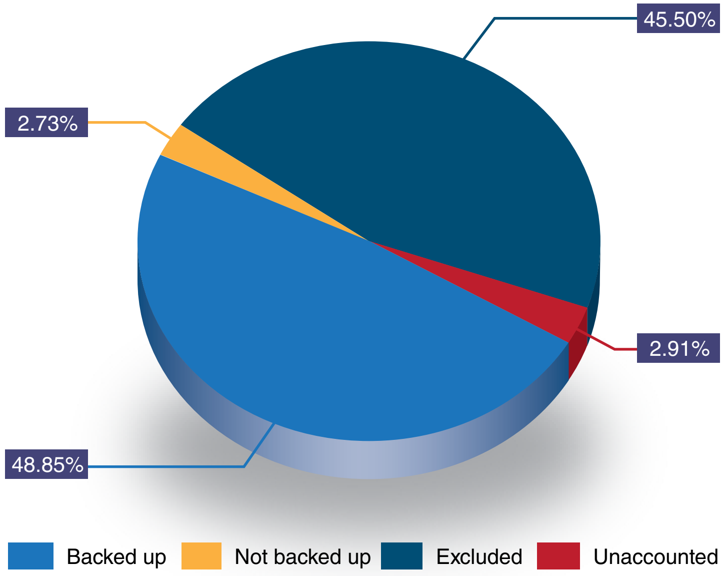

BLAST provides system administrators with a Web-based console featuring the tools and metrics they need to quickly and easily monitor the backup operations of the systems in their domains. This chart shows backup statistics for systems in the Director's Office.

BRETT CORNELL, BCORNELL@LLNL.GOV

which the tasks would be executed on their systems. A limited set of users is piloting Maintenance Manager before the tool is made available across the Laboratory.

Another feature of Tier 0 support is RightAnswers, a commercial IT knowledge base provided by the Laboratory's Institutional Service Desk to assist computer users in quickly finding answers to their IT questions without having to contact support services. RightAnswers is a vast central repository of IT information that includes specific articles that address various system issues. The RightAnswers knowledge base contains more than 75,000 IT solutions with easy-to-follow steps. Most of these solutions are provided by the vendor, but LLNL IT system administrators and technicians have written more than 1,600 articles on topics that pertain to the Laboratory. If users cannot find a solution to a problem, they can submit support requests from within the application to the Institutional Service Desk. During 2012, an average of 450 people per month used RightAnswers.

Tier 0 support also incorporates a novel approach to user customer service called BRING-IT (Build, Repair, Image, Next-Generation IT). BRING-IT is a computer systems support service center staffed by senior technicians who provide a broad range of technical services. Users can schedule an appointment via phone, e-mail, or the Web, or simply drop in at the center for help. Although aimed primarily at assisting users with mobile systems (e.g., laptops), the BRING-IT staff can help with most computer-related problems. The most common incidents include encryption issues, hardware repair requests, system upgrades, and computer trade-ins.

The BRING-IT center staff work with users face-to-face in the customer service area to resolve their IT issues. If a BRING-IT staff member is unable to resolve an issue directly, the staff member acts as the user's advocate and coordinates additional support with Tier 2 and Tier 3 resources, as needed, until the issue is resolved.

The BRING-IT service center offers customers a place to receive in-person, immediate assistance with various computer issues.

MATTHEW MYRICK, MYRICK3@LLNL.GOV

Master Block List Protects Against Cyberthreats

In an increasingly technical world, people have the ability to communicate and instantly connect with others at any time. Mobile devices and social networking have become the norm rather than the exception for both business and personal interactions. However, the data that is shared and stored online as a result of our virtual interconnectedness is also a target for cyber criminals who seek to exploit technological and human vulnerabilities for personal gain.

To combat growing cyber threats, LLNL's Cybersecurity Program created the Master Block List (MBL) in collaboration with the LLNL-led Department of Energy (DOE) Focused Advanced Persistent Threat Group. MBL allows DOE-affiliated laboratories and plants to share cyber threat information in real time and strengthen the cyber posture of the entire DOE complex. Currently, 11 DOE institutions use MBL, and the number continues to rise. In recognition of the unique approach to cybersecurity and MBL's success, the project was awarded a SysAdmin, Audit, and Network Security National Cybersecurity Innovation Award on November 5, 2012.

Progress in 2012

Initially, MBL was created for LLNL to easily and automatically share malicious Uniform Resource Locators (URLs) and Internet Protocol (IP) addresses with Sandia National Laboratories. The steady increase in cyber attacks had made it difficult to share cyber threat data among trusted partners using traditional methods, prompting a need for a new tool that could restore security between the laboratories. The working prototype between LLNL and Sandia quickly garnered the attention of other laboratories and evolved into a much larger framework for sharing complex-wide. MBL originally aggregated malicious URLs and IPs into a single list. Today, MBL has expanded to include three lists that are also used to share MD5 hashes—cryptographic checksums for fingerprinting malicious files—and e-mail addresses or e-mail server IPs for sharing spear phish details. Spear phishing is the preferred mechanism of skilled, persistent attackers.

The heart of MBL leverages a lightweight and agile custom threat-sharing protocol that was developed based on command and control techniques LLNL honed after dealing with its persistent adversaries. This protocol allows virtually any application to be easily modified to automatically share threat intelligence despite the disparate tools and sophisticated networks used throughout the DOE complex. While other sharing models within DOE use complicated XML schemas and structures, MBL's success is largely attributed to its simplicity.

MBL makes it extremely easy to share data. To submit data, DOE complex personnel can use almost any network tool at their disposal to perform "GET" requests in a pre-specified format over HTTPS to a Web server at LLNL. Data is retrieved in a similar manner, where users perform "GET" requests over HTTPS for statically named lists containing easily separated values. By leaving the mechanism for data submission and retrieval open to the interests of the participating party, MBL facilitates sharing.

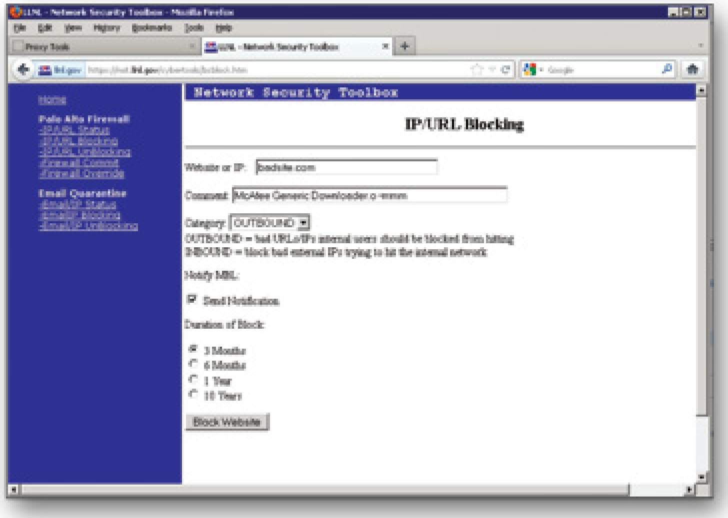

Network Security Toolbox

IP/URL Blocking

Website or IP: badsite.com

Comment: McAfee Generic Downloader.o-mmm

Category: OUTBOUND

OUTBOUND = bad URLs/IPs internal users should be blocked from hitting
INBOUND = block bad external IPs trying to hit the internal network

Notify MBL:

☑ Send Notification

Duration of Block:

● 3 Months
○ 6 Months
○ 1 Year
○ 10 Years

The Master Block List is a Department of Energy-wide service and data aggregations tool that allows organization to share in real time domain names that are known or suspected to be untrustworthy. Cyber experts at any of the 11 DOE partner institutions can log into the database, enter the information pertaining to the malicious source, and notify the other DOE entities of the block.

MBL will continue to evolve, but its focus will remain on providing a mechanism where threat data can be easily and automatically shared among DOE institutions. The MBL framework continues to expand within the complex as more DOE laboratories and plants share data. Other government agencies have also expressed interest in participating in the framework. Recently, MBL was implemented to facilitate LLNL sharing cyber threat information with private industry partners participating in the Bay Area Advanced Persistent Threat Special Interest Group.

ROBYNE TESLICH, TESLICH2@LLNL.GOV

LLNL Leads the Mobility Charge

Mobile computing, such as that offered by smartphones and tablets, is changing the way people access information. The ability to send and receive data from virtually anywhere at anytime is driving technological changes in the workplace and in our personal lives. LLNL is entering a new paradigm in information security, providing employees with the technology they want and need to facilitate their work activities while improving how the data that is shared, stored, and viewed on these devices is protected.

As part of a multiyear effort, LLNL's CIO organization has delivered mobile solutions to employees by expanding the wireless network, providing employees with government-owned BlackBerry phones, and more recently offering support for iOS devices such as iPhones and iPads. Policies and guidelines have also been updated to enable employees to use personal devices on-site and for accessing LLNL resources, including e-mail.

LLNL is leading the mobile computing transition within the National Nuclear Security Administration (NNSA) complex. Others laboratories have started leveraging LLNL policies and approaches to deliver similar capabilities to their employees. Additionally, LLNL CIO staff are lead authors on an NNSA whitepaper on mobility management for the complex.

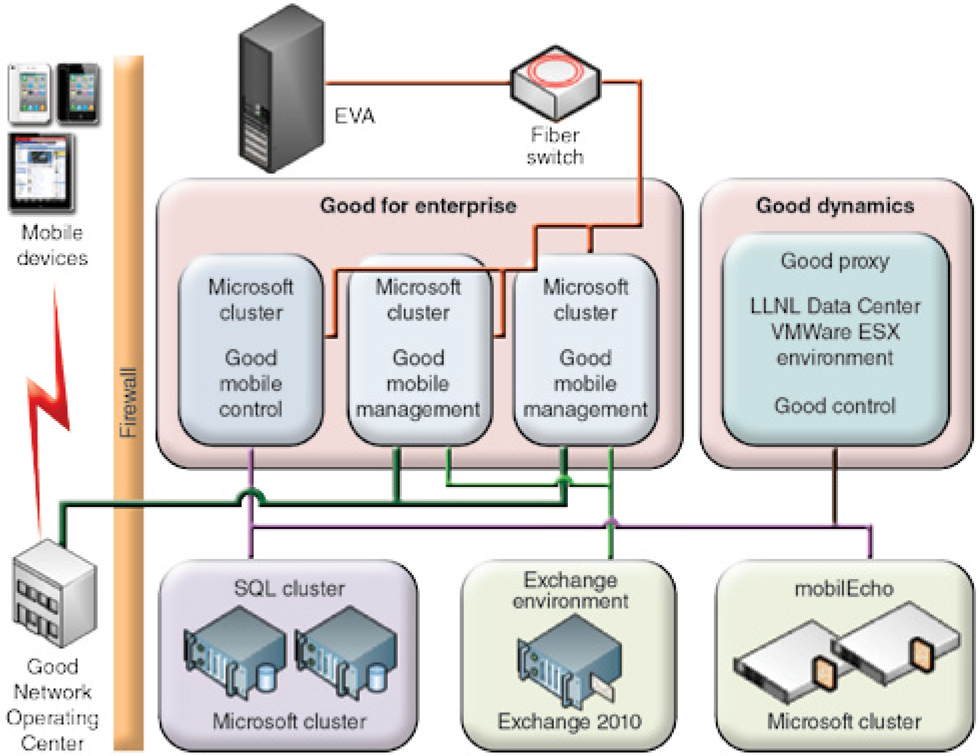

This schematic illustrates the Good for Enterprise infrastructure implemented at LLNL for mobile device management.

Progress in 2012

This year, LLNL took significant steps to expand mobile computing capabilities throughout the Laboratory. As an example, senior managers and employees recognized the iPad as a potentially valuable tool for documenting work tasks and accessing needed information while away from physical workstations. However, information security could not be sacrificed for the sake of mobility. As a result, the CIO organization developed a service to support the technology using a risk-based approach to mobile device management that minimized potential data loss. LLNL selected Good Technology's "Good for Enterprise" product as the Mobile Device Management (MDM) solution. MDM provides NNSA-compliant encryption for sensitive information on iPads and enables basic management capabilities including password enforcement. The market-leading product supports iOS and Android devices. MDM was initially deployed on LLNL-owned devices only.

The Laboratory also amended policies and procedures to allow personally and LLNL-owned devices in Property Protected and Limited Areas. LLNL was the first laboratory in the NNSA complex to implement such broad sweeping changes to policy.

A pilot program called Bring Your Own Device (BYOD) will launch in January 2013 to test the feasibility of using personally owned devices for LLNL business. The BYOD Program also addressed the many legal challenges related to information security as well as potential litigation and contamination issues. As part of the program, employees are required to accept all costs associated with using their own device. They must also agree to certain operating requirements and permit legal access to their devices.

LLNL continues to explore solutions and options for developing and supporting mobile applications and for restructuring how data is stored, accessed, and protected in the new mobile world. As the mobile computing landscape evolves, LLNL will strive to leverage current technologies and integrate them into the existing network infrastructure, helping to maximize employee productivity and creativity.

Improving InfiniBand Effectiveness through Best Practices

JIM SILVA, SILVA50@LLNL.GOV

InfiniBand (IB) has rapidly gained popularity within the high performance computing community. The low-latency, high-bandwidth, inexpensive hardware, and open standard of IB have made it attractive for industry and academia. For the last five years, LLNL has deployed many IB fabrics including interconnects within our commodity clusters, and subsequently, Storage Area Networks (SANs). However, IB remains a relatively "unfamiliar" feature to the Laboratory's computing center and presents major challenges. To resolve these issues, LLNL works closely with IB hardware vendors and focuses efforts on developing a common software stack. Commodity solutions help LLNL to use similar processes for configuring and monitoring each of our fabrics regardless of vendors and use cases.

Progress in 2012

The IB component set encompasses both hardware and software. LLNL does not produce hardware but maintains close relationships with hardware vendors. In addition, our system administrators test new hardware during preproduction. We provide feedback from large deployments to vendors during procurements and through our membership in the InfiniBand Trade Association and OpenFabrics Alliance (OFA). This feedback is especially useful to vendors when dealing with issues of scale.

In 2012, LLNL deployed a new fabric topology for the Sequoia SAN. Overall, the topology is functioning well, providing approximately 850 gigabytes per second of data transfer.

The IB specification has also allowed greater flexibility and better network performance with less expensive hardware. However, these advances also rely on more complex software. LLNL partnered with the IB industry through OFA to use the alliance's OpenFabrics Enterprise Distribution (OFED)—an open IB software stack. LLNL no longer uses OFED directly, but Red Hat (our base Linux distributor) draws directly from this work. Additionally, LLNL integrates its software changes into the OpenFabrics stack to leverage the testing and maintenance provided by the community.

The main component of IB management software is the subnet manager. LLNL uses an open-source subnet manager called OpenSM as a platform to leverage the latest in routing algorithms. It also allows for rapidly modifying source to fix critical production issues before vendors release commercial software, and provides a single common interface to all of our fabrics for configuration. Through OpenSM and additional diagnostic utilities, our system administrators can monitor physical connections to ensure a consistent topology, periodically check and record hardware error counters, and remedy link errors swiftly to prevent unnecessary burden on the subnet manager.

During 2012, LLNL also added three major features to our IB management software. Our new master configuration file allows management software to compare the current operating fabric with the desired fabric as specified in the file. By modifying only a few command line tools, the master file flags improper connections, invalid speeds, and missing nodes on a fabric. We also integrated congestion control configuration management to OpenSM, which helps slow ingress of flows that share congested links. Research shows congestion control can substantially ease fabric traffic and results in higher overall throughput. However, further analysis of configuration parameters and workloads is needed before this feature can function as expected. Additionally, LLNL implemented a plug-in to OpenSM that allows unprecedented access to the data available within the subnet manager, providing a wealth of previously inaccessible data. The plug-in was first deployed to display link errors on the SKUMMEE cluster-monitoring tool.

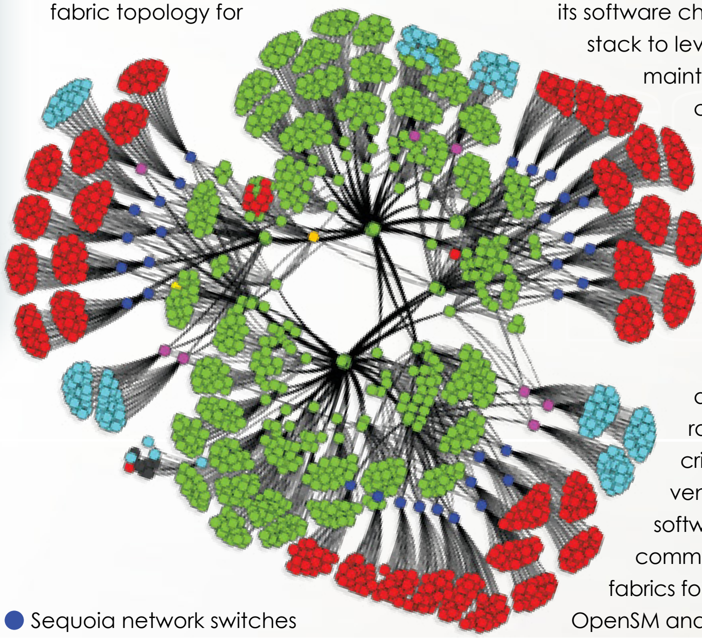

The layout of Sequoia's Lustre InfiniBand (IB) fabric is visualized using the Fruchterman–Reingold algorithm. The circles represent nodes on the IB fabric, and the lines depict network connectivity.

● Sequoia ● Vulcan ● Grove (file system) ● Sequoia network switches
● Vulcan network switches ● Lustre gateways ● NFS gateways

Appendices

Academic Outreach

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
Arizona State University	Stephen Johnson	Joint research	Peptide microarrays	DHS	Tom Slezak
Brigham Young University	Bryan Morse	Joint research	Mosaics and super-resolution of unmanned aerial vehicle-acquired video using locally adaptive warping	LDRD	Mark Duchaineau and Jon Cohen
California Institute of Technology	Michael Ortiz	ASC Predictive Science Academic Alliance Program Center	Center for the Predictive Modeling and Simulation of High-Energy-Density Dynamic Response of Materials	ASC	Dick Watson
California Polytechnic State University, San Luis Obispo	Ignatios Vakalis	Joint research	Cybersecurity research; joint proposals	SMS	Celeste Matarazzo
Cambridge University	Nikos Nikiforkis	Joint research	Simulation and modeling using Overture	ASCR Base	Bill Henshaw
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
Chalmers University of Technology, Sweden	Sally McKee	Collaboration	Leveraging OpenAnalysis for alias analysis within ROSE	ASC	Dan Quinlan
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	Donald Estep	Subcontract	Posteriori error analysis for hydrodynamic systems	LDRD	Carol Woodward
Colorado State University	Stephen Guzik	Subcontract	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
Columbia University	Mark Adams	Joint research	FASTMath: Frameworks, Algorithms, and Scalable Technologies for Mathematics	ASCR SciDAC	Lori Diachin
Columbia University	Ian Lipkin	Joint research	Viral discovery and microarrays	DTRA	Tom Slezak
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

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
Dresden University of Technology	Wolfgang Nagel	Joint research	Semantic debugging of message-passing interface applications; trace-based performance analysis	ASCR, ASC	Bronis de Supinski
ETH Zürich	Thorsten Hoefler	Joint research	Message-passing interface forum and advanced message-passing interface usage, power modeling for Blue Gene/Q	ASC	Martin Schulz
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	Richard Fujimoto	Subcontract	Research in reverse computation	LDRD	David Jefferson
Georgia Institute of Technology	Jarek Rossignac	Subcontract	Compact streamable mesh formats	ASCR	Peter Lindstrom
Georgia Institute of Technology	Richard Vuduc	Subcontract	Compiler support for reverse computation	ASCR	Dan Quinlan
Georgia Institute of Technology	Dan Campbell and Mark Richards	Collaboration	Data-intensive applications	WFO/DARPA	Maya Gokhale
Imperial College	Paul Kelly and José Gabriel de Figueiredo Coutinho	Collaboration	Field-programmable gate arrays research	ASCR	Dan Quinlan
Indiana University	Jeremiah Wilcock	Joint research	Binary analysis	ASCR	Dan Quinlan
Johns Hopkins University	Allan Boyles	Collaboration	Seismoacoustic modeling for defense-related efforts	DOE	Shawn Larsen
Karlsruhe Institute of Technology	Wolfgang Karl	Joint research	Hardware transactional memory	ASC	Martin Schulz
Louisiana State University	Lu Peng, Lide Duan, and Sui Chen	Joint research	Characterizing the propagation of soft faults through numeric applications	ASCR	Greg Bronevetsky
Ludwig Maximilian University of Munich	Dieter Kranzlmueller	Joint research	Detecting communication patterns to optimize applications	ASCR, ASC	Bronis de Supinski
Ludwig Maximilian University of Munich	Dieter Kranzlmueller	Joint research	Message-passing interface tool infrastructure and performance analysis	ASC	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

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
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	P. Sadayappan and Christophe Alias	Collaboration	Optimizing compiler program analysis	ASCR	Dan Quinlan
Ohio University	Yusu Wang	Joint research	Analysis and visualization of high-dimensional function	LDRD	Timo Bremer
Pennsylvania State University	Ludmil Zikatanov	Subcontract	Fast solvers for discrete Hodge–Laplacians	LDRD	Van Emden Henson
Pennsylvania State University	Ludmil Zikatonov	Subcontract	Multilevel methods for graph Laplacians and piece-wise constant approximations	ASCR Base	Panayot Vassilevski
Pennsylvania State University	Jinchao Xu and James Brannick	Subcontract	Multigrid methods for systems of partial differential equations	ASCR	Robert Falgout
Polytechnic University of Puerto Rico	Alfredo Cruz	Joint research	Cybersecurity research; joint proposals	—	Celeste Matarazzo
Princeton University	Adam Burrows	Joint research	Computational Astrophysics Consortium	ASCR SciDAC	Louis Howell
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	Saurabh Bagchi	Joint research	Statistical debugging tools, fault tolerance, scalable checkpointing	ASCR, ASC, LDRD	Bronis de Supinski
Purdue University	Ziqiang Cai	Summer faculty	A posteriori error estimates for partial differential equations	ASC	Robert Falgout
Purdue University	Jayathi Murthy	ASC Predictive Science Academic Alliance Program Center	Center for Prediction of Reliability, Integrity, and Survivability of Microsystems (PRISM)	ASC	Dick Watson
Purdue University	Jennifer Neville	Joint research	Modeling behavior in dynamic networks	LDRD, SMS	Brian Gallagher
Purdue University	Mithuna Thottethodi	Joint research	Optimized node mapping techniques	ASC	Martin Schulz
Purdue University	Dongbin Xiu	Subcontract	Intrusive and nonintrusive polynomial chaos methods for the quantification of uncertainties in multiphysics models	ASC	Charles Tong

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
Queens University Belfast and Virginia Polytechnic Institute and State University	Madhav Marathe and Dimitris Nikolopoulos	Joint research	Power optimization for hybrid codes	ASC	Martin Schulz
Rensselaer Polytechnic Institute	Don Schwendeman	Subcontract	Development of numerical methods for mathematical models of high-speed reactive and nonreactive flow	ASCR Base	Bill Henshaw
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 tools and tool infrastructures	ASC	Martin Schulz
Rice University	John Mellor-Crummey	Joint research	Performance tools, OpenMP tool interface, performance evaluation	ASCR	Bronis de Supinski
Rice University	Vivek Sarkar	Subcontract	Data abstractions for portable high performance computing performance	LDRD	James McGraw
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 Micheal Burke	Joint research	Development of a static single assignment-based dataflow compiler framework for ROSE	ASCR	Greg Bronevetsky
Rice University	Vivek Sarkar, Jisheng Zhao, Vincent Cave, and Micheal Burke	Subcontract	Message-passing interface producer–consumer program analyses for ROSE	ASCR	Greg Bronevetsky
Rochester Institute of Technology	Kara Maki	Collaboration	Droplet flows	—	Bill Henshaw
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
RWTH Aachen University	Matthias Mueller	Joint research	Message-passing interface correctness checking	ASC	Martin Schulz
RWTH Aachen University	Felix Wolf	Joint research	Message-passing interface performance analysis tools	ASCR, ASC	Bronis de Supinski
RWTH Aachen University and German Research School for Simulation Sciences	Felix Wolf	Joint research	Performance tools and tool infrastructures	ASC	Martin Schulz
Southern Methodist University	Thomas Hagstrom	Joint research	High-order structure grid methods for wave propagation on complex unbounded domains	ASCR Base	Bill Henshaw

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
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
Stanford University	Juan Alonso	Subcontract	Analysis of shear-layer/wake interaction for drag reduction of heavy vehicles	DOE	Kambiz Salari
Stanford University	Olav Lindtjorn	Collaboration	Reverse-time seismic imaging for hydrocarbon exploration	CRADA	Shawn Larsen
Stanford University	Subhasish Mitra	Joint research	Quantifying the accuracy of fault injection tools that operate at difference system abstraction levels	ASCR	Greg Bronevetsky
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
Technical University of Denmark	Robert Read	Collaboration	Water waves and wave energy generation	—	Bill Henshaw
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	UCOP	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	Bayesian uncertainty quantification in prediction of flows in highly heterogeneous media	ASCR Base	Panayot Vassilevski
Texas A&M University	Bjarne Stroustrup and Lawrence Rauchwerger	Joint research	Compiler construction and parallel optimizations	ASCR	Dan Quinlan
Texas State University	Byron Gao	Collaboration	Search user interfaces; clustering	LDRD	David Buttler
Tufts University	Scott MacLachlan	Joint research	Parallel multigrid in time	ASCR	Robert Falgout

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
UC Berkeley	Domagoj Babic	Joint research	Formal models of communication protocols	ASCR	Greg Bronevetsky
UC Berkeley	Doug Dreger	Collaboration	Earthquake hazard	IGPP	Shawn Larsen
UC Berkeley and Ohio State University	Alper Atamturk and Simge Kucukyavuz	Joint research	Stronger formulations for optimization problems in electrical power generation	Strategic hire	Deepak Rajan
UC Davis	François Gygi	Subcontract	Algorithms for electronic structure and first-principles molecular dynamics simulations using large-scale parallel computers	ASC	Art Mirin and Erik Draeger
UC Davis	Bernd Hamann	Joint research	Analysis and visualization of performance data	UC fee shared project	Timo Bremer
UC Davis	Bernd Hamann	Joint research	Performance analysis and visualization	ASC	Martin Schulz
UC Davis	Ken Joy	Joint research	Visualization and Analytics Center for Enabling Technologies	ASCR SciDAC	Jon Cohen
UC Davis	Ken Joy	Subcontract	Improving accuracy and efficiency of 3D aerial video	DoD	Jon Cohen
UC Davis	Ken Joy	Collaboration, Lawrence Scholar Program	Discrete multimaterial interface reconstruction for volume fraction data	UCOP	Jon Cohen
UC Davis	Ken Joy	Collaboration, Lawrence Scholar Program	Video-processing research for the VidCharts and Persistics projects	UCOP	Jon Cohen
UC Davis	Louise Kellogg	Joint research	Topological analysis of geological data	LDRD	Timo Bremer
UC Davis	Kwan-Liu Ma	Subcontract	Interactive tomographic reconstruction of aerial motion imagery	DoD	Jon Cohen
UC Davis	Nelson Max	Joint research	3D from wide-area aerial video	DOE Nonproliferation	Jon Cohen
UC Davis	Zhendong Su	Subcontract	ROSE support for rewrapping macro calls	ASCR	Dan Quinlan
UC Davis	Matt Bishop and Sean Peisert	Joint research	Cybersecurity research; joint proposals; cyber defenders	—	Celeste Matarazzo
UC Riverside	Michalis Faloutsos	Joint research	Cybersecurity research; joint proposals	—	Celeste Matarazzo
UC San Diego	Ido Akkerman	Subcontract	Adding isogeometric analysis capability in a high-order curvilinear research hydrocode	LDRD	Tzanio Kolev
UC San Diego	Randy Bank	Subcontract	Solvers for large sparse systems of linear equations	ASCR	Robert Falgout
UC San Diego	Laura Carrington	Joint research	Institute for Sustained Performance, Energy, and Resilience	ASCR SciDAC, ASC	Bronis de Supinski
UC San Diego	Falko Kuester	Joint research	3D from video	DoD	Jon Cohen
UC San Diego	Falko Kuester	Joint research	Large-scale atomistic simulation visualization	ASC	Jon Cohen

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
UC San Diego	Laura Carrington	Collaboration	Data-intensive architectures	LDRD	Maya Gokhale
UC San Diego	Steve Swanson	Collaboration	Persistent memory emulator	—	Maya Gokhale
UC San Diego	Erik Gartzke and Jon Lindsey	Joint research	Cybersecurity research; joint proposals	—	Celeste Matarazzo
UC San Diego, Scripps Institution of Oceanography	Julie McClean	Collaboration	Ultrahigh-resolution coupled climate simulations	BER	Art Mirin
UC Santa Cruz	Steve Kang	Collaboration	Persistent memory devices	LDRD	Maya Gokhale
UC Santa Cruz	Carlos Maltzahn	Collaboration, Lawrence Scholar Program	Semantic file systems	LDRD	Maya Gokhale
UC Santa Cruz	Stan Woosley	Joint research	Computational Astrophysics Consortium	ASCR SciDAC	Louis Howell
United States Army Medical Research Unit, Kenya	John Waitumbi	Joint research	Pathogen diagnostics for biosurveillance in Kenya	—	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 Bath	Robert Scheichl	Joint research	Multilevel Markov chain Monte Carlo methods	ASCR Base	Panayot Vassilevski
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 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	ASCR, ASC	Robert Falgout
University of Colorado, Denver	Andrew Knyazev	Subcontract	Improving efficiency and accuracy of O(N) solvers in finite-difference density functional theory calculations	LDRD	Jean-Luc Fattebert

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
University of Delaware	Richard Braun	Collaboration	Models of the eye	ASCR Base	Bill Henshaw
University of Delaware	John Cavazos	Subcontract	ROSE compiler project	ASCR	Dan Quinlan and Chunhua Liao
University of Houston	Yuriy Fofanov	Joint research	Metagenomic analysis	DTRA	Tom Slezak
University of Illinois at Urbana-Champaign	William Gropp	Collaboration, Lawrence Scholar Program	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	Joint research	Node mapping optimizations for high performance computing systems, optimization of CHARM++ codes	ASC, LDRD	Martin Schulz
University of Illinois at Urbana-Champaign	Laxmikant Kale and Esteban Meneses	Joint research	Scalable check pointing and message logging for fault-tolerant high performance computing systems	ASCR	Greg Bronevetsky
University of Illinois at Urbana-Champaign	Rakesh Kumar and Joseph Sloan	Joint research	Algorithm-specific fault tolerance for sparse linear algebra applications	ASCR	Greg Bronevetsky
University of Illinois and IBM	Hormozd Gahvari, William Gropp, Luke Olson, and Kirk Jordan	Collaboration	Modeling algebraic multigrid performance on multicore architectures	LDRD	Ulrike Yang
University of Louisville	Yongsheng Liam	Collaboration	Micro-air vehicles	—	Bill Henshaw
University of Maryland	Jeffrey Hollingsworth	Joint research	Autotuning and tool infrastructures	ASCR	Martin Schulz
University of Maryland	Jeffrey Hollingsworth	Joint research	Institute for Sustained Performance, Energy, and Resilience	ASCR SciDAC	Bronis de Supinski
University of Massachusetts, Amherst	Andrew McCallum	Joint research	Cross-language topic models	LDRD	David Buttler
University of Michigan	R. Paul Drake	ASC Predictive Science Academic Alliance Program Center	Center for Radiative Shock Hydrodynamics (CRASH)	ASC	Dick Watson
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	Scalable tool infrastructures	ASC, ASCR	Bronis de Supinski
University of New Mexico	Dorian Arnold	Joint research	Tool infrastructures	ASC	Martin Schulz

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
University of North Carolina	Robert Fowler	Joint research	Institute for Sustained Performance, Energy, and Resilience	ASCR SciDAC	Bronis de Supinski
University of North Carolina	Jan Prins	Joint research	Efficient OpenMP runtimes for tasking	ASC	Bronis de Supinski
University of North Carolina	Jan Prins	Joint research	OpenMP task scheduling	ASC	Martin Schulz
University of San Francisco	Jeff Buckwalters	Joint research	Performance modeling	ASC	Martin Schulz
University of Southern California	Gerard Medioni	Subcontract	Activity analysis in wide-area overhead video	DOE Nonproliferation	Jon Cohen
University of Southern California	Robert Lucas and Jacqueline Chame	Joint research	Institute for Sustained Performance, Energy, and Resilience	ASCR SciDAC	Bronis de Supinski
University of Tennessee	Jack Dongarra	Joint research	Empirical tuning	ASCR	Dan Quinlan
University of Tennessee	Jack Dongarra	Joint research	Institute for Sustained Performance, Energy, and Resilience	ASCR SciDAC	Bronis de Supinski
University of Texas, Austin	Robert Moser	ASC Predictive Science Academic Alliance Program Center	Center for Predictive Engineering and Computational Sciences (PECOS)	ASC	Dick Watson
University of Texas, San Antonio	Shirley Moore	Joint research	Institute for Sustained Performance, Energy, and Resilience	ASCR SciDAC	Bronis de Supinski
University of Turabo, puerto Rico	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	Semantic debugging of message-passing interface applications	ASCR, ASC	Bronis de Supinski
University of Utah	Valerio Pascucci	Joint research	Performance analysis and visualization	ASC	Martin Schulz
University of Utah	Valerio Pascucci	Subcontract	Performance analysis and visualization	LDRD	Timo Bremer

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
University of Utah	Chris Johnson, Valerio Pascucci, Chuck Hansen, Claudio Silva, Lee Myers, Allen Sanderson, and Steve Parker	Joint research	Visualization and Analytics Center for Enabling Technologies	ASCR SciDAC	Jon Cohen
University of Washington	Carl Ebeling and Scott Hauck	Collaboration	Coarse-grain processor architectures	—	Maya Gokhale
University of Waterloo	Hans de Sterck	Subcontract	Numerical methods for large-scale data factorization	LDRD	Van Emden Henson
University of Wisconsin	Ben Liblit	Joint research	Scalable debugging	ASCR, ASC	Bronis de Supinski
University of Wisconsin	Bart Miller	Joint research	Scalable debugging	ASCR, ASC	Bronis de Supinski
University of Wisconsin	Karu Sankaralingam	Joint research	Fault tolerant computing models for high performance computing	ASC	Martin Schulz
University of Wisconsin	Karu Sankaralingam	Joint research	Resilient computing	ASCR, ASC	Bronis de Supinski
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	Mathematical and computational foundations of network sciences	SMS	Celeste Matarazzo
Virginia Institute of Technology	Madhav Marathe	Subcontract	Research in network sciences	SMS	Evi Dube and James Brase
Virginia Polytechnic Institute	John Burns	Collaboration	Building simulations, reduced order models, and control	GPIC	Bill Henshaw
Voorhees College	Tim Kentopp	Joint research	Malware intelligence harvesting for greater cyber defense	NNSA	Matt Myrick

Publications

Book Chapters

Childs, H., et al., "VisIt: An End-User Tool for Visualizing and Analyzing Very Large Data," in *High Performance Visualization*, E. W. Bethel, H. Childs, and C. Hansen (Chapman and Hall/CRC), 357–372.

Childs, H., et al., "Visualization at Extreme Scale Concurrency," in *High Performance Visualization*, E. W. Bethel, H. Childs, and C. Hansen (Chapman and Hall/CRC), 291–306.

Childs, H., et al., "In Situ Processing," in *High Performance Visualization*, E. W. Bethel, H. Childs, and C. Hansen (Chapman and Hall/CRC), 171–198.

Journal Papers

Ames, S., M. Gokhale, and C. Maltzahn, "QMDS: A File System Metadata Management Service Supporting a Graph Data Model-Based Query Language," *International Journal of Parallel, Emergent, and Distributed Systems*, 1–25.

Appelo, D., et al., "Numerical Methods for Solid Mechanics on Overlapping Grids: Linear Elasticity," *J. Comput. Phys.* **231** (18), 6012–6050.

Appelo, D. and N. A. Petersson, "A Fourth-Order Accurate Embedded Boundary Method for the Wave Equation," *SIAM J. Sci. Comput.* **34** (6), A2982–A3008 (LLNL-JRNL-417163).

Arsenlis, A., et al., "A Dislocation Dynamics Study of the Transition from Homogeneous to Heterogeneous Deformation in Irradiated Body-Centered Cubic Iron," *Acta Mater.* **60** (9), 3748–3757.

Baker, A.H., et al., "Scaling *hypre*'s Multigrid Solvers to 100,000 Cores," *High Performance Scientific Computing: Algorithms and Applications*, 261–279 (LLNL-JRNL-479591).

Banks, J. W. and W. D. Henshaw, "Upwind Schemes for the Wave Equation in Second-Order Form," *J. Comput. Phys.* **231** (17), 5854–5889.

Banks, J. W., W. D. Henshaw, and D. W. Schwendeman, "Deforming Composite Grids for Solving Fluid Structure Problems," *J. Comput. Phys.* **231** (9), 3518–3547.

Banks, J. W., et al., "Numerical Error Estimation for Nonlinear Hyperbolic PDEs via Nonlinear Error Transport," *Comput. Method. Appl. M.* **213**, 1–15 (LLNL-JRNL-478411).

Benedetti, L. R., et al., "Crosstalk in X-Ray Framing Cameras: Effect on Voltage, Gain, and Timing," *Rev. Sci. Instrum.* **83** (10).

Bhatia, H., et al., "Flow Visualization with Quantified Spatial and Temporal Errors Using Edge Maps," *IEEE T. Vis. Comput. Gr.* **18** (9), 1383–1396.

Bihari, B. L., "Transactional Memory for Unstructured Mesh Simulations," *J. Sci. Comput.* **54** (2), 311–332 (LLNL-JRNL-484451).

Borglin, S., et al., "Application of Phenotypic Microarrays to Environmental Microbiology," *Curr. Opin. Biotech.* **23** (1), 41–48.

Brezina, M., et al., "Relaxation-Corrected Bootstrap Algebraic Multigrid (rBAMG)," *Numer. Linear Algebr.* **19** (2), 178–193.

Brezina, M., P. Vanek, and P. S. Vassilevski, "An Improved Convergence Analysis of Smoothed Aggregation Algebraic Multigrid," *Numer. Linear Algebr.* **19** (3), 441–469.

Broering, T. M., Y. S. Lian, and W. Henshaw, "Numerical Investigation of Energy Extraction in a Tandem Flapping Wing Configuration," *AIAA J.* **50** (11), 2295–2307.

Budanur, S., F. Mueller, and T. Gamblin, "Memory Trace Compression and Replay for SPMD Systems Using Extended PRSDs," *Comput. J.* **55** (2), 206–217.

Cadag, E., et al., "Computational Analysis of Pathogen-Borne Metallo-ß-Lactamases Reveals Discriminating Structural Features Between B1 Types," *BMC Research Notes* **5**, 96.

Casas, M., et al., "Extracting the Optimal Sampling Frequency of Applications Using Spectral Analysis," *Concurr. Comp.-Pract. E.* **24** (3), 237–259.

Cohen, R. H., M. Dorf, and M. Dorr, "Reduced Electron Models for Edge Simulation," *Contrib. Plasm. Phys.* **52** (5–6), 529–553.

Connors, J. M. and J. S. Howell, "A Fluid–Fluid Interaction Method Using Decoupled Subproblems and Differing Time Steps," *Numer. Meth. Part. D. E.* **28** (4), 1283–1308.

Connors, J. M., J. S. Howell, and W. J. Layton, "Decoupled Time Stepping Methods for Fluid–Fluid Interaction," *SIAM J. Numer. Anal.* **50** (3), 1297–1319.

Dennis, J. M., et al., "CAM-SE: A Scalable Spectral Element Dynamical Core for the Community Atmosphere Model," *Int. J. High Perform. C.* **26** (1), 74–89 (LLNL-JRNL-469111).

Dennis, J. M., et al., "An Application-Level Parallel I/O Library for Earth System Models," *Int. J. High Perform. C.* **26** (1), 43–53 (LLNL-JRNL-469337).

Dennis, J. M., et al., "Computational Performance of Ultra-High-Resolution Capability in the Community Earth System Model," *Int. J. High Perform. C.* **26** (1), 5–16 (LLNL-JRNL-484456).

Di Martino, B., et al., "Graphical Processing Units and Scientific Applications," *Int. J. High Perform. C.* **26** (3), 189–191.

Dobrev, V. A., T. V. Kolev, and R. N. Rieben, "High-Order Curvilinear Finite Element Methods for Lagrangian Hydrodynamics," *SIAM J. Sci. Comput.* **34** (5), B606–B641.

Doppner, T., et al., "Direct Measurement of Energetic Electrons Coupling to an Imploding Low-Adiabat Inertial Confinement Fusion Capsule," *Phys. Rev. Lett.* **108** (13).

Dorf, M. A., et al., "Progress with the COGENT Edge Kinetic Code: Collision Operator Options," *Contrib. Plasm. Phys.* **52** (5–6), 518–522.

Dougherty, M. J., et al., "Glycoside Hydrolases from a Targeted Compost Metagenome, Activity-Screening and Functional Characterization," *BMC Biotechnol.* **12** (38).

Epperly, T. G. W., et al., "High-Performance Language Interoperability for Scientific Computing through Babel," *Int. J. High Perform. C.* **26** (3), 260–274 (LLNL-JRNL-465223).

Fattebert, J.-L., D. F. Richards, and J. N. Glosli, "Dynamic Load-Balancing Algorithm for Molecular Dynamics Based on Voronoi Cells Domain Decompositions," *Comput. Phys. Commun.* **183** (12), 2608–2615.

Fiore, A. M., et al., "Global Air Quality and Climate," *Chem. Soc. Rev.* **41** (19), 6663–6683 (LLNL-JRNL-563696).

Fry, M. M., et al., "The Influence of Ozone Precursor Emissions from Four World Regions on Tropospheric Composition and Radiative Climate Forcing," *J. Geophys. Res.-Atmos.* **117** (D7).

Gaither, K. P., et al., "Visual Analytics for Finding Critical Structures in Massive Time-Varying Turbulent-Flow Simulations," *IEEE Comput. Graph.* **32** (4), 34–45.

Glenn, S. M., et al., "Extracting Core Shape from X-Ray Images at the National Ignition Facility," *Rev. Sci. Instrum.* **83** (10).

Glenzer, S. H., et al., "Cryogenic Thermonuclear Fuel Implosions on the National Ignition Facility," *Phys. Plasmas* **19** (5).

Glenzer, S. H., et al., "First Implosion Experiments with Cryogenic Thermonuclear Fuel on the National Ignition Facility," *Plasma Phys. Contr. F.* **54** (4).

Graziani, F. R., et al., "Large-Scale Molecular Dynamics Simulations of Dense Plasmas: The Cimarron Project," *High Energ. Dens. Phys.* **8** (1), 105–131 (LLNL-JRNL-463244).

Grout, R. W., et al., "A Direct Numerical Simulation Study of Turbulence and Flame Structure in Transverse Jets Analyzed in Jet-Trajectory-Based Coordinates," *J. Fluid Mech.* **706**, 351–383.

Gyulassy, A., P.-T. Bremer, and V. Pascucci, "Computing Morse–Smale Complexes with Accurate Geometry," *IEEE T. Vis. Comput. Gr.* **18** (12), 2014–2022.

Hadjadj, A., H. C. Yee, and B. Sjogreen, "LES of Temporally Evolving Mixing Layers by an Eighth-Order Filter Scheme," *Int. J. Numer. Meth. Fl.* **70** (11), 1405–1427.

Hakim, A. H., et al., "Coupled Core-Edge Simulations of H-Mode Buildup Using the Fusion Application for Core-Edge Transport Simulations (FACETS) Code," *Phys. Plasmas* **19** (3).

Harmand, M., et al., "Plasma Switch as a Temporal Overlap Tool for Pump-Probe Experiments at FEL Facilities," *J. Instrum.* **7**.

Harrison, C. D., et al., "The XMM Cluster Survey: The Stellar Mass Assembly of Fossil Galaxies," *Astrophys. J.* **752** (1).

Henson, V. E. and G. Sanders, "Locally Supported Eigenvectors of Matrices Associated with Connected and Unweighted Power-Law Graphs," *Electron T. Numer. Anal.* **39**, 353–378.

Hilton, M., et al., "The XMM Cluster Survey: Evidence for Energy Injection at High Redshift from Evolution of the X-Ray Luminosity–Temperature Relation," *Mon. Not. R. Astron. Soc.* **424** (3), 2086–2096.

Hysom, D. A., et al., "Skip the Alignment: Degenerate, Multiplex Primer, and Probe Design Using K-mer Matching Instead of Alignments," *PLoS ONE* **7** (4) (LLNL-JRNL-491362).

Iglesias, C. A. and V. Sonnad, "Partially Resolved Transition Array Model for Atomic Spectra," *High Energ. Dens. Phys.* **8** (2), 154–160.

Jones, O. S., et al., "A High-Resolution Integrated Model of the National Ignition Campaign Cryogenic Layered Experiments," *Phys. Plasmas* **19** (5).

Khan, S. F., et al., "Methods for Characterizing X-Ray Detectors for Use at the National Ignition Facility," *Rev. Sci. Instrum.* **83** (10).

Khudyakov, J. I., et al., "Global Transcriptome Response to Ionic Liquid by a Tropical Rain Forest Soil Bacterium, *Enterobacter lignolyticus*," *P. Natl. Acad. Sci. USA* **109** (32), E2173–E2182 (LLNL-JRNL-491427).

Kolev, Tz. and P. S. Vassilevski, "Regular Decompositions for *H*(div) Spaces," *Comput. Method. Appl. M.* **12** (4), 437–447 (LLNL-JRNL-520451).

Kolev, Tz. and P. S. Vassilevski, "Parallel Auxiliary Space AMG Solver for *H*(div) Problems," *SIAM J. Sci. Comput.* **34** (6), A3079–A3098.

Kreiss, H. O., O. E. Ortiz, and N. A. Petersson, "Initial Boundary Value Problems for Second-Order Systems of Partial Differential Equations," *ESAIM-Math. Model. Num.* **46** (3), 559–593.

Kreiss, H. O. and N. A. Petersson, "Boundary Estimates for the Elastic Wave Equation in Almost Incompressible Materials," *SIAM J. Numer. Anal.* **50** (3), 1556–1580 (LLNL-JRNL-482152).

Landen, O. L., et al., "Progress in the Indirect-Drive National Ignition Campaign," *Plasma Phys. Contr. F.* **54** (12).

Landge, A. G., et al., "Visualizing Network Traffic to Understand the Performance of Massively Parallel Simulations," *IEEE T. Vis. Comput. Gr.* **18** (12), 2467–2476 (LLNL-CONF-543359).

Lang, D. M., A. T. Zemla, and C. L. Zhou, "Highly Similar Structural Frames Link the Template Tunnel and NTP Entry Tunnel to the Exterior Surface in RNA-Dependent RNA Polymerases," *Nucleic Acids Res.* **41** (3), 1464–1482 (LLNL-JRNL-602932).

Lashuk, I. V., et al., "A Massively Parallel Adaptive Fast Multiple Method on Heterogeneous Architectures," *Commun. ACM* **55** (5), 101–109.

Lashuk, I. V. and P. S. Vassilevski, "Element Agglomeration Coarse Raviart–Thomas Spaces with Improved Approximation Properties," *Numer. Linear Algebr.* **19** (2), 414–426.

Lauritzen, P. H., et al., "Implementation of New Diffusion/Filtering Operators in the CAM-FV Dynamical Core," *Int. J. High Perform. C.* **26** (1), 63–73 (LLNL-JRNL-467302).

Levine, J. A., et al., "A Quantized Boundary Representation of 2D Flows," *Comput. Graph. Forum* **31** (3), 945–954.

Liao, Z. M., et al., "Predictive Modeling Techniques for Nanosecond-Laser Damage Growth in Fused Silica Optics," *Opt. Express* **20** (14), 15569–15579.

Lott, P. A., et al., "An Accelerated Picard Method for Nonlinear Systems Related to Variably Saturated Flow," *Adv. Water Res.* **38**, 92–101.

Mackinnon, A. J., et al., "Assembly of High-Areal-Density Deuterium–Tritium Fuel from Indirectly Driven Cryogenic Implosions," *Phys. Rev. Lett.* **108** (21).

Madan, D., et al., "Enhanced Performance of Dispenser Printed MA n-type Bi_2Te_3 Composite Thermoelectric Generators," *ACS Appl. Mater. Interfaces* **4** (11), 6117–6124.

Mehrtens, N., et al., "The XMM Cluster Survey: Optical Analysis Methodology and the First Data Release," *Mon. Not. R. Astron. Soc.* **423** (2), 1024–1052.

Mirin, A. A. and P. H. Worley, “Improving the Performance Scalability of the Community Atmosphere Model,” *Int. J. High Perform. C.* **26** (1), 17–30 (LLNL-JRNL-436196).

Mohror, K. and K. L. Karavanic, “Trace Profiling: Scalable Event Tracing on High-End Parallel Systems,” *Parallel Comput.* **38** (4–5), 194–225.

Moody, J. D., et al., “Multistep Redirection by Cross-Beam Power Transfer of Ultrahigh-Power Lasers in a Plasma,” *Nat. Phys.* **8** (4), 344–349.

Negres, R. A., et al., “Probability of Growth of Small Damage Sites on the Exit Surface of Fused Silica Optics,” *Opt. Express* **20** (12), 13030–13039 (LLNL-JRNL-546011).

Neville, J., et al., “Correcting Evaluation Bias of Relational Classifiers with Network Cross Validation,” *Knowl. Inf. Syst.* **30** (1), 31–55 (LLNL-JRNL-455699).

Norgard, G. and P.-T. Bremer, “Second Derivative Ridges are Straight Lines and the Implications for Computing Lagrangian Coherent Structures,” *Physica D-Non Lin. Phen.* **241** (18), 1475–1476.

Ozlem, M., et al., “A Numerical Study of Shock-Induced Cavity Collapse,” *Shock Waves* **22** (2), 89–117.

Park, J. I., et al., “A Thermophilic Ionic Liquid-Tolerant Cellulase Cocktail for the Production of Cellulosic Biofuels,” *PLoS ONE* **7** (5).

Petersson, N. A. and B. Sjogreen, “Stable and Efficient Modeling of Anelastic Attenuation in Seismic Wave Propagation,” *Commun. Comput. Phys.* **12** (1), 193–225 (LLNL-JRNL-460239).

Potter, D. and N. Antipa, “Visualization of Target Inspection Data at the National Ignition Facility,” *Fusion Eng. Des.* **87** (12), 2136.

Rauscher, S., et al., “Exploring a Global Multiresolution Approach Using Aquaplanet Simulations,” *Journal of Climate*, (LLNL-JRNL-538911).

Reddy, A. P., et al., “Thermophilic Enrichment of Microbial Communities in the Presence of the Ionic Liquid 1-Ethyl-3-Methylimidazolium Acetate,” *J. Appl. Microbiol.* **113** (6), 1362–1370.

Robey, H. F., et al., “Shock Timing Experiments on the National Ignition Facility: Initial Results and Comparison with Simulation,” *Phys. Plasmas* **19** (4).

Robey, H. F., et al., “Precision Shock Tuning on the National Ignition Facility,” *Phys. Rev. Lett.* **108** (21).

Scheichl, R., P. S. Vassilevski, and L. T. Zikatanov, “Multilevel Methods for Elliptic Problems with Highly Varying Coefficients on Nonaligned Coarse Grids,” *SIAM J. Numer. Anal.* **50** (3), 1675–1694.

Schleife, A., et al., “Plane-Wave Pseudopotential Implementation of Explicit Integrators for Time-Dependent Kohn–Sham Equations in Large-Scale Simulations,” *J. Chem. Phys.* **137** (22).

Schneider, M. B., et al., “Soft X-Ray Images of the Laser Entrance Hole of Ignition Hohlraums,” *Rev. Sci. Instrum.* **83** (10).

Schroder, J. B., “Smoothed Aggregation Solvers for Anisotropic Diffusion,” *Numer. Linear Algebr.* **19** (2), 296–312.

Schwendeman, D. W., A. K. Kapila, and W. D. Henshaw, “A Comparative Study of Two Macroscale Models of Condensed-Phase Explosives,” *IMA J. Appl. Math.* **77** (1), 2–17.

Sjogreen, B., “High Order Finite Difference and Finite Volume Methods for Advection on the Sphere,” *J. Sci. Comput.* **51** (3), 703–732.

Sjogreen, B. and N. A. Petersson, “A Fourth-Order Accurate Finite Difference Scheme for the Elastic Wave Equation in Second-Order Formulation,” *J. Sci. Comput.* **52** (1), 17–48 (LLNL-JRNL-483427).

Strozzi, D. J., et al., “Threshold for Electron Trapping Nonlinearity in Langmuir Waves,” *Phys. Plasmas* **19** (11), 112306 (LLNL-JRNL-574812).

Sun, Y., et al., “Combining Simulation and Emulation for Calibrating Sequentially Reactive Transport Systems,” *Transport Porous Med.* **92** (2), 509–526.

Turney, J. M., et al., “PSI4: An Open-Source *Ab Initio* Electronic Structure Program,” *Wires Comput. Mol. Sci.* **2** (4), 556–565.

Wang, W., et al., “High Order Finite Difference Methods with Subcell Resolution for Advection Equations with Stiff Source Terms,” *J. Comput. Phys.* **231** (1), 190–214.

Wild, O., et al., “Modeling Future Changes in Surface Ozone: A Parameterized Approach,” *Atmos. Chem. Phys.* **12** (4), 2037–2054.

Wolf, J., et al., “On the Optimization of Schedules for MapReduce Workloads in the Presence of Shared Scans,” *VLDB J.* **21** (5), 589–609.

Wolf, J. L., et al., “CIRCUMFLEX: A Scheduling Optimizer for MapReduce Workloads with Shared Scans,” *Operating Systems Review* **46** (1), 26–32.

Zhang, Y. Y., et al., “Regional Assessment of the Parameter-Dependent Performance of CAM4 in Simulating Tropical Clouds,” *Geophys. Res. Lett.* **39** (14).

Zhao, C. F., et al., “Toward Understanding of Differences in Current Cloud Retrievals of ARM Ground-Based Measurements,” *J. Geophys. Res.-Atmos.* **117**.

Conference Proceedings

Baker, A., et al., “Scaling Algebraic Multigrid Solvers: On the Road to Exascale,” *Competence in High Performance Computing*, (Springer), (LLNL-PROC-463941).

Beeler Jr., R., et al., “Shot Planning and Analysis Tools on the NIF Project,” *8th IAEA Technical Meeting on Control, Data Acquisition, and Remote Participation for Fusion Research*, (Fusion Eng. Des., Elsevier), **87**, 2020 (LLNL-CONF-491575).

Benedetti, L. R., et al., “Crosstalk in X-Ray Framing Cameras: Effect on Voltage, Gain, and Timing,” *19th Topical Conference on High-Temperature Plasma Diagnostics*, (Rev. Sci. Instrum., Amer. Inst. Physics), **83**.

Bhatele, A., et al., “Novel Views of Performance Data to Analyze Large-Scale Adaptive Applications,” *International Conference for High Performance Computing, Networking, Storage and Analysis (SC12)*, (ACM/IEEE), (LLNL-CONF-554552).

Bhatele, A., et al., "Mapping Applications with Collectives Over Subcommunicators on Torus Networks," *International Conference for High Performance Computing, Networking, Storage and Analysis (SC12)*, (ACM/IEEE), (LLNL-CONF-556491).

Bihari, B. L., et al., "A Case for Including Transactions in OpenMP II: Hardware Transactional Memory," *International Workshop on OpenMP*, (Lect. Note Comput. Sci., Springer), **7312**, 44–58, (LLNL-PROC-528852).

Böhme, D., et al., "Scalable Critical-Path-Based Performance Analysis," *26th International Parallel and Distributed Processing Symposium (IPDPS 2012)*, (IEEE), (LLNL-CONF-522431).

Bronevetsky, G., et al., "Automatic Fault Characterization via Abnormality-Enhanced Classification," *42nd Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, (Int. C. Depend. Systems Networks, IEEE), (LLNL-CONF-464273).

Correa, C. D. and P. Lindstrom, "Locally Scaled Spectral Clustering Using Empty Region Graphs," *ACM SIGKDD Conference on Knowledge Discovery and Data Mining 2012*, (ACM), 1330–1338, (LLNL-CONF-513768).

Glenn, S. M., et al., "Extracting Core Shape from X-Ray Images at the National Ignition Facility," *19th Topical Conference on High-Temperature Plasma Diagnostics*, (Rev. Sci. Instrum., Amer. Inst. Physics), **83**.

Henderson, K., et al., "RolX: Structural Role Extraction and Mining in Large Graphs," *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, (LLNL-CONF-559231).

Hilbrich, T., et al., "GTI: A Generic Tools Infrastructure for Event-Based Tools in Parallel Systems," *26th International Parallel and Distributed Processing Symposium (IPDPS 2012)*, (IEEE), (LLNL-CONF-522031).

Hutton, M. S., et al., "Experiment Archive, Analysis, and Visualization at the National Ignition Facility," *8th IAEA Technical Meeting on Control, Data Acquisition, and Remote Participation for Fusion Research*, (Fusion Eng. Des., Elsevier), **87**, 2087.

Khan, S. F., et al., "Methods for Characterizing X-Ray Detectors for Use at the National Ignition Facility," *19th Topical Conference on High-Temperature Plasma Diagnostics*, (Rev. Sci. Instrum., Amer. Inst. Physics), **83**.

Khandekar, R., et al., "Scheduling with Setup Costs and Monotone Penalties," *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2012)*, (Leibniz International Proceedings in Informatics), **18**, 185–198.

Langer, S., et al., "Simulating Laser–Plasma Interaction in Experiments at the National Ignition Facility on a Cray XE6," *Cray User Group Meeting 2012*, (LLNL-PROC-547711).

Liebman, J., et al., "Correcting Raw Diagnostic Data for Oscilloscope Recording System Distortions at the National Ignition Facility," *8th IAEA Technical Meeting on Control, Data Acquisition, and Remote Participation for Fusion Research*, (Fusion Eng. Des., Elsevier), **87** (12), 2125–2130.

Pilla, L., et al., "Hierarchical Approach for Load Balancing on Parallel Multicore Systems," *International Conference on Parallel Processing (ICPP12)*, (LLNL-CONF-536171).

Protze, J., et al., "Holistic Debugging of MPI Derived Datatypes," *26th International Parallel and Distributed Processing Symposium (IPDPS 2012)*, (IEEE), (LLNL-CONF-522011).

Rodgers, A., et al., "Progress on Seismic Ground Motion Simulations for Nuclear Explosion Monitoring Using Hydrodynamic and Anelastic Modeling," *Monitoring Research Review 2012*, (NNSA), (LLNL-PROC-562892).

Rossi, R., et al., "Role-Dynamics: Fast Mining of Large Dynamic Networks," *International Workshop on Large Scale Network Analysis*, (LLNL-CONF-537171).

Rountree, B., et al., "Beyond DVFS: A First Look at Performance Under a Hardware-Enforced Power Bound," *26th International Parallel and Distributed Processing Symposium (IPDPS 2012)*, (IEEE), 947–953.

Sanderson, A. R., et al., "A System for Query-Based Analysis and Visualization," *EuroGraphics International Workshop on Visual Analytics (EuroVA 2012)*, 21–25.

Schneider, M. B., et al., "Soft X-Ray Images of the Laser Entrance Hole of Ignition Hohlraums," *19th Topical Conference on High-Temperature Plasma Diagnostics*, (Rev. Sci. Instrum., Amer. Inst. Physics), **83** (LLNL-CONF-555011).

Scogland, T. R. W., et al., "Heterogeneous Task Scheduling for Accelerated OpenMP," *26th International Parallel and Distributed Processing Symposium (IPDPS 2012)*, (IEEE), (LLNL-CONF-522311).

Van Essen, B., et al., "DI-MMAP: A High Performance Memory-Map Runtime for Data-Intensive Applications," *The International Workshop on Data-Intensive Scalable Computing Systems (DISCS)*, (LLNL-CONF-583953).

Van Essen, B., et al., "Accelerating a Random Forest Classifier: Multicore, GP-GPU, or FPGA?," *20th IEEE International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, (IEEE), 232–239.

Van Essen, B., et al., "On the Role of NVRAM in Data-Intensive Architectures: An Evaluation," *26th International Parallel and Distributed Processing Symposium (IPDPS 2012)*, (IEEE), 703–714.

Vo, H. T., et al., "HyperFlow: A Heterogeneous Dataflow Architecture," *Eurographics Symposium on Parallel Graphics and Visualization 2012*, (Eurographics), 1–10 (LLNL-CONF-468431).

Wolf, J., et al., "CIRCUMFLEX: A Scheduling Optimizer for MapReduce Workloads with Shared Scans," *LADIS*, (ACM SIGOPS Operating Systems Review).

Industrial Collaborations

Company	Topic	LLNL contact
3M Cogent	Linux solution to running on vendor hardware	Anthony Ghilarducci, Kim Ferrari, and Matt Dralle
Adaptive Computing Enterprises, Inc.	Moab workload manager	Don Lipari
Allinea Software	Scalable debugging infrastructure	Matt Legendre
AMD	Performance modeling	Martin Schulz
AMD	Power and energy	Barry Rountree
AMD	Resilience	Kathryn Mohror
AOSense, Inc.	Gravity gradiometry in the detection of anomalous mass distribution in vehicles	Vijay Sonnad
Argo Navis Technologies	Automated cache performance analysis and optimization in Open|SpeedShop	Kathryn Mohror
Arista Networks	Low-latency Ethernet networks	Matt Leininger
Battelle	Terrorism risk assessments	Amy Waters and Lisa Belk
Cisco Systems, Dell, DataDirect Networks, Intel, NetApp, Mellanox Technolgoies, QLogic, Red Hat, Oracle, and Supermicro	Hyperion collaboration	Matt Leininger
Commissariat à l'Énergie Atomique	SLURM resource management software	Don Lipari
Cray	Scalable capacity clusters	Matt Leininger and Trent D'Hooge
Cray	Exploring the Chapel programming language using LULESH	Abhinav Bhatele
DataDirect Networks	RAID 6 research and development for I/O systems	Mark Gary
Dell Computers	Scalable capacity clusters	Matt Leininger and Trent D'Hooge
Électricité de France	Aeroacoustics	Bill Henshaw
Energy Exemplar	PLEXOS power grid optimization	Jeff Wolf
EOG Resources	Seismic processing	Shawn Larsen
ExxonMobil	Cooperative Research and Development Agreement: computational mathematics	Rob Falgout and Panayot Vassilevski
GAMS	Solvers	Barry Rountree
GE Energy Consulting	hpc4energy incubator: improving positive sequence load flow simulation performance and capability	Steve Smith
GE Global Research	Wind power	Bill Henshaw
IBM	Deep Computing Solutions	Fred Streitz
IBM	Blue Gene/Q: the evaluation of transactional memory and the modeling and optimization of algebraic multigrid	Martin Schulz
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	Heart modeling	Art Mirin
IBM	High performance storage system	Jerry Shoopman
IBM	Scalable systems, multiple areas	Kim Cupps and others

Company	Topic	LLNL contact
IBM	Tool interface for OpenMP	Martin Schulz
IBM and Energy Exemplar	Solvers	Barry Rountree
IBM Research	Improvements to CPLEX optimization software geared toward use cases for California Energy Systems for the 21st Century (CES-21)	Deepak Rajan
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, Stony Brook University, and Carnegie Mellon University	Mining large time-evolving graphs for proximity queries	Deepak Rajan
InfiniBand Trade Association	InfiniBand specifications body	Pam Hamilton
Intel Corporation	Many integrated core	Greg Lee
Intel Corporation	Research and development for I/O systems	Mark Gary, Robin Goldstone, and Ned Bass
Intel: High Performance Data Division	Lustre file system development and support	Kim Cupps, Chris Morrone, and Marc Stearman
ION Geophysical Corporation	Oil exploration	Shawn Larsen
ISO New England	hpc4energy incubator: evaluation of robust unit commitment	Barry Smith
Krell	Small business technology transfer	Barry Rountree
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
Mellanox	Long haul InfiniBand	Trent D'Hooge
MITRE Corporation	Subsurface modeling	Shawn Larsen
National Instruments	Object-oriented methods applied to controls	Mike Flegel
NetApp	High performance I/O systems	Marc Stearman and Mark Gary
NSTec	Instrument calibration techniques	Steve Glenn
OCZ Technology	Solid-state drive technology use	Trent D'Hooge
OpenFabrics Alliance, Mellanox, and QLogic	OpenFabrics enterprise distribution	Matt Leininger
OpenSFS	Lustre file system development and deployment	Terri Quinn, Pam Hamilton, and Chris Morrone
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 cybersecurity	Dan Quinlan
Red Hat	Operating systems	Mark Grondona and Jim Foraker

Company	Topic	LLNL contact
Red Hat, Appro, Intel, and AMD	Hardware performance counters	Barry Rountree
Roche NimbleGen	Next-generation peptide microarrays	Tom Slezak
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 and others
ScaleMP	Large memory architectures	Robin Goldstone
SchedMD	SLURM resource management software	Kim Cupps and Don Lipari
Tennessee Valley Authority and Applied Communication Sciences	Robust adaptive topology control for Advanced Research Projects Agency-Energy project	Barry Rountree
United Technologies Research Center	hpc4energy incubator: improving simulations of advanced internal combustion engines	Steve Smith