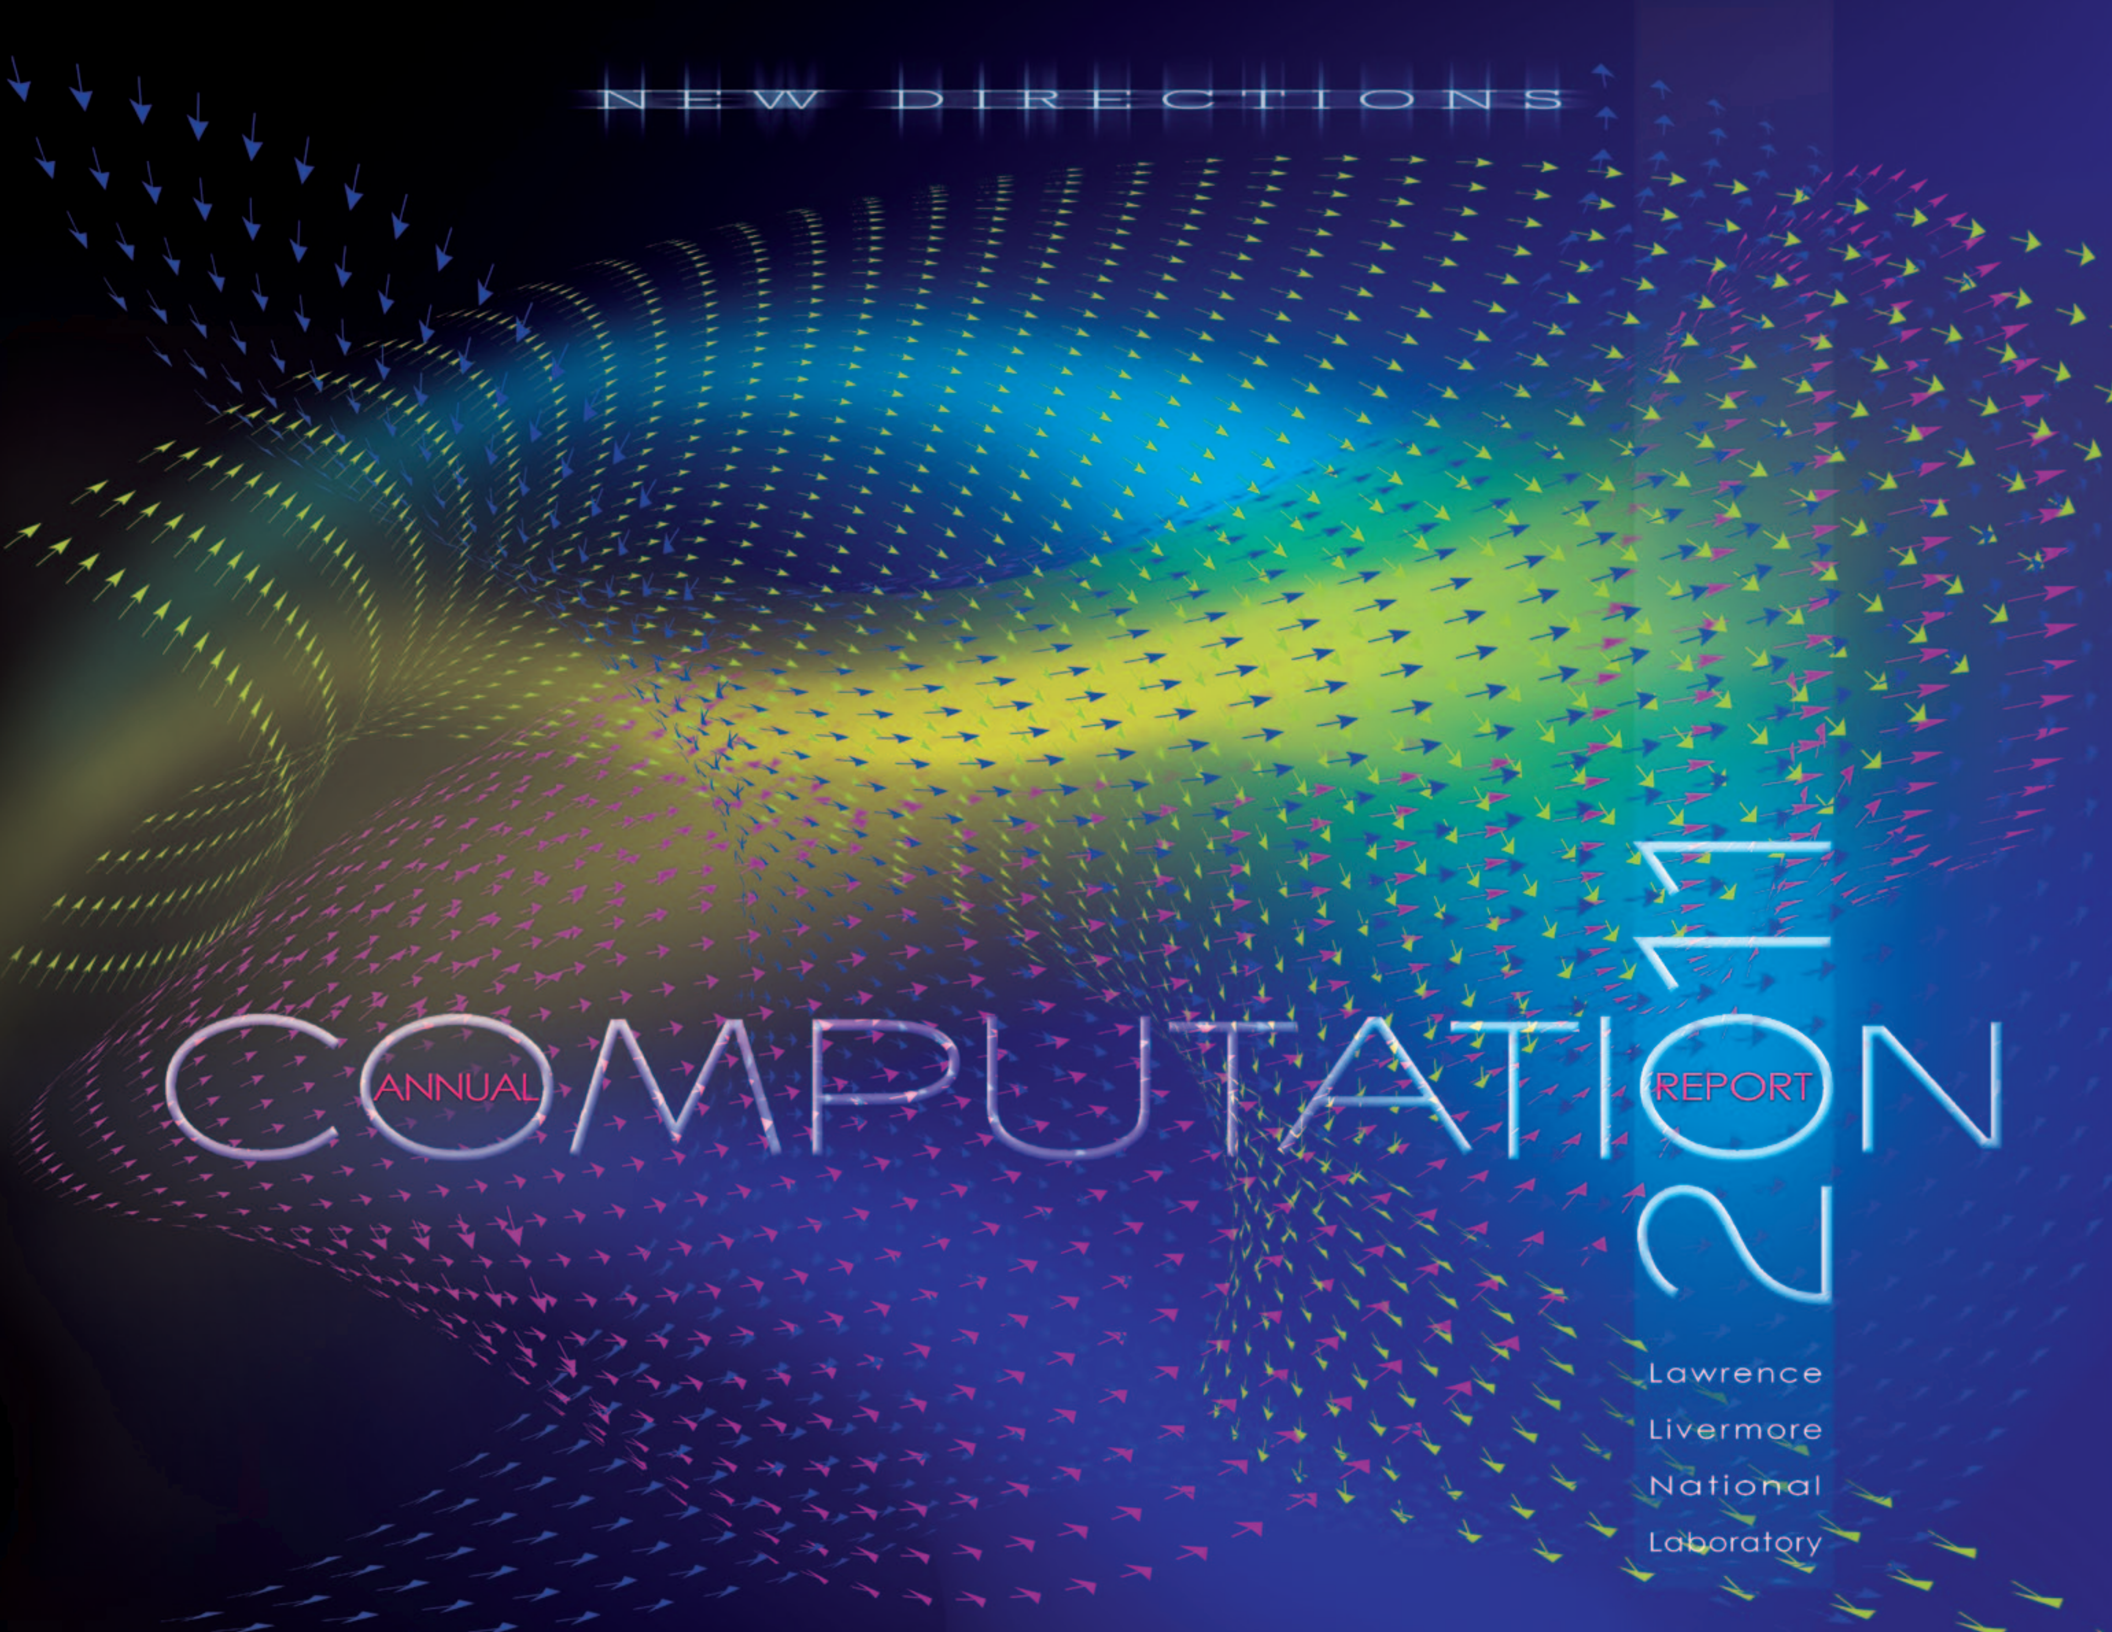

NEW DIRECTIONS

COMPUTATION ANNUAL REPORT 2011

Lawrence Livermore National Laboratory

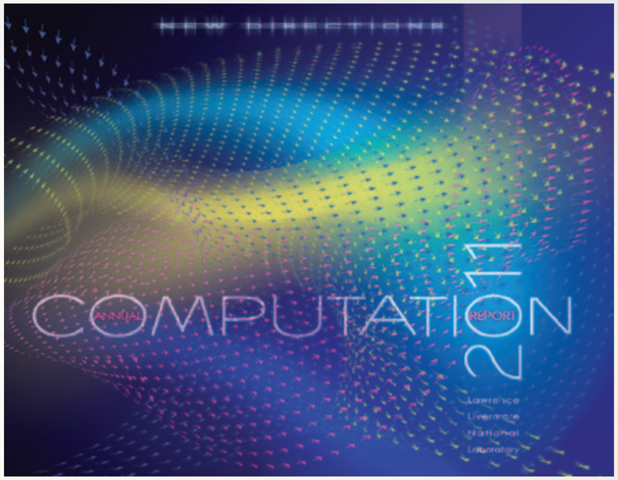

The cover is an artist's conception of a vector field, meant to signify the speed and direction with which the Computation Directorate is expanding its portfolio. Computation is using its expertise in high performance computing as a catalyst to form unique, innovative partnerships that will improve national security and advance U.S. economic competitiveness. (Cover design by Daniel Moore.)

National Nuclear Security Administration

Production Team

LLNL Associate Director for Computation
Dona Crawford

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

Scientific Editors
Susan Springer and Marc Stearman

Art Director
Daniel Moore

Production Designer
Amy E. Henke

Production Editor
Deanna Willis

Editorial Assistance
Caryn Meissner

Proofreader
Pamela MacGregor

Photographers
Michael Anthony and Nancy Rutter

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

LLNL-TR-548352

Contents

Message from the Associate Director

DONA CRAWFORD
ASSOCIATE DIRECTOR, COMPUTATION

From its founding in 1952 until today, Lawrence Livermore National Laboratory (LLNL) has made significant strategic investments to develop high performance computing (HPC) and its application to national security and basic science. Now, 60 years later, the Computation Directorate and its myriad resources and capabilities have become a key enabler for LLNL programs and an integral part of the effort to support our nation's nuclear deterrent and, more broadly, national security. In addition, the technological innovation HPC makes possible is seen as vital to the nation's economic vitality. LLNL, along with other national laboratories, is working to make supercomputing capabilities and expertise available to industry to boost the nation's global competitiveness.

LLNL is on the brink of an exciting milestone with the 2012 deployment of Sequoia, the National Nuclear Security Administration's (NNSA's) 20-petaFLOP/s resource that will apply uncertainty quantification to weapons science. Sequoia will bring LLNL's total computing power to more than 23 petaFLOP/s—all brought to bear on basic science and national security needs. The computing systems at LLNL provide game-changing capabilities. Sequoia and other next-generation platforms will enable predictive simulation in the coming decade and leverage industry trends, such as massively parallel and multi-core processors, to run petascale applications. Efficient petascale computing necessitates refining accuracy in materials property data, improving models for known physical processes, identifying and then modeling for missing physics, quantifying uncertainty, and enhancing the performance of complex models and algorithms in macroscale simulation codes.

Nearly 15 years ago, NNSA's Accelerated Strategic Computing Initiative (ASCI), now called the Advanced Simulation and Computing (ASC) Program, was the critical element needed to shift from test-based confidence to science-based confidence. Specifically, ASCI/ASC accelerated the development of simulation capabilities necessary to ensure confidence in the nuclear stockpile—far exceeding what might have been achieved in the absence of a focused initiative. While stockpile stewardship research pushed LLNL scientists to develop new computer codes, better simulation methods, and improved visualization technologies, this work also stimulated the exploration of HPC applications beyond the standard sponsor base. As LLNL advances to a petascale platform and pursues exascale computing (1,000 times faster than Sequoia), ASC will be paramount to achieving predictive simulation and uncertainty quantification. Predictive simulation and quantifying the uncertainty of numerical predictions where little-to-no data exists demands exascale computing and represents an expanding area of scientific research important not only to nuclear weapons, but to nuclear attribution, nuclear reactor design, and understanding global climate issues, among other fields.

Aside from these lofty goals and challenges, computing at LLNL is anything but "business as usual." International competition in supercomputing is nothing new, but the HPC community is now operating in an expanded, more aggressive climate of global competitiveness. More countries understand how science and technology research and development are inextricably linked to economic prosperity, and they are aggressively pursuing ways to integrate HPC technologies into their native industrial and consumer products. In the interest of the nation's economic security and the science and technology that underpins it, LLNL is expanding its portfolio and forging new collaborations. We must ensure that HPC remains an asymmetric engine of innovation for the Laboratory and for the U.S. and, in doing so, protect our research and development dynamism and the prosperity it makes possible.

One untapped area of opportunity LLNL is pursuing is to help U.S. industry understand how supercomputing can benefit their business. Industrial investment in HPC applications has historically been limited by the prohibitive cost of entry, the inaccessibility of software to run the powerful systems, and the years it takes to grow the expertise to develop codes and run them in an optimal way. LLNL is helping industry better compete in the global market place by providing access to some of the world's most powerful computing systems, the tools to run them, and the experts who are adept at using them. Our scientists are collaborating side by side with industrial partners to develop solutions to some of industry's toughest problems. The goal of the Livermore Valley Open Campus High Performance Computing Innovation Center is to allow American industry the opportunity to harness the power of supercomputing by leveraging the scientific and computational expertise at LLNL in order to gain a competitive advantage in the global economy. HPC, modeling, and simulation can reduce the costly and lengthy cycle of developing and testing multiple prototypes, making it possible to design and manufacture materials and products faster and at a lower cost. As the Council on Competitiveness maintains, "The country that wants to out-compete must out-compute."

LLNL is expanding the national-security application of computational science to include research in energy and climate, network and cyber security, nuclear counterterrorism, defense, and space situational awareness. For example, several Computation Directorate experts are engaged in a partnership, called the Carbon Capture Simulation Initiative, with other national laboratories, industry, and academic institutions. This partnership helps the private energy sector use state-of-the-art computational modeling and simulation tools to accelerate the commercialization of carbon capture technologies from discovery to development, demonstration, and ultimately the widespread deployment to hundreds of power plants. By developing a comprehensive, integrated suite of validated science-based computational models, this initiative will provide simulation tools that will increase confidence in designs, thereby reducing the risk associated with incorporating multiple innovative technologies into new carbon capture solutions.

In a project aimed at reducing the United States' dependency on fossil fuels, Livermore computational experts teamed with Navistar, Inc., NASA's Ames Research Center, and the U.S. Air Force to develop and test devices for reducing the aerodynamic drag of semitrucks. The devices were designed by considering the tractor and trailer as an integrated system and taking into account operational requirements. If these simple devices were deployed across the U.S. trucking fleet, they would increase fuel efficiency by as much as 12 percent and could prevent 36 million tons of CO_2 from being released into the atmosphere annually—approximately the same amount of CO_2 that is emitted from four 1-GW power plants every year.

To facilitate these and future collaborations with industry, Computation made a significant change to its unclassified HPC network (or Open Computing Facility) and the resources associated with it, creating a separate network segment off the Laboratory's ESnet connection. The new unclassified HPC network segment, called the HPC Enclave, is historic in that it allows formerly prohibited sensitive country foreign nationals (SCFNs) access to LLNL's HPC resources. This project required LLNL to reconsider the longstanding policy that prohibited SCFN access to HPC resources. With the appropriate safeguards in place, SCFNs can now access LLNL's HPC resources in the same manner as other authorized users.

LLNL has spent 60 years taking its in-house expertise and basic science concepts and rapidly translating them into new technologies or scientific solutions that address real-world problems. It is what we do best. If the U.S. is to continue leading the $25.6 billion HPC industry, it is urgent that we push forward, develop next-generation technologies, and expand the application base to new areas of research. By leading the way in these new directions, LLNL and the Computation Directorate will be instrumental in helping the U.S. maintain its global science and technology leadership and the economic advantage that comes with it.

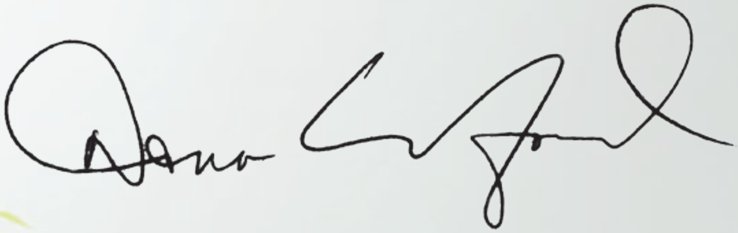

An Award-Winning Organization

The stories in this annual report present a cross-section of Computation's accomplishments in research, high performance computing (HPC), software applications, and information technology and security. In addition to the projects highlighted in the report, several Computation personnel and projects received prestigious external recognition in 2011. Some of the notable accomplishments are featured in this section.

Top Technological Innovation

A team of computer scientists won an R&D 100 Award for developing the Stack Trace Analysis Tool (STAT), a highly scalable debugging tool for identifying errors in computer codes running on supercomputers of 100,000 processor cores and more. STAT allows users to quickly locate in their code the most challenging bugs that emerge only at extreme scales and to get critical applications back up and running.

Dubbed the "Oscars of Invention," R&D 100 Awards are given each year for the development of cutting-edge scientific and engineering technologies with commercial potential.

Presidential Early Career Award

Computer scientist Greg Bronevetsky received a Presidential Early Career Award for Scientists and Engineers (PECASE) for helping advance the state of the art in HPC. 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. Bronevetsky was one of 94 early career scientists and engineers to be recognized in 2011.

Greg Bronevetsky

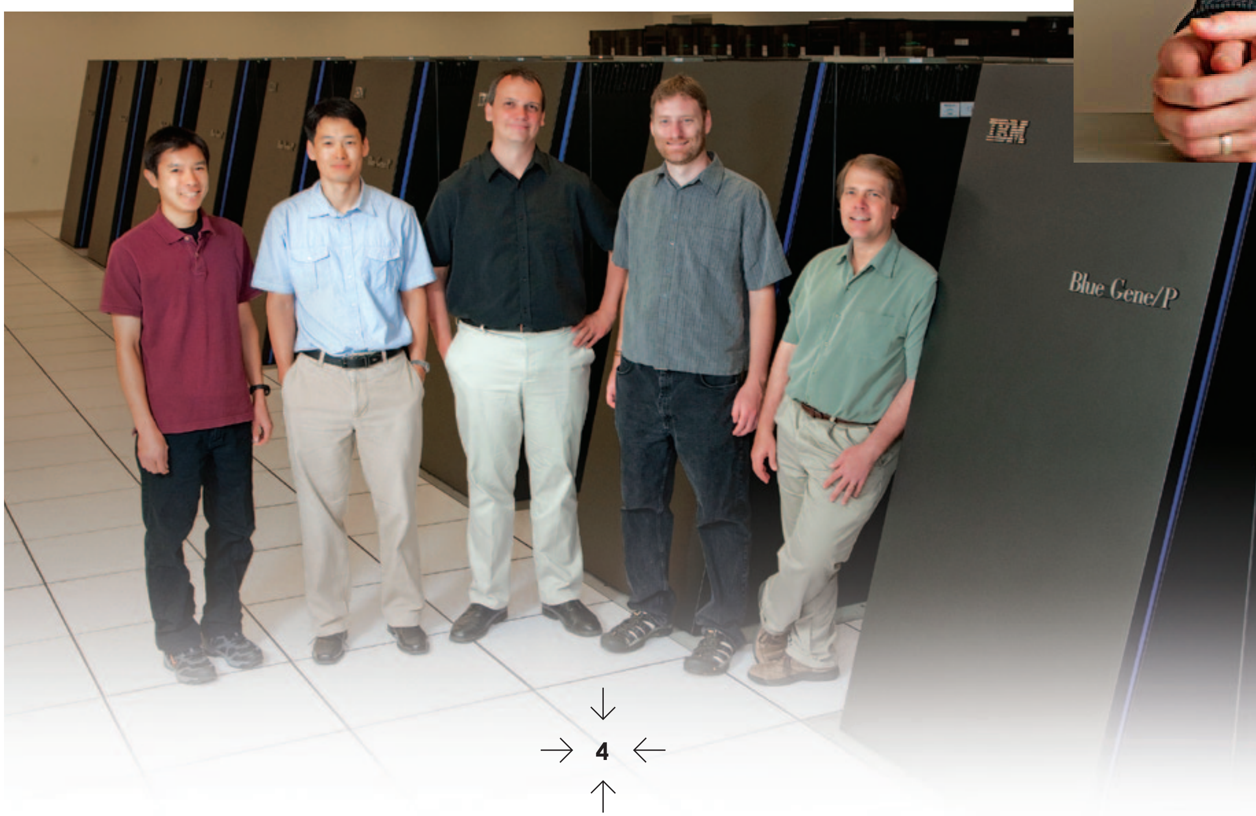

The Livermore development team for the Stack Trace Analysis Tool (STAT): (from left) Greg Lee, Dong Ahn, Martin Schulz, Matthew LeGendre, and Bronis de Supinski.

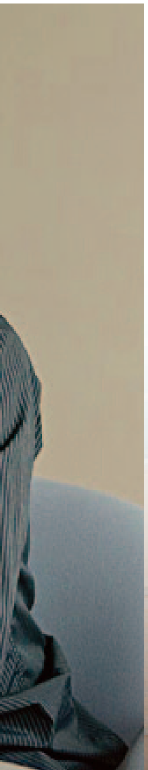

Bronevetsky has dedicated his scientific career to ensuring that the increasing power, size, and complexity of the supercomputers critical to national security research and scientific discovery do not come at the expense of reliability. The methodologies he is developing to study the effects of the hardware failures that are inevitable on supercomputers with millions of components are likely to influence the design of next-generation high performance computers and the software applications that run on them. (See page 28.)

A Record-Setting Scholar Program

The Computation Directorate hosted a record number of students (137) and faculty (27) representing 79 colleges and universities and involving more than 70 different projects. The number of Computation students is 26% of the overall LLNL student population. Students represent the future of computing at LLNL; their success will be our legacy.

Computation's Institute for Scientific Computing Research hosted 137 students in the 2011 scholar program.

Truck Aerodynamics Innovations

The Heavy Vehicle Aerodynamic Drag Project, led by Kambiz Salari, was awarded an HPC Innovation Excellence Award by the International Data Corporation for its use of modeling and simulation to find practical ways to reduce aerodynamic drag and improve the fuel efficiency of the tractor trailers ubiquitous on America's highways. (See pages 32–33.)

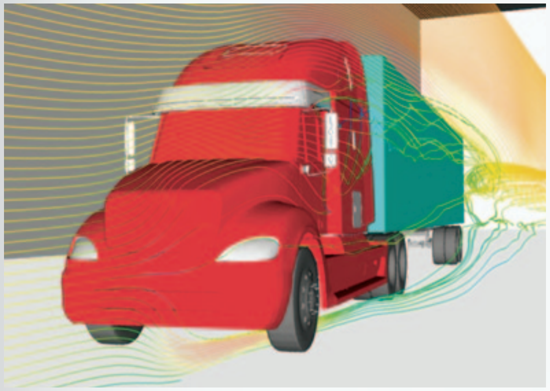

High performance simulation codes help accurately capture critical flow regions around a Class 8 tractor-trailer.

Early Career Research Grant

Peter Lindstrom was one of only 67 researchers selected to receive a U.S. Department of Energy Office of Science Early Career Award. He will receive $2.5M over five years to conduct his research, which will focus on alleviating the data-movement bottleneck in extreme-scale computing to accelerate numerical simulation and data analysis. Lindstrom's goal is to develop tools that will greatly reduce data movement while increasing performance and reducing power consumption on next-generation massively multicore computer architectures.

Peter Lindstrom

"LEEDing" the Way in Energy and Environmental Design

The Computation Directorate is committed to continuously reducing energy costs and finding operational efficiencies in its computing facilities. Computation's Building 451 received a Leadership in Energy and Environmental Design (LEED) silver certification under the U.S. Green Building Council rating system. LEED is an internationally recognized green building certification system. It provides third-party verification that a building or community was designed and built using strategies aimed at improving performance in energy savings, water efficiency, carbon dioxide emissions reduction, and other factors.

Building 451, which was built in 1980 and upgraded in 1998, houses some of the Laboratory's largest HPC clusters. It is the second Computation facility to become LEED-certified; the TSF received gold certification in December 2009.

The "Other" 500 Lists

The BlueGene/Q Prototype II, which will become the Advanced Simulation and Computing Program's Sequoia supercomputer when it is delivered to LLNL in early 2012, won first place on the Graph 500 list in November. BlueGene/Q was able to traverse more than 254 billion graph edges per second, more than two and one-half times more edges per second than the next machine on the list.

LLNL had multiple entries on this year's Graph 500 lists, including several entries submitted by Maya Gokhale and Roger Pearce. As yet lacking the TOP500's visibility and glory, the data intensive computing capabilities represented by the Graph 500 are of growing importance to LLNL and the HPC community. A machine on the top of this list can quickly and efficiently analyze huge quantities of data to find the proverbial needle in the data haystack.

The BlueGene/Q system was also ranked number 1 on the Green500 list of the world's most energy-efficient computers in both June and November. Energy efficiency remains one of the greatest challenges for next-generation exascale supercomputers. It was a critical factor in selecting Sequoia and will continue to be of paramount importance as LLNL plans for more powerful machines in the future.

New Directions

and New Vision for High Performance Computing

JOHN GROSH GROSH1@LLNL.GOV

Lawrence Livermore National Laboratory (LLNL) is developing novel approaches in high performance computing (HPC) to improve national security and advance U.S. economic competitiveness. The Laboratory, and in particular the Computation Directorate, will focus on partnering with industry, foreign governments, and nontraditional government sponsors to achieve this goal. These relationships will drive technology advancement, speed development of new applications, and help shape an "ecosystem" (i.e., the environment, facilities, and processes) to cultivate further external collaborations.

As part of the Department of Energy's (DOE's) Exascale Initiative, LLNL plays a key role in advancing and deploying the next generation of high-end computing architectures and science and engineering software applications. LLNL will continue to lead the development of enhanced supercomputer architectures by conducting research in computer tools and algorithms and creating new predictive simulation and analytic applications. The Computation Directorate will also conduct research in programming tools for million-way concurrent systems, uncertainty quantification, and distributed data analytics to enable improved predictive simulation capabilities.

LLNL plans to expand the reach of its computing applications and to communicate the value of HPC on a national scale. Over the past three years, the Laboratory has significantly increased its involvement in applied energy applications, such as those for enabling carbon capture and sequestration, modeling and simulating the electric grid, researching nuclear energy, and building energy efficiency. In the future, LLNL intends to leverage the technologies developed under the National Nuclear Security Administration's (NNSA's) Advanced Simulation and Computing Program and the DOE Advanced Scientific Computing Research Office to create applications focused on energy technologies and industrial engagements. The Laboratory will also continue to promote HPC as a critical element for U.S. economic competitiveness, with partners such as the Council of Competiveness, through workshops, conferences, and seminars. Finally, the Laboratory has hosted open competitions for computing resources to spur development of HPC applications and promote industrial innovation.

Infrastructure is key to technological advancement and economic growth. Computation is leading an effort to improve external collaborators' access to the Laboratory through the Livermore Valley Open Campus (LVOC). This area will include the High Performance Computing Innovation Center for building advanced modeling, simulation, and analytics for HPC systems. In addition, the Laboratory's computing and networking infrastructure has been modified to support unclassified processing with industry and other external organizations, while preserving cyber security for critical intellectual property and other protected information.

New approaches in computing will be essential to the Laboratory for addressing important national needs in science and technology. This section of the Computation Annual Report highlights LLNL's initial efforts toward realizing these advances.

HPC Innovation Center Sharpens America's Competitive Edge

The High Performance Computing Innovation Center (HPCIC) opened in June 2011 as an outreach by LLNL to U.S. industry for developing, proving, and deploying high-impact HPC solutions. The HPCIC offers industrial partners a formidable combination of computing resources and technical expertise from across the Laboratory and other institutions, organized in an accessible, business-friendly environment. Partner engagements focus on using HPC to solve complex science and technology problems and help companies gain a competitive edge in global markets.

FRED STREITZ STREITZ1@LLNL.GOV
HPCINNOVATIONCENTER.LLNL.GOV

Progress in 2011

The global economy rewards ingenuity, and aptly applied HPC technologies drive innovation and improve competitiveness. Using HPC, companies can shorten design cycles and reduce expenses. However, barriers such as the high cost of hardware, unavailability of software, and lack of accessible expertise have impeded U.S. industry's adoption of HPC.

To help partners systematically overcome these significant barriers to using supercomputers, the HPCIC orchestrates projects and high-value resources to address three areas of common need:

- Rapid assembly and leveraging of world-class talent.
- Cost-effective access to supercomputing resources, services, and support.
- HPC software applications that solve high-impact, real-world business problems.

Thanks to decades of significant achievements by LLNL personnel and U.S. government investment, HPCIC partners can access hardware and computational science expertise without the time delays, costs, and management overhead that would be necessary to develop broad in-house capabilities. Additionally, partners gain access to world-class Livermore Computing resources and services 24/7/365, without incurring the front-end expenses of procuring, siting, staffing, and supporting highly complex supercomputing environments. The HPCIC also enables project teams to develop software solutions by leveraging open source codes, common frameworks, libraries, processes, and methodologies developed and assembled by the Laboratory over time.

The HPCIC pairs external collaborators with some of the world's foremost HPC practitioners to define and devise actionable solutions for industry's toughest problems. Today, LLNL business developers are tailoring

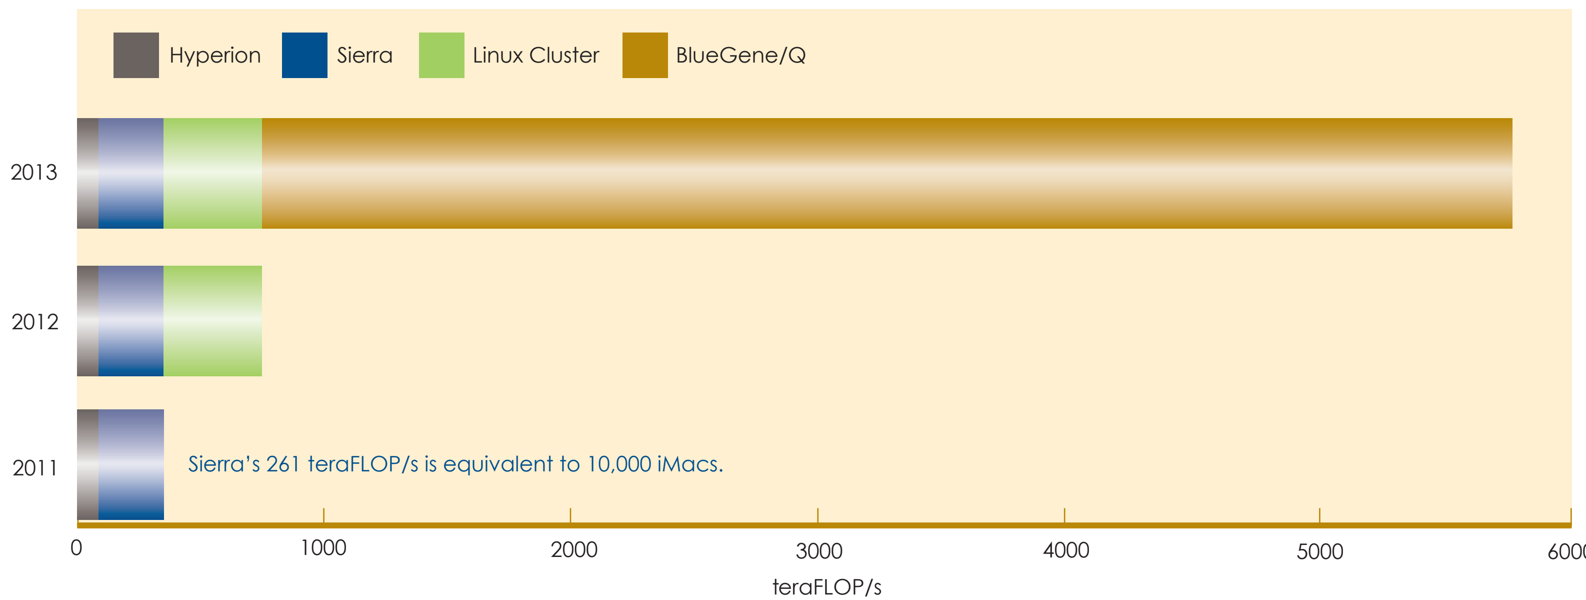

The High Performance Computing Innovation Center (HPCIC) provides access to multiple high performance computing solutions and will house one of the world's largest supercomputers in 2013. This graph shows the current and future computing resources for HPCIC partners.

engagements with the HPCIC to an individual partner's needs, based on the work to be performed and sources of funding. Principal investigators and program managers are working with partners to develop and execute mutually agreeable scopes of work to employ the expertise of the Laboratory's scientists and engineers, along with appropriate HPC resources. The HPCIC is ready to offer a range of services spanning the entire software development lifecycle, including:

- Application development and enhancement, algorithm optimization, and scaling of existing codes to improve fidelity, resolution, or performance.
- Coding infrastructure development to address data movement, visualization, debugging, testing, usability, and other HPC software best practices.
- Training and education to address partners' workforce development needs.

The HPCIC is LLNL's first facility in LVOC, the new general access area designed for expanding interactions and facilitating collaboration with partners from industry and academia. Its facilities include office space, networked conference rooms, classrooms, and open collaboration areas. Large enough to host multiple project teams and events simultaneously, the HPCIC uses key protective measures to allow project work groups to share data during collaborative work while protecting against unauthorized access.

Since June 2011, the HPCIC has hosted more than 450 events involving more than 5,200 people. These events have included discussions with vendors and potential employees, workshops, conferences, and training classes.

The HPCIC is designed for remote access as well. A large classroom is equipped with distance-learning technologies, and conference rooms and offices include videoconferencing systems to maximize face-to-face communications. Remote users benefit from more than two decades of LLNL's experience providing a remotely accessible, secure, and balanced HPC computational environment suitable for industrial partnerships.

The High Performance Computing Innovation Center is located in the Livermore Valley Open Campus. Shown here are the collaboration areas in use.

The HPCIC is a means of delivering educational resources to U.S. industry in parallel programming, parallel application tools, and use of HPC systems. Through lectures and hands-on training, introductory-level workshops help new users improve computer productivity and minimize challenges typically encountered with complex supercomputing systems. Other workshops are targeted toward experienced users and cover a range of topics related to new technologies, performance and programming tools, and cross-platform training. HPCIC partners have access to online training resources provided by the Computation Directorate.

Looking forward, the Laboratory plans to host many events at the HPCIC to foster industry collaborations, HPC community development, ecosystem expansion, and, most importantly, innovation.

HPC Enclave Improves Partner Access to Capabilities

DOUGLAS EAST EAST1@LLNL.GOV

LLNL recently took significant steps to increase its unclassified collaborations with industry. In concert with Sandia National Laboratories in California, LLNL created LVOC, an open campus area between the two institutions. One of the campus's centerpiece facilities, the HPCIC, opened in June 2011. The HPCIC's express purpose is to enable collaborations with U.S. industry partners and to provide them advanced HPC and scientific simulation capabilities by which to boost U.S. economic competitiveness.

One challenge to this initiative is that many leading scientists and engineers are not U.S. citizens but rather foreign nationals, including citizens of countries on DOE's sensitive country list. These sensitive country foreign nationals (SCFNs) have historically been prohibited from accessing the Livermore Computing (LC) unclassified HPC environment. It was recognized that prohibition is no longer tenable if expanded partnerships with industry are to be realized.

To allow SCFNs access to LC's unclassified HPC resources, LC launched an effort in early 2010 that resulted in the design of the HPC Enclave. Its implementation began in earnest in early 2011. Key features of the HPC Enclave provide the necessary protections for sensitive unclassified information and also improve the overall security posture of the LLNL Enterprise (or "Yellow") Network. The Laboratory Director formally approved SCFN use of LC's unclassified HPC resources, consistent with NNSA computer security policy, in November 2011.

Progress in 2011

This year, LC focused on implementing the HPC Enclave. At the heart of its architecture are zones that correspond to data sensitivity and other access considerations. Three zones are defined, but more can be created depending on future needs. The majority of HPC resources exist in the collaboration zone (CZ), which SCFNs are allowed to access. Most forms of unclassified data, including unclassified controlled information, are allowed in this zone. However, Unclassified Controlled Nuclear Information (UCNI) and Naval Nuclear Propulsion Information (NNPI) are prohibited. All forms of unclassified data, including UCNI and NNPI, are allowed in the restricted zone (RZ), which is primarily intended for application development purposes. Infrastructure, such as file system and authentication servers, and many other systems that provide services for the HPC environment are located in the infrastructure zone. Only LC staff has access to this zone.

Principal features of the HPC Enclave include interzone sharing of major infrastructure services, independent zone authentication methods, and Enterprise Network autonomy. Interzone sharing allows expensive resources to be shared between the CZ and the RZ. In particular, archive storage is cost prohibitive to implement consistently in multiple zones. To avoid these high costs, LLNL developed a common underlying archive storage infrastructure for the RZ and the CZ but deployed separate access name spaces for each zone. LLNL also considered interzone sharing for the Lustre parallel file systems, but ultimately determined that vulnerabilities existed in the Lustre shared model that constituted too much risk.

Independent zone authentication methods for the CZ and the RZ ensure that a security compromise of one zone will not lead to a compromise of the other. The CZ retained the existing RSA-based one-time password (OTP) method. A new CryptoCard-based OTP method provides access to the RZ via a gateway service (the RZ hosts still use the RSA OTP method). Eventually, the RZ gateway service will be eliminated and the CryptoCard OTP will be implemented directly on all RZ hosts.

Enterprise Network autonomy was achieved by relocating the HPC Enclave outside the core Enterprise Network, while retaining equivalent protection mechanisms that are relevant to HPC network traffic. Thus, it can be considered an extension of the Enterprise Network. The chief advantage of this relocation is that the HPC Enclave can more rapidly exploit the availability of increased wide-area network bandwidth. The relocation also removes HPC collaborators' visibility into the Enterprise Network, which they obtained through the LLNL Virtual Private Network (VPN) service. Remote access to the CZ does not require LLNL VPN access.

The HPC Enclave was implemented in three phases, the first of which was completed on June 29, and primarily involved implementing Enterprise Network autonomy. Phase 2, completed on August 30, established the zones including interzone sharing, independent authentication methods, and other features. The time frame between Phase 2 and Phase 3, when SCFNs would have been allowed to access the CZ, was intended to ensure the technical implementations were working as designed. During this time, issues were identified that had to be resolved before Phase 3 could be completed, resulting in an approximate one-month delay before SCFNs were finally allowed access to the CZ on November 10. Access was provided after a review of the effort, risks, and mitigations, and approval by the Laboratory Director.

The HPC Enclave was a highly complex effort. Its design and validation involved staff from LC, Cyber Security, and other LLNL programs, and the majority of LC staff took part in the overall implementation. The HPC Enclave is an innovative and comprehensive approach to addressing the vitally important issue of allowing SCFN collaborators from industrial partners and Laboratory programs access to LC's preeminent unclassified HPC environment. As a result of this work, LLNL is now positioned to engage in substantive partnerships through the HPCIC that are aimed at advancing the use of supercomputing in U.S. industry.

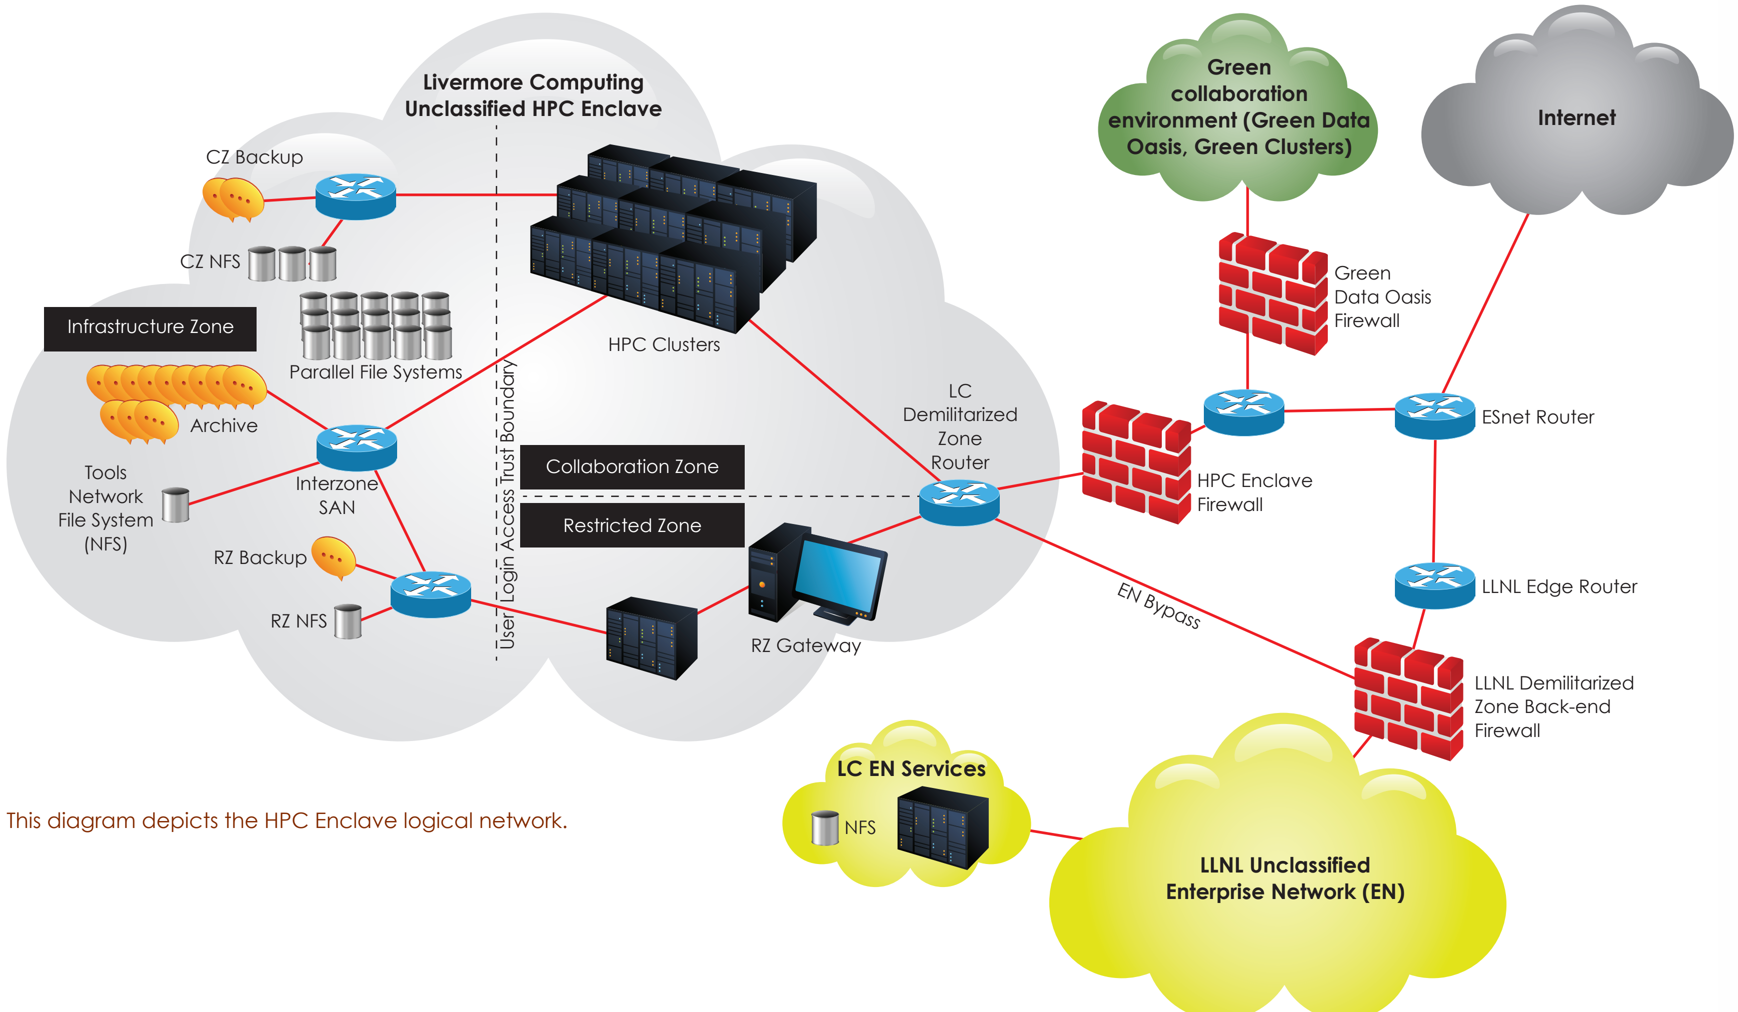

This diagram depicts the HPC Enclave logical network.

New Collaborations Focus High Performance Computing on Energy

JOSH GROSH GROSH1@LLNL.GOV

Over the past two years, the Computation Directorate has embraced energy systems as an expanding application area for HPC and modeling and simulation. The Directorate is using computational technologies developed in the DOE Advanced Scientific Computing Research and NNSA Advanced Simulation and Computing programs to create capabilities for a variety of energy applications. These innovative computing ventures involve new collaborations with industry that require the Laboratory to think "out of the box" to strategize engagement approaches. Three projects are particularly notable: Using Computational Fluid Dynamics to Improve Building Energy Efficiency, Promoting High Performance Computing as an Enabler for Industrial Innovation, and Clean Emissions through Advanced Simulation.

Progress in 2011

More than 40% of all U.S. energy is consumed by commercial and residential buildings. Despite significant investments and research, the energy efficiency of buildings has shown minimal improvement over the last several decades. To a large extent, this is due to the highly fragmented nature of the building industry, which prevents a whole-system approach to design, retrofit, and operation. In 2010, DOE funded a research "hub" to develop and demonstrate technologies that dramatically improve the energy efficiency of buildings by as much as 50%. The Greater Philadelphia Innovation Cluster for Energy Efficiency Buildings (GPIC) hub, led by the Pennsylvania State University, focuses on retrofits of buildings that are less than 100,000 square feet, which comprise a significant portion of the commercial building stock in the U.S. The goal is to provide the building community—from architecture engineering firms to builders to operators—with the analytic tools they need to design and build more energy-efficient structures.

The GPIC team is developing computational tools to understand airflow, heat transfer, and lighting in buildings and vastly improve building control systems. LLNL's role in the project is to develop high-fidelity computational fluid dynamics simulations that model highly resolved airflow and heat transfer within buildings. Widely used building energy simulations codes, such as EnergyPlus, combine physical systems characteristics to enable rapid analysis. Unfortunately, this approach does not provide a detailed look at spatial and temporal issues associated with the air and heat dynamics of a building or the complex interactions of the control systems. A team of LLNL computer scientists and mathematicians is enhancing the incompressible Navier–Stokes solver CGINS (a program that can solve incompressible fluid-flow problems in complex geometry in two and three dimensions using composite overlapping grids) to provide grid generation for buildings from computer-aided design programs and integrate sensor models into complex flow simulations. The scientists are working with United Technologies Corporation, Purdue University, and Virginia Polytechnic Institute and State University to create approaches for using CGINS to develop fast reduced-order models for building design and operations.

Computation researchers are also participating in a project to promote HPC as an enabler for industrial innovation. HPC, modeling and simulation, and advanced analytics are considered to be a—sometimes *the*—key enabler for commercial competiveness. Many countries, including China, are rapidly adopting HPC as a means of increasing their gross national product, thus sparking an international race where a computing advantage is a main differentiator for solving problems of great complexity that require less time to solution. In the U.S., industry has historically lagged behind academia, DOE, and other parts of government in using HPC. However, a new initiative, the hpc4energy incubator, is poised to help DOE and LLNL motivate industry and introduce companies to the benefits of HPC.

The Laboratory issued a call for proposals in October 2011 for energy companies to collaborate with LLNL experts and HPC specialists on clean-energy projects. This initiative marks a new approach to developing public–private partnerships and creating an ecosystem of energy technology and HPC to help the nation pioneer innovative solutions. The Laboratory solicited proposals in five areas: Building Energy Efficiency; Carbon Capture, Utilization, and Sequestration; Liquid Fuels Combustion; Nuclear Energy; and Smart Grid and Power Storage. Of the 30 letters of intent that were received, the Laboratory awarded six proposals in March 2012.

Computation is also involved in a project to use advanced simulation to reduce harmful elements in emissions and improve carbon capture technologies, which are key to producing clean energy. For example, carbon capture technologies are used to substantially reduce carbon emissions in coal-burning electric power plants. Two challenges are hindering wider adoption and deployment: First, carbon capture technologies can be expensive, increasing the cost of a power plant's electricity up to 80%. Second, it can take 30 years to develop new processes, scale the technologies from the laboratory to pilot facilities, and then deploy them. To address these challenges, DOE launched the Carbon

Applications

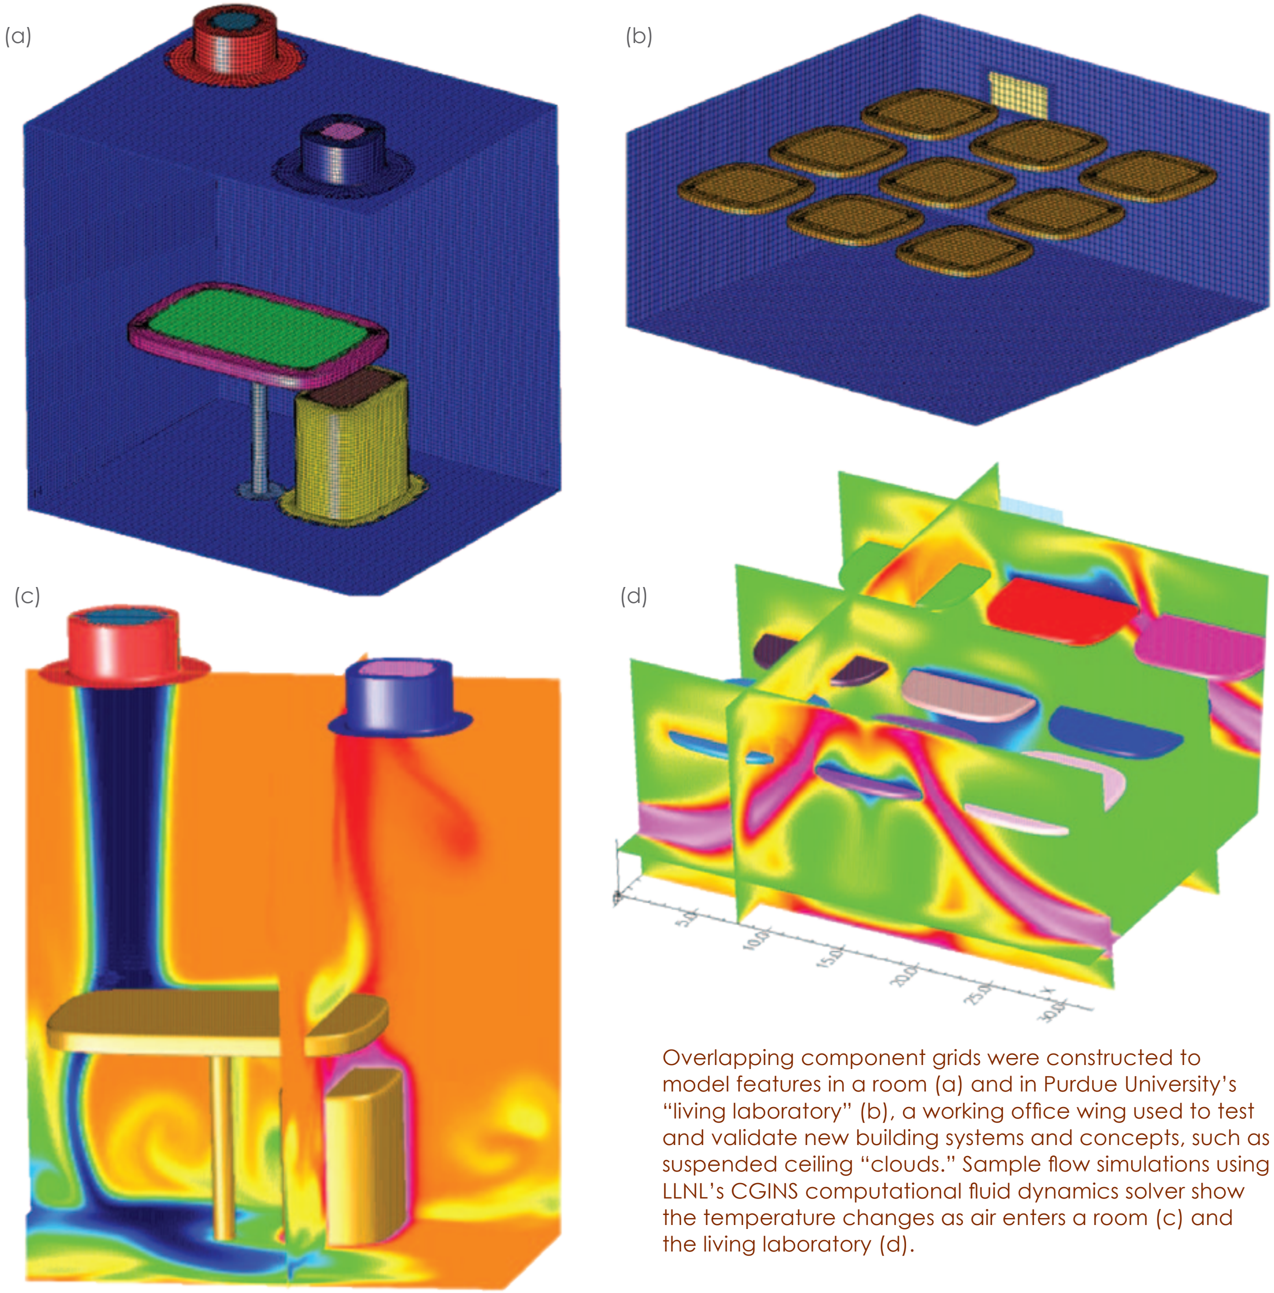

Overlapping component grids were constructed to model features in a room (a) and in Purdue University's "living laboratory" (b), a working office wing used to test and validate new building systems and concepts, such as suspended ceiling "clouds." Sample flow simulations using LLNL's CGINS computational fluid dynamics solver show the temperature changes as air enters a room (c) and the living laboratory (d).

Capture Simulation Initiative (CCSI), a partnership between five DOE laboratories, academia, and industry. CCSI researchers are building a simulation code toolset that will accelerate new carbon capture technologies to market. Once realized, this initiative will enable rapid screening of devices and processes and reduce time to design, technical risk, and cost.

The CCSI Toolset is an integrated modeling environment that uses a science-based approach to build predictive models for carbon capture. LLNL is contributing to this project in the areas of uncertainty quantification, software quality assurance and testing, modeling tools, and risk analysis. This year, LLNL computational scientists integrated the uncertainty quantification tool PSUADE (Problem Solving environment for Uncertainty Analysis and Design Exploration) and Amazon cloud computing into the toolset. These enhancements have allowed CCSI researchers to run thousands of computing tests to improve sensitivity studies and model calibrations. This project requires substantial effort to integrate software with commercial tools such as Aspen Plus (used for process modeling) and represents a new venture in software development at LLNL.

Computation Pursues a Transition to Exascale

The Laboratory's many successes in high performance computing have given LLNL scientists a deep understanding and appreciation of the potential contributions computing can make to scientific and engineering challenges. However, hardware technologies in the next decade will require architectures that are very different than the architectures that exist today. Without an immediate and clear strategy and investment, the next generation of computers will be unaffordable and arduous to program and operate. To solve these problems, Computation experts are collaborating with other national laboratories and DOE headquarters on various projects funded by the Laboratory Directed Research and Development Program, the Advanced Scientific Computing Research Program, and NNSA. The common goal is to plan and pursue the nation's transition to exascale.

TERRI QUINN QUINN1@LLNL.GOV
LORI DIACHIN DIACHIN2@LLNL.GOV

Progress in 2011

Scientists from seven DOE national laboratories formed a team, called E7, to enable high performance scientific computing on the new generation of computers and achieve exascale computing by 2020. LLNL brings to the E7 team a unique expertise gained from the series of BlueGene research and development (R&D) and acquisition projects (BlueGene/L, BlueGene/P, and BlueGene/Q). In 2011, LLNL's team of computer scientists and procurement specialists helped E7 prepare proposals for a national exascale plan. The proposals included implementation plans, budgets, risk analyses, R&D investment plans, and procurement strategies. LLNL also helped analyze responses to an E7 exascale Request for Information in the areas of system technologies, programming models, tools, data management, processors, interconnects, and contracting terms and conditions (e.g., management of intellectual property). LLNL's long-term partnership with IBM on BlueGene is proving to be a model for DOE and industry collaborations related to exascale.

In parallel with the E7 venture, LLNL computer scientists, applied mathematicians, and data scientists are working to better understand the intricacies, challenges, and opportunities associated with exascale computing. As with every major change in computer architecture, transitioning to exascale poses several unavoidable scientific and technological challenges. Developing and running application software will be complicated by the potential for billion-way parallelism. New programming models that account for inter- and intra-node parallelism will be needed as will the ability to handle faults at multiple levels. Although exascale computing comes with many challenges, it will also provide great opportunities, such as the ability to model physical phenomenon at unprecedented levels of fidelity and with increased confidence and accuracy. In anticipation of these opportunities, Computation's R&D portfolio has expanded to include application-level technologies such as runtime libraries, math libraries, programming models, tools, frameworks, algorithms, as well as operations-level requirements such as parallel file systems, resource management, system management, and facility engineering.

Computation exascale research highlights of 2011 include:

- An R&D 100 Award for collaborative work to develop statistical debugging techniques. (See page 4.)
- International chair of the OpenMPI forum, leading the development of new standards and technologies that will ease application programming on multicore architectures.
- A DOE Office of Science Early Career Award for research in large-scale data streaming and compression techniques. (See page 6.)
- Advanced research in collaboration with LLNL programs to understand code performance at the exascale, develop algorithms that will minimize communication and data movement, explore new paradigms for using hybrid techniques for uncertainty quantification, and experiment with persistent memory paradigms for data-intensive applications in intelligence and biosecurity.

Through these and related activities, LLNL computational scientists are positioning DOE and the nation for a successful transition to exascale computing.

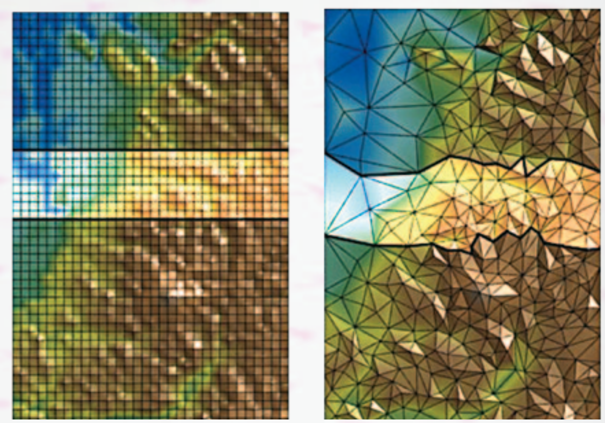

Windowed stream processing of structured (left) and unstructured (right) grids reduces latency and keeps data movement to a minimum, both of which will be critical in analyzing large data sets associated with exascale computing. This work was instrumental in demonstrating proof of concepts that led to Peter Lindstrom's Early Career Award from the Office of Science. (See page 6.)

High Performance Computing

Solutions Remain Flexible and Agile

TERRI QUINN QUINN1@LLNL.GOV

Computing is integral to all of Lawrence Livermore National Laboratory's (LLNL) missions. As the Laboratory's mission needs change and evolve, the Computation Directorate must respond by predicting and developing the high performance computing (HPC) advances necessary to meet our users' demands. Computation's strategy is to simultaneously harness today's technology while aggressively pursuing technologies for the future. The HPC organization within Computation provides leading-edge, high-end simulation capabilities for modeling complex systems, but we anticipate that LLNL researchers will need computer architectures that can model even larger complex systems with greater accuracy. Moreover, we anticipate a demand for computer architectures that can quickly search and analyze massive amounts of data.

One path begins with the Sequoia supercomputer. Delivery of the hardware began in January 2012 and will take several months to complete. Sequoia is unprecedented in terms of speed, size, and power. Thus, operating and programming Sequoia will be an immense challenge, and we will be confronted with problems the computing community has never seen before.

Even while Sequoia is being integrated, Computation researchers are working on projects a generation beyond Sequoia. We have developed a portfolio of research projects to prepare for the inevitability of computers capable of performing one quintillion (10^{18}) calculations per second (50 times faster than Sequoia). We will add projects to our portfolio in 2012. In partnership with world-class HPC companies, we are exploring potential computer technologies and testing prototype hardware in LLNL environments and with LLNL simulations.

LLNL researchers also need data-centric computing solutions to search and analyze large data sets. To this end, we are collaborating with scientists to design computing systems to solve problems where data is the primary challenge, whether it is the complexity, size, or rate of the data being processed. By leveraging our HPC expertise, we were able to deploy a few exploratory data-centric solutions such as computers with large amounts of memory and file systems optimized for data analysis. These systems are now being used for bioinformatics and national security applications. For LLNL's Global Security organization, we are developing a shared distributed computing system for large data sets. These new systems are just the beginning. As requests for data-centric computing increase and evolve, we will continue to advance and adapt to fulfill our users' needs.

Computation embraces the idea that HPC propels scientific innovation. The Directorate's efforts are and will continue to focus on delivering superior HPC solutions today while growing and nurturing an expanded constituency of HPC users.

Computation Prepares for a Petascale System

KIM CUPPS CUPPS2@LLNL.GOV

Predictive simulation requires scientists to build better science models and quantify the margins of uncertainty in their integrated calculations. The Sequoia supercomputer, which began arriving at LLNL in January 2012 and is deployed for the Advanced Simulation and Computing (ASC) Program, provides the capability to address these challenges. Sequoia is the first advanced technology platform that will tackle production environment simulations. Previously, advanced architectures were used for science simulations but not for large-scale multiphysics codes run in a production environment. Sequoia will have 20-petaFLOP/s of peak performance, 1.6 petabytes of memory, and more than 1.5 million cores. This proliferation of cores is two orders of magnitude greater than any other LLNL production simulation HPC resource and presents a daunting code-porting challenge. Significant work was accomplished in 2011 to prepare LLNL codes to move to Sequoia, and the application performance benchmark results are encouraging.

In addition to ongoing collaboration on all facets of the Sequoia system procurement, 99% of the Sequoia computer room floor construction was completed in 2011. This project was a significant undertaking and required installation of electrical wiring and consolidation of receptacles; a liquid-cooling infrastructure; an entire subfloor infrastructure to support 96 4,000-pound racks; and steel trays in the north–south and east–west directions to hold the network cabling.

Progress in 2011

Livermore Computing (LC) had a very busy year preparing for the arrival of Sequoia. Progress was significant both locally and with our partners at IBM Watson and IBM Rochester. The Sequoia system poses several challenges for applications to achieve optimal performance. In addition to the Message Passing Interface scalability challenges posed by previous systems, this system will extend scalability requirements to more than 1.5 million cores. The individual nodes of the system will feature higher numbers of processors and virtual threads of execution, which will necessitate the use of programming techniques such as OpenMP, Pthreads, or other means to exploit the full capability of the machine. Hence, much of the code porting and performance benchmarking emphasis is on analyzing OpenMP overhead and performance and investigating transactional memory and speculative execution.

Efforts this year focused on three principle areas. First, various usability issues surfaced via the porting and scaling work on Dawn, the BlueGene/P initial delivery system. Second, efforts were made to gain a detailed understanding of Sequoia hardware and software characteristics through benchmarking on system simulators and early BlueGene/Q hardware. Third, a broad and deep collaboration with IBM focused IBM's efforts on items of importance to LLNL code teams. Educating IBM about why certain elements are important to LLNL ensured that design decisions and trade-offs were made as partners, which was essential to the co-design process for Sequoia.

A team of LC computer scientists focused on the day-to-day requirements for building the latest code version and streamlining regression tests on Dawn. LC assisted the KULL team's efforts to port the LLNL-developed regression Automatic Testing System (ATS). ATS is used to run a suite of approximately 1,000 KULL regression tests. By addressing input/output issues, system scheduling, and software constraints, the completion time for the test suite was improved from 19 hours to less than two hours.

Early BlueGene/Q porting and tuning efforts used several simulators provided by IBM for the platform. Though this effort was limited and tedious, the results provided valuable insight into the potential performance of existing Sequoia applications and technologies presented for the first time on Sequoia. Actual Sequoia hardware, available for application performance work during the latter half of 2011, ranged from 4-rack (4,096-node) systems with limited accessibility to a small but dedicated 32-node machine at IBM Rochester.

The Sequoia benchmark applications, as well as some newer micro-applications, were used to understand the performance characteristics of BlueGene/Q hardware. These applications were designed to serve as proxies for existing or future applications or algorithms. Early results were generated using cycle or feature-accurate simulators provided by IBM. These efforts enabled the porting of the applications and validation of the results projected by IBM and LLNL developers. In addition, early work with CLOMP (an LLNL-developed OpenMP benchmark) and similar benchmarks enabled significant progress in debugging and understanding completely new hardware technologies, transactional memory, and speculative execution.

Sequoia required significant under-floor construction work, including stands, electrical receptacles, and piping for liquid cooling.

The availability of actual hardware helped validate small-scale simulation results as well as larger scaling studies. Two target figure-of-merit (FOM) values are desired for the full-scale Sequoia platform: a base (nonoptimized) and optimized FOM. At the end of 2011, the results of a 4,096-node scaling test indicated that the base run achieved 148% of the projected goal, and the optimized run was 124% of the target for optimized performance. Validation of these results, as well as larger scale runs, will take place when the complete Sequoia hardware is made available. In summary, the porting and optimization of the Sequoia benchmark applications proceeded on schedule with encouraging results. All of the applications reached at least 98% of their minimum targets and substantially more when considering the stretch goals.

Significant effort was also made to prepare LLNL's TSF for Sequoia. Sequoia is the first liquid-cooled machine to be installed in the facility and only the second TSF machine powered with 480 volts. The system is 96 racks in 4,000 square feet of space. Weighing more than 4,000 pounds per rack, the load on the floor is 192 tons—equivalent to the weight of 30 adult elephants.

Considerable space planning was required to design an efficient under-floor infrastructure capable of containing liquid cooling, electrical conduit, network cable trays, and rack structural supports. Ninety-six reinforced steel structural stands were installed to support the racks. The design of the stands had to be changed to accommodate an increased weight of the BlueGene/Q racks due to several late design changes. The facility construction plan had to be completely reworked to minimize the delay and meet scheduled rack installation dates.

LLNL made two significant changes to IBM's original facility specifications. First, LLNL designed innovative in-floor power distribution units to minimize under-floor congestion and reduced the conduit distribution requirements from the first floor to the second floor by a factor of 4. Second, although IBM specified stainless steel pipe for the cooling infrastructure, LLNL pursued an approved polypropylene piping system that saved $1.5M in material costs. The piping will also reduce ongoing operational costs via increased flow capability and reduced losses.

The Sequoia procurement posed many challenges. LC overcame these challenges with a combination of ingenuity, resourcefulness, and fortitude.

IBM employees install the first eight Sequoia racks at LLNL.

Lorenz Simplifies High Performance Computing for Users

JEFF LONG LONG6@LLNL.GOV

Historically, users of HPC centers such as LC required expert knowledge of nearly every aspect of the computer center operation to take full advantage of the immense resources available. This knowledge can take years to develop and, without it, the world of HPC can be intimidating for new users. The goal of the Lorenz project is to improve the accessibility and productivity of HPC users by leveraging modern Web technologies. Users now have more access to information, and tedious or difficult tasks are simplified. The Lorenz Web application suite helps coordinate workflow and facilitates the proper and optimal use of HPC resources.

The Lorenz project has three initial phases: (1) a system dashboard that shows information about the computing center at a glance; (2) a job management element, including batch job submission and monitoring capabilities; and (3) application portals, which provide a high-level interface for defining simulation input.

A new generation of Web technologies made it possible for the Lorenz team to develop applications with the richness usually associated with standalone applications. A modern standards-based Web browser is all that is needed to access the multifaceted Lorenz application suite.

The Lorenz "MyLC" system dashboard offers users a variety of portlets from which to choose, or they can create their own.

Progress in 2011

Lorenz is a sophisticated Web-based framework that serves as a front end to LC resources. Its primary features include a personalized dashboard for each user that provides an at-a-glance overview of all computing center resources; a front end job submission and management capability; and a software framework to extend the convenience of Web interfaces to simulation codes, including support for setting up, launching, steering, and monitoring simulations.

Because the breadth of desktop operating systems used both inside and outside LLNL spans all flavors of Linux, Macintosh, and Windows, the decision was made to use exclusively Web-based technologies in Lorenz. Recent advances in JavaScript programming, along with innovations in the form of HTML5 and CSS3, made the browser the preferred development platform to provide users with a robust application interface that requires nothing more than a modern, standards-based Web browser. This approach provides LLNL and non-LLNL users convenient access to Lorenz and offers significant advantages over conventional applications in terms of rapid deployment and updates.

The first phase of the Lorenz project was to create a system dashboard, known as MyLC. The dashboard gives users a rich set of portal applications, or "portlets," from which to choose. Information that was previously buried behind dozens of commands, files and Web pages is now available in one central location. This information includes account details, bank and disk usage, quotas, cluster utilization, news items, and active jobs on LC systems. Additionally, due to modular and extensible software design patterns, the dashboard allows users and staff members to develop their own portlets so developers can easily share their efforts with other users or work groups.

The second phase of Lorenz was built on the idea that HPC centers are managed around batch processing and facilitating simulation jobs' execution life cycles through several discrete stages: setup, submission, queuing/waiting, running, completion, analysis, and archiving. Lorenz ties all of these pieces together and even integrates them with the dashboard and application portals. In addition to guiding users through the process of defining batch parameters, monitoring job progress, and interacting with jobs and job output, Lorenz allows access to tools and capabilities that were otherwise only available to "power users." Ultimately, Lorenz will help users choose the appropriate resources, such as cluster and file system, given the requirements of a job.

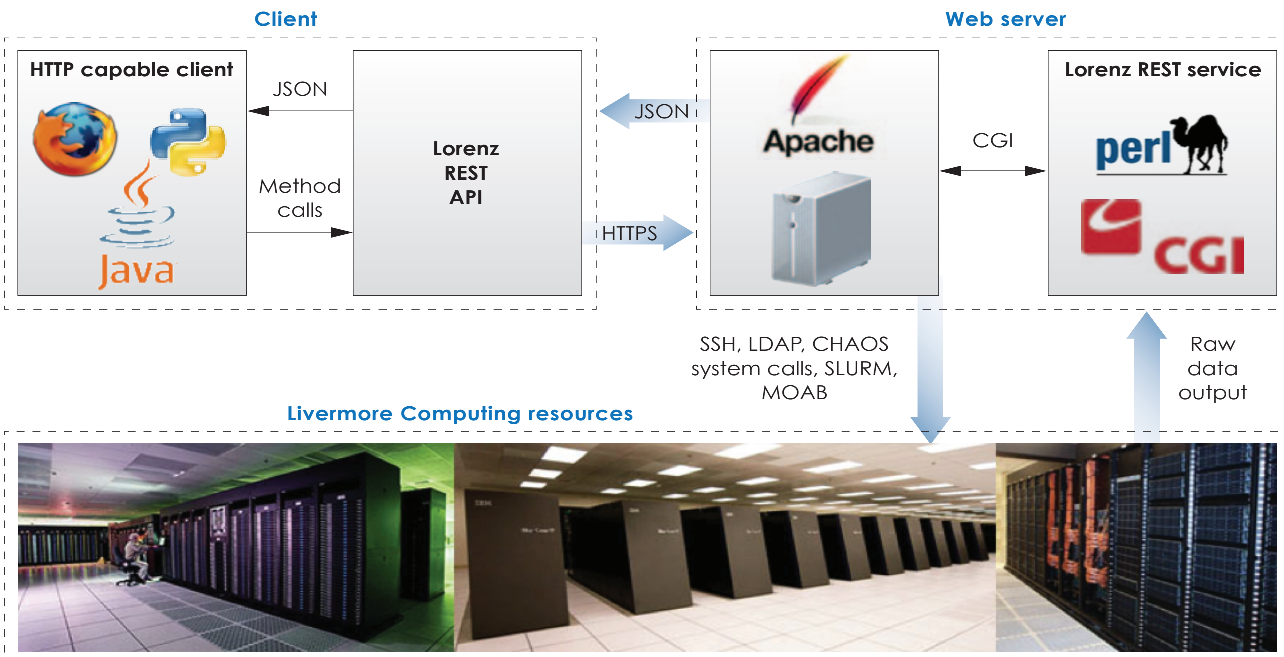

Each Lorenz component is built on top of a Web application programming interface (API) that provides an abstraction layer to hide the details of how the computing center is implemented. This API is called LORA, or Lorenz REST API.

The third major phase of Lorenz is the development of an application portal. Lorenz developers are working with application code groups to provide a simple graphical interface that seamlessly couples job submission and monitoring of their simulations. Users will be able to directly define simulation parameters and launch applications using a powerful Web front end. This capability will benefit current users of LC resources as well as future users of the Livermore Valley Open Campus computing center resources.

Because complex multiphysics codes can have thousands of inputs, providing a general graphical interface for building an input deck from scratch is impractical. However, most users of large applications start with an existing input file and make small changes to it. The Lorenz application portal will feature a process that combines trusted and validated input files with a graphical interface that guides the user to make the relatively minor changes needed for their particular problem. The input files will be in template form and will contain most of the data needed for a particular problem along with parameterization of changeable inputs. Expert users will be able to edit the final input file directly if input changes are needed that were not captured in the graphical interface.

The Lorenz dashboard MyLC was made available in 2011 as a production tool for all LC users. The job management component was released in beta form, and an initial prototype of the application portal was developed. Lorenz allows disparate areas of the computing center to aggregate information in a single user experience, which has allowed additional capabilities to be introduced into Lorenz, including a set of Web-based utilities for performing services such as file transfers and diagnostics.

Projects Target Computational Efficiency and Sustainability

ANNA MARIA BAILEY BAILEY31@LLNL.GOV

HPC computational efficiency is an important component of delivering LLNL mission excellence. In the last decade, Computation has led many successful sustainability projects aimed at optimizing the efficiency of HPC. Since 2004—when the ASC Program's Purple supercomputer was brought online—until today, through multiple generations of HPC platforms, the computing power in the TSF has increased five-fold (in one quarter of the space) while using 2.4 times less electricity. In addition, two computer buildings have been certified as "green" facilities by the Leadership in Energy and Environmental Design (LEED) rating system—the TSF achieved a Gold-level rating in 2009, and Building 451 achieved a Silver-level rating in 2011.

Computation is involved in several new efforts to reduce the energy use of LLNL computers and computing facilities and to promote new standards of quantifying efficiency gains beyond gross energy use. One of the organization's goals is to develop measurable advances in sustainable HPC stewardship that can be implemented throughout the DOE complex.

Progress in 2011

Computation's strategy for pursuing sustainability projects is documented in the HPC Sustainability Master Plan, which is a subset of the LLNL Site Sustainability Plan. According to the plan, Computation will pursue optimal HPC efficiency by focusing on the following core areas: benchmarking, computational fluid dynamics, LEED certifications, leveraging existing HPC capabilities, HPC capability gap analysis, free cooling, liquid cooling, innovative electrical distribution, sustainable HPC solutions, HPC platform power budgets, and power management.

The Sequoia system, which will be deployed in 2012, will combine several innovative approaches to address sustainability issues. More than 91% of Sequoia will be cooled using advanced liquid-cooling techniques. Its efficient design also includes a novel 480-V electrical distribution system, which provides improved voltage optimization to reduce energy losses.

Computation is also pursuing power management solutions, which will be critical to the success of exascale computing. System engineers created and implemented a centralized real-time data management system for all data sources, from the individual computer racks to the entire LLNL site. This effort presented many challenges: understanding how different types of hardware and software affect power utilization, correlating multiple data sources, coordinating with multiple owners of the data, accessing the data, selecting the best interface, comparing and viewing the data on a common platform, and creating various dashboards. Once complete, the infrastructure could be broadened to include all LLNL data centers and used throughout the DOE complex.

Another way to improve an HPC center's energy efficiency is to implement "free cooling," a technique that uses the outside air to drive the machine-cooling process. Although it will cost approximately $5.5M to implement free cooling in the TSF, studies show that free cooling would save an estimated 16M kWh per year and would pay for itself in four years. A free-cooling design would also increase the TSF's computational capacity from 30MW to 45MW and improve the overall facility power usage effectiveness from 1.27 to 1.15.

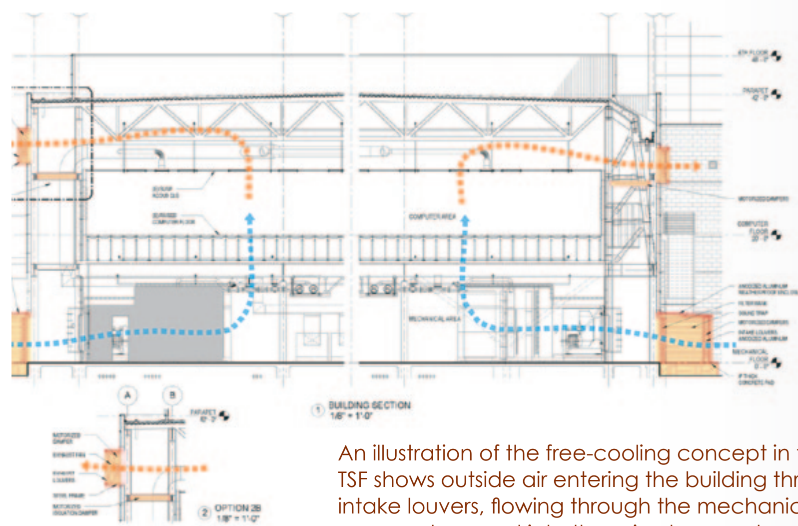

An illustration of the free-cooling concept in the TSF shows outside air entering the building through intake louvers, flowing through the mechanical area and upward into the raised computer room floor before exiting the building through exhaust louvers near the top of the building.

Grove Sets Down Roots as Sequoia's File System

MARK GARY GARY1@LLNL.GOV
MARC STEARMAN STEARMAN2@LLNL.GOV

The 55-petabyte (PB) Sequoia file system hardware and its collaboratively retooled Lustre file system software (collectively called Grove) will serve as the first stop for the deluge of data coming from the 20-petaFLOP/s Sequoia platform. The availability, accessibility, and reliability of the massive amounts of simulation data on Grove will be critical to the success of the missions that rely on this world-class computer.

In 2011, significant research was conducted and contracts established for all aspects of Grove hardware. To scale file system capabilities to meet the requirements of Sequoia, LC personnel collaborated with third-party developers to make major changes to the Lustre code base. Late in the calendar year, LC began installing and testing the first components of Grove.

The installation of Grove began in December 2011.

Progress in 2011

The rapid pace of innovation often requires that final technology be selected as close to the time of implementation as possible. In 2011, LC personnel finished researching file system hardware and software technologies and authored a series of procurements and supporting benchmark documents that would ultimately establish Grove, the Sequoia file system. Key requirements of these contracts included providing a minimum of 50PB of storage capacity with a sustained transfer rate of at least 500GB/s.

The final selection was a combination of NetApp, Inc. Redundant Array of Independent Disks (RAID) devices coupled with a Mellanox-based InfiniBand Storage Area Network (SAN) infrastructure, and a Lustre software development contract with challenging milestones. The hardware infrastructure will yield 55PB of usable storage capable of being written at a minimum rate of 750GB/s. When it is complete, Grove and its SAN infrastructure will span 56 racks of equipment.

The delivery of the Grove hardware was structured in three phases. During the initial phase, which was completed in September, hardware was provided for a test system sited at IBM Rochester where the Sequoia machine was built and where end-to-end functionality was performed. During the second phase, approximately half of the Grove system was delivered to LLNL in December. The third and final phase is scheduled for the first quarter of 2012.

The scale of both the bandwidth and the capacity of Grove are unprecedented. The existing Lustre file system could not handle the requirements of Sequoia, largely due to limitations of its underlying EXT-based file system implementation. LC software developers have worked for several years to replace the Lustre underpinnings with ZFS, an alternate and more scalable file system. Supported by a development contract with Whamcloud, Inc., developers collaboratively retooled Lustre and provided an initial file system software infrastructure for Grove. The performance requirements for the software are extreme and are expected to proceed iteratively between now and mid-2012 to achieve peak transfer rates of 750 GB/s. This software infrastructure has the added benefit that it will allow for a wide-range of alternate file systems to exist under Lustre.

Satisfying the requirements of Sequoia is challenging and requires a complex combination of state-of-the-art hardware and software. In 2011, LC established a framework capable of supporting Sequoia and began to install and test the combination, thus planting the seed for a successful deployment environment for Sequoia and the applications and missions it will support.

GPU Cluster Gives Scientists an Edge

BECKY SPRINGMEYER SPRINGMEYER1@LLNL.GOV
COMPUTING.LLNL.GOV/VIS/GPU_TECHNOLOGY.SHTML

For more than a decade, Computation researchers and developers have leveraged graphics processing units (GPUs) to accelerate the performance of their applications. GPUs are highly specialized parallel processors that evolved from graphics cards designed to accelerate three-dimensional graphics. As newer generations advanced in programmability and speed, developers applied GPUs to a larger application set. LLNL's 216-node cluster, called Edge, makes GPUs readily available to scientists who want to accelerate the performance of their simulations with parallel GPU algorithms. Several LLNL code teams have had breakthroughs in scientific computations and data analysis using Edge. As a test bed, Edge is helping researchers assess the path forward to exascale computing on heterogeneous architectures.

Progress in 2011

The Edge visualization and data analysis cluster, with 2,592 Intel Westmere processors and 412 NVIDIA Tesla M2050 GPUs with "Fermi" chips, offers excellent floating-point performance, nearly 20TB of total compute memory, and access to data via high-speed connections to the Lustre file system. As highly specialized, low-cost engines, GPUs have been successfully programmed for general-purpose science applications. Computation developers are experienced in accelerating applications such as image processing, bioinformatics, georegistration, seismic exploration, and hyperspectral analysis.

LLNL's Persistics team uses Edge to prototype their algorithms and test concepts and scalability for system design. Their project delivers scalable algorithms running on high performance, heterogeneous computing architectures to solve problems in wide-area persistent video surveillance. They were among the early developers who harnessed GPUs to solve intelligence community problems. Developers fluent in the graphics application programming interfaces (APIs) harnessed the GPUs, but they programmed in an environment on the bleeding edge. The current development environment on Edge includes APIs such as OpenCL and NVIDIA's CUDA (Compute Unified Device Architecture), which has enabled success for a larger set of application developers.

In an LLNL Laboratory Directed Research and Development project, scientists are researching exascale multiphysics simulations using arbitrary Lagrangian–Eulerian hydrodynamics and diffusion. The team uses Edge to focus on what kinds of code or data transformations can take advantage of heterogeneous nonuniform memory access architectures. Because future architectures will be constrained by power requirements, there is a premium on data motion and memory bandwidth. According to the project leader, early, regular, and sustained access to accelerator-equipped hardware, including GPUs and MICs (Many Integrated Cores), will be needed to advance the codes.

In the field of lattice quantum chromodynamics, LLNL researchers use Edge to perform novel calculations of the forces between nucleons and between exotic baryons. The latter calculations resolved conflicting claims about the force between two omega particles. Using GPU accelerators was essential for producing timely and adequate statistics. The team estimates that equivalent calculations on homogeneous clusters would have taken at least three times longer.

The Edge machine represents a decade of evolution, stretching back to the 2001 cluster that began LC's history of leveraging commodity graphics cards to accelerate Linux clusters. The success for data analysis is apparent through the interactive analysis of ever-larger and more complex problems. But GPUs are no longer only for visualization. This year marks the beginning of concerted efforts to port several LLNL application codes to heterogeneous architectures in pursuit of exascale computing, enabled by the Edge cluster.

Livermore postdoctoral researchers (left to right) Michael Buchoff, Joseph Wasem, and Christopher Schroeder used Edge to run novel lattice quantum chromodynamics calculations.

Foundational Research

Prepares the Laboratory for High Performance Computing in the Next Decade

LORI DIACHIN DIACHIN2@LLNL.GOV

High performance computing and simulation has greatly influenced the science, technology, and engineering research performed at Lawrence Livermore National Laboratory (LLNL). Recent strategic planning activities for areas such as stockpile stewardship and energy security highlight how simulating the evolution of complex systems with quantifiable uncertainty can improve decision-making processes. Moreover, the ability to create actionable information from massive data sets remains critical to current and emerging mission areas, including nuclear threat reduction, biosecurity, cyberspace, and intelligence.

An increasing number of sponsors ranging from the Department of Energy's (DOE's) Office of Science to the Laboratory's industrial partners are recognizing the value of high performance computing to their various research initiatives. To meet the needs of these sponsors and advance the missions of the Laboratory, LLNL must remain at the forefront of high performance computing. This goal drives the Computation research portfolio in computer science, mathematics, and data sciences.

Over the next several years, computer architectures will have nodes with hundreds to thousands of possibly heterogeneous computing cores, deep memory hierarchies, and million- to billion-way parallelism in capability systems. Using these systems effectively will require scalable computing tools and environments, new programming models, advanced mathematical algorithms, and robust simulation codes that are resilient to faults and sensitive to power consumption. Computation has engaged in significant research in these areas to increase application programmer productivity, improve simulation fidelity, reduce overall time to solution, and create knowledge from disparate data types.

As part of this effort, computer scientists are creating novel techniques for performance analysis and optimization. They are also providing tools to help scientists more easily develop and use mixed-programming models, statistical debugging procedures, and strategies for addressing resilience and fault tolerance. Mathematics research focuses on improved modeling and discretization processes, stable multiphysics coupling methods, and algorithms that will perform well on next-generation architectures. Research that combines intrusive and nonintrusive uncertainty quantification techniques for error estimation will further enhance understanding of simulation error. Computation is also exploring techniques for visualizing scientific simulation data and high-dimensional data associated with uncertainty quantification, as well as new analysis methods for very large graphs, streaming data, text analytics, and anomaly detection.

Through these and many additional research activities, Computation scientists are acquiring a deeper understanding of the issues application developers will face in the near future. This understanding is critical to maintaining the Laboratory's international leadership position in high performance computing and its ability to deliver new capabilities for LLNL's national security mission.

Optimization Tools Resolve Challenges with the Electrical Grid

THOMAS EPPERLY EPPERLY2@LLNL.GOV

The electrical grid, an extremely complex system of power plants, transmission lines, and distribution systems, is facing many challenges in the 21st century. Environmental concerns are driving the transition to more renewable energy sources, such as wind farms and photovoltaic solar systems, which are more variable and fluctuate with changing weather conditions. Smart grid technology is providing vast amounts of data about the current state of the grid that require enhanced processing capabilities. These technologies are also promoting new forms of demand response. In addition, plug-in electric vehicles are testing the grid distribution system and offer opportunities as a form of distributed power storage.

LLNL is implementing an innovative, parallel suite of modeling, analysis, and optimization tools to resolve issues with electrical grid planning and operations and to provide situational awareness. In 2011, the Laboratory continued collaborating with Energy Exemplar and IBM to show how parallel computing can provide game-changing information to decision makers. In addition, LLNL began a project to develop formulations for more accurate modeling of the Western Electric Coordinating Council (WECC) grid, enabling increases in network and time resolution. These larger models will require supercomputers to solve a new generation of algorithms and will ultimately provide decision makers with more reliable and accurate models of the electrical system.

Progress in 2011

Several factors are driving the need for improved simulation, optimization, and analytical tools for modeling the electrical grid. First, California is to derive 33% of its total electrical energy from renewable energy sources by 2020. Because the energy these sources produce is intermittent, energy suppliers must have better tools for gauging how much energy can be produced to keep up with demand. Second, the introduction of small to midsized solar plants and plug-in electric vehicles is creating a more complex distribution network. As these systems become more common, there may be localized excesses or shortages that exceed the grid infrastructure's operating constraints. Third, smart grid technology and future smart appliances will require new approaches to demand response. For example, an improved refrigerator might be able to sub-cool the freezer at night, when power is abundant, to reduce demand during the daytime. Ultimately, technologies such as these will have an affect on grid operators' ability to adjust power requirements.

LLNL continued working with Energy Exemplar and IBM to solve one of the fundamental problems in electrical grid modeling and optimization. The unit commitment problem is a mixed-integer linear program (MIP) that involves continuous and discrete variables. A solution to this problem provides an operating schedule for power plants in which the predicted demand is coupled with the lowest-cost combination of power plants, while satisfying the operational and maintenance constraints on the equipment. LLNL helped improve the performance of PLEXOS, Energy Exemplar's unit commitment software, by launching PLEXOS jobs on LLNL supercomputers instead of a desktop computer or PC. A PLEXOS–PC graphical user interface (GUI) was recently added to provide supercomputer speed with the convenience and usability of a PC GUI. The speed was also increased by relaxing the solution tolerance to a more practical value.

LLNL is also developing high-fidelity unit commitment models of the WECC grid for planning and operational modeling. Industry has focused on simplifying models to be solvable on workstation-sized computers. For example, some statewide planning exercises use a 42-node PLEXOS model of an electrical network that involves more than 10,000 nodes. LLNL is increasing the network and time resolution of these models. Traditionally, WECC traded power in hour increments. With the addition of variable power generation, WECC operators are now experimenting with mid-hour changes. Models with increased time resolution can help predict the effects of switching the power market to a sub-hour trading format.

The high-fidelity unit commitment models will require novel, parallel MIP solvers. MIP solvers combine linear programming techniques such as the simplex algorithm or interior-point methods and parallel branch and bound techniques (a systematic combinatorial search). Research is needed to achieve scalable results. Toward this end, LLNL plans to parallelize unit commitment algorithms from its academic partners at Princeton and the University of California at Berkeley. These collaborators have identified sources of parallelism in their algorithms. LLNL will implement, evaluate, and adapt these algorithms to maximize performance and scalability. In addition, LLNL is evaluating whether problem decomposition techniques can be used to discover other sources of parallelism for achieving scalable performance.

LLNL can leverage its existing capabilities in uncertainty quantification (UQ) to address the challenges of modeling and optimization under uncertainty. UQ techniques can generate ensembles of simulations to calculate the

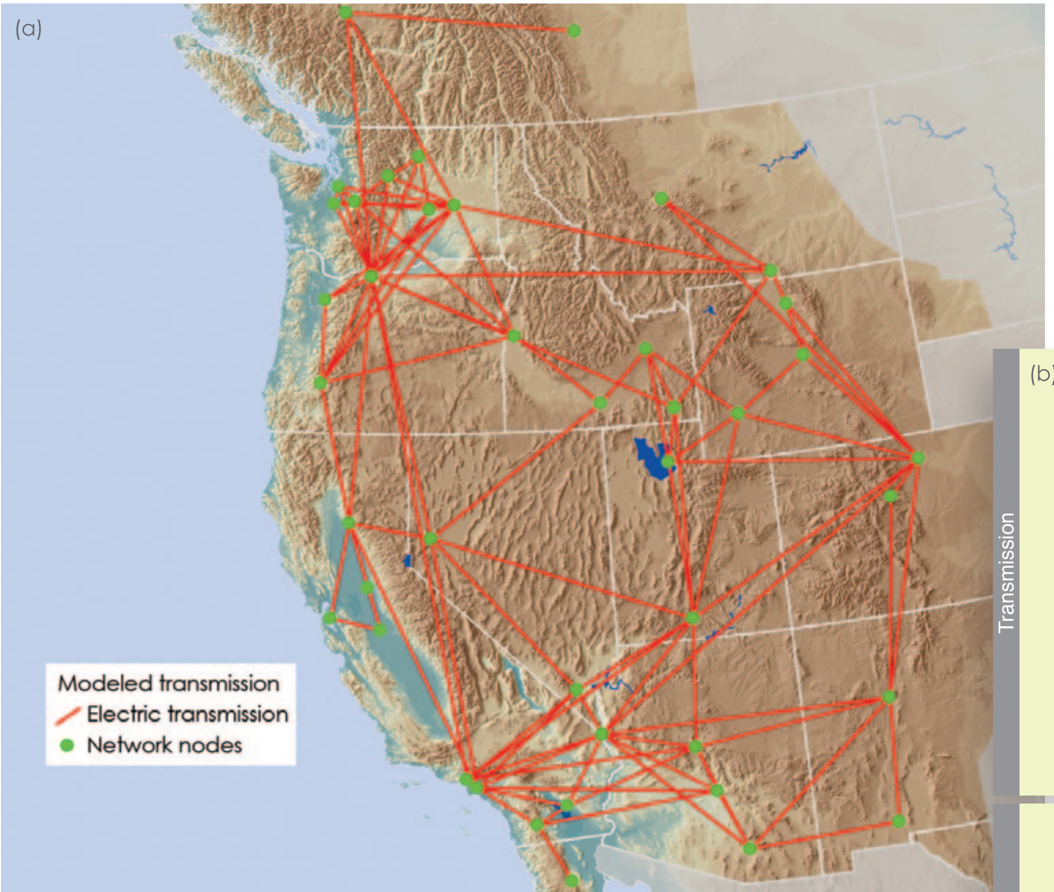

(a) PLEXOS uses 42 nodes (green dots) to model the Western Electric Coordinating Council grid. LLNL will develop models with orders of magnitude more detail.

(b) This conceptual drawing illustrates the vast, complex resources in the electric grid of tomorrow. Livermore's expertise in high performance computing is advancing the development of new technologies to secure the nation's energy supply for years to come.

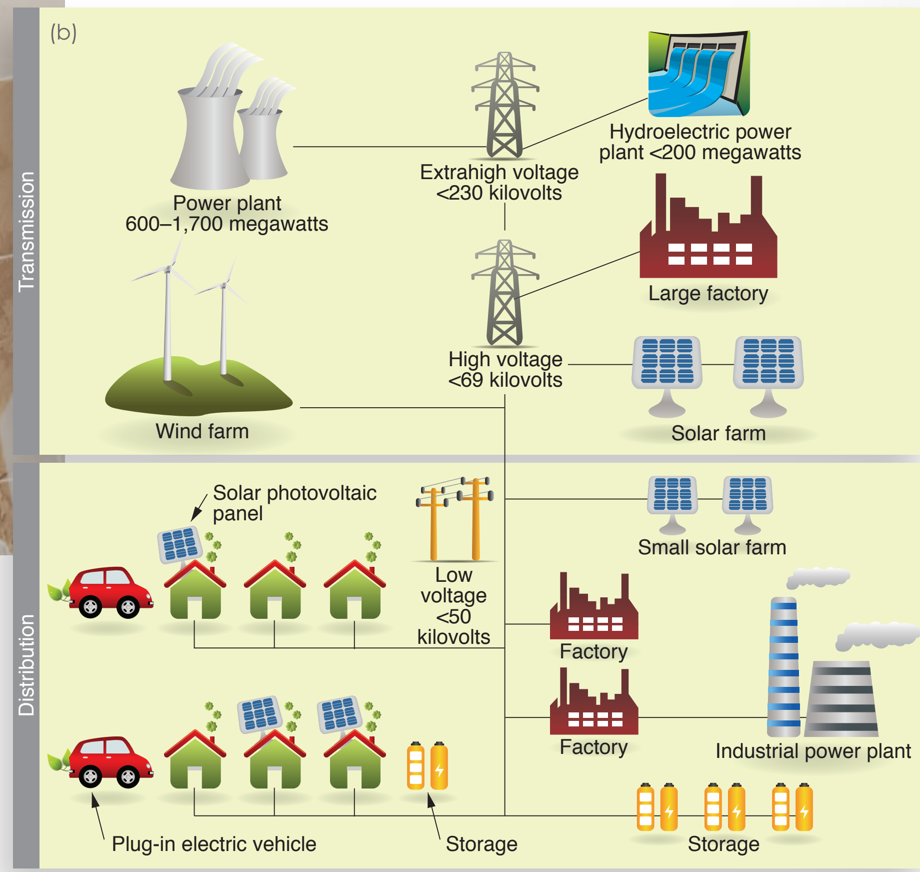

expected value for various wind and solar scenarios. Similarly, UQ techniques can identify the sources of uncertainty that have the largest impact on model solutions.

Working with the California Independent System Operators (CAISO), LLNL is also investigating smart-grid data processing techniques for handling the influx of data from smart technologies. Adoption of smart electrical meters is rapidly increasing the grid data available for analysis, but analysis tools have not kept pace. LLNL is creating new algorithms to provide better situational awareness for operators and regulators.

LLNL is working with collaborators to bring cutting-edge technology to bear on the many challenges arising from changes to the nation's electrical grid. Working with regulators and utilities, LLNL will partner with California and other states to achieve their goals of increased use of renewable energy sources in the 21st century. An improved grid infrastructure will enable the U.S. to grow its economy while safeguarding the environment.

Researchers Prepare for Exascale Systems with ExaCT

BRONIS R. DE SUPINSKI BRONIS@LLNL.GOV

The Exascale Computing Technologies (ExaCT) Strategic Initiative has prepared LLNL for addressing current trends in computer architecture, in particular, large-scale systems with many more nodes, each of which uses multicore chips. The Sequoia system, which will arrive at LLNL in 2012, will have more than 1.5 million cores. Exascale systems expected in the 2020 time frame are anticipated to have concurrency on the order of billions. Systems at this extraordinary scale will also have less memory and memory bandwidth per core, and reduced relative parallel-file system bandwidth.

ExaCT has four main thrust areas: (1) scalable math algorithms and associated programming models to address the memory restrictions; (2) advanced checkpointing solutions and application-level fault tolerance techniques that adapt to emerging architecture trends; (3) new debugging paradigms that reduce the effort required to produce correct code for exascale systems; and (4) performance methodologies that automate critical aspects of application development. Highlights in 2011 included an R&D 100 Award for the Stack Trace Analysis Tool (STAT), a novel lightweight correctness tool, as well as significant progress in all four thrust areas.

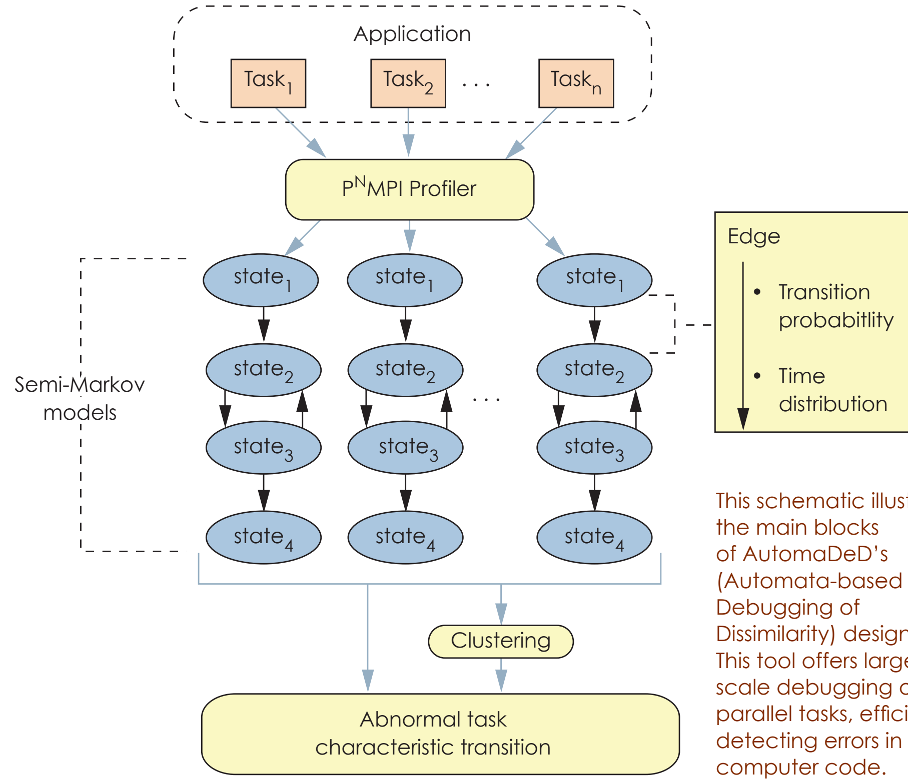

This schematic illustrates the main blocks of AutomaDeD's (Automata-based Debugging of Dissimilarity) design. This tool offers large-scale debugging of parallel tasks, efficiently detecting errors in computer code.

Progress in 2011

Through ExaCT, LLNL is developing three complementary code correctness paradigms. First, lightweight mechanisms, such as those embodied in STAT, efficiently gather data from application processes at the system's full scale and identify behavior equivalence classes—sets of processes with similar behavior. By narrowing the search for root-cause analysis to a few nodes, these classes reduce the problem to a scale suitable for traditional interactive techniques. Second, LLNL is designing methods to analyze how applications can use parallelization mechanisms, such as Message Passing Interface (MPI) processes, to identify semantic errors automatically. Third, statistical and machine learning approaches are being created to identify the source and type of coding error for algorithmic defects at scale. These new paradigms will improve application scientists' productivity by reducing the time it takes to determine which module, whether from the application itself or from the system (e.g., the MPI library), contains the root cause of a problem that arises at scale.

As part of the third paradigm, LLNL developed the AutomaDeD (Automata-based Debugging of Dissimilarity) tool to detect errors based on runtime information about the control paths that the parallel application follows and the times spent in each control block. AutomaDeD can suggest possible root causes by pinpointing, in a probabilistic rank-ordered manner, the erroneous processes (which form a small minority of the total number of processes) and the code regions in which the errors occurred. The processes form an outlier when they are clustered based on control flow and timing. Further, executions in the initial iterations are more likely to be correct than in later iterations, which are leveraged to determine correct or erroneous labels.

AutomaDeD models control flow and timing information of each task as semi-Markov models (SMMs). Through nonintrusive runtime monitoring, AutomaDeD

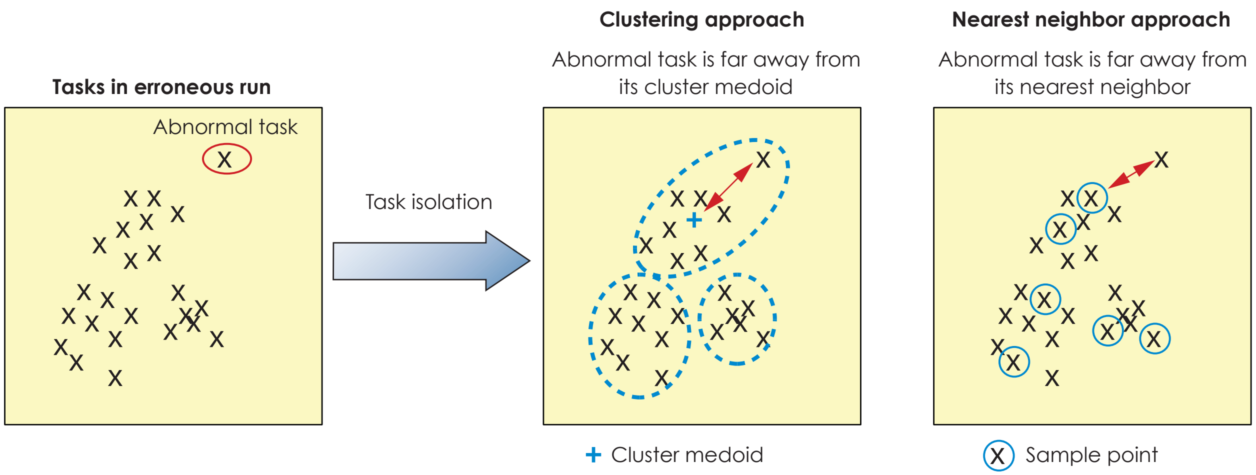

AutomaDeD uses clustering and nearest neighbor methods to isolate abnormal processes from normal ones in a run.

creates an SMM for each of the application's processes. SMM states represent communication code regions (i.e., MPI communication routines) and computation code regions (i.e., code executed between two MPI communication routines). A state consists of call stack information, such as the module name and an offset of the function calls that are currently active in the MPI process. This call stack information is obtained using the Dyninst tool (or by using the backtrace application programming interface from the GNU C library when Dyninst is not available) with a P^NMPI profiler at the beginning and end of MPI routines. Every edge is a tuple that represents a transition between the source state and the destination state. Two attributes are assigned to each edge: a transition probability that captures the frequency of the transition, and a time probability distribution that supports modeling the time spent in the source state conditioned on the destination state.

The application's execution is divided into a series of time periods, called phases. Within each phase, the application repeatedly exhibits the same execution pattern. First, AutomaDeD clusters SMMs to find the phase in which the error first occurs. AutomaDeD then clusters the SMMs of the different processes within this phase to determine the erroneous process(es). The natural number of clusters in an application can be provided by the developer from clusters of previous phases in the same run (before a fault is manifested), or can be inferred from traces of previous normal runs. Faulty tasks are those that deviate from the normal number of clusters, for example, by creating a separate cluster with few elements. Finally, AutomaDeD performs an activity called edge isolation, through which it identifies the SMM edge that most contributed to the distance between the normal processes and the erroneous one. In one mode of operation, multiple SMM edges, ranked by their contributions to the difference, are provided to the user. Developers can then target these code regions for debugging.

Recent work has created novel scalable mechanisms that allow online execution of AutomaDeD in large-scale systems. Other techniques that compress AutomaDeD's SMMs reduce the tool's overhead and increase its accuracy by eliminating noise that complicates detection of real dissimilarities between processes. Scalable outlier detection uses distributed sampling approaches to find faulty processes among many parallel ones with low overhead. AutomaDeD uses two mechanisms, including a scalable clustering technique, previously developed through ExaCT, and a novel nearest-neighbor method to find outliers efficiently. Results demonstrate that AutomaDeD can automatically detect the source code region that corresponds to a variety of errors in jobs with thousands of processes.

Novel Techniques Reduce System Error Rates

As HPC systems grow more capable, they also become larger and more complex. The increasing numbers of system components increase the probability that one of them will fail or that multiple components will interact in unexpected ways. Soft faults in chip circuitry can corrupt the state of transistors and cause applications to return incorrect results. The complex connections between components can induce small variations in system behavior that can cascade into significant degradations in application performance. For instance, if a network wire becomes frayed, the resulting packet losses and retransmissions can create a bottleneck for communication-intensive HPC applications. LLNL researchers are analyzing the effects of faults and performance anomalies on the behavior of HPC applications and systems. In doing so, the Laboratory has improved its ability to detect, localize, and tolerate system errors.

GREG BRONEVETSKY BRONEVETSKY@LLNL.GOV

Progress in 2011

Scientists used statistical techniques to analyze the behavior of HPC applications, creating models that capture an application's behavior during normal and faulty execution. These models make it possible to detect faults and characterize their location and properties. The Laboratory's AutomaDeD tool monitors the amount of time spent in each application code region, as well as other metrics, and computes the probability of the observed behavior given the probability distribution. In this way, AutomaDeD builds a probability distribution of normal behavior. A sequence of low-probability observations indicates a fault. By comparing the pattern of the unusual events to previously observed failures, application developers can identify the type of fault that is occurring. As a result, they can more easily pinpoint the root cause of a bug and alert system administrators to possible failures or performance degradation.

Precise measurements are needed to accurately model application behavior. However, in many cases, developers cannot obtain fine-grained measurements because information is only available at coarse granularity, or the measurements themselves perturb the application. Statistical profiling provides a solution to this problem by inferring fine-grained application properties using coarse-grained observations. LLNL researchers have developed a technique called diffractive profiling, which uses information about the application's control flow to improve the precision of the measurements. Diffractive profiling can also quantify the dependence of an application's behavior on the properties of its input data.

LLNL is collaborating with the University of Illinois Urbana-Champaign to develop algorithm-specific techniques to detect data corruptions in sparse matrix-vector multiplication, the fundamental operation of sparse linear algebra. Specifically, researchers are verifying the algorithm's output by checking the identity $c^T(Ax)=(c^TA)x$ for a range of carefully selected check vectors c. This approach significantly reduces the cost of fault detection and makes the technique applicable to a wider range of matrices than possible with prior techniques.

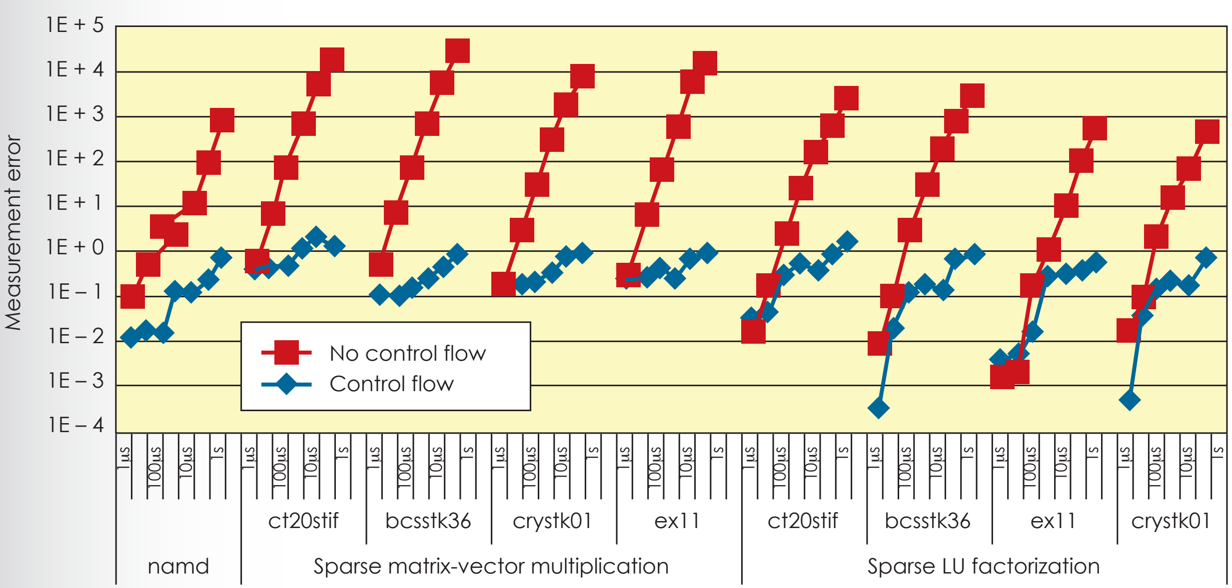

Diffractive profiling is significantly more accurate than traditional statistical profiling and, thus, can converge faster to the same result. The data show the prediction error of the two techniques on three different applications (bottom x axis), running on different inputs (middle x axis), as the period at which samples are taken varies from 1µs to 1s (top x axis).

NetWarp Provides Time-Accurate Network Simulations

DAVID JEFFERSON JEFFERSON6@LLNL.GOV
STEVE SMITH SMITH84@LLNL.GOV

NetWarp is a scalable parallel network simulator targeted at real-time and time-sensitive network scenarios. It combines the qualities of a virtualized network emulator with those of a classic discrete event network model. Each network node is represented as a full virtual machine running an actual operating system (OS) and application binaries. However, a parallel discrete event simulation layer also controls all of the virtual clocks so that they are causally consistent with one another and read the same as the real-time clocks on the target network being simulated. As a result, time-dependent and real-time software behaves in the simulation as it would on the target network.

Progress in 2011

Many networked systems, such as those for mobile phones, multimedia distribution, networked infrastructure control, military and cyber defense, and stock exchanges, have critical real-time performance requirements. Unfortunately, the behavior of these systems is notoriously difficult to simulate accurately. Their dynamics, especially their resource utilization, are often sensitive to minute details of the code at all levels of their software stacks. In addition, they are sensitive to the specific timing and rates of events in the scenario. Classical discrete event simulations may incorporate accurate network performance models, but they can only crudely approximate the network's logical behavior because they run models of the software rather than the actual code. Conversely, standard virtualized emulations run the software and accurately reproduce at least one of its legal executions, but their timing is off because their real-time clocks are synchronized with wall clocks instead of with those of the target system being simulated.

NetWarp is a parallel network simulator designed to combine the advantages of both classical and virtualized network simulation. It represents the network nodes as full virtual machines (using the virtualizer QEMU) running an actual OS, routing server, and application binaries. NetWarp also incorporates detailed performance models of the nodes, devices, and communication links in the target network. It uses those models along with a parallel discrete event simulator to synchronize all the virtual machines and their real-time clocks, which are controlled to read what the actual clocks on the target network would read. NetWarp accurately reproduces the timing of the simulated scenario so that its synchronization, resource utilization, and behavior match those of the target network.

NetWarp can currently simulate a virtualized network of mixed x86 and ARM nodes running a combination of Windows and Linux operating systems, along with HTTP and DNS servers and routing software. NetWarp has been tested at a scale of 1,024 virtualized nodes but could feasibly run at a larger scale.

Recently, LLNL began modifying NetWarp to exchange packets with nonvirtualized network simulations running on NS-3, a discrete-event network simulator. LLNL will then be able to run mixed simulations in which the critical network nodes requiring the most accurate detail are virtualized, while less central nodes are modeled as ordinary discrete event processes. As a result, very large simulations can be run using much less memory and fewer nodes of the underlying hardware cluster, without sacrificing speed or accuracy.

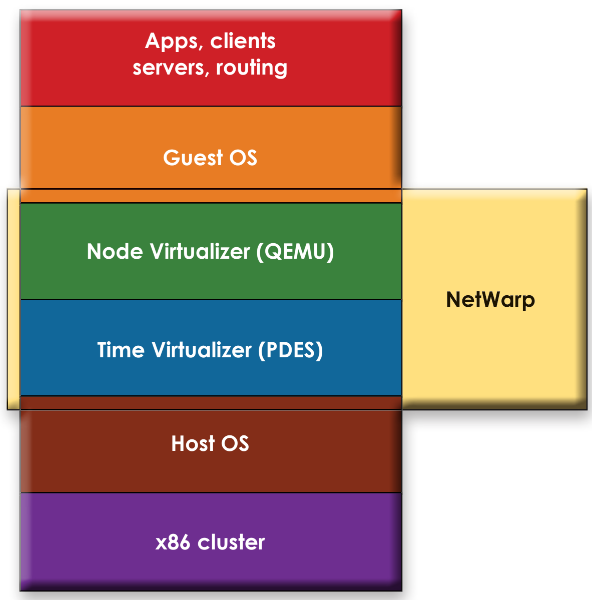

NetWarp runs as a parallel user application on a cluster. It virtualizes both network nodes (using QEMU) and their clocks (using a parallel discrete event simulator).

Error Estimation Builds Simulation Confidence

CAROL S. WOODWARD WOODWARD6@LLNL.GOV
JEFFREY A. F. HITTINGER HITTINGER1@LLNL.GOV

Simulation is relied on as a scientific and engineering tool, and has become increasingly important for making informed policies and decisions in areas such as stockpile stewardship and climate science. Improving and maintaining confidence in simulation results is therefore essential for DOE and Laboratory missions.

One important source of uncertainty in simulation results is discretization error, which arises from the approximation of a continuous model with a discrete one. Discretization error is typically the dominant numerical error and the most difficult to characterize. As part of a UQ strategic initiative, researchers in the Computation Directorate are developing adjoint and error transport techniques to estimate discretization errors with the goal of incorporating such estimates into larger uncertainty assessments. In the past year, both techniques were successfully applied to models relevant to hydrodynamics. LLNL also obtained theoretical results supporting the techniques' use as reliable error estimators.

Progress in 2011

Discretization error obeys partial differential equations (PDEs) similar to the original continuous PDEs, but with forcing due to the local residual, that is, the discrepancy obtained by applying the original PDEs to the approximate solution. In the error transport method, this auxiliary set of error equations is discretized and advanced with the approximate solution. The result is a field estimate of the discretization error on the problem domain.

The nonlinear error transport method was considered for finite difference and finite volume discretizations of nonlinear hyperbolic equations. However, further analysis indicated possible performance constraints depending on the type of solution. For smooth solutions, the theory relates the order of accuracy for the error estimate to the rates of convergence of the discretizations for the governing equations, the error equations, and the residual. In this case, the technique is asymptotically correct—meaning it provides reliable estimates. For discontinuous solutions, the theoretical results do not guarantee an accurate error estimate. Yet, numerical results demonstrate effective error estimates, provided nonlinear transport terms are included. Theory indicates that without these terms the accuracy of the estimate degrades, and the resulting inconsistent approximations can lead to numerical instability.

Using the adjoint error estimation technique, the error in a specific quantity of interest (QOI), computed from the approximate solution, is evaluated directly. An error representation formula is derived that relates the error in QOI to the approximate solutions of the governing PDEs and an auxiliary adjoint problem. The adjoint problem, which evolves backwards in time, is a set of PDEs specified by the adjoint operator of the governing "forward" PDEs and initial data that are specific to QOI.

LLNL has developed a theory that provides greater flexibility for implementing the adjoint solver by allowing separate adjoint and forward solvers and grids. This work bridges a gap in the adjoint error estimation theory for finite volume methods, which was previously developed in a variational framework. The theory accommodates the nonlinear polynomial reconstruction of the forward and adjoint solutions from local cell averages. Asymptotic convergence rates for the error estimates are obtained for smooth solutions. Numerical results demonstrate the effectiveness of the adjoint-based approach on smooth and discontinuous solutions of linear hyperbolic problems.

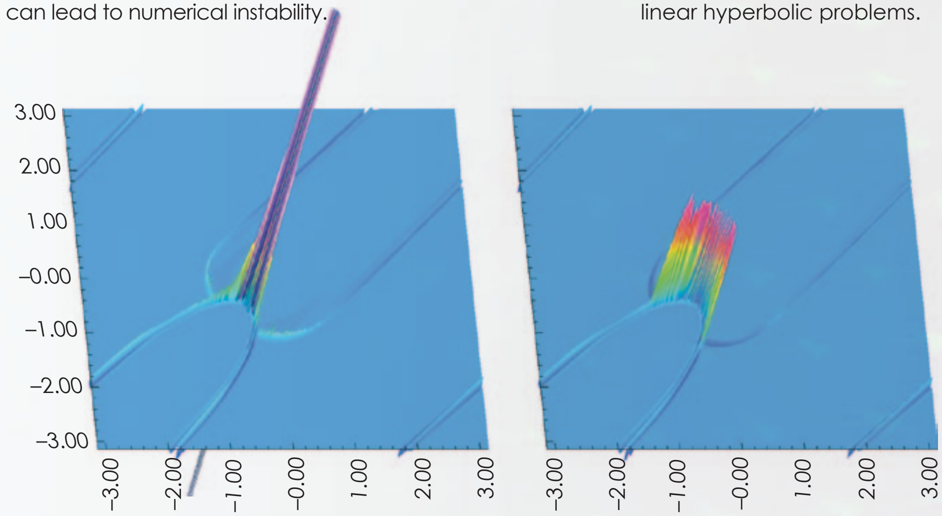

Unstable growth of the error estimate occurs at a curved shock (left) if nonlinear terms are omitted from the error transport equations. When these terms are retained, error estimates remain bounded (right).

Applications

Help Ensure National Safety and Security

TRISH DAMKROGER DAMKROGER1@LLNL.GOV

The Computing Applications and Research (CAR) Department in the Computation Directorate reaches across Lawrence Livermore National Laboratory (LLNL) to ensure mission success for all Laboratory programs. Scientists in CAR also collaborate with universities, industry, and other national laboratories to perform research and development in computer science, computational science, and mathematics. They use their expertise to develop software that enables the programs to fulfill the mission requirements of LLNL's sponsors in the National Nuclear Security Administration (NNSA), the Department of Energy (DOE), the Department of Defense, the Department of Homeland Security, other federal and state agencies, and industry.

CAR scientists collaborate with multidisciplinary project teams that range from small two-or-three person teams to larger teams of 100 or more scientists. As a result, CAR personnel develop deep domain expertise in a far-reaching set of fields, such as materials modeling, fluid dynamics, genomic sequencing, climatology, geophysics, astrophysics, scalable algorithms, parallel processing, cyber security, and real-time control systems. This extraordinarily broad "corporate" expertise enables CAR scientists to become an integral part of the various Laboratory programs while ensuring that the very best computational research and software engineering practices are applied to each project.

As the world's technological capabilities evolve at an increasingly rapid pace, so do threats to the U.S., particularly in the areas of energy, environmental security, domestic and international security, and nuclear security. A national laboratory has the depth and breadth of talent to confront these complex challenges. In addition, solving the technical aspects of these problems requires the ability to quickly transform research into usable application software. A large, dynamic organization such as Computation has the ability to rapidly deliver from first concept to hardened software product.

Future software applications will help model, simulate, and analyze systems of increasing complexity, and more attention will be given to the quality of the answers provided by the applications software. As a result, there will be a demand for more precise analytical tools, and progress will need to be made in uncertainty quantification and verification and validation research. This expectation is driving current improvements and enhancements, which, in turn, is resulting in more complex software.

Many challenges face the computer and computational scientists working in the Laboratory programs. CAR scientists often provide the tools the programs need to be successful, whether it is an exquisite control system, an integrated visualization tool, or other unique capability. The computer science discipline is an integral part of nearly every project team at the Laboratory, and our scientists continue to provide the expertise to solve many of the nation's problems.

Simulations and Experiments Improve the Fuel Economy of Heavy

KAMBIZ SALARI SALARI1@LLNL.GOV

Commercial trucks in the U.S. vary significantly in size and load-carrying capacity. The U.S. Department of Transportation classifies these vehicles as Light Duty (1–3), Medium Duty (4–6), or Heavy Duty (7–8) based on their gross vehicle weight rating. At highway speeds, a Class 8 tractor-trailer uses approximately 50% of the usable energy produced by the engine to overcome aerodynamic drag, while rolling resistance consumes approximately 30% of the usable energy.

This project focuses on improving the aerodynamics of Class 8 tractor-trailers to achieve better fuel economy and satisfy regulation and industry operational constraints. For 12 years, LLNL has collaborated with industry to use high performance computational modeling and simulations to conduct aerodynamic research on heavy vehicles. LLNL's large-scale computer platforms and advanced fluid dynamics codes have helped identify aerodynamic improvements that boost the fuel efficiency of Class 8 tractor-trailers by as much as 12–13%. Low-rolling-resistance, wide-base single tires can add an additional 4–5% fuel efficiency. This 17% fuel-economy improvement could save approximately 6.2 billion gallons of diesel fuel and reduce CO_2 emissions by 63 million tons per year. These significant results are motivating the modern trucking industry to pursue aerodynamic improvement technologies.

This project is sponsored by DOE's Vehicle Technologies Program. An important part of this effort is to expand and coordinate industry collaborations and participate with industry in designing the next generation of highly aerodynamic Class 8 heavy-duty vehicles.

Progress in 2011

There are several critical aerodynamic flow regions around a typical tractor-trailer that can significantly influence the overall vehicle drag and, therefore, must be accurately captured with high performance computing (HPC) simulation codes. The flow regions include the trailer base, trailer underbody, trailer axle and wheel assembly, tractor-trailer gap, around the tractor body, and through the engine. Drag-reducing devices, such as trailer boat tails, skirts, and gap sealers, can be added to the truck to treat the drag produced by the critical flow regions. Boat tails are made of three or four flat plates that attach to the trailer base at a slight inward angle. The plates turn the flow field more sharply into the trailer wake, which increases the pressure on the trailer base and improves fuel economy by 4–7%. Trailer skirts are flat plates that extend beneath the trailer and span the distance between the rear tractor wheels and the trailer wheels. The skirts shield the trailer wheels and axles from crosswinds, thereby streamlining the trailer underbody and improving the fuel efficiency by about 5–7%. Some designs include skirts that also cover the trailer wheels and extend to the trailer base. Gap sealers are devices that reduce the amount of cross flow in the tractor-trailer gap, a phenomenon that occurs when the tractor-trailer is operating in a crosswind and a portion of the flow field is drawn into the gap. Gap sealers can improve the fuel economy by about 1–2%. Additional aerodynamic improvements can be achieved through vehicle streamlining, tractor-trailer integration, and low-rolling-resistance, wide-base single tires. In 2011, the LLNL team published a document that provides comprehensive design guidance for trailer-base devices to achieve optimal aerodynamic performance.

Also in 2011, the LLNL team collaborated with Navistar International Corporation to run and analyze full-scale wind tunnel test results at NASA Ames National Full-Scale Aerodynamics Complex facility. This initiative coupled

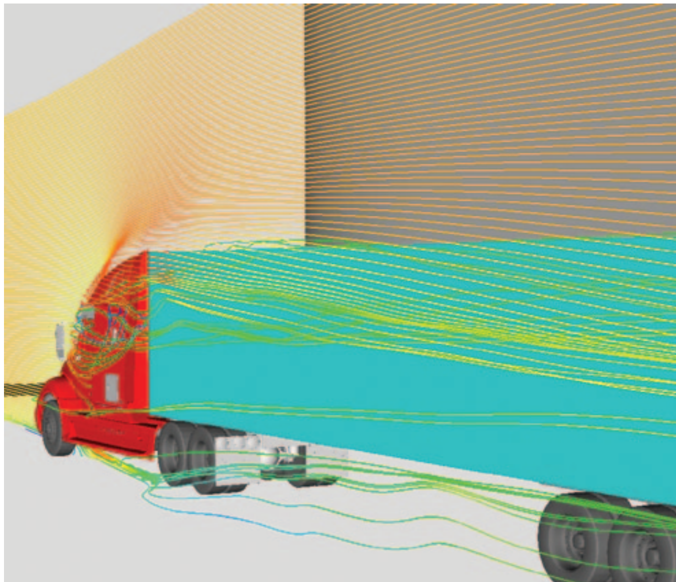

High performance computing simulation codes help accurately capture critical flow regions around a Class 8 tractor-trailer.

Trucks

HPC with experimental validation by using an 80'×120' wind tunnel and track tests to accelerate the design, prototyping, and deployment of energy-saving technology. The three-month test included 23 aerodynamic drag-reduction devices and concepts with four combinations of tractors and trailers: a long-sleeper tractor with a 53'-straight-frame trailer, a long-sleeper tractor with a 28'-straight-frame trailer, a day-cab tractor with a 53'-straight-frame trailer, and a day-cab tractor with a 53'-drop-frame trailer. Approximately 140 wind tunnel runs were completed during the study.

The LLNL team conducted computational fluid dynamics simulations that improved several drag-reduction add-on devices. The simulations were run on a full-scale heavy vehicle geometry in a 6-degree crosswind (7 mph) at highway speed (65 mph) using a finite-volume code with polyhedral meshes up to approximately 700 million cell faces. The Reynolds-averaged Navier–Stokes equations were solved for the turbulent flow over the vehicle with various drag reduction devices installed in the tractor-trailer gap, trailer underbody, and trailer base. The results of the simulations demonstrated several important aerodynamic trends that impact the performance of add-on devices.

LLNL partners with several key members of industry in this effort, including Navistar (tractor manufacturer); Kentucky Trailer and Wabash National (trailer manufacturers); Frito-Lay and Safeway (large fleet operators); Freight Wing, Inc. and ATDynamics (developer and manufacturer of aerodynamic devices); and Michelin (manufacturer of low-rolling-resistance, wide-base single tires). The research and development has resulted in two DOE research grants in collaboration with Navistar: Fleet Evaluation and Factory Installation of Aerodynamic Heavy Duty Truck Trailers, and Supertruck and Advanced Technology Powertrains. The first grant will introduce the best available drag reduction devices and concepts to the public within 2.5 years. The second grant will introduce a completely new highly aerodynamic integrated tractor-trailer that will be designed from the ground up. This next-generation tractor-trailer will demonstrate a 50% increase in vehicle freight efficiency measured in ton-miles per gallon.

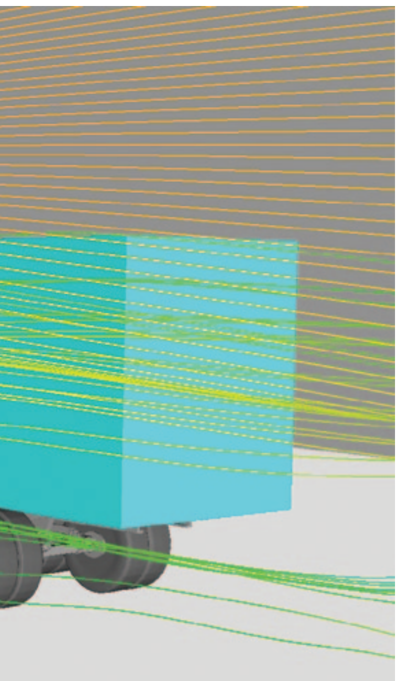

A track test vehicle is equipped with a gap fairing, trailer skirts, and a trailer boat tail.

Climate Science Gets Boost from Ultrascale Visualization Analysis

CHARLES DOUTRIAUX DOUTRIAUX1@LLNL.GOV

For more than two decades, LLNL has developed data-driven technology that allows simulation and observation data to be shared, reproduced, and reused. In 2011, LLNL added a new project to its visualization capabilities—the Ultra-scale Visualization Climate Data Analysis Tools (UV-CDAT) framework. UV-CDAT integrates five analytical and visualization tools—CDAT, VisTrails, ParaView, VisIt, and R—into one application that allows scientists to explore and analyze large-scale data sets. The UV-CDAT framework, which is based on Python, includes parallel streaming statistics, analysis and visualization pipelines, optimized parallel input/output (I/O), remote interactive execution, and automatic data-provenance processing and capturing. The framework's novel graphical user interface (GUI) includes workflow data analysis, visualization construction tools, and customizable functions.

UV-CDAT is led by Computation's Atmospheric and Earth Science Computer Applications software team and is a cross-institutional effort of integrated computational and science teams from Kitware, Inc.; Lawrence Berkeley National Laboratory; NASA; Oak Ridge National Laboratory; Los Alamos National Laboratory; Polytechnic Institute of New York University; and University of Utah. The project, which is funded by DOE's Office of Biological and Environmental Research, directly supports DOE's climate science mission.

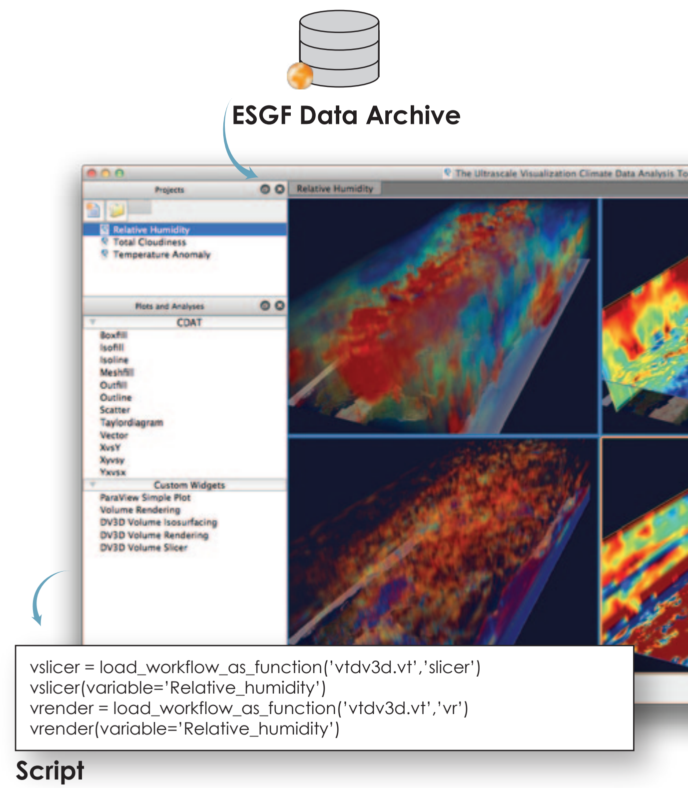

UV-CDAT accesses distributed data from the Earth System Grid Federation (ESGF) data archive. Parallel processing and data reduction and analysis takes place via ParaView, and then several views of the data are displayed. A workflow captures the entire process for reproducibility and knowledge sharing.

Progress in 2011

At its core, UV-CDAT builds on four key open-source technologies: (1) the Climate Data Analysis Tools (CDAT) framework previously developed at LLNL and widely used by the climate community for analyzing, visualizing, and managing large-scale climate data; (2) ParaView, a multiplatform parallel-capable visualization tool with new capabilities that better support the specific needs of the climate science community; (3) VisTrails, a scientific workflow and provenance management system for data exploration and visualization; and (4) VisIt, a parallel-capable visual data exploration and analysis tool. These combined tools, along with glues to general-purpose tools (e.g., R) and custom packages (e.g., vtDV3D), provide a synergistic approach to diagnosing climate models.

The UV-CDAT team's original goal was to integrate three technologies: CDAT, ParaView, and VisTrails. Their initial effort resulted in a new build system; a GUI that seamlessly integrates all components, including remote visualization and parallel functionality; provenance capture functionality; and a parallel I/O infrastructure. The team also developed vtDV3D, a package of high-level ParaView modules, which allows users to create and integrate custom visualization tools through a series of user-friendly workflow interfaces.

The UV-CDAT team was determined to deliver a product that offered an integrated, easy to use, more intuitive user experience than competing technologies. To this end, the team developed a GUI front end that far surpassed their initial expectations. The new interface is attractive, intuitive, and exposes all the available technologies in a simplistic way. The interface allows users to quickly explore, discover, analyze, and visualize data. In addition, the interface is designed so that components, such as VisIt three-dimensional visualization and R statistical analysis, can easily be incorporated. Currently, there are more than 40 components included in UV-CDAT.

The unceasing exponential growth in computation power and storage capability presents scientists with the daunting task of analyzing petascale data with tools conceived for much smaller data sets. As a result, the simple task of reading data can consume most, if not all, of

Tools

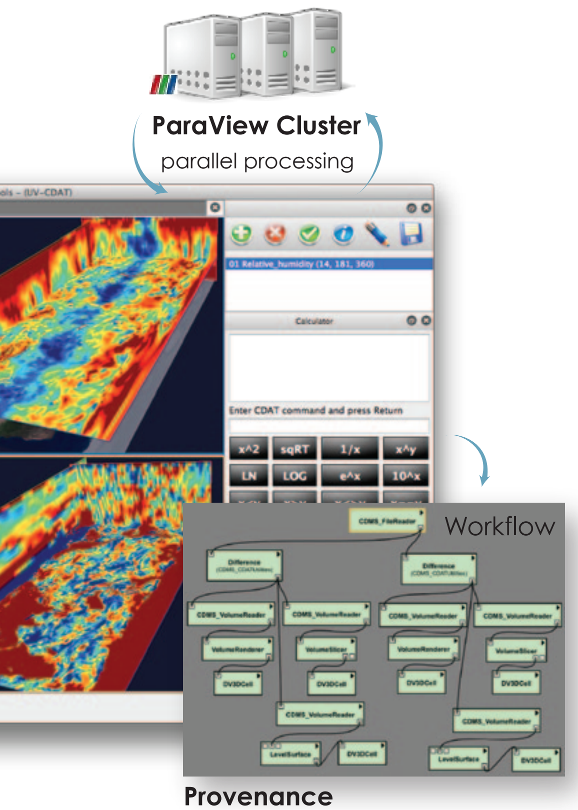

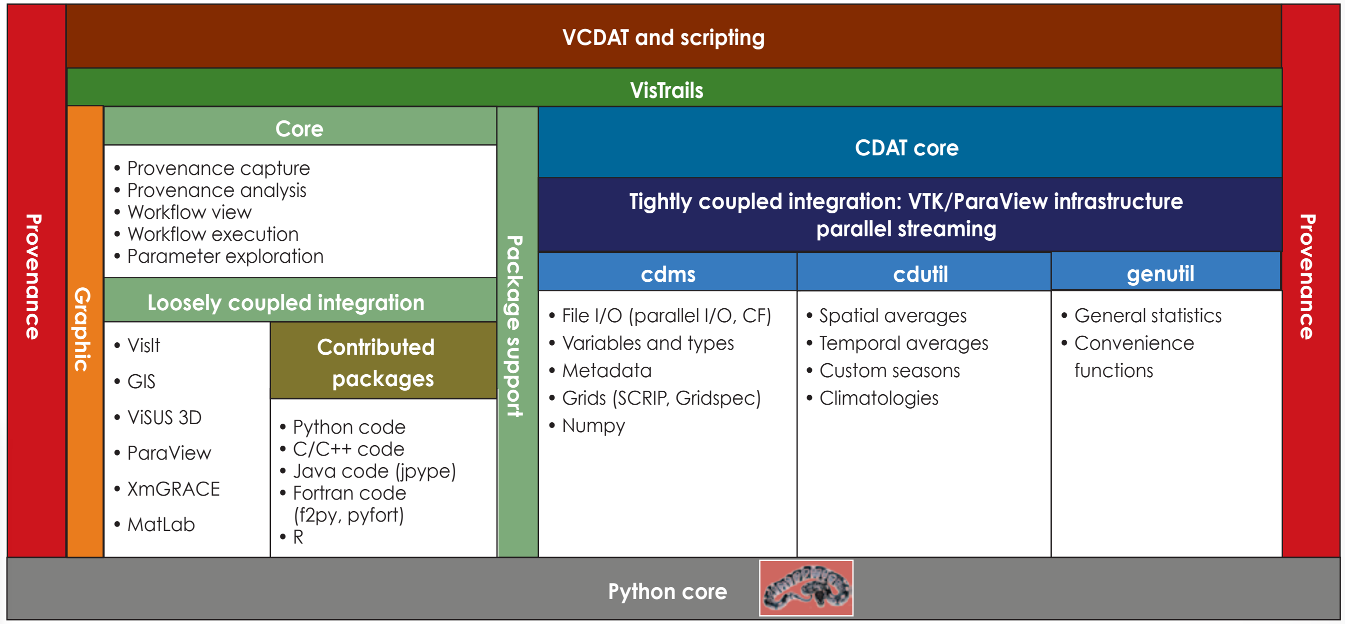

The UV-CDAT architectural layer shows the framework's components, either integrated tightly or loosely coupled. From this design, other packages can easily be integrated.

a scientist's time, hence reducing the scope of work. UV-CDAT addresses this issue by taking advantage of the newer multicore and multiprocessor technologies. In particular, the UV-CDAT team made remarkable progress in developing parallel-capable I/O tools. They optimized and generalized existing NetCDF-vtk filters, which significantly reduced the time it takes to assimilate data. At the same time, while focusing on common spatiotemporal dimensions, the team altered ParaView in such ways that it could align some problems to the existing parallel hardware, which drastically improved processing time. For more general problems, the team set up a parallel streaming interface.

Another challenge UV-CDAT addresses is the inability of a scientist to harvest all the data needed for a study onto a single local system because of the large numbers and massive sizes of available data sets. UV-CDAT solves this problem by connecting directly to the community's Earth System Grid Federation (ESGF) system of federated peer-to-peer (P2P) nodes. ESGF hosts remote global data and metadata services, which allow users to search semantic repositories, browse ultrascale data holdings, and visually inspect pregenerated plots. By connecting to ParaView servers, selected data sets can be visualized at multiple scales and compared across centers via standard visualization algorithms or interactive visual analytics.

The world is looking at climate science with increased scrutiny, thus UV-CDAT must allow for easily reproducible and transparent analyses of data. The integration of VisTrails into the framework fulfills both goals. A scientist can now issue a request to the product server for a standard analysis process or upload a custom algorithm that is run server-side onto distributed data sets, and then receive processed data on their desktop for further analysis and inspection. Through VisTrails, provenance metadata is recorded during each step of the process and archived as a workflow co-located with the data product. Other scientists can run the same analysis from the workflow descriptor at any point in the future to confirm the results. They can even expand on the findings by running variants of the processing algorithm or using different input data sources. All data products and workflow descriptors can be automatically archived to an ESGF P2P node.

New Target Diagnostics Support the Path to Ignition

RANDY SHELTON SHELTON1@LLNL.GOV

The physics results derived from the National Ignition Facility (NIF) experimental campaigns require a wide variety of target diagnostics. Several software development projects implement the configuration, controls, data analysis, and visual representation of most of these diagnostics. To date, more than 40 target diagnostics have been developed and commissioned to support NIF experiments.

Progress in 2011

In 2011, diagnostics were developed and enhanced to measure ignition performance in a high-neutron-yield environment. CAR's Integrated Computer Control System team works closely with NIF scientists to control, configure, and ensure successful delivery of the diagnostics. Performance is optimized around four key variables: adiabat, which is the strength and timing of four shocks delivered to the target; mix, which is the uniformity of the burn; velocity of the imploding target; and the shape of the imploding deuterium–tritium (DT) hot spot.

Adiabat is measured using the Velocity Interferometer System for Any Reflector (VISAR) diagnostic. VISAR consists of three streak cameras, instruments that measure the variation in a pulse of light's intensity with time. The enhanced VISAR streak cameras use a 10-comb fiducial signal controller to allow for postshot correction of the streak camera's sweep nonlinearity.

Mix is measured by the Neutron Time of Flight (NTOF) and Radiochemical Analysis of Gaseous Samples (RAGS) diagnostics. To accommodate high-neutron-yield shots, NTOF diagnostic controls are being modified to use Mach–Zehnder interferometer signals to allow the digitizers to be moved from near the target chamber to the neutron-shielded diagnostic mezzanine. The first phase of RAGS diagnostic commissioning was completed in December 2011. RAGS analyzes the trace gases added to the NIF target capsules that undergo nuclear reactions during the shot. These gases are collected and purified for nuclear counting by the RAGS system. Three new instrument controllers were developed and commissioned to support RAGS: a residual-gas analyzer that measures the gas content at various points in the system, a digital gamma spectrometer that measures the radiological spectrum of the decaying gas isotopes, and an instrument controller that interfaces to a programmable controller-based gas-collection system.

A third Gated X-ray Detector (GXD) diagnostic, which supports the implosion velocity measurements, was tested and commissioned. This GXD views the target through a slit contained in its "snout," which allows the instrument to measure the implosion shape on the same shot.

Two diagnostics, Active Readout in a Neutron Environment (ARIANE) and Neutron Imaging (NI), were commissioned to measure the implosion shape in a high-neutron environment. The controls for ARIANE, a fixed-port gated x-ray imager, contain a neutron-shielded camera and a micro-channel plate-pulser with neutron-sensitive electronics located in the diagnostic mezzanine. The NI diagnostic is composed of two Spectral Instruments SI-1000 cameras located 20m from the target. The cameras provide neutron images of the DT hot spot for high-yield shots.

These new or enhanced diagnostics provide insight into NIF's experimental campaigns and facilitate the optimization of the four key ignition variables. In 2012, three new diagnostics will be added and four diagnostics will be enhanced to support the continuing optimization series of campaigns.

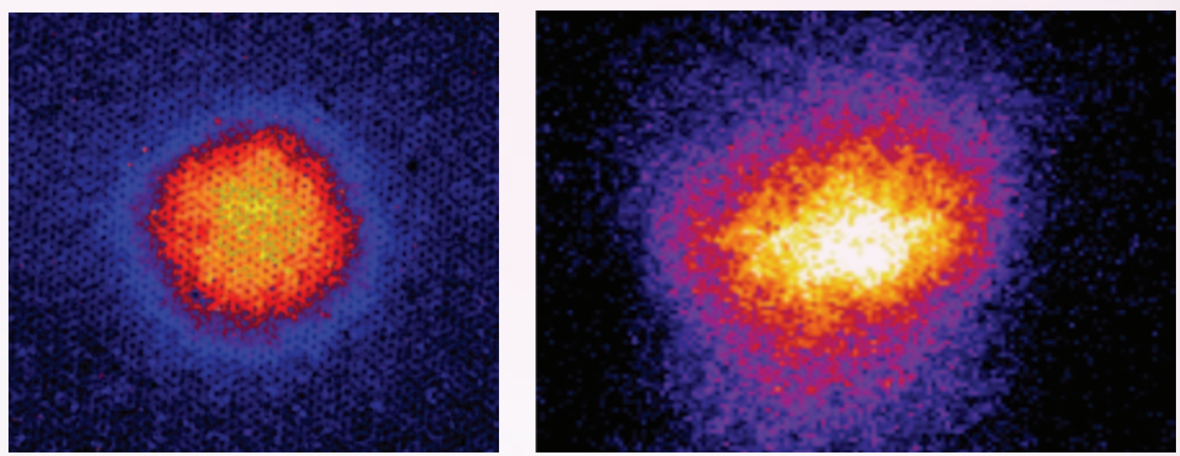

Output data from the Active Readout in a Neutron Environment (left) and Neutron Imaging (right) diagnostics show the shape of the imploding deuterium–tritium hot spot in a high-neutron environment.

Information Technology

Improvements Support a Modern, Mobile Workforce

GREG HERWEG HERWEG1@LLNL.GOV

Lawrence Livermore National Laboratory (LLNL) missions require highly productive scientists, researchers, and support personnel who can collaborate effectively and have access to the proper information technology (IT) tools for their particular mission space. Computation's IT staff provide support and technology solutions to users in a broad range of program and mission areas. One challenge for the IT staff is to make the tools that are commonplace outside the LLNL perimeter available to employees. As the popularity of mobile wireless devices and other gadgets continues to grow, Computation IT staff work diligently to make sure the institution's IT posture is not outpaced by technology. However, LLNL must contend with a unique set of cyber-security and fiscal constraints that are not present in many commercial organizations.

One fundamental principle of the IT strategy is to help end-users be more productive in a cost-effective and secure manner. To help meet this goal, Computation's IT team began a Laboratory-wide rollout of new desktop and laptop computers on an unprecedented scale. A bulk procurement of up-to-date equipment resulted in cost savings for the institution while providing the right tools for each LLNL program. Deploying new desktop and laptop equipment produced many benefits: higher productivity for users through faster machines and less downtime, more efficient IT staff due to a decreased number of trouble calls, and a more robust response to cyber-security threats by configuring standard platforms with the latest software.

Computation's IT staff is committed to providing software tools that help users work effectively on a day-to-day basis, whether they are research scientists or operational staff. The IT team is pursuing ways to give users a convenient, modern environment with anywhere–anytime connectivity while complying with necessary mandates and regulations. Several IT improvements have been made to meet the needs of a mobile workforce, including allowing BlackBerrys in many Limited Access areas, expanding WiFi availability in buildings across the Livermore site, launching a pilot program for office-to-office video conferencing, rolling out "LabBook" for dynamic peer-to-peer collaboration, and conducting a pilot test for LLNL-managed iPads and iPhones.

Computation employees are critical to planning, deploying, and operating the LLNL-wide IT services that enable a productive and modern workforce. The articles in this section provide a sample of important improvements that are under way at the Laboratory.

Cyber Security Event Turns Up the Heat on Cyber Attacks

ED MATSCHE MATSCHE1@LLNL.GOV

Cyber attackers are targeting the Department of Energy (DOE), the National Nuclear Security Administration (NNSA), contractors, and other federal agencies with increasing frequency and intensity. NNSA has formed a team of elite cyber responders, called Cyber Tracers, so agencies can coordinate cyber-security responses and respond to threats and attacks more quickly and effectively.

In 2011, LLNL's Cyber Security Program (CSP) hosted a first-ever exercise that convened 26 cyber security experts from around the country for three days of cyber combat to repel a simulated intruder. The event, called Tracer INFERNO (Intensive Network Forensic Exercise on Real Network Operations), focused on building relationships and sharing expertise across the DOE complex. In addition, LLNL contributes to the NNSA cyber security posture by hosting and maintaining a Master Block List service for nine DOE organizations. This service provides a set of domain names that are known or suspected to be untrustworthy. Sites use the list to create their filters and blocks and keep their users safe from attack. CSP also increased their efforts to train employees to recognize and prevent potential cyber threats through Laboratory-wide presentations and mock spear-phishing e-mail attacks. This multi-pronged approach to cyber security is critical to preventing and responding to attacks on DOE computers, networks, and sites.

Progress in 2011

The Tracer INFERNO scenario was based on a real cyber attack on several NNSA laboratories that occurred in early 2011. The exercise provided an opportunity for experts from around the complex to meet and work together, which ultimately creates a more unified response capability. Through exercises like INFERNO, NNSA is developing a multi-institutional team of incident responders who share information and are better equipped to recognize patterns and cohesively respond to cyber attacks.

Although training cyber security experts is a critical component of a successful cyber security program, employee awareness is the first and most important line of defense. The LLNL e-mail system receives 53 million e-mails per month, 70% of which are identified as spam and filtered before they reach the users. LLNL's comprehensive cyber strategy includes teaching employees to recognize a potential threat and to instinctively not open e-mail attachments or follow links provided by unknown or untrusted sources.

Instead of sending thousands of e-mails randomly, hoping to trick a few victims, today's adversaries are more often targeting select groups of people who have something in common. In one newer tactic, called spear-phishing, adversaries send e-mails to a specific user or group, and the apparent source of the e-mail is an individual within the recipients' own company. These

A cyber security exercise hosted by Lawrence Livermore National Laboratory in September 2011 brought together the National Nuclear Security Administration's elite cyber responders.

virtual traps often lure users to fake Web sites and trick them into revealing personal information, or the Web sites are equipped to download malicious code to the user's computer without the user's knowledge.

The figure on this page shows an actual spear-phishing e-mail (sanitized for anonymity) that was sent to an LLNL employee. The scenario is typical of these types of attacks: The attacker sends an e-mail to a small targeted group claiming to be from the group manager's personal e-mail address. The e-mail states that a report has been completed and the manager needs feedback. The victim receives the e-mail, clicks on the link provided, and downloads, in this case, a zip archive. The zip archive contains a malicious file that, when executed, installs a backdoor on the LLNL computer and displays an unfinished report. The attacker now has an entry point to the LLNL computer. In this case, the attack was unsuccessful; however, it highlights the importance of employees using vigilance and caution every time they open an e-mail.

For spear phishing to be successful, criminals need inside information about their targets to convince them the e-mails are legitimate. They often obtain this information by hacking into an organization's computer network or by coming through other Web sites, blogs, and social networking sites. The Laboratory's CSP counters this threat by using comprehensive Web proxy systems to scan Internet traffic for malicious code and block known, infected Web sites, thereby preventing LLNL users from unknowingly downloading malicious code to their computers.

LLNL's CSP training and awareness campaigns throughout the year give employees information regarding threats, risks, and practical measures to keep LLNL computers and networks secure. These forums show employees the front lines of the cyber security effort. In turn, employees continue to do their part to report suspicious activity to CSP. In addition, CSP works with local cyber investigators from the Federal Bureau of Investigation and partners with other agencies to stay aware of the latest trends and threats. These information exchanges are helping keep LLNL computers, networks, and users safe from cyber attacks.

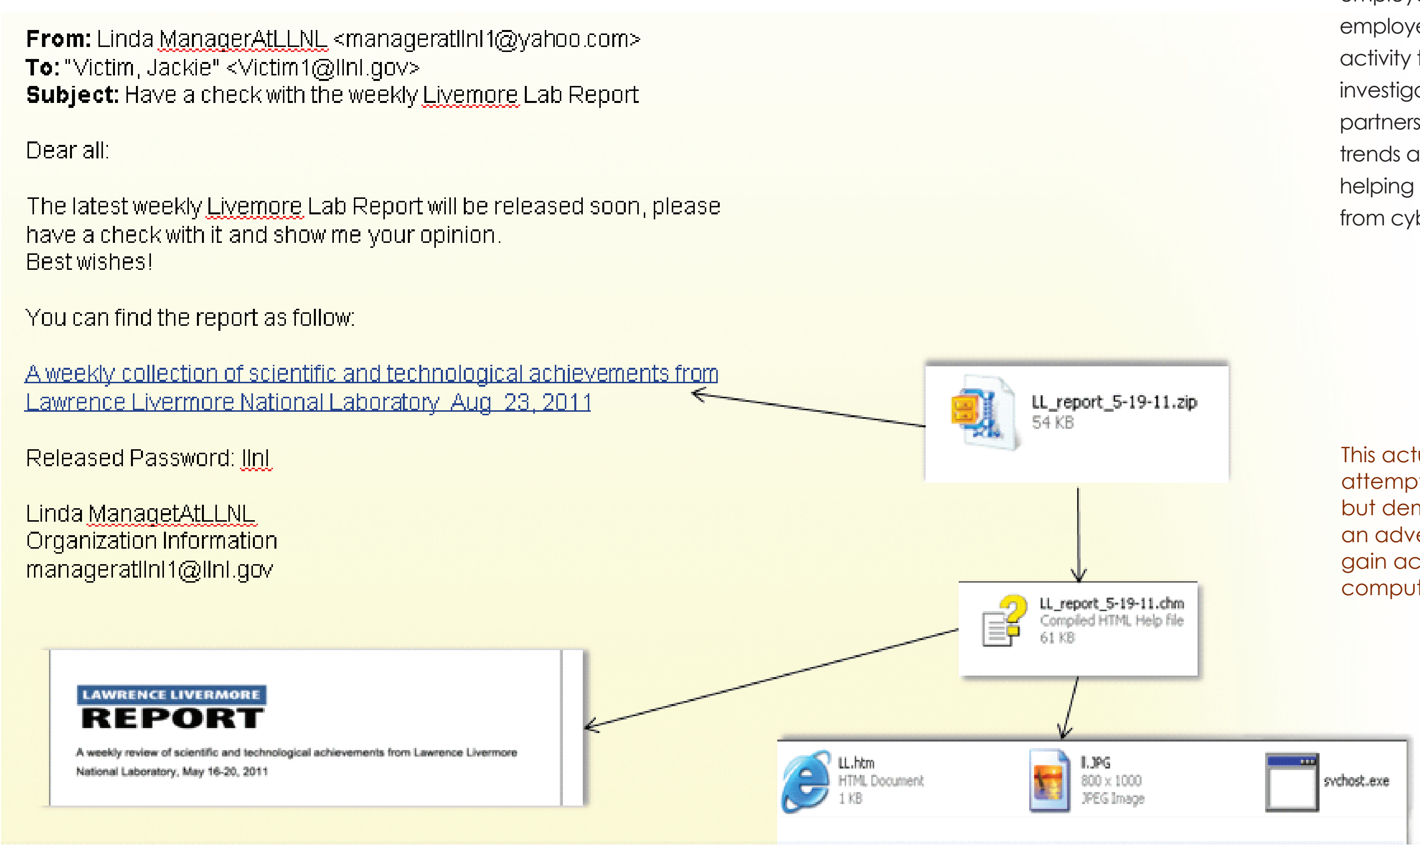

This actual spear-phishing attempt was unsuccessful but demonstrates how an adversary tried to gain access to an LLNL computer.

Unified Communications Capabilities Improve Productivity

ROBYNE TESLICH TESLICH2@LLNL.GOV

Collaboration and mobility are fundamental to an organization's productivity, competitiveness, and ability to attract and retain employees. In 2011, the LLNL Chief Information Officer (CIO) Program defined a strategy to enhance individual, workgroup, and organizational productivity by providing an integrated set of communication capabilities, including telephony and voice, conferencing, presence information, and instant messaging across diverse clients and endpoints (e.g., handsets, mobile phones, laptops, tablets, and desktops). Several multiyear projects were initiated to improve the wireless and wired network infrastructure and provide a foundation on which to expand LLNL's collaboration capabilities. Significant progress has already been made on improving the infrastructure and offering additional services. This unified communications plan will substantially improve employee productivity and give LLNL a competitive edge in hiring and retaining a quality workforce.

Progress in 2011

In the summer of 2011, the CIO Program worked with organizations around LLNL to develop a three-year strategy to deploy integrated communication and collaboration capabilities. The resulting road map defined an approach to deploying Internet Protocol-enabled video, voice, and data services, while adhering to industry standards. The strategy is geared toward device independence and is designed with service delivery and security in mind. It also makes use of partnerships that already exist between LLNL and vendors such as Microsoft Corporation and Cisco Systems, Inc. The new capabilities, collectively called unified communications, connect people, information, and teams in real-time and independent of device or location. Unified communications capabilities enhance productivity, improve agility, and allow a faster response to mission needs and opportunities.

Unified communications services depend on a robust, modern network infrastructure and broad client support to enable collaboration and communication anytime, from anywhere, using varied technology devices. In 2011, LLNL launched two major projects to improve the site's network infrastructure and create a foundation on which to build the new collaboration services. LLNL submitted a proposal and was awarded $1 million from NNSA to expand the wireless network service for employees and guests at select locations at Site 200 and Site 300. With this funding, WiFi was installed in approximately 22 buildings in both Property Protected and Limited Access areas, increasing wireless coverage to 60 buildings by the end of 2011. Buildings were selected and prioritized by LLNL program representatives. Employees and guests can now work in various locations away from their primary office. Additional funding has been requested to expand the wireless network coverage.

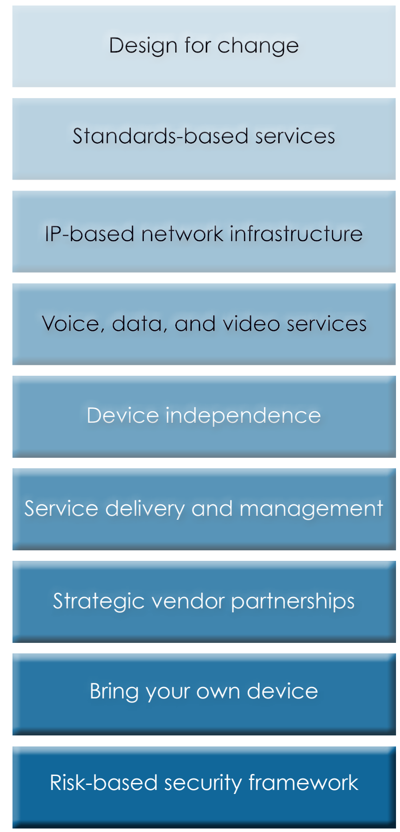

LLNL's unified communications plan will take several years to implement but will significantly improve employees' collaboration and mobility options.

The second project was aimed at improving the security, bandwidth, and capabilities of LLNL's wired network infrastructure to support modern services. In 2011, the network in 16 buildings was upgraded and approximately 188 of the 600+ end-of-life network switches that comprise LLNL's Yellow Network were replaced. With additional funding, network teams will continue to remove outdated network infrastructure by replacing approximately 140 switches. In addition, a single-building pilot will be conducted to demonstrate the feasibility of using wireless technology as the primary office network for employees. Using wireless rather than wired infrastructure provides an option for modernizing building networks at a significantly reduced cost.

Collaboration service enhancements planned in 2012 include desktop video conferencing and Web conferencing capabilities for internal users and external collaborators. Video conferencing will enable users to communicate with an individual or be present at multiperson conferences via full video from their desktop systems. Web conferencing is similar to video conferencing, but it also allows for shared views of desktops, shared presentations, and group document markups to enable users to fully participate in meetings they cannot physically attend.

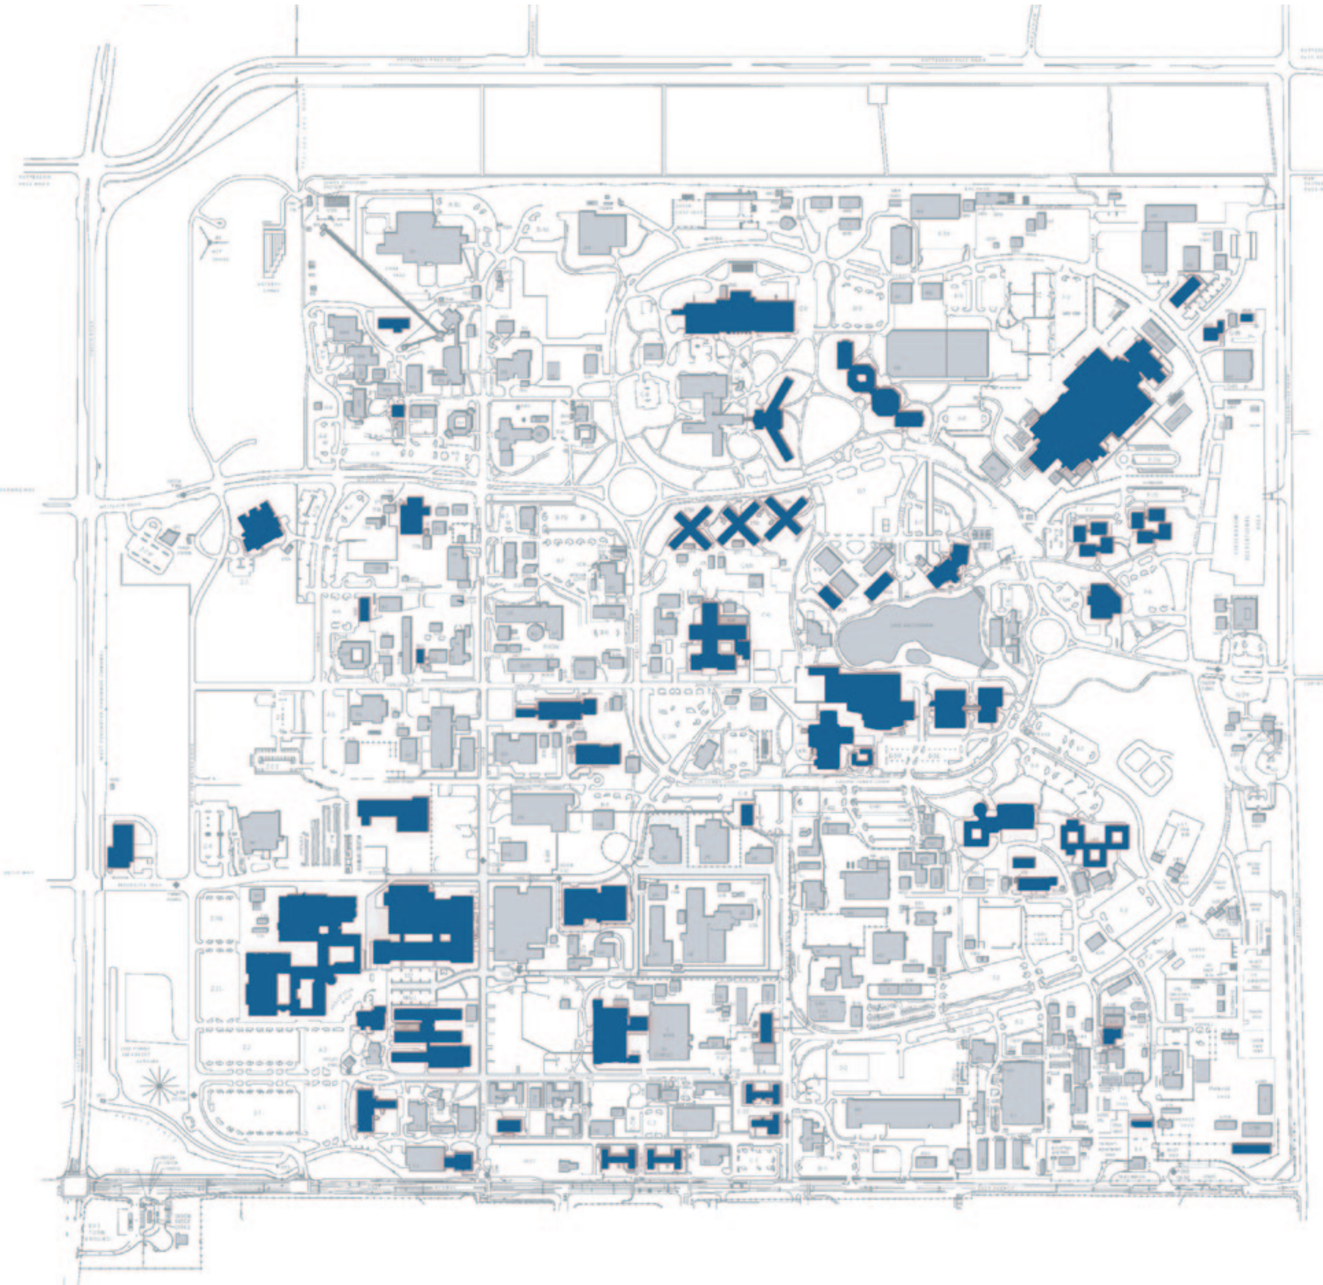

In 2012, LLNL will expand mobile device options beyond BlackBerrys and laptops by initiating a project to support iPhones and iPads. Initially, employees will be limited to accessing e-mail, calendar, and internal Web-based applications on Laboratory-managed versions of the devices, and they will only be able to use the devices in Property Protected areas. However, plans for the short-term include investigating whether or not these devices can be used in Limited Access areas and whether or not personal iPhones/ iPads can be used for business in Property Protected areas. Eventually, Web conferencing and other collaboration services will also be available on iPhones and iPads.

Over the next several years, the CIO Program, with support from the Computation Directorate, will continue to add to the Laboratory's communication and collaboration capabilities. Desktop video conferencing, Web conferencing, and modern mobile device support will be followed by integrated voice, Internet Protocol-based video services, and state-of-the-art teleconferencing solutions. Implementing the vision and strategy defined in 2011 will produce capabilities that set LLNL apart from and ahead of its competitors and allow employees to collaborate and share information.

WiFi is available around the LLNL site in approximately 60 buildings (blue) in 2011.

4Help Service Desk Improves Issue Resolution Rate

MERRY CARTER CARTER11@LLNL.GOV

One key to ensuring a company's efficiency and effectiveness is to keep its users' IT functions available and operational as much as possible. LLNL's 4Help service desk helps meet this need for Laboratory employees and off-site collaborators by providing tier-1 support of desktop systems, mobile devices, software applications, and several enterprise services (e.g., remote access, e-mail, encryption, networks, and password tokens). The goal of 4Help is to offer customers a high level of quality support, through multiple channels, delivered by knowledgeable and courteous staff.

4Help supports a large, secure, heterogeneous computing environment of Windows, Macintosh, and Linux systems; classified and unclassified networks; and myriad scientific applications and hardware. In response to a Lean Six Sigma project, the service desk made several efficiency improvements over the past two years. 4Help now resolves more computer support requests, and survey results show an increase in customer satisfaction. Since October 2010, the percentage of incidents resolved by 4Help technicians increased from 50% to 70%, an industry best-practice level.

Progress in 2011

A Lean Six Sigma project identified several key areas for improvement at the 4Help service desk, including implementing a multilevel service desk with an automated call tree, documenting work procedures and scope of support, improving training programs for new technicians, and offering a remote assistance tool to resolve desktop support issues.

A multilevel call distribution system was implemented to route customers to the service desk technician most experienced in handling the nature of the call. The call system was designed to assign tasks categorically to take advantage of each technician's particular skills (e.g., expertise in Windows, Mac OS X, or Microsoft Office applications). The multilevel structure lets entry-level technicians perform basic password resets and resolve other common issues using a script and, in turn, allows senior technicians to become shift leads, troubleshoot complex issues, and review incidents before they are escalated to tier-2 or tier-3 field support. Shift leads also monitor service desk calls to provide technicians with real-time feedback.

4Help deployed a commercial online knowledge base, called RightAnswers, to document procedures and provide customers with an easy, self-help approach to answering common questions. 4Help staff added LLNL-specific IT solutions and service desk operating procedures to the canned solutions provided by the vendor. RightAnswers also serves as a training resource and reference for new service desk technicians.

4Help began using a product called Bomgar this year to improve LLNL's overall remote-assistance experience. A remote-assistance capability allows a technician to log in and "share a screen" with the customer to troubleshoot and resolve some of the more complicated issues without physically visiting the customer's desk. With Bomgar, the user initiates the remote session, resulting in connections through firewalls that the previous product could not penetrate. Bomgar has increased the successful connection rate between service desk technicians and customers from 40% to more than 95%.

4Help also established a standardized process improvement method where metrics are collected and targets are established and reviewed on a monthly basis. Providing better tools to track, log, and resolve issues has significantly improved the overall efficiency of the 4Help service desk. The cumulative effect of these improvements is a higher resolution rate and increase in customer satisfaction.

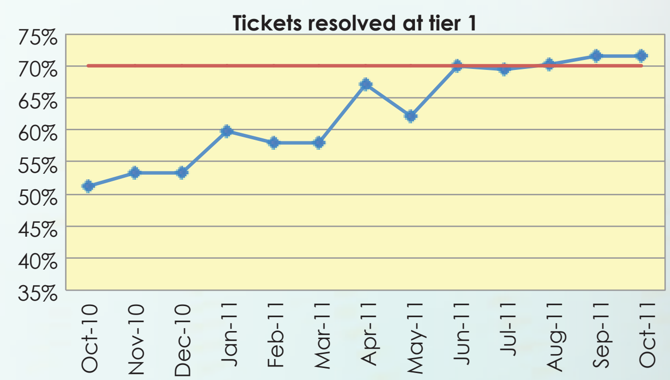

Changes implemented at the 4Help service desk since October 2010 have substantially increased the resolution rate of tier-1 incidents while increasing customer satisfaction.

Appendices

Academic Outreach

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
Ball State University	Irene Livshits	Subcontract	Bootstrap algebraic multigrid methods for solving time-dependent wave and heat equations	ASCR	Robert Falgout
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	Cyber security research: joint proposals	SMS	Celeste Matarazzo
California Polytechnic State University, San Luis Obispo	David Clague	Joint Research	Low-cost microarrays	LDRD	Tom Slezak
Cambridge University	Nikos Nikiforkis	Joint Research	Simulation and modeling using Overture	ASCR Base	Bill Henshaw
Carnegie Mellon University	Christos Faloutsos	Joint Research	Role discovery in networks	LDRD	Brian Gallagher
Carnegie Mellon University	Christos Faloutsos	Subcontract	Mining large, time-evolving data for cyber domains	LDRD	Celeste Matarazzo
Centers for Disease Control, Kenya	Barry Fields	Joint Research	Global biosurveillance	—	Tom Slezak
Chalmers University, Sweden	Sally McKee	Collaboration	Leveraging OpenAnalysis for alias analysis within ROSE	ASC	Dan Quinlan
Colorado State University	Donald Estep	Subcontract	A posteriori error analysis for hydrodynamic systems	LDRD	Carol Woodward
Colorado State University	Michelle Strout and Sanjay Rajopadhye	Collaboration	Program analysis	ASCR	Dan Quinlan
Colorado State University	Donald Estep	Subcontract	Adjoint-based methods for uncertainty quantifications	ASC	Xabier Garaizar

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
Columbia University	Mark Adams	Joint Research	FASTMath: Frameworks, Algorithms, and Scalable Technologies for Mathematics	ASCR SciDAC	Lori Diachin
Columbia University	David Keyes	Joint Research	Towards optimal petascale simulations	ASCR SciDAC	Robert Falgout
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 DOE systems and applications	ASCR	Greg Bronevetsky
Dresden University of Technology	Matthias Mueller	Joint Research	Message-passing interface correctness checking	ASC	Martin Schulz
Dresden University of Technology	Matthias Mueller and Wolfgang Nagel	Joint Research	Semantic debugging of message-passing interface applications	ASCR, ASC, LDRD	Bronis de Supinski
Duke University	Herbert Edelsbrunner	Joint Research	Discrete methods for computing continuous functions	LDRD	Timo Bremer
Électricité de France	Phillippe Lafon	Collaboration	Aeroacoustics	—	Bill Henshaw
Georgia Institute of Technology	Richard Vuduc	Subcontract	Compiler support for reverse computation	ASCR	Dan Quinlan
Georgia Institute of Technology	Richard Fujimoto	Subcontract	Research in reverse computation	LDRD	David Jefferson
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
Indiana University	Thorsten Hoefler	Joint Research	Message-passing interface forum and advanced message-passing interface usage	ASC	Martin Schulz

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
Institute of Computer Science of the Foundation for Research and Technology—Hellas (FORTH), Greece	Angelos Bilas and Dimitris Nikolopoulos	Joint Research	Power optimization for hybrid codes	ASC	Martin Schulz
Institute of Computer Science of the Foundation for Research and Technology—Hellas (FORTH) and University of Crete	Demitrios Nikolopoulos	Joint Research	Power-aware computing for OpenMP programs	ASCR, ASC	Bronis de Supinski
Johns Hopkins University	Allan Boyles	Collaboration	Seismoacoustic modeling for defense-related efforts	DOE	Shawn Larsen
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-Maximilians University of Munich	Dieter Kranzlmueller	Joint Research	Detecting communication patterns to optimize applications	ASCR, ASC	Bronis de Supinski
Ludwig-Maximilians University of Munich	Dieter Kranzlmueller	Joint Research	Message-passing interface tool infrastructure and performance analysis	ASC	Martin Schulz
Norfolk State University	Aftab Ahmad	Joint Research	Cyber security research: joint proposals	SMS	Celeste Matarazzo
North Carolina State University	Frank Mueller	Joint Research	Compressing message-passing interface traces	ASCR, ASC	Bronis de Supinski
North Carolina State University	Vincent Freeh	Joint Research	Power-aware computing for message-passing interface programs	ASCR, ASC	Bronis de Supinski
Northern Arizona University	Paul Keim	Joint Research	Microbial forensics	DTRA	Tom Slezak
Ohio State University	P. Sadayappan and Christophe Alias	Collaboration	Optimizing compiler program analysis	ASCR	Dan Quinlan
Ohio State University	Krishna Kandalla	Joint Research	Use of nonblocking AllReduce in conjugate gradient	ASC CSSE	Ulrike Yang
Ohio University	Yusu Wang	Joint Research	Analysis and visualization of high-dimensional function	LDRD	Timo Bremer

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
Pennsylvania State University	Ludmil Zikatonov and Robert Schechl	Subcontract	Construction of coarse scale numerical models based on energy minimization and applications to upscaling techniques	ASCR Base	Panayot Vassilevski
Pennsylvania State University	Ludmil Zikatonov	Subcontract	Multilevel methods for graph Laplacians and piece-wise constant approximations	ASCR	Panayot Vassilevski
Pennsylvania State University	Jinchao Xu and James Brannick	Subcontract	Algebraic multigrid methods and fast auxiliary space preconditioners	ASCR	Robert Falgout
Princeton University	Adam Burrows	Joint Research	Computational Astrophysics Consortium	ASCR SciDAC	Louis Howell
Purdue University	Jennifer Neville	Joint Research	Correcting bias in statistical tests for network classifier evaluation	LDRD	Brian Gallagher
Purdue University	Saurabh Bagchi	Joint Research	Statistical debugging tools, fault tolerance, scalable checkpointing	ASCR, ASC, LDRD	Bronis de Supinski
Purdue University	Dongbin Xiu	Subcontract	Intrusive and non-intrusive polynomial chaos methods for the quantification of uncertainties in multiphysics models	ASC	Charles Tong
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	Saurabh Bagchi	Joint Research	Anomaly detection and tracking in high performance computing	ASC	Martin Schulz
Purdue University	Mithuna Thottethodi	Joint Research	Low-level message-passing interface	ASC	Martin Schulz
Purdue University	Ziqiang Cai	Summer Faculty	A posteriori error estimates for partial differential equations	ASC	Robert Falgout
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	Joint Research	FASTMath: Frameworks, Algorithms, and Scalable Technologies for Mathematics	ASCR SciDAC	Lori Diachin
Rensselaer Polytechnic Institute	Onkar Sahni	Joint Research	FASTMath: Frameworks, Algorithms, and Scalable Technologies for Mathematics	ASCR SciDAC	Lori Diachin
Rice University	John Mellor-Crummey	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
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	Richard Baraniuk	Subcontract	New compression techniques for high-resolution airborne sensors	LDRD	Yang Liu
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	Capturing node-level behavioral structure in static and dynamic networks	LDRD	Brian Gallagher
Rutgers University	Tina Eliassi-Rad	Subcontract	Cyber situational awareness through host and network analysis	LDRD	Celeste Matarazzo
RWTH Aachen University	Felix Wolf	Joint Research	Message-passing interface performance analysis tools	ASCR, ASC, LDRD	Bronis de Supinski
Southern Methodist University	Thomas Hagstrom	Joint Research	High-order structure grid methods for wave propagation on complex unbounded domains	ASCR Base	Bill Henshaw
Southern Methodist University	Dan Reynolds	Joint Research	Implicit solvers and preconditioning techniques for simulations of magnetohydrodynamics	ASCR SciDAC	Carol Woodward
Southern Methodist University	Dan Reynolds	Subcontract	New time integration methods and support for multiscale solution methods in the LLNL SUNDIALS software suite	ASCR SciDAC	Carol Woodward
Southern Methodist University	Dan Reynolds	Joint Research/ Subcontract	FASTMath: Frameworks, Algorithms, and Scalable Technologies for Mathematics	ASCR SciDAC	Lori Diachin
Stanford University	Jurij Leskovec	Subcontract	Modeling networks with rich node-level attributes	LDRD	Brian Gallagher
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

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
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	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/ SLAC National Accelerator Laboratory	Stuart Marshall and Jacek Becla	Joint Research	Large Synoptic Survey Telescope camera data output hardware and imaging formats	LDRD	Celeste Matarazzo
Technical University of Denmark	Robert Read	Collaboration	Water waves and wave energy generation		Bill Henshaw
Texas A&M University	Nancy Amato	Joint Research	Load balance optimizations	ASC	Martin Schulz
Texas A&M University	Nancy Amato	Collaboration, Lawrence Scholar Program	Parallel graph algorithms	UCOP	Maya Gokhale
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	Bjarne Stroustrup and Lawrence Rauchwerger	Joint Research	Compiler construction and parallel optimizations	ASCR	Dan Quinlan
Texas A&M University	Raytcho Lazarov and Yalchin Efendiev	Joint Research	Algebraic multigrid and Bayesian uncertainty quantification for Darcy and Brinkman problems	ASCR Base	Panayot Vassilevski
Texas State University	Byron Gao	Collaboration	Search user interfaces; clustering	LDRD	David Buttler
Tufts University	Scott MacLachlan	Joint Research	Algebraic multigrid algorithms	ASC	Robert Falgout
UC Berkeley	Domagoj Babic	Joint Research	Formal models of communication protocols	ASCR	Greg Bronevetsky
UC Berkeley	James Demmel	Joint Research	Towards optimal petascale simulations	ASCR SciDAC	Robert Falgout
UC Berkeley	Doug Dreger	Collaboration	Earthquake hazard	IGPP	Shawn Larsen

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
UC Davis	François Gygi	Subcontract	Algorithms for electronic structure and first-principles molecular dynamics simulations using large-scale parallel computers	ASC	Art Mirin
UC Davis	Matt Bishop and Sean Peisert	Joint Research	Cyber security research: usable security in the face of insider threats	—	Celeste Matarazzo
UC Davis	Zhendong Su	Subcontract	ROSE support for rewrapping macro calls	ASCR	Dan Quinlan
UC Davis	François Gygi	Subcontract	Improvements in scalability and capability to the Qbox code	ASC	Erik Draeger
UC Davis	Ken Joy	Joint Research	Visualization and Analytics Center for Enabling Technologies	ASCR SciDAC	Mark Duchaineau
UC Davis	Ken Joy	Subcontract	Improving accuracy and efficiency of 3D aerial video	DoD	Mark Duchaineau
UC Davis	Ken Joy	Collaboration, Lawrence Scholar Program	Discrete multimaterial interface reconstruction for volume fraction data	UCOP	Mark Duchaineau
UC Davis	Ken Joy	Collaboration, Lawrence Scholar Program	Video processing research for the VidCharts and Persistics projects	UCOP	Mark Duchaineau
UC Davis	Nelson Max	Joint Research	3D from wide-area aerial video	DOE Nonproliferation	Mark Duchaineau
UC Davis	Kwan-Liu Ma	Subcontract	Interactive tomographic reconstruction of aerial motion imagery	DoD	Mark Duchaineau
UC Davis	Bernd Hamann	Joint Research	Analysis and visualization of scientific data using topology-based methods	LDRD	Timo Bremer
UC Davis	Bernd Hamann	Joint Research	Performance analysis and visualization	ASC	Martin Schulz
UC Davis	Louise Kellogg	Joint Research	Topological analysis of geological data	LDRD	Timo Bremer
UC Riverside	Michalis Faloutsos	Subcontract	GraphWare: using graphs to identify malware traffic on the Internet based on their community structure and their pack-level behavioral characteristics	LDRD	Celeste Matarazzo
UC San Diego	Allan Snavely and Laura Carrington	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
UC San Diego	Erik Gartzke and Jon Lindsey	Joint Research	Cyber security research: large-scale risks in cyberspace	—	Celeste Matarazzo

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
UC San Diego	Falko Kuester	Joint Research	Large-scale atomistic simulation visualization	ASC	Mark Duchaineau
UC San Diego	Falko Kuester	Joint Research	3D from video	DoD	Mark Duchaineau
UC San Diego	Allan Snavely	Collaboration	Data intensive architectures	LDRD	Maya Gokhale
UC San Diego	Steve Swanson	Collaboration	Persistent memory emulator		Maya Gokhale
UC San Diego	Randy Bank	Subcontract	Solvers for large sparse systems of linear equations arising from discretizations of partial differential equations	ASC	Robert Falgout
UC San Diego	Randy Bank	Subcontract	Numerical solutions of partial differential equations, multilevel iterative methods, and adaptive grid generations	ASC	Robert Falgout and Panayot Vassilevski
UC San Diego	Ido Akkerman	Subcontract	Adding isogeometric analysis capability in a high-order curvilinear research hydrocode	LDRD	Tzanio Kolev
UC San Diego, Scripps Institution of Oceanography	Julie McClean	Collaboration	Ultrahigh-resolution coupled climate simulations	BER	Art Mirin
UC Santa Barbara	Fred Chong and Diana Franklin	Joint Research	Reduction of memory footprints	ASC, LDRD	Bronis de Supinski
UC Santa Cruz	Stan Woosley	Joint Research	Computational Astrophysics Consortium	ASCR SciDAC	Louis Howell
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
United States Army Medical Research Unit, Kenya	John Waitumbi	Joint Research	Pathogen diagnostics	—	Tom Slezak
Universität Karlsruhe	Wolfgang Karl	Joint Research	Hardware transactional memory	ASC	Martin Schulz
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	Martin Schulz
University of British Columbia	Carl Olivier-Gooch	Subcontract	FASTMath: Frameworks, Algorithms, and Scalable Technologies for Mathematics/Mesquite	ASCR SciDAC/ASCR Base	Lori Diachin

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
University of Colorado	Tom Manteuffel	Joint Research	Solution methods for transport problems	ASC	Peter Brown
University of Colorado	Tom Manteuffel and Steve McCormick	Joint Research	Towards optimal petascale simulations	ASCR SciDAC	Robert Falgout
University of Colorado	Steve McCormick, Tom Manteuffel, John Ruge, Marian Brezina, Minho Park, and Toby Jones	Subcontract	Geometric and algebraic multigrid methods for quantum chromodynamics, magnetohydrodynamics, elasticity, transport, and other Department of Energy applications	ASCR, ASC	Robert Falgout
University of Colorado at Boulder	Ken Jansen	Joint Research	FASTMath: Frameworks, Algorithms, and Scalable Technologies for Mathematics	ASCR SciDAC	Lori Diachin
University of Delaware	Richard Braun	Collaboration	Models of the eye	ASCR Base	Bill Henshaw
University of Houston	Yuriy Fofanov	Joint Research	Genomic algorithms	DHS	Tom Slezak
University of Illinois, Urbana-Champaign	William Gropp	Joint Research	Optimization for AMG	LDRD	Martin Schulz
University of Illinois, Urbana-Champaign	Laxmikant Kale	Joint Research	Node mapping optimizations for high performance computing systems	ASC	Martin Schulz
University of Illinois, Urbana-Champaign	William Gropp	Joint Research	Message-passing interface, hybrid programming models	ASCR, ASC, LDRD	Bronis de Supinski
University of Illinois, Urbana-Champaign	Rakesh Kumar	Subcontract	Algorithmic error correction techniques for sparse linear algebra	ASCR	Greg Bronevetsky
University of Illinois, Urbana-Champaign	Rakesh Kumar and Joseph Sloan	Joint Research	Algorithm-specific fault tolerance for sparse linear algebra applications	ASCR	Greg Bronevetsky
University of Illinois, 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/ IBM	Hormozd Gahvari, William Gropp, Luke Olson, and Kirk Jordon	Collaboration	Modeling algebraic multigrid performance on multicore architectures	LDRD	Ulrike Yang
University of Louisville	Yongsheng Liam	Collaboration	Micro-air vehicles	—	Bill Henshaw

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
University of Maryland	Jeff Hollingsworth	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
University of Maryland	Jeffrey Hollingsworth	Joint Research	Autotuning and tool infrastructures	ASCR	Martin Schulz
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 North Carolina	Robert Fowler	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
University of North Carolina	Jan Prins	Joint Research	Efficient OpenMP runtimes for tasking	ASC, LDRD	Bronis de Supinski
University of North Carolina	Jan Prins	Joint Research	OpenMP task scheduling	ASC	Martin Schulz
University of Southern California	Robert Lucas, Mary Hall, and Jacqueline Chame	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
University of Southern California	Gerard Medioni	Subcontract	Activity analysis in wide-area overhead video	DOE Nonproliferation	Mark Duchaineau
University of Tennessee	Jack Dongarra and Shirley Moore	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
University of Tennessee	Jack Dongarra	Joint Research	Empirical tuning	ASCR	Dan Quinlan
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, Austin	Omar Ghattas	Joint Research	Towards optimal petascale simulations	ASCR SciDAC	Robert Falgout

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
University of Turabo, Puerto Rico	Jeffrey Duffany	Joint Research	Cyber security research: analysis and defense of large-scale smart meter networks	—	Celeste Matarazzo
University of Utah	Ganesh Gopalakrishnan	Joint Research	Semantic debugging of message-passing interface applications	ASCR, ASC, LDRD	Bronis de Supinski
University of Utah	Ganesh Gopalakrishnan	Collaboration	Message-passing interface optimizations	ASCR	Dan Quinlan
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	Mark Duchaineau
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 of Utah	Ross Whitaker	Joint Research	Analysis and visualization of high-dimensional function	LDRD	Timo Bremer
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 Henson
University of Wisconsin	Bart Miller	Joint Research	Scalable debugging	ASCR, ASC, LDRD	Bronis de Supinski
University of Wisconsin	Ben Liblit	Joint Research	Scalable debugging	ASCR, ASC	Bronis de Supinski
University of Wisconsin	Karu Sankaralingam	Joint Research	Resilient computing	ASCR, ASC, LDRD	Bronis de Supinski
University of Wisconsin	Bart Miller	Joint Research	Performance tools and tool infrastructures	ASCR	Martin Schulz
University of Wuppertal	Karsten Kahl	Visiting Researcher	Algebraic multigrid algorithms	ASC	Robert Falgout
Utah State University	Steena Monteiro	Subcontract	Test case selection and statistical modeling of the behavior of the Ceph parallel file system	ASCR	Greg Bronevetsky

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
Utah State University	Renée Bryce and Steena Monteiro	Joint Research	Statistical modeling of data-driven applications	ASCR and Lawrence Scholar Program	Greg Bronevetsky
Vienna University of Technology	Markus Schordan	Collaboration	Compiler construction	ASCR	Dan Quinlan
Virginia Institute of Technology	Madhav Marathe	Subcontract	Research in network sciences	SMS	Evi Dube and James Brase
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, LDRD	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	Kirk Cameron	Joint Research	Power optimization and scheduling	ASC	Martin Schulz
Virginia Polytechnic Institute	John Burns	Collaboration	Building simulations, reduced order models, and control	GPIC	Bill Henshaw
Wake Forest University	Errin Fulp	Joint Research	Beyond snorting: extending malicious activity detection with support vector machines over auxiliary session summary features	SMS	Celeste Matarazzo
Worcester Polytechnic University	Homer Walker	Joint Research	Nonlinear solvers and subsurface simulation methods	ASCR SciDAC	Carol Woodward

Publications

Book Chapter

Correa, C.D. and K.-L. Ma, "Visualizing Social Networks," in *Social Network Data Analytics*, C. Aggarwal (Springer), 307–326.

Journal Papers

Abdulla, G., et al., "Practical Data Mining and Machine Learning for Optics Applications: Introduction to the Feature Issue," *Appl. Optics* **50** (22), PDM1–PDM2.

Abt, B., et al., "Complete Genome Sequence of Leadbetterella byssophila Type Strain (4M15)," *Stand Genomic Sci.* **4** (1), 2–12.

Ahrens, J.P., et al., "Data-Intensive Science in the U.S. DOE: Case Studies and Future Challenges," *Comput. Sci. Eng.* **13** (6), 14–23.

Bacuta, C., P.S. Vassilevski, and S.Y. Zhang, "A New Approach for Solving Stokes Systems Arising from a Distributive Relaxation Method," *Numer. Meth. Part. D. E.* **27** (4), 898–914.

Baker, A.H., et al., "Multigrid Smoothers for Ultraparallel Computing," *SIAM J. Sci. Comput.* **33** (5), 2864–2887 (LLNL-JRNL-473191).

Bakker, M., et al., "Computational Challenges in the Solution of Water Resources Problems," *Adv. Water Res.* **34** (9), 1059–1061.

Banks, J.W., "A Note on Compressive Limiting for Two-Material Flows," *Int. J. Numer. Meth. Fl.* **65** (5), 602–608 (LLNL-JRNL-407488).

Banks, J.W. and B. Sjogreen, "A Normal Mode Stability Analysis of Numerical Interface Conditions for Fluid/Structure Interaction," *Commun. Comput. Phys.* **10** (2), 279–304 (LLNL-JRNL-415710).

Banks, J.W., et al., "Two-Dimensional Vlasov Simulation of Electron Plasma Wave Trapping, Wavefront Bowing, Self-Focusing, and Sideloss," *Phys. Plasmas* **18** (5) (LLNL-JRNL-44392).

Beketayev, K., et al., "Topology-Based Visualization of Transformation Pathways in Complex Chemical Systems," *Comput. Graph. Forum* **30** (3), 663–672.

Bennett, J.C., et al., "Feature-Based Statistical Analysis of Combustion Simulation Data," *IEEE T. Vis. Comput. Gr.* **17** (12), 1822–1831 (LLNL-JRNL-515471).

Bhatia, H., et al., "Flow Visualization with Quantified Spatial and Temporal Errors Using Edge Maps," *IEEE T. Vis. Comput. Gr.* **17** (12), (LLNL-JRNL-485511).

Brancaccio, R., et al., "Real-Time Reconstruction for 3D CT Applied to Large Objects of Cultural Heritage," *IEEE T. Nucl. Sci.* **58** (4), 1864–1871.

Bremer, P.T., et al., "Interactive Exploration and Analysis of Large-Scale Simulations Using Topology-Based Data Segmentation," *IEEE T. Vis. Comput. Gr.* **17** (9), 1307–1324.

Brunner, T.A. and T.V. Kolev, "Algebraic Multigrid for Linear Systems Obtained by Explicit Element Reduction," *SIAM J. Sci. Comput.* **33** (5), 2706–2731 (LLNL-JRNL-440891).

Burkhart, S.C., et al., "National Ignition Facility System Alignment," *Appl. Optics* **50** (8), 1136–1157.

Chang, B., A. Greenbaum, and E. Machorro, "Global Error Bounds for the Petrov–Galerkin Discretization of the Neutron Transport Equation," *Numer. Linear Algebr.* **18** (1), 141–154.

Colella, P., et al., "High-Order, Finite-Volume Methods in Mapped Coordinates," *J. Comput. Phys.* **230** (8), 2952–2976 (LLNL-JRNL-422807).

Connors, J.M., E.W. Jenkins, and L.G. Rebholz, "Small-Scale Divergence Penalization for Incompressible Flow Problems via Time Relaxation," *Int. J. Comput. Math.* **88** (15), 3202–3216 (LLNL-JRNL-465054).

Correa, C.D. and K.-L. Ma, "Visibility Histograms and Visibility-Driven Transfer Functions," *IEEE T. Vis. Comput. Gr.* **17** (2), 192–204.

Correa, C.D. and P. Lindstrom, "Towards Robust Topology of Sparsely Sampled Data," *IEEE T. Vis. Comput. Gr.* **17** (12), 1852–1861 (LLNL-JRNL-479881).

Correa, C.D., P. Lindstrom, and P.T. Bremer, "Topological Spines: A Structure-Preserving Visual Representation of Scalar Fields," *IEEE T. Vis. Comput. Gr.* **17** (12), 1842–1851 (LLNL-JRNL-479882).

Correa, C.D., R. Hero, and K.-L. Ma, "A Comparison of Gradient Estimation Methods for Volume Rendering on Unstructured Meshes," *IEEE T. Vis. Comput. Gr.* **17** (3), 305–319.

de Sterck, H., V. Henson, and G. Sanders, "Multilevel Aggregation Methods for Small-World Graphs with Application to Random-Walk Ranking," *Comput. Inform.* **30** (2), 225–246.

Dewald, E.L., et al., "Experimental Demonstration of Early Time, Hohlraum Radiation Symmetry Tuning for Indirect Drive Ignition Experiments," *Phys. Plasmas* **18** (9).

Dobrev, V.A., et al., "Curvilinear Finite Elements for Lagrangian Hydrodynamics," *Int. J. Numer. Meth. Fl.* **65** (11–12), 1295–1310 (LLNL-JRNL-422302).

Dreger, D., et al., "Near-Field Across-Fault Seismic Ground Motions," *B. Seismol. Soc. Am.* **101** (1), 202–221.

Erlandsson, L., et al., "The Microbial Detection Array Combined with Random Phi29-Amplification Used as a Diagnostic Tool for Virus Detection in Clinical Samples," *PLoS ONE* **6** (8) (LLNL-JRNL-480852).

Glenzer, S.H., et al., "Demonstration of Ignition Radiation Temperatures in Indirect-Drive Inertial Confinement Fusion Hohlraums," *Phys. Rev. Lett.* **106** (8).

Gopalakrishnan, G., et al., "Formal Analysis of MPI-Based Parallel Programs," *Commun. ACM* **54** (12), 82–91 (LLNL-JRNL-489171).

Gronow, S., et al., "Complete Genome Sequence of *Bacteroides salanitronis* Type Strain (BL78T)," *Stand Genomic Sci.* **4** (2), 191–199.

Gurung, T., et al., "LR: Compact Connectivity Representation for Triangle Meshes," *ACM T. Graphic.* **30** (4).

Gurung, T., et al., "SQuad: Compact Representation for Triangle Meshes," *Comput. Graph. Forum* **30** (2), 355–364 (LLNL-JRNL-461651).

Haan, S.W., et al., "Point Design Targets, Specifications, and Requirements for the 2010 Ignition Campaign on the National Ignition Facility," *Phys. Plasmas* **18** (5).

Harris, D.B. and D.A. Dodge, "An Autonomous System for Grouping Events in a Developing Aftershock Sequence," *B. Seismol. Soc. Am.* **101** (2), 763–774.

Hauser, J., et al., "A Probabilistic Seismic Model for the European Arctic," *J. Geophys. Res.-Sol. Ea.* **116**.

Hauser, L.J., et al., "Complete Genome Sequence and Updated Annotation of *Desulfovibrio alaskensis* G20," *J. Bacteriol.* **193** (16), 4268–4269.

Hoefler, T., et al., "The Scalable Process Topology Interface of MPI 2.2," *Concurr. Comp.-Pract. E.* **23** (4), 293–310 (LLNL-JRNL-438133).

Hsu, W.-H., K.-L. Ma, and C.D. Correa, "A Rendering Framework for Multiscale Views of 3D Models," *ACM T. Graphic.* **30** (6).

Jiao, Y.Q., et al., "Identification of Biofilm Matrix-Associated Proteins from an Acid Mine Drainage Microbial Community," *Appl. Environ. Microb.* **77** (15), 5230–5237 (LLNL-JRNL-506873).

Kramer, K.J., et al., "Fusion–Fission Blanket Options for the LIFE Engine," *Fusion Sci. Technol.* **60** (1), 72–77.

Landen, O.L., et al., "Capsule Implosion Optimization During the Indirect-Drive National Ignition Campaign," *Phys. Plasmas* **18** (5).

Lang, E., et al., "Complete Genome Sequence of *Weeksella virosa* Type Strain (9751T)," *Stand Genomic Sci.* **4** (1), 81–90.

Lindl, J.D., et al., "Progress Towards Ignition on the National Ignition Facility," *Nucl. Fusion* **51** (9).

Liu, B., "Dislocation Interactions and Low-Angle Grain Boundary Strengthening," *Acta Mater.* **59** (19), 7125–7134.

Lloyd, S. and Q.O. Snell, "Accelerated Large-Scale Multiple Sequence Alignment," *BMC Bioinformatics* **12**, 466.

Lloyd-Davies, E.J., et al., "The XMM Cluster Survey: X-Ray Analysis Methodology," *Mon. Not. R. Astron. Soc.* **418** (1), 14–53.

Lott, A.P. and E. Howard, "Fast Iterative Solver for Convection-Diffusion Systems with Spectral Elements," *Numer. Meth. Part. D. E.* **27** (2), 231–254.

Maxwell, R.M., et al., "Development of a Coupled Groundwater-Atmosphere Model," *Mon. Weather Rev.* **139** (1), 96–116.

McClean, J.L., et al., "A Prototype Two-Decade Fully Coupled Fine-Resolution CCSM Simulation," *Ocean Model.* **39** (1–2), 10–30 (LLNL-JRNL-459231).

McFadden, G.B., S.R. Coriell, and P.A. Lott, "Onset of Morphological Instability in Two Binary Liquid Layers," *Phys. Fluids* **23** (4).

McLoughlin, K., "Microarrays for Pathogen Detection and Analysis," *Briefings in Functional Genomics* (LLNL-JRNL-482671).

Ndoye, M., et al., "A Recursive Multiscale Correlation-Averaging Algorithm for an Automated Distributed Road-Condition-Monitoring System," *IEEE T. Intell. Transp.* **12** (3), 795–808.

Negres, R.A., et al., "Exploration of the Multiparameter Space of Nanosecond-Laser Damage Growth in Fused Silica Optics," *Appl. Optics* **50** (22), D12–D20 (LLNL-JRNL-480354).

Neville, J., et al., "Correcting Evaluation Bias of Relational Classifiers with Network Cross Validation," *Knowl. Inf. Syst: Int. J.* (LLNL-JRNL-455699).

Olson, L.N. and J.B. Schroder, "Smoothed Aggregation Multigrid Solvers for High-Order Discontinuous Galerkin Methods for Elliptic Problems," *J. Comput. Phys.* **230** (18), 6959–6976.

Olson, L.N., J.B. Schroder, and R.S. Tuminaro, "A General Interpolation Strategy for Algebraic Multigrid Using Energy Minimization," *SIAM J. Sci. Comput.* **33** (2), 966–991.

Paschek, D., R. Day, and A.E. Garcia, "Influence of Water–Protein Hydrogen Bonding on the Stability of Trp-cage Miniprotein: A Comparison Between the TIP3P and TIP4P-Ew Water Models," *Phys. Chem. Chem. Phys.* **13** (44), 19840–19847.

Patra, P.K., et al., "TransCom Model Simulations of CH4 and Related Species: Linking Transport, Surface Flux and Chemical Loss with CH4 Variability in the Troposphere and Lower Stratosphere," *Atmos. Chem. Phys.* **11**, 12813–12837 (LLNL-JRNL-485321).

Pon, R.K., et al., "Measuring the Interestingness of Articles In a Limited User Environment," *Inform. Process. Manage.* **47** (1), 97–116 (LLNL-JRNL-407783).

Pukall, R., et al., "Complete Genome Sequence of *Deinococcus maricopensis* Type Strain (LB-34T)," *Stand Genomic Sci.* **4** (2), 163–172.

Sanchez, P.D., et al., "Study of $B \rightarrow \varpi l \nu$ and $B \rightarrow \rho \lambda \nu$ decays and determination of $|V_{ub}|$," *Phys. Rev. D* **83** (3).

Sanchez-Lavega, A., et al., "Long-Term Evolution of the Aerosol Debris Cloud Produced by the 2009 Impact on Jupiter," *Icarus* **214** (2), 462–476.

Santer, B.D., et al., "Separating Signal and Noise in Atmospheric Temperature Changes: The Importance of Timescale," *J. Geophys. Res.-Atmos.* **116**.

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

Shalf, J., D. Quinlan, and C. Janssen, "Rethinking Hardware-Software Codesign for Exascale Systems," *Computer* **44** (11), 22–30.

Sjogreen, B. and N.A. Petersson, "A Fourth Order Accurate Finite Difference Method for the Elastic Wave Equation in Second Order Formulation," *J. Sci. Comput.* (LLNL-JRNL-483427).

Sonnad, V. and C.A. Iglesias, "Efficient Algorithm for Generating Spectra Using Line-by-Line Methods," *High Energ. Dens. Phys.* **7** (2), 43–46 (LLNL-JRNL-461678).

Stengle, G. and J. McEnerney, "A Semialgebraic Closure for Commutative Algebra," *J. Pure Appl. Algebra* **215** (9), 2257–2261.

Summa, B., et al., "Interactive Editing of Massive Imagery Made Simple: Turning Atlanta Into Atlantis," *ACM T. Graphic.* **30** (2).

Tang, M., J.F. Belak, and M.R. Dorr, "Anisotropic Phase Boundary Morphology in Nanoscale Olivine Electrode Particles," *J. Phys. Chem. C* **115** (11), 4922–4926 (LLNL-JRNL-457673).

Torres, C., et al., "LAVA: An Open-Source Approach to Designing LAMP (Loop-Mediated Isothermal Amplification) DNA Signatures," *BMC Bioinformatics* **12** (LLNL-JRNL-413646).

Town, R.P.J., et al., "Analysis of the National Ignition Facility Ignition Hohlraum Energetics Experiments," *Phys. Plasmas* **18** (5).

Ulmer, C., et al., "Massively Parallel Acceleration of a Document-Similarity Classifier to Detect Web Attacks," *J. Parallel Distr. Com.* **71** (2), 225–235 (LLNL-JRNL-422668).

Wang, W., et al., "Construction of Low Dissipative High-Order Well-Balanced Filter Schemes for Non-Equilibrium Flows," *J. Comput. Phys.* **230** (11), 4316–4335.

Williams, S., et al., "Adaptive Extraction and Quantification of Geophysical Vortices," *IEEE T. Vis. Comput. Gr.* **17** (12), 2088–2095.

Wilson, B.G. and V. Sonnad, "A Note on Generalized Radial Mesh Generation for Plasma Electronic Structure," *High Energ. Dens. Phys.* **7** (3), 161–162.

Yelton, A.P., et al., "A Semi-Quantitative, Synteny-Based Method to Improve Functional Predictions for Hypothetical and Poorly Annotated Bacterial and Archaeal Genes," *PLoS Comput. Biol.* **7** (10).

Zemla, A.T. and B.W. Segelke, "Practical Applications of Structural Genomics Technologies for Mutagen Research," *Mutat. Res.-Gen. Tox. En.* **722** (2), 165–170.

Zemla, A.T., et al., "StralSV: Assessment of Sequence Variability within Similar 3D Structures and Application to Polio RNA-Dependent RNA Polymerase," *BMC Bioinformatics* **12** (226) (LLNL-JRNL-463671).

Zhang, W., "CASTRO: A New Compressible Astrophysical Solver II: Gray Radiation Hydrodynamics," *Astrophys. J. Suppl. S.* **196** (2).

Conference Proceedings

Ames, S., M.B. Gokhale, and C. Maltzahn, "QMDS: A File System Metadata Service Supporting a Graph Data Model-Based Query Language," *International Conference on Networking Architecture and Storage*, (IEEE), **6**, 268–277, (LLNL-CONF-454941).

Anastasiu, D.C., B.J. Gao, and D. Buttler, "A Framework for Personalized and Collaborative Clustering of Search Results," *20th International Conference on Information and Knowledge Management*, (ACM), 573–582, (LLNL-CONF-461652).

Anastasiu, D.C., B.J. Gao, and D. Buttler, “ClusteringWiki: Personalized and Collaborative Clustering of Search Results,” *34th International ACM SIGIR Conference on Research and Development in Information*, (ACM), 1263–1264 (LLNL-CONF-461652).

Andrzejewski, D. and D. Buttler, “Latent Topic Feedback for Information Retrieval,” *17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, (ACM), 600–608, (LLNL-CONF-471258).

Awwal, A.A.S., et al., “Image Processing and Control of a Programmable Spatial Light Modulator for Optic Damage Protection,” *SPIE*, **7916**.

Baker, A.H., M. Schulz, and U.M. Yang, “On the Performance of an Algebraic Multigrid Solver on Multicore Clusters,” *VECPAR 2010: 9th International Meeting on High Performance Computing for Computational Science*, (Lect. Notes Comput. Sc., Springer), **6449**, 102–115, (LLNL-CONF-429864).

Bennett, J.C., et al., “Feature-Based Statistical Analysis of Combustion Simulation Data,” *IEEE Visualization Conference (Vis)/IEEE Information Visualization Conference (InfoVis)*, (IEEE T. Vis. Comput. Gr.), **17**, 1822–1831.

Bhatele, A. and L.V. Kale, “Heuristic-Based Techniques for Mapping Irregular Communication Graphs to Mesh Topologies,” *Workshop on Extreme Scale Computing Application Enablement: Modeling and Tools at HPCC (ESCAPE 2011)*, (IEEE), (LLNL-CONF-491311).

Bhatele, A., et al., “Avoiding Hot Spots on Two-Level Direct Networks,” *International Conference for High Performance Computing, Networking, Storage and Analysis (Supercomputing)*, (ACM), (LLNL-CONF-491454).

Biswas, S., et al., “Exploiting Data Similarity to Reduce Memory Footprints,” *25th International Parallel and Distributed Processing Symposium (IPDPS 2011)*, (IEEE), (LLNL-CONF-468331).

Brunton, G., et al., “The Shaping of a National Ignition Campaign Pulsed Waveform,” *IAEA 2011*.

Correa, C.D. and P. Lindstrom, “Towards Robust Topology of Sparsely Sampled Data,” *IEEE Visualization Conference (Vis)/IEEE Information Visualization Conference (InfoVis)*, (IEEE T. Vis. Comput. Gr.), **17**, 1852–1861.

Correa, C.D., P. Lindstrom, and P.T. Bremer, “Topological Spines: A Structure-Preserving Visual Representation of Scalar Fields,” *IEEE Visualization Conference (Vis)/IEEE Information Visualization Conference (InfoVis)*, (IEEE T. Vis. Comput. Gr.), **17**, 1842–1851.

de Sterck, H., V. Henson, and G. Sanders, “Multilevel Aggregation Methods for Small-World Graphs with Application to Random-Walk Ranking,” *International Workshop on Information Retrieval by Matrix Methods on Supercomputer Systems (IRRM 2010)/24th ACM International Conference on Supercomputing (ICS 2010)*, (Comput. Inform., Slovak Acad. Sciences Inst. Informatics), **30**, 225–246.

Dobrev, V.A., et al., “Curvilinear Finite Elements for Lagrangian Hydrodynamics,” *Conference on Numerical Methods for Multimaterial Fluids and Structures*, (Int. J. Numer. Meth. Fl., Wiley–Blackwell), **65**, 1295–1310.

Fedorov, M.A., et al., “Control System for Cryogenic THD Layering at the National Ignition Facility,” *13th International Conference on Accelerator and Large Experimental Physics Control Systems (ICALEPCS2011)*, (ICALEPCS/JACoW), 1236–1239, (LLNL-CONF-505491).

Glenn, S., et al., “Advanced Gated X-Ray Imagers for Experiments at the National Ignition Facility,”*SPIE Penetrating Radiation Systems and Applications XII*, (SPIE), **8144**, (LLNL-PROC-494311).

Harrison, C.D., H. Childs, and K.P. Gaither, “Data-Parallel Mesh Connected Components Labeling and Analysis,” *Eurographics Symposium on Parallel Graphics and Visualization*, (Eurographics Association), (LLNL-CONF-466493).

Heebner, J., et al., “Programmable Beam Spatial Shaping System for the National Ignition Facility,” *SPIE*, (SPIE), **7916**, (LLNL-PROC-469093).

Hilbrich, T., et al., “Order Preserving Event Aggregation in TBONs,” *EuroMPI 2011*, (Springer–Verlag), (LLNL-CONF-490684).

Kale, V., A. Bhatele, and W.D. Gropp, “Weighted Locality Sensitive Scheduling for Mitigating Noise on Multicore Clusters,” *18th Annual IEEE International Conference on High Performance Computing (HiPC)*, (IEEE), (LLNL-CONF-492091).

Kramer, K.J., et al., “Fusion–Fission Blanket Options for the LIFE Engine,” *19th Topical Meeting on the Technology of Fusion Energy (TOFE-19)*, (Fusion Sci. Technol., Amer. Nuclear Soc.), **60**, 72–77.

Li, D., et al., “Scalable Memory Registration for High Performance Networks Using Helper Threads,” *ACM International Conference on Computing Frontiers (CF 2011)*, (ACM), (LLNL-CONF-476531).

Lagin, L., et al., “Status of the National Ignition Campaign and National Ignition Facility Integrated Computer Control System,” *ICALEPCS 2011*, (LLNL-CONF-475299).

Laguna, I., et al., “Large-Scale Debugging of Parallel Tasks with AutomaDeD,” *Supercomputing 2011 (SC11)*, (ACM), (LLNL-CONF-486911).

Moody, A., D.H. Ahn, and B.R. de Supinski, “Exascale Algorithms for Generalized MPI_Comm_split,” *EuroMPI 2011*, (Springer–Verlag), (LLNL-CONF-484653).

Pirkelbauer, P., D. Dechev, and B. Stroustrup, “Support for the Evolution of C++ Generic Functions,” *3rd International Conference on Software Language Engineering*, (Lect. Notes Comput. Sc., Springer–Verlag Berlin), **6563**, 123–142.

Pirkelbauer, P., et al., “Runtime Detection of C-Style Errors in UPC Code,” *Partitioned Global Address Space Models (PGAS)*, (ACM), (LLNL-CONF-502592).

Rodgers, A.J., et al., “Improving Ground Motion Simulation Capabilities for Underground Explosion Monitoring: Coupling Hydrodynamic-to-Seismic Solvers and Studies of Emplacement Conditions,” *Monitoring Research Review*, (NNSA/NA-22), (LLNL-PROC-491834).

Rountree, B., et al., “Practical Performance Prediction Under Dynamic Voltage Frequency Scaling,” *2nd International Green Computing Conference (IGCC11)*, (IEEE), (LLNL-CONF-476493).

Solomonik, E., A. Bhatele, and J. Demmel, “Improving Communication Performance in Dense Linear Algebra via Topology Aware Collectives,” *International Conference for High Performance Computing, Networking, Storage and Analysis (Supercomputing)*, (ACM), (LLNL-CONF-491442).

Szebenyi, Z., et al., “Reconciling Sampling and Direct Instrumentation for Unintrusive Call-Path Profiling of MPI Programs,” *25th International Parallel and Distributed Processing Symposium (IPDPS 2011)*, (IEEE), **57**, (LLNL-CONF-468224).

Totoni, E., et al., “Simulation-Based Performance Analysis and Tuning for a Two-Level Directly Connected System,” *17th IEEE International Conference on Parallel and Distributed Systems*, (IEEE), (LLNL-CONF-500821).

Town, R.P.J., et al., “Analysis of the National Ignition Facility Ignition Hohlraum Energetics Experiments,” *52nd Annual Meeting of the APS Division of Plasma Physics*, (Phys. Plasmas, Amer. Inst. Physics), **18**.

Van Essen, B.C., et al., “Energy-Efficient Specialization of Functional Units in a Coarse-Grained Reconfigurable Array,” *19th ACM/SIGDA International Symposium on Field Programmable Gate Arrays*, (ACM), 107–110, (LLNL-CONF-463475).

Vo, A., et al., “Large-Scale Verification of MPI Programs Using Lamport Clocks with Lazy Update,” *20th International Conference on Parallel Architectures and Compilation Techniques (PACT-2011)*, (ACM), (LLNL-CONF-476497).

Whitlock, B.J., J.M. Favre, and J.S. Meredith, “Parallel In Situ Coupling of Simulation with Fully Featured Visualization System,” *Eurographics Symposium on Parallel Graphics and Visualization*, (Eurographics Association), (LLNL-CONF-466494).

Wilhelmsen, K., et al., “2011 Status of the Automatic Alignment System for the National Ignition Facility,” *IAEA 2011*, (LLNL-CONF-491497).

Williams, S., et al., “Adaptive Extraction and Quantification of Geophysical Vortices,” *IEEE Visualization Conference (Vis)/IEEE Information Visualization Conference (InfoVis)*, (IEEE T. Vis. Comput. Gr.), **17**, 2088–2095.

Wong, M., et al., “A Case for Including Transactions in OpenMP,” *International Workshop on OpenMP*, (Springer), **6132**, 149–160, (LLNL-CONF-422888).

Industrial Collaborators

COMPANY	TOPIC	LLNL CONTACT
Adaptive Computing Enterprises, Inc.	Moab workload manager	Don Lipari
Allinea Software	Scalable debugging infrastructure	Greg Lee
Appro	Scalable capacity clusters	Matt Leininger and Trent D'Hooge
Arista Networks	Low-latency Ethernet networks	Matt Leininger
Cisco Systems, Dell, DataDirect Networks, Intel, NetApp, Mellanox Technologies, QLogic, Red Hat, Oracle, and Supermicro	Hyperion collaboration	Matt Leininger
Commissariat à l'Énergie Atomique	SLURM resource management software	Don Lipari
Cray	Scalable systems, open-source strategy	Tom Spelce
DataDirect Networks	RAID 6 research and development for I/O systems	Mark Gary
Dell Computers	Scalable capacity clusters	Matt Leininger
IBM	BlueGene/Q: the evaluation of transactional memory and the modeling and optimization of algebraic multigrid	Martin Schulz
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	I/O forwarding for Linux clusters using the 9P protocol	Jim Garlick
IBM	Scalable systems, multiple areas	Kim Cupps and others
IBM Research	Task scheduling with setup times	Deepak Rajan
IBM Research and Carnegie Mellon University	Mining large time-evolving graphs for proximity queries	Deepak Rajan
IBM Research, HP Labs, and Bank of America	Scheduling heterogenous jobs in MapReduce environments	Deepak Rajan
InfiniBand Trade Association	InfiniBand specifications body	Ira Weiny
Intel Corporation	Many integrated core	Greg Lee
Krell Institute and Argo Navis Technologies	Open | SpeedShop development and support and the component-based tool framework	Martin Schulz
NetApp	High performance I/O systems	Marc Stearman and Mark Gary
Obsidian Strategics	BGFC—Next-generation InfiniBand subnet manager	Ira Weiny
OpenFabrics Alliance, Mellanox Technologies, and QLogic	OpenFabrics enterprise distribution	Ira Weiny and Matt Leininger
Open SFS, Inc.	Lustre file system development and deployment	Terri Quinn, Pam Hamilton, and Chris Morrone
OpenWorks	Valgrind memory tool and threading tool development	John Gyllenhaal
Red Hat	Operating systems	Mark Grondona
Rogue Wave Software	TotalView parallel debugger scalability and enhanced memory tools	Dong Ahn and Scott Futral
Samplify Systems, Inc.	Data compression	Peter Lindstrom
SchedMD	SLURM resource management software	Don Lipari
Siemens	Wind power	Bill Henshaw
WhamCloud	Lustre file system development and support	Kim Cupps, Chris Morrone, and Marc Stearman