
COMPUTATION DIRECTORATE
2008 ANNUAL REPORT
LAWRENCE LIVERMORE NATIONAL LABORATORY

ABOUT THE COVER

In a unique collaboration, Lawrence Livermore National Laboratory has teamed with 10 computing industry leaders to accelerate the development of powerful next-generation Linux clusters in a project dubbed Hyperion. Hyperion brings together Dell, Intel, Supermicro, QLogic, Cisco, Mellanox, DDN, Sun, LSI and RedHat to create a large-scale test bed for high-performance computing technologies critical to the National Nuclear Security Administration's work to maintain the aging U.S. nuclear weapons stockpile without underground testing. Hyperion also advances industry's ability to make petaFLOPS (quadrillion floating-point operations per second) computing and storage more accessible for commerce, industry, and research and development. (Cover design by Daniel Moore)

PRODUCTION TEAM

LLNL Associate Director for Computation
Dona Crawford

Deputy Associate Directors
Ronald Schmucker (Principal Deputy), Trish Damkroger, and Michel McCoy

SCIENTIFIC EDITORS
James Guse and Benjamin Grover

ART DIRECTOR
Daniel Moore

PRODUCTION EDITOR
Deanna Willis

PROOFREADERS
Carolin Middleton and Pamela MacGregor

PHOTOGRAPHERS
Nancy Rutter and Jacquelin McBride

PRINT PRODUCTION
Charlie Arteago, Jr., and Monarch Print Copy and Design Solutions

LLNL-TR-411728

Contents

Message from the Associate Director 2

Section 1: High-Performance Computing 7

1.00 The Three-Curve Strategy Pays Off 7

1.01 Hyperion Project Begins Collaboration to Develop an Advanced Technology Cluster Test Bed 9

1.02 Operational Efficiencies Lead to Increased Return on HPC Investments 10

1.03 Common Capacity Computing Reduces Cluster Deployment Time and Cost 12

1.04 Open|SpeedShop Provides Insight into the Performance of Parallel Applications 13

1.05 Sequoia Grows from Conception to Contract 14

1.06 HPC Storage Infrastructure Expands in Preparation for Dawn 16

1.07 Stack Trace Analysis Tool Enables Debugging on More than 200,000 Processors 18

Section 2: Research and Advanced Development 19

2.00 Research Activities Typify Scientists' Wide-Ranging Expertise........ 19

2.01 Phase-Field Models Solve Medium-Scale Materials Simulation Problems 21

2.02 ISCR Charts a New Course for Scientific Computing 22

2.03 Overture Software Gains New Parallel Capabilities and Widens Its Application Base 24

2.04 DocEx Increases the Speed and Ease of Text Analysis 26

2.05 New Morse Theory-Based Technique Brings Clarity to Turbulent Combustion 27

2.06 Machine-Learning Algorithms Aid Cyber-Security Defense 28

2.07 New Parallel Methods Help Solve Inherently Sequential Problems 30

Section 3: Applications Software Development 31

3.00 Computer Scientists Dramatically Impact Laboratory Programs One Software Application at a Time 31

3.01 Shot Data Analysis Engine Automates Analysis of NIF Diagnostics 33

3.02 EMSolve Leads the Way in Electromagnetic Analysis...................... 34

3.03 Computation Helps Russia Secure Nuclear Materials...................... 36

3.04 Biodefense Knowledge Management System Addresses Knowledge Discovery Challenges 37

3.05 VisIt Project Partners on Visualization and Analysis of Massive Data 38

3.06 Software Quality Transforms into a Valued Business Asset 40

Section 4: Information Technology and Security 41

4.00 Information Technology Evolution Helps Modernize and Protect the Institution 41

4.01 Consolidation Establishes a Foundation to Deliver Next-Generation Network Capabilities 43

4.02 IT Support Centralization Lowers Costs and Improves Service 44

4.03 Major Cyber-Security Incident Drives Transformational Improvements in Security System Management............................ 46

4.04 Facility Team Creates Energy Efficiencies and Turns Megawatts into PetaFLOPS 48

Appendices 49

5.01 Industrial Collaborators 50

5.02 Academic Outreach 51

5.03 Publications 57

5.04 Acronyms 61

DONA CRAWFORD,
ASSOCIATE DIRECTOR,
COMPUTATION DIRECTORATE

Computation: Mobilizing Science and Action for National and Global Change

In his book Pasteur's Quadrant: Basic Science and Technological Innovation (1997: Brookings), Donald Stokes describes why national laboratories, like Lawrence Livermore National Laboratory (LLNL), are best suited to move the nation and world forward. Stokes argues that research can help fulfill our societal needs and enhance global economic health by establishing balance between basic science and technological innovation. The Computation Directorate, with its myriad expertise in algorithms research, high-performance computing, applications development, information technology, and security infrastructure is one of the most relevant and able organizations of its kind and strikes an ideal balance, as defined by Stokes, between the quest for fundamental scientific understanding and practical use.

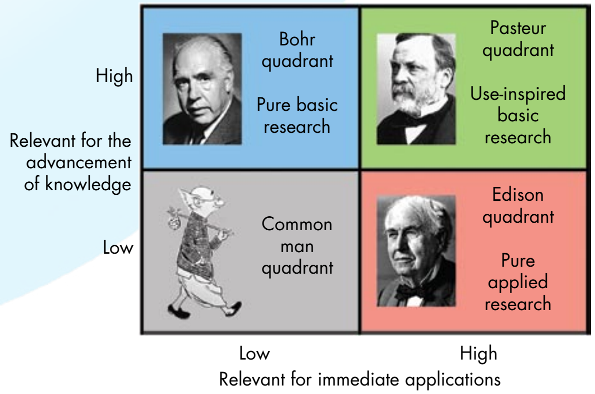

Donald Stokes contends that the quest for fundamental understanding can coincide with consideration of the practical use for research results. Some research is undertaken to improve understanding of a phenomenon, with no regard to specific use (e.g., Neils Bohr's work on atomic structure). Other research focuses strictly on developing applied uses (e.g., Thomas Edison's invention of the phonograph). Stokes argues for a hybrid type of research, typified by the work of Louis Pasteur, as an important target of public funds. The work of the Computation Directorate fits most accurately into this category.

Whether a computer is simulating the aging and performance of a nuclear weapon, the folding of a protein, or the probability of rainfall over a particular mountain range, the necessary calculations can be enormous. Our computers help researchers answer these and other complex problems, and each new generation of system hardware and software widens the realm of possibilities. Building on Livermore's historical excellence and leadership in high-performance computing, Computation added more than 331 trillion floating-point operations per second (teraFLOPS) of power to LLNL's computer room floors in 2008. In addition, Livermore's next big supercomputer, Sequoia, advanced ever closer to its 2011–2012 delivery date, as architecture plans and the procurement contract were finalized. Hyperion, an advanced technology cluster test bed that teams Livermore with 10 industry leaders, made a big splash when it was announced during Michael Dell's keynote speech at the 2008 Supercomputing Conference. The *Wall Street Journal* touted Hyperion as a "bright spot amid turmoil" in the computer industry. Computation continues to measure and improve the costs of operating LLNL's high-performance computing systems by moving hardware support in-house, by measuring causes of outages to apply resources asymmetrically, and by automating most of the account and access authorization and management processes. These improvements enable more dollars to go toward fielding the best supercomputers for science, while operating them at less cost and greater responsiveness to the customers.

Computation led diverse efforts across many LLNL mission areas and made important strides in application development. New large-text-corpus analysis technologies reduce from weeks to hours the time needed by key customers in the intelligence community to exploit large data sets. LLNL staff specified and validated codes used to track and monitor nuclear materials in the Russian Federation, thus fulfilling key U.S. treaty goals. A new workflow-based analysis engine automatically processes data from National Ignition Facility (NIF) shot diagnostics, which enables the National Ignition Campaign goal of three shots per day. LLNL's premier data analysis and visualization tool (VisIt) supplanted other less-capable tools in the research community and became "open source" while preserving traditional support for "funded" customers.

Computation's research this year led to new insights in hydrogen fuel efficiency, improved laser-optics maintenance, and the most effective prediction of intruder/malicious traffic on our networks yet achieved by automated analysis. In addition, new algorithms are enabling scientists to more effectively apply the massively parallel computing platforms of today and tomorrow to the radiation-transport problem—a key capability to securing the nation through radiation detector networks and portal monitoring.

This myriad activity is not accomplished in a bubble. The Laboratory, and Computation as a microcosm, is affected by and must respond and react to several outside factors, including directives from our parent organization, Lawrence Livermore National Security, LLC, and the fluctuations and uncertainties of a flailing economy and federal budget. This year, LLNL formally reduced its workforce—an action that had been avoided since the 1970s. Computation suffered a 16.8% reduction in staff. This change demanded even stronger adherence to one of our Computation and Laboratory values: simultaneous excellence in science, technology, and operations. Our employees have been asked to do more with less—more work with fewer people and less financial resources—and we have risen to the challenge.

Given the Laboratory's fiscal conservativism, Computation sought to implement operational efficiencies with renewed vigor. The automation and centralization of network and cyber-security practices across LLNL's 20,000+ computers supported an intrusion response and recovery and laid the groundwork for prevention of future incidents—all of which was achieved with an impressively constrained resource commitment and limited mission impact. Energy-efficiency improvements in the supercomputing data centers exceeded both an institutional goal and a presidential directive for energy savings and more than doubled the design megawatt capacity of the high-performance computing data centers. Full deployment of a common operating system and tool sets to operate the capacity-computing clusters at the National Nuclear Security Administration's (NNSA's) laboratories (Lawrence Livermore, Los Alamos, and Sandia) allowed new supercomputing capability to be fielded much more quickly and cheaply across the complex. Computation's institutional software quality assurance staff partnered with programs across all LLNL

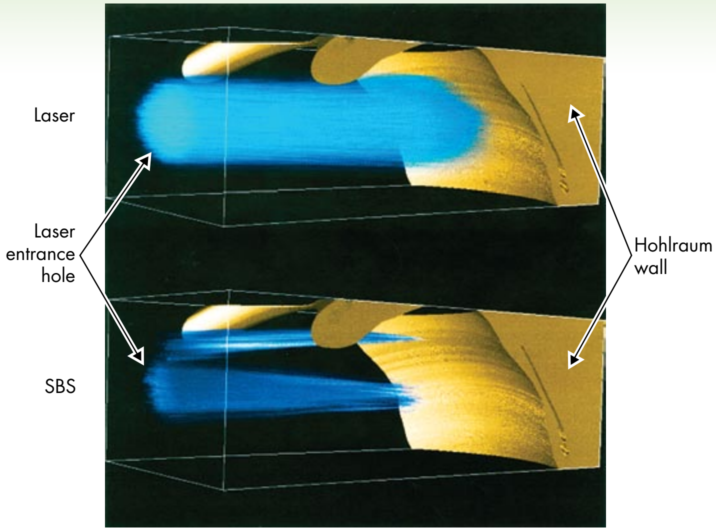

This parallel F3D (pF3D) simulation of a NIF 50-degree beam shows the incident laser stimulating Brillouin scatter (SBS) from two distinct regions. The simulation ran for about 10 days and 13 hours on 196,608 processors of BlueGene/L; this was the first time a simulation modeled the full cross-section of a NIF beam from the laser entrance hole to the hohlraum wall. The level of SBS reflectivity was less than 1% in this simulation.

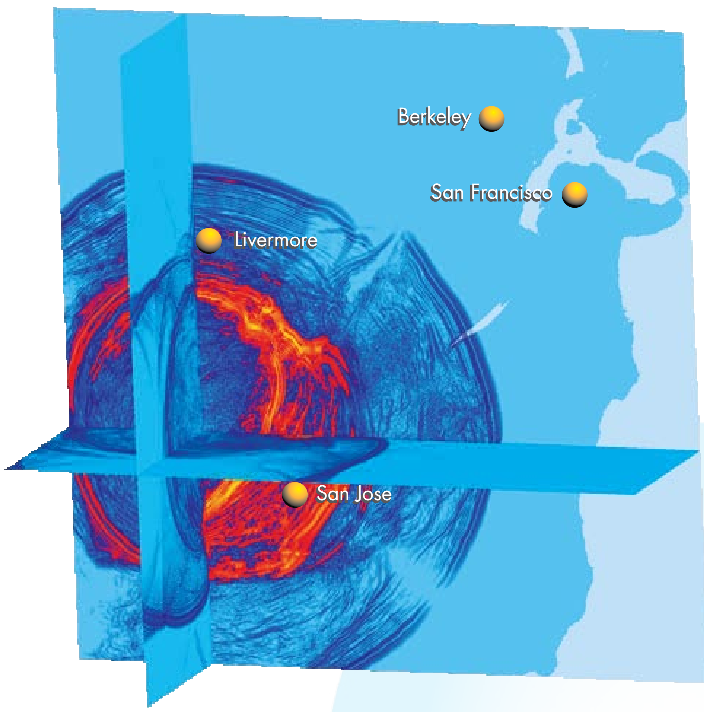

This three-dimensional representation of the October 2007 Alum Rock magnitude-5.5 earthquake was created using the Wave Propagation Program (WPP) code on BlueGene/L to simulate the shear waves. (The viewpoint is below the topography to visualize both the inside of the earth and the surface in the same image.) The colors indicate the strength of the shaking at one instant in time, with orange being the strongest and blue the weakest. The lightest blue indicates deep water. The Alum Rock earthquake was the strongest to hit the Bay Area since the 1989 Loma Prieta earthquake.

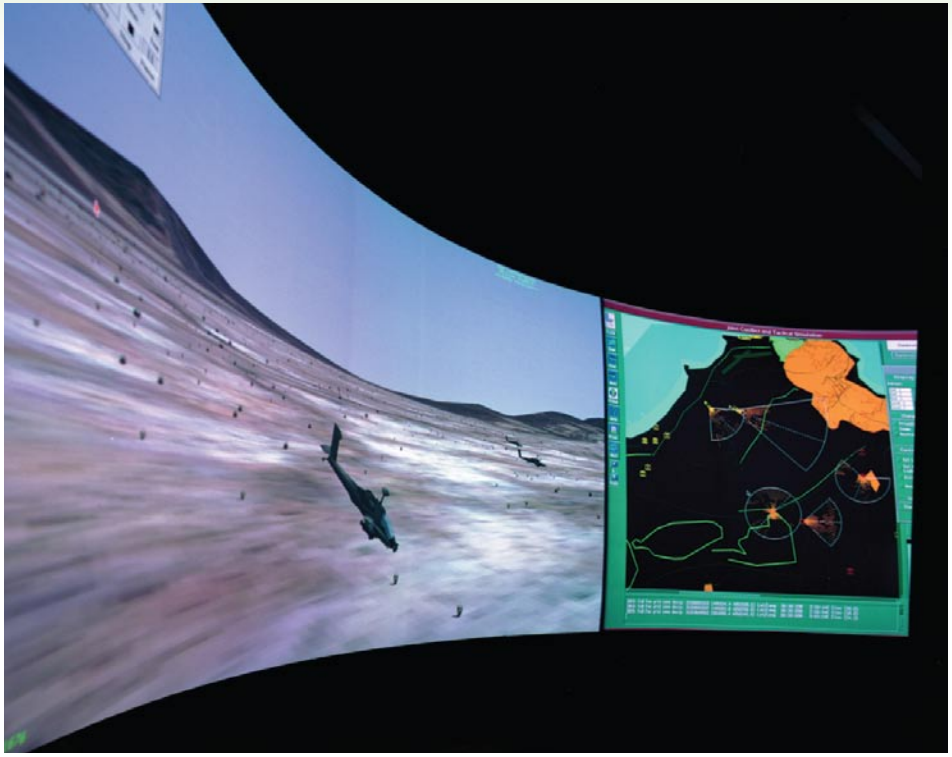

Computation personnel play an essential role in developing the Joint Conflict and Tactical Simulation (JCATS) software suite. JCATS is a multisided, interactive, high-resolution, entity-level, joint conflict simulation used by more than 100 organizations around the world, including all four U.S. military services, to plan, train, and evaluate combat missions. Besides war-fighting scenarios, JCATS has also simulated exercises for drug interdiction, disaster relief, peacekeeping, counterterrorism, hostage rescue, and site security. LLNL released JCATS version 8.0 in 2008, one of largest releases of the software suite to date with more than 90 new features and higher resolution modeling in and around buildings.

organizations to deliver on an institutional goal to implement full software quality assurance practices on all LLNL safety-relevant software.

The Laboratory has embarked on a new integrated science and technology planning effort that will guide institutional investments and workforce strategy during the next five years. Laboratory leaders have identified seven strategic mission thrust areas and several "big audacious goals" where the Laboratory can develop new capabilities that advance scientific discovery and enable future missions. Not surprisingly, Computation fits into all seven strategic thrust areas from cyber and space security to regional climate change to stockpile stewardship to the Laser Inertial Confinement Fusion–Fission Energy project.

I am confident we will continue to face these and other new challenges with creativity and perseverance and remain at the cutting edge of use-inspired research in computation and information technology. Our commitment to realizing the Laboratory's strategic goals of information systems technology and high-performance computing and simulation, combined with the incentives of private-enterprise management, help focus our efforts on efficient and effective mission-driven, use-inspired delivery. President Obama's new administration is advocating for an increase in science and technology research and considers it a key element in addressing the energy, economic, and security challenges facing the nation. When House Speaker Nancy Pelosi was asked about the priorities of a stimulus plan to revive the national economy, she replied, "Science, science, science, and science." I believe this sentiment bodes well for the national laboratories, and our combined focus on effectiveness and efficiency in science will prepare us for the unforeseen challenges of the future.

AN AWARD-WINNING ORGANIZATION

The stories in this annual report present a cross-section of Computation's work in high-performance computing, research, applications and software, and information technology and security. In addition to the projects highlighted in the articles, several external awards were received by Computation personnel and projects in 2008. Some of the honors are featured in this section.

AAPLF R&D 100 Award-winning development team: (back row, from left) Wilbert McClay, James Candy, **Christopher Estes**, **Lawrence Lagin**, **David McGuigan**, Erlan Bliss, Sean Lehman, Abdul Awwal, **Haiyan Zhang**, and **Suzanna Townsend**; (middle row) Paul Van Arsdall, Charles Orth, Thad Salmon, Scott Burkhart, Allan Casey, and **Charles Reynolds**; (front row) Roger Lowe-Webb, **Robert Carey**, **Ben Horowitz**, Mark Bowers, Michael Flegel, Mark Miller, Walter Ferguson, **Eric Stout**, Karl Wilhelmsen, Victoria Miller Kamm, and Richard Leach, Jr. Not pictured: **Stephanie Daveler**, **Holger Jones**, and Karl Pletcher. (Computation employees in bold.)

TOP TECHNOLOGICAL INNOVATION

Each year, *R&D Magazine* recognizes the 100 most technologically significant products of the year. Eleven Computation scientists were part of a multidisciplinary team that won a 2008 R&D 100 Award for the Autonomous Alignment Process for Laser Fusion Systems (AAPLF), which automatically aligns beams for laser fusion systems, such as NIF. Using software controls and signal image processing of sensor and camera data, AAPLF performs 26 separate optical adjustments on each of the 192 NIF beams prior to the initiation of the laser. Each optical adjustment is managed by a control loop that coordinates device movements and image processing tasks while mediating all shared resources. Individual beams are further organized into three parallel segments that are independently aligned. In total, AAPLF is responsible for 3,800 closed-loop optical adjustments involving more than 35,000 devices. AAPLF completes the autonomous alignment process for the entire NIF laser system in 15 minutes and can be scaled to any size system.

NNSA WEAPONS AWARDS OF EXCELLENCE

Fourteen Computation employees were members of teams that received NNSA/Defense Programs' Weapons Awards of Excellence in November. The W80 SS21 Project Team was awarded Technical Excellence in Seamless Safety for analyzing and documenting the various possible hazard scenarios at the Pantex Plant to ensure that nuclear explosive operations can be conducted in a demonstrably safe and reliable manner. The Advanced Simulation and Computing Code Development Team was awarded for making crucial contributions to the W80 Life-Extension Program through code development, problem setup, and integrated physics simulations.

INVENTORS ADVANCE OUR S&T FRONTIERS

Patents represent key technological achievements of scientists and researchers. They also provide economic protection for the invention and generate revenue for the Laboratory. In 2008, one patent was issued to a scientist in the Computation Directorate, and seven applications for patents were filed (still in process).

Patent No. 7339584 (issued March 2008)

Method of Generating a Surface Mesh: Benjamin Grover, et al.

A method and machine-readable medium provides a technique to generate and modify a quadrilateral finite-element surface mesh using dual creation and modification. After generating a dual of a surface (mesh), a predetermined algorithm may be followed to generate and modify a surface mesh of quadrilateral elements. The predetermined algorithm may include the steps of generating two-dimensional cell regions in dual space, determining existing nodes in primal space, generating new nodes in the dual space, and connecting nodes to form the quadrilateral elements (faces) for the generated and modifiable surface mesh.

AND THE OASCR GOES TO…

Computation personnel were part of a NIF-led team that won an Office of Advanced Scientific Computing Research (OASCR) award for visualization work at the Department of Energy's annual Scientific Discovery through Advanced Computing (SciDAC) meeting in Seattle this year. The winning visualization used the electromagnetic code EMSolve (see article 3.02) and Computation's VisIt software (see article 3.05) to show a burst of electrons coming off the target after a short-pulse laser hits the target inside the Titan laser's two-meter-diameter target chamber. The experiment was designed to help determine the best placement and shielding of diagnostic instruments located inside or near laser chambers by changing the geometry of the chamber and analyzing the resulting changes. Attendees of the SciDAC meeting were shown 52 different simulations during a special Visualization Night, but only 10 received OASCRs, which were awarded based on votes by the attendees.

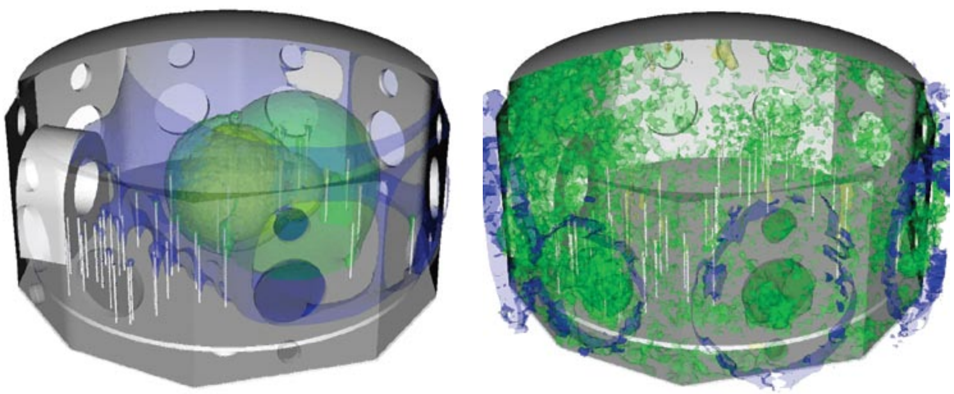

Two frames from the EMSolve simulation of electromagnetic pulses inside the Titan chamber are shown at 1.4 and 8.8 nanoseconds. The laser beams enter the chamber at the far right and then hit the large curved object at the left, which represents a mirror that directs the laser beams onto the target. The upright stalks represent stands for diagnostic instruments and mirrors. The yellow, green, and blue circles show the contours of the resulting electrical fields.

RECOGNITION FROM SCIENTIFIC AND ACADEMIC PUBLICATIONS

Computation scientists and researchers are some of the most prolific in the world (see section 5.03), and their work is published and recognized in various ways.

Computation-led research on data-intensive computing landed the cover article of April's Institute of Electrical and Electronics Engineers *Computer* magazine. This research tackles the data-intensive problems that challenge conventional computing architectures with demanding processor, memory, and input and output requirements. LLNL experiments with three benchmarks suggest that emerging hardware technologies can significantly boost performance of a wide range of applications by increasing compute cycles and bandwidth and reducing latency.

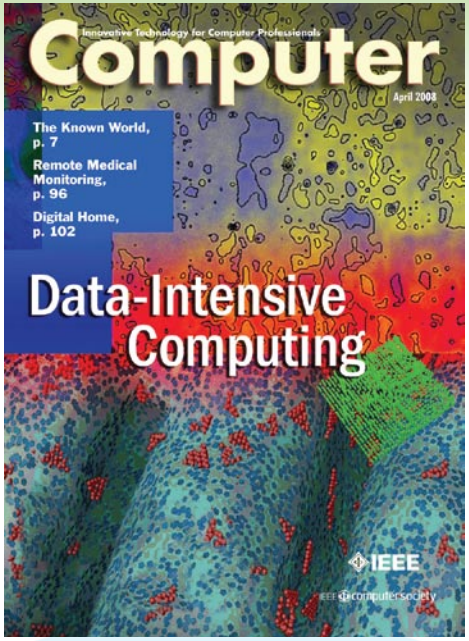

April 2008 issue of *Computer* magazine.

Two Computation-authored papers appeared among the "most accessed" in reputable online scientific journals. Jonathan Allen's paper, "DNA Signatures for Detecting Genetic Engineering in Bacteria," was one of the most accessed files in *Genome Biology*'s Web site, cracking the top 10 and rising as high as number 2. A paper by Shea Gardner and former LLNL summer student Gordon Lemmon (currently in a Ph.D. program at Vanderbilt University) reached "highly accessed" status in the *Annals of Clinical Microbiology and Antimicrobials*. The title of their paper is "Predicting the Sensitivity and Specificity of Published Real-Time PCR Assays."

The BlueGene/L supercomputer still receives significant interest from the scientific and academic community, despite being dethroned from its first place reign on the TOP500 list in June. The 360-teraFLOPS machine was referenced as an example of state-of-the-art supercomputing in the newly published college textbook, *Understanding Computers: Today and Tomorrow* (2008: Cengage Learning).

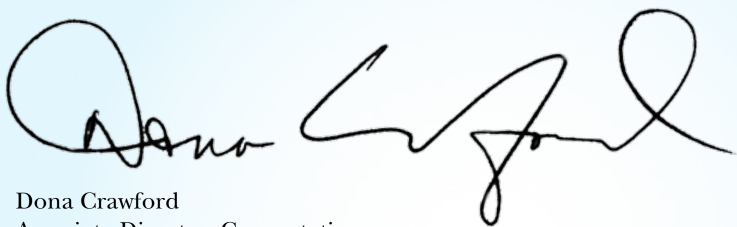

Dona Crawford
Associate Director, Computation

HIGH-PERFORMANCE COMPUTING

1

1.00 The Three-Curve Strategy Pays Off

Since the beginning of this decade, the Three-Curve Strategy has been formalized within the Advanced Simulation and Computing (ASC) Program as a way of thinking about cost-effective approaches to prepare for and meet the spectrum of program computing requirements. The figure on the next page highlights several Lawrence Livermore National Laboratory (LLNL) systems (and a few from Los Alamos and Sandia national laboratories) superimposed on the original viewgraph we prepared in 2001 to describe the approach. This year has enjoyed accomplishments on all three branches, demonstrating the maturation of the strategy and its remarkable impact on the ASC Program.

The ASC Purple computer system has maintained an exceptionally high level of utilization as the National Nuclear Security Administration (NNSA) user facility for capability computing. Utilization rates on the machine consistently exceed 90%; some weeks are as high as 99%, which speaks to the reliability of the resource and the attractiveness of its scheduling and service models. Users across the nuclear weapons complex have been availing themselves of dedicated application time slots to run at the 8,192 core count. For example, personnel from Livermore's Weapons and Complex Integration Principal Directorate ran multiple weekends on Purple to meet an upcoming Level-1 NNSA milestone on energy balance—one of the four areas of physics that must be understood better and modeled more accurately in the quest for predictive simulation. Purple has set a standard for reliability and customer service, demonstrating that computing at a distance is cost-effective and advantageous for the ASC Program.

The second branch, capacity computing, also made notable headway. LLNL coordinated the Tri-Laboratory Linux Capacity Cluster (TLCC)

NEXT PAGE

procurement (see article 1.03), and the LLNL Clustered High Availability Operating System software stack was selected as the core of the Tripod Operating System Software. These systems, comprising more than 500 trillion floating-point operations per second (teraFLOPS) across the three laboratories, are in production and have already demonstrated both the cost effectiveness of common procurements and the coming of age of Linux clusters as reliable production systems. As an example of the impact of these computers, the National Ignition Campaign (NIC) used approximately five million cluster-class processing hours in 2008 and is requesting 10 million in 2009 to support ignition in 2010. The 2009 request is roughly equivalent to running continuously on a 250-teraFLOPS cluster for the entire year. Our ability to support NIC—an addition to our core mission—at this level was due primarily to the extraordinarily low cost of the cluster procurement that provided enough capacity to meet the substantial needs. The preferred machines to run the laser–plasma interaction calculations are the clusters; however, if this preference cannot be met in full, it will require a strategic split between the clusters and the BlueGene systems.

This brings us to the third curve, advanced architectures (AA). This branch looks toward the future; it is here that the Three-Curve Strategy has most dramatically demonstrated its potency. The notion of AA is to explore technologies at scale with the intent that matured technology can be used as a basis to design future production resources. BlueGene will be the first example of a complete cycle, from exploration to production systems. BlueGene/L, the first system in the series, was sited at LLNL in 2005, having benefited from a $30-million multiyear research and development ASC contract. This machine dispelled doubts in the high-performance computing community, went rapidly into production, and contributed to the resolution of the plutonium-aging controversy. We have since invested, with the Department of Energy (DOE) Office of Science, in research and development for BlueGene/P and BlueGene/Q systems. As article 1.05 details, LLNL will take delivery of a 0.5-petaFLOPS (or quadrillion floating-point operations per second) BlueGene/P system and a 20-petaFLOPS BlueGene/Q system between January 2009 and 2011. BlueGene technology, with its multicore chip, presents significant challenges to the integrated weapons codes, which must be adapted to utilize threading

MCCOY

CONTACT INFORMATION

MICHEL MCCOY,
MCCOY2@LLNL.GOV

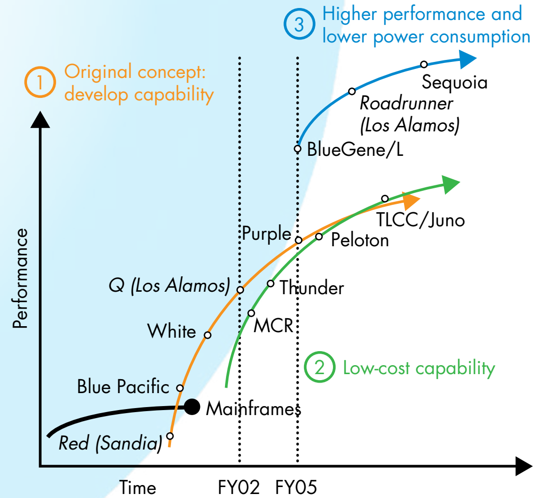

effectively. BlueGene/P was procured to give the ASC Program a platform on which to evolve our target programming model. The mission of BlueGene/Q will be uncertainty quantification (UQ) using integrated weapons codes and weapons science using single-package science codes. The UQ mission will require scaling the weapons codes to 8,192 tasks (while possibly using threads to employ more cores), which is well within the realm of possibility. The science mission will require running the science codes across very large blocks of the machine. These challenges provide for a reasonable evolution without requiring revolution.

It is useful to delineate the complementary missions of the capability and AA systems. The capability machines are necessary to run highly resolved examples of calculations done in great numbers in the UQ suites to demonstrate convergence (or calibrate the error in running at lower resolution). The capability machines are also used for discovery science using the integrated weapons codes. These two computer types are complementary—both are necessary to advance toward predictive simulation.

1.01

Hyperion Project Begins Collaboration to Develop an Advanced Technology Cluster Test Bed

SIGNIFICANCE

One of the challenges associated with developing, managing, and utilizing high-performance computing (HPC) systems is a lengthy and costly testing phase that must be completed to ensure the new hardware, middleware, and application software works well at scale. LLNL has teamed with 10 industry partners—Cisco Systems, DataDirect Networks, Dell, Intel, LSI, Mellanox, QLogic, Red Hat, Sun Microsystems, and Supermicro—to develop an innovative solution to this problem. Hyperion is a large-scale Linux cluster test environment, which will speed the development and reduce the cost of powerful HPC clusters by making a petascale cluster available solely for development and testing infrastructural components. It will also be used to evaluate system hardware and software critical to maintaining the aging U.S. nuclear weapons stockpile. Hyperion, too large for any one organization to deploy, represents an innovative approach for industry-leading collaborations and exemplifies LLNL's leadership in HPC.

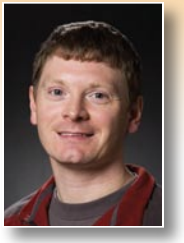

LEININGER

CONTACT INFORMATION

MATT LEININGER,
MATT@LLNL.GOV

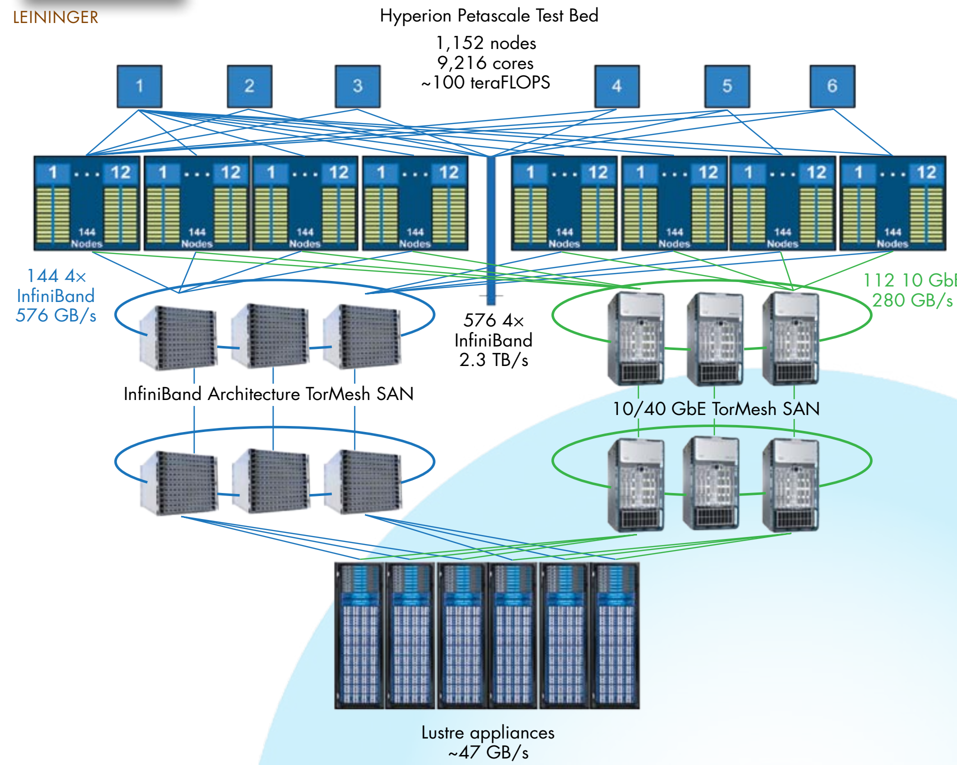

The Hyperion cluster architecture includes eight scalable units (1,152 nodes), the InfiniBand and Ethernet TorMesh SAN, and high-performance storage. Hyperion connects to one SAN via 144 4x-InfiniBand network links (blue lines) and to the other SAN by 112 10-GbE links (green lines). There are also 576 4x-InfiniBand links, totaling 2.3 terabytes (TB) of bandwidth, that are not currently in use.

PROGRESS IN 2008

Michael Dell's announcement of the Hyperion Project during his 2008 Supercomputing Conference (SC08) keynote address was the culmination of more than a year and a half of work between LLNL and industry collaborators to field an advanced technology cluster as a testing and scaling resource. "Storage, connectivity, management software: these are all challenges that we are going to be dealing with as we implement [hyperscale] systems," Dell said. "Hyperion is a test bed big enough to really test scale, and it will share those breakthroughs with the entire open-source community."

Hyperion is a dedicated development and test environment for advanced Linux cluster technologies, an evaluation test bed for new hardware and software technologies, and an input and output (I/O) test bed to prepare for LLNL's future ASC Sequoia platform. The first phase of Hyperion was installed at LLNL in September 2008 and consists of 576 nodes comprised of two-socket quad-core Intel Harpertown processors, eight gigabytes (GB) of memory per node, and a double data rate (DDR) InfiniBand interconnect. The first-phase cluster is testing the Lustre parallel file system and the Red Hat-based ASC Tripod Operating System Software (TOSS). It is also testing Message Passing Interface (MPI), OpenMP, and the OpenFabrics high-performance networking software for InfiniBand and low-latency Ethernet. DataDirect Networks, LSI, and Sun Microsystems each provided high-performance storage for Hyperion, which is connected through two separate storage area networks (SANs), one based on DDR InfiniBand and another based on low-latency 10-gigabit Ethernet (GbE).

The second phase of Hyperion, scheduled for March 2009, will double the system size to 1,152 nodes, with the additional nodes using the new Intel Nehalem processors (two-socket quad-core), 12 GB of memory per node, and a DDR InfiniBand interconnect. All Hyperion nodes have Peripheral Component Interconnect Express 2.0, which enables future upgrades to quad data rate (QDR) InfiniBand.

Operational Efficiencies Lead to Increased Return on HPC Investments

SIGNIFICANCE

Operating a world-class computer center 24 hours a day, 365 days a year, requires a great deal more than fielding best-in-class computers and leading-edge software. The Livermore Computing (LC) Center continuously strives to improve the user experience and reduce the cost of delivering exceptional service. This year, as in years past, LC developed and implemented several efficiency-increasing processes, including a streamlined parts-return process, a maintenance cost-avoidance plan, platform-specific hardware and software metrics reporting, and an electronic account request process. Continuous operational improvement is a cultural norm at LC that we nurture and protect to maintain our ability to provide a computational capability worthy of addressing vexing national security issues.

CUPPS

CONTACT INFORMATION

KIM CUPPS,
CUPPS2@LLNL.GOV

PROGRESS IN 2008

Several years ago, LC achieved one of the largest cost reductions in its history by training a few staff to "self-repair" LC Linux clusters rather than contracting with an outside vendor to provide support. This change was instituted with the purchase of commodity Linux clusters in 2003 to avoid the considerable up-front costs of on-site repair in the purchase contract. What began as a hardware repair team of two people for a single 1,000-node cluster has become a five-person team that provides Tier-1 diagnosis, repair, and replacement on more than 115,000 nodes, 44,000 disks, and 344 disk controllers in the center. This year, we returned more than 3,240 parts to various vendors for replacement.

The return maintenance agreement process that LLNL traditionally uses requires matching a purchase order number to a particular part's serial number before the part can be shipped back to the supplier. We followed this procedure for many years, but as our installed base of nodes grew and the number of failures increased, the burden of the return process became unreasonably high. LC worked with LLNL's procurement department, shipping team, and vendors to negotiate a new, streamlined hard-drive return policy that eliminates the purchase-order association and allows us to group hundreds of returns into a single shipment. Since hard disks constitute the majority of our parts failures, the return process can now be accomplished in one-tenth of the time.

LC also reduced costs in 2008 by eliminating several maintenance agreements after the original term expired. We generally buy three years of hardware maintenance from a vendor in the initial procurement with the understanding that there will be a small on-site spare-parts cache, and failed parts will be returned for replacement/diagnosis when security rules allow. Once the maintenance contract expires, we typically harvest parts from a smaller, identical system to replenish failed parts on the larger system until it reaches its end of life. As an example, this year LC did not renew the maintenance agreement for the 1,532 nodes of the Purple computer, which saved $4 million in annual maintenance costs. Instead, LC designated a smaller system as a "hot spare cluster" and replaced

Computer support technologist Orlando Grant replaces a memory board on a node of Hera, a large-capacity Linux resource.

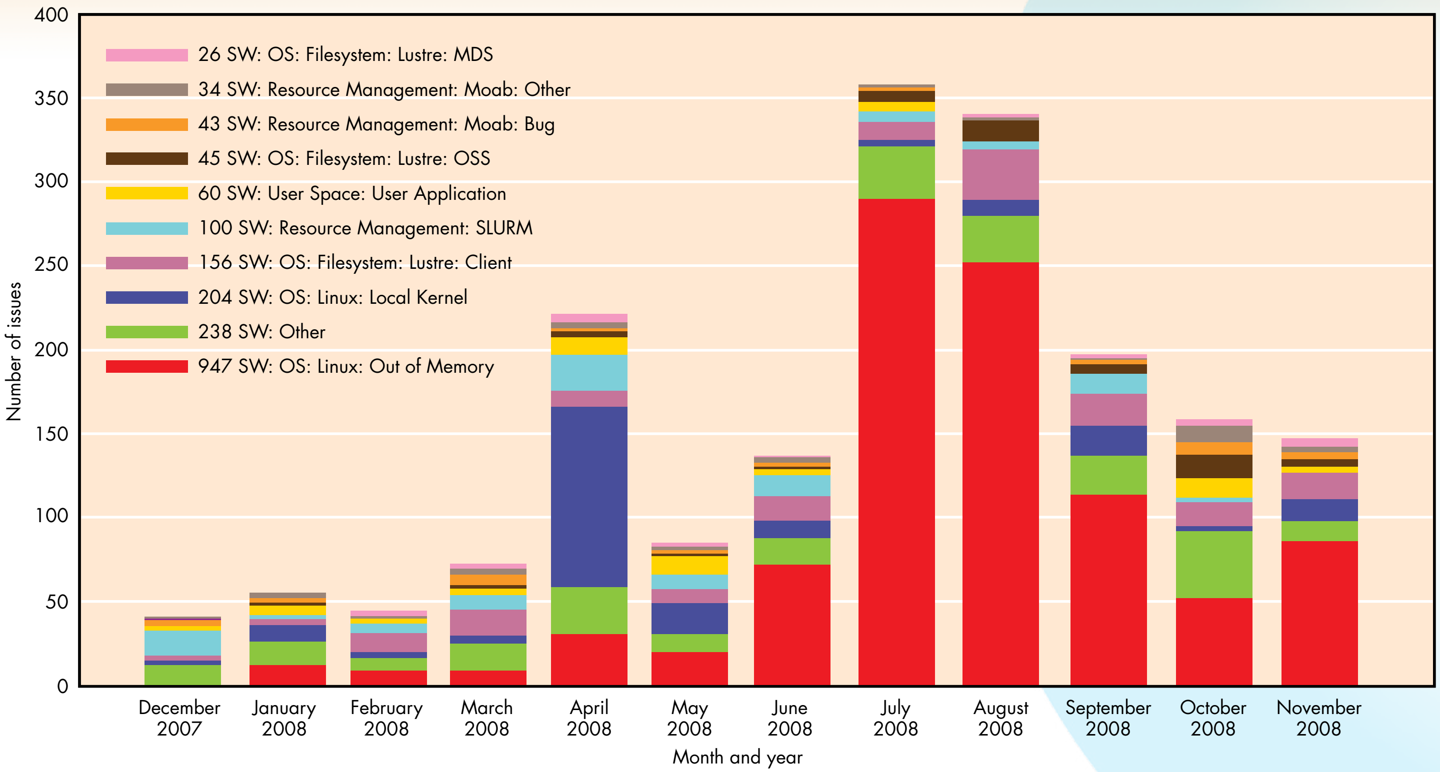

Top software (SW) issues as measured throughout the year. Red Hat Enterprise Linux version 5 was released to production machines in June 2008 and caused an increase in "out-of-memory" issues.

26 nodes of Purple with the smaller system's parts. As Purple's size is reduced in its last 18 months of service, at this failure rate no additional node purchases from outside vendors will be required.

Because of the size of modern clusters, we can no longer rely on one or two people to have a firm grasp on all equipment failure nodes. In 2008, LC began generating monthly reports on the Top 10 hardware and software failures for each platform type. The first report highlighted a large number of node failures caused by an "out-of-memory" condition on the Linux clusters. Red Hat's proactive out-of-memory cleanup process, OOM Killer, was not performing well. Although the issue was known, the number of failures resulting from the rollout of the Red Hat Enterprise Linux version 5 operating system was unexpected. Once the magnitude of the issue was understood, we diverted two kernel developers to identify a fix. Within six weeks, we deployed a temporary solution that decreased the number of node failures by 50%, while we waited for a permanent solution in Red Hat version 6.

In another case, the metrics report identified an ongoing systemic power supply issue in sufficient enough detail that the vendor agreed to diagnose and fix all installed power supplies and extend maintenance coverage from three years to five years at no additional cost.

The final cost efficiency implemented in 2008 was a commercial account management system. This system provides an electronic account provisioning and approval workflow process. LC has 3,800 user accounts on 21 production clusters, and we handle approximately 20,000 account modify/create/delete requests per year via a paper-based system. These requests generally take three to five days to complete from the time a user makes the request until the account is provisioned. During beta testing of the new electronic workflow system, an account was created from beginning to end in an hour. The new system is not yet in full production, but it promises a significant increase in productivity for our users and staff, which is one of our top priorities.

Common Capacity Computing Reduces Cluster Deployment Time and Cost

SIGNIFICANCE

The purpose of the Tri-Laboratory Linux Capacity Cluster (TLCC) procurement is to build a common capacity hardware environment at the three NNSA laboratories: Livermore, Los Alamos, and Sandia. In the past, each laboratory could field only one or two large clusters per year using locally maintained operating systems. This year, 13 clusters were installed at the three laboratories in only nine months. By deploying a common hardware environment multiple times at all three sites, the laboratories reduced the time and cost to deploy each cluster by approximately an order of magnitude.

The common hardware environment also facilitated the adoption of a common software environment, called the Tripod Operating System Software (TOSS). TOSS is based on the Red Hat Linux operating system and includes Livermore-developed additions that make it a clustered high-availability system. TOSS runs across all newly procured TLCC machines. After a year of successful cluster installations, TLCC and TOSS have proven to reduce the total cost of ownership associated with system administration and software development and enable application portability within the nuclear weapons complex.

HAMILTON

CONTACT INFORMATION

PAM HAMILTON,
PGH@LLNL.GOV

PROGRESS IN 2008

The delivery of the cluster hardware was significantly delayed because of two issues related to the transition from dual- to quad-core processors. Once the problems were resolved, the systems were delivered quickly (see figure), and they are now operational at the three laboratories.

LLNL's part of the TLCC procurement included four clusters—Hype, Eos, Juno, and Hera—that provide more than 280 teraFLOPS of peak computing power to the laboratories' programs. In March, Hype became the first TLCC cluster deployed at LLNL. At 144 nodes, Hype is the main test resource for the other clusters and is being used for file system, operating system, and resource management software testing.

Eos, comprising 288 nodes, was installed and ready for use in September. Eos will primarily run simulations in support of the National Ignition Campaign.

Juno, the largest TLCC system at 1,152 nodes, was made available to the ASC Program in October. Early success on Juno included a National Ignition Facility (NIF) laser–plasma simulation that ran continuously for four days on 1,024 nodes using 16 tasks per node. This is comparable to running concurrently on 16,384 desktops for four consecutive days.

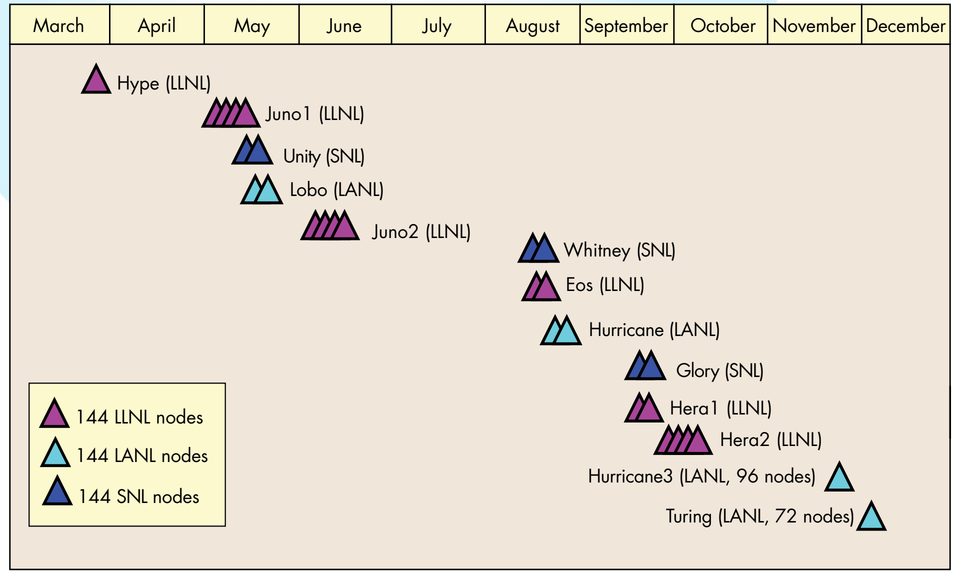

TLCC deliveries occurred at Livermore, Los Alamos (LANL), and Sandia (SNL) from March to December 2008. Each triangle represents a 144-processor scalable unit, the building block of TLCC hardware.

The final TLCC cluster deployed at LLNL was Hera, an 864-node cluster. Hera is shared by the Multiprogrammatic and Institutional Computing and ASC programs.

TOSS delivers a fully functional cluster operating system (kernel, Linux distribution, InfiniBand stack and related libraries, and resource manager) capable of running MPI jobs at scale on TLCC hardware. The TOSS environment is a complete product with full life-cycle support. Its well-defined processes for release management, packaging, quality assurance testing, configuration management, and bug tracking ensure that production-quality software is deployed at the three laboratories in a consistent and manageable way.

TOSS version 1 was released in early March. Soon after, the three laboratories used TOSS to pass the synthetic workload acceptance tests (a suite of benchmarks and application codes representative of ASC workloads) for their respective TLCC systems. A revision of TOSS (1.0.6) was released in August and contained numerous patches and a workaround solution for a particularly impactful out-of-memory issue.

Open|SpeedShop Provides Insight into the Performance of Parallel Applications

SIGNIFICANCE

Performance analysis is a critical step in the development cycle of a high-performance application. To help with this task, users and developers require sophisticated, yet easy-to-use tools that provide the necessary insight into an application's performance. The Open|SpeedShop (O|SS) Project, initiated in 2004 by DOE/NNSA as part of the ASC PathForward Program, provides such a tool set. O|SS is an open-source performance analysis tool set that targets sequential codes, large-scale parallel applications based on MPI, and multithreaded codes. It relies on the concept of dynamic instrumentation and is built on top of open-source software. Although O|SS was initially intended for use on DOE's capacity computing systems, we are taking steps to adapt the tool for capability systems as well.

The O|SS Project is jointly sustained by the Krell Institute and the NNSA Tri-Laboratory community. It supports a wide range of performance measurements within a single framework and covers the most typical performance analysis needs, including statistical run-time sampling, hardware counter analysis, MPI, and I/O profiling and tracing.

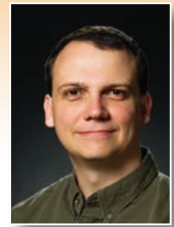

SCHULZ

CONTACT INFORMATION

MARTIN SCHULZ,
SCHULZM@LLNL.GOV
WEB SITE: WWW.OPENSPEEDSHOP.ORG

PROGRESS IN 2008

2008 marked a set of important milestones for the O|SS Project. We transitioned the tool set from a research project to a production capability installed on TLCC machines at the three NNSA laboratories and several partner sites. In addition, by leveraging portable and scalable underlying libraries, we opened the door for the tool to be used on capability computing systems, such as BlueGene/L, BlueGene/P, XT5, and Roadrunner, and eventually on ultrascale machines such as LLNL's future petaFLOPS system, Sequoia.

Figure (a) shows the typical workflow for an O|SS user. After selecting an experiment (i.e., specifying which performance data will be collected), the user controls the application from within O|SS. The tool then collects the data, stores it with the necessary metadata in a structured query language (SQL) database, and makes results available using multiple analytical views. O|SS offers the user multiple interfaces in which to view the results, including a comprehensive graphical user interface with performance analysis wizards, a command-line interface to enable batch processing, and a Python interface for scripting support.

Among the technical changes to O|SS in 2008 was the replacement of the collection infrastructure. O|SS was originally built on top of the aging, unscalable, and hard-to-maintain Dynamic Probe Class Library, but it now offers two new modes of operation (figure b): a flexible and easy-to-port off-line data collection mechanism, and a scalable online collection overlay network based on a reduction tree combined with dynamic instrumentation. The latter is built using the MRNet infrastructure, which also forms the basis for other LLNL/ASC tools such as the Stack Trace Analysis Tool (STAT). The user chooses between the off-line mode, which provides a fast and easy-to-deploy mechanism for data collection on entire runs, and the online mode, which allows the user to monitor the data acquisition at run time and to attach to or detach from running jobs.

In addition to the run-time changes, O|SS now features updated performance wizards, more options for performance comparisons, and single-command start options. Furthermore, O|SS distributions contain a comprehensive setup script for convenient installation, which ensures that the required underlying libraries are installed and functional. The new script also configures and compiles O|SS with the required options and creates the appropriate run-time setup for running the tool set.

During SC08, we released version 1.9 of the O|SS tool set and made it available from the project Web page, along with extensive documentation and a CD for easy evaluation. We engaged the end-user community through a series of presentations and a keynote address, tutorials, and workshops. We also visited potential user sites (Boeing and the National Aeronautics and Space Administration) and gave well-attended seminars and hands-on workshops at the three NNSA laboratories. The workshops gave code teams the chance not only to learn about O|SS but also to apply the tool set and its capabilities directly to their codes.

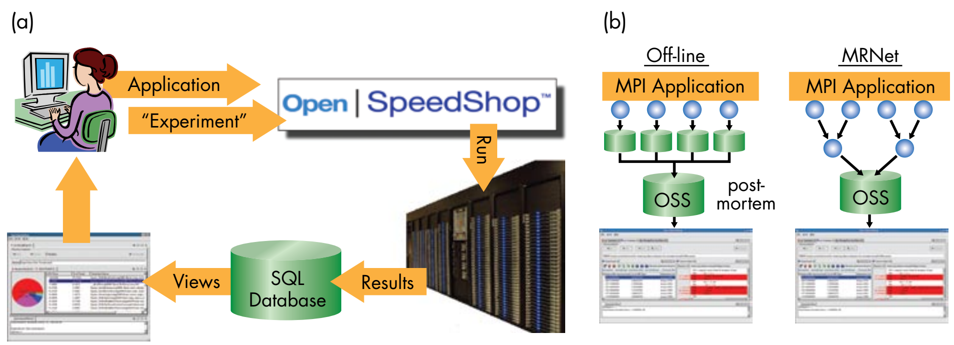

(a) In a typical O|SS workflow, a user selects an experiment, starts, and controls the application. O|SS collects performance data, stores it in a database, and presents the results back to the user. (b) Enhanced data collection mechanisms in O|SS include robust and portable off-line data collection (left) and highly scalable online data collection using hierarchical overlay networks and dynamic instrumentation (right).

Sequoia Grows from Conception to Contract

SIGNIFICANCE

In the summer of 2006, the 100-teraFLOPS ASC Purple supercomputer delivered the first three-dimensional, full-system, full-physics nuclear weapons simulation, and the 360-teraFLOPS BlueGene/L supercomputer claimed the number one spot on the TOP500 list for the fourth consecutive time. During this prolific period, the ASC Program realized that to further advance the nation's Stockpile Stewardship mission would require the most audacious leap in computing power ever attempted: 12 to 24 times the level of performance that can be achieved on Purple for uncertainty quantification suites of integrated coupled physics calculations, and 20 to 50 times the level of performance that can be achieved on BlueGene/L with materials science codes. Thus began the nearly three-year odyssey from conception to contract for the Sequoia computer system.

To achieve this computing power, Sequoia would need to be a breakthrough system with 20 petaFLOPS peak performance speed, 1.68 petabytes (PB) of memory, and, because of the multicore revolution, unprecedented parallelism of six million threads exploiting approximately 1.5 million cores. Sequoia was not just a platform challenge; it was also an applications development challenge. In addition, a machine of this size would require significant enhancements to the other simulation environment components, such as the Global Parallel File System, SAN, and archive resources.

SEAGER

CONTACT INFORMATION

MARK SEAGER,
SEAGER@LLNL.GOV

PROGRESS IN 2008

To solve the daunting challenges associated with creating a breakthrough system, the Sequoia procurement team developed a vendor-neutral Sequoia target architecture, which included everything from the programming model to scalable system hardware and software and the SAN and Lustre parallel file systems. The Sequoia target architecture builds on the idea of the "Livermore Model," originally developed in the early 1990s to extend both the programming paradigm and the systems paradigm to scale through petascale system generations.

The original Livermore Model assumed a full Unix operating system (OS) and local disk/file system on every node, with a few nodes designated for login and system functions. The systems model was flat, meaning every node was configured the same except for the login nodes, which had external networking. The programming model leveraged the workstation programming environment by providing OpenMP/Pthreads parallelism on node, and it extended the workstation by providing MPI parallelization at the highest level for communications between nodes over a high-speed, low-latency network.

ACTIVITY	START
Market Survey	Summer 2006
RFP Development	June 2007
Draft RFP Release	May 2008
Formal RFP Release	July 16, 2008
Response Due Date	August 21, 2008
Best Value Evaluation	August 26–28, 2008
Vendor Selection	September 6, 2008
Contract Negotiations	September 9, 2008
Contract Approval	October 28, 2008
Contract Award	January 9, 2009
Delivery of Dawn	January 22, 2009
Delivery of Sequoia	2nd Quarter, 2011

The Sequoia RFP development and DOE approval process took more than a year and a half to complete. Going from RFP to a signed contract took an additional year.

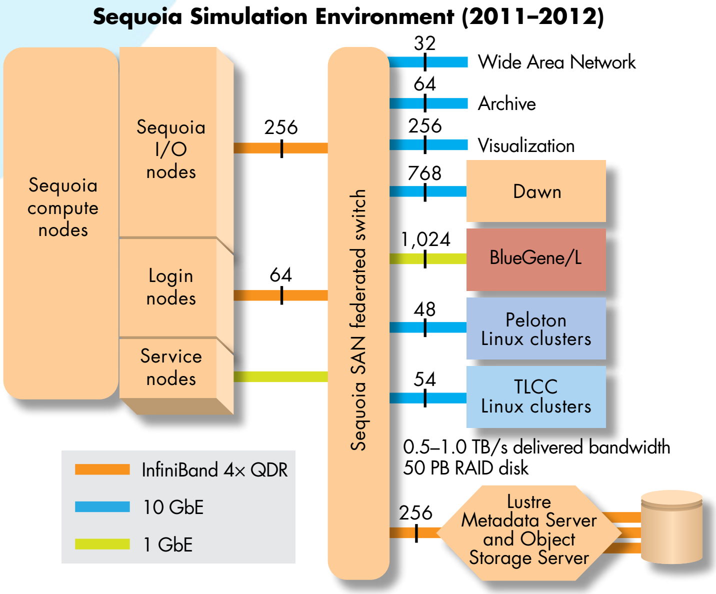

The Sequoia target systems architecture is a hierarchical systems model with a very large number of compute nodes that access the Lustre parallel file system via a smaller number of I/O nodes and SAN. Other elements of the simulation environment include capacity Linux clusters, BlueGene/L, Dawn (BlueGene/P), and visualization and archive resources.

The Sequoia target architecture extends the Livermore Model concept by specifying four kinds of nodes: compute nodes (diskless with a very lightweight OS), I/O nodes (diskless with a fully functional Linux OS providing remote OS functions to multiple compute nodes and directly attached to the SAN), login nodes, and service nodes (both with local disk and fully functional Linux OS). The hardware and software systems model is hierarchical to simplify the hardware and software of the large number of compute nodes. The model requires that compute nodes be associated with a smaller number of I/O nodes and that I/O nodes be associated with the even fewer service nodes and login nodes for system services.

The Sequoia procurement team engaged the Tri-Laboratory community, ASC Program Office, and industry to develop the procurement strategy and detailed request for proposal (RFP) requirements for the Sequoia target architecture as well as programmatic requirements such as representative benchmarks and workload characterization. The team's discussions with industry provided valuable feedback on how the Laboratory's requirements would affect companies' target architecture schedule and pricing. Interactions with the Tri-Laboratory community helped the team develop the benchmarking strategy and representative benchmarks, and refine requirements and usage models. The benchmarks had to realistically represent and measure Sequoia's program goals. They also had to represent the intended workload, scale to the size of the Sequoia system, and not be export controlled.

In addition, the team constructed a strategy for migrating the applications running on Purple and BlueGene/L to Sequoia, while still delivering results to the ASC Program. This migration will require more than just porting and scaling the codes. The purpose of Sequoia is to increase the scientific simulation predictivity of these codes, so new physics models must be developed as well as the mathematical representation of those models, algorithms for solving the new (and existing) mathematical representations at much higher scalability, implementations of those algorithms, and verification and validation of the whole process. Thus, a key part of the Sequoia strategy is an early delivery system, called Dawn, which will allow the ASC Program to start the development process two years before Sequoia is deployed.

In 2008, DOE approved the final draft of the Sequoia RFP. Soon after, several vendors responded with competitive proposals. A 27-member Tri-Laboratory group evaluated the technical excellence of the proposals, feasibility, cost (including total cost of ownership factors such as power consumption, schedule, development, and risk mitigation plan), and qualifications of the offeror. In the end, IBM was selected to build Sequoia. Contract negotiations were completed within 45 days, and the final contract approval occurred 122 days later.

1.06

HPC Storage Infrastructure Expands in Preparation for Dawn

SIGNIFICANCE

The success of every mission supported by LC depends on the data generated by applications that run on LC platforms. The accuracy, availability, and speed of data access are essential, as these data represent the entirety of the programs' investment in platforms, applications, facilities, networks, and people. Even a small failure in the storage infrastructure can idle multimillion-dollar compute platforms within seconds.

LC's data storage infrastructure consists of large, global parallel file systems (Lustre) and world-class digital archives (High-Performance Storage System, or HPSS). Lustre provides high-speed global access to disk-based data from LC's compute resources, and HPSS provides high-speed access and long-term retention of data products representing more than four decades of scientific investigation and simulation. Both Lustre and HPSS exemplify successful DOE/NNSA, industry, and open-source community software collaborations. In 2008, the Lustre and HPSS infrastructures were expanded and enhanced to support the TLCC and Dawn platforms.

GARY

SHOOPMAN

CONTACT INFORMATION

MARK GARY, MGARY@LLNL.GOV
JERRY SHOOPMAN, SHOOPMAN2@LLNL.GOV

PROGRESS IN 2008

Storage bandwidth and capacity requirements are based on a number of factors, the most important of which is a platform's memory capacity. LC storage deployments in 2008 were dictated by the immediate requirements of TLCC and anticipated requirements of a 2009 platform, called Dawn. Dawn's memory capacity will be 4.5 times greater than its predecessor's. Because platform advances continue to exceed those of I/O devices, Lustre and HPSS architectures are built to deliver the bandwidth of multiple storage devices in parallel.

Robotics expert Geoff Cleary calibrates a new robotic handbot in one of four expanded SL8500 tape libraries with 10,000 tape cartridge slots.

Lustre system administrator Jim Harm works on one of the new SAS-based Lustre SSUs.

The expansion of the Lustre parallel file system involved selecting a new storage scalable unit (SSU), which is a building block consisting of servers, redundant array of independent disks controllers, and disk drives that can be easily replicated, deployed, and managed. In a departure from past deployments, LC selected lower-capacity, more-reliable serial attach SCSI (SAS) disk drives rather than higher-capacity, less-reliable serial ATA drives. This choice was made under the principle that, at the SSU level, requests are truly random; therefore, performance and reliability should be the most important factors in hardware selection. To directly leverage the TLCC procurement, technology, and maintenance investments, LC chose SSU servers, network adapters, and InfiniBand interconnects based on TLCC hardware. LC also made two major upgrades to the Lustre software, installed an additional 5.7 PB of storage, and increased off-platform storage bandwidth by 104 GB per second.

The enhancement of HPSS included increasing the capacity of the SL8500 tape library by 20,000 slots to a total of 130,000. The tape libraries now have the potential to hold up to 130 PB of data. Dual robotic handbots, which provide redundancy and increased performance, were deployed in all SL8500s, and more than 200 high-bandwidth/capacity tape drives were put into service behind the Clustered High Availability Operating System mover nodes via new optical fiber. More than 30,000 new tape cartridges were also added to the libraries. Because higher file-create rates are expected with the Dawn platform, HPSS core services were migrated to newer platforms, metadata disk capacity was quadrupled, and the underlying fabric connecting core servers to metadata was modernized and made more resilient.

Architecturally, LC transitioned to a new file-transfer model, enabling the majority of data transfers between Lustre and HPSS to take place via the Storage Lustre Interface Cluster (SLIC) machines. Transfer statistics show that more than 95% of data moved by LC's network file-transport user interface is being routed via this approach. Using the SLIC methodology reduces the cost of installing jumbo-frame networks and provides users global access to Lustre and HPSS data even when machines are not functioning.

Impressive hardware expansions are only as good as the software running on them. HPSS is the product of an ongoing, highly successful, 16-year collaborative development effort between five DOE laboratories and IBM Corporation. In 2008, HPSS collaborators finalized and released HPSS R7.1, which features file-performance enhancements, parallelized background media repack, and more intelligent use of disk/tape resources—functionality critical to supporting TLCC and Dawn. HPSS R7.1 will be deployed in 2009.

Lustre gained popularity in HPC centers in large part because of DOE PathForward investments. Today, Lustre is an open-source development project led by Sun Microsystems, with participation from LLNL and other HPC centers.

Stack Trace Analysis Tool Enables Debugging on More than 200,000 Processors

SIGNIFICANCE

Large-scale simulations that utilize hundreds of thousands of processing cores are an increasingly important part of LLNL's programmatic requirements. As a result, developers face significant challenges in resolving run-time errors that emerge only at scale. LLNL users identified "a means of debugging at scale" as their number one requirement for petascale software development tools. The Stack Trace Analysis Tool (STAT) responds to this critical need. STAT is a lightweight tool that samples and merges code execution states from the run-time stack of an MPI application when it executes incorrectly. STAT presents the user with equivalence classes of MPI tasks, a key aspect of the anomalous execution. Our progress in 2008 dramatically improved STAT's scalability and usability; the tool now scales to all 212,992 cores of BlueGene/L.

AHN

LEE

CONTACT INFORMATION

DONG AHN, AHN1@LLNL.GOV

GREGORY LEE, LEE218@LLNL.GOV

PROGRESS IN 2008

This year, we scaled STAT to the world's largest systems and tested a core count two orders of magnitude higher than we tested in our initial experiments. We discovered and solved various scalability bottlenecks that were difficult to address, particularly because little or no performance data exists at the requisite scale. For example, we found that ad hoc sequential launching of remote STAT daemons was too expensive, so we developed LaunchMON, a scalable daemon launcher that improves the scalability and portability of any tool. We then found that bit vectors, a highly compact encoding scheme used to represent lists of MPI ranks, required too much bandwidth at extreme scales to function correctly. Thus, we developed a hierarchical representation of MPI ranks that dramatically reduced data-transfer requirements on MRNet, the tree-based overlay network that STAT employs to reduce the data and processing loads on its front end. We also developed a file-relocation mechanism that scalably distributes and relocates application binaries to remote nodes to avoid swamping the file system when too many daemons simultaneously retrieve static information necessary to decode stack data.

We complemented our scalability breakthroughs with usability enhancements to ready STAT for end users. In collaboration with the University of Wisconsin, we transformed STAT from a proof-of-concept prototype to a production-ready tool. We hardened software layers by defining precise failure modes and increasing code coverage through enriched test cases. Additionally, we augmented the tool with many user-friendly features, including STATview, a graphical user interface that helps users navigate through STAT graphs represented as call graph prefix trees with color-coded process equivalence classes.

We began deploying STAT on HPC machines within LC and used it to diagnose some of the most challenging real-world errors. This production experience is beginning to help unlock the common patterns of "errors at scale" and, perhaps more importantly, further innovate debugging techniques.

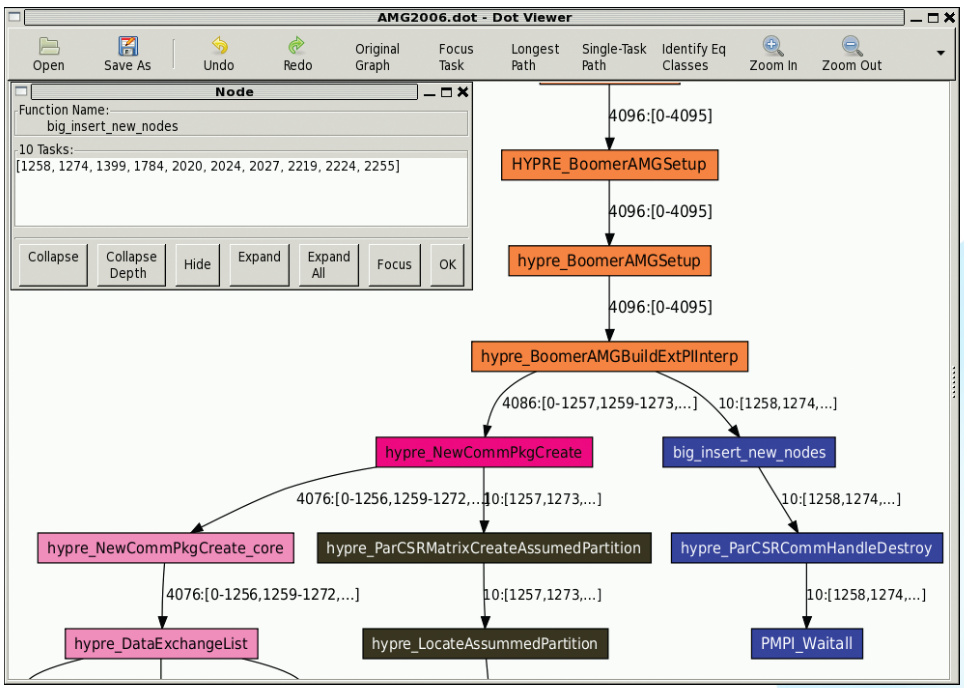

STATview diagnoses a hang in the Algebraic Multigrid (AMG) code at 4,096 tasks. The scalable call graph prefix tree with color-coded equivalence classes helped a developer identify an anomalous class of the 10 tasks in the right-most branch. Subsequent code review in this execution flow revealed a type coercion problem on a function parameter at the big_insert_new_nodes function.

RESEARCH AND ADVANCED DEVELOPMENT

2

2.00 Research Activities Typify Scientists' Wide-Ranging Expertise

Research and development in the Computation Directorate is focused on preparing Lawrence Livermore National Laboratory (LLNL) for a new era in computational science. With the acquisition of the Laboratory's first petascale high-performance computing system in 2011, research is required to develop algorithms and applications that effectively run on supercomputers of this massive scale. The new class of architectures—many-core processors and million-core systems—is driving novel approaches in scalable algorithms, tools, and libraries. If successful, we will be able to pursue simulations orders of magnitude greater in capability and capacity. In addition to exploiting petascale architectures, we continue to develop capabilities in simulation and data analysis to enable discoveries in science and novel applications for national security. Our research in knowledge discovery, machine learning, and data-intensive computing is focusing on computer network defense and text-analysis capabilities to support intelligence analysis.

During the past several years, our experiences winning three Gordon Bell prizes for leading-edge applications of high-performance computing have forged our thinking into how to pursue research for petascale science applications. For example, the 2007 Gordon Bell Prize involved

NEXT PAGE

developing a large-scale simulation of Helmholtz instabilities at the atomic scale on 212,000 cores of BlueGene/L. Based on these successes, the Institute for Scientific Computing Research (ISCR) expanded its scope in 2008 from facilitating interactions with academia to demonstrating leadership in large-scale simulations. ISCR has become the Computation Directorate's focal point for ultra-high-performance scientific computing, serving as a magnet for recruiting talent and a hub for discoveries in computational science that will drive future research in the Directorate. In addition to expanding ISCR, we are conducting new research in algorithm scalability, focusing on serial bottlenecks involved with modeling radiation and neutron transport, which is critical to many important LLNL stockpile stewardship and inertial confinement fusion codes. This research requires linear solvers that exploit the mathematical structure of solved physics equations.

During the past decade, the Computation Directorate developed several libraries and frameworks that we use to rapidly create breakthrough modeling and simulation capabilities. In our current research on phase-field modeling, we seek to develop computer models for understanding melting and solidification processes in materials under extreme conditions. This work requires new modeling methods that fill a length- and time-scale gap between molecular dynamics and continuum dynamics. At the beginning of the project, the approach required for algorithms and problem formulation was unclear. To quickly explore a range of algorithmic possibilities, we employed a rapid software development approach for this research, using LLNL libraries such as *hypre* for linear solvers, Structured Adaptive Mesh Refinement Application Infrastructure (SAMRAI) for adaptive mesh refinement, and Suite of Nonlinear and Differential/Algebraic Equation Solvers (SUNDIALS) to solve differential-algebraic equations.

In another project, we built new mission applications using the Overture framework. Until recently, Overture had been used primarily as a tool for exploring and prototyping new applications of composite gridding (also known as overset gridding) for applications in computational fluid dynamics and electromagnetic modeling. We also developed applications to model laser propagation problems


GROSH

CONTACT INFORMATION

JOHN GROSH,
GROSH1@LLNL.GOV

for the National Ignition Facility (NIF), centrifuge modeling for nonproliferation, and high explosives for munitions development.

In the data sciences domain, we embarked on an extremely broad range of efforts to understand and extract information contained in large data sets. We are developing new topological approaches based on Morse theory to find features in large-scale simulation data. These quantitative measures allow scientists to better understand and draw conclusions from complex phenomena. One application of this work is understanding and quantifying turbulent mixing of gases. We also combined our recent research in text processing (e.g., co-reference algorithms to identify similar text entities) with new methods of categorizing and organizing information, allowing analysts to rapidly scan large sets of documents. Computer network defense emerged as one of the most important and timely cyber-security research efforts in the Directorate. Specifically, we are exploring the application of algorithms in machine learning to identify Web addresses that may harbor cyber attacks. In addition, we are examining how behavioral-based approaches can be used to detect intrusion attempts.

To accomplish our research and advanced development activities, the Computation Directorate carries out a program that includes internal funding (via the Laboratory Directed Research and Development program), programmatic support (via the National Nuclear Security Administration's Advanced Simulation and Computing Program), and extramural programs (e.g., the Department of Energy Office of Science). In addition, Computation partners extensively with other national laboratories, academia, and industry. We engage in strategic planning exercises and research projects listed in this section (and others not included in this report) that enable us to maximize focus and personnel for both research and mission programs.

2.01

Phase-Field Models Solve Medium-Scale Materials Simulation Problems

SIGNIFICANCE

The need to computationally predict the behavior of materials, often under extreme conditions, arises in many LLNL programs. Such predictions typically require researchers to consider a wide range of spatial and temporal scales. Phase-field models (PFMs) represent one approach for bridging the gap between microscale atomistic and macroscale (e.g., hydrodynamic) simulations. We are collaborating with the Physical and Life Sciences Directorate and the Weapons and Complex Integration (WCI) Principal Directorate to develop a new PFM capability to solve materials science problems, including a PFM of microstructure evolution in polycrystalline materials.

DORR

CONTACT INFORMATION

MILO DORR,
MILODORR@LLNL.GOV
WEB SITE: COMPUTATION.LLNL.GOV/CASC/PHASE_FIELD_MODELING

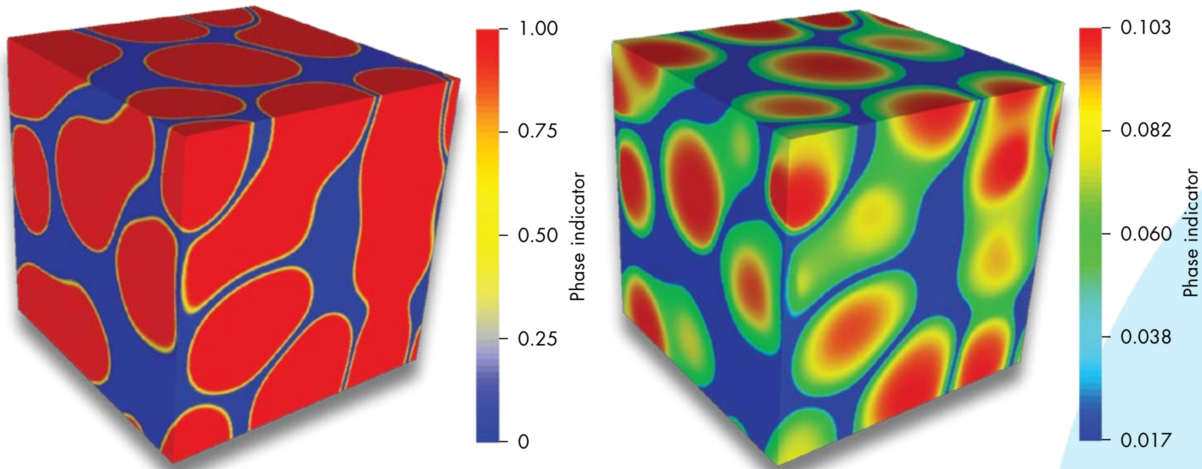

AMPE prediction of grain growth in a binary alloy. On the left is phase (blue is body-centered cubic, and red is face-centered cubic). On the right is the species-relative composition.

PROGRESS IN 2008

We have developed a new code, called Adaptive Mesh Phase Evolution (AMPE), to solve a PFM of microstructure evolution in polycrystalline materials. The model used in AMPE evolves the boundaries between phases, grains, and species as diffuse interfaces in such a way as to reduce the total energy of the system in the quickest manner. In addition to local order and species parameters, the model includes a quaternion representation of the orientation of individual grains. AMPE utilizes an implicit time integrator from the SUNDIALS library to numerically solve the resulting coupled nonlinear diffusion equations. A Newton–Krylov algorithm is employed to solve the nonlinear system at each time step, which ultimately requires the solution of a block preconditioner system. AMPE is built using the SAMRAI library, which provides adaptive mesh refinement. This capability enables the moving interfaces to be adequately resolved without requiring a fine mesh of the entire problem domain.

An example of the results predicted by AMPE is shown in the figure. The simulation was initialized by 20 nonoverlapping grains in face-centered cubic phase distributed randomly in a background body-centered cubic lattice. Each grain was assigned one of three possible orientations and an initial species composition. As the temperature of the system is lowered, the face-centered cubic grains grow. Eventually, grains with the same orientation begin to coalesce, and the misoriented grains form interfaces as the system follows a more favorable energy minimization path. Because of the relatively slow diffusion of species composition in grains, a feature called coring develops, which is indicated by the color gradient within the grains in the right side of the figure.

The fully implicit, adaptive solution of a PFM system including grain orientation required the development and application of advanced numerical methodologies, and it has furthered the state of the art in PFM simulation. The creation of AMPE depended on software, such as SUNDIALS, *hypre*, and SAMRAI, that was developed by scientists in the Center for Applied Scientific Computing. Future versions of AMPE will include models of elasticity and heat conduction as well as phase-order and composition parameters to address a variety of materials science problems at LLNL.

2.02

ISCR Charts a New Course for Scientific Computing

SIGNIFICANCE

This year, the central mission of ISCR shifted from catalyzing interactions between academia and the LLNL computer science and applied mathematics communities to demonstrating and extending LLNL's leadership in simulation capability. ISCR teams with programs and directorates across the Laboratory to extend the boundaries of what is computationally possible and delivers game-changing scientific simulations in support of Laboratory missions. With its new emphasis on the scientific application of computational innovation, ISCR becomes the focal point for ultra-high-performance scientific computing at LLNL and a magnet for computing talent worldwide.

The Institute continues its mission of outreach and education, with a goal of broadening the base of scientists and engineers who are effective practitioners of high-performance computing. To fulfill this goal at the Laboratory, ISCR works with Computation's Multiprogrammatic and Institutional Computing Program to administer LLNL's Computing Grand Challenge (CGC) Program. The CGC Program allocates dedicated capability computing time to projects through a competitive, peer-reviewed process. The Institute also sponsors the Masterworks Lecture Series, which enriches the intellectual atmosphere of LLNL's large simulation community through visits from leaders in diverse areas of computation.

STREITZ

MCGRAW

CONTACT INFORMATION

FRED STREITZ, STREITZ1@LLNL.GOV
JAMES MCGRAW, MCGRAW1@LLNL.GOV
WEB SITE: ISCR.LLNL.GOV

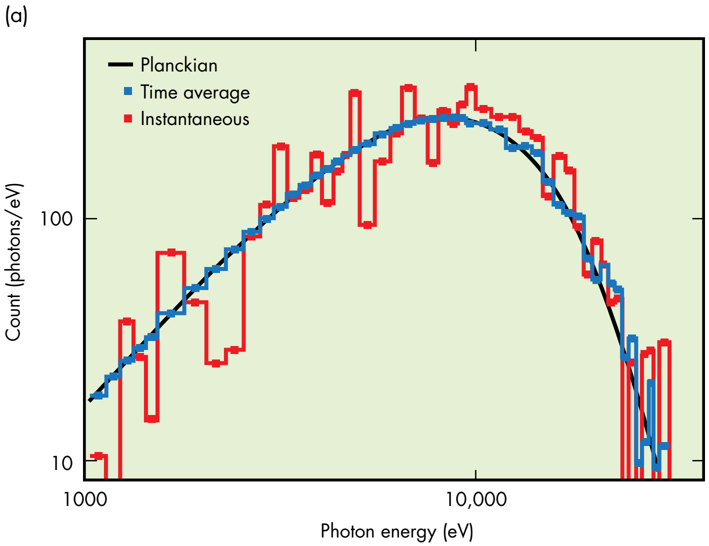

Radiation spectrum developing in a molecular dynamics simulation of hydrogen plasma at 3 kiloelectron volts (keV). The red line represents the instantaneous value of the spectrum after 2 picoseconds of equilibration; the blue line represents a time average. The black line indicates theoretical blackbody spectrum at 3 keV.

PROGRESS IN 2008

In May, Fred Streitz was appointed Director of ISCR. (Streitz is also the Modeling and Simulations Group Leader in the Physical and Life Sciences Directorate.) Since joining the Laboratory in 1999, he has twice led multi-institutional teams that were awarded the Gordon Bell Prize for significant contributions to supercomputing—in 2005 for a simulation of metal solidification, which was the first to break the 100-trillion floating-point operations per second barrier, and in 2007 for a landmark simulation of fluid instability forming at the atomic scale.

Streitz and ISCR Deputy Director Jim McGraw assembled a core team (including veteran members of three Gordon Bell Prize-winning teams: Bor Chan, Milo Dorr, Eric Draeger, Jean-Luc Fattebert, Jim Glosli, Liam Krauss, and David Richards) to partner with scientists in the WCI and NIF and Photon Science principal directorates to develop a highly scalable, multiphysics, N-body computational capability to simulate hot, dense radiative plasmas undergoing thermonuclear burn. The team, in support of a WCI-led Strategic Initiative, retooled the highly scalable ddcMD code to model a hydrogen plasma; their results were published earlier this year in *Physical Review E.* They also succeeded in incorporating simple radiation (brehmsstrahlung emission and absorption) and two-body collisions (figure a).

The CGC Program continues to produce significant results across a broad spectrum of scientific disciplines, as it has for several years. In February, an LLNL symposium showcased 16 technical presentations on CGC results; these presentations are available on the ISCR Web site. Key science highlights include:

- Three-dimensional NIF simulations of nearly whole beams captured the effects of hydrodynamic gradients, refraction, and laser speckles. These simulations are unprecedented in size and the physics incorporated, and they provide LLNL scientists with the capability to simulate beam propagation in ignition targets (figure b).
- The first-ever approximation-free exact dynamic mean field theory calculations of plutonium proved that plutonium is a strongly correlated Fermi liquid.
- Large-scale dislocation dynamics simulations contributed a better understanding of dislocation interactions and patterning. Results showed that the strength of the dislocation interactions changes with dislocation structure and loading conditions.
- A novel hybrid quantum molecular dynamics (QMD)/molecular dynamics algorithm was designed and successfully demonstrated with a prediction of the atomic structure of liquid molybdenum.
- The first full-scale simulations of a proposed Laser Wakefield Accelerator demonstrated the feasibility of 10-gigaelectron volt gain and explained the anomalous electron transport inhibition seen in earlier experiments.

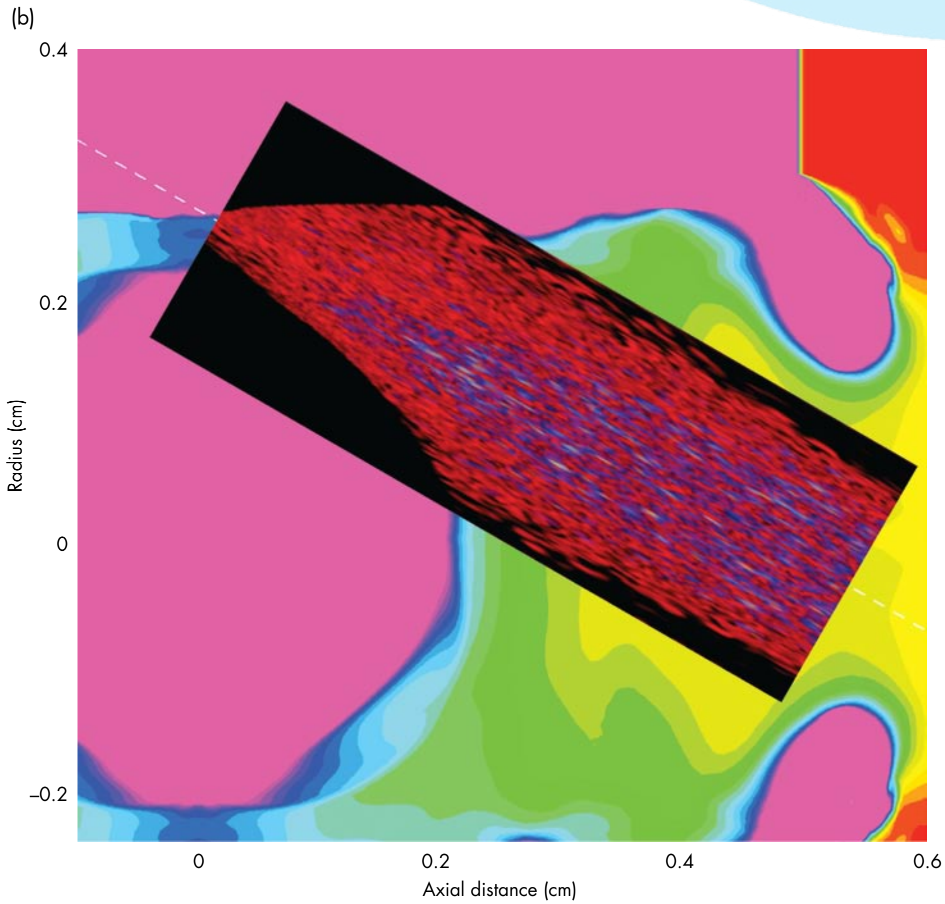

This propagation simulation of the 285-electron volt (eV) point design (30-beam, circa peak power) shows low reflectivity. The calculation used 4,096 Atlas processors and ran for approximately 35,000 processor days.

In June, the results of the most recent CGC computer allocations were announced. Seven proposals received Tier-1 allocations (80–90% of the requested computing time), and 13 received Tier-2 allocations (30–70% of the requested computing time). In total, almost 1.2-million hours of computing time per week were allocated on the Atlas system and almost 420,000 hours on the Thunder system.

The ISCR Masterworks Lecture Series is a reincarnation of the highly regarded Institute for Terascale Simulation Lecture Series, which has been organized by ISCR for the past six years on behalf of the Advanced Simulation and Computing (ASC) Program Office. The 2008 speakers were Luis von Ahn (Carnegie-Mellon University), William Gropp (University of Illinois), and Linda Petzold, (UC Santa Barbara). Saul Teukolsky (Cornell University) will begin the 2009 series in January.

Overture Software Gains New Parallel Capabilities and Widens Its Application Base

SIGNIFICANCE

The Overture team is developing the mathematical theory, numerical algorithms, and software to solve partial differential equations (PDEs) in complex moving geometry. The Overture framework uses composite overlapping grids to obtain accurate results in an efficient and flexible manner. Accuracy is obtained using smooth boundary-fitted grids, and efficiency is achieved using structured and Cartesian grids. The Overture software tool kit includes support for representing geometry, computer-aided design cleanup, grid generation, adaptive mesh refinement (AMR), moving grids, graphics, and discretization of equations and boundary conditions.

The composite-grid (CG) suite of PDE solvers is built using Overture. This suite includes capabilities to solve low- and high-speed fluid flows, heat transfer, solid mechanics, and electromagnetics. A multiphysics solver is being developed to address coupled fluid–solid problems, such as conjugate heat-transfer and fluid-structure interactions. Scientists and students at LLNL, the Department of Energy, and around the world are using these algorithms and tools to solve a variety of scientific and industrial problems. The CG suite is being used at LLNL to simulate high-speed rotating flows, detonations in condensed phase explosives, conjugate heat transfer in the fuel assembly of a nuclear reactor, blood flow in veins with implanted blood-clot traps, and the propagation of lasers through lens surfaces in NIF. The CG solvers are also being used at universities, along with new Overture-based physics solvers developed by students and professors, to model flapping airfoils, aircraft icing, flow on the surface of the eye, shock focusing, friction-stir welding, and relativistic hydrodynamics.

HENSHAW

CONTACT INFORMATION

BILL HENSHAW,
HENSHAW@LLNL.GOV

PROGRESS IN 2008

The Overture team developed several significant new capabilities this year, including the first parallel AMR capability for overlapping grids (figure a). AMR can efficiently model problems that contain fine-scale features such as shocks and detonations. Refinement grids are automatically generated every few time steps based on an error estimator. On an overlapping grid, refinement grids that exist on the different component grids are connected through interpolation.

(a)

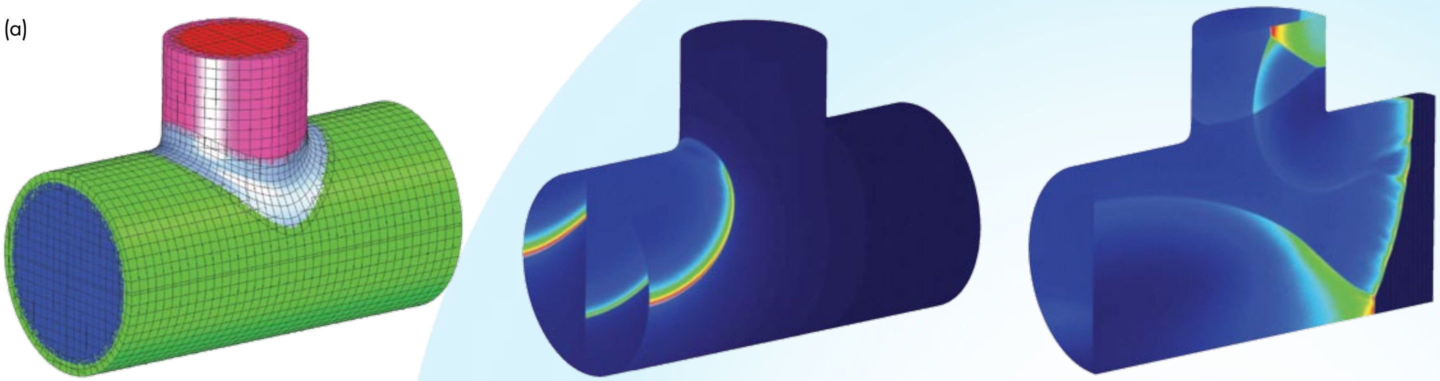

Overlapping grid — Pressure at t=2.0 units — Pressure at t=2.8 units

Parallel AMR is used to perform high-resolution simulations of a gas-phase detonation in a pipe with a T-junction. Refinement grids are automatically generated every few time steps to provide higher resolution where it is needed.

The team also developed new multidomain algorithms to simulate conjugate heat-transfer problems that couple fluid flow to heat transfer in solids. A mathematical analysis of some model problems led to more accurate and stable interface approximations for conjugate heat transfer. This capability was used to model a nuclear reactor fuel assembly (figure b) and convection in the NIF hohlraum target.

The team developed adaptive algorithms to solve high-speed reactive flow used in high-resolution simulations of explosives (figure c). This capability resolved a controversy related to the formation of dead zones (unexploded material) in corner-turning problems and led to an improved macroscale model that includes desensitized effects. The new model predicts the formation of dead zones as seen in experiments.

High-order accurate algorithms for simulating electromagnetics and material interfaces on overlapping grids are being used to study a high-intensity laser as it propagates through the defects on the surface of a NIF lens (figure d). This capability will help quantify how the defects perturb the laser. Such information will be used to prevent damage to the optics.

(b)

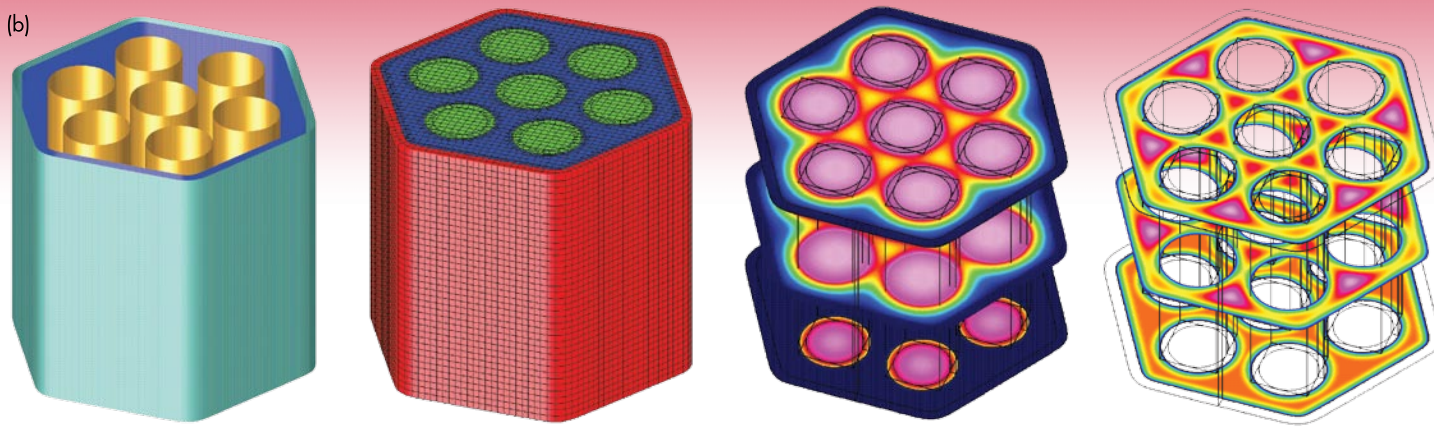

Fuel assembly geometry | Composite grid | Temperature | Flow speed

The incompressible fluid flow and heat transfer in a nuclear reactor fuel assembly are modeled with new multiphysics conjugate heat-transfer capabilities based on overlapping grids. New approximations were developed in 2008 to treat the interfaces between solids and fluids.

(c)

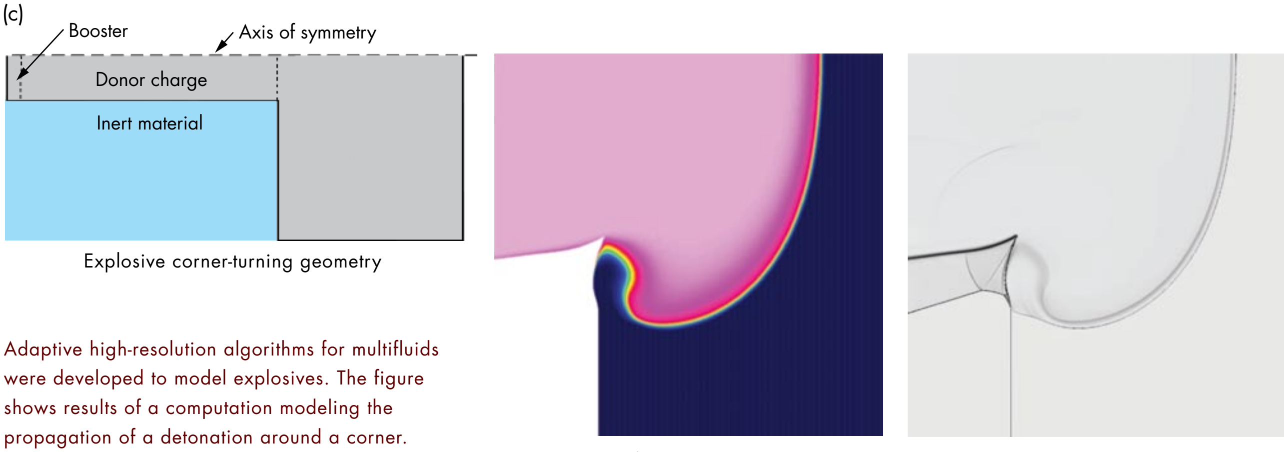

Schlieren plot of inert and reactive material | Reaction progress variable

Adaptive high-resolution algorithms for multifluids were developed to model explosives. The figure shows results of a computation modeling the propagation of a detonation around a corner.

(d)

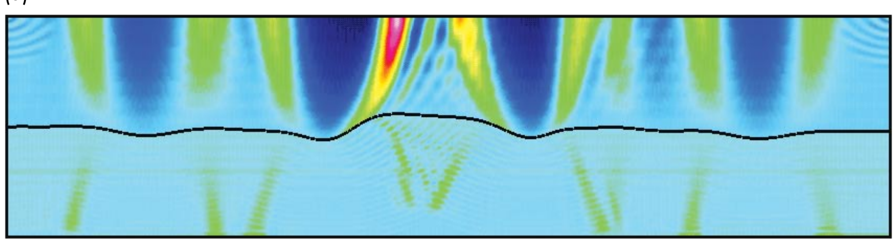

Air

Glass

The laser intensity near the surface of a NIF lens is determined by solving Maxwell's equations with a fourth-order accurate method that treats the interface between the glass lens and air.

DocEx Increases the Speed and Ease of Text Analysis

SIGNIFICANCE

In late 2007, a U.S. government agency sought LLNL assistance with triaging a large number of high-value foreign language documents. It was taking the agency weeks to catalog, summarize, and prioritize the document assets using its outdated tools. In response to the request for help, LLNL began a project called DocEx (for document exploitation) to develop a proof-of-concept prototype triage system by integrating commercial tools with existing Laboratory analysis software. Initial document exploitation exercises using DocEx on the sponsor's data suggested a triage time reduction from weeks to hours for a typical data set. Based on these initial results, the sponsor is funding additional DocEx technology development and deployment of the DocEx tools at the sponsor's site.

The DocEx approach can be applied to various problems involving the triage and rapid processing of large data collections. These tools are being used in two other projects to analyze large collections of structured and unstructured records, and two years of additional funding are expected next year.

POLAND

CONTACT INFORMATION

DOUGLAS POLAND,
POLAND1@LLNL.GOV

PROGRESS IN 2008

In its first year, the DocEx project team made tremendous progress toward developing an integrated system that leverages commercial technology, open-source software, and internal LLNL research products. DocEx features custom ingest pipelines that convert documents; extract blocks of text, tagged entities, and images; and extract keywords and untagged entities from the unstructured text fields. The results are indexed using a commercial faceted search engine (Endeca), and the analyst can navigate the index using a custom DocEx interface through a standard Web browser. The project team also incorporated commercial machine translation tools into the pipeline and provided basic export capabilities to the LLNL Everest (graph visualization and analysis) and FactWeave (multiple hypothesis and timeline tracking) platforms.

The DocEx team is researching how best to use active real-time user feedback to incrementally model user interests. The model scores and ranks new documents and assists analysts in identifying the documents that are "most interesting" for their particular task. The model includes adaptive fine-grained topic models, source reputation, writing style, and other features that can be extracted from the document text and metadata. The features are combined in an online machine-learning framework that continually adapts to analyst feedback.

The team applied cutting-edge research in probabilistic topic models to generate corpus-specific topic sets. This feature allows faster recognition of the common themes in a group of documents. The team is also researching techniques to supplement human-generated concept hierarchies with automatically generated, domain-specific hierarchies for specialized corpora. Effective hierarchies form the basis for efficient, intuitive faceted browsing, which allows analysts to drill down to specific portions of a corpus without having to formulate queries.

One challenge associated with exploiting large data collections is that specific named entities may be referred to in many different ways, and multiple entities may have similar or identical names. The DocEx team is developing approaches to disambiguate entities based on their relationships, which will provide a more precise view of the entities in the data set and will generate canonical forms for the entity names.

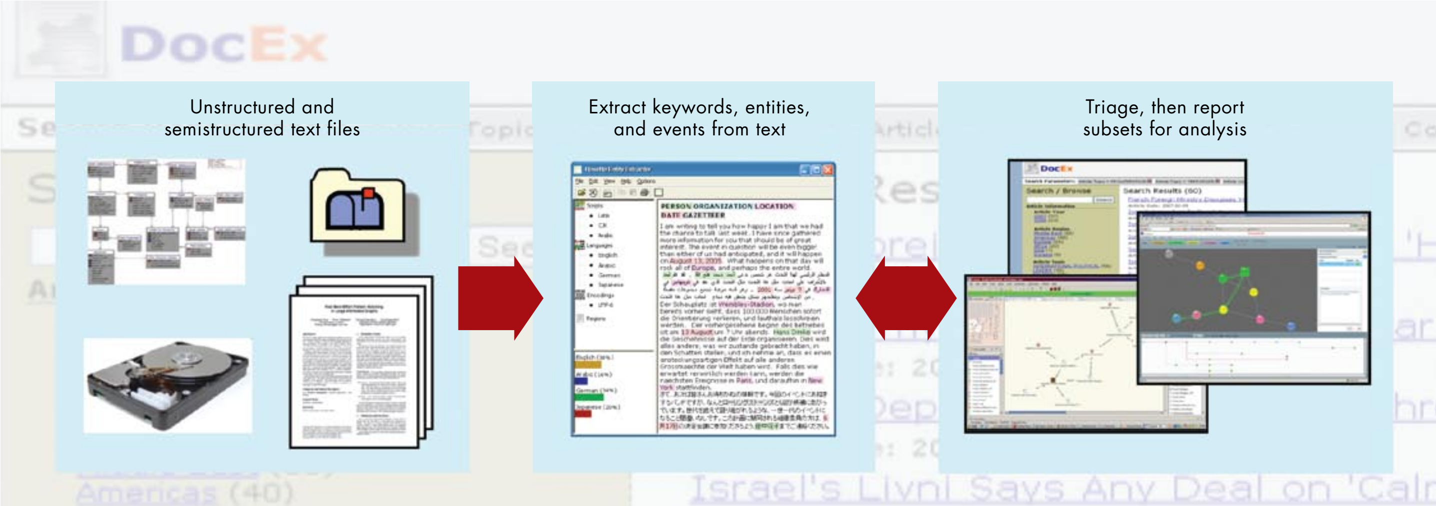

DocEx provides an interface and a system for performing document triage and exploitation. An analyst can search and browse collections of text databases, e-mails, reports, and more. Relevant results are located via a faceted search and then are exported directly to graph or hypothesis-based analytic tools.

New Morse Theory-Based Technique Brings Clarity to Turbulent Combustion

SIGNIFICANCE

Computation scientists are developing a suite of general-purpose data analysis tools based on Morse theory to reliably and efficiently analyze scientific data sets. By applying Morse theory to convert large scientific data sets into hierarchically encoded topological graphs, we can reduce the size of the data sets by several orders of magnitude. The hierarchical models provide complete families of structural segmentations, which enable the data to be explored at multiple scales and with extensive parameter ranges. Our framework is built on provably robust and streaming algorithms that enable us to process large-scale scientific data on commodity hardware. This technique provides scientists with insight into several application areas, including combustion simulations, hydrodynamic instabilities, and the structural analysis of porous media.

PROGRESS IN 2008

In 2008, we extended our framework by incorporating augmented threshold-based hierarchies, analyzing the temporal evolution of features in time-dependent data sets, and customizing

BREMER

CONTACT INFORMATION

PEER-TIMO BREMER,
BREMER5@LLNL.GOV

the system to support different applications. We applied the enhanced framework to topological analyses of several simulations designed to study the effects of turbulence on the combustion process. Understanding combustion over a broad range of operational regimes is important to many applications, including engine or power plant design. Of particular interest are premixed burners that utilize ultra-lean hydrogen–air fuel mixtures because they minimize exhaust gas temperatures and reduce emissions. However, lean premixed flames, and hydrogen–air mixtures in particular, are subject to a variety of flame-induced hydrodynamic and combustion instabilities that render consistent flame stabilization difficult. The flames burn in a cellular mode that is highly nonuniform, time-dependent, and difficult to characterize.

To study the combustion process at different levels of turbulence, application scientists used three numerical simulations of lean premixed hydrogen flames in three dimensions to demonstrate no turbulence, weak turbulence, and strong turbulence, respectively. The burning behavior is characterized by the number and size of burning cells. Flames are represented as isotherms of constant temperature, and thresholds on the local fuel consumption divide the surfaces into burning cells separated by nonburning regions. However, no unique correct threshold exists, so a primary goal was

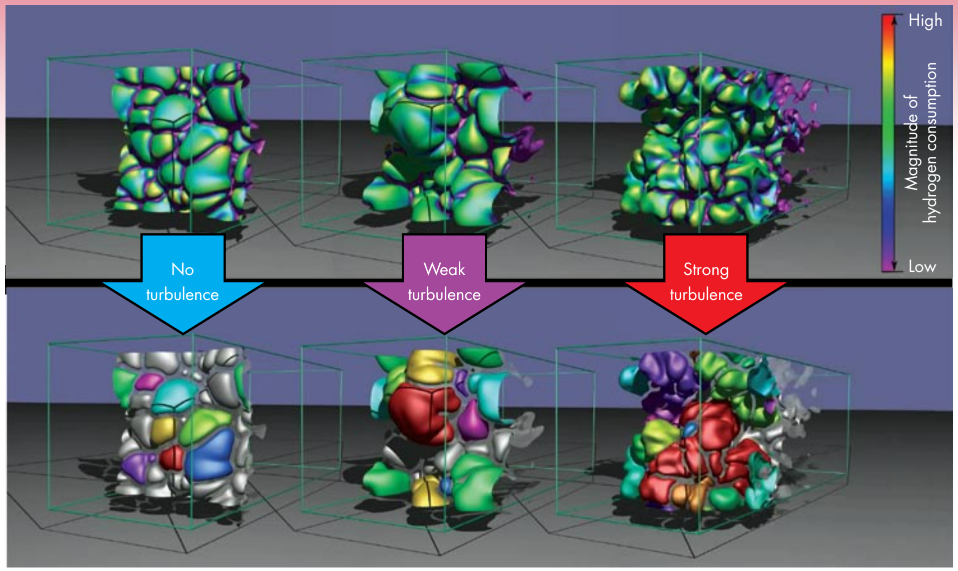

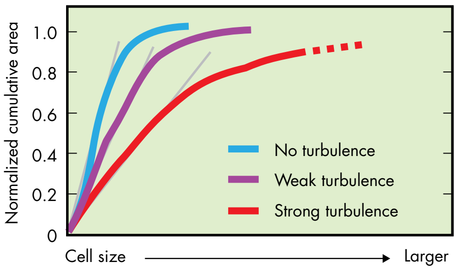

The top three-dimensional (3D) diagrams show the hydrogen consumption on the flame surfaces at different turbulence levels, with darker colors indicating less fuel consumption. For each turbulence level in the bottom 3D diagrams, a small set of burning cells are randomly colored to show the irregularity of the more turbulent cells. In the graph, the corresponding cumulative densities of cell area distributions show that more turbulence creates larger and more irregular cells with wider distribution of normalized surface areas.

to determine and verify the choice of parameters.

For each time step of the simulation, we extracted the isotherms in a streaming fashion and computed merge trees of the fuel-consumption function defined on the isotherms. Each tree encodes the cell configurations for all possible thresholds. We augmented the trees by storing cell areas as a function of the threshold along their corresponding branches. We used the augmented trees to perform extensive parameter studies, which determined a viable threshold and demonstrated the consistency of the results using varying parameters. Once a threshold had been determined, we tracked the resulting burning cells through time and created a tracking graph that encoded their temporal evolution.

The methods we applied allow, for the first time, a quantitative analysis of the cellular burning structures and yield important scientific insight. Although it seems counterintuitive, higher turbulence levels lead to larger cell structures, which also burn more intensely than simple theories of flame propagation indicate. These results suggest that premixed hydrogen flames could be stabilized at much leaner conditions than previously believed.

2.06

Machine-Learning Algorithms Aid Cyber-Security Defense

SIGNIFICANCE

In 2008, the National Academy of Engineering listed securing cyber space among its 14 grand challenges. Today's cyber defense tends to be reactive and signature-based and occurs only at the network's perimeter. Tomorrow's cyber defense must be proactive, dynamic, and ubiquitous. We are developing techniques that will facilitate a better cyber-security posture. For example, we are applying machine-learning algorithms to build statistical models of micro- and macro-level behaviors based on observations of activities on a network over time. Our efforts—supported by LLNL research and development funds—have been instrumental in demonstrating the benefits of using machine-learning algorithms for various cyber-security tasks, such as detecting attacks in Hypertext Transfer Protocol (HTTP) requests and analyzing Internet Protocol (IP) address behavior in network traffic data (figure a) for malicious indicators. Our early successes attracted the attention of a key sponsor from the intelligence community, and both of the algorithms described in this article are part of programmatic software deliverables for FY09.

ELIASSI-RAD

CONTACT INFORMATION

TINA ELIASSI-RAD,
ELIASSIRAD1@LLNL.GOV

PROGRESS IN 2008

In 2008, we developed approaches to detect attacks in HTTP requests. Cyber attacks on the World Wide Web are increasing in frequency and severity as the world relies more on Web applications for commercial, financial, and medical transactions. Web applications provide an attractive alternative to traditional desktop applications because they are accessible and easy to deploy. Unfortunately, the distributed nature of Web applications and the complexity of configuring application servers have led to the proliferation of Web-based attacks, in which attackers surreptitiously inject code into HTTP requests. The attackers then execute arbitrary commands on remote systems and perform malicious activities, such as reading, altering, or destroying sensitive data.

One approach to protect against HTTP-based attacks is to identify malicious code in incoming requests and eliminate bad requests before they are processed. To this end, we developed a simple but effective HTTP attack classifier, using machine-learning techniques, that automatically labels requests as "valid" or "attack." By basing the classifier on the vector-space model commonly used for information retrieval, the classifier not only separates attacks from valid requests but also identifies the specific type of attack (e.g., Structured Query Language injection, path traversal, cross-site scripting). Moreover, because the computational complexity of our approach is linear in the size of a request (i.e., proportional to the time needed to read the data), our approach can be used on data streams.

We demonstrated the effectiveness of our HTTP attack classifier through experiments on the discovery challenge data set from the 2007 European Conference on Machine Learning and Principles

(a)

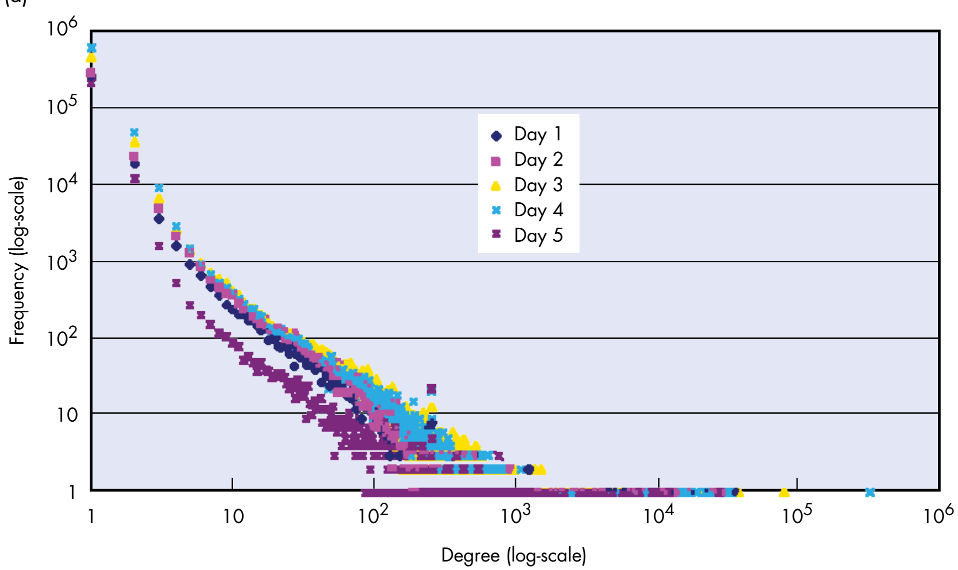

Degree distribution of the IP-to-IP communication graph for each of five days. There are 2.9 million IP addresses and 6.6 million communication links over the five-day period.

and Practice of Knowledge Discovery in Databases. Our approach achieved higher precision (98%) and recall (88%) than other published results. Compared to the nearest competing algorithm, our approach produced 19% better precision and 13% better recall. In addition, our approach has several unique and desirable characteristics, including robustness to missing contextual information, interpretability of models, and scalability.

We also investigated algorithms to analyze IP address behavior in network traffic data. Cyber-security analysts often have access to IP traffic over a network (trace data). However, detecting cyber threats (e.g., compromised IPs) among the traffic is a challenge because the volume and diversity of legitimate traffic makes it difficult to detect the skew of malicious traffic. Moreover, a single IP address can dynamically map to several physical machines during a given period, and the behavior of a machine can change (e.g., after a virus infection). While there are some promising methods for classifying network traffic, most of them require analyzing the packet payloads (which are computationally costly) and/or prior knowledge of malicious behavior (which makes them vulnerable to novel malevolent activity).

Given network trace data that evolves over time, we built a dynamic Latent Dirichlet Allocation for Graphs (LDA-G) behavior analysis algorithm that uses only the topology of the dynamic IP-to-IP communication graph. Dynamic LDA-G is a nonparametric Bayesian model that discovers latent group structure in time-evolving graphs and produces two matrices: (1) the IP × Group Membership Matrix, where each entry is the probability of an IP belonging to a specific group; and (2) the Group × Group Likelihood Matrix, where each entry is the probability of a connection between two given groups. The statistical models inferred by LDA-G can perform link prediction between a pair of IPs and discover "unstable" IPs with volatile communication patterns.

Our experimental study on trace data collected over five days at an access link demonstrated that dynamic LDA-G is highly predictive of communications (i.e., links) between IP addresses. The average area under the Receiver Operating Characteristic curve of our algorithm was 0.91 compared to the default performance of 0.51. (Average area under the curve, or AUC, is equal to the probability of ranking a randomly chosen positive instance higher than a randomly chosen negative instance. An area of 1 represents perfect performance; an area of 0.5 represents a random performance.) LDA-G performs well even when observed data are scarce. Our experiments showed that the slope of the curve representing average AUC on link prediction (y axis) versus the percentage of observed links (x axis) is 0.0023. In other words, LDA-G's performance decreases very slowly as the number of observed communications decreases sharply. This robustness makes LDA-G a desirable approach in real-world settings. Furthermore, our study showed that dynamic LDA-G can accurately flag various IP addresses based on their volatility, a measure of how much an IP address changes its group membership (or role) over time (figure b).

(b)

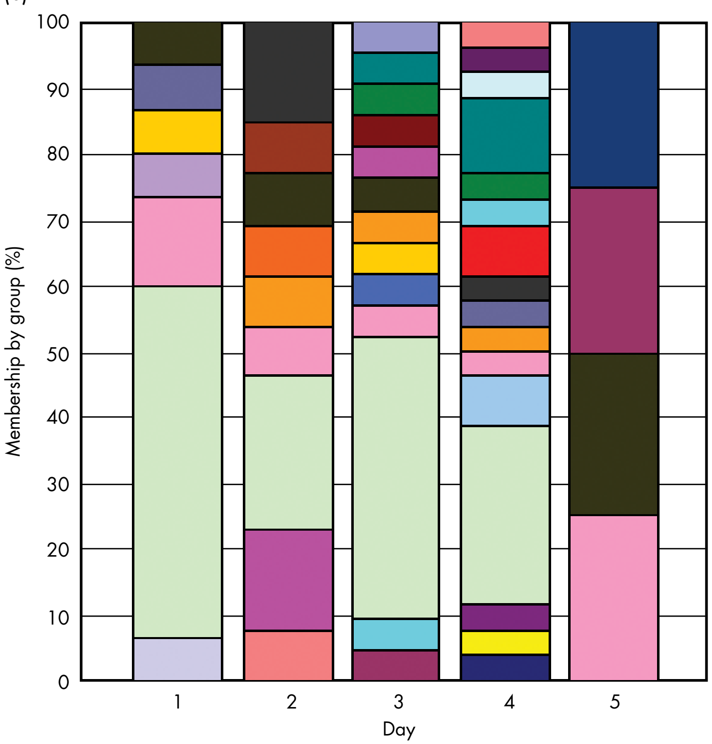

Group memberships of a compromised IP address over five days. The colors represent different groups. This IP address has relatively high volatility and unstable group memberships.

2.07

New Analysis Helps Answer Parallel Scalability Questions

FALGOUT

CONTACT INFORMATION

ROB FALGOUT,
FALGOUT2@LLNL.GOV

SIGNIFICANCE

The scalable and robust solution of the radiative transfer and neutron transport equations is an important issue in many scientific codes. A key part of these codes is the "sweep," which involves an inherently sequential component that fundamentally limits the amount of parallelism that can be achieved. Although various overloading techniques have been used to amortize the cost of these sweeps, the practicality of scaling to massively parallel machines with tens of thousands of processors has been unclear until now. We derived an expression for the minimum stages possible in the sweep algorithm and demonstrated an algorithm that achieves this minimum. We also performed parallel scaling tests to verify the theory and showed that sweeps can scale acceptably well to huge processor counts.

PROGRESS IN 2008

A potential bottleneck when solving Boltzmann transport equations in parallel is the inversion of the streaming operator. The discretized form of this operator is a lower triangular matrix or block lower triangular matrix with small blocks. The solution of these triangular systems by direct methods involves an inherently sequential component. We have derived new theoretical scaling models for several sweep algorithms, paying particular attention to what we refer to as "collisions," which occur when multiple angles are being swept at once in parallel. The existing literature either ignores this important issue or overestimates its effect on the overall performance of the algorithm. Our theory shows that these algorithms have the potential to scale proportionally to a root (given by the number of spatial dimensions) of the number of processors plus the number of angular directions. When the number of directions is fairly large, the dependence on the number of processors is masked, thereby delaying the poor scaling behavior. In many cases, this delay is adequate to achieve practical performance, even up to tens or hundreds of thousands of processors.

As part of this work, we derived an expression for the minimum number of stages possible in a sweep algorithm and demonstrated a new algorithm that achieves this minimum. We also showed that the choice of which angle to sweep in a given processor at some stage of the algorithm affects the number of overall stages required, especially in three dimensions (3D) at large processor counts. In addition, we showed that transport sweeps can scale acceptably well to huge processor counts. Figure (a) shows that an S10 discretization would be only two times slower on 100,000 processors than on 27 processors. The figure also shows that increased angular discretization accuracy can further improve scaling, assuming that there is enough memory to store the necessary angle data on each processor. These scaling predictions have been verified in parallel on up to 8,000 processors of Atlas, a 2.4GHz AMD Opteron cluster (with eight processors per node) at LLNL (figure b).

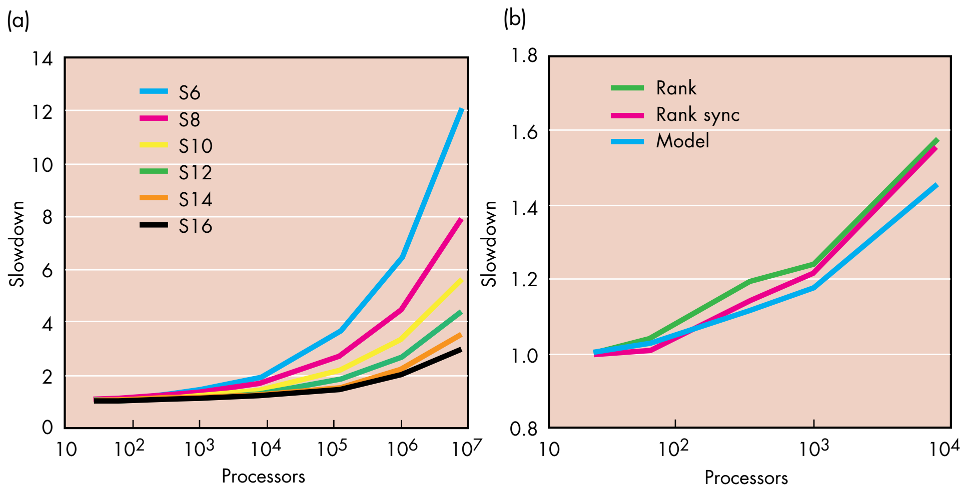

Scaling results for increasing problem sizes on proportionally more processors. Slowdown is defined as the run time for a sweep on one processor configuration divided by the run time on 27 processors. (a) Sweeps can scale acceptably well to huge processor counts in 3D, and increased angular accuracy further improves scaling. This plot shows the slowdown predicted for our minimal-stage data-driven sweep algorithm. (b) The slowdown predicted by our model compares well to actual slowdowns obtained from the ARGES code on LLNL's Atlas parallel computer.

APPLICATIONS SOFTWARE DEVELOPMENT

3

3.00 Computer Scientists Dramatically Impact Laboratory Programs One Software Application at a Time

The Computing Applications and Research (CAR) Department stewards Lawrence Livermore National Laboratory's (LLNL's) cutting-edge capabilities in software technologies and applications. Our staff collaborates with universities, industry, and other national laboratories to perform research and development in computer science, mathematics, and scientific computing. We focus on advancing the disciplines within the international scientific community and supporting the long-term programmatic needs of LLNL and Department of Energy (DOE) programs.

In addition to defining and advancing the state of the art, CAR also embeds staff in Laboratory programs—computer scientists who develop the software technologies and applications required to execute LLNL's national security mission. The work and culture vary widely across the programs, creating a diverse set of rich opportunities in which employees apply and broaden their skills and contribute immediately and directly to mission outcomes. To truly appreciate the work and contributions of the computer scientists embedded in the programs, an understanding of the mission and customer is required. This section will highlight examples

NEXT PAGE

DAMKROGER

CONTACT INFORMATION

TRISH DAMKROGER,
DAMKROGER1@LLNL.GOV

from a few Laboratory programs and what we refer to as "principal directorates."

The Weapons and Complex Integration Principal Directorate is charged with ensuring the safety, security, and reliability of the nuclear stockpile. In the absence of underground testing, stockpile stewardship relies on computational models and simulations to understand the effects of aging, defects, and reliability-based modifications on weapon performance. The Advanced Simulation and Computing (ASC) Program partners computer scientists and mathematicians with nuclear physicists to develop the multiphysics, multidimensional codes required to simulate nuclear tests. One example of an ASC-developed code is EMSolve, which is used to model electromagnetic fields. VisIt is another tool used throughout the nuclear weapons complex to visualize the output of ASC models and codes.

The Global Security Principal Directorate focuses on critical problems in the areas of energy, environment, military applications, biodefense, and nonproliferation, and also supports the nation's intelligence community. This diverse and dynamic organization creates a challenging set of requirements for computer scientists. Some projects are long-term endeavors with broad responsibilities, requiring a strategic balance of capability development and application. Other projects arise overnight and require immediate engagement and intense, short-duration efforts to produce crucial results. The Laboratory's Russian Monitoring Program is a long-term, sustained effort, where the need to balance transparency and Russian security concerns places unique demands on the computational infrastructure. We are addressing this challenge by creating the infrastructure needed to verify the presence of weapons-grade material in containers, without compromising sensitive design details. In another area of Global Security, the Biodefense Knowledge Center is charged with maintaining a comprehensive, up-to-the-minute database of the nation's biodefense information—in a form that is immediately and easily accessible for responding to emergency situations. CAR computer scientists have overcome this challenge and created a usable data store from vast amounts of heterogeneous and unstructured data.

The National Ignition Facility (NIF) is home to the world's largest laser. When completed in 2009, NIF will focus 192 laser beams on a BB-sized target filled with hydrogen fuel. The intense energy of the combined beams will ignite the hydrogen atoms' nuclei, simulating what happens at the core of a star. CAR's computer scientists are developing the Integrated Command and Control System, an intricate combination of hardware and software that controls all 192 beams of NIF. An equally complex challenge facing computer scientists working in NIF is creating and maintaining the accuracy of the diagnostics necessary to ensure the lasers are functioning properly, focused, and on target.

CAR also supports a wide range of institutional functions, from the Laboratory's physical and cyber-security programs to management and quality assurance. We staff these projects to develop the software technologies and applications required for the needed functionality. For instance, CAR initiated the Laboratory's Software Quality Assurance (SQA) Program in response to a growing need for consistency and high professional standards, and we now provide all Laboratory programs both subject matter experts and the templates required to meet SQA requirements.

Shot Data Analysis Engine Automates Analysis of NIF Diagnostics


BETTENHAUSEN EDWARDS

CONTACT INFORMATION

RITA BETTENHAUSEN, BETTENHAUSEN1@LLNL.GOV
PHILIP EDWARDS, EDWARDS60@LLNL.GOV

SIGNIFICANCE

Scientists working in NIF must quickly analyze and interpret data obtained from more than 20 specialized diagnostic instruments located at the NIF target chamber. While preparing for an experiment, the analyses help guide laser alignment, focus, and other tuning parameters necessary to accomplish experiment goals. After an experiment, the analyses help interpret the raw data collected from the diagnostic instruments. In combination, these analyses constitute the physical measurements on which all NIF experiments are based—experiments that span critical mission areas of stockpile stewardship, high-energy–density physics, astrophysics, and clean fusion energy.

The NIF Shot Data Analysis Engine is the first production-ready automated analysis system of its kind. It integrates commercial workflow tools and messaging technologies within a scientific software architecture to provide shot data analysis. The highly parallel, scalable, and flexible architecture currently meets the analysis requirements for five NIF diagnostics. It will support additional diagnostic instruments for the National Ignition Campaign and is the foundation for automatically analyzing the diagnostic data of future NIF experiments.

PROGRESS IN 2008

The Shot Data Analysis Engine was deployed in 2008, and it now supports instrument correction and diagnostic analysis for five NIF diagnostics. The highly flexible and scalable analysis framework (see figure) features a parallel architecture that provides automatic triggering of analysis upon arrival of new data, analysis workflow sequencing, data provisioning from various sources, data mapping to analysis functions written in interactive data language (IDL), and results archiving with pedigree (a record of the data inputs and analysis software version).

The architecture's scalability leverages two technologies: (1) message queues with Java messaging that dynamically schedule and balance analysis tasks across all available resources—processes and processors; and (2) a commercial workflow processor based on the industry standard Business Process Execution Language (BPEL), which contributes an "orchestration" paradigm that integrates and coordinates external systems through Web services.

A Java-based analysis director application processes data arrival events (both raw and intermediate analysis results). When data arrives, the analysis director schedules the appropriate analysis workflow on the analysis engine in a queue. The BPEL workflow processor then takes the analysis workflow request and performs three functions. It (1) orchestrates the sequence of analysis steps; (2) integrates, transforms, and transfers data between the various data repositories and the analysis modules; and (3) schedules the appropriate analysis on the analysis servers and retrieves the results from the adapter-in and adapter-out queues. The results are stored in a content management system archive.

At the end of the analysis pipeline is a "pool" of analysis servers, which monitor the adapter-in queue to receive "analysis invocation" messages. This processing pool can contain multiple computer processors, with multiple analysis servers running on each processor. Each analysis server utilizes IDL and an instance of the analysis adapter. The analysis adapter provides the generic framework that interprets (adapts) an analysis request into an algorithm-specific IDL invocation.

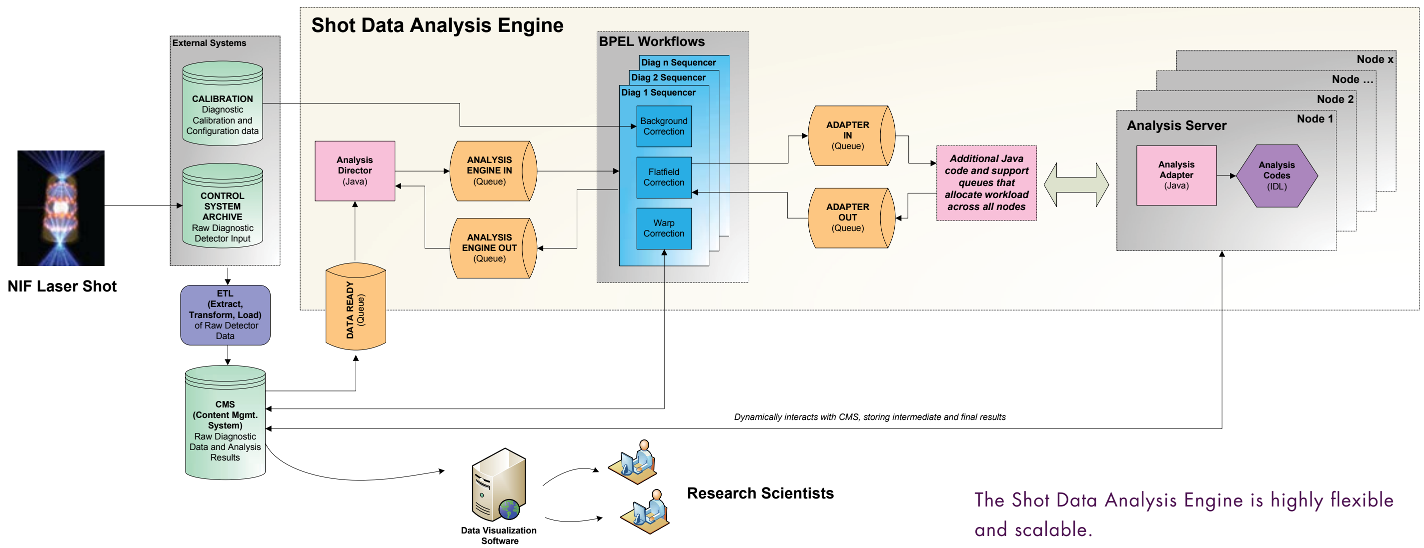

The Shot Data Analysis Engine is highly flexible and scalable.

3.02

EMSolve Leads the Way in Electromagnetic Analysis

SIGNIFICANCE

EMSolve is a suite of codes designed to model electromagnetic fields in many situations. The codes use various finite element basis functions and yield provably stable, charge- and energy-conserving solutions to the Maxwell equations. These basis functions support various standard two- and three-dimensional element types, including high-order elements for accurately representing curved geometries. The basis functions are constructed using arbitrary order polynomials to provide high-fidelity solutions on coarse meshes.

The EMSolve code suite consists of separate applications that share a common framework and have been tailored to efficiently address the requirements of several types of electromagnetic phenomena. The suite includes applications to model the second-order Helmholtz wave equation, the coupled first-order wave equation, and the electromagnetic diffusion equation (known as the eddy current approximation). EMSolve also includes frequency domain and eigenvalue codes. Because of its versatility and robust mathematical underpinnings, EMSolve is gaining interest within the scientific community. Its core finite element library has been incorporated into LLNL's ALE3D and HYDRA codes and has also been licensed to a commercial finite element analysis company to provide limited support for electromagnetic phenomena.

STOWELL

CONTACT INFORMATION

MARK STOWELL,
STOWELL@LLNL.GOV

"The benefits of EMSolve will continue to be manifest as computer architectures increase in capabilities. Not only will EMSolve remain one of the most capable computational electromagnetic tools for the size of analysis it enables, it may well be the most accurate capability in the world due to the high-order numerical techniques that it employs."

Robert Sharpe, Engineering Directorate Deputy Associate Director for Science and Technology

PROGRESS IN 2008

In 2008, the EMSolve team added two complex-valued frequency domain codes to their suite. Both codes can be used to model structures containing lossy or imperfect conductors, but they differ in the types of source functions that can be applied and boundary conditions that are available. The first code, developed in collaboration with Ohio State University, provides a new second-order absorbing boundary condition. The second code is formulated in terms of the magnetic field and provides a more accurate representation of surface currents and surface charge densities.

EMSolve was used by members of LLNL's Center for Meso, Micro, and Nanotechnologies (MMNT) working on a Defense Advanced Research Projects Agency initiative on Surface Enhanced Raman Scattering (SERS). SERS techniques can be used for molecular identification, especially as a sensor for biological molecules. SERS detection relies on specially manufactured substrates, which have nanoscale features coated in gold or silver. At optical frequencies, these metallic coatings have a negative relative permittivity, similar to a plasma, which can support plasmonic resonant modes. Effective SERS devices rely on local electric field enhancement from these plasmonic resonances on the surface of the substrate, with the overall SERS signal being proportional to the magnitude of the electric field to the fourth power. The quality of a SERS substrate can be measured by the amplification of the local electric field and the area over which this enhancement exists. The frequency-domain version of EMSolve was used

NEXT PAGE

to study a variety of SERS substrates. Structures such as nanopillars, bow ties, nanospheres, and inverted pyramids were examined at optical wavelengths between 400–800 nanometers (μm). These EMSolve studies helped develop insight into results reported in literature, and they guided MMNT team members in their fabrication and analysis of nanostructures. The simulations represent a new realm for EMSolve due to the very small-length scales and the rather exotic material properties involved.

The EMSolve team also developed a scheme for modeling transient current pulses in their full-wave, time-domain simulation code. The new formulation allows the team to represent current densities using H(div)-conforming finite element basis functions, which, in turn, allows the code to maintain charge conservation as the pulse moves through an unstructured mesh. Consequently, EMSolve can now accurately model problems involving electrostatic discharge or charged particle beams without introducing nonphysical charge buildup and without performing divergence cleaning. This new capability was instrumental in modeling the electromagnetic pulses generated by highly energetic electrons produced in the Titan and NIF laser target chambers.

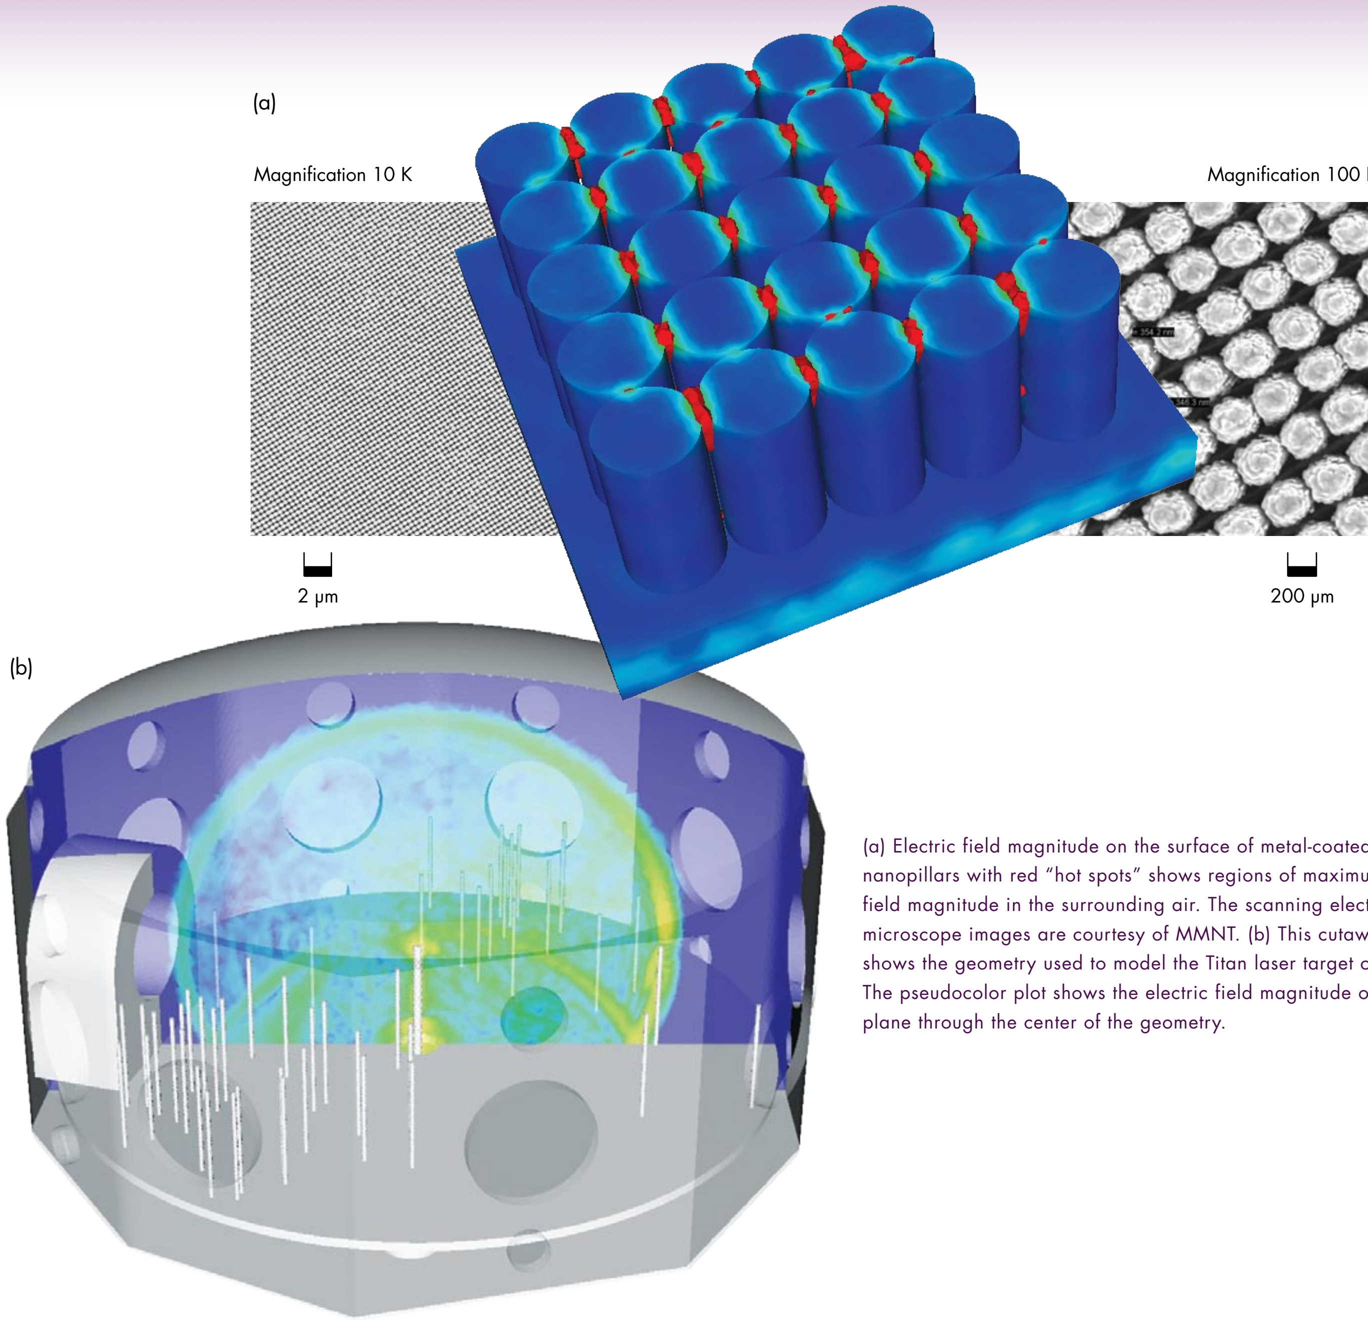

(a) Electric field magnitude on the surface of metal-coated, silicon nanopillars with red "hot spots" shows regions of maximum field magnitude in the surrounding air. The scanning electron microscope images are courtesy of MMNT. (b) This cutaway view shows the geometry used to model the Titan laser target chamber. The pseudocolor plot shows the electric field magnitude on a plane through the center of the geometry.

3.03

Computation Helps Russia Secure Nuclear Materials

SIGNIFICANCE

One of the gravest threats that resulted from the dissolution of the Soviet Union was the proliferation of nuclear materials. The security surrounding these materials diminished as the Soviet Union split into individual states in 1991. Five years later, DOE established a cooperative agreement with Russia to secure nuclear materials from theft and diversion, reduce stockpiles of attractive nuclear material, and assist the Russian government in retaining scientists trained to use and manage the material.

Computation employees working in the Global Security Principal Directorate provide project management and technical expertise to a multilaboratory team tasked with helping the Russia Federation improve their security posture. With guidance from DOE, Department of State, and the Defense Threat Reduction Agency, the team implemented a nuclear material accounting system and nuclear material treaty verification instruments with integrated software, which protect classified details of the material being verified and provide a high degree of assurance that the type and quantity of material is as it is claimed.

MCANINCH

WHITE

CONTACT INFORMATION

CONSTANCE MCANINCH, MCANINCH2@LLNL.GOV

GREGORY WHITE, WHITE6@LLNL.GOV

PROGRESS IN 2008

In 1996, LLNL began leading a multilaboratory project to develop, implement, and operate the Russian Federal Information System (FIS), a national-level automated nuclear material accounting system. This system, in operation since 2002, processes nuclear material inventory reports from all civilian nuclear sites in Russia. In the six years since its inception, FIS has undergone several phases of enhancements, the most recent of which was completed in 2008. This development phase, the last to be supported with funding from the U.S., enables the system to collect inventory information from the nuclear facilities in enough detail to qualify accurate nuclear material management at the national level, while imposing minimal additional burden on facilities.

To support a decade-long series of nonproliferation regimes between the U.S. government, Russian Federation, and International Atomic Energy Agency, LLNL has made significant contributions to the design and development of nuclear material monitoring systems that verify the contents of containers carrying special nuclear materials, such as plutonium. This type of system is certified by the "host party" (the owner of the material) to protect the classified information detectable by the nuclear monitoring system (e.g., detailed gamma spectra and neutron energy that reveal isotopic composition and quantity) from inadvertent disclosure to the "monitoring party" (the inspectors). The monitoring party must authenticate that the system correctly calculates the quantity and type of nuclear material.

Computation staff created software design requirements for C and C++ to address these conflicting needs. The team also developed tools, based on LLNL's ROSE software suite, to authenticate the software developed for nuclear material monitoring systems, enforce adherence to the design requirements, and detect vulnerable code that might produce incorrect results. In addition, the team increased the capabilities of the premiere LLNL gamma isotopics code (Pu600) for these nonproliferation regimes by translating the code from Fortran to C (while maintaining the code's integrity) and porting it to an open-source, embedded operating system (eCos).

Computation employees have participated in nuclear nonproliferation projects at civilian and military nuclear sites throughout the Russian Federation.

3.04

Biodefense Knowledge Management System Addresses Knowledge Discovery Challenges

SIGNIFICANCE

The Department of Homeland Security Biodefense Knowledge Center (BKC) is developing a knowledge discovery system, called the Biodefense Knowledge Management System (BKMS), to help the homeland security community analyze and characterize biological threats posed by terrorists.

The seemingly simple task of searching several data sources for warning indicators, people, and places to produce actionable information belies a daunting set of challenges facing computer scientists and the homeland security community. To avoid the shortcomings of past attempts to solve these challenges, a team of computer scientists in BKC shifted the paradigm from helping the analyst find a needle in the haystack to systematically reducing the haystack(s) for clarity and focus. In 2008, Computation staff helped BKC deliver the first production release of BKMS on both unclassified and classified networks. The release included an integrated core search engine, an advanced search interface, and a large-scale data processing pipeline.

WEISS

CONTACT INFORMATION

ERIC WEISS, WEISS7@LLNL.GOV
WEB SITE: BKC.LLNL.GOV

PROGRESS IN 2008

The challenges facing BKC include entity extraction (i.e., who, what, and where), named-entity recognition (i.e., does “Georgia” refer to the U.S. state or to the Republic of Georgia?), large-scale data processing, rapid query response, and data fusion. In addition, BKC faces more mundane, understated challenges, such as building a production quality system, deploying and maintaining installations on multiple classified and unclassified networks, and maintaining high-quality and timely data.

Although the long-term vision for BKC is to address all of the looming challenges, the Computation team focused on building and deploying the foundational elements in 2008. The team’s first step was to identify and integrate a full-text search engine. The commercial, off-the-shelf package from Endeca Technologies, Inc., provides BKMS with an immediate full-text indexing capability for the vast amounts of structured and unstructured textual data employed by analysts. Like most modern search engines, Endeca uses an index (similar to an index in a book) at its core to provide fast data ingest and querying capability for large data sets. The ingest of 20 million documents now takes four days rather than four weeks (research indicates this time could yet be improved to less than 10 hours), and query times have decreased from minutes or hours to milliseconds. These radical improvements allowed the team to widen the search net and implement more complex processing algorithms.

The BKC team also implemented an open-source Hadoop/MapReduce cluster used for large-scale, distributed data processing. The MapReduce framework divides jobs into many small blocks of work, places them on compute nodes around the cluster, and then processes the data where it is located, making it massively scalable. Hadoop is a vast improvement over previous BKC capabilities; it is instrumental for all manner of data processing, including entity extraction, disambiguation, clustering, data fusion, and normalization.

The initial release of BKMS was deployed to classified and unclassified networks, and users from several U.S. government agencies and partners in Australia are accessing the system.

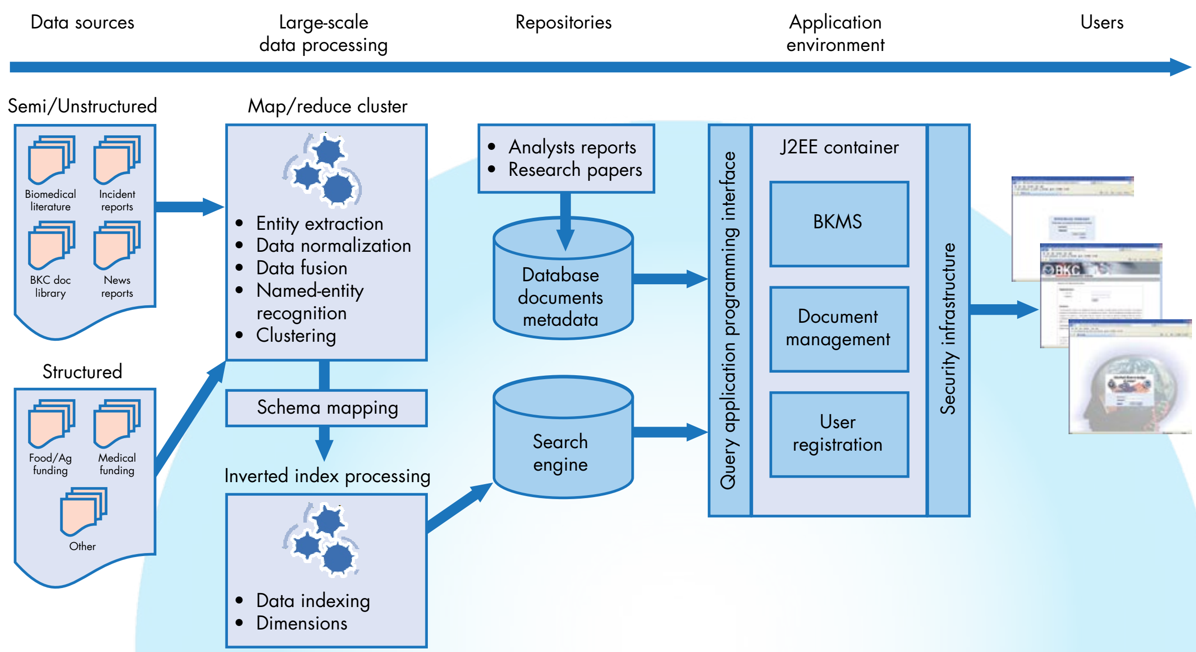

BKMS System Overview. This system extracts entities (e.g., who, what, and when) and relationships from structured and unstructured textual data sources, indexes the entire corpus and provides analysts with an advanced search capability to “discover” knowledge and information.

3.05

VisIt Project Partners on Visualization and Analysis of Massive Data

SIGNIFICANCE

VisIt is an open-source tool for visualizing and analyzing extremely large-scale data. Originally developed by the ASC Program, VisIt is now supported by three DOE computing programs—ASC, Scientific Discovery through Advanced Computing (SciDAC), and Advanced Fuel Cycle Initiative (AFCI). VisIt is used by hundreds of analysts at LLNL and throughout DOE, the Department of Defense, and academia.

VisIt has a rich feature set for visualizing and analyzing scientific data and a flexible data model that accommodates a wide range of simulation data, including particle data, structured adaptive meshes, and unstructured meshes. Because of its distributed architecture and fully scalable, parallel, end-to-end data processing pipeline, VisIt is able to process very large data sets without moving data. In 2008, the VisIt team enhanced the tool with several significant capabilities, completed its transformation to a distributed, open-source development effort, and positioned the code base to solve upcoming data analysis and visualization challenges.

BRUGGER

CHILDS

CONTACT INFORMATION

ERIC BRUGGER, BRUGGER1@LLNL.GOV
HANK CHILDS, CHILDS3@LLNL.GOV
WEB SITE: VISIT.LLNL.GOV

"By leveraging our efforts on a common tool, we are collectively able to deliver a superior product at a reduced cost, better than any of us could do individually."

Jim Rathkopf, ASC Integrated Codes Program Leader

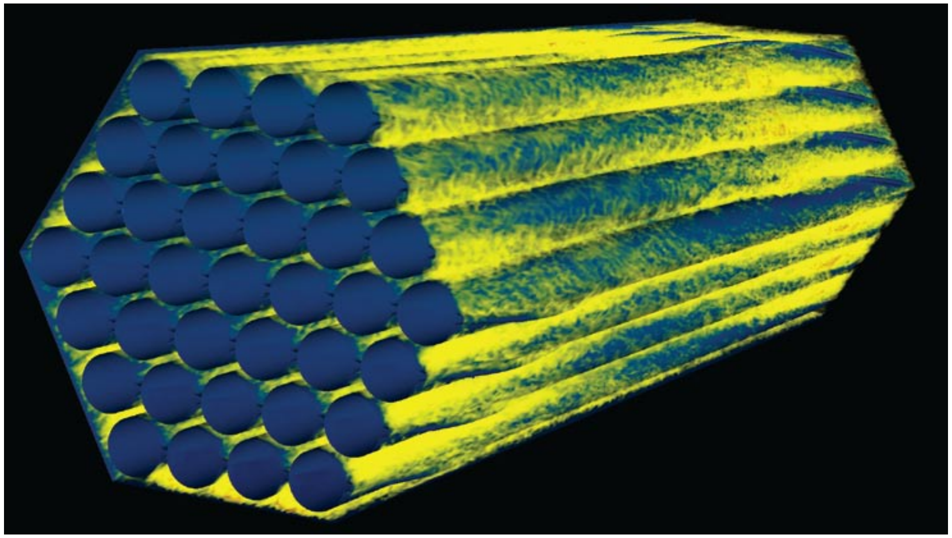

A volume rendering of a reactor subassembly with 37 wire-wrapped fuel pins. The simulation was produced using the NEK code by ANL's Paul Fischer on the BlueGene/P machine. The image was produced by LLNL's Hank Childs, in support of the AFCI Program.

PROGRESS IN 2008

Several state-of-the-art third-party visualization capabilities were integrated into VisIt in 2008, including the Image Composition Engine for Tiles (IceT), developed by Sandia National Laboratories; the SCIRun Library for Interactive Volume Rendering (SLIVR), developed at the University of Utah; and FastBit, a 2008 R&D 100 Award-winning technology developed by Lawrence Berkeley National Laboratory (LBNL), which implements fast query operations using compressed bitmap indexes.

VisIt's support for structured adaptive meshes from Adaptive Mesh Refinement (AMR) simulations was improved by incorporating curvilinear adaptive meshes and new AMR-specific interaction modalities. These enhancements, combined with significant performance tuning, convinced the SciDAC team that develops the Chombo AMR infrastructure to stop development on their dedicated visualization tool, ChomboVis, in favor of VisIt.

The VisIt team embarked on several research efforts this year, one of which culminated in the addition of a new streamline algorithm. Streamlines are commonly used for visualizing vector

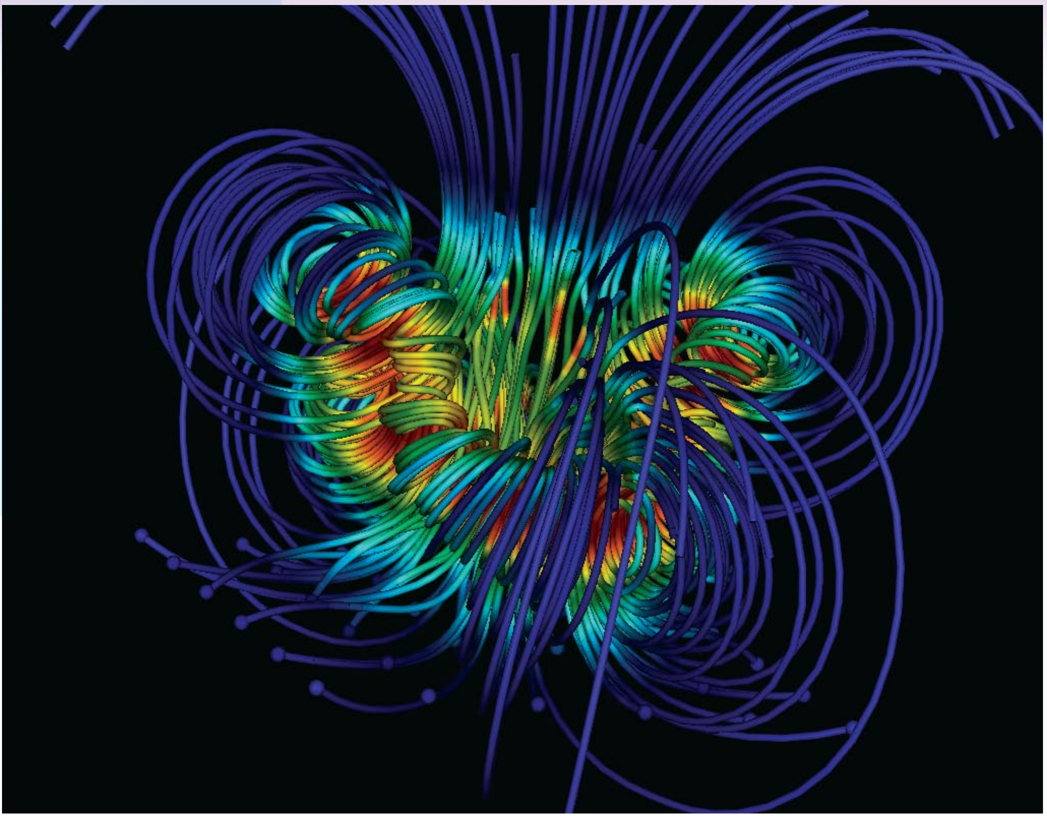

This streamline plot of a vector field of a converging vortex core solves the Incompressible Navier–Stokes Equations in three dimensions. The simulation was produced using the Chombo AMR code from LBNL and rendered using VisIt. The image was produced by ORNL's Dave Pugmire.

fields but are difficult to implement in a distributed parallel memory setting because the streamline path can pass from processor to processor and necessitate complex and inefficient communication schemes. This problem was solved by a collaboration of researchers from UC Davis, Oak Ridge National Laboratory (ORNL), and LLNL, who implemented several schemes inside of VisIt that adapt to the properties of the input data set.

VisIt adopted a hybrid project management approach this year by combining conventional- and community-based support models. Supported customers obtain assistance via telephone hotlines or dedicated e-mail lists. Unsupported customers seek help through a "vistor" e-mail list, which reaches more than 200 subscribers and receives 200–400 e-mails per month, or through a user's Wiki forum at www.visitusers.org, which has almost 200 pages of content. The VisIt team transitioned development of VisIt from a proprietary version control system only accessible to LLNL developers to a Subversion repository located at the National Energy Research Scientific Computing Center, which is accessible to all collaborators. Twenty-two developers from nine institutions have permission to write to the VisIt repository, including staff at LBNL, ORNL, Argonne National Laboratory (ANL), the Atomic Weapons Establishment, Commissariat à l'Énergie Atomique (CEA), and the National Center for Supercomputing Applications.

The next challenge facing the VisIt team is to keep pace with the increasing size of data sets as industry moves toward exascale supercomputing. Supercomputing trends show input and output (I/O) rates getting slower relative to compute power, while visualization tools already spend the large majority of time reading data off disk. Thus, the current paradigm for production visualization tools, where the full resolution of the data is read off disk and into memory, will likely change. The VisIt team is pursuing multiple strategies that deemphasize I/O and memory footprint, including in situ processing (where the visualization and analysis are done in place with the simulation), multiresolution processing (adaptively coarsening the data to maintain interactivity), and streaming (where pieces are processed one at a time). None of these techniques is a panacea, but together they form a viable strategy for meeting the challenge without requiring large computing resources. LLNL is partnering with CEA to develop multiresolution techniques that will allow for quick browsing of data in a way that maintains data integrity. An LLNL Laboratory Directed Research and Development project has already contributed innovative techniques for streaming, in addition to a module that was deployed in VisIt.

3.06

Software Quality Transforms into a Valued Business Asset

SIGNIFICANCE

The basic principle of assuring software quality is simple: the higher the risk associated with the software, the greater the care that must be taken. The Institutional Software Quality Assurance (ISQA) office provides guidance on the quality of software used at the Laboratory, helps identify software-related risks, and characterizes the risks for likelihood and consequences of failure. ISQA's objective is to create a high degree of confidence in the performance, cost, and schedule of the Laboratory's software products. This year, ISQA was tasked by the National Nuclear Security Administration Livermore Site Office (LSO) to implement the SQA Program in all nuclear and radiological facilities by September 30, 2008. ISQA completed this task and several other activities, all of which helped SQA become a valued business asset at LLNL rather than a mere compliance activity.

PROGRESS IN 2008

ISQA developed a formal interpretation of what constitutes "safety software" at LLNL and used it to review software from every organization at the Laboratory. The review identified approximately 60 candidate safety software applications. After further evaluation and consultations with experts in each organization, 11 pieces of software were added to the official list of LLNL safety software. Each piece was then risk-graded using an ISQA-developed Web-based tool.

LSO's SQA Program Manager Adeliza Cordis, NMTP Quality Assurance Manager Mark Accatino, and LLNL SQA Manager Darrel Whitney review SQA plans and implementation documents.

WHITNEY

POPE

CONTACT INFORMATION

DARREL WHITNEY, WHITNEY1@LLNL.GOV

VICKI POPE, POPE13@LLNL.GOV

> *"The efforts of the ISQA office were critical to NMTP's ability to navigate the process of LSO's assessment of SQA implementation in the Superblock."*
>
> Mark Accatino, Quality Assurance Manager, Nuclear Materials Technology Program (NMTP)

To further improve the posture of the Laboratory's safety software, ISQA enhanced their best-practices documentation and provided IEEE-based templates and checklists to project teams, so they could perform SQA self-assessments and gap analyses and quickly implement missing documents. SQA experts helped personnel from Chemistry, Materials, Earth, and Life Sciences (CMELS); NIF; Radioactive and Hazardous Waste Management; and the Laboratory's defense plutonium research and development facility (the Superblock) close out all SQA issues that had been reported and tracked by LSO in FY08.

Also in 2008, ISQA made additional SQA processes and tools available to LLNL users through the ISQA Web site and the Software Improvement Networking Group Wiki. More than 200 pages and 100 downloadable files were published to these sites, including descriptions of ISQA's recommended practices and guidance on tailoring them to specific projects. Upon recommendation from the tri-annual ASC SQA assessment, ISQA also added best-practices information garnered from previous work in the ASC Verification and Validation Program.

ISQA assisted many LLNL programs either by facilitating or performing SQA assessments. Examples of their effort include:

- Leading the tri-annual ASC SQA assessment, which included teams from Lawrence Livermore, Los Alamos, and Sandia national laboratories.
- Performing an SQA audit for the Electrostatic Discharge Project at the request of the project leader, in preparation for their next software release.
- Leading the LSO assessment of CMELS's Radioactivity Allowance Tracking System safety software.
- Facilitating the LSO SQA assessment of the Superblock.

INFORMATION TECHNOLOGY AND SECURITY

4

4.00 Information Technology Evolution Helps Modernize and Protect the Institution

In 2004, the Computation Directorate published a strategy for advancing information technology (IT) services at Lawrence Livermore National Laboratory (LLNL). This strategy put forth as the top priority the implementation of several best practices aimed at modernizing the Laboratory's IT environment. Interestingly, our strategy has seen only modest changes since its publication four years ago, and the main goal of the strategy has remained mostly unchanged: we seek to serve the Laboratory with a modern, efficient IT environment—one that is viewed by management and employees as an enabler of science, a partner in helping LLNL meet its mission and business commitments, and an asset that gives the Laboratory a competitive advantage.

We identified two foundational best practices in our strategy: centralized desktop management and centralized network management. Although it is clear that Laboratory employees (our customers) are in need of some new services, such as Blackberry support and desktop back-up features, we recognize that we cannot meet those needs until we modernize the Laboratory's network and desktop environment. A prerequisite for modernization is the adoption of a new approach to managing the networks and desktops—one that will endow all organizations at the entire Laboratory with a reliable and efficient IT infrastructure. Historically, each program and department had the power to design its piece of the Laboratory's IT infrastructure and pick its

NEXT PAGE


QUINN

CONTACT INFORMATION

TERRI QUINN,
QUINN1@LLNL.GOV

favorite technology to deploy—a classic example of local optimization at the expense of the common good. When IT was in its infancy, this was tolerable, and even advantageous, as each organization got what it needed and wanted without affecting the whole. However, IT has matured, and the advantages of this practice no longer hold true. The assumption today in industry and in security circles is that an organization (such as LLNL) has a vital need for a common underlying architecture and common technologies. Without control over LLNL's business network and desktops, the Laboratory would never acheive a common underlying architecture for IT.

During the last few years, the Computation Directorate made some progress toward realizing our vision of a centrally managed network and desktop environment at LLNL, but it was not nearly as much progress as we had hoped for. Fortunately, events in late 2007 and 2008 created an opening that accelerated our progress. Specifically, supportive customers and a heightened concern for computer security created this opportunity, which we were well-positioned to exploit with ready-to-implement plans and an eagerness to put the plans into practice. One of LLNL's largest organizations, the Operations and Business (O&B) Principal Directorate, approached Computation in the fall of 2007, enthusiastic to partner with us to change their desktop and network service model. Their request aligned exactly with our strategy—not surprising, as the management of that organization is keenly interested in improving the efficiency of its operation. We were delighted to work with them, and in doing so, put into practice our centralized desktop and network models. Following in the footsteps of the O&B organization, the Science and Technology (S&T) Principal Directorate heard about the work we were doing for O&B and asked us to do the same for them. In this section, we report on our progress in serving both these customers.

The second event that helped create an opportunity to advance our IT vision occurred in 2008. At this time, LLNL senior management decided that improving the Laboratory's computer security posture needed to become a priority and, in particular, requested that we tighten the security of the unclassified network and improve the protection of unclassified desktops. Again, we were prepared to meet their request. With our pre-existing desktop management capability and significant time and dedication from our staff, we were able to implement a new desktop patching process capable of patching the Laboratory's desktops in a very short timeframe. This accomplishment and others that resulted from this effort had the dual benefit of not only improving security, but also furthering our movement toward a centrally managed network and desktop environment.

4.01

Consolidation Establishes Foundation to Deliver Next-Generation Network Capabilities

SIGNIFICANCE

Historically, LLNL unclassified networks were funded and managed in a decentralized manner, where each program provided the network hardware and wiring in their buildings. The independent networks, in turn, connected to each other and the Internet through an institutionally funded and managed network backbone. This model has led to widely disparate investments and heterogeneity in network hardware and configuration, as well as levels of staffing and expertise. Approximately 70% of LLNL building network equipment has reached its end of life, which increases the risk of failure and security breaches. Additionally, older electronics and nonstandard network designs have created a barrier to delivering next-generation network-based services to LLNL users. Network consolidation will position LLNL to deliver enterprise-wide network capabilities both securely and cost effectively.

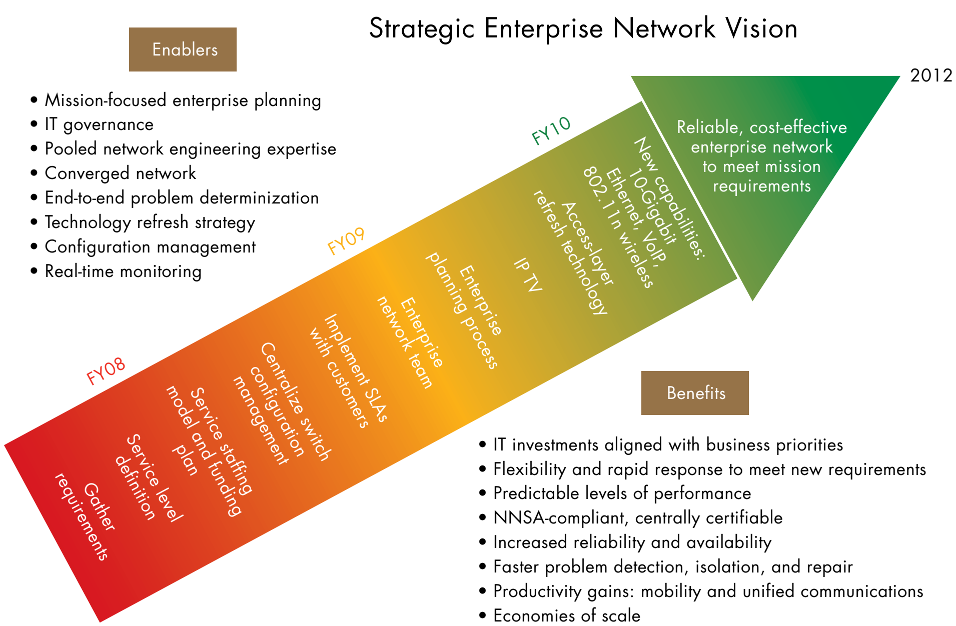

The Chief Information Officer's FY09 planning process established a strategic vision for a consolidated LLNL enterprise network.

GOLDSTONE

CONTACT INFORMATION

ROBIN GOLDSTONE,
GOLDSTONE1@LLNL.GOV
WEB SITE: BNT.LLNL.GOV

PROGRESS IN 2008

Network consolidation is part of a broader IT transformation that is occurring at LLNL. As the Laboratory migrates toward centralized IT, it is necessary to define a framework upon which IT can be delivered as a service, with well-defined expectations for performance, availability, cost, support, and response, and with metrics to ensure that expectations are being met. The IT Solutions Division adopted the widely used Information Technology Infrastructure Library (ITIL) framework for IT service management. During 2008, several IT Solutions Division staff members obtained ITIL Practitioner certification and used their knowledge and training to develop key service management components, including an IT service catalog, Service Level Agreement (SLA) templates, and an ITIL-compliant change management process. With this framework established, ITSD negotiated SLAs with the Global Security, S&T, and O&B principal directorates, which allowed the Enterprise Network Solutions (ENS) team to begin network consolidation.

Consolidation accomplishments in 2008 included documenting the physical and logical topology of customers' access-layer networks, integrating switches and routers into ENS's network management and monitoring tools, and ensuring devices were configured according to industry best practices and LLNL standards. When a customer's networks are well-documented and correctly configured, ENS can assess the state of the networks with regard to reliability, security, and ability to support new capabilities, such as Internet Protocol (IP) video and Voice over IP (VoIP).

Users in programs with established SLAs now obtain support from the Institutional Service Desk (4HELP) rather than from ad hoc sources. The Service Desk provides accurate issue characterization and prioritization, Tier-1 troubleshooting, incident management from "cradle to grave," and gathering and reporting of metrics to validate the effectiveness of the service. Additionally, ENS designed and developed a new software tool to give Tier-2 field support personnel the ability to perform common office configuration and troubleshooting activities on access-layer switches in a safe and auditable manner.

In the spring of 2008, LLNL was victimized by a significant cyber attack. Nearly every IT and Computer Security staff member was involved in a highly reactive and costly incident response. In the future, network consolidation will provide the integrated framework required to deploy network admission control (NAC) solutions, including preadmission health check, quarantine, and remediation. NAC technology, in concert with enhanced desktop security tools, will dramatically improve LLNL's cyber-security posture and reduce costs associated with incident response and compliance.

IT Support Centralization Lowers Costs and Improves Service

SIGNIFICANCE

The Computation Directorate's Computer System Support and IT Solutions Divisions are responsible for providing LLNL users with IT support. Historically, Computation's IT support model consisted primarily of providing skilled personnel to each directorate, with each directorate then utilizing these people at their discretion to provide computer support. Based on industry benchmarking and best practices, Computation developed a new IT support model and implemented it in three Principal Directorates (PDs) at the Laboratory—the Director's Office PD, O&B PD, and S&T PD. These three PDs comprise approximately 75% of LLNL staff. The new IT support model employs common methods and technologies, but it also has enough flexibility to meet the unique support needs of diverse customers.

The IT support consolidation in the three PDs has produced positive results that clearly demonstrate the benefits of a more centralized IT support model for LLNL. The changes have increased the effectiveness of field support staff, lowered customer support costs, and increased the overall security posture of the network through improved configuration management.

TESLICH

CORNELL

CONTACT INFORMATION

ROBYNE TESLICH, TESLICH2@LLNL.GOV
BRETT CORNELL, CORNELL5@LLNL.GOV

PROGRESS IN 2008

Under the new support model, each PD better leverages LANDesk for systems management tasks. Rather than assigning field technicians to perform routine desktop management tasks, LANDesk is being used more frequently for systems management tasks such as patching, software distribution, and system certification and accreditation. LANDesk increases both the efficiency of technicians and the configuration consistency of the desktops and servers on the network, thereby improving security.

Computation also implemented a common enterprise call-tracking system and directed all support requests from the PDs to the Institutional Service Desk (4HELP). The Service Desk provides Tier-1 support for all IT-related calls and resolves 45% to 60% of service requests without dispatching a field technician to a user's office. The Service Desk helps reduce support costs and provides users with a faster resolution to their problems. In the future, the remote administration capabilities in LANDesk will be expanded, which should enable the Service Desk to handle approximately 70% of all support requests.

Computation helped the three PDs transfer responsibility for unclassified network management to the ENS team (see article 4.01). Centralizing network management ensures that the PDs' unclassified computer networks are managed in a consistent and efficient manner by computer networking professionals.

In each of the three PDs, Computation reorganized its computer support staff into a single computer support unit (CSU) dedicated to serving the entire PD rather than to individual organizations within the PD. This support model provides greater flexibility and better utilizes the IT staff. In the previous model, if one organization had a heavy call volume, there was no way to leverage resources to address the spike in support needs. In addition, resources can now be directed to handle the highest priority issues within the PD at any given time.

The Director's Office CSU team serves a diverse set of users, including LLNL senior management, doctors, lawyers, accountants, cyber security experts, and scientists. The team implemented the four initiatives outlined above and adopted a common set of metrics to measure computer support performance across the PD. The various IT managers within the Director's Office review these metrics monthly to ensure that each organization is being properly served. The new streamlined processes permitted the Director's Office to reduce its number of support staff by nearly half, resulting in the largest relative cost savings among the three PDs that use the new support model.

> *"O&B is already realizing significant value from consolidating IT support services within the PD. In particular, the reduction in deployment time for critical patches from months to days and the implementation of tools to allow tracking these deployments have been critical to achieving and sustaining a heightened cyber security posture during the past several months."*
>
> Frank Russo,
> O&B Principal Associate Director

NEXT PAGE

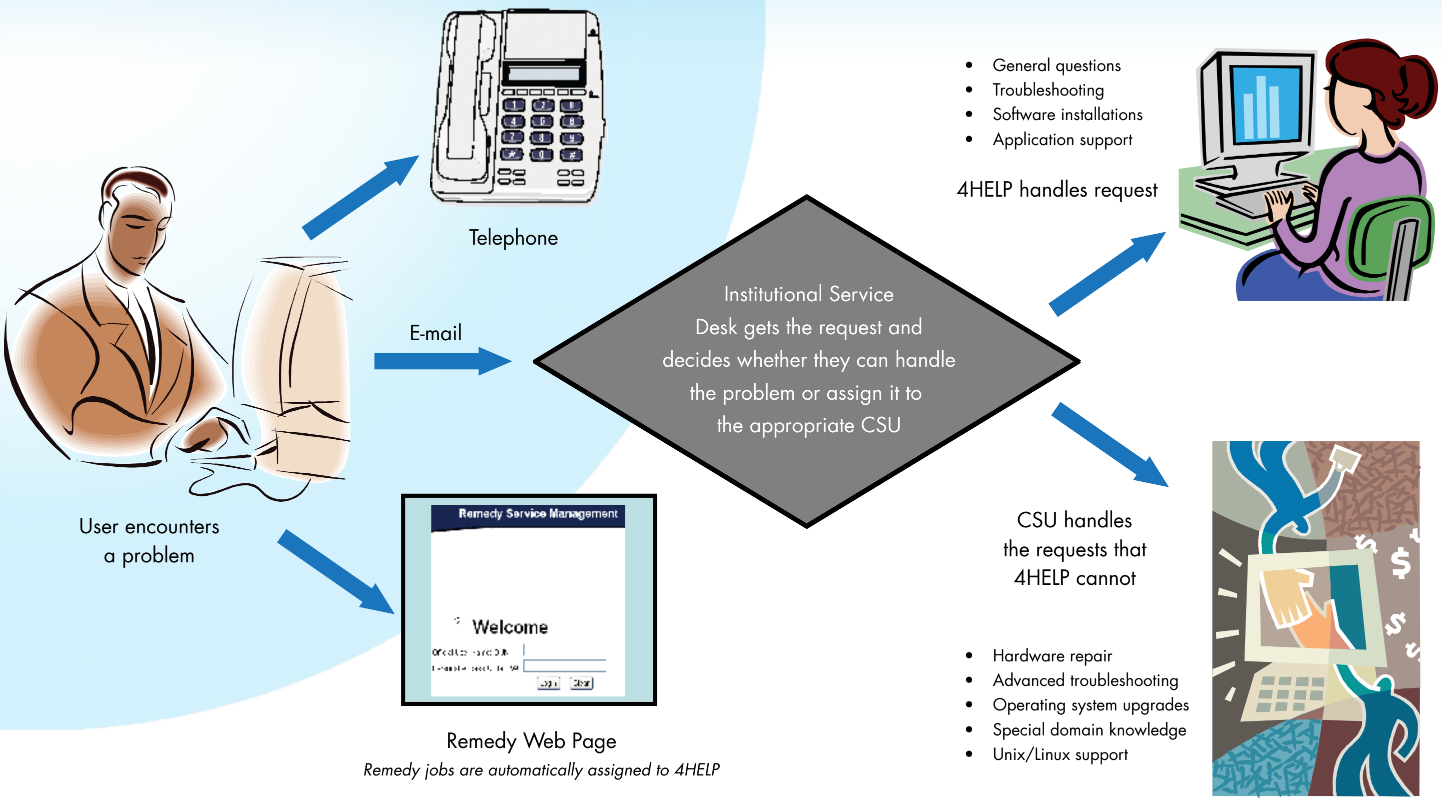

The new IT model calls for filtering all user requests through a centralized help desk.

The Director's Office PD also consolidated its infrastructure and reduced the total number of servers by approximately 40%.

The O&B CSU, which provides computer support to organizations responsible for maintaining LLNL's business operations, implemented the aforementioned initiatives and introduced greater rigor and formalism in defining IT support by implementing ITIL-based processes within the PD. The O&B CSU worked closely with Computation and O&B PD managers to negotiate a formal Service Level Agreement that documents the specific IT services the PD will receive and the metrics that will be used to monitor and measure the success of service delivery. In addition, the O&B CSU centralized configuration management in a group outside of the CSU. IT Solutions Division staff now provide management of all patching, Active Directory, antivirus, and certification and accreditation for the general-purpose O&B PD desktops. The new support model allowed the O&B CSU to reduce staff by approximately 20%, while completing more customer requests in less time.

The S&T CSU supports a PD that is home to nearly half of the users at LLNL, making it the largest and most diverse CSU at the Laboratory. The S&T CSU reorganized its support personnel into three teams—an operations team, a projects team, and a security team. The operations team handles day-to-day support and manages the various organizations' large infrastructure of servers and other specialized devices. The projects team implements specialized IT initiatives across the organizations and provides advanced technical support as needed. The security team addresses cyber security and compliance issues across the PD. The S&T CSU consolidation will continue in 2009, and the results are expected to yield significant cost savings.

Major Cyber-Security Incident Drives Transformational Improvements in Security-Related System Management

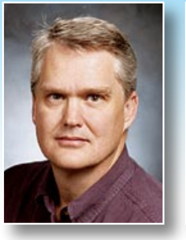

EAST

MATSCHE

CONTACT INFORMATION

DOUG EAST, DRE@LLNL.GOV
EDWARD MATSCHE, MATSCHE1@LLNL.GOV

SIGNIFICANCE

In the spring of 2008, LLNL came under cyber attack from a very sophisticated adversary. LLNL's response and recovery actions implemented a transformational change in how LLNL secures and updates its Internet-connected systems and networks. The key challenge of the recovery effort was to dramatically increase security while maintaining a highly functional environment for mission-driven research and collaboration. By adding security improvements incrementally each month, consistent progress was made in increasing LLNL's cyber-security posture while attempting to not overwhelm the resources and undermine users' ability to work. This approach led to significantly more centralized coordination and execution of the improvements and removed historical barriers to centrally driven changes in LLNL's IT environment.

At the same time, a clear path forward to implementing NAPs (the National Nuclear Security Administration's [NNSA's] new National Institute of Standards and Technology-based computer security policies) emerged, as only through centralized, automated means could the requirements be met in a cost-effective manner. In addition, a 2008 audit by the Department of Energy's (DOE's) Office of Health, Safety, and Security identified LLNL's "one security plan" approach on the unclassified network as one of its key successes. In this method, all IT services are documented at the institutional level. This security plan architecture enabled an institutional approach to security-related automation that was implemented later in the year.

PROGRESS IN 2008

Among the cyber attacks LLNL has encountered over the years, none penetrated more broadly or deeply as the one that began in late spring 2008. The attack exploited a well-known application software vulnerability triggered by users clicking on a link in a phishing e-mail. The attack was successful because the adversaries used targeted phishing to lead LLNL users to believe they received a business-related message from a sender with whom they were familiar. Once launched, the malware "phoned home" to establish a control channel, targeted both cached and typed passwords, and rapidly and intelligently attacked sensitive parts of the cyber infrastructure, including manager's systems, major databases, and authentication servers.

An initial cyber-response effort, led by the Laboratory's Safeguards and Security Organization, applied traditional emergency management principles to respond to the attacks. This response effort focused on containing the attack: disrupting attackers' control channels, closing or tightly monitoring data-exfiltration paths, and isolating and rebuilding potentially compromised systems. Soon after the attack, LLNL established new intersite and interagency relationships to share information about cyber attacks in real-time, including intrusion/detection monitoring aggregation through the DOE Cyber Incident Response Capability. Computation's Information Protection Support Organization incident management team found fragments of the malicious code during forensic analysis and shared this information with members of LLNL's Information Operations and Analytics Center. The latter was able to apply cutting-edge research in the area of malicious network traffic detection to the programmatic research effort, using real-world sample data from LLNL's cyber incident. This collaboration produced valuable research results and assisted in the detection aspect of the cyber response and recovery, positioning LLNL as a leader in automated malicious network traffic detection and analysis.

The longer-term cyber-recovery effort, led by Computation's Doug East, will continue in 2009 and is focused on post-response cleanup, including improving detection and containment, as well as

NEXT PAGE

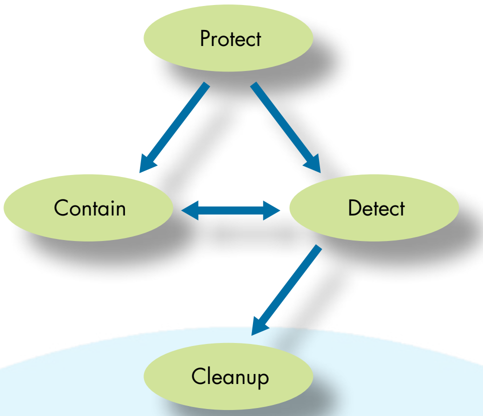

The cyber response and recovery model shows how LLNL protects via patching, quarantine, and security certification; detects via intrusion detection, known-bad-site monitoring, and increased intersite/interagency information sharing; contains via system blocking and isolation, remediation, and autobeaconing disruption; and cleans up through system rebuilds, forensic analysis, and password changes.

Significant and sustained changes in LLNL's cyber environment in 2008

- Allowed only Web-based ports Internet access by default, and required authentication before granting access.
- Enforced recent, known-secure versions of remote access clients.
- Required two-factor authentication for privileged users. Applied this technology to the most privileged users first with the ultimate goal of applying it to all users.
- Rebuilt all Active Directory domain controllers, and required all users to change passwords.
- Reviewed Active Directory and Domain Name System with Microsoft, and produced a set of recommendations to work through, although generally LLNL's implementation was considered good. Active Directory now reliably uses Kerberos rather than the less secure Windows-based protocols.
- Mangled links in e-mails that originate off site to ensure users take manual steps to fix the actual link string, rather than "just clicking" it.
- Blocked out-of-date/vulnerable Web browers and Web components to limit desktop/laptop exposure to Internet-based threats.
- Deployed tool (SLAM) to centrally manage one-time passwords for local account(s) on desktops/laptops, and ensured that authenticators are managed/controlled at the same level as institutional authenticators.

implementing new preventative security countermeasures to preclude a similar recurrence. A major effort involved bringing all Windows and Mac operating systems under LANDesk management and patching all systems with an institutionally defined patch-set according to an institutionally defined schedule. Users are now asked to leave their computers powered on and connected to the network permanently to enable rapid, automated system patching. Other preventative and containment-oriented measures that were implemented include Web authentication and Web application blocking. Web authentication requires users to authenticate using their "official user name" and "personal access code" when first accessing a non-Laboratory network. Authenticating opens a time window during which subsequent accesses do not require authentication. Web application blocking is aimed at ensuring that only current (and likely least vulnerable) versions of Web-oriented desktop applications (e.g., browsers) are able to access non-Laboratory networks.

Because of the success of the cyber-recovery effort, a new model of security compliance assurance has emerged. Computation's Computer Systems Support and IT Solutions Divisions are responsible for meeting cyber-security requirements using their existing system management tools (e.g., Active Directory and LANDesk); these tools will implement cyber policy as well as measure and control compliance of systems, including driving the quarantine/blocking actions for noncompliant systems. The Cyber Security Program then uses complementary tools, separate from those implementing cyber requirements, to audit and independently verify compliance. This model allows Cyber Security Program personnel to provide quality assurance information that Computer Systems Support and IT Solutions personnel can use to improve their ongoing implementation plans.

4.04

Facility Team Creates Energy Efficiencies and Turns Megawatts into PetaFLOPS

SIGNIFICANCE

As a major contractor to NNSA, LLNL is required by DOE Order 430.2B to have a documented energy management program and plan to reduce energy intensity (energy consumption per unit area) by at least 30% by 2015 relative to the energy intensity rates of 2003. LLNL is committed to meeting this challenge through innovative design, re-engineered maintenance and operations practices, enhanced performance requirements for operational equipment and systems, and alternative financing. Computation made significant energy-saving improvements within the Terascale Simulation Facility (TSF) and its other buildings (B451, B439, B115, and B117) to help LLNL exceed its FY08 incremental target of 9% reduction relative to 2003. Computation is also motivated to improve infrastructure efficiencies to ensure that ample power is available for the anticipated computational load growth.

BAILEY

CONTACT INFORMATION

ANNA MARIA BAILEY,
BAILEY31@LLNL.GOV

PROGRESS IN 2008

In conjunction with the DOE Industrial Technologies Program and Lawrence Berkeley National Laboratory, TSF facility staff used a benchmarking tool to identify prospective energy-saving initiatives within the computer rooms in TSF. The facility staff ran a comprehensive computational fluid dynamic (CFD) model to analyze airflow patterns and identify leaks. After considering the results of the CFD model, both the supply-air temperature and chilled-water-supply temperature in the large computer rooms were increased by 7°F. These temperature changes will save nearly 26,000 megawatt hours (MWh) annually, and no measurable increase in component failure rates has been detected.

> *"LLNL's energy-efficiency efforts in their high-performance computing facilities are setting the pace for DOE/NNSA in terms of implementing one of the pillars of the President's New Energy for America Plan: Deploy the cheapest, cleanest, fastest, energy sources."*
>
> Karin King, Livermore Site Office, Leadership in Energy and Environmental Design Accredited Professional

The CFD modeling also identified approximately 450,000 cubic feet per minute (CFM) of air leaking from the closed louvers on the perforated and grated tiles in the computer rooms. The tiles were leaking at a rate of 10% while in the closed position, which was much higher than the manufacturer's published data. Consequently, facility staff replaced approximately 50% of the perforated and grated tiles with solid tiles, which allowed two 80,000CFM air handlers to be taken off line and saved an additional 1,000MWh.

The TSF is a fairly energy-efficient building, especially considering the massive architectural, electrical and mechanical systems it contains. Power usage effectiveness (PUE) is one industry-standard measure that is often used to quantify and compare the energy efficiency of data centers. PUE is a ratio expressed as: PUE=(total computational plus mechanical facility power) divided by total computational power. PUE can range from 1.0 to infinity, with the ideal value of 1.0 indicating 100% efficiency. The TSF's PUE is 1.34, based on the original facility design plus the improvements implemented in 2008. In comparison, industrywide studies to date indicate that most data centers have PUE values ranging from 1.3 to 3.0.

In FY09, facility staff will perform additional audits and CFD modeling to consider the effects of further increasing the air and water supply temperatures in all Computation buildings. The current models indicate that, given the 30˚F to 50˚F temperature rise from the supply temperature to the hottest part of a computer rack, the supply-air temperatures could be increased another 6˚F and still keep the average room-return temperature well within the desired operating parameters. These results will be evaluated in relation to the potential increase in component failure rates.

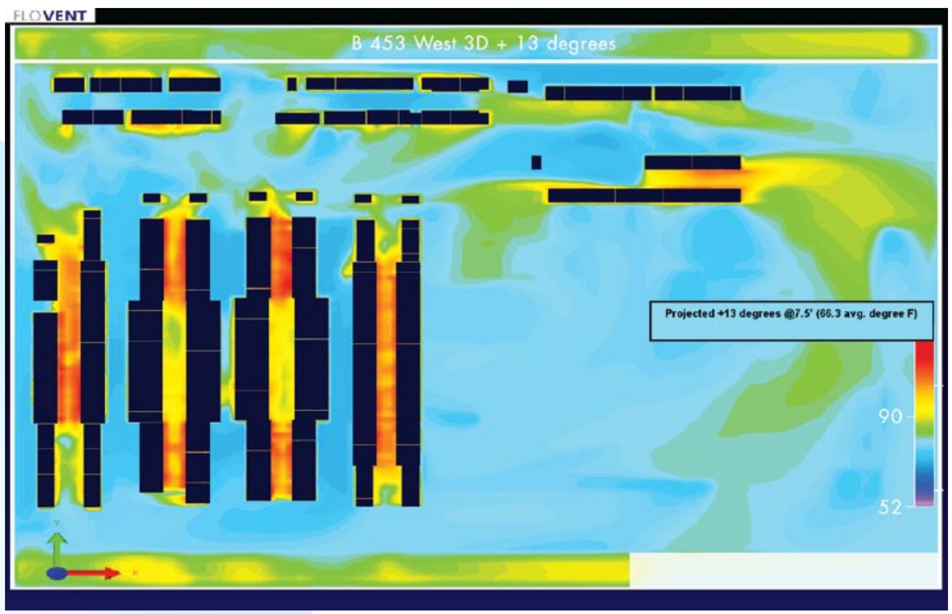

This figure shows the projected air temperatures at 7.5 feet above the floor with the supply-air temperature raised 13°F above the original design temperature in the TSF computer room. Although the average room temperature is relatively cool, the warmest areas in the exhaust aisles at the tops of the racks help determine if the hottest components are adequately cooled.

5

APPENDICES

NEXT PAGE

Computation Directorate Industrial Collaborators

Company	Topic	LLNL Contact
AbsInt	Worst-case execution time analysis	Dan Quinlan
Ambric	Massively parallel processor array	Maya Gokhale
Appro	Scalable capacity clusters	Mark Seager
Aramco	Oil well reservoir visualization	Eric Brugger
Arista Networks	Low-latency Ethernet networks	Matt Leininger
Barcelona Supercomputer Center, Bull, Commissariat à l'Énergie Atomique, Hewlett-Packard, Sicortex, Swedish National Computer Centre	SLURM resource management software	Morris Jette
Bosch	Worst-case execution time analysis	Dan Quinlan
Cisco Systems, Dell, DataDirect Networks, Intel, LSI Industries, Mellanox, Qlogic, Red Hat, Sun Microsystems, Supermicro	Hyperion collaboration	Mark Seager
Cluster Resources, Inc.	MOAB workload manager	Don Lipari
Commissariat à l'Énergie Atomique, NNSA (LLNL, SNL, LANL)	Deploying high-performance computing to produce more precise and reliable simulations	Kim Cupps
Cray	Programming environments	Bronis de Supinski
Cray	Scalable systems, open-source strategy	Mark Seager
DataDirect Networks	Raid 6 R&D for I/O systems	Brent Gorda
Exegy	Regular expression accelerator	Maya Gokhale
Fusion I/O	NAND flash memory	Maya Gokhale
IBM	High-performance storage system	Jerry Shoopman
IBM	Phase-change memory	Maya Gokhale
IBM	Scalable systems, multiple areas	Mark Seager and others
IBM	Seismic modeling	Shawn Larsen
IBM Haifa	Software assurance and testing	Dan Quinlan

Company	Topic	LLNL Contact
Imigene	DNA signatures for blood-borne diseases	Tom Slezak
Impulse Accelerated Technologies	Field-programmable gate array C compiler	Maya Gokhale
Intel Corporation	Optimization of seismic algorithms	Shawn Larsen
Intel Corporation	Solid-state drive	Maya Gokhale
Microfluidic Systems	Pathogen signatures	Tom Slezak
Netezza	Database engine	Maya Gokhale
NVIDIA	3D-mesh compression	Jon Cohen
OpenMP Consortium	Shared-memory programming models	Bronis de Supinski
OpenWorks	Valgrind memory tool and threading tool development	John Gyllenhaal
Red Hat	Operating systems	Mark Seager
Saber	Binary analysis	Dan Quinlan
Schlumberger	Synthetic seismic data set generation	Shawn Larsen
Spansion, Inc.	NOR flash memory	Maya Gokhale
SRI International	Acoustic and meteorological modeling	Shawn Larsen
Sun Microsystems	Lustre file system deployment	Brent Gorda and Mark Gary
Sun Microsystems	Seismic modeling on grid systems	Shawn Larsen
Terascala	High-performance I/O subsystems	Brent Gorda
TotalView Technologies	Parallel debugger scalability and enhanced memory tools	Dong Ahn and Scott Futral
URS Corporation	Earthquake hazard estimation	Shawn Larsen
Violin Memory, Inc.	NAND flash memory	Maya Gokhale
Voltaire	OFED: Open-source software stack for InfiniBand networks	Matt Leininger
XtremeData, Inc.	Field-programmable gate array accelerator	Maya Gokhale

Academic Outreach

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
UC Berkeley	Chris McKee	Joint Research	Computational Astrophysics Consortium	ASCR SciDAC	Louis Howell
UC Berkeley	Doug Dreger	Subcontract	Earthquake hazard	IGPP	Shawn Larsen
UC Berkeley	James Demmel	Joint Research	Towards optimal petascale simulations	ASCR SciDAC	Robert Falgout
UC Berkeley	Jonathan Wurtele	Subcontract	Theoretical and numerical investigations of Raman backscatter	LLNL Overhead	Richard Berger
UC Berkeley	Kurt Miller and Michael Jordan	Subcontract	Latent variable models	LDRD	Tina Eliassi-Rad
UC Davis	Bernd Hamann	Joint Research	Analysis and visualization of scientific data using topology-based methods	LDRD	Peer-Timo Bremer
UC Davis	Bernd Hamann	Collaboration, Lawrence Scholar Program	Topological methods for visualization	UCOP	Valerio Pascucci
UC Davis	Dave Wittman	Joint Research	Large Synoptic Survey Telescope long-term data analysis	LDRD	Celeste Matarazzo
UC Davis	Francois Gygi	Subcontract	Algorithms for electronic structure and first-principles molecular dynamics simulations using large-scale parallel computers	ASC	Art Mirin
UC Davis	Ken Joy	Joint Research	Visualization and Analytics Center for Enabling Technologies	ASCR SciDAC	Hank Childs
UC Davis	Ken Joy	Subcontract	Semantic analysis of large-scale aerial video	LDRD	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	Robust methods to cross parameterize meshes of arbitrary unequal genus	UCOP	Valerio Pascucci
UC Davis	Kwan-Liu Ma	Joint Research	Large-graph data visualization	LDRD	Peer-Timo Bremer
UC Davis	Soheil Ghiasi	Collaboration	Network security	LDRD	Maya Gokhale
UC Davis	Zhendong Su	Subcontract	Software security analysis research	LDRD	Dan Quinlan
UC Los Angeles	Alfonso Cardenas	Collaboration, Lawrence Scholar Program	iScore: measuring the interestingness of articles in a limited user environment	UCOP	David Buttler
UC Los Angeles	Kendall Houk	Subcontract	Using a combined quantum mechanics/molecular mechanics approach with Car–Parrinello molecular dynamics rare-event methods to study enzymatic reaction mechanisms	LLNL Overhead	James McGraw
UC Los Angeles	Thomas Grismayer	Collaboration	Comparison of particle and continuum kinetic simulations of laser–plasma interaction		Jeff Hittinger
UC Los Angeles and Weizman Institute (Israel)	Achiezer Brandt	Consultant	Geometric and algebraic multigrid techniques	ASCR Base	Robert Falgout
UC Merced	Qinghua Guo	Joint Research	Learning from presence-only data	UCOP	Tina Eliassi-Rad
UC San Diego	Allan Snavely and Laura Carrington	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
UC San Diego	Charles Elkan	Collaboration, Lawrence Scholar Program	Learning from presence-only data	UCOP	Tina Eliassi-Rad
UC San Diego	David Benson	Subcontract	A predictive model of fragmentation using adaptive mesh refinement and hierarchical material model	LDRD	Alice Koniges
UC San Diego	Joel Conte	Joint Research	Sensitivity analysis for structural engineering and uncertainty quantification	LDRD	Charles Tong

Academic Outreach (continued)

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
UC San Diego	Randy Bank	Consultant	Numerical solutions of partial differential equations, multilevel iterative methods, and adaptive grid generations	ASC	Robert Falgout and Panayot Vassilevski
UC San Diego	Scott Baden	Subcontract	Data-driven execution of latency-tolerant algorithms	LLNL Overhead	Dan Quinlan
UC San Diego	Tim Barnett, Michael Norman, and Randy Bank	Subcontract	LLNL–UCSD scientific data management	UC Management Contract	James McGraw
UC Santa Cruz	Ethan Miller and Carlos Maltzahn	Collaboration	Semantic file systems	LDRD	Maya Gokhale
UC Santa Cruz	Stan Woosley	Joint Research	Computational Astrophysics Consortium	ASCR SciDAC	Louis Howell
Texas A&M University	Joseph Pasciak	Subcontract	Stability and solution techniques for the Brinkman and Stokes equations	ASCR Base	Panayot Vassilevski
Texas A&M University	Marv Adams	Joint Research	Numerical methods for radiation transport	ASC	Peter Brown
Texas A&M University	Raytcho Lazarov	Joint Research	Algebraic multigrid for problems for H(curl)	ASCR Base	Panayot Vassilevski
Texas A&M University	Bjarne Stroustrup and Lawrence Rauchwerger	Joint Research	Compiler construction and parallel optimizations	ASCR	Dan Quinlan
Atomic Weapons Establishment	Matt Wheeler and Paul Selby	Joint Research	Development and use of VisIt on Atomic Weapons Establishment applications	ASC	Hank Childs
Atomic Weapons Establishment	Jim Andrews	Joint Research	Debris and shrapnel experiments and modeling	LDRD	Alice Koniges
Ball State University	Irene Livshits	Subcontract	Algebraic multigrid algorithms for finding many eigenpairs of partial differential operators	ASCR	Robert Falgout
Brigham Young University	Bryan Morse	Subcontract	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
Cambridge University	Nikos Nikiforkis	Joint Research	Simulation and modeling using Overture	ASCR Base	Bill Henshaw
Carnegie Mellon University	Christos Faloutsos	Subcontract	Mining large dynamic weighted graphs	LDRD	Tina Eliassi-Rad
Carnegie Mellon University	Hanghang Tong	Collaboration	Mining large dynamic weighted graphs	LDRD	Tina Eliassi-Rad
Colorado State University	Donald Estep	Subcontract	A posteriori error calculation of hydrodynamics simulations using adjoint methodologies	LDRD	Carol Woodward
Colorado State University	Michelle Strout and Sanjay Rajopadhye	Collaboration	Program analysis	ASCR	Dan Quinlan
Columbia University	David Keyes	Consultant	Algorithms for the solution of partial differential equations on massively parallel computers; general high-performance computing	LLNL Overhead	James McGraw
Columbia University	David Keyes	Subcontract	Acting Director for the Institute for Scientific Computing Research	LLNL Overhead	James McGraw
Columbia University	David Keyes	Joint Research	Towards optimal petascale simulations	ASCR SciDAC	Robert Falgout
Commissariat à l'Énergie Atomique (CEA)	Alain Geille	Joint Research	Debris and Shrapnel Modeling for National Ignition Facility and Laser MegaJoule, CEA	NIF	Alice Koniges
Commissariat à l'Énergie Atomique (CEA)	Thierry Carrard	Joint Research	Multiresolution visualization techniques	ASC	Hank Childs

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
Cornell University	Claire Cardie	Joint Research	Coreference resolution	DHS	David Buttler
Cornell University	Sally McKee	Subcontract	Leveraging OpenAnalysis for alias analysis within ROSE	ASC	Dan Quinlan
Duke University	Herbert Edelsbrunner	Joint Research	Discrete methods for computing continuous functions	LDRD	Peer-Timo Bremer
Imperial College	Paul Kelly and Jose Gabriel de Figueiredo Coutinho	Collaboration	Field-programmable gate arrays research	ASCR	Dan Quinlan
Indiana University	Jeremiah Wilcock	Joint Research	Binary analysis	ASCR	Dan Quinlan
Johns Hopkins University	Allan Boyles	Collaboration	Seismoacoustic modeling for defense-related efforts		Shawn Larsen
Krell Institute	Lucille Kilmer	Subcontract	Department of Energy High-Performance Computer Science Fellowship Program: LLNL portion	ASC	John May
Krell Institute	Tom Brennan	Subcontract	Investigative work on cooperative parallelism	LLNL Overhead	James McGraw
National Optical Astronomy Observatory	Chris Smith	Joint Research	Large Synoptic Survey Telescope data management and high-speed data acquisition and analysis	LDRD	Don Dossa
National University of Defense Technology	Hongjia Cao	Joint Research	SLURM resource management software	LC	Morris Jette
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
Ohio State University	P. Sadayappan and Christophe Alias	Collaboration	Optimizing compiler program analysis	ASCR	Dan Quinlan
Ohio State University	Joel Saltz and Tahsin Kurc	Collaboration	Data science algorithms		Ghaleb Abdulla
Penn State University	Jinchao Xu, James Brannick, and Ludmil Zikatanov	Subcontract	Multigrid methods for numerical models arising in plasma simulations and quantum field theories	ASCR	Robert Falgout
Portland State University	Karen Karavanic	Joint Research	Performance analysis infrastructure for petascale applications	ASC CSSE	John May
Princeton Plasma Physics Laboratory	Ravi Samtaney	Collaboration	Implicit methods for magnetohydrodynamics	ASCR SciDAC	Lori Diachin
Princeton University	Adam Burrows	Joint Research	Computational Astrophysics Consortium	ASCR SciDAC	Louis Howell
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	Ziqiang Cai	Summer Faculty	A posteriori error estimates for partial differential equations	ASC	Barry Lee
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	Ken Jansen	Joint Research	Interoperable technologies for advanced petascale simulation	ASCR SciDAC	Lori Diachin
Rensselaer Polytechnic Institute	Mark Shephard	Joint Research	Terascale simulation tools and technologies	ASCR SciDAC	Lori Diachin
Rice University	John Mellor-Crummey	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski

Academic Outreach (continued)

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
Royal Institute of Technology, Sweden	Heinz-Otto Kreiss	Consultant	Adaptive methods for partial differential equations	ASCR Base	Lori Diachin and Anders Petersson
Rutgers University	Fred Roberts and Paul Kantor	Joint Research	DHS: University Affiliates Centers	DHS	James McGraw
San Diego State University	Calvin Johnson	Subcontract	Frontier computations in the structure of atomic nuclei	LLNL Overhead	Erich Ormand
Southern Methodist University	Dan Reynolds	Joint Research	Implicit solvers and preconditioning techniques for simulations of magnetohydrodynamics and core-collapse supernovae	ASCR Base	Carol Woodward
Southern Methodist University	Thomas Hagstrom	Joint Research	High-order structure grid methods for wave propagation on complex unbounded domains	ASCR Base	Bill Henshaw
Stanford University	Greg Beroza	Collaboration	Simulation of historical and scenario earthquakes in California		Shawn Larsen
Stanford University	Parvis Moin	ASC Predictive Science Academic Alliance Program Center	Center for Predictive Simulations of Multiphysics Flow Phenomena with Application to Integrated Hypersonic Systems	ASC	Dick Watson
Stanford University/SLAC National Accelerator Center	Stuart Marshall and Jacek Becla	Joint Research	Large Synoptic Survey Telescope camera data output hardware and imaging formats	LDRD	Celeste Matarazzo
State University of New York, Stony Brook	Doug Swesty	Joint Research	Solvers for supernova simulation	ASCR SciDAC	Carol Woodward
State University of New York, Stony Brook	Jim Glimm	Joint Research	Terascale simulation tools and technologies	ASCR SciDAC	Lori Diachin
State University of New York, Stony Brook	Michael Zingale and Eric Myra	Subcontract	Verification and validation of radiation hydrodynamics for astrophysical applications	LDRD	Louis Howell
State University of New York, Stony Brook	Xiao-Lin Li	Joint Research	Interoperable technologies for advanced petascale simulation	ASCR SciDAC	Lori Diachin
Swiss Federal Institute of Technology	Stephan Brunner	Subcontract	Efficient numerical algorithms for Vlasov simulation of laser–plasma interactions	LDRD	Jeff Hittinger
Swiss Federal Institute of Technology	Stephan Kerkemeier	Joint Research	Global Nuclear Energy Partnership visualization		David Bremer
Technical University of Vienna	Markus Schordan	Collaboration	Compiler construction	ASCR	Dan Quinlan
Tufts University	Scott MacLachlan	Joint Research	Algebraic multigrid algorithms	ASC	Robert Falgout
University of Nevada, Reno	John Louie	Collaboration	Seismic modeling in the Basin and Range Region		Shawn Larsen
University of Arizona	Tim Axelrod	Joint Research	Large Synoptic Survey Telescope data management and high-speed data acquisition and analysis	LDRD	Celeste Matarazzo
University of British Columbia	Carl Olivier-Gooch	Subcontract	Interoperable technologies for advanced petascale simulation	ASCR SciDAC	Lori Diachin
University of Chicago	Don Lamb	ASC Academic Strategic Alliance Program Center	Center for Astrophysical Thermonuclear Flashes	ASC	Dick Watson

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
University of Colorado	John Ruge	Joint Research	Algebraic multigrid algorithms	ASC	Robert Falgout
University of Colorado	Joshua Nolting	Joint Research	Multipass interpolation for algebraic multigrid	ASC	Ulrike Yang
University of Colorado	Marian Brezina	Joint Research	Algebraic multigrid, spectral AMGe, FOSPACK, applications to ALE3D problems	ASC	Robert Falgout
University of Colorado	Steve McCormick	Joint Research	Towards optimal petascale simulations	ASCR SciDAC	Robert Falgout
University of Colorado	Steve McCormick	Subcontract	Geometric and algebraic multigrid methods for quantum chromodynamics, magnetohydrodynamics, elasticity, transport, and other DOE applications	ASCR, ASC	Robert Falgout
University of Colorado	Tom Manteuffel	Joint Research	Solution methods for transport problems	ASC	Peter Brown
University of Colorado	Tom Manteuffel	Joint Research	Geometric and algebraic multigrid techniques; first-order system least-squares approach to solving partial differential equations	ASC	Robert Falgout
University of Colorado	Tom Manteuffel	Joint Research	Towards optimal petascale simulations	ASCR SciDAC	Robert Falgout
University of Delaware	Calvin Keeler	Subcontract	Avian influenza resequencing microarray		Tom Slezak
University of Delaware	Richard Braun	Collaboration	Models of the eye	ASCR Base	Bill Henshaw
University of Georgia	David Lowenthal	Joint Research	Power-aware computing for message-passing interface programs; scalable performance models	ASCR, ASC	Bronis de Supinski
University of Houston	Yuriy Fofanov	Joint Research	Genomic algorithms	DHS	Tom Slezak
University of Illinois/ National Center of Supercomputing Applications	Dave Semeraro	Joint Research	VisIt technology applications on National Center of Supercomputing Applications Blue Water system	ASC	Hank Childs
University of Illinois/ National Center of Supercomputing Applications	Ray Plante	Joint Research	Large Synoptic Survey Telescope data management and high-speed data acquisition and analysis	LDRD	Celeste Matarazzo
University of Illinois, Urbana-Champaign	Dan Roth, David Forsythe, Jiawei Han, and Cheng Xiang Zhai	Joint Research	DHS: University Affiliates Centers	DHS	James McGraw
University of Illinois, Urbana-Champaign	Michael Heath	ASC Academic Strategic Alliance Program Center	Center for Simulation of Advanced Rockets	ASC	Dick Watson
University of Maryland	Jeff Hollingsworth	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
University of Maryland	Lise Getoor	Subcontract	Role discovery in evolving semantic graphs	LDRD	Tina Eliassi-Rad
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 Minnesota	Douglas Arnold	Participating Institution	Institute for Mathematics and Its Applications	ASC	James McGraw
University of Munich	Dieter August Kranzlmueller	Joint Research	Detecting communication patterns to optimize applications	ASCR, ASC	Bronis de Supinski
University of North Carolina	Robert Fowler	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski

Academic Outreach (continued)

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
University of Pittsburgh	Janyce Wiebe	Joint Research	DHS: University Affiliates Centers	DHS	David Buttler
University of Southern California	Eduard Hovy and Patrick Pantel	Joint Research	DHS: University Affiliates Centers	DHS	James McGraw
University of Southern California	Robert Lucas, Mary Hall, and Jacqueline Chame	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
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	Omar Ghattas	Joint Research	Towards optimal petascale simulations	ASCR SciDAC	Robert Falgout
University of Texas, Austin	Robert Moser	ASC Predictive Science Academic Alliance Program Center	Center for Predictive Engineering and Computational Sciences (PECOS)	ASC	Dick Watson
University of Texas, San Antonio	Qing Yi	Subcontract	Program analysis and optimization for the empirical tuning of scientific applications	ASCR	Dan Quinlan
University of Utah	Charles Hansen	Consultant	Data exploration, multiresolution scientific data visualization, and algorithm design	ASC CSSE	Mark Duchaineau
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	Hank Childs
University of Utah	Claudio Silva	Joint Research	Studying the topology of point-set surfaces	LDRD	Peer-Timo Bremer
University of Utah	David Pershing	ASC Academic Strategic Alliance Program Center	Center for the Simulation of Accidental Fires and Explosions	ASC	Dick Watson
University of Utah	Ellen Riloff	Subcontract	Coreference resolution research	LDRD	David Buttler
University of Utah	Gannesh Gopalakrishnan	Collaboration	Message-passing interface optimizations	ASCR	Dan Quinlan
University of Utah	Valerio Pascucci	Subcontract	Algorithms for geometric processing, image segmentation, and data streaming	LDRD	Peer-Timo Bremer
University of Washington	Carl Ebeling	Subcontract	Coarse-grain processor architectures		Maya Gokhale
University of Waterloo	Hans de Sterck	Joint Research	Long-distance interpolation for algebraic multigrid	ASCR SciDAC	Ulrike Yang
University of Wisconsin	Bart Miller	Joint Research	Support for enhanced Dyninst testing and initial steps toward open binary editing environment	ASC CSSE	Martin Schulz
University of Wisconsin	Ben Liblit	Joint Research	Scalable debugging	ASCR, ASC	Bronis de Supinski
University of Wisconsin	Dan Negrut	Joint Research	Implicit integration methods for molecular dynamics	ASCR	Radu Serban
University of Wisconsin	Jason Kraftcheck	Joint Research	Mesquite software development	ASCR Base	Lori Diachin
University of Wuppertal	Karsten Kahl	Visiting Researcher	Algebraic multigrid algorithms	ASC	Robert Falgout
Virginia Institute of Technology	Demitrios Nikolopoulos	Joint Research	Power-aware computing for OpenMP programs	ASCR, ASC	Bronis de Supinski
Virginia Institute of Technology	Kirk Cameron	Joint Research	Power-aware computing for hybrid systems	ASCR, ASC	Bronis de Supinski

Publications

Books and Book Chapters

Koniges, A. E., et al. "Development Strategies for Modern Predictive Simulation Codes," in *Parallel Computing: Architectures, Algorithms and Applications Volume 15 Advances in Parallel Computing*, C. Bischof, M. Bücker, P. Gibbon, G. Joubert, T. Lippert, B. Mohr, and F. Peters (IOS Press, Amsterdam, The Netherlands), pp. 697–704 (UCRL-PROC-235391).

Vassilevski, P. S., *Multilevel Block Factorization Preconditioners: Matrix-Based Analysis and Algorithms for Solving Finite Element Equations* (Springer, New York), 530 pp.

Journal Papers

Aagaard, B., et al. "Ground-Motion Modeling of the 1906 San Francisco Earthquake, Part II: Ground-Motion Estimates for the 1906 Earthquake and Scenario Events," *B. Seismol. Soc. Am.* **98** (2), 1012–1046 (UCRL-JRNL-228005).

Aamodt, K., et al. "The ALICE Experiment at the CERN LHC," *J. Instrum.* **3**.

Abazov, V., et al. "Evidence for Production of Single Top Quarks," *Phys. Rev. D* **78** (1).

Allen, J. E., S. N. Gardner, and T. R. Slezak, "DNA Signatures for Detecting Genetic Engineering in Bacteria," *Genome Biol.* **9** (3).

Amendt, P., et al. "Rugby-Like Hohlraum Experimental Designs for Demonstrating X-Ray Drive Enhancement," *Phys. Plasmas* **15** (1).

Anderson, J., C. Garth, M. Duchaineau, and K. Joy, "Discrete Multimaterial Interface Reconstruction for Volume Fraction Data," *Comput. Graph. Forum* **27** (3), 1015–1022.

Armesto, N., et al. "Heavy-Ion Collisions at the LHC—Last Call for Predictions," *J. Phys. G-Nucl. Partic.* **35** (5).

Aubert, B., et al. "Study of *B* Meson Decays with Excited η and η' Mesons," *Phys. Rev. Lett.* **101** (9).

Aubert, B., et al. "Observation of the Bottomonium Ground State in the Decay $Y(3S) \rightarrow \gamma\eta_b$," *Phys. Rev. Lett.* **101** (7).

Babonneau, D., et al. "Efficient Multi-keV X-Ray Sources from Laser-Exploded Metallic Thin Foils," *Phys. Plasmas* **15** (9).

Bala, G., et al. "Simulated Climate Near Steep Topography: Sensitivity to Numerical Methods for Atmospheric Transport," *Geophys. Res. Lett.* **35** (14) (UCRL-JRNL-400217).

Bala, G., et al. "Evaluation of a CCSM3 Simulation with a Finite Volume Dynamical Core for the Atmosphere at 1 Degrees Latitude × 1.25 Degrees Longitude Resolution," *J. Climate* **21** (7), 1467–1486 (UCRL-JRNL-228921).

Banks, J., T. Aslam, and W. Rider, "On Sublinear Convergence for Linearly Degenerate Waves in Capturing Schemes," *J. Comput. Phys.* **227** (14), 6985–7002 (LLNL-JRNL-402386).

Banks, J., W. Henshaw, D. Schwendeman, and A. Kapila, "A Study of Detonation Propagation and Diffraction with Compliant Confinement," *Combust. Theor. Model.* **12** (4), 769–808 (UCRL-JRNL-233735).

Barber, J., K. Kadau, T. Germann, and B. Alder, "Initial Growth of the Rayleigh–Taylor Instability via Molecular Dynamics," *Eur. Phys. J. B* **64** (2), 271–276.

Barnett, T., et al. "Human-Induced Changes in the Hydrology of the Western United States," *Science* **319** (5866), 1080–1083 (UCRL-JRNL-236198).

Barton, N., et al. "Embedded Polycrystal Plasticity and Adaptive Sampling," *Int. J. Plasticity* **24** (2), 242–266 (UCRL-JRNL-226715).

Bennett, J., V. Pascucci, and K. Joy, "A Genus-Oblivious Approach to Cross Parameterization," *Comput. Aided Geom. D* **25** (8), 592–606 (LLNL-JRNL-405252).

Bogomjakov, A., C. Gotsman, and M. Isenburg, "Distortion-Free Steganography for Polygonal Meshes," *Comput. Graph. Forum* **27** (2), 637–642.

Bond, E., et al. "Preparation of ^{235m}U Targets for $^{235}(\eta, \eta')^{235m}$U Cross-Section Measurements," *J. Radioanal. Nucl. Ch.* **276** (2), 549–554.

Bonfils, C., et al. "Identification of External Influences on Temperatures in California," *Climatic Change* **87**, S43–S55 (UCRL-JRNL-221904).

Bonfils, C., et al. "Detection and Attribution of Temperature Changes in the Mountainous Western United States," *J. Climate* **21** (23), 6404–6424 (UCRL-JRNL-237327).

Bowler, D., et al. "Introductory Remarks: Linear Scaling Methods—Preface," *J. Phys. Condens. Mat.* **20** (29).

Brezina, M., et al. "A Generalized Eigensolver Based on Smoothed Aggregation (GES-SA) for Initializing Smoothed Aggregation (SA) Multigrid," *Numer. Linear Algebr.* **15** (2–3), 249–269 (UCRL-JRNL-231442).

Chand, K., et al. "Toward Interoperable Mesh, Geometry, and Field Components for PDE Simulation Development," *Eng. Comput. Germany* **24** (2), 165–182 (UCRL-JRNL-213577).

Chatrchyan, S., et al. "The CMS Experiment at the CERN LHC," *J. Instrum.* **3**.

Cosman, M., et al. "Identification of a Thermo-Regulated Glutamine-Binding Protein from *Yersinia pestis*," *Protein Peptide Lett.* **15** (9), 887–894 (UCRL-JRNL-213767).

Day, S., et al. "Model for Basin Effects on Long-Period Response Spectra in Southern California," *Earthq. Spectra* **24** (1), 257–277.

De Sterck, H., R. Falgout, J. Nolting, and U. Yang, "Distance-Two Interpolation for Parallel Algebraic Multigrid," *Numer. Linear Algebr.* **15** (2–3), 115–139 (UCRL-JRNL-230844).

de Supinski, B., et al. "BlueGene/L Applications: Parallelism on a Massive Scale," *Int. J. High Perform. Comput.* **22** (1), 33–51 (UCRL-JRNL-224370).

Delle Monache, L., et al. "Bayesian Inference and Markov Chain Monte Carlo Sampling to Reconstruct a Contaminant Source on a Continental Scale," *J. Appl. Meteorol. Clim.* **47** (10), 2600–2613.

Dewald, E., et al. "NIF-Scale Re-Emission Sphere Measurements of Early-Time Tr=100 eV Hohlraum Symmetry," *Rev. Sci. Instrum.* **79** (10).

Di, L., et al. "The Development of a Geospatial Data Grid by Integrating OGC Web Services with Globus-Based Grid Technology," *Concurr. Comp.-Pract. E.* **20** (14), 1617–1635.

Edelsbrunner, H., et al. "Time-Varying Reeb Graphs for Continuous Space–Time Data," *Comput. Geom.-Theory Appl.* **41** (3), 149–166 (LLNL-JRNL-403147).

Publications (continued)

Eliasson, V., W. Henshaw, and D. Appelo, "On Cylindrically Converging Shock Waves Shaped by Obstacles," *Physica D-Non Lin. Phen.* **237** (14–17), 2203–2209 (UCRL-JRNL-232832).

Falgout, R. D., "Multigrid Methods," *Numer. Linear Algebr.* **15** (2–3), 85–87.

Fallis, J., et al. "Determination of the Proton Separation Energy of Rh-93 from Mass Measurements," *Phys. Rev. C* **78** (2) (LLNL-JRNL-405613).

Fang, H., et al. "Optimization of the Cooling Profile to Achieve Crack-Free Yb:S-FAP Crystals," *J. Cryst. Growth* **310** (16), 3825–3832 (UCRL-JRNL-233948).

Fattebert, J.-L., "Adaptive Localization Regions for O(N) Density Functional Theory Calculations," *J. Phys.-Condens. Mat.* **20** (29), 294210 (UCRL-JRNL-400527).

Gatti-Bono, C., "Coupling and Decoupling of the Acoustic and Gravity Waves through Perturbational Analysis of the Euler Equations," *J. Comput. Phys.* **227** (3), 1609–1612 (UCRL-JRNL-232445).

Gleckler, P., K. Taylor, and C. Doutriaux, "Performance Metrics for Climate Models," *J. Geophys. Res.-Atmos.* **113** (D6).

Gokhale, M., et al. "Hardware Technologies for High-Performance Data-Intensive Computing," *Computer* **41** (4), 60–68.

Graves, D., D. Trebotich, G. Miller, and P. Colella, "An Efficient Solver for the Equations of Resistive MHD with Spatially Varying Resistivity," *J. Comput. Phys.* **227** (10), 4797–4804.

Gyulassy, A., P.-T. Bremer, B. Hamann, and V. Pascucci, "A Practical Approach to Morse–Smale Complex Computation: Scalability and Generality," *IEEE T. Vis. Comput. Gr.* **14** (6), 1619–1626.

Haas, B. J., et al. "Automated Eukaryotic Gene Structure Annotation Using EVidenceModeler and the Program to Assemble Spliced Alignments," *Genome Biol.* **9** (UCRL-JRNL-237232).

Henderson, P., et al. "Towards Personalized Chemotherapeutics: Exploring Variations in DNA Damage/Repair in Cells and Tumors Treated with Carboplatin Using Accelerator Mass Spectrometry," *Environ. Mol. Mutagen.* **49** (7), 552.

Henshaw, W. D., and D. W. Schwendeman, "Parallel Computation of Three-Dimensional Flows Using Overlapping Grids with Adaptive Mesh Refinement," *J. Comput. Phys.* **227** (16), 7469–7502 (UCRL-JRNL-236681).

Hindson, B. J., et al. "Diagnostic Evaluation of Multiplexed Reverse Transcription-PCR Microsphere Array Assay for Detection of Foot-and-Mouth and Lookalike Disease Viruses," *J. Clin. Microbiol.* **46** (3), 1081–1089 (UCRL-JRNL-233231).

Ipek, E., et al. "Efficient Architectural Design Space Exploration via Predictive Modeling," *ACM Trans. Archit. Code Optim.* **4** (4) (UCRL-JRNL-227222).

Jaing, C., et al. "A Functional Gene Array for Detection of Bacterial Virulence Elements," *PLoS One* **3** (5), e2163.

Jefferson, D. R., and A. D. Rubin, "New Research Results for Electronic Voting," *IEEE Secur. Priv.* **6** (3), 12–13.

Jones, I., K. Burkhart-Schultz, C. Strout, and D. Nelson, "Studies of Thioguanine-Resistant Lymphocytes Induced by In Vivo Irradiation of Mice," *Environ. Mol. Mutagen.* **49** (5), 343–350.

Kadau, K., J. Barber, T. Germann, and B. Alder, "Scaling of Atomistic Fluid Dynamics Simulations," *Phys. Rev. E* **78** (4).

Keasler, J., "Performance Portable C++: Taking Full Advantage of New Architectures," *Dr. Dobbs J.* **33** (6), 40 (LLNL-JRNL-402527).

Kilbourne, K. H., et al. "Paleoclimate Proxy Perspective on Caribbean Climate since the Year 1751: Evidence of Cooler Temperatures and Multidecadal Variability," *Paleoceanography* **23** (3).

King, J. S., et al. "Ultralow Loading Pt Nanocatalysts Prepared by Atomic Layer Deposition on Carbon Aerogels," *Nano Lett.* **8** (8), 2405–2409 (UCRL-JRNL-403291).

Knap, J., et al. "Adaptive Sampling in Hierarchical Simulation," *Int. J. Numer. Meth. Eng.* **76** (4), 572–600.

Kolev, T. V., J. E. Pasciak, and P. S. Vassilevski, "H(curl) Auxiliary Mesh Preconditioning," *Numer. Linear Algebr.* **15** (5), 455–471 (UCRL-JRNL-224227).

Lagin, L., et al. "Status of the National Ignition Facility Integrated Computer Control System (ICCS) on the Path to Ignition," *Fusion Eng. Des.* **83** (2–3), 530–534.

Lashuk, I., and P. Vassilevski, "On Some Versions of the Element Agglomeration AMGe Method," *Numer. Linear Algebr.* **15** (7), 595–620 (UCRL-JRNL-233772).

Law, R. M., et al. "TransCom Model Simulations of Hourly Atmospheric CO_2: Experimental Overview and Diurnal Cycle Results for 2002," *Global Biogeochem. Cy.* **22** (3) (UCRL-JRNL-231758).

Lemmon, G. H., and S. N. Gardner, "Predicting the Sensitivity and Specificity of Published Real-Time PCR Assays," *Ann. Clin. Microbiol. Antimicrobials* **7**, 18.

Lenhoff, R. J., et al. "Multiplexed Molecular Assay for Rapid Exclusion of Foot-and-Mouth Disease," *J. Virol. Meth.* **153** (1), 61–69 (UCRL-JRNL-232289).

Martinez-Frias, J., S. Aceves, J. Smith, and H. Brandt, "A Coal-Fired Power Plant with Zero-Atmospheric Emissions," *J. Eng. Gas Turb. Power* **130** (2).

May, M. J., et al. "Measuring the Ionization Balance of Gold in a Low-Density Plasma of Importance to Inertial Confinement Fusion," *Can. J. Phys.* **86** (1), 251–258.

Notay, Y., and P. Vassilevski, "Recursive Krylov-Based Multigrid Cycles," *Numer. Linear Algebr.* **15** (5), 473–487.

Patel, M., et al. "Complexities Associated with Moisture in Foamed Polysiloxane Composites," *Polym. Degrad. Stabil.* **93** (2), 513–519.

Patra, P. K., et al. "TransCom Model Simulations of Hourly Atmospheric CO_2: Analysis of Synoptic-Scale Variations for the Period 2002–2003," *Global Biogeochem. Cy.* **22** (4) (UCRL-JRNL-231740).

Paul, B., et al. "Chromium Renderserver: Scalable and Open Remote Rendering Infrastructure," *IEEE T. Vis. Comput. Gr.* **14** (3), 627–639.

Pierce, D. W., et al. "Attribution of Declining Western U.S. Snowpack to Human Effects," *J. Climate* **21** (23), 6425–6444 (UCRL-JRNL-236255).

Rao, S. I., et al. "Athermal Mechanisms of Size-Dependent Crystal Flow Gleaned from Three-Dimensional Discrete Dislocation Simulations," *Acta Mater.* **56** (13), 3245–3259 (UCRL-JRNL-222934).

Rodgers, A., et al. "Broadband Waveform Modeling of Moderate Earthquakes in the San Francisco Bay Area and Preliminary Assessment of the USGS 3D Seismic Velocity Model," *B. Seismol. Soc. Am.* **98** (2), 969–988 (UCRL-JRNL-228251).

Saini, S., et al. "Performance Evaluation of Supercomputers Using HPCC and IMB Benchmarks," *J. Comput. Syst. Sci.* **74** (6), 965–982.

Sanderson, M. G., et al. "A Multimodel Study of the Hemispheric Transport and Deposition of Oxidized Nitrogen," *Geophys. Res. Lett.* **35** (17) (UCRL-JRNL-235952).

Savin, D. W., et al. "Analog and Digital Simulations of Maxwellian Plasmas for Astrophysics," *Can. J. Phys.* **86** (1), 209–216.

Scrofano, R., M. Gokhale, F. Trouw, and V. Prasanna, "Accelerating Molecular Dynamics Simulations with Reconfigurable Computers," *IEEE Trans. Parall. Distrib. Sys.* **19** (6), 764–778.

Sen, P., et al. "Collective Classification in Network Data," *AI Mag.* **29** (3), 93–106.

Shindell, D. T., et al. "A Multimodel Assessment of Pollution Transport to the Arctic," *Atmos. Chem. Phys.* **8** (17), 5353–5372 (LLNL-JRNL-402271).

Singer, S. W., et al. "Characterization of Cytochrome 579, an Unusual Cytochrome Isolated from an Iron-Oxidizing Microbial Community," *Appl. Environ. Microb.* **74** (14), 4454–4462.

Stahlberg, A., et al. "Multiway Real-time PCR Gene Expression Profiling in Yeast *Saccharomyces cerevisiae* Reveals Altered Transcriptional Response of ADH-Genes to Glucose Stimuli," *BMC Genomics* **9**.

Stout, E. A., et al. "CORBA-Based Distributed Software Framework for the NIF Integrated Computer Control System," *Fusion Eng. Des.* **83** (2–3), 535–538.

Stowell, M. L., B. J. Fasenfest, and D. A. White, "Investigation of Radar Propagation in Buildings: A 10-billion Element Cartesian-Mesh FETD Simulation," *IEEE T. Antenn. Propag.* **56** (8), 2241–2250.

Sun, J., et al. "Two Heads Better Than One: Pattern Discovery in Time-Evolving Multiaspect Data," *Data Min. Knowl. Discov.* **17** (1), 111–128.

Suratwala, T., et al. "Effect of Rogue Particles on the Subsurface Damage of Fused Silica during Grinding/Polishing," *J. Non-Cryst. Solids* **354** (18), 2023–2037 (UCRL-JRNL-230858).

Tamura, M., and P. D'haeseleer, "Microbial Genotype–Phenotype Mapping by Class Association Rule Mining," *Bioinformatics* **24** (13), 1523–1529.

Trebotich, D., G. Miller, and M. Bybee, "A Penalty Method to Model Particle Interactions in DNA-Laden Flows," *J. Nanosci. Nanotechnol.* **8** (7), 3749–3756 (UCRL-JRNL-225206).

Tripp, J. L., M. B. Gokhale, and A. A. Hansson, "A Case Study of Hardware/Software Partitioning of Traffic Simulation on the Cray XD1," *IEEE Trans. Very Large Scale I* **16** (1), 66–74.

Vassilevski, P. S., and L. T. Zikatanov, "Advanced Multigrid Methods for Systems of PDEs," *Numer. Linear Algebr.* **15** (5), 393–393.

Zemla, A., and C. Zhou, "Structural Realignment in an Immunogenic Surface Region of Ricin A Chain," *Bioinform. Biol. Insights* **2**, 5–13 (UCRL-JRNL-233290).

Zhang, Y. Y., et al. "On the Diurnal Cycle of Deep Convection, High-Level Cloud, and Upper Troposphere Water Vapor in the Multiscale Modeling Framework," *J. Geophys. Res.-Atmos.* **113** (D16).

Zolensky, M., et al. "Comparing Wild 2 Particles to Chondrites and IDPs," *Meteorit. Planet. Sci.* **43** (1–2), 261–272.

Refereed Conferences

Ahn, D., et al. "Overcoming Scalability Challenges for Tool Daemon Launching," *2008 International Conference on Parallel Processing (ICPP-08)* (LLNL-CONF-401480).

Anderson, J., C. Garth, M. Duchaineau, and K. Joy, "Discrete Multimaterial Interface Reconstruction for Volume Fraction Data," *10th Eurographics/IEEE VGTC Symposium on Visualization (EuroVis 08)* (Comput. Graph. Forum, Blackwell Publishing), **27**, 1015–1022 (LLNL-CONF-400287).

Benson, D., et al. "Hierarchical Material Models for Modeling Fragmentation of Targets in Inertial Confinement Fusion Research," *International Symposium on Structures under Earthquake, Impact, and Blast Loading.*

Bogomjakov, A., et al. "Distortion-Free Steganography for Polygonal Meshes," *29th Annual Conference of the European Association for Computer Graphics* (Comput. Graph. Forum, Blackwell Publishing), **27**, 637–642.

Bronevetsky, G., et al. "CLOMP: Accurately Characterizing OpenMP Application Overheads," *4th International Workshop on OpenMP in New Era of Parallelism* (Lect. Notes Comput. Sc., Springer–Verlag Berlin), **5004**, 13–25 (LLNL-CONF-401267).

Bronevetsky, G., et al. "Error Vulnerability of Iterative Linear Algebra Methods," *22nd International Conference on Supercomputing (ICS 2008)* (ACM) (LLNL-CONF-400659).

Carr, A., et al. "Defect Classification Using Machine Learning in Laser-Induced Damage in Optical Materials," *Laser-Induced Damage in Optical Materials: 2008,* (SPIE), **7132**, (LLNL-POST-407099).

Cary, J., et al. "First Results from Core-Edge Parallel Composition in the FACETS Project," *SciDAC 2008* (J. Phys. Conf. Ser., IOP Publishing Ltd.), **125**.

Chau, D. H., et al. "GRAPHITE: A Visual Query System for Large Graphs," *ICDM'08: 8th IEEE International Conference on Data Mining,* (IEEE), 963–966 (LLNL-CONF-407302).

Childs, H., "Why Petascale Visualization and Analysis Will Change the Rules," *8th International Conference on Computational Science* (Lect. Notes Comput. Sc., Springer–Verlag Berlin), **5101**, 32.

Curtis-Maury, M., et al. "Prediction Models for Multidimensional Power-Performance Optimizations on Many Cores," *17th International Conference on Parallel Architectures and Compilation Techniques (PACT-2008)* (IEEE/ACM) (LLNL-CONF-400453).

de Supinski, B. R., et al. "An Open Infrastructure for Scalable, Reconfigurable Analysis," *International Workshop on Scalable Tools for High-End Computing (STHEC)* (LLNL-CONF-403954).

Dahlgren, T. L., et al. "Gaining Confidence in Scientific Applications through Executable Interface Contracts—Art. No. 012086," *4th Annual Scientific Discovery through Advanced Computing Conference (SciDAC 2008)* (J. Phys. Conf. Ser. IOP Publishing LTD), **125**, 12086 (LLNL-CONF-404758).

Damkroger, T., T. Eliassi-Rad, M. Gokhale, J. Grosh, and S. Kohn, "Rethinking Supercomputing for Data Intensive Problems," *ISC'08: 2008 International Supercomputing Conference.*

Publications (continued)

Dewald, E. L., et al. "NIF-Scale Re-Emission Sphere Measurements of Early-Time Tr=100 eV Hohlraum Symmetry," *17th Topical Conference on High-Temperature Plasma Diagnostics* (Rev. Sci. Instrum. Amer. Inst. Physics), **79**.

Gallagher, B., et al. "Using Ghost Edges for Classification in Sparsely Labeled Networks," *KDD'08: 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (ACM), 256–264 (LLNL-CONF-404625).

Gallagher, B., and T. Eliassi-Rad, "Leveraging Label-Independent Features for Classification in Sparsely Labeled Networks: An Empirical Study," *SNA–KDD'08: 2nd ACM SIGKDD Workshop on Social Network Mining and Analysis* (ACM) (LLNL-CONF-405028).

Gamblin, T., et al. "Scalable Load Balance Measurement for SPMD Codes," *Supercomputing 2008* (IEEE/ACM) (LLNL-CONF-406045).

Gezahegne, A., and C. Kamath, "Tracking Nonrigid Structures in Computer Simulations," *IEEE International Conference on Image Processing* (IEEE), 1548–1551 (LLNL-CONF-400422).

Gyulassy, A., P.-T. Bremer, B. Hamann, and V. Pascucci, "A Practical Approach to Morse–Smale Complex Computation: Scalability and Generality," *IEEE Visualization 2008* (IEEE), **14**, 1619–1626.

Hanghang T., et al. "Fast Mining of Complex Time-Stamped Events," *CIKM'08: 17th ACM Conference on Information and Knowledge Management* (ACM), 759–768 (LLNL-CONF-405539).

Henderson, P. T., et al. "Towards Personalized Chemotherapeutics: Exploring Variations in DNA Damage/Repair in Cells and Tumors Treated with Carboplatin Using Accelerator Mass Spectrometry," *39th Annual Meeting of the Environmental Mutagen Society* (Wiley–Liss), **49**, 552.

Kamath, C., "Sapphire: Experiences in Scientific Data Mining—Art. No. 012094," *4th Annual Scientific Discovery through Advanced Computing Conference (SciDAC 2008)* (J. Phys. Conf. Ser., IOP Publishing LTD), **125**, 12094.

Keasler, J., et al. "TALC: A Simple C Language Extension for Improved Performance and Code Maintainability," *9th LCI International Conference on High-Performance Clustered Computing* (LLNL-CONF-400261).

Koop, M., T. Jones, and D. Panda, "MVAPICH–Aptus: Scalable High-Performance Multitransport MPI over InfiniBand," *IEEE International Parallel and Distributed Processing Symposium (IPDPS 08)* (IEEE), 1–12 (LLNL-CONF-400676).

Koutsourelakis, P.-S., et al. "Finding Mixed Memberships in Social Networks," *AAAI 2008 Spring Symposia* (AAAI) (LLNL-CONF- 400644).

Lee, G., et al. "Lessons Learned at 208K: Towards Debugging Millions of Cores," *Supercomputing 2008* (IEEE/ACM) (LLNL-CONF-402967).

Lindstrom, P., and M. Isenburg, "Lossless Compression of Hexahedral Meshes," *Data Compression Conference 2008* (IEEE), 192–201 (UCRL-CONF-236522).

Masters, N. D., and W. J. Ye, "Multiscale Modeling of Heat Transfer in a Thermal-Sensing Atomic Force Microscope," *1st ASME Micro/Nanoscale Heat Transfer International Conference* (Amer. Soc. Mechanical Engineers), 1375–1376.

May, J., "Pianola: A Script-Based I/O Benchmark," *Petascale Data Storage Workshop* (LLNL-CONF-407121).

Panas, T., D. Quinlan, and R. Vuduc, "Tool Support for Inspecting the Code Quality of HPC Applications," *ICSE: High-Performance Workshop* (UCRL-PROC-229217).

Panas, T., D. Quinlan, and R. Vuduc, "Visualization of Program Architecture Using Whole-Program Analysis," *ICSE: Aerospace Workshop.*

Pon, R. K., et al. "Improving Naive Bayes with Online Feature Selection for Quick Adaptation to Evolving Feature Usefulness," *SIAM SDM Text Mining Workshop* (UCRL-CONF-235295).

Pon, R. K., et al. "Online Selection of Parameters in the Rocchio Algorithm for Identifying Interesting News Articles," *10th ACM International Workshop on Web Information and Data Management (WIDM)* (LLNL-CONF-405224).

Preissl, R., et al. "Using MPI Communication Patterns to Guide Source Code Transformations," *8th International Conference on Computational Science* (Lect. Notes Comput. Sc., Springer–Verlag Berlin), **5103**, 253–260.

Quinlan, D., G. Barany, and T. Panas, "Shared and Distributed Memory Parallel Security Analysis of Large-Scale Source Code and Binary Analysis," *Scalable Program Analysis Workshop.*

Ratn, P., et al. "Preserving Time in Large-scale Communication Traces," *22nd International Conference on Supercomputing (ICS 2008)* (ACM) (LLNL-CONF-400703).

Schulz, M., et al. "On the Performance of Transparent MPI Piggyback Messages," *15th European Parallel Virtual Machine and Message Passing Interface (PVM/MPI) Users Group Meeting* (Lect. Note Comput. Sc., Springer–Verlag Berlin), **5205**, 194–201.

Sjogreen, B., and H. Yee, "Smoothness Monitors for Compressible Flow Computation," *5th International Conference on Computational Fluid Dynamics* (*ICCFD5*) (LLNL-CONF-406735).

Stowell, M. L., et al. "Investigation of Radar Propagation in Buildings: A 10-billion Element Cartesian-Mesh FETD Simulation," *IEEE Antennas and Propagation Society International Symposium 2008* (IEEE), 1–9, 496–499 (UCRL-CONF-400561).

Sun, J., C. Tsourakakis, E. Hoke, C. Faloutsos, and T. Eliassi-Rad, "Two Heads Better than One: Pattern Discovery in Time-Evolving Multiaspect Data," *2008 European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases.*

Thomas, P., "Signature Visualization of Software Binaries," *ACM Conference on Software Visualization* (ACM) (LLNL-CONF-405088).

Trebotich, D., et al. "Performance of Embedded Boundary Methods for CFD with Complex Geometry—Art. No. 012083," *4th Annual Scientific Discovery through Advanced Computing Conference (SciDAC 2008)* (J. Phys. Conf. Ser., IOP Publishing LTD), **125**, 12083.

Williams, D. N., et al. "Data Management and Analysis for the Earth System Grid—Art. No. 012072," *4th Annual Scientific Discovery through Advanced Computing Conference (SciDAC 2008)* (J. Phys. Conf. Ser., IOP Publishing LTD), **125**, 12072 (LLNL-CONF-404832).

Acronyms

AAPLF........... Autonomous Alignment Process for Laser Fusion

AMPE .... Adaptive Mesh Phase Evolution

AMR............. adaptive mesh refinement

ASC .... Advanced Simulation and Computing

BKC.............. Biodefense Knowledge Center

BKMS .... Biodefense Knowledge Management System

CAR .... Computing Applications and Research

CFD.............. computational fluid dynamic

CGC .... Computing Grand Challenge

CSU .... computer support unit

DOE .... Department of Energy

ENS .... Enterprise Network Solutions

HPC .... high-performance computing

HPSS .... High-Performance Storage System

HTTP............. Hypertext Transfer Protocol

I/O .... input and output

IDL .... interactive data language

IP .... Internet Protocol

ISCR .... Institute for Scientific Computing Research

ISQA .... Institutional Software Quality Assurance

IT .... information technology

ITIL .... Information Technology Infrastructure Library

ITPD............. Information Technology Protection Division

LC .... Livermore Computing

LDA-G.......... Latent Dirichlet Allocation for Graphs

LLNL .... Lawrence Livermore National Laboratory

LSO .... Livermore Site Office

MPI .... Message Passing Interface

NIF............... National Ignition Facility

NNSA .... National Nuclear Security Administration

O&B .... Operations and Business

O|SS............ Open|SpeedShop

PD .... Principal Directorate

petaFLOPS..... quadrillion floating-point operations per second

PFM.............. phase-field model

PUE .... power usage effectiveness

RFP............... request for proposal

S&T .... Science and Technology

SAN .... storage area network

SERS............. Surface Enhanced Raman Scattering

SQA .... software quality assurance

SSU .... storage scalable unit

STAT .... Stack Trace Analysis Tool

teraFLOPS ..... trillion floating-point operations per second

TLCC .... Tri-Laboratory Linux Capacity Cluster

TOSS............ Tripod Operating System Software

TSF............... Terascale Simulation Facility