



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

2006 Computation Directorate Annual Report

M. Zosel

March 7, 2007

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This work was performed under the auspices of the U.S. Department of Energy by University of California, Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

Computation Directorate

ANNUAL REPORT

2006



Lawrence Livermore National Laboratory



About the Cover

The Atlas supercomputer shown on the cover is the largest of six new commodity Linux clusters being deployed at Lawrence Livermore. The clusters will provide more than 100 trillion floating-point operations per second of computing power to help achieve the Laboratory's important and diverse science-based simulation missions. Shown in the background are various simulations run on the high-performance computers maintained by Livermore's Computation Directorate. The simulations are: (clockwise from left, front cover) atomistic polymer ordering, dislocation dynamics, nanoporous metal material properties, water molecules in a carbon nanotube, and (back cover) protein-folding thermodynamics. (Cover design by Daniel S. Moore)

LLNL Computation Directorate Annual Report 2006

UCRL-TR-228768

LLNL Associate Director for Computation Dona L. Crawford
 Deputy Associate Directors Ronald F. Schmucker (Principal Deputy), Steven F. Ashby, and Michel G. McCoy
 Scientific Editors Mary E. Zosel and Van Emden Henson
 Art Director Daniel S. Moore
 Production Editor Deanna K. Willis
 Photographer Joseph Martinez
 Proofreader Alane L. Alchorn
 Print Production Charlie M. Arteago, Jr. and Monarch Print Copy and Design Solutions

About the Computation Directorate

Our Vision

Computation aspires to be the preeminent high-performance computing and computer science organization in order to enable scientific discovery and Laboratory missions.

Our Mission

Computation ensures Laboratory mission and program goals are attained by delivering outstanding computer science expertise, world-class high-performance computing capabilities, and creative technology and software solutions.



This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This work was performed under the auspices of the U.S. Department of Energy by University of California, Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

Contents

Awards and Recognition	2
-------------------------------------	----------

Section 1: A Blueprint for the Future: Anticipate, Innovate, and Deliver	4
---	----------

Section 2: High-Performance Computing	7
--	----------

Negotiating the Transition to Routine Capability Computing	
2.01 — The Evolution of High-Performance Computing Resources	9
2.02 — Peloton Procurement Employs Novel Strategy	10
2.03 — Thunder Supports Breakthrough Science	11
2.04 — BlueGene/L Makes New Science Applications Possible	12
2.05 — User Environment Benefits from Productivity Improvements	14
2.06 — Metrics to Mastery: Using Data to Achieve Excellence	16

Section 3: Research and Advanced Development	17
---	-----------

Defining the Path to Future Applications of Computing	
3.01 — Qbox: Understanding High-Z Metals at Record Speed	19
3.02 — AMG Scalability Improvements Lead to Faster Simulations	20
3.03 — Quantifying Uncertainties in Multiphysics Applications	21
3.04 — ALE-AMR Brings Adaptivity to Multiscale Applications	22
3.05 — Wave Propagation Project Models 1906 Earthquake	23
3.06 — Babel Sets the Standard for High-Performance Middleware	24
3.07 — ROSE: Source-to-Source Optimization of Scientific Software	25
3.08 — Gaining Insight from Large-Scale Dynamic Semantic Graphs	26
3.09 — Image Viewer Analyzes Streaming Data in Real Time	27
3.10 — Cooperative Parallelism Paves the Way for Petascale Computing	28

Section 4: Developing Applications Software	29
--	-----------

Coding for Success: Delivering Software to the Programs	
4.01 — Testing the Scalability of Supercomputers	31
4.02 — Visit: Beyond Pretty Pictures	32
4.03 — Designing Mirrors for the Linac Coherent Light Source	33
4.04 — Bringing Out the Best in the Big Lasers	34
4.05 — Beyond Pathogen Identification: Detecting Mechanisms of Bioterrorism	35
4.06 — Sonoma: A Leader in Broad-Area Surveillance	36
4.07 — Discovering Information in Unstructured Text	37
4.08 — ChemTrack: End-to-End Chemical Logistics	38

Section 5: Institutional Activities	39
--	-----------

Impacting Institutional Infrastructure, Security, and Service	
5.01 — Assuring the Quality of Safety Software at LLNL	41
5.02 — Configuration Management Database Aids Accreditation Process	42
5.03 — LANDesk Enhances IT Efficiency and Security	43
5.04 — Blue Network Improves Computer Security Posture	44
5.05 — Cyber Security Protects Against Evolving Threats	45
5.06 — Academic Partnerships Create Pipelines for People and Ideas	46
5.07 — Techbase Investments Advance Core Competencies	48

Section 6: Appendices	49
------------------------------------	-----------

6.01 — Industrial Collaborators	50
6.02 — Academic Outreach	51
6.03 — Book Chapters, Journal Papers, and Conference Papers	56
6.04 — Acronyms and Abbreviations	60

Awards and Recognition

In addition to the directorate-wide accomplishments presented in this annual report, I want to recognize several of the awards received by Computation personnel and projects during the past year. On behalf of Lawrence Livermore National Laboratory (LLNL) and the directorate, I congratulate the award winners. They embody the resolve, tenacity, and excellence that continue to inspire all of us.



Dona Crawford
Associate Director, Computation

2006 Gordon Bell Prize

For the second year in a row, Computation scientists were part of a team that received the Gordon Bell Prize for pioneering simulations on the world's fastest supercomputer—the IBM BlueGene/L (BG/L) at Livermore. Led by François Gygi, formerly of LLNL and currently at UC Davis, the team included Erik Draeger, Martin Schulz, and Bronis de Supinski of LLNL; Franz Franchetti of Carnegie Mellon University; Stefan Kral, Juergen Lorenz, and Christoph Ueberhuber of Vienna University of Technology; and Vernon Austel, John Gunnels, and James Sexton of IBM's Thomas J. Watson Research Center. The peak performance team was honored for their large-scale electronic structure simulation of the heavy metal, molybdenum. They ran their simulation using Qbox, a code that uses innovative communication efficiencies and node-mapping improvements to take advantage of massively parallel systems like BG/L. A sustained peak performance of 207.3 trillion floating-point operations per second, representing 56.5% of the theoretical full-machine peak, was measured using all 131,072 processors on the 65,536 nodes of the machine. (See Section 3.01.)



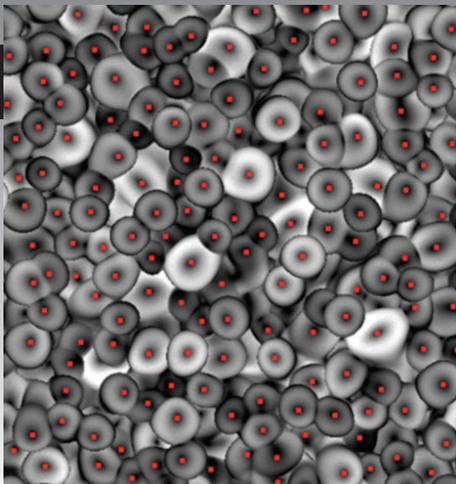
Gordon Bell Prize winners Franz Franchetti, Martin Schulz, Erik Draeger, Bronis de Supinski, and François Gygi collect their award from Gordon Bell presenter William Gropp (third from left).

Top Technological Inventions

Each year, *R&D Magazine* recognizes the 100 most technologically significant products of the year. Computation scientists played a significant role in three R&D 100 Awards for 2006.

Babel: The High-Performance Language Interoperability Tool provides a high-performance solution to the programming language–interoperability problem. It enables software pieces written in different programming languages to seamlessly call each other. Co-led by Tom Epperly and Gary Kumfert, other current team members include Tamara Dahlgren and James Leek. (See Section 3.06.)

The Sonoma Persistent Surveillance System is the first integrated, broad-area, high-resolution, real-time motion imagery system for surveillance applications. It is unique in its ability to provide continuous real-time video imagery of an area the size of a small city, with resolutions sufficient enough to track up to 8,000 moving objects within the field of view. Computation's David Bloom, Mike Kartz, and Aaron Wegner are members of the four-directorate, multidisciplinary Sonoma team. (See Section 4.06.)



Sapphire software was used to characterize and track bubbles and spikes in an 80-TB data set from a three-dimensional, high-fidelity simulation of the Rayleigh–Taylor instability. This image shows the bubble counts (in red) highlighted on the magnitude of the x–y velocity at the bubble boundary.

Sapphire: Scientific Data Mining Software is a Computation-developed software toolkit for the analysis of massive, complex data sets that arise from scientific experiments, observations, and computer simulations. Sapphire technology is being applied to a variety of disciplines, including plasma physics experiments and simulations, remote sensing imagery, video surveillance, climate simulations, astronomy, and fluid-mix experiments and simulations. The LLNL team has six patents on Sapphire technology to go with its R&D 100 Award; two of the six patents were obtained this year. Led by Chandrika Kamath, the Sapphire team includes Abel Gezahegne, Cyrus Harris, and Nu Ai Tang, as well as former LLNL employees Erick Cantú-Paz and Samson Cheung.

Using Information Technology to Benefit Society

LLNL and the National Nuclear Security Administration’s Advanced Simulation and Computing (ASC) Program were recognized by the Computerworld Honors Program in 2006 for developing the world’s fastest supercomputer, BG/L. The Computerworld awards are presented to men, women,

organizations, and institutions that have contributed to the betterment of society through the exceptional—if not heroic—use of information technology. LLNL was one of 220 laureates worldwide to be formally inducted into the Computerworld Honors Program’s international archives and was one of the top five laureates in the Science category.



Dean Williams

Value-Added Project Management

LLNL’s new 253,000-square-foot Terascale Simulation Facility (TSF) houses the ASC Program’s capability machines, Purple and BG/L. In 2006, the TSF project was awarded the Department of Energy (DOE) Secretary’s Project Management Award of Achievement for demonstrating significant results in completing a project within cost and on schedule. The \$94.1-million project demonstrated outstanding accomplishments in performance, innovation, teamwork, and safety, and was completed eight months ahead of schedule and \$1.2 million under budget. The TSF features several innovative design elements that combine to make it a computing facility that is second to none.



Excellence in Leadership

Computer scientist Dean Williams received *Science Spectrum* magazine’s Trailblazer Award for his distinguished leadership in the computational aspects of DOE’s Program for Climate Model Diagnosis and Intercomparison. Williams leads a variety of software projects aimed at data analysis and visualization of climate models. He is also active in LLNL’s recruiting efforts and participates in the summer student program to help ensure that the Laboratory continues to hire qualified minorities into scientific fields.

The TSF project team in front of the TSF.

A Blueprint for the Future: Anticipate, Innovate, and Deliver

Dona Crawford,
Associate Director
Computation Directorate



In 2006, the people of Lawrence Livermore National Laboratory's (LLNL's) Computation Directorate made impressive programmatic achievements of which I am very proud. Our focus on Laboratory missions and the partnerships with our customers, collaborators, and stakeholders enables these successes and feeds into the ideas that help the directorate anticipate long-term needs and deliver innovative solutions that will keep LLNL at the forefront of the nation's computing organizations.

As a directorate, we strive to find a balance between providing exceptional customer support to Laboratory programs today while simultaneously preparing to support them tomorrow. As we plan for the future, we must confront two critical issues at the forefront of our industry: how to plan for petaFLOPS computing and how to deal with the massive data sets that are already being produced and will grow even more daunting as computing power increases.

Our effort to bring petaFLOPS systems to LLNL is well underway. A Department of Energy-funded research and development contract is now in place that forms a partnership between IBM, LLNL, and Office of Science's Argonne

National Laboratory to develop BlueGene/P and BlueGene/Q—the follow-on systems to BlueGene/L (BG/L). We are also partnering on a bid with universities and institutions in Georgia, New York, and California to build the world's fastest computer for the National Science Foundation. If this University of California-led team wins, the system would be housed in LLNL's Terascale Simulation Facility. The formal planning cycle was also initiated for Sequoia—a procurement project to bring a multi-petaFLOPS Advanced Simulation and Computing (ASC) system to the Laboratory in 2011. This would include an initial delivery in 2008 of a scaling platform at approximately 10% of the final system's peak performance

to allow application development targeted at effectively using the 2011 full system.

It's important to remain mindful that the purpose of computing is not to generate data, but to gain insight into problems we are studying. Yet finding answers from the massive amount of data our computing systems produce can be as elusive as locating the proverbial needle in the haystack. The challenge is not unique to the intelligence community. More and more disciplines are creating bigger and more complex data sets from varied sources: data from large amounts of text, from large relational data sets or graphs, from statistics such as climate records, and from sensors and surveillance, to name just a few.

This annual report describes a number of innovative approaches that are in progress at LLNL. Two of the three R&D 100 Awards that Computation personnel were involved with this year are related to data analysis. Several more of the annual report articles discuss data-related projects that use visualization, innovative algorithms, automation, and hardware. It is vital that we continue to deliver solutions in this area to take full advantage of our large computing systems today and the even larger systems on the horizon.

Computation is further demonstrating our record of excellent innovation through our participation and leadership in the Department of Energy Office of Science Scientific Discovery through Advanced Computing (SciDAC) Program. SciDAC research projects are collaborative, multidisciplinary efforts that advance major software and algorithm development for

LLNL systems appearing in the top 100 on the Top500 Supercomputers list (November 2006).

2006 Rank	2005 Rank	Linpack Speed	System Name	Description
1	1	280.6 TF	ASC BlueGene/L	65,536-node system built in partnership with IBM.
4	3	75.7 TF	ASC Purple	1,532-node system that was also built in partnership with IBM. Deployment of this system had been a NNSA milestone since 1996.
19	11	19.9 TF	M&I Thunder	1,024-node Intel IA-64 Linux cluster system used for unclassified capability computing.
81	New	8.1 TF	ASC Zeus	288-node Linux cluster system for unclassified capacity computing.
88	45	7.6 TF	M&I MCR	1,152-node Intel IA-32 Linux cluster system for capacity computing.
90	47	7.3 TF	ASC White	512-node IBM Power 3 system decommissioned in 2006.

problems related to the Office of Science's mission. As a result of the 2006 competition for SciDAC funding, more than 30 Computation researchers are contributing to projects during the next three to five years. Two of the new multiorganization projects are being led by Computation personnel: Lori Diachin is leading the Center for Interoperable Technologies for Advanced Petascale Simulations, and Dean Williams is leading the project on Scaling the Earth System Grid to Petascale Data Center for Enabling Technologies. Through our leadership and involvement in SciDAC, Computation is helping deliver the next generation of scientific computing software infrastructure and developing new data management and knowledge discovery tools for large data sets resulting from both experiments and simulation.

Livermore Computing (LC) continued its leadership role in high-performance computing (HPC) by delivering impressive computing power for both the ASC Program and institutional programs. Completing the ASC Purple Level 1 milestone was one of the top delivery priorities for the National Nuclear Security Administration.

Before the milestone could be deemed complete, we had to demonstrate the Purple system's general availability to perform Stockpile Stewardship Program calculations.

As part of the deliverables for this milestone, full-system runs were performed by Lawrence Livermore, Los Alamos, and Sandia national laboratories to demonstrate Purple's readiness for production. The LLNL run studied

an inquiry related to upcoming National Ignition Facility capsule design decisions. Part of the calculation involved a billion-zone mesh and produced a 66-TB data set. The run produced an informative result and uncovered scaling issues that were fixed in input/output libraries and message-passing interface-tuning parameters that had not been detected on smaller problems. The run was a good test of Purple's

readiness for Grand Challenge problems. The formal report from the Purple milestone review in December voiced high praise for the usefulness of the system and for the "exceptional service" provided by LC.

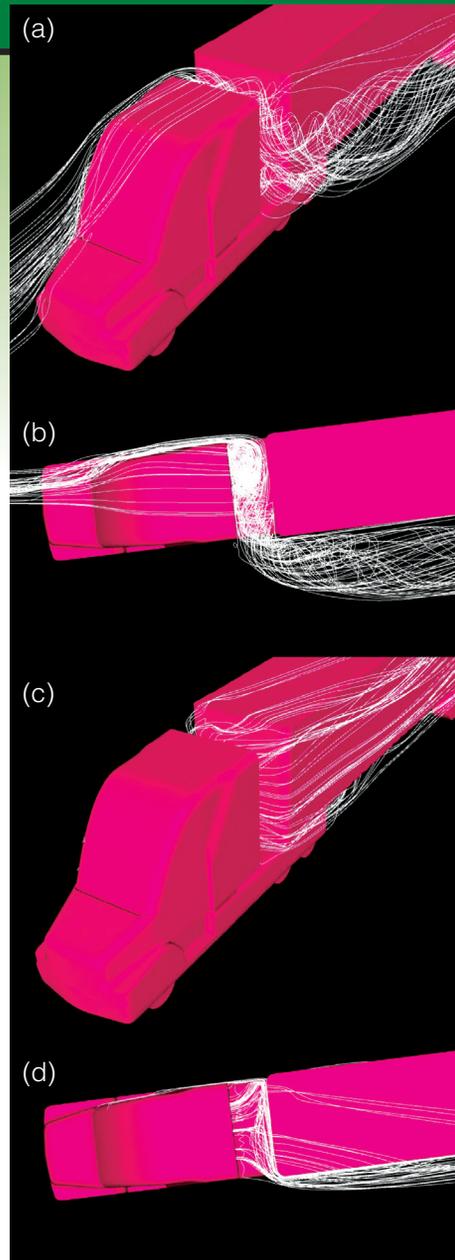
LLNL again made a strong showing in the biannual Top500 Supercomputers lists, but a comparison of the 2006 lists with the 2005 lists shows that the Laboratory's six entries in the top



The ASC Purple supercomputer provides the computing capability to routinely run simulations for the Stockpile Stewardship Program.

50 are now six entries in the top 100—an indication of how fast the HPC community is progressing. It is notable that the sustained Qbox application performance on the top-ranked BG/L system was more than twice the Linpack performance of the number 2-ranked system. This adds to BG/L's already impressive track record. Before its installation, there were several predictions that the sheer number of parts and processors would make the system too unreliable to perform any real calculations. But during the 2006 plutonium-aging production runs, the mean-time-to-failure of the system was twice the contracted value—around two weeks. BG/L has gone beyond a research project to become a productive part of LLNL's computing resources.

All these things are evidence of the strength, depth, and breadth of our directorate, but it is really our people who make things happen. In this report, we do not list the names of all the contributors to the work described; there would be many long, multidisciplinary lists if we did. Instead, we list one person to contact for additional information. Similarly, there are dozens of other interesting efforts and innovative results that are



New developments in the heavy vehicle aerodynamics drag reduction project include a "bleeding flow" device. These simulations show velocity streamlines in the gap without (a and b) and with (c and d) the tractor base device.

not detailed in this report. For example, Computation personnel were involved in a first-of-its-kind project to sequence the DNA of a poplar tree and another effort to study drag-reduction concepts for heavy vehicles that will improve truck fuel economy. In another project not detailed here, Computation employees worked on a Department of Homeland Security rail-security pilot field test that will help scientists and federal officials evaluate the effectiveness of explosives detection technologies and operating procedures for rail rapid transit.*

In closing, I would like to call special attention to and welcome two new key members of the Computation management team who are contributing to the annual report this year. Ron Schmucker assumed the position of principal deputy associate director for Computation, and John Grosh was named director of the Center for Applied Scientific Computing. Grosh came to LLNL from the Department of Defense and is highly regarded in the HPC community, even making his second appearance on HPCwire's "People to Watch" list this year.

I invite you to review this compendium of directorate projects, visit the Web sites listed for more information, and contact the investigators to find out more about their work. Thank you for your interest in the Computation Directorate at Lawrence Livermore.

Dona Crawford
Associate Director, Computation

* To learn about additional Computation efforts, visit <http://www.llnl.gov/comp>.

Negotiating the Transition to Routine Capability Computing

Routine capability computing is a holy grail that has both tantalized and eluded the high-performance computing (HPC) community for the better part of two decades. The norm has been to publicly proclaim the procurement of capability computers, intended for the largest calculations, but to quietly employ them as capacity computers as soon as the headlines fade. Two formidable adversaries, one technical and one economic, impeded the realization of true capability computing. The first was the challenge of developing scalable applications and balanced computing machines that could be used effectively together for science of scale. The second was the dearth of reliable, low-cost computing power to address the myriad smaller scale calculations required to do defensible science.

Negotiating the Transition to Routine Capability Computing (cont.)



CONTACT INFORMATION

Michel McCoy, mymccoy@llnl.gov

Capability studies rely on a foundation of capacity calculations, and, in the past, equilibrium was often achieved by relegating larger machines to capacity use. National security programs, like the Advanced Simulation and Computing (ASC) Program, and unclassified science programs, like Office of Science, struggled to contrive ad hoc solutions. Today, due to a number of developments, the national security laboratories are positioned to evolve into true capability sites. In fact, Lawrence Livermore National Laboratory (LLNL) is well into this transition in both its classified and unclassified environments. Low-cost Linux clusters, well-balanced capability computers, the maturation of powerful three-dimensional (3D) applications codes, and strategic allocation policies made the transition possible. How the change from sporadic special calculations to routine capability computing is being effected is one of Computation's highlights in 2006.

Beginning with the Multiprogrammatic Capability Resource (MCR) computing system in 2002, Livermore Computing (LC) has worked with industry to build Linux cluster systems that scale to levels that challenge many vendor-integrated, proprietary solutions. The Multiprogrammatic and Institutional Computing (M&IC) initiative could not afford proprietary capability solutions, so a strategy to enable routine capability computing was formed. The strategy was to transition existing generations of Linux computers to capacity use and employ the newest generation as the capability computer, maintaining a 50–50 capacity/capability ratio. This approach was first employed when the Thunder system displaced MCR. We are now employing the strategy with several powerful Linux clusters that were obtained through the Peloton procurement. Atlas, a 44-trillion-floating-point-operations-per-second (TF) system, will soon become the capability resource, while Thunder and an 11-TF system, Zeus, will serve the capacity needs of the Laboratory. Access to Atlas, like Thunder before it, is granted through a competitive proposal process managed by the

Director's Office. Few projects are enfranchised on the capability machine to facilitate the routine running of large problems. This is possible because the capacity environment addresses the demands from the remainder of the community.

ASC has taken this concept to the next level this year by introducing Capability Computing Campaigns (CCCs) on the Purple system to the tri-laboratory community. Each campaign represents a body of work that is made up of some very large runs and some smaller supporting runs. To be successful, a campaign must require capability cycles and fulfill a serious mission need. This approach is moving ASC (and LC) into a national user facility model. The first six-month CCC suite began on October 1, 2006. The success of this strategy is due to four factors. The first is the existence of verified and validated weapons codes (representing 10 years of systematic code development), which can use 4,000 to 8,000 processors effectively; the second is the balance and scalability of Purple (and BlueGene/L) which permit the resource to be used effectively for science of scale; the third is the existence of the classified capacity infrastructure at LLNL, which is adding at least 50 TF of additional capacity; the fourth is that, without the CCCs to control access and monitor progress, the machine might be inundated with small calculations and would flounder in a sea of capacity work.

Such developments may seem trivial and obvious, but they have, in fact, eluded the HPC community due to historically intractable technical and cost factors. Although M&IC is the smaller partner in computing at LLNL, it has often been the leader because it benefits from the leverage of ASC but enjoys the advantage of extreme flexibility. However, by employing a variant of the allocation model on Purple and BlueGene/L (BG/L), ASC has already taken capability computing to a new level; some of the scientific work that has resulted is highlighted in this annual report.

2.01 — The Evolution of High-Performance Computing Resources



CONTACT INFORMATION

Pam Hamilton, pgh@llnl.gov
 Web site: www.llnl.gov/icc/lc/

OVERVIEW

LC continued to provide HPC resources to the nuclear weapons complex and beyond. LC meets the computational needs of the Laboratory's M&IC initiative and also makes compute resources available to the entire National Nuclear Security Administration (NNSA) ASC community in both classified and unclassified environments. LC strives to provide the right balance of capacity computing (systems that run many small jobs at a time) versus capability computing (systems that run only a few large jobs but help solve the most demanding scientific problems) through the careful planning and coordination of system retirements and new system procurements.

PROGRESS IN 2006

On July 27, we permanently decommissioned the ASC White computer system, once described as “the foundation of science-based stockpile stewardship.” For more than five years, White had been a production workhorse. The first full-system, 3D simulations of a nuclear weapons explosion were run on the machine in March 2002. Until just before its retirement, White was being used by Lawrence Livermore, Sandia, and Los Alamos national laboratories. However, it was not cost-effective to maintain and power a five-year-old system when new technology was available.

The retirement of White was offset by the planned availability

of the ASC Purple computer system. Purple became available for limited production use in May and was made generally available in early November. It has been extremely productive for a new system, boasting an average utilization of 85% during the last eight months of 2006. December's utilization was an astounding 94%.

Multidisciplinary teams supporting both the M&IC and ASC programs began a half-year effort in August to deploy six new commodity Linux clusters, which will provide more than 100 TF of additional computing power to Laboratory programs. The new clusters are part of the Peloton procurement. (See Section 2.02.) Four of the clusters are Infiniband-interconnected and are built from multiples of identical scalable units. Two, four, or eight



Linux System Administrators Jeff Cunningham (front) and Ellsworth Wente take a closer look at Zeus, one of six new Linux clusters.

scalable units are aggregated into a cluster, thus ensuring hardware and software commonality among platforms. This strategy is expected to result in operational efficiencies on many levels, thereby reducing the total cost of ownership.

Yana, the first of the six system deployments, is an 80-node serial cluster (no interconnect) to be used by M&IC customers for porting and serial jobs. Yana became available for users in November. Zeus, a 288-node M&IC capacity cluster, became available in a limited mode in December. Rhea, a 576-node ASC capacity cluster was transitioned to the classified network and became available to a limited set of users for further scaling and integration

in December. Three more systems will be deployed in early 2007, including Atlas, a 1,152-node cluster that is the largest of the six systems.

SIGNIFICANCE

LC continues to strive for excellence in the HPC environment by successfully integrating systems with cutting-edge architectures, such as BG/L and Purple, while planning and procuring systems for the future. The six new systems will provide substantial additional capacity and capability for achieving important and diverse science-based simulation missions.

2.02 — Peloton Procurement Employs Novel Strategy



CONTACT INFORMATION

Mark Seager, seager@llnl.gov
 Web site: www.llnl.gov/linux/peloton/rfp/

OVERVIEW

The two previous M&IC Linux cluster procurements, MCR in 2002 and Thunder in 2004, integrated emerging open-source software technologies with commodity hardware advances. Both procurements resulted in breakthrough Linux cluster systems that have helped accomplish outstanding science. Since the MCR and Thunder procurements, we have gained a better understanding of how the commodity-computing ecosystem works. LC used this knowledge to devise a new Linux cluster procurement and deployment strategy aimed at reducing the total cost of ownership (TCO) of world-class Linux clusters. The strategy had three components: (1) leverage the commodity market's ability to build a large number of server racks quickly and cheaply, (2) leverage scalable unit building blocks based on the first stage of a federated fat-tree interconnect, and (3) partner with the entire supply chain to deploy the clusters in a more timely and cost-efficient manner.

PROGRESS IN 2006

This strategy was first developed by LC for the ASC Program's planned FY06 Tri-Laboratory

Linux Capacity Cluster (TLCC) procurement. Unfortunately, the TLCC procurement never came to fruition due to budget constraints and shifting program priorities. Instead, LLNL capitalized on the

novel approach and agreed to fund the procurement under a different name, calling it "Peloton" for the pack of riders who set the pace in bicycle races.

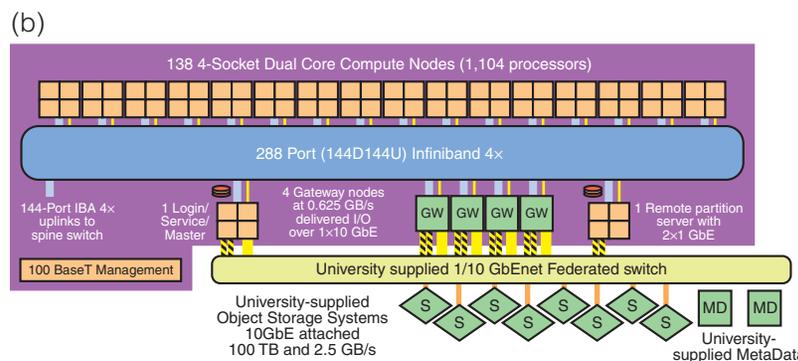
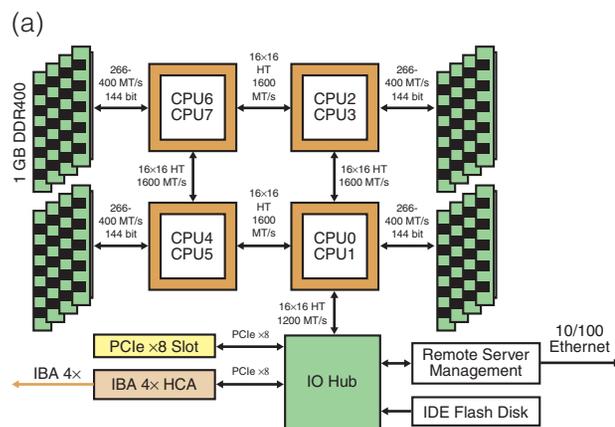
We estimate that the new procurement strategy resulted in a TCO savings of 30–50% compared to previous Linux cluster procurement practices. The savings were obtained by modeling our procurement after the contract manufacturing supply chain. For example, to improve pricing and assess margins and markups, we negotiated directly with the key component providers in the commodity supply chain for things like motherboards, processors, memory, interconnect, and disks. We also defined a scalable unit of 144 nodes based on a 288-port

Infiniband 4x double data rate first- and second-stage switch that was the highly replicated unit of manufacture. We allowed bidders to propose a delivery schedule that optimized their supply chain to minimize costs.

From the Peloton procurement, LLNL purchased and deployed 14 scalable units with an aggregate peak of 77.4 TF. In 2007, we will add six more scalable units, for a total of 110.6 TF.

SIGNIFICANCE

The Peloton procurement, deployment, and integration used a novel strategy that minimizes the TCO of Linux by purchasing a large number of scalable units and deploying them in multiple world-class clusters. Several vendors across the country have since adopted this method as their model for Linux-cluster product positioning. In addition, several HPC sites have started using it to manage their large-scale procurements. The ASC Program will likely purchase additional scalable units for Linux capacity clusters using this strategy in FY07.



(a) Generic node block diagram dual-core Opteron. (b) Example of a 4x-socket dual-core Opteron scalable unit system architecture for 144 nodes, 5.1 TF peak.

2.03 — Thunder Supports Breakthrough Science



CONTACT INFORMATION

Brian Carnes, carnes1@llnl.gov
 Web site: www.llnl.gov/icc/tc/mic/

OVERVIEW

M&IC supports several unclassified, mission-relevant projects by giving them substantial time allocations on the Thunder supercomputer to tackle problems for breakthrough science. These Grand Challenge projects range from research in climate change, to nanofluidics, to dislocation dynamics to protein folding. All of the projects involve collaboration with other leading research institutions, bringing expertise and experience to LLNL and further developing LLNL's already formidable reach in basic science.

PROGRESS IN 2006

Two of the five 2006 M&IC Grand Challenge projects are highlighted in this section. A description of the complete set of projects is covered in the Laboratory's November 2006 issue of *Science and Technology Review*.

Doug Rotman, in LLNL's Carbon Management and Climate Change Program, is leading a Grand Challenge project to model regional

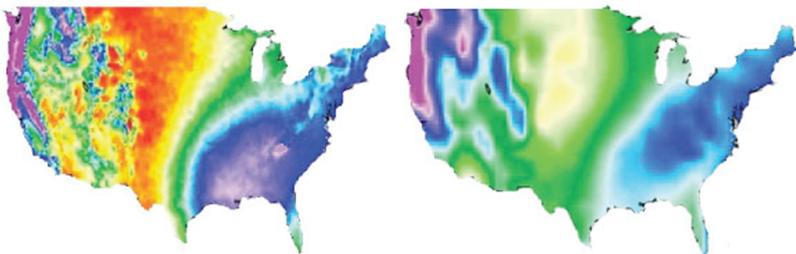
climate change at resolutions of approximately 10 km. The calculations will allow scientists to evaluate whether changes in certain regional phenomena, such as river flow, are caused by humans or by other factors. The team will perform multicentury, global simulations at a resolution of 100 km, which is finer than that of any other century-long comprehensive climate run to date. The Community Climate System Model achieves approximately 10 simulated years per computing day using 472 of

Thunder's processors and includes a calculation that simulates more than 1,100 years. Results of the global simulations then serve as initial conditions and boundary data for 10-km-scale atmospheric calculations focusing on areas the size of California.

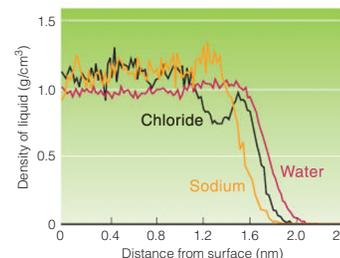
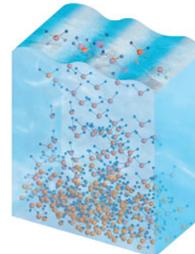
In a second Grand Challenge project, chemist Chris Mundy leads a team that has combined Monte Carlo methods with first-interaction potentials to demonstrate the vapor-liquid coexistence curve for water. Even a seemingly basic molecule, such as water, has unresolved mysteries because researchers have, until now, lacked the computational power combined with first-principles codes that are necessary to do the work. Understanding water is important for biological, physical, atmospheric, and

planetary sciences. "You can see water's vapor-liquid coexistence and, from this, determine the thermodynamic properties," explained Mundy. The results provide a better understanding of heterogeneous chemistry at aqueous interfaces. Mundy's team also redefined the predicted phase diagram of water in the interior of large planets, confirming a theory that a superionic phase of water exists in large planets. In this phase, oxygen atoms are fixed, or frozen, while hydrogen atoms move freely. "There has been a great debate about the existence of superionic water," said Mundy. "Our calculations reveal that this exotic phase may exist."

(a)



(b)



(a) Winter precipitation patterns from the Community Climate System Model (left) agree with observed precipitation from 1971–2000 (right), particularly in the western U.S. (b) Hydrogen atoms in water molecules tend to orient themselves into the air at the surface. This orientation may increase reactivity with other types of molecules. Simulations of sodium chloride and water molecules show greater interaction at the surface and decreased interaction for molecules in bulk liquid.

SIGNIFICANCE

Thanks to substantial institutional support and leverage from the ASC Program, M&IC initiated Grand Challenge efforts on institutional machines for the first time. The success of the 2006 Grand Challenge projects prove that investing in institutional computing will help further LLNL's national security and basic science missions.

2.04 — BlueGene/L Makes New Science Applications Possible

OVERVIEW

The 2001 contract between LLNL and IBM to develop BG/L emphasized the importance of science applications in guiding the design and assessing the effectiveness of the BlueGene architecture. While BG/L was being developed and built, key science application codes were ported to the supercomputer and then tuned to allow scaling to extreme levels of parallelism (more than 100,000 processing units). Running this set of first-wave applications on progressively larger versions of BG/L as it was being built provided powerful confirmation of the design and implementation of BlueGene's hardware and software architectures.

This coordinated "computer + applications" approach resulted in LLNL scientists winning the 2005 Gordon Bell Prize for best HPC performance. Using BG/L, the team achieved a sustained computation rate of more than 100 TF on the stockpile science domain decomposition molecular-dynamics code. The focus on applications continued in 2006, and new performance thresholds were crossed in both unclassified and classified computing for problems of importance to NNSA's national security mission. This year's effort culminated with the 2006 Gordon Bell Prize for the Qbox application's sustained rate of more than 200 TF on BG/L—double the previous record. (See Section 3.01.)

PROGRESS IN 2006

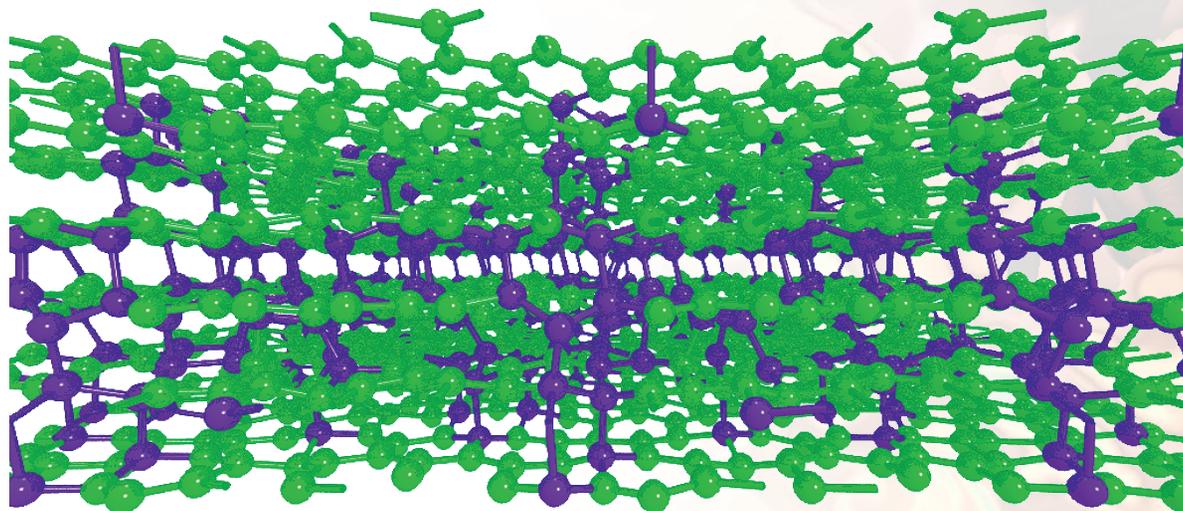
In early 2006, four ASC Program partners (University of Illinois, University of Chicago, Caltech, and Stanford University) received time allocations on BG/L, allowing them to run largest-ever simulations with their flagship application codes. The resulting data sets were treated as experimental results from a large instrument and were analyzed and mined over a period of many months, in one case (Chicago's Flash code) becoming the focus of a scientific workshop.

LLNL also provided a large BG/L allocation to one of our research partners, the IBM research center

in Zurich, Switzerland, to run the widely used Car-Parinello molecular dynamics code to simulate subjects of interest to LLNL, including high explosives, super-ionic water, and graphite-to-diamond experiments. The graphite-to-diamond simulation provided insight that confirmed the results of earlier experiments conducted at LLNL.

Personnel in LLNL's Defense and Nuclear Technologies Directorate used BG/L to perform the first instability simulation with a Reynolds number large enough to determine the nature of turbulence beyond the mixing transition,

The compression of graphite into diamond shows the low-strength P3 bonds (in green) being converted into the significantly higher-strength P4 bonds of diamond.





CONTACT INFORMATION

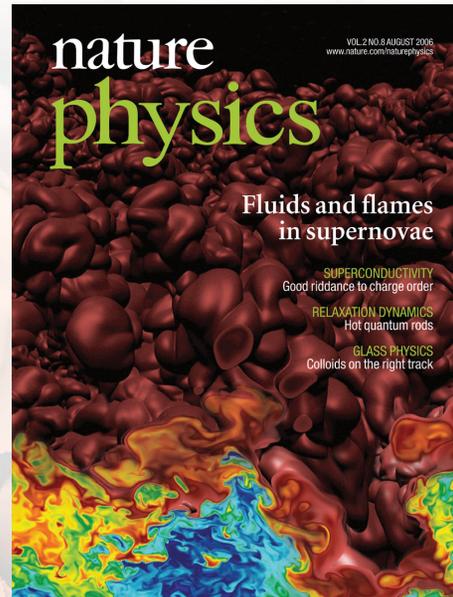
Dale Nielsen, dnielsen@llnl.gov

which warranted a *Nature Physics* cover article in August 2006.

Since March 2006 when it was certified for classified work, BG/L has continuously performed simulations designed to answer key questions about plutonium aging. The primary concern of plutonium aging is the damage that results from plutonium's natural nuclear decay process. This process occurs at a constant rate that will, over time, alter nuclear weapon performance. Because changes in performance shorten the service life of a stockpile weapon system, they must be well understood in order to prudently plan for timely

refurbishment or replacement, in accordance with national policy.

On November 29, 2006, then-NNSA Administrator Linton Brooks announced the results of a multiyear research program to which BG/L provided critical insight. "These studies show that the degradation of plutonium in our nuclear weapons will not affect warhead reliability for decades,"



The August 2006 *Nature Physics* cover highlights a BG/L simulation of turbulently mixing flow fields of a Rayleigh–Taylor instability.

said Brooks. "It is now clear that although plutonium aging contributes, other factors control the overall life expectancy of nuclear weapons systems."

BG/L simulations helped complete *ab initio* plutonium equation-of-state (EOS) calculations using smaller concentrations of self-irradiation-induced impurities than ever before. The simulations allowed more accurate estimates of the changes in EOS as a function of age. We also proved that BG/L is capable of performing such calculations at a sustained rate of more than 200 TF.

SIGNIFICANCE

While BG/L's most important programmatic deliverable during 2006 was its contribution to the plutonium aging program, its importance in determining a strategic direction for computing also advanced significantly during 2006. Exploratory efforts to port integrated weapons design codes to BG/L have shown that the BlueGene architecture is suitable for a much broader scope of applications than originally envisioned.

For simulations running on up to one-quarter of the full machine, the problem-wide interprocessor communication network in BG/L performs better than so-called full bisection bandwidth networks that serve as the measure for communication quality in modern, high-performance supercomputers. These smaller, one-quarter-sized problems on BG/L use more processors than are available on any other supercomputer, showing that smaller is relative where BG/L is concerned.

The pioneering, low-power approach to extreme performance that produced BG/L, developed in close partnership between LLNL and IBM, will help the United States maintain HPC pre-eminence in the next decade.

2.05 — User Environment Benefits from Productivity Improvements

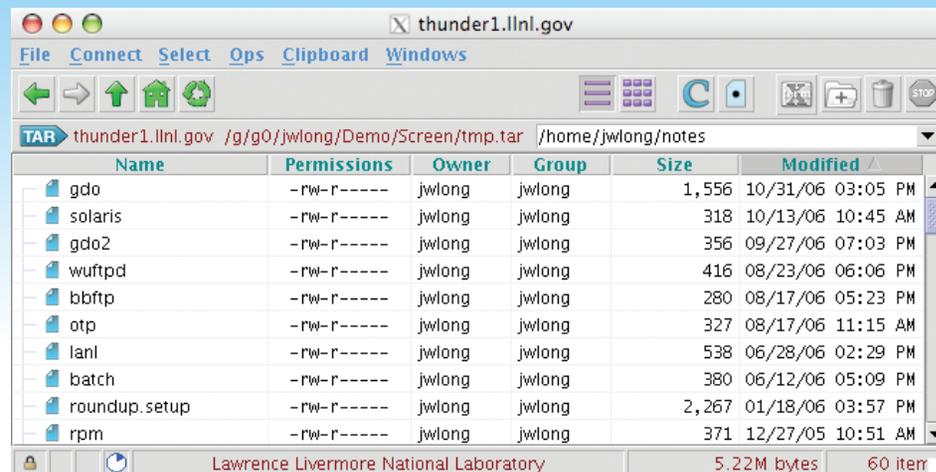
OVERVIEW

As LC pushes the frontiers in HPC, we must not let progress in the user environment lag. Regular deployments of high-performance computers substantially increase the complexity of the computing environment. Additional smaller computers are needed to provide support, and larger and more numerous file systems are needed to offload and manage the huge data sets that are produced on the high-performance machines. Users will see their productivity drop unless synchronous enhancements are made to the computing environment. LC has three strategies for improving user productivity: simplify the user's compute workflow, hide complexity, and force the environment to be as consistent as possible across the different computer architectures.

PROGRESS IN 2006

A few years ago, a typical user's compute workflow included moving data sets that were generated on the disk cache of each computer to the visualization server to analyze the data, while keeping a copy on the computer that generated the data. Today, the deployment of a parallel file system shared by compute platforms, visualization servers, and archival storage has greatly simplified the workflow. Users can now employ visualization servers to analyze data as it is generated, without moving the data sets and managing copies.

Additionally, if users need to continue the job on another computer that is connected to the same file system, they can restart the job without moving or copying the data set. LC uses the Lustre file system, from Cluster File Systems, as our solution. We deployed Lustre in 2003 on a single system. We now have unclassified computers sharing a Lustre file system with other computers, visualization servers, and the archival storage. Lustre has provided a much simpler data environment and has also resulted in cost savings. LC requires less hardware because the shared architecture allows multiple computers to leverage a common file system.



In this example of the Hopper graphical interface, a remote Hopper window is viewing a TAR file on a remote computer (Thunder) with the listing sorted by date.

In today's HPC environment, managing files can be a complicated and time-consuming task that distracts from scientific work. Users must deal with multiple file transfer protocols, repeated authentication, massive amounts of data, and other tedious but necessary tasks. Accomplishing these tasks requires several different tools, each with its own interface and idiosyncrasies. Our goal is to develop tools that help automate file management and substantially improve users' ability to transfer, organize, search, and operate on collections of files on multiple computers and storage systems located in LC and elsewhere.

One such tool is LC's Hopper, an easy-to-use graphical interface that handles multiple protocols, runs on nearly all platforms, and performs concurrent operations. Hopper helps users move, copy, find, delete, and otherwise operate on files between local and remote hosts or between remote hosts. The drag-and-drop interface is familiar to most users because it mimics interfaces used in popular Web applications. Hopper's popularity has grown since its introduction two years ago; in 2006, it was used between 1,000 and 2,000 times each month by more than 150 users. LC is currently developing a



CONTACT INFORMATION

Terri Quinn, quinn1@llnl.gov
Web sites: www.llnl.gov/hopper,
www.clusterresources.com/, www.lustre.org/

Jobs Moab Access Portal™ eval5@clusterresources.com

Refresh Logout

List Jobs Basic Submission Advanced Submission View Job

Show System Workload

Job ID	Job Name	State	User	Group	Account	Partition	Priority	Start Time	Submit Time	Est. Time Remaining	Processors	Nodes	Messages
14067		Idle	eva5	RedRock	Research	ALL	650		-10:29:49		16		
14051		Idle	eva5	RedRock	Research	ALL	688		-14:27:29		16		
14043		Idle	eva5	RedRock	Research	ALL	914		-14:53:50		16		
14027		Idle	eva5	RedRock	Research	ALL	1194		-19:33:52		16		
14019		Idle	eva5	RedRock	Research	ALL	1372		-22:32:07		16		
13977		Idle	eva5	RedRock	Research	ALL	2651		-46:21:05		32		
13905		Idle	eva5	RedRock	Research	ALL	4491		-74:31:06		32		
13602		Idle	eva5	RedRock	Research	ALL	7001		-116:20:33		16		

Powered by Moab Access Portal™ v1.2
 Copyright © 2001-2005 Cluster Resources, Inc. All Rights Reserved.

The Moab Access Portal empowers users with a Web-based job submission capability and an intuitive graphical interface to track and manage their own workload.



new client/server version, which will introduce a major usability feature—the ability to detach the client so that long-running jobs can continue even after the local workstation has been shut down.

LC helps users stay productive by maintaining consistent software environments regardless of the underlying hardware. We limit the types and versions of operating systems and maintain common access tools, common home

directories, common file transfer tools, and a common set of environment variables. We also provide users with a common resource management system. In 2006, LC led an effort with Los Alamos and Sandia national laboratories to implement a resource management system that gives users one tool for launching and managing their jobs from any of the three sites. The team selected the Moab Cluster Suite

from Cluster Resources, Inc., and all three laboratories announced their intentions to deploy Moab on the ASC computers. The Moab interface features self-help capabilities, such as start-time estimates and visual reservation tools.

SIGNIFICANCE

The modern user of HPC systems runs jobs on multiple computers with many processors and large memory capabilities. Jobs often take several days to complete and generate an extraordinary amount of data. Although the environment is increasingly complicated, LC helps ensure user productivity and allows users to make the most of these expensive scientific tools.

2.06 — Metrics to Mastery: Using Data to Achieve Excellence

OVERVIEW

The MCR supercomputer will be decommissioned in early 2007, after serving the M&IC large-scale computing program for nearly five years. MCR burst into the HPC spotlight in 2002 as the first Linux parallel-processing machine to achieve worldwide recognition by placing number 5 on the Top500 Supercomputers list. Why, then, will the machine be retired? Metrics and data, obtained from LC databases, show that hardware failures are occurring more frequently, maintenance costs are increasing, and the machine is available less often.

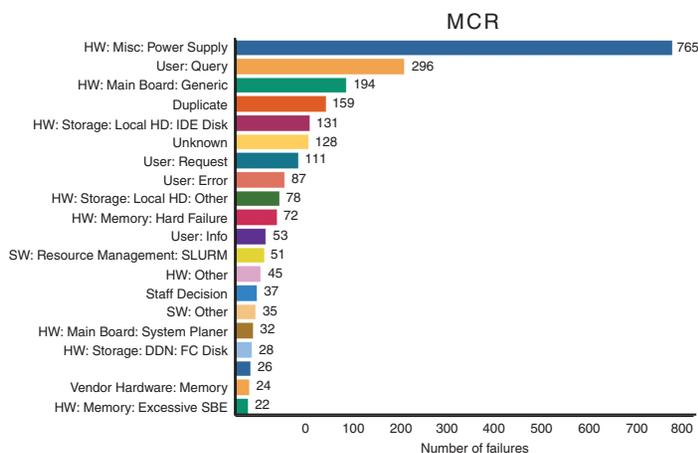
Volumes of computer-generated and human-entered data have been stashed away in various LC databases and archives waiting to be put to good use. In 2006, we mined important nuggets of this data to uncover emerging trends, identify potential problems, and alert us to potential system failures that were ultimately averted.

PROGRESS IN 2006

Managing LC's increasingly complex environment requires an analytical approach. We focus on

using our databases, especially the problem-tracking database, to help manage operations. Efforts are underway to improve the quality of the data by educating staff on the importance of the database and developing a common language

Hardware failures are a major factor in the decision to retire MCR.



CONTACT INFORMATION

Jean Shuler, jshuler@llnl.gov
 Barbara Herron, bherron@llnl.gov



to describe problems. In 2006, we began publishing reports for our staff to track problems by computer, by root cause, and by priority. LC management now tracks trouble-ticket resolution rate and time-to-resolution to monitor the organization's ability to respond to customers. LC personnel resolved 97% of the 19,000 trouble tickets submitted in 2006; 54% were resolved in fewer than five hours.

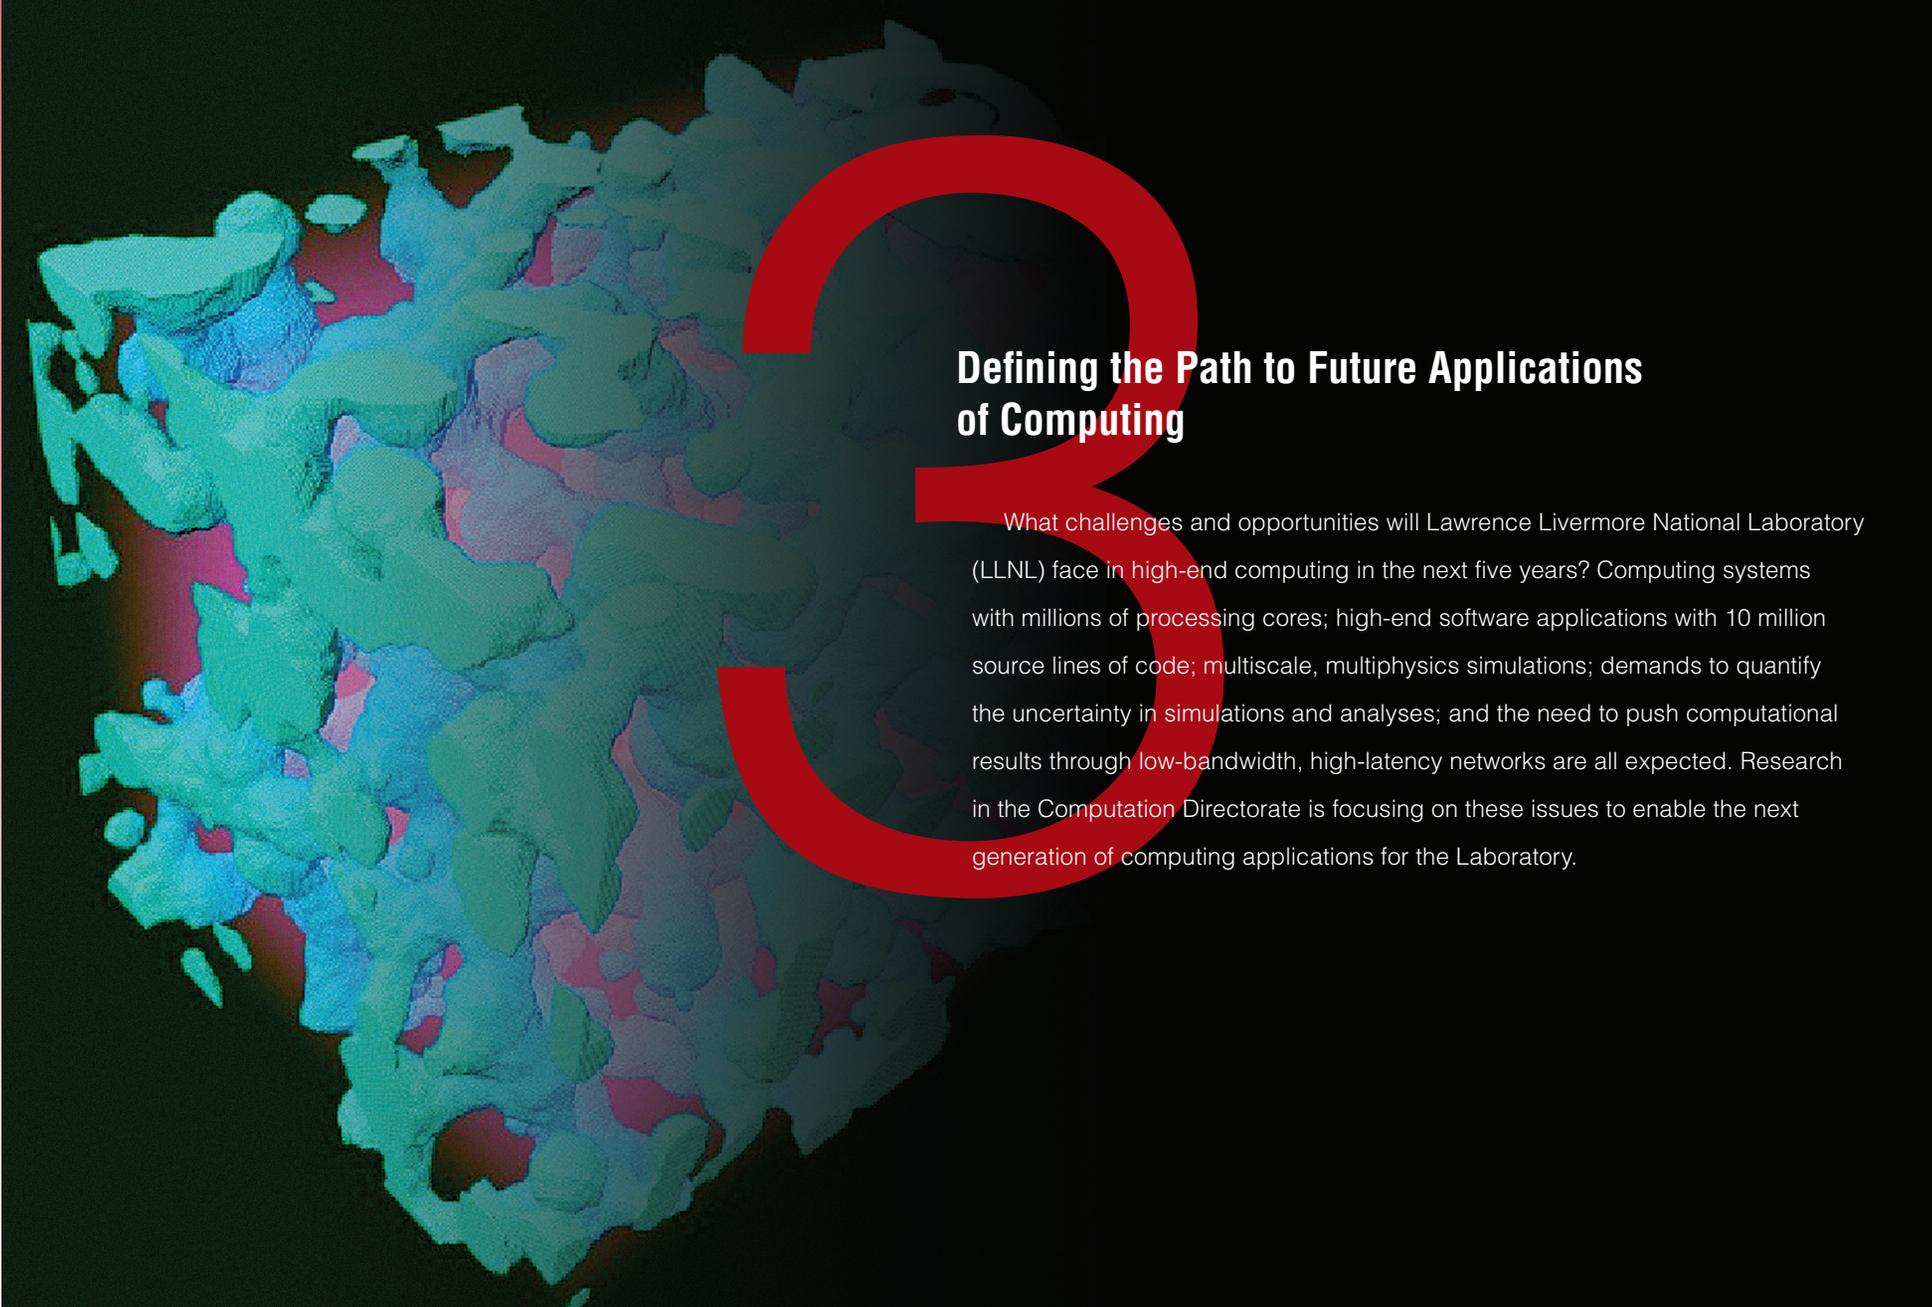
Knowledge gleaned from the computer accounting database helped significantly improve the utilization of an important HPC resource. In this case, machine-generated accounting records exposed the fact that the Thunder system was being underutilized and that there was an immediate need to improve the parallel computer job-scheduling parameters. The resource-scheduling algorithm was improperly set to favor large node-count jobs. We altered the algorithm to fully utilize the nodes, and the throughput of jobs increased immediately.

We also programmed the Remedy database to automatically generate a standard set of help desk reports. Remedy trouble-ticket reports revealed that, in

2006, LC's Customer Service Group responded to a significant increase in hotline requests for "best settings" of environment variables. Once we identified this as a hot topic among users, we developed an online environment-variable manual. The manual helped reduce the number of hotline calls by providing customers with round-the-clock access to information that helps them run codes more efficiently on the LC computers. Since its release in January 2006, users have printed nearly 600 copies. In November alone, it was viewed online 3,835 times.

SIGNIFICANCE

Delivering more than 600 TF of computing power to more than 3,000 scientists requires a business-like approach to managing the operation. This includes a strong reliance on metrics. In 2006, we increased our use of data and metrics to display trends, anticipate outcomes, and make decisions to help maximize productivity. We plan to develop more new metrics and methods of analysis that will help us better support our HPC customers.



Defining the Path to Future Applications of Computing

What challenges and opportunities will Lawrence Livermore National Laboratory (LLNL) face in high-end computing in the next five years? Computing systems with millions of processing cores; high-end software applications with 10 million source lines of code; multiscale, multiphysics simulations; demands to quantify the uncertainty in simulations and analyses; and the need to push computational results through low-bandwidth, high-latency networks are all expected. Research in the Computation Directorate is focusing on these issues to enable the next generation of computing applications for the Laboratory.

Defining the Path to Future Applications of Computing (cont.)



CONTACT INFORMATION

John Grosh, grosh1@llnl.gov

Computational science and engineering research in the directorate is focused on two major areas. The first is predictive simulation, where research is conducted to develop the tools, techniques, and methodologies for building scalable application codes that simulate complex phenomena with quantifiable accuracy. The second area is the data sciences, where we develop novel approaches for ingesting, processing, and understanding data to extract and synthesize actionable information and knowledge. Embedded in both areas is an emphasis on time-to-solution, where we balance maximizing the execution performance of applications software with our ability to develop and test large, complex software in a timely and cost-effective manner. This section highlights work that advances the directorate's agenda in these areas.

The Computation Directorate continues to demonstrate the leap-ahead capability of BlueGene/L (BG/L) by making enhancements to the molecular-dynamics code, Qbox. In addition to achieving unprecedented levels of performance and winning the 2006 Gordon Bell Prize for peak performance, Qbox was a significant programmatic simulation run on a system with more than 100,000 processing cores. This new modeling capability is being used to understand the behavior of high-Z (high atomic number) materials that are important to the stockpile stewardship effort. From a discipline perspective, simulations such as these are critical in driving the research and development of parallel programming tools and techniques, which can be applied to other simulations within the Laboratory.

Over the past decade, LLNL has moved towards simulations that involve multiple length and time scales and multiple physics components. Current approaches to developing multiscale, multiphysics software packages have been application-specific, with little reuse of approaches and software infrastructure. New high-performance computing (HPC) architectures, with large number of processors coupled with maturing component technology, offer an opportunity to glue numerous domain codes together as components and launch them across the massive

amounts of processors available on next-generation systems. From a software engineering perspective, this pushes the Laboratory towards extremely large and complex applications codes. To realize this concept, the Laboratory is developing a new multiple-program, multiple-data model and software libraries. The model offers a framework for developers to scale multicomponent applications with a reasonable amount of effort. The framework schedules coordination of data transfer and execution of the components. An example of a new application enabled by this research is the coupling of ALE3D, a multiphysics hydrocode developed at LLNL, to a subfield model.

One of the major challenges in predictive simulation, specifically for the Advanced Simulation and Computing (ASC) Program, is uncertainty quantification. Typically, given exact input parameters and boundary conditions, application codes will simulate phenomena and provide a precise output. Unfortunately, errors in measurement and representation produce uncertainties in simulations that require quantification, especially if they are to be used in weapons design or policy decisions. The PSUADE code was developed to address this challenge.

For LLNL's many highly complex simulation codes, it is critical to develop approaches that reduce coding complexity, improve code correctness, and ease the burden on the software developer while achieving high levels of execution performance. Several of our projects, such as the Babel library and ROSE compiler tool, address these challenges.

The directorate also supports important work in data analysis and visualization under constrained environments. We are applying techniques originally developed within the ASC Program for visualization of large-scale data sets to problems in the intelligence, reconnaissance, and surveillance domains. Specifically, we are adapting these techniques and developing innovative approaches to data compression, analysis, and visualization using new algorithms for storage and new polygonal representations for images. The intelligence community will benefit from the development and application of these techniques.

3.01 — Qbox: Understanding High-Z Metals at Record Speed



CONTACT INFORMATION

Erik Draeger, draeger1@llnl.gov

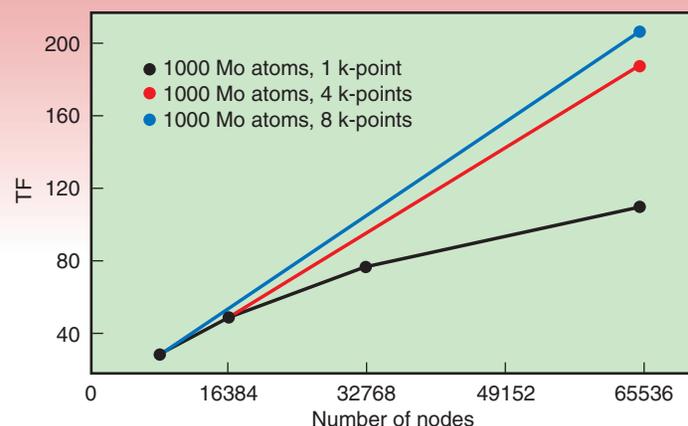
OVERVIEW

We continued to develop and optimize Qbox, a first-principles molecular-dynamics (FPMD) code that was developed at the Center for Applied Scientific Computing (CASC) to enable the large, predictive simulations of high-Z materials that are needed for stockpile stewardship. We achieved a record-breaking performance on the full 65,536-node BG/L supercomputer by extending the capabilities of the code to allow complex metallic systems to be represented as accurately as possible while simultaneously building on previous optimization efforts. This accomplishment earned the Qbox team the 2006 Gordon Bell Prize for peak performance.

PROGRESS IN 2006

Predicting the properties of metals under extreme temperature and pressure conditions has been a long-standing challenge for researchers in materials science and high-energy-density physics. FPMD is an approach that enables predictive materials simulations, without a priori knowledge of a given system of atoms, through a quantum-mechanical description of electrons. However, the method is computationally expensive, particularly in the case of complex metallic systems, which require a large number of atoms and high numerical accuracy to represent the electrons correctly.

We expanded Qbox to simultaneously calculate multiple Brillouin zone points (or k-points), each of which requires a separate solution of the quantum-mechanical equations for the electrons, thus providing the required accuracy. A three-dimensional (3D) grid of k-points must be fine enough to ensure convergence, although symmetry-equivalent points can be combined for efficiency. The plane-wave basis of the electrons must also be converged; typical high-Z metals require millions of plane waves for each k-point. For systems that require a large number of atoms to minimize finite-size effects, the combination of these factors can result in an extremely large problem size. Therefore, efficient data layout and careful data-flow management



Strong scaling results for Qbox simulations of 1,000 molybdenum atoms on BG/L. Blue and red lines indicate a connection between multiple k-point calculations and the corresponding single k-point (black line) results and demonstrate the importance of performance optimization across a large range of node numbers.

are critical to achieving performance on massively parallel machines like BG/L.

In addition to Qbox's new ability to exploit the natural parallelism of multiple k-points, we made several other enhancements that contributed to the 207.3-trillion-floating-point-operations-per-second (TF) record-setting sustained performance. For instance, we optimized the complex number-matrix multiplication that is required for most individual k-point calculations. We also attained a better understanding of the application's communication performance on BG/L and made improvements in task layouts. Finally, we optimized the use of the message-passing interface within the ScaLAPACK library.

SIGNIFICANCE

The ability to perform predictive simulations of metals under extreme conditions has significant implications for stockpile stewardship and computational materials science. The latest Qbox results represent a new record in sustained application performance and a leap forward in our capability to perform predictive materials simulations of large, complex high-Z systems.

3.02 — AMG Scalability Improvements Lead to Faster Simulations



OVERVIEW

The Scalable Linear Solvers project develops algorithms and software for solving large, sparse systems of linear equations on massively parallel machines. The primary software product is *hypr*, a collection of multigrid methods that is used extensively in simulation codes at LLNL and other research institutions. Multigrid solvers are ideally suited for large-scale computing because their convergence rate is optimally independent of the problem size. The algebraic multigrid method (AMG) is one of the most efficient algorithms for solving large, unstructured-grid problems because it does not require knowledge of the problem geometry. Although many AMG components can be parallelized with ease, others require more innovation. The recent advent of machines with hundreds of thousands of processors, such as BG/L, has presented new challenges to obtaining scalability and has forced the re-examination of parallel codes that were once suitable for fewer processors. Three major modifications were made to *BoomerAMG*, *hypr*'s AMG code, that significantly improve scalability on a large number of processors: a reduced-storage description of the data distribution; new, low-complexity coarsening algorithms; and long-distance interpolation schemes.

PROGRESS IN 2006

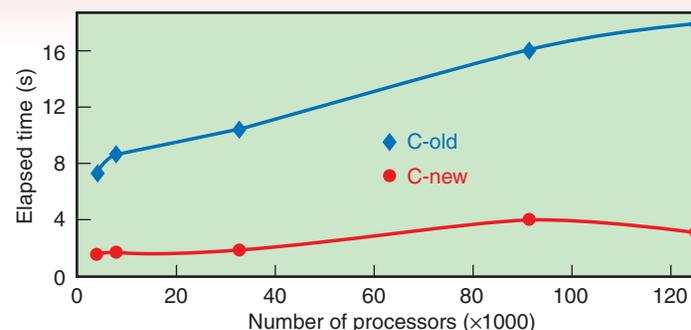
AMG (or any linear solver) requires processors to retrieve nearby data from other processors to complete the solution. However, on a machine with a large number of processors, storing and querying a global description of the data is either too costly or, in some cases, impossible. Instead, we developed a new algorithm that employs an assumed partition to answer queries through a type of “rendezvous” algorithm. This method significantly reduces storage and computational costs.

As problem sizes increase, traditional AMG-coarsening schemes frequently suffer from the growth in computational complexity. The increase in memory usage and execution time reduces scalability, particularly in 3D problems. To overcome this complication, we developed new, low-complexity parallel-coarsening techniques that are based solely on enforcing a maximum independent-set property, resulting in sparser grids than the traditional heuristics allow. We reduced the complexity further by applying an aggressive coarsening strategy and achieved scalable results for up to 125,000 processors on BG/L.

CONTACT INFORMATION

Ulrike Meier Yang, umyang@llnl.gov
 Allison Baker, abaker@llnl.gov
 Web site: www.llnl.gov/CASC/linear_solvers

Algorithms	Global partition		Assumed partition (new)	
	C-old	C-new	C-old	C-new
No. of processors				
4,096	12.42	3.06	12.32	2.86
64,000	67.19	10.45	19.85	4.23



Total runtimes (in seconds) when AMG is applied to a 3D, 7-point Laplace problem on BG/L using different coarsening algorithms (C-old: traditional coarsening algorithm; C-new: new, low-complexity coarsening scheme with aggressive coarsening strategy). The results in the table were achieved in co-processor mode using 30x30x30 points per processor; the results in the graph were obtained using virtual mode with 25x25x25 points per processor.

When the new coarsening methods are applied to more complex problems, such as changes in material properties or rotated anisotropies, convergence scalability degrades. For these problems, traditional interpolation methods that are based solely on immediate neighbors are not sufficiently accurate. Therefore, we developed and implemented long-distance interpolation schemes that use information from points that are farther away than the nearest neighbors. Combining the new interpolation and coarsening algorithms restored the desirable convergence factors of traditional AMG schemes while retaining the memory reduction that was

previously achieved. The end result is a scalable AMG method that runs up to six times faster for some problems.

SIGNIFICANCE

Linear solvers are a critical, time-consuming component of many application codes. Several simulation codes at LLNL will benefit from the improvements in scalability and memory usage that *BoomerAMG* offers. The AMG improvements allow for faster and larger simulations of physical phenomena, including radiation transport, elasticity, and magnetohydrodynamics.

3.03 — Quantifying Uncertainties in Multiphysics Applications



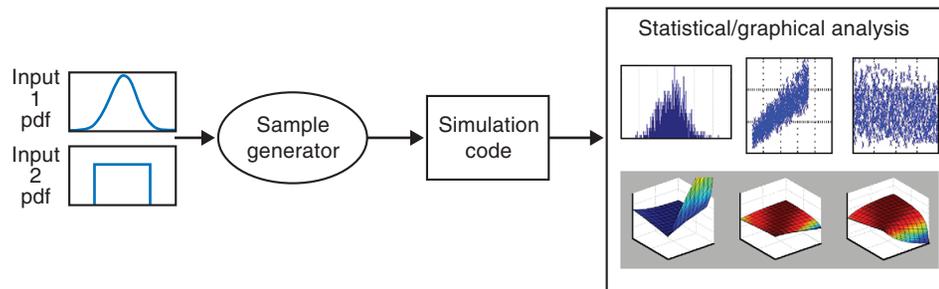
CONTACT INFORMATION

Charles Tong, chtong@llnl.gov

OVERVIEW

One of the challenges associated with building a credible predictive simulation capability for scientific applications is to accurately characterize all the relevant physics that are involved and their interactions. In practice, it is impossible to know everything about a physical system, which means there are many uncertainties in the corresponding simulation models. Therefore, it is important to quantify these uncertainties to establish the integrity of the simulation models.

PSUADE is a problem-solving environment developed for performing uncertainty quantification (UQ) and sensitivity analysis (SA) on simulation codes in a nonintrusive, or “black box,” manner. PSUADE was initially developed to implement a global sensitivity analysis methodology for multiphysics applications that have a large number of uncertain parameters. Three basic modules make up the problem-solving environment: a collection of sampling strategies, an automated job-launching environment on the ASC computer platforms, and an enriched set of deterministic and statistical analysis tools. PSUADE has been applied to various UQ/SA activities at the Laboratory, including nuclear cross-section studies in the Physics and Advanced Technologies Directorate, hydrodynamic instability and material property studies in the Defense and Nuclear Technologies (DNT) Directorate, verification and validation (V&V) efforts in the Weapons Engineering Program, and optimal sensor-placement analysis in a multidirectorate Laboratory Directed Research and Development (LDRD) project.



The PSUADE execution model consists of three stages: sampling design, job launching and data collection, and analysis.

PROGRESS IN 2006

In 2006, we contributed to the completion of two ASC Level 2 V&V milestones for DNT. Both milestones investigated parameter sensitivities of several DNT applications. We implemented and applied a global sensitivity analysis methodology via the PSUADE framework, which consists of three stages: First, uncertain parameter ranges are identified using expert judgment; literature searches; or, for multiphysics applications and given experimental data, a recursive application of the three-stage methodology. Second, screening methods are used that do not rely on assumptions (for example, linearity) about the behavior of the system that is being studied. Third, exotic sampling methods, various response surface methodologies, and variance decomposition techniques are exploited to quantify parameter sensitivities. Graphical techniques are routinely used to provide further insights into the model behaviors.

A complete sensitivity analysis of our selected simulation models

involved up to 64 uncertain parameters and encompassed uncertainties in hydrodynamics, material strengths, and radiation transport, among others. The studies consumed more than one million processor hours on ASC Purple.

We also made enhancements to PSUADE in 2006. We designed and implemented the ability to handle arbitrary input and output constraints, multiresolution sampling and analysis capabilities for improving robustness, and multilevel numerical optimization capabilities.

SIGNIFICANCE

The success of the ASC Level 2 milestones has generated interest from several Laboratory scientists who are now considering using UQ tools for their investigations. Our success charted new territory in sensitivity studies for large-scale multiphysics models using petascale computers. We are planning a seminar series to introduce these methods to LLNL scientists and engineers, and we will continue to collaborate with scientists to meet their UQ needs.

3.04 — ALE-AMR Brings Adaptivity to Multiscale Applications

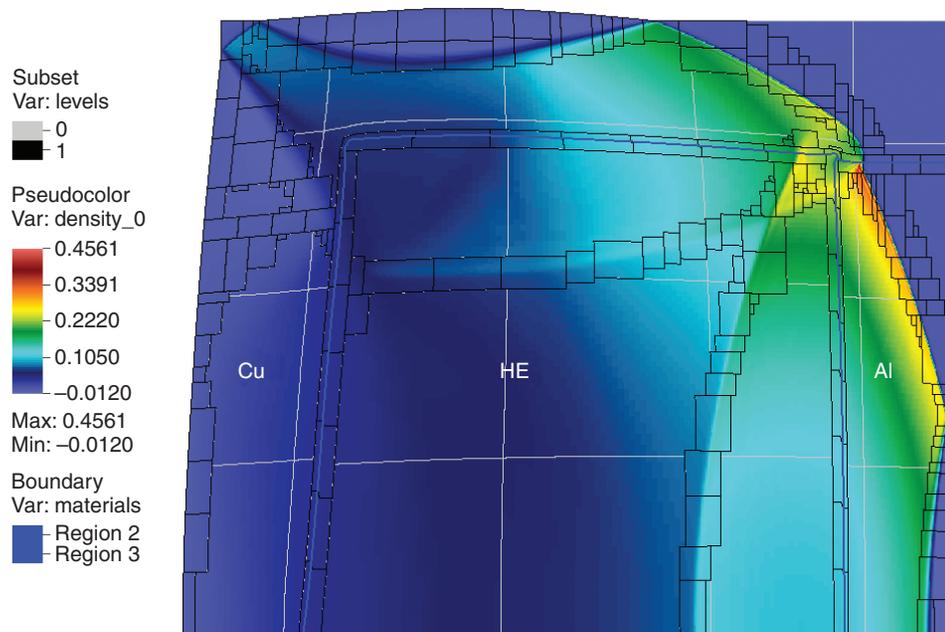


CONTACT INFORMATION

Robert Anderson, anderson110@llnl.gov
 Alice Koniges, koniges@llnl.gov

OVERVIEW

New numerical methods and suites of codes are being developed that use adaptive mesh refinement (AMR) combined with the more standard LLNL methods based on Arbitrary Lagrangian-Eulerian (ALE) hydrodynamics. The applications of this research range from high-energy-density physics to National Ignition Facility (NIF) target-debris modeling. This unique combination of methods allows dynamic tracking of phenomena such as particles and wave fronts using Lagrangian methods while maintaining an accurate and efficient mesh description based on an underlying AMR framework. The codes are already parallel because both are built using the scalable Structured AMR Application Interface (SAMRAI) framework that was developed in CASC.



Adaptive mesh refinement is used to simulate the detonation of a high explosive in a metal container.

PROGRESS IN 2006

The integration of an existing multimaterial, high-explosives code and the SAMRAI AMR library was completed in 2006 for use in the Lagrangian mode, where the computational elements remain fixed with the material as the material moves in space. There were significant software engineering challenges to overcome in the code-development phase of the project. The difficulties arose when an existing physics code that was not designed for use in an AMR context was adapted for just such a use. We designed the integration of the physics code with AMR to cause as little disruption in the original code as possible so that the new AMR infrastructure could exist outside the core physics code. In this way, code physicists can continue to refine the physics of the algorithms without being required to become AMR experts.

The ALE-AMR technology is being applied to a new stand-alone modeling code that is designed to predict the behavior of debris and shrapnel generated by NIF experiments. Significant progress was made on the two-dimensional version of the code,

including multimaterial interface reconstruction, material failure models, and radiation diffusion in Lagrangian mode. This code will enable the evaluation of NIF targets and the effects of target debris and shrapnel on the diagnostics and optics in the target chamber.

SIGNIFICANCE

Adding an AMR capability allows for higher fidelity solutions at significantly lower computational costs by focusing computational resources where they are required. In the case of NIF chamber simulations, the unique gridding capability the method affords will enable us to model an enormous multiscale problem, where the scale lengths range from targets of less than a centimeter to chamber walls holding optics that lie meters away. Additionally, the availability of AMR levels will allow us to introduce different material models at different scales of the problem. ALE codes have myriad applications, including structural analysis and vehicle efficiency. As simulations grow in sophistication, it will be increasingly common to find phenomena occurring simultaneously at dramatically different physical scales. In these settings, using the integrated ALE-AMR approach will pay off handsomely.

3.05 — Wave Propagation Project Models 1906 Earthquake



OVERVIEW

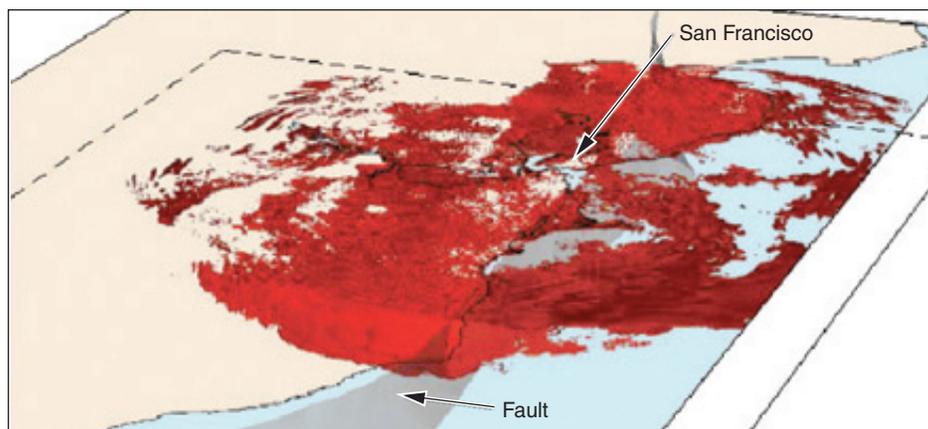
The Wave Propagation Project (WPP) is developing a new 3D elastic wave propagation code that can be used for modeling in various applications, such as the prediction of ground motion during earthquakes, nondestructive evaluation of materials, underground facility detection, acoustic noise generation, and oil and gas exploration.

We developed capabilities for modeling large earthquakes so that we could participate in an ambitious multi-institutional effort coordinated by the U.S. Geological Survey (USGS) to study the 1906 San Francisco earthquake. Our goals were to re-examine historical data to better understand the 3D structure of the earth's crust in northern California, explore the source characteristics of the fault rupture, and use the information to simulate the resulting ground motion.

PROGRESS IN 2006

Historical data indicate that the 1906 earthquake rupture began offshore near San Francisco and propagated along a 480-km stretch of the San Andreas Fault, from the Mendocino triple junction (near Eureka) to Hollister. One of the mysteries of this earthquake is why the Santa Rosa area experienced more severe shaking than San Francisco even though it is located much farther from the San Andreas Fault. Seismologists have determined that the fault rupture speed has a strong influence on the ground motion, especially when the rupture speed

approaches the shear velocity in the earth. One of our goals was to explore how the rupture velocity changes the ground motion response farther from the fault.



Snapshot of a secondary (S-) wave propagation through the earth as visualized by an isosurface of the curl of the velocity field. The red regions near Santa Rosa (bottom left) indicate strong shearing ground motions, which can cause significant structural damage.

CONTACT INFORMATION

Anders Petersson, andersp@lnl.gov

Arthur Rodgers (Energy and Environment Directorate), rodgers7@lnl.gov

Web site: www.lnl.gov/casc/serpentine/

The WPP team evaluated several rupture scenarios by calculating ground motions over a 550x200-km area of northern California, measuring from the earth's surface to a depth of 40 km. To perform these calculations, we wrote software that queried the USGS high-resolution geologic/seismic-material database on a massively parallel machine. We also developed an accurate technique to place earthquake sources on the 3D shape of the San Andreas surface. Once we accomplished these tasks, we partnered with visualization experts in the Computation

Directorate to generate a 3D seismic wavefront movie. The simulation used to create the movie captured the first 300 seconds of shaking, generated 2.2 TB of data, and required approximately 40 hours on 1,024 processors. We presented our results at the Seismological Society of America annual meeting in April, which commemorated the centennial of the earthquake.

SIGNIFICANCE

The 1906 calculation was the first large-scale application of the WPP code. Our next step is to improve the efficiency of WPP by implementing a mesh refinement technique that will allow shorter waves to be resolved near the surface. Capturing the shorter, higher frequency waves will enable us to couple the ground motion simulation to structural simulations, which can then be used to evaluate the seismic impact on bridges, dams, levees, and large buildings. We are also implementing an embedded boundary technique to incorporate topography, which will further improve the realism in our simulations.

3.06 — Babel Sets the Standard for High-Performance Middleware



CONTACT INFORMATION

Gary Kumfert, kumfert@llnl.gov

Web site: www.llnl.gov/CASC/components

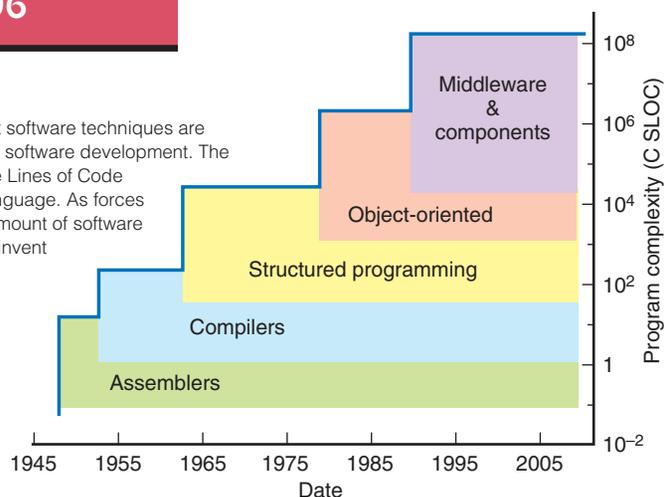
OVERVIEW

Middleware is a class of software infrastructure that is needed for extremely large or complex software development projects. This extreme scale of software is often called “enterprise software” because it exceeds the capacity of a single team or business unit to produce. Instead, the software is a shared product that requires the commitment and coordination of multiple teams across the enterprise. Several middleware technologies exist for industry and finance, and their efficacy in improving productivity and reducing the cost of developing enterprise software is well established. In scientific computing, however, the effective use of middleware has been hampered because standard middleware offerings do not meet the performance demands of HPC.

Babel was invented at LLNL to fill the need for HPC middleware. Using Babel, library vendors can support customers in many programming languages. Communities are designing discipline-specific application programming interfaces and using Babel to deal with the implementation-specific details. Framework designers are using our middleware to make their systems more flexible. Babel is also the basis for new component-based software systems that are being researched.

PROGRESS IN 2006

The right tool for the job: Different software techniques are appropriate for different scales of software development. The vertical axis is normalized Source Lines of Code (SLOC) in the C programming language. As forces continually push the limit of the amount of software that can be produced, scientists invent new technologies to increase the scalability of the software development process.



During an invited talk at the Scientific Discovery through Advanced Computing (SciDAC) Program 2006 conference, we presented 20 success stories of Babel being used in HPC. Application areas included radio astronomy, fusion, accelerator beam dynamics, materials science, quantum chemistry, cell biology, subsurface transport, sparse linear algebra, linear solvers, meshing, performance monitoring, and source code refactoring. Babel was even used in a nuclear power plant training simulator developed in Spain.

Babel is a foundational element of the SciDAC center that is developing the Common Component Architecture (CCA), a component standard for HPC. CCA promotes interoperability between components developed by teams at different institutions. Babel is also part of a strategic application partnership with a SciDAC fusion application to assist in developing a next-generation multiscale Tokamak simulation framework.

Babel earned one of LLNL’s seven R&D 100 Awards in 2006 for delivering superior performance over industry middleware alternatives. Babel achieves this by mixing languages the hard way—in the single address space of compiled code rather

than relying on network-based schemes or interpreters to broker between languages. Babel was also featured prominently in two LLNL LDRD projects in 2006. One, the Cooperative Parallelism Project (see Section 3.10), contributed significant resources to accelerate the development of Babel’s distributed computing capabilities.

SIGNIFICANCE

Babel is the leader in high-performance middleware, and we expect this to be a growth market for HPC. Software component technology and middleware is useful in many contexts, but it is absolutely essential for very large-scale, multidisciplinary development efforts. Our primary success metric is Babel’s impact on applications. In addition to more customers, downloads, e-mail list subscribers, and science publications citing Babel, we have seen a significant shift in our customer base. Babel is being taken into account more often when application teams consider overhauling the design of their codes. Babel is becoming less of an optional feature and more of a critical dependency because application teams see the benefits that middleware presents to their long-term productivity.

3.07 — ROSE: Source-to-Source Optimization of Scientific Software



CONTACT INFORMATION

Dan Quinlan, dquinlan@llnl.gov
 Web site: www.llnl.gov/casc/rose/

OVERVIEW

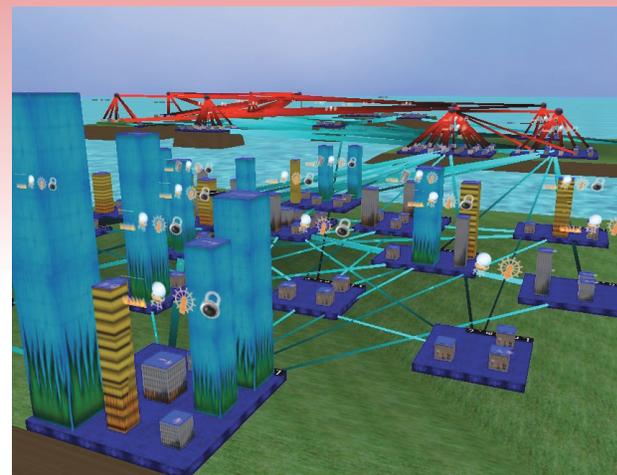
The ROSE project addresses the increasing complexity of developing and optimizing software for large-scale scientific software applications that target state-of-the-art parallel computer architectures. Compiler-based tools automate the introduction of performance optimizations. ROSE, a new compiler technology, makes it easy to develop robust source-to-source translators for C, C++, and Fortran scientific applications. Our focus is on developing domain-specific optimizations that are not addressed by common vendor compilers and making them available to Department of Energy (DOE) applications. The ROSE project also contributes to research on the optimization of high-level abstractions, with the goal of significantly raising or customizing the level of abstraction at which scientific applications are developed. ROSE provides a common set of optimizations that are oriented toward existing legacy applications. In effect, ROSE permits an application that uses a library to be optimized as a domain-specific language.

PROGRESS IN 2006

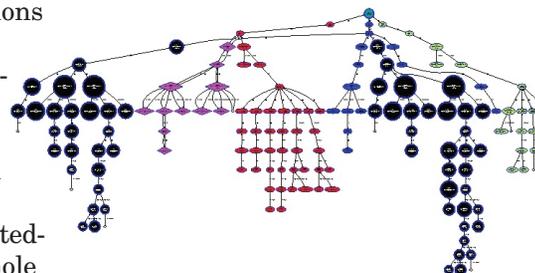
Our research during the past year focused on increasing ROSE's robustness and utility to support LLNL's optimization research and collaborations with more than two dozen research groups around the world. ROSE is regularly tested on six different million-line C and C++ applications and a commercial compiler test suite. In the case of C++, the applications are among the most complex scientific C++ codes within DOE. We also worked with Rice University to develop Fortran support.

Our focus on optimization led us to use aggressive loop optimizations for recent automated-tuning research that targets whole applications. This technique was demonstrated on the SMG2000 benchmark and showed a 50% performance improvement by empirical evaluation of more than 400 different types of loop optimizations. The University of Tennessee used ROSE on similar work and improved performance on a DOE SciDAC application by 300%. In collaboration with the University of Houston, ROSE was used to develop a parallel OpenMP compiler.

We also added program visualization capabilities to better



Postdoctorate researcher Thomas Panas's novel program visualization research is represented in this visualization of two directories of a larger-scale LLNL application. The dark green represents the two directories, the blue squares represent the files in the directories, and the bars (appearing as city buildings) represent the functions. The height of the bars indicates the size of the functions, and the colors and shapes are used to encode additional properties that are specific to either performance or software engineering details. The lines between the functions indicate the call graph within the program.



A simple example of the program-analysis capability to read multiple project files, merge the common parts as required to support large-scale million-line applications, and perform custom analysis to evaluate global properties of the whole program.

understand complex software and communicate security flaws and general program-analysis results to application developers. Most recently, the team focused on whole-program representation to support global analysis and optimization of extremely large-scale scientific applications.

SIGNIFICANCE

ROSE provides powerful mechanisms to support custom optimization and analysis tools that operate automatically on the source code of large-scale DOE software projects. Our research supports the development of domain-specific analysis and optimizations that are required to improve productivity and robustness in the development of scientific software for unique DOE state-of-the-art parallel computer architectures. We are also conducting research in telescoping languages, which provides domain-specific compiler support for user-defined abstractions in existing languages as an alternative to the expensive development and time-consuming standardization of new languages.

3.08 — Gaining Insight from Large-Scale Dynamic Semantic Graphs



CONTACT INFORMATION

Tina Eliassi-Rad, eliassirad1@lml.gov

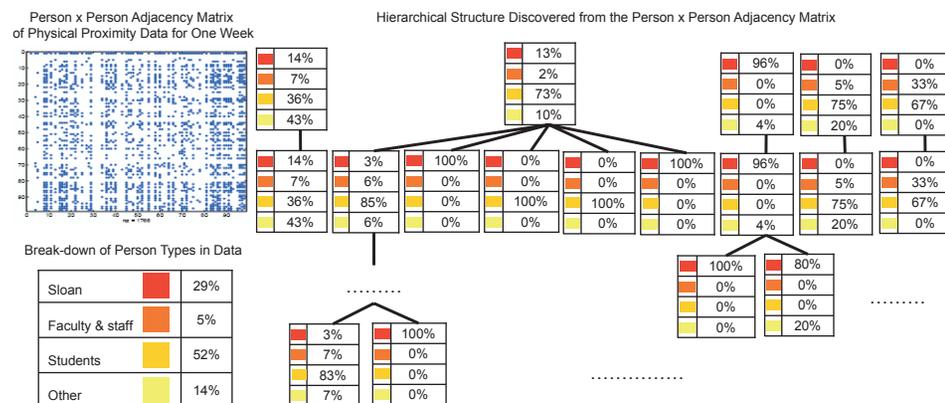
OVERVIEW

Dynamic semantic graphs are directed graphs that evolve over time. The term “semantic” refers to the fact that such graphs contain multiple concepts and relationships. Dynamic semantic graphs provide one way to represent large amounts of data from multiple sources. Our goal is to develop efficient algorithms that provide insights into large-scale dynamic semantic graphs. Specifically, we focused on approaches that disambiguate redundant nodes, detect communities or groups, and characterize the graph based on semantic graph-theoretic measures. The underlying theme of our work is the similarity between any two semantic graphs, especially when no a priori output is available.

PROGRESS IN 2006

In 2006, we addressed the disambiguation problem, which consists of detecting and consolidating redundant data. Disambiguation is essential

for maintaining data quality. Because disambiguation is a form of inexact pattern matching, we produced an extensive survey on state-of-the-art technology for graph-based pattern matching. Our survey was extremely well-received by our sponsors and the academic community. Next, we



Our probabilistic model discovered the above hierarchical structure in people based on their proximity (via Bluetooth connections) during one week. Data source: <http://reality.media.mit.edu/>.

demonstrated the advantages of applying consecutive filtering functions—ranging from low to high computational costs—to detect redundant data. We produced results with 99% precision and recall rates on the Digital Bibliography and Library project bibliographic data with three stages of filtering: blocking, attribute matching, and structure matching. By using simple binning techniques for blocking, we were able to reduce the set of potential matches to 1% of the original data. This work was performed as part of our participation in the Research Industrial Projects for Students program at the UCLA Institute for Pure and Applied Mathematics, in which we mentored four outstanding undergraduate students from institutions across the country.

We also investigated various techniques for detecting communities or groups in dynamic semantic graphs. In particular, we compared approaches based on probability theory, compression theory, and network flow theory. Despite their theoretical differences, we saw significant coherence among their resultant communities. Currently, we are examining hybrid approaches that

take advantage of the speed of compression-based approaches, the robustness of probability-based approaches, and the structural information of network flow approaches.

Finally, we demonstrated that standard graph measures are not sufficient for capturing local and global properties of dynamic semantic graphs. This result led to the development of a suite of new measures for such graphs. We illustrated the usefulness of these metrics in visual analysis of large semantic graphs by incorporating them into a system called OntoVis, which was developed in collaboration with colleagues at UC Davis.

SIGNIFICANCE

Data is only valuable when insight can be gained from it. Our efforts—funded by the Department of Homeland Security and the Predictive Knowledge Systems Strategic Initiative—have been instrumental in demonstrating the benefits of going beyond search and visualization to devise algorithms that capture similarities with respect to semantics and structure in dynamic semantic graphs.

3.09 — Image Viewer Analyzes Streaming Data in Real Time



CONTACT INFORMATION

John R. Johnson, jjohnson@llnl.gov

OVERVIEW

The national security community faces a challenge to analyze massive volumes of imagery and associated spatiotemporal data. With advancements in acquisition platforms and computing technologies, the generation of images is increasing exponentially while the capability to interact with data and extract knowledge is not.

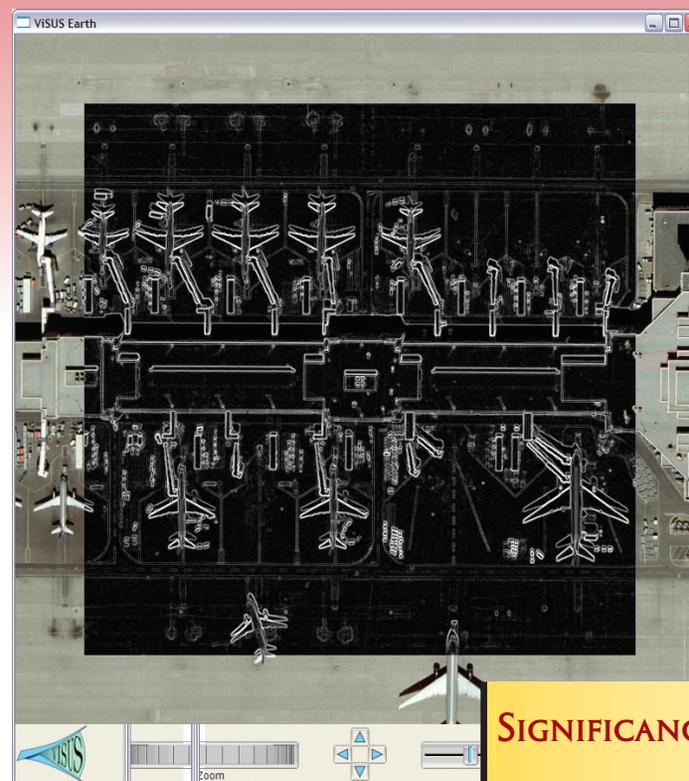
The Novel Architecture and Progressive Algorithms (NAPA) project is using progressive algorithm techniques to advance the real-time analysis of massive image data sets. Progressive algorithms subdivide input data into a hierarchy, from coarse to fine, where finer levels of detail are dynamically constructed to improve the result. The algorithm can present partial (coarse) results while the complete hierarchy is being developed. Data is processed as it is read, which is known as stream processing. Good data layouts are critical to efficient stream processing. NAPA uses cache-oblivious layouts that don't optimize for a specific block size but are beneficial in any block-based data transfer (i.e., disk-to-memory or remote network).

PROGRESS IN 2006

The NAPA team implemented a progressive image viewer to demonstrate progressive algorithm techniques on imagery stored in cache-oblivious layouts. The progressive image viewer provides interactive, real-time processing of multiple terabytes of image data over bandwidth-limited networks, including wireless networks, and scales to hand-held Windows PCs and even to the Apple iPod.

The NAPA viewer allows both browsing and interactive processing, which makes it more powerful than popular image-

navigation tools, like Google Earth. Progressive techniques allow significantly larger data sets and higher resolution imagery than existing commercial tools, and the progressive techniques enable interactive processing. The user can draw a window over an image and apply an algorithm (e.g., edge detection) to that region. Sliding the window over the imagery reapplies the algorithm to the new region in real time. If the machine has a programmable graphics processing unit (GPU), the image-processing algorithms are pipelined to the GPU, further increasing the NAPA viewer's performance. Data can be viewed and processed temporally and spatially.



Results of the NAPA progressive image viewer with a real-time edge-detection filter being applied to an image of an airport.

SIGNIFICANCE

The NAPA viewer also allows spatial and temporal region-based queries that can overlay arbitrary data represented in Keyhole Markup Language, a standard used by Google Earth. This feature makes the NAPA viewer interoperable with third-party software that uses Google Earth and demonstrates how progressive techniques can be used to facilitate multimodal data fusion in a geospatial/temporal context.

The NAPA project significantly advanced the ability to interact with large repositories of image data and for multiple modalities of data to be combined in resource-constrained environments. By demonstrating operation over wireless channels and scalability to hand-held devices, the NAPA team proved that progressive techniques can deliver information to individuals both on a workstation and in the field. These techniques benefit the defense and intelligence communities and are applicable to other disciplines, such as medical imaging.

3.10 — Cooperative Parallelism Paves the Way for Petascale Computing



OVERVIEW

Future supercomputers will gain more performance from increased parallelism than from faster processors. To use these machines effectively, application developers must find ways to exploit massive parallelism. At present, most parallel applications use a single-program, multiple-data (SPMD) model, in which each processor simultaneously applies the same algorithm to a different portion of the workload. This approach is conceptually simple, but balancing the workload to maximize efficiency can be tricky, especially if the workload changes as the computation proceeds. Moreover, load balancing becomes more difficult as parallelism increases.

The Cooperative Parallelism Project (initiated in 2004 as the LDRD Petascale Simulation Initiative) developed a programming model that lets applications manage and distribute work more flexibly. It is a multiple-program, multiple-data model, meaning that different parts of a parallel application work concurrently on different tasks running different executables. Cooperative parallelism is designed to complement the SPMD model, so existing parallel applications can be combined or augmented to form larger federations.

PROGRESS IN 2006

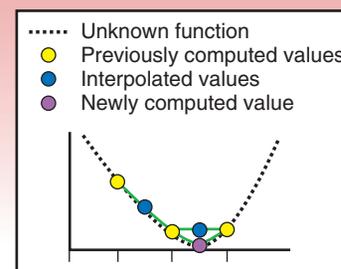
In collaboration with team members from the Engineering and Chemistry, Materials, and Life Sciences directorates, we adapted the ALE3D materials-modeling code for use in a cooperative parallelism federation. We also developed and implemented a technique called “adaptive sampling” that can dramatically improve the efficiency of multiscale modeling.

Multiscale applications use a coarse-scale model to compute the physical behavior of a

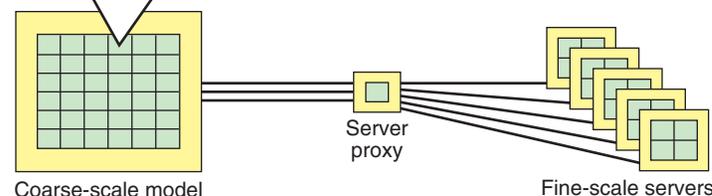
system over the entire problem domain, and they supplement these computations in certain regions with data from a more computationally intensive fine-scale model. With adaptive sampling, an application saves the fine-scale data it computes and then interpolates or extrapolates subsequent fine-scale values as needed. If the estimated error of a value is within a specified tolerance, the value can be used; otherwise, the application completes a full fine-scale computation and stores the result for future use. Our implementation runs the coarse-scale model as an ordinary SPMD job. When a process in this job needs to compute fine-scale data,

CONTACT INFORMATION

John May, johnmay@llnl.gov
 David Jefferson, drjefferson@llnl.gov
 Web site: www.llnl.gov/casc/coopParallelism/



Adaptive sampling uses interpolation to reduce the number of fine-scale model calculations needed in a multiscale simulation. When an interpolated value is not accurate enough, the coarse-scale model requests a full computation from the server proxy, which forwards the request to an available fine-scale server.



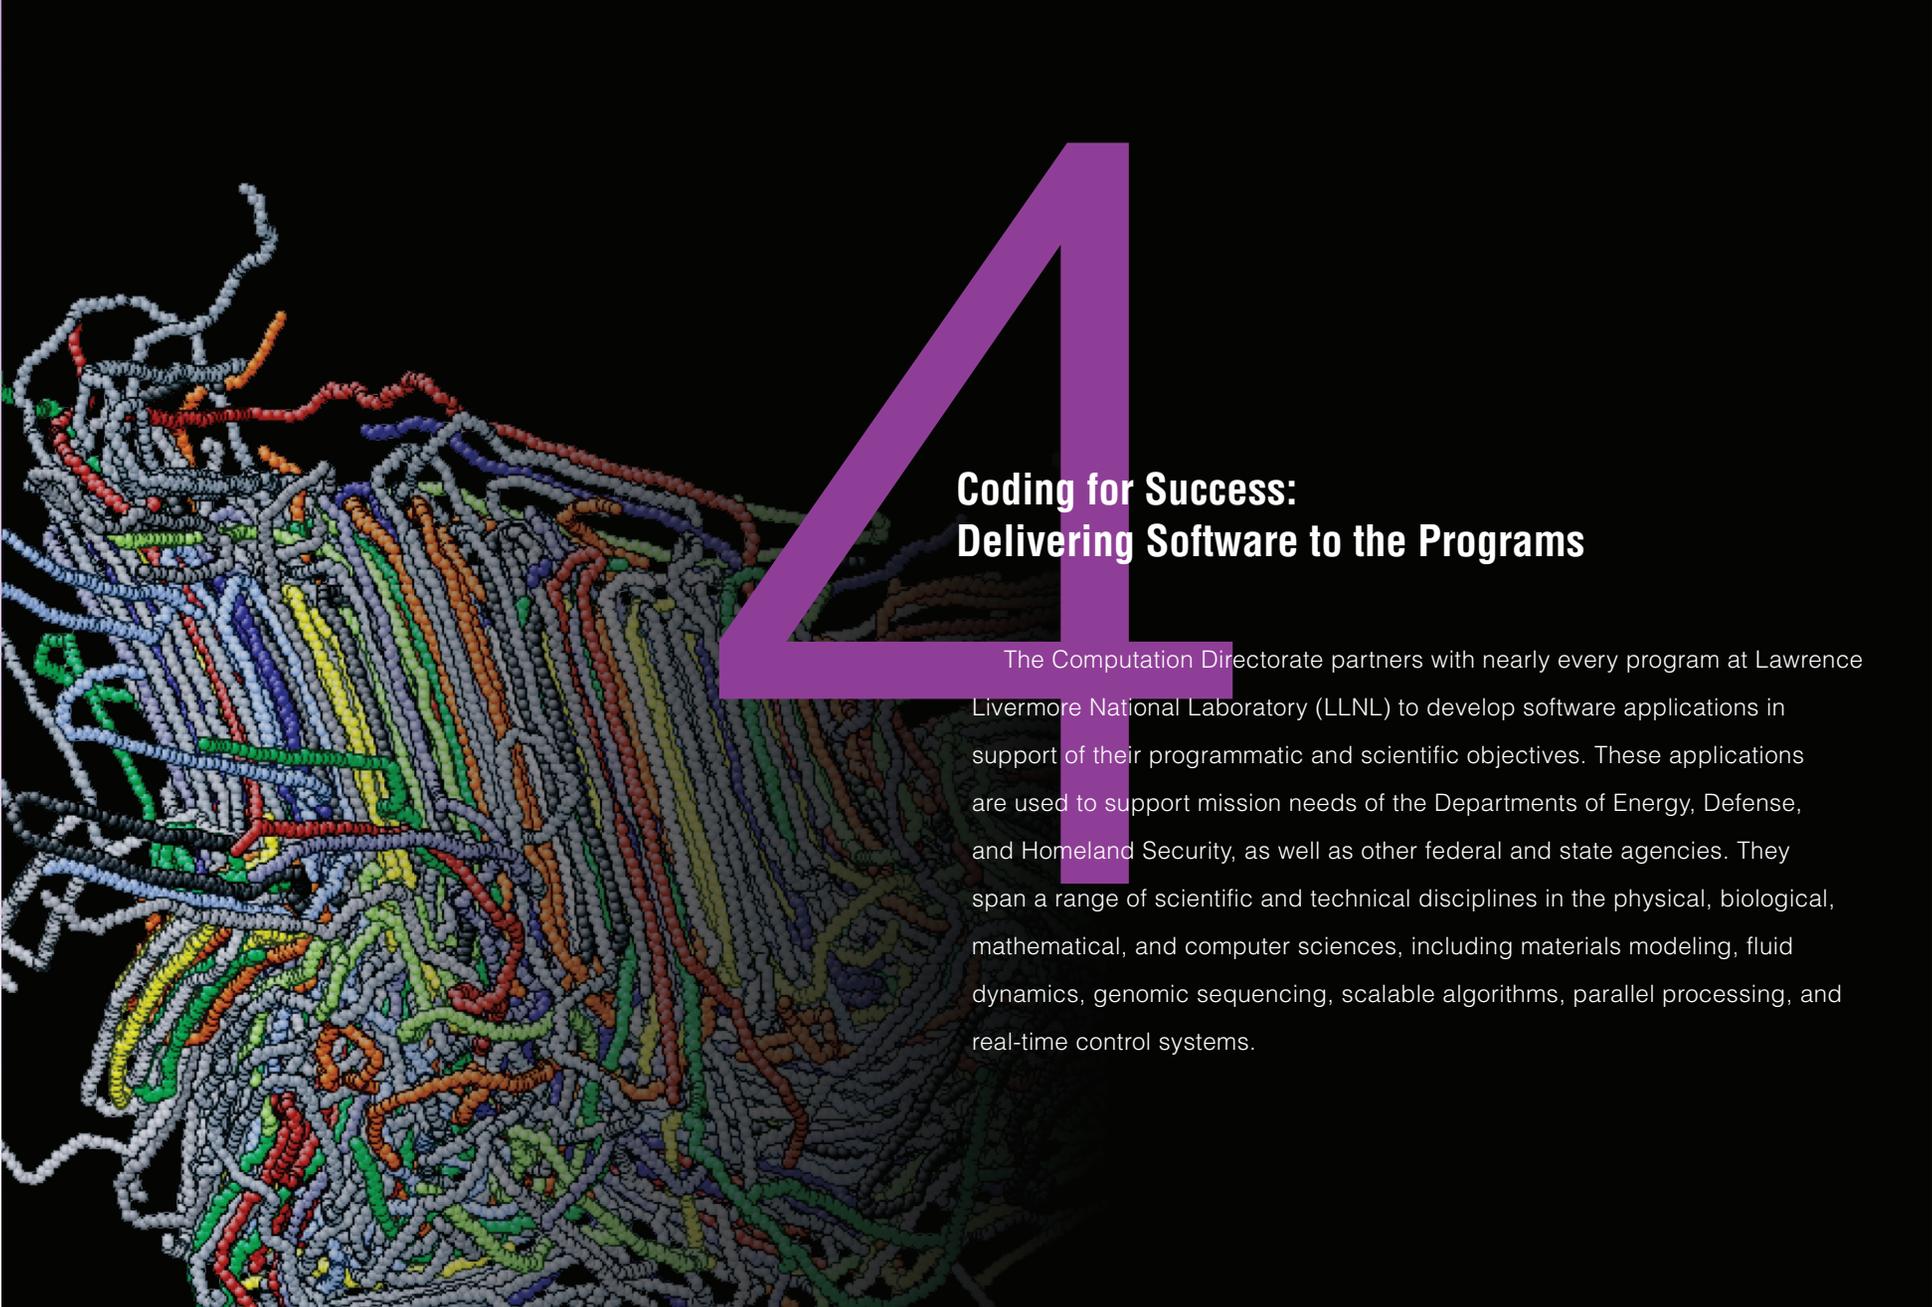
it sends a request to an available fine-scale server by proxy. (The proxy balances the workload among the servers.)

In September, our team completed its first large-scale demonstration by running the federated application on 680 nodes (1,360 processors) of the Multiprogrammatic Capability Resource cluster. This demonstration showed the benefits of adaptive sampling and, to a lesser extent, the load-balancing effects of cooperative parallelism. Without these features, our simulation would have taken an estimated 262 hours to complete; with the features, it finished in only seven hours, producing a 37-fold speed-up.

We expect larger systems to experience even greater benefits from load balancing.

SIGNIFICANCE

Cooperative parallelism will benefit large-scale codes as developers seek efficient ways to use massive parallelism. Its task-parallel model lets applications exploit multiple dimensions of parallelism at once. Moreover, because cooperative parallelism can coexist with current SPMD parallelism, applications can adopt it incrementally. We are discussing cooperative parallelism with groups internal and external to LLNL, and we expect to apply our model to several new applications in 2007.



Coding for Success: Delivering Software to the Programs

The Computation Directorate partners with nearly every program at Lawrence Livermore National Laboratory (LLNL) to develop software applications in support of their programmatic and scientific objectives. These applications are used to support mission needs of the Departments of Energy, Defense, and Homeland Security, as well as other federal and state agencies. They span a range of scientific and technical disciplines in the physical, biological, mathematical, and computer sciences, including materials modeling, fluid dynamics, genomic sequencing, scalable algorithms, parallel processing, and real-time control systems.

Coding for Success: Delivering Software to the Programs (cont.)



CONTACT INFORMATION

Steven Ashby, ashby1@llnl.gov

The Computation Directorate employs approximately 500 application programmers, making it one of the largest applied computer science organizations in the world. These highly skilled professionals are matrixed to the Laboratory's programs and work closely with domain experts to design, implement, and run software applications. The teaming arrangements are tailored to meet the programs' needs, but they share common features. For example, programs often rely on a small number of computer codes that grow and evolve over time to meet changing mission needs. In most cases, Computation personnel stay with a particular project for many years—sometimes for an entire career—and thus gain considerable domain knowledge. In other cases, programmers are assigned to a project to meet short-term needs. Projects range from small teams to large, interdisciplinary groups with more than 20 computer scientists, engineers, and physicists. They work together to define the software requirements and architecture, and to implement, test, optimize, and run the code.

Our employees possess a multitude of technical, leadership, and administrative skills, as well as highly specialized domain knowledge. This combination makes them extraordinarily valuable to the programs they support. The required computer science expertise varies from one project to the next. For example, stockpile stewardship applications require individuals with expertise in parallel computing and numerical algorithms. The National Ignition Facility (NIF) control system software, on the other hand, requires expertise in real-time, distributed computing. A variety of programmatic and scientific applications, including those in support of counterterrorism, depend on scalable information-analysis capabilities.

In addition to meeting today's needs, the directorate recently launched a science and technology planning activity that will identify the skills needed to ensure future programmatic successes. Our employee development activities will be focused on these and other critical skills. We are developing a training program with technical and leadership tracks, and we also started a successful mentoring program and two group leader internship programs.

Software quality engineering is increasingly important to our customers and their sponsors. In response, the directorate initiated three complementary activities. First, we led an effort to establish an institutional software quality assurance policy and program. Second, we worked with our sister National Nuclear Security Administration laboratories and the Department of Energy (DOE) to define appropriate software engineering practices. Finally, we worked with the programs to tailor these practices for their needs. Computation matrixes software quality and software engineering experts throughout the Laboratory and actively promotes good practices through various mechanisms, including the Software Improvement Networking Group.

Our employees have applied their skills to a remarkable breadth of projects around the Laboratory and have done so with great success. They ran large-scale simulations in support of key stockpile stewardship milestones, and calculations critical to the Livermore reliable replacement warhead design. A team of computer scientists working on the NIF control system enabled the successful demonstration of cluster operations using 48 laser beams. Computer scientists supporting homeland security projects developed and deployed several new technologies, including improved radiation and biological threat detectors. Our employees' accomplishments were recognized externally, too. In 2006, Computation personnel received three prestigious R&D 100 Awards.

This section highlights a few of the directorate's many contributions in software applications development. Two articles describe work in support of the weapons program—application scaling on massively parallel computers and the visualization of the resulting large data sets. Two articles discuss how modeling is used in two large-scale experimental facilities, the Linac Coherent Light Source and NIF. Our varied work in data management and information analysis is highlighted in four applications: bioinformatics for pathogen detection, airborne tracking of ground objects, text analytics for counterterrorism, and institutional chemical tracking. Each article highlights how we apply our skills and expertise to institutional, programmatic, or scientific objectives.

4.01 — Testing the Scalability of Supercomputers



CONTACT INFORMATION

Michael Collette, collette1@llnl.gov

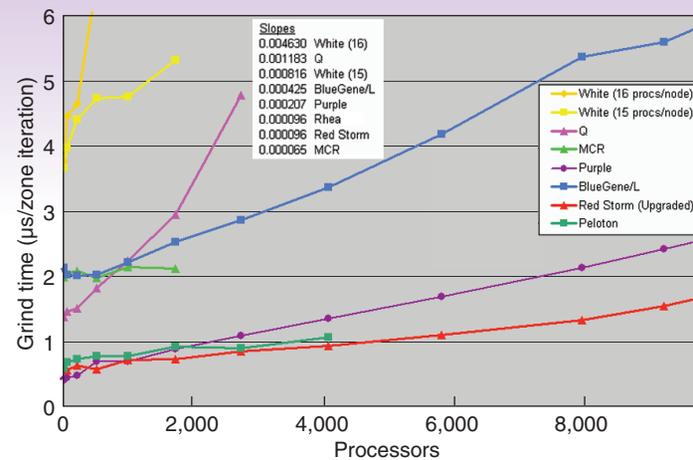
OVERVIEW

During the past several years, supercomputers have used more processors to attain high performance. But now the computing community is questioning how effectively applications take advantage of increasingly large numbers of processors. To answer this question, we tested the scalability of a large multiphysics two-dimensional (2D) and three-dimensional (3D) Advanced Simulation and Computing (ASC) code on the major ASC parallel platforms. Weak scaling studies, in which the work per processor is held constant while the number of processors is varied, measure the communications and overhead costs imposed by adding processors. Problems that predominantly exhibit intensive collective communications have historically had the highest additional costs. Scalability improvements reduce these costs. Comparing weak scaling results against older generations of ASC machines highlights the progress in network efficiency that is necessary to make additional processors beneficial.

PROGRESS IN 2006

We performed weak scaling studies on the latest ASC parallel platforms, BlueGene/L (BG/L), Purple, Red Storm, and Rhea (one of the systems obtained from the Peloton procurement, see Section 2.02). We used a 3D thermal radiation diffusion problem because it exhibits intensive collective communication properties that scale poorly on the older generations of ASC machines (White and Q). This radiation problem stresses the implicit radiation solver and is initialized with a radiation profile that causes a large number of

iterations in the matrix solution. A collective reduction operation is performed in each iteration, making the problem especially sensitive to the performance of collective messaging. The problem is uniformly distributed across the processors. Ideally, the grind time (time per zone iteration) will be identical in each processor configuration. An increase in grind time is attributed to communication and other additional overhead costs. The results are recorded in a line graph; the slope of the line indicates each machine's scalability. When comparing slopes, flatter is better, and horizontal is ideal. Our results show that the slopes for the



Weak scaling results of the 3D radiation problem's grind time per machine. Flatter slopes are better, and horizontal slopes are ideal.

newer machines are flatter than the older machines, and that Red Storm, Rhea, and Multiprogrammatic Capability Resource scale the best.

However, factors other than scalability influence which new generation machine is more appropriate for a job. Purple has the fastest processors and the most memory per processor of the ASC machines. BG/L has the most processors but a very small memory. Sandia's Red Storm machine, which was tested after its major hardware upgrade in November, showed impressive scalability. Finally, Rhea was tested, and, although it also scaled well, it has the fewest processors of the new machines.

SIGNIFICANCE

The application and problem used in our scaling study is similar to actual ASC user jobs. The scaling of the newer ASC platforms exhibits a 10-fold improvement over early generations of ASC machines in handling collective communications. This impressive gain shows that excellent progress has been made over historically limiting scaling issues, allowing us to run problems that were previously impossible. Overall, the newest platforms are very competitive. The "best" machine to use will depend on the application, the communication pattern, and the problem size.

4.02 — VisIt: Beyond Pretty Pictures



CONTACT INFORMATION

Hank Childs, childs3@llnl.gov
 Web site: www.llnl.gov/visit

OVERVIEW

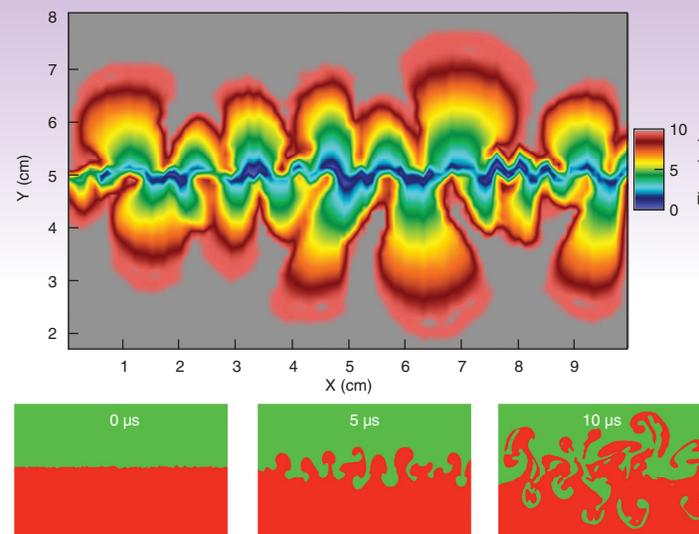
As LLNL transitions from terascale to petascale computing, simulation codes are producing finer mesh resolutions and more data. As a result, developers of postprocessing tools are confronted with two challenges. The first is to support the sheer scale of data; the second challenge is more subtle. Analysts cannot personally examine every byte of data. Instead, they need tools that allow them to manage, navigate, and present their data in ways that facilitate a conclusive interpretation of the phenomena being studied. The second challenge, therefore, is to facilitate the legibility of the massive data sets. This is often addressed through quantitative, rather than visual, means. In 2006, the VisIt team made significant progress toward providing an end-user tool for analysts that meets both of these challenges.

VisIt is an open-source visualization tool that supports more than 20 simulation codes, hundreds of users at LLNL, and many more at other institutions. It has many assets for data analysis already built into it: a scalable infrastructure; several file format readers; support for advanced data models, such as mixed-material elements; methods for slicing, dicing, clipping, and culling data; and, most importantly, a flexible and extensible software architecture. In addition, VisIt contains many common data-analysis techniques that span scientific domains, such as integration over a region, volume and surface area calculations, histograms, picks, and line-outs.

PROGRESS IN 2006

In 2006, the VisIt team introduced several features to help users better understand simulations. For example, a new query examines a computer simulation and then calculates what a detector at an aboveground experimental facility would observe, thereby enabling a

direct comparison of simulation and experiment. Many shape-characterization metrics were also added to support ASC's verification and validation efforts. Among them are techniques for calculating chord-length distributions (the probability that a random line over a shape will be a certain length) and ray-length distributions (the probability that a randomly placed and oriented particle will travel a certain distance before



Time-evolution results of a 2D Rayleigh–Taylor instability, in which a heavy fluid (green) and a light fluid (red) mix. This plot was created using VisIt's new comparative techniques. The field was generated by surveying all of the time slices produced by the simulation and determining at what time mixing occurred for each point in space. Blue areas mixed early in the simulation, red areas mixed later, and gray areas did not mix.

exiting a shape). Comparative techniques were added that allow analysts to go beyond side-by-side comparisons. Analysts can now map a set of field values from one simulation onto the corresponding mesh of another and then study the differences. This feature allows subtle changes to be discovered and makes more quantitative queries possible, such as computing a measure of the distance between the function on the two meshes. These comparative techniques permit multiple simulations to be incorporated, meaning hundreds of simulations can be studied, not just two. All of the improvements benefit from VisIt's built-in scalable architecture.

SIGNIFICANCE

The successful efforts of the last year demonstrate that VisIt can serve as a delivery vehicle for analysis capabilities for petascale data well into the future. This new direction has several benefits. First, it is easy for current customers to begin using the analysis capabilities. Second, there is a significant financial benefit associated with coupling visualization and analysis capabilities; future improvements will only have to be written into the code once. Finally, extensions can be implemented as custom data analysis arrives because an extensible and flexible code base was used.

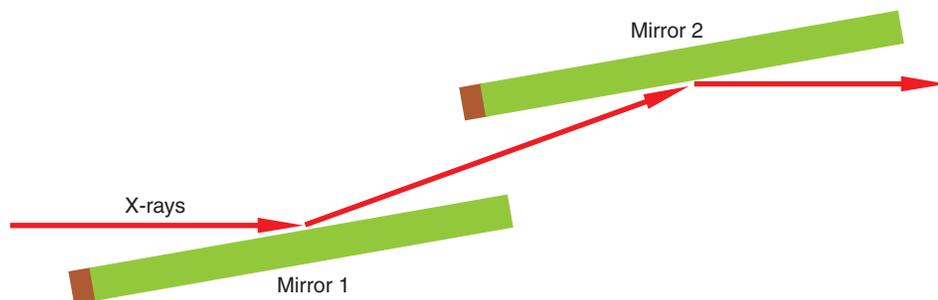
4.03 — Designing Mirrors for the Linac Coherent Light Source



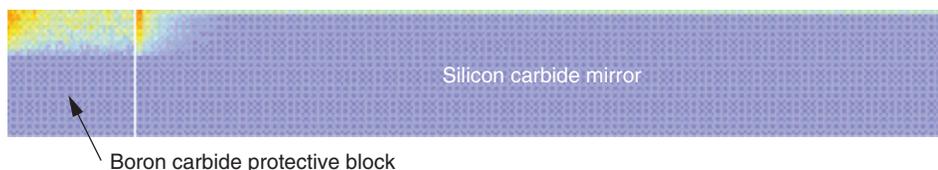
OVERVIEW

LLNL is collaborating with other national laboratories and universities to construct the world's first x-ray free-electron laser at the Stanford Linear Accelerator Center. The Linac Coherent Light Source (LCLS) will be operational in 2009. Characterized by high intensities, short wavelengths, and short pulse durations, the laser will enable new research, such as investigating new states of matter, understanding and following chemical reactions and biological processes in real time, and imaging nanoscale and noncrystalline biological materials.

LCLS is a pulsed laser; each pulse consists of electrons that are accelerated through a tuned magnetic undulator, producing desired coherent x-rays but also producing undesired wide-spectrum spontaneous radiation. A pair of offset x-ray mirrors acts as a low-pass filter and absorbs the radiation wavelengths, which are shorter than the laser wavelengths. Preserving the properties of the laser beam requires mirrors that are state-of-the-art in size, flatness, and smoothness.



Mirrors are made of silicon carbide, 1 m long, 2 cm wide, and 5 mm thick. Protective end blocks are made of 1 cm of boron carbide. Drawing is not to scale.



The sideways-projected false color contour plots of heat in the first mirror and its protective block are separately normalized to their hottest points. Drawing is not to scale.

CONTACT INFORMATION

Kirby Fong, fong5@llnl.gov
 Linda Ott, ott2@llnl.gov
 Web site: www-ssrl.slac.stanford.edu/lcls

PROGRESS IN 2006

Our first step in designing the mirrors was to determine how much heat the mirrors and mirror supports were absorbing from the radiation. This is important because the mirror supports have to be strong enough to conduct the heat away in addition to holding the mirrors in position. We prepared a model of the mirrors and a Monte Carlo simulation to determine the distribution of energy that is deposited in the mirrors by each laser pulse.

Although the top figure (at left) depicts the central path of the spontaneous radiation, the radial distribution is broad enough that some photons will pass over the top of the mirror and others will hit the front end. The mirror tilt is highly exaggerated in the figure; in reality it is no more than 1.5 mrad. Boron carbide blocks protect the mirrors from stray hits by the laser.

Our simulation begins with photons in the undulator that have a distribution of positions, directions, and energies analogous to the spontaneous radiation. The

photons are propagated out of the undulator through a fixed mask and collimator and into the mirror assembly. The simulation takes into account Compton scattering as well as photoelectric effect to determine if the photons will stop, where they will stop, and how much energy they will deposit. The output is a database of photon positions and energies.

SIGNIFICANCE

Heating of the LCLS mirrors from the laser can be calculated theoretically, but heating from the spontaneous radiation cannot be determined theoretically or experimentally until LCLS is built. As there are no empirical data about the spontaneous radiation from which to begin the thermal analysis of the LCLS mirrors, our simulations provide the only basis for the thermal design. While a worst-case estimate of the heat load might dictate expensive refrigerant-cooled mirror supports, our results show that the amount and spatial distribution of heat can be handled passively by radiation in the vacuum to the containment walls.

4.04 — Bringing Out the Best in the Big Lasers



CONTACT INFORMATION

Ronald House, house6@llnl.gov

OVERVIEW

The primary mission of NIF at LLNL is to attain fusion ignition. NIF will ultimately consist of 192 laser beams designed to deliver 1.8 million joules of ultraviolet laser energy and 500 terawatts of power to millimeter-sized targets. The Laser Performance Operations Model (LPOM) supports NIF operations by automatically determining how to accurately set up the 48 injection laser systems and 192 beam lines to meet the goals of each individual shot. LPOM utilizes laser propagation calculations and a living performance model of each beam line to predict the performance of NIF lasers. The predictions are then compared to diagnostic feedback and used to verify and adjust the system performance during the shot lifecycle. In addition, LPOM supports NIF shot planning and postshot reviews.

PROGRESS IN 2006

LPOM is the first software of its kind. The initial LPOM prototype was deployed several years ago to support the NIF Early Light campaign, which allowed the LPOM development team to gain operational experience. Since then, the LPOM team has followed a rapid prototyping spiral strategy, with an emphasis on supporting commissioning operations, automating processes to eliminate the need for an LPOM operator in the control room, improving feedback using a Web-based data browser interface, and scaling to support full NIF operations.

Without concurrent processing, one full shot set-up calculation on NIF could take eight hours or more to complete. LPOM has been designed to take advantage of the independent nature of the NIF quad-based architecture by running calculations in parallel on a distributed computing cluster. The initial LPOM prototype consisted of two separate clusters: a large cluster used for NIF shot planning and “what-if” analysis, and a small cluster to support NIF operations and shot verification. During 2006, LPOM was redesigned to allow a single computational cluster to support both functions concurrently and efficiently. Since this single LPOM cluster became operational, it has



LPOM compute cluster, comprised of 112 hyperthreaded compute nodes, provides computational capabilities to simultaneously support NIF shot planning, shot reviews, and shot operations.

supported large power balance analysis calculations and the first full cluster (48 beam lines) shots of NIF.

LPOM was interfaced with the virtual beamline (VBL) code for its physics modeling. VBL is a state-of-the-art propagation and laser-energy extraction code, written in Java, designed to support NIF. By using a detailed physics model for its setup calculations, LPOM is able to routinely deliver energies within 3% of what is requested and provide energy balance among the four beam lines in a quad to within 1%.

SIGNIFICANCE

LPOM is designed to be scalable and to efficiently share computational resources in support of NIF shot planning, shot operations, and shot reviews. Successful NIF operations will depend on bringing out the best of the big lasers; this means delivering precise high-power waveforms from 192 laser beams for a variety of pulse lengths and temporal shapes with the minimum possible optical damage. LPOM will provide this critical capability for NIF.

4.05 — Beyond Pathogen Identification: Detecting Mechanisms of Bioterrorism



CONTACT INFORMATION

Tom Slezak, slezak1@llnl.gov
 Crystal Jaing (CMLS Directorate), jaing2@llnl.gov

OVERVIEW

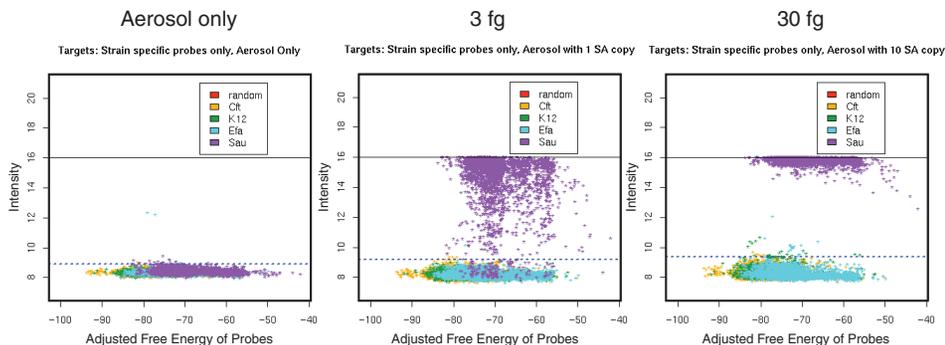
Since 2000, LLNL has led the effort to design DNA signatures for pathogen identification, working closely with the Department of Homeland Security and other agencies. Many of these signatures have been in use around the country for more than five years, including during the 2002 Winter Olympics, as part of the Biological Aerosol Sentry and Information System and BioWatch programs.

LLNL researchers recently began designing higher-resolution signatures that can identify known mechanisms of pathogen virulence and antibiotic resistance, as well as evidence of bacterial genetic engineering. This is the first step toward combating the threat of engineered organisms.

PROGRESS IN 2006

A multidisciplinary team of employees was assembled to build a prototype of a microarray to detect known mechanisms. We used comparative genomics

software on the Thunder supercomputer to compare more than 700 hidden Markov models against the microbial sequence information that was available, a calculation that would not have been feasible without access to LLNL's high-performance computers.



Limit-of-detection test results for Chip 2A. Three levels of pure DNA from a test organism were spiked into DNA from a BioWatch aerosol sample. Each blue spot represents a probe to a known virulence mechanism. The x-axis is a measure of free energy, and the y-axis indicates microarray intensity on a log scale. The dotted line is a dynamically calculated threshold for detection. Clear evidence of detection is seen at 3 fg. At 30 fg, nearly all mechanism probes for the target organism are pegged at maximum intensity.

With our first prototype, Virulence Chip 1, we learned how to design optimal signature probes for the NimbleGen synthesizer, a machine capable of making microarrays with 390,000 DNA signatures. Ten known virulence genes that were present in each of four common bacteria were selected. We then designed thousands of signature probes that tested numerous design parameters, such as probe length and DNA melting temperature. Many thousands of deliberately mismatched probes were also included. The results of the Chip 1 tests showed that signature probe lengths of 55–70 base pairs provided the robustness that was needed to deal with natural variation.

We designed the next prototype, Chip 2A, to include the known virulence and antibiotic-resistance genes in the four test organisms as well as several genetic engineering vectors. Lab testing attempted to identify the limit of detection of the test organisms in realistic conditions by spiking pure DNA into BioWatch aerosol samples. The results of the test showed an ability to detect the mechanisms of the spiked organism and the genetic engineering vectors against the complex environmental

background at a level of approximately 1–10 genomic copies.

In late 2006, we completed the Chip 2B design, which contains the known mechanisms for current BioWatch organisms. The first tests of the 2B chips matched the performance of Chip 2A. The 2B chips are now being tested against a large set of pathogen and near-neighbor strains. These tests are expected to prove the ability of Chip 2B, which is now called the BioWatch Mechanisms Chip, to provide forensic-level resolution between different strains of the same pathogen.

SIGNIFICANCE

Our work is the first large-scale attempt to detect, in parallel, many thousands of known mechanisms of pathogens and evidence of bacterial genetic engineering. The BioWatch Mechanisms Chip will provide actionable information about public health threats and the efficacy of potential countermeasures days before genome sequencing can be performed. It is also the first detection mechanism designed to deal with engineered threats.

4.06 — Sonoma: A Leader in Broad-Area Surveillance



CONTACT INFORMATION

Aaron Wegner, wegner2@llnl.gov
 Deanna Pennington (NHI Directorate),
pennington1@llnl.gov

OVERVIEW

The R&D 100 Award-winning Sonoma project is a broad-area persistent surveillance system designed to continuously monitor an area for an extended period of time. The system is able to survey a city-sized area at relatively high resolutions, both spatially and temporally, from an airborne platform. As data is collected, it is processed using algorithms specially designed for modern graphics processing units (GPUs) to facilitate compression, transmission, and analysis in real time. The data often contain hundreds or thousands of moving objects and other points of interest. The Sonoma data analysis and visualization tools allow users to analyze movements in the field of view and identify relationships between previously unassociated objects based on traffic patterns. This information is relevant to a variety of mission types, including nuclear nonproliferation, border security, and special-events monitoring.

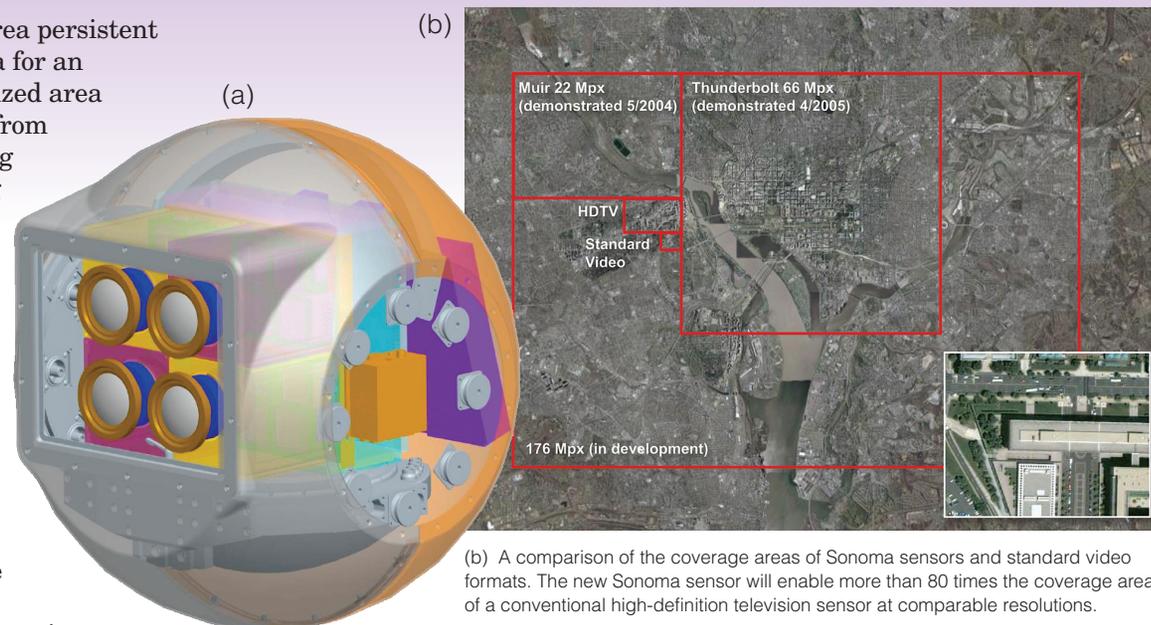
PROGRESS IN 2006

2006 was a year of intense development for the Sonoma team. Major design advancements were made to both the sensor and supporting software, dramatically improving the system's capabilities in terms of coverage area and data handling by a factor of approximately 3. Improvements to sensor quality and data postprocessing algorithms were also realized.

A new 16-camera, 176-megapixel sensor system was designed and is currently under construction. Using techniques similar to those used with the 6-camera, 66-megapixel

sensor that was deployed in 2005, the new design yields a significant improvement in focal plane array size. It takes advantage of a modified charge-coupled device (CCD) that improves the signal-to-noise ratio of imagery, particularly at the edges of the field of view.

A next-generation software architecture was needed to handle the larger volumes of data produced by the new sensor. Development began on new data-acquisition and processing software that will take full advantage of the parallelism offered by multicore processors and modern GPUs. The Sonoma software programmers conducted research on cutting-edge



(a) The latest Sonoma sensor consists of four extremely low-distortion, large-format lenses and sixteen 11-megapixel CCDs, which form a 176-megapixel focal plane array. It takes considerable processing power to organize, store, and analyze the data produced by this sensor.

(b) A comparison of the coverage areas of Sonoma sensors and standard video formats. The new Sonoma sensor will enable more than 80 times the coverage area of a conventional high-definition television sensor at comparable resolutions.

compression techniques and took cues from ASC Program personnel on large-volume data visualization. The new generation of Sonoma software defines a distributed processing model that enables highly efficient image processing, compression, and visualization, all in real time. In addition, the architecture is versatile enough to use with other sensor modalities.

SIGNIFICANCE

Sonoma is a leader in the field of persistent video surveillance and is the only effort of its kind that provides a complete, integrated end-to-end solution. Each component of the system continues to push the technological envelope, including the sensor design, data handling, data visualization, and exploitation. The success of Sonoma has had an appreciable effect on design and concepts of operation of similar systems throughout the surveillance community.

4.07 — Discovering Information in Unstructured Text

OVERVIEW

The difficulty of analyzing unstructured data is a major obstacle for a variety of national security and scientific applications. In 2005, the Counterproliferation Analysis Section of the Nonproliferation, Homeland, and International Security Directorate addressed this difficulty by prototyping a software solution that discovers relevant information in unstructured text. The software examines the text, determines context, applies criteria, and then displays a unified set of results to the analyst.

The prototype was integrated with commercial, off-the-shelf software to leverage the customer's existing toolset for categorization and indexing and was enhanced by adding a government-developed clustering tool and custom filtering features. Through use of the prototype, it was determined that the existing toolset was insufficient to filter the data to a manageable amount; furthermore, additional integration with other tools would not be feasible in the long term. It became apparent that a customizable software solution to facilitate technology reuse and rapid integration of new analytical tools was needed. This need led to the creation of a tool that utilizes the Unstructured Information Management Architecture (UIMA) as an application framework for the rapid development of collaborative scalable solutions for processing unstructured text.

PROGRESS IN 2006

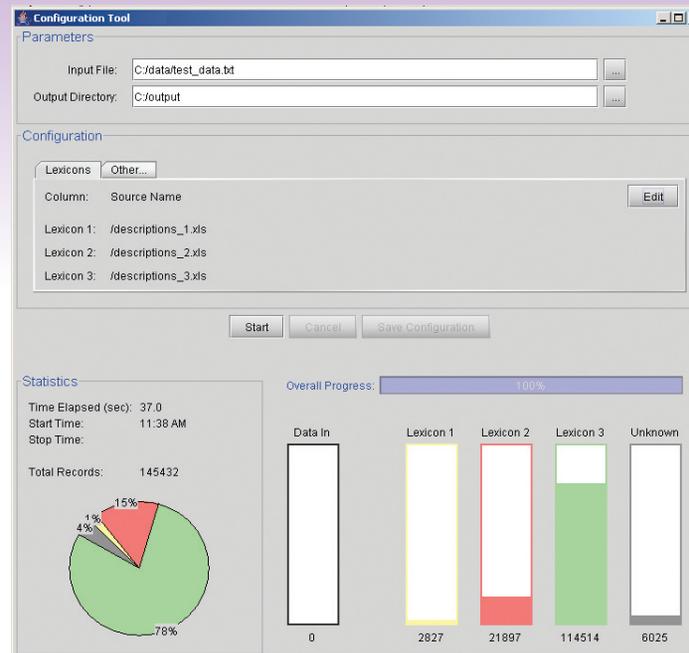
The new tool was specifically created to solve the data overload problem that occurs in the assessment of commodities data and proliferation. It enables the categorization and prioritization of commodities based on context-sensitive criteria defined by the analyst. Prior to creation of the tool, data preparation was done manually. Now, the process is automated, which allows a time-constrained user more

opportunities for deeper analysis of relevant data.

The UIMA framework provides the foundation to “plug and play” analytical tools and gives the end user a unified view of the results. The process begins by preprocessing the data set through customizable, data-independent components. Once the data set has been processed, multiple text-based analytics (software components designed to parse and interpret text) are applied in a problem-specific configuration. These analytics include open-source and homegrown components that are

CONTACT INFORMATION

Marisol Gamboa, gamboa2@lnl.gov
Steve Suppe, ssuppe@lnl.gov



Analyst's automated processing tool for filtering document sets.

implemented independently of the data set, making them useful in any context. Each analytic then produces a relevance rating defined by a user-configurable system of weighted criteria. Along with the relevance rating, a vote/confidence value is provided that weighs similar analytic results. These results help users identify the key findings and validate the accuracy of the analytics. Further analysis can then be performed on the data set via a feedback mechanism that allows users to reprocess results to leverage previous findings.

SIGNIFICANCE

A standard, unified view of otherwise heterogeneous analytical solutions enhances users' ability to perform data processing and information discovery. In one instance, the tool reduced the time it took an analyst to filter data from approximately two hours to 28 seconds. In the future, we will improve and expand our analytical toolset to support a broader customer base. The ability to aggregate multiple analytic technologies and provide a unified view of the results will create a more effective environment for analyzing unstructured information.

4.08 — ChemTrack: End-to-End Chemical Logistics



CONTACT INFORMATION

Craig Cothren, cothren2@llnl.gov

Web site: chemtrack.llnl.gov

OVERVIEW

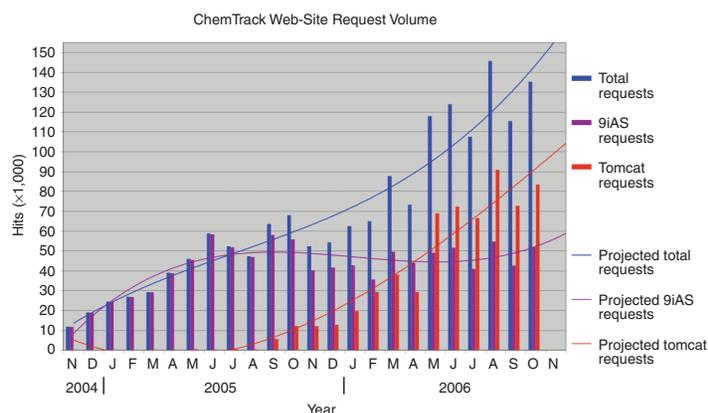
Safety and regulatory compliance have long been top priorities at LLNL, especially when it comes to managing hazardous material. ChemTrack helps the institution meet regulatory requirements by providing a custom end-to-end chemical inventory and logistics system that is capable of threshold maintenance and safety data correlation. ChemTrack is a Web-based application with a database back-end that captures the movement of chemicals, tracks their aggregate quantities, and maintains the corresponding material safety data sheets across LLNL. The data are used to fulfill the regulatory community right-to-know reporting requirements of Alameda and San Joaquin counties, state and federal environmental protection agencies, the University of California and DOE. ChemTrack is available to LLNL employees around the clock. Approximately 2,000 registered users can update the data either directly or via integration with other systems.

PROGRESS IN 2006

Prior to 2006, ChemTrack was synched with the Laboratory's electronic ordering system (the system through which chemical and

other purchase transactions are made), but information only flowed one way—from the ordering system to ChemTrack. In early 2006, we integrated ChemTrack more tightly with the electronic ordering system via a loose coupling so that information could flow both ways.

ChemTrack usage has grown 11-fold during the past 24 months.



This may seem like an oxymoron, but the loose coupling allows both systems to interoperate while maintaining completely different timetables; at the same time, tightened integration is achieved because data moving between the systems use a defined synchronous protocol.

In late 2006, we completed ChemTrack's Safety Basis notification and authorization feature. Safety Basis is a set of features that enable authorized program personnel to set and manage the maximum amount of chemicals allowed within their facility. Previously, ChemTrack sent notification and authorization for movement of chemicals within the Laboratory but did not do the same for purchases of new chemicals. Now, the tighter integration with the electronic ordering system enables ChemTrack to warn buyers that the purchase of an additional quantity of a chemical will approach or exceed a predetermined threshold. If buyers choose to continue with the purchase, a notification is sent to the appropriate facility manager. The facility manager may then reject the purchase prior to the transaction with the vendor

to avoid the consequences of a chemical overage.

These enhancements were made possible by ChemTrack's adoption of Web 2.0 design patterns and architecture. The model-view-controller design allows reuse and shorter development life cycles. Asynchronous JavaScript and Extensible Markup Language provides a state-of-the-art user experience that is pleasing and intuitive. Using these elements, ChemTrack achieves greater agility to create systems that provide safety and regulatory compliance for LLNL programs, allowing them to concentrate on their work in support of Laboratory missions.

SIGNIFICANCE

The demand for the tools that ChemTrack provides to accomplish safety and regulatory compliance has been steadily increasing. Customer usage has increased 1,100% since 2004 and more than 150% since 2005.



5 Impacting Institutional Infrastructure, Security, and Service

The Computation Directorate supports institutional activities at Lawrence Livermore National Laboratory (LLNL) in several ways. Some activities help the institution operate smoothly and meet technical mandates, others move the institution forward on a technical level, and some do both. Computation personnel are often involved in setting the direction of LLNL's institutional activities in addition to implementing them. As a multifaceted interface across LLNL, Computation takes on the role of integrating, leveraging, and sometimes de-conflicting efforts to maximize success.

5.00 — Impacting Institutional Infrastructure, Security, and Service (cont.)



CONTACT INFORMATION

Ronald Schmucker, schmucker1@llnl.gov

Among the institution's myriad Information Technology (IT) needs are institutional software quality assurance (ISQA) planning and implementation, desktop computing management, large data storage environments, network backbone operations, and cyber security—initiatives that span virtually the entire Laboratory. Our work is made more complex because IT requirements come from several sources: LLNL's Safeguards and Security Organization, the Cyber Security Program (CSP), the office of the Chief Information Officer (CIO), and the Department of Energy (DOE). Computation is also home to the Institute for Scientific Computing Research (ISCR), an organization that advances external outreach, collaboration, and visibility efforts on behalf of the Laboratory.

One of the issues we face that is becoming increasingly acute is how to deal with the capabilities and services being developed in the marketplace that challenge the security requirements of a national security laboratory. (Try buying a cell phone without a built-in camera or a laptop without built-in wireless.) Many of LLNL's security requirements impact the entire institution, not just programs that are located in classified areas. To compound this difficulty, in our intense pursuit of mission, we utilize a wide range of platforms and operating systems, which make seamless interoperability doubly difficult.

Supporting LLNL's cyber security organization is one of our most important responsibilities; a few examples of our efforts are included in this section. Following a "defense-in-depth" philosophy and adding new defenses as cyber security threats evolve has proven very successful. In addition to regular improvements, a new Web proxy and more rigorous network protocol egress port blocking drastically improved our security posture. Although the cyber threat volume and sophistication continue to grow, we can claim another year without a major service disruption or

compromise. Computation also helped implement new DOE requirements to protect information contained on devices (i.e., laptops or handheld organizers) or media that leave the Livermore site. We are completing our institution-wide effort to ensure that sensitive information (personally identifiable information, for example) is encrypted when it is taken off site.

Two additional Computation-led institutional efforts that are elucidated in this section are the Configuration Management Database (CMDB) and LANDesk. The projects are related, and both aim to improve LLNL's IT best-practice management and reporting capabilities. LANDesk will be a data source into the CMDB and replace the Hewlett-Packard (HP) Radia and BigFix data sources. Inventory data that is harvested off institutional computer systems is provided via these tools and fed into CMDB. The inventory data includes operating system information and system configuration data based on certification and accreditation (C&A) requirements.

Computation's Green Data Oasis project was implemented to facilitate the sharing of huge amounts of information with external collaborators. This project enables LLNL scientists and researchers to share large volumes of data with other institutions so that we can fully participate in national and global science efforts. The Green Data Oasis is moving to a 10-GB-per-second network connection delivering more than 600 TB of data files. This capability allows our climate modeling data, for example, to be moved rapidly to other institutions for comparison with their climate models. We expect the Green Data Oasis to be utilized by our FY07 Grand Challenge science efforts and collaborations that are beginning with the State of California on environmental issues and alternate fuels.

The seven articles in this section highlight our continuing success in pursuit of best practices and service enhancements to the benefit of LLNL.

5.01 — Assuring the Quality of Safety Software at LLNL

OVERVIEW

In 2005, DOE issued a new set of software quality assurance (SQA) requirements to the National Nuclear Security Administration (NNSA) complex. The requirements pertained to safety-related software associated with nuclear or radiological facilities that falls into one of three categories: safety system software, safety and hazard analysis and design software, or safety management and administrative controls software. At LLNL, this description includes software that controls emergency response systems in certain facilities, software that determines the risk of exposure to harmful substances, and software that determines when it is safe to move materials.

Although the Laboratory has had a documented ISQA plan in place since 2004, we broadened its scope and made specific modifications to address the 2005 DOE requirements.

PROGRESS IN 2006

The original ISQA implementation plan included a framework and graded approach by which all software development efforts were evaluated to determine their potential risk and associated software quality practices. The ISQA risk-based graded approach helps identify high-consequence software in use at the Laboratory. Software that has higher risks associated with its use, such as physical safety dangers or negative environmental consequences, receives greater oversight and guidance.

To address the 2005 DOE requirements, we adjusted the grading process to include a

consequence-of-failure level that reflects the potential life-and-death consequences of safety software. We also reworked our likelihood-of-failure calculation using Sunita Denani-Chulani's COCOMO II-based



LLNL's certified software quality engineers: (left to right, standing) Robert Gleason, Uma Reddi, Greg Pope, Thomas Crook, and Darrel Whitney; (sitting) Juan Hernandez, Sanni Schreck, Vicki Pope, Barbara Campbell, and Jeremy McCarthy (not pictured: Trish Damkroger).

CONTACT INFORMATION

Barbara Campbell, campbell2@llnl.gov
 Darrel Whitney, whitney1@llnl.gov



empirical defect-insertion cost-estimation model. The changes we made to the implementation plan were also made to the Web-based tool that customers use to grade the risk of their software.

Next, we researched and selected an industry standard for developing or acquiring safety software. Safety software concerns were then integrated into the Laboratory's procurement process by updating the procurement worksheet and the contract template. The NNSA Livermore Site Office and NNSA headquarters approved the newly revised ISQA program description in the summer of 2006. In addition to assuring Laboratory compliance

with the new DOE requirements, the LLNL ISQA staff participated in assessments at other NNSA sites to help them demonstrate their compliance.

To assure the quality of LLNL's safety software, we must first identify which software is considered "safety software" by reviewing the safety analysis documentation for nuclear and radiological facilities. In 2006, we reviewed 75% of that documentation. The next steps are to evaluate the current SQA practices for each safety software application, identify potential gaps between current practices and desired practices, and resolve those gaps.

SIGNIFICANCE

The basic principle of assuring software quality is a simple one: The higher the risk associated with the project, the greater the care that must be taken. Providing specific guidance on the quality of our safety software is one way we can be stewards of public health and better protect Laboratory employees so they can continue to safely serve the scientific interests of our country.

5.02 — Configuration Management Database Aids Accreditation Process

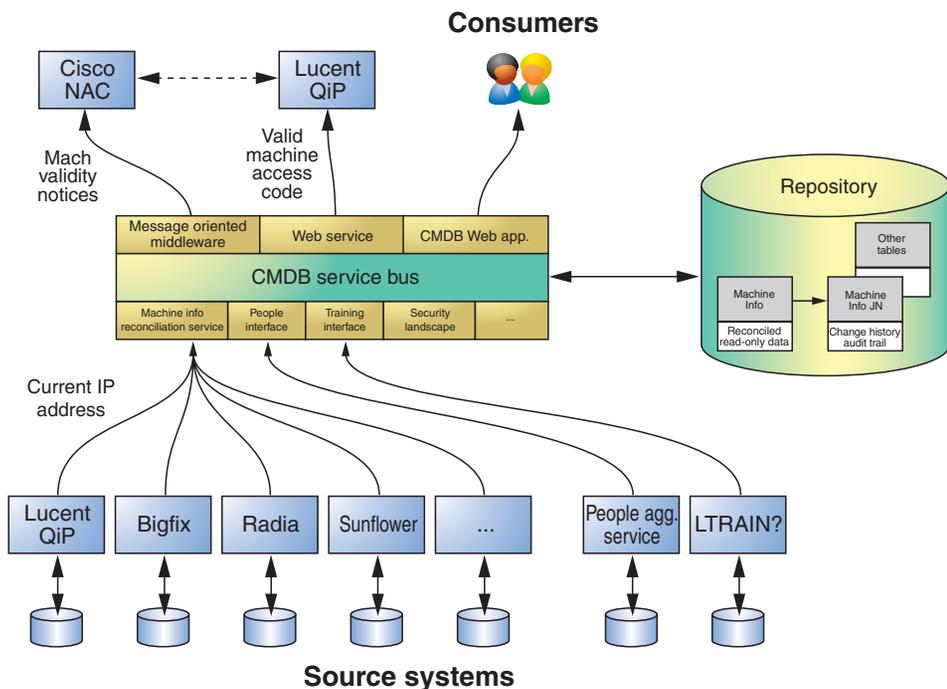


CONTACT INFORMATION

Robyne Teslich, teslich2@llnl.gov
 Amanda Goldner, goldner2@llnl.gov
 Web site: www.cmdb.llnl.gov

OVERVIEW

The Computation Directorate supports the majority of desktop, server, and IT services at LLNL. For the past several years, our focus has been on implementing and automating best-practice management capabilities. Best practices require that the Laboratory be more accountable not only for the financial tracking of IT assets, but also for their configuration state throughout the IT asset lifecycle. Our goals include improving the Laboratory's ability to respond to data calls and compliance requirements and to support business decisions about the IT environment. To this end, the CIO sponsored an initiative in 2005 to develop an institutional CMDB. We implemented the CMDB in 2006 and used it to help LLNL achieve C&A to continue to operate.



The CMDB architecture.

PROGRESS IN 2006

Prior to the CMDB initiative, LLNL had several authoritative sources of information about IT assets, including an asset-tracking system, a network-device registration system, and a configuration management tools inventory. When it came time to gather information and make decisions about IT, we had to consolidate and reconcile these sources manually, which was a very time-consuming and error-prone process.

In 2006, we made significant progress toward consolidating data sources using the CMDB. The CMDB creates a data-reconciliation service by gathering information from multiple sources and establishing authoritative sources for each element of information. Following standardized rules, the CMDB then merges the information from all the sources into a single, seamless view.

Our primary objective in the initial release of the CMDB was to use it to deliver information that would help the Laboratory through the computer security C&A process. During this process, we account for our IT assets and ensure that a

certain percentage of those assets meets established configuration criteria. We used the reconciled data from the CMDB to create the reports we needed to satisfy the C&A requirements. Although there was a significant amount of data clean-up required to create the initial reports, if we continue to use the common processes in the CMDB, over time we will acquire more uniform and accurate data with less effort. Automated configuration management tools, such as HP Radia and Bigfix, were used to collect information from systems across the institution.

SIGNIFICANCE

The CMDB initiative was established to provide LLNL an authoritative source of information about IT resources. The database provides the foundation for sound business decisions, security analysis, and compliance measurement. This year marked the first time LLNL used a common system and metric process to respond to C&A requirements. The CMDB also provides a consistent process for IT configuration management. In the future, it will enable security plan management, and it will be expanded to support configuration management for classified systems.

5.03 — LANDesk Enhances IT Efficiency and Security



CONTACT INFORMATION

Neal Mackanic, mackanic2@llnl.gov
 Web sites: <https://smsg.llnl.gov>,
<http://www.llnl.gov/icc/netdiv/>

OVERVIEW

LLNL operates in an environment in which urgent demands concerning our IT infrastructure are continually placed on us. These demands often involve providing information to outside auditors, responding to security threats, complying with rules and regulations, or making sound business and technical decisions. Our ability to respond depends on having a robust systems-management infrastructure in place. Our support infrastructure must provide the following capabilities: hardware and software inventory management, security update management (patching), automated software distribution, remote administration, reporting, software usage monitoring, and a configuration management database. In addition, these capabilities must be cost effective and implemented with as much automation as possible. To provide the necessary capabilities, LLNL decided to purchase and deploy the LANDesk Enterprise Management Suite. When fully implemented, LANDesk will manage more than 10,000 Windows, Macintosh, and Linux desktop and server computers on the unclassified network.

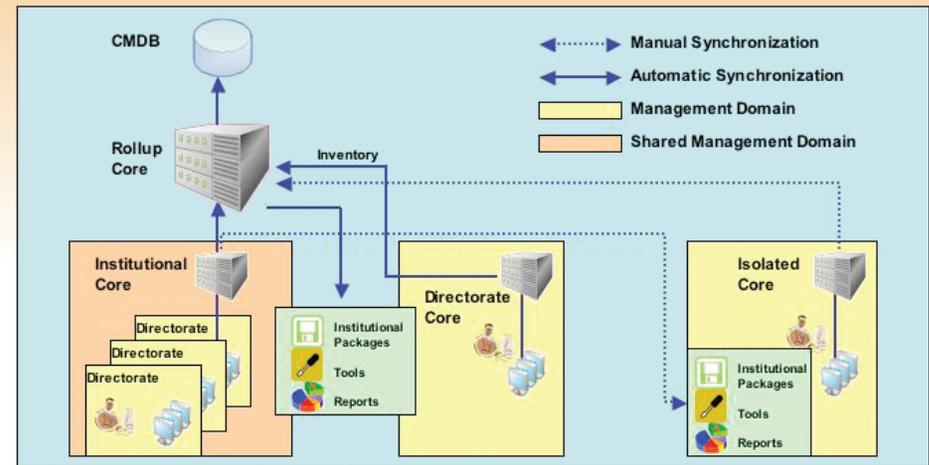
PROGRESS IN 2006

As progress is made toward automated desktop and server management, the awareness of new requirements has increased. Computation customers have expressed interest in expanding LLNL's capabilities for automated desktop management. Specifically, there was a strong interest in investigating the market leaders' enterprise-management suites based upon the assertion that a single tool with multiple administrative functions would increase the efficiency of desktop support staff and improve customer support. In addition,

LLNL needed a solution that could architecturally accommodate diverse security requirements and could be deployed in classified or highly sensitive environments while also supporting institutional data collection requirements.

In 2005, we conducted a small pilot of Microsoft's System Management Suite. In general, we found that the suite provided enough of an increase in functionality and improved efficiency to warrant a broader investigation of the other market leaders, specifically LANDesk and Altiris. After comparing the three vendors, we selected the LANDesk system in 2006.

LLNL is in the planning stages for the LANDesk deployment. We



LANDesk architecture.

have defined a logical architecture that meets the various security needs of the directorates. We will focus initially on replacing the current tools, Bigfix and HP Radiia, and will then move to leveraging the product's additional functionality, such as remote assistance, software usage metering, and port locking.

In addition to the basic enterprise-management functionality, LANDesk provides an automated business-process manager that LLNL plans to employ. This component will automate IT processes, such as patch management and software distribution, and will allow us to standardize the various processes for IT management while enabling customers to maintain control of their desktop environment.

SIGNIFICANCE

The LANDesk suite will provide a single administrative interface to effectively manage more than 10,000 computers at LLNL. LANDesk will enable us to respond quickly to the ever-changing computing environment, including security and system configurations, compliance auditing, and software applications. The LANDesk process manager will automate patch and software distribution, thereby resulting in a higher percentage of systems meeting baseline configuration requirements and reducing exposure to vulnerabilities.

5.04 — Blue Network Improves Computer Security Posture



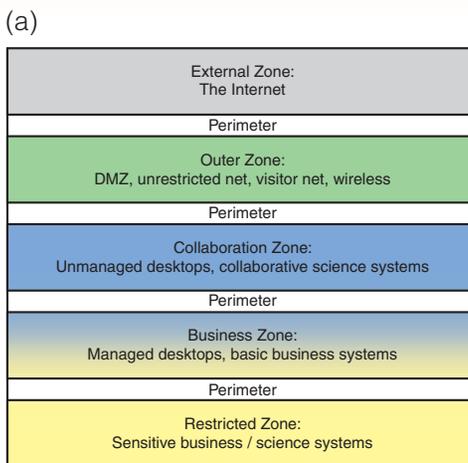
CONTACT INFORMATION

Robin Goldstone, robing@llnl.gov

OVERVIEW

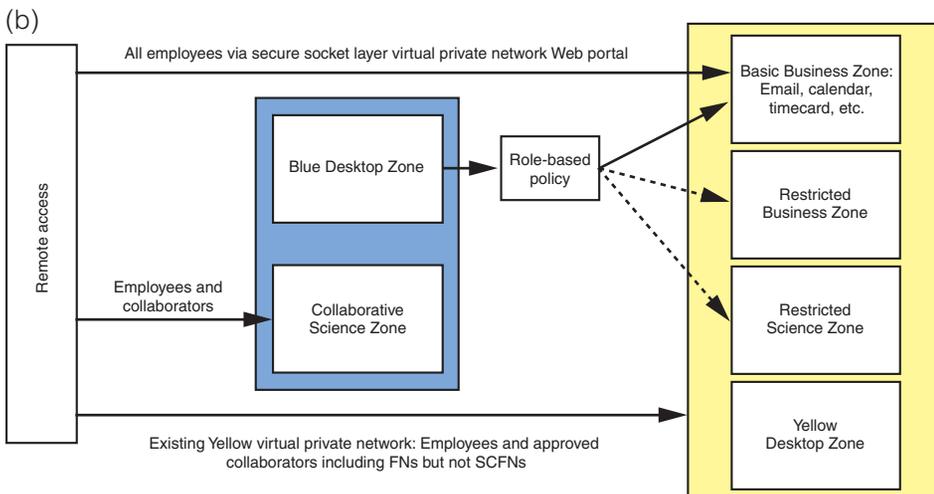
The Blue Network project is the first phase of a multiyear effort to deploy an integrated network access architecture (NAA) at LLNL. This project is sponsored by the CIO with funding from the Institutional Information Technology Program. The NAA provides a framework for compartmentalizing data, services, and users, taking into consideration business needs and the associated security exposure. By identifying

security zones and perimeter mechanisms between these zones, we will ensure that the desired security policies can be enforced. Additional goals of the NAA include implementing a two-factor authentication of devices to the network and detecting and remediating out-of-compliance devices.



(a) Network Security Zone Model.

(b) Blue Network Access Model.



PROGRESS IN 2006

The Blue Network project was initiated in August 2006. Since then, our focus has been on defining requirements, establishing security policy, and evaluating candidate network hardware.

Today, the majority of LLNL unclassified scientific and business systems reside on what is known as the Yellow network. The Yellow network is a restricted network that contains varying levels of sensitive but unclassified data, including Unclassified Controlled Information. Nearly all LLNL employees and many collaborators have local access to the Yellow network, including foreign nationals (FNs) and sensitive country foreign nationals (SCFNs). While many users of the Yellow network also have remote access, SCFNs are currently prohibited from having remote access. Within the Yellow network, access control is enforced to varying degrees at the host and/or application level. The Blue Network project will introduce a new network for SCFNs, FNs, and LLNL employees and collaborators who do not need

to access sensitive unclassified data. The Blue network will include end-user workstations and collaborative scientific resources. Users on the Blue network may gain access to select business and scientific resources on the Yellow network based on their identity and role. Local and remote access will be provided. This will serve as a significant step toward protecting our computing resources while improving our flexibility to provide access to information that our employees, sponsors, and collaborators need. In early 2007, we expect to finalize the architecture and begin deploying Blue network to beta users.

SIGNIFICANCE

The Blue Network project addresses long-standing issues surrounding how LLNL accommodates foreign national employees, who are valuable contributors to our Laboratory mission. To this end, we are creating a general-purpose network-access framework that will deliver an improved security posture as well as enhanced services and productivity for network users.

5.05 — Cyber Security Protects Against Evolving Threats



CONTACT INFORMATION

Edward Matsche, matsche1@llnl.gov

OVERVIEW

Information Technology Protection Division (ITPD) personnel support LLNL's CSP by providing technical and administrative expertise to enhance the security of LLNL's computing assets and ensure compliance with DOE orders and other requirements related to classified and unclassified computer and telecommunication security. In addition, ITPD staff manages LLNL's intrusion-detection and response capabilities, provides programming support for network authentication and Entrust encryption, and conducts cyber security awareness training.

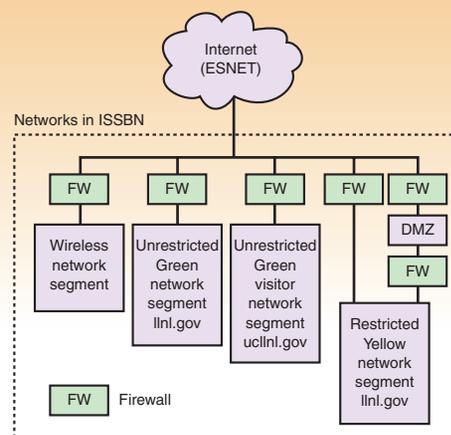
In 2006, LLNL's suite of network protection software was expanded to defend against the increasing threats associated with Web access. According to a recent security intelligence report from Microsoft, more than 43,000 new variants of malicious remote control software were detected in the first half of 2006 alone.

PROGRESS IN 2006

ITPD participated in three notable activities. First, the site-wide master security plan for the unclassified backbone network was accredited for three years. The goal of the security plan is to provide a secure and unclassified backbone network that enables scientific and technological missions without sacrificing intellectual property or compromising sensitive information. The plan establishes a baseline for the cyber security policies and practices that are required for LLNL's entire unclassified environment, particularly for the Institutional Support Services Backbone Network (ISSBN), general

desktop and server computing resources, and major applications and large-scale computing facilities. Many of Computation's software service elements were involved in this effort, including active directory infrastructure, configuration and patch management software, remote access services and firewalls, and various intrusion-detection and scanning mechanisms. Potential weaknesses within the security plan were documented and are being mitigated.

The second notable activity improved the suite of network protection software that guards against the increasing threats associated with Web access, including malicious remote control software. Blue Coat, a Web proxy service, deployed in August.



A simplified diagram of the four unclassified network segments and a demilitarized zone (DMZ) segment contained in the ISSBN.

It provides a mechanism for protecting LLNL users from booby-trapped sites that automatically exploit vulnerabilities and execute arbitrary code on a system without the user knowing. Blue Coat blocks access to bad sites (those known to host malware) and scans incoming Web content for viruses.

The third notable activity improved remote access to LLNL's unclassified networks by deploying a clientless virtual private network solution using secure socket layer (SSL) technology. This solution requires fewer institutional support resources and is enhancing security. Employees on business travel who need remote access will now be able to use any Web-based browser that supports Java and SSL. Information that is downloaded during the user's session is encrypted and automatically removed from the computer after logout.

SIGNIFICANCE

The new Web proxy service protects LLNL users from accessing potentially unsafe Web sites. Between August and November, Blue Coat blocked 245 attempts to download malicious code, 32,615 attempts to access Web sites that contain spyware, and more than two million attempts to visit sites that host malicious content.

The new site-wide master security plan reduced the number of accredited unclassified security plans at LLNL. Now, programs gain approval to operate subnets at LLNL through a documented process that aligns with National Institute of Standards and Technology requirements. These improvements help protect LLNL's networks and computer infrastructure from compromise and attack.

5.06 — Academic Partnerships Create Pipelines for People and Ideas

OVERVIEW

Computational science evolves so rapidly and along so many fronts that, to define the leading edge, the Laboratory engages researchers at many academic centers of excellence. Through its ISCR, jointly operated with the University Relations Program, Computation builds and maintains bridges to the academic community with a program of managed research, collaborative subcontracts, faculty visits, student internships, workshops, and several active seminar series. We host both short- and long-term visits, with activity peaking in the summer months.



Krzysztof Niski, a graduate student from Johns Hopkins University, demonstrates his summer research project, an image-rendering system operating in a distributed setting that displays extremely large data sets. The image behind him, rendered over 16 million triangles, is a portion of the U.S. Geological Survey earth data set, comprising 22 billion data points.

PROGRESS IN 2006

The ISCR supervised the Livermore–UC San Diego Scientific Data Study (LUSciD), an effort in which scientists from LLNL and UC San Diego created mechanisms for collaborative scientific research involving the massive data sets required for extremely large-scale scientific simulation. LUSciD is funded by the UC Office of the President. Three topics are included in this effort: climate modeling, astrophysics, and scientific data management.

The University Education Partnerships Program, which is operated through the ISCR, funded four educational projects at UC campuses in Berkeley, Los Angeles, and San Diego, as well as at San Diego State University. On each project, the faculty principal investigator and a graduate student worked closely with an LLNL collaborator.

Sixty-nine students from 48 colleges and universities and one high school participated in the 2006 Summer Scholar Program. The ISCR also hosted 10 students on graduate fellowships who were in residence the entire year. For this major outreach activity, Computation provides administrative support, facilities, programs that promote student-

mentor interaction, and lectures that enrich the educational experience of the students. Typically in residence for 12 weeks, each student works closely with LLNL scientists to develop software or conduct research that addresses Computation's missions—not academic exercises with known outcomes. Just over half of the summer students are doctoral-degree candidates. Most work with mentors in the Computing Applications and Research Department; however, their projects span many areas, from the National Ignition Facility to biological sciences, physics, electrical engineering, chemistry, aeronautics and astronautics, mechanical engineering, information science and systems, and, of course, computer science and applied mathematics. Interns from Computation accounted for approximately 20% of the posters presented at the annual LLNL Student Research Symposium. While primarily here for research, students are encouraged to attend as many seminars as possible.

The Institute for Terascale Simulation Distinguished Lecturer program brings international leaders in simulation to LLNL to deliver general lectures and engage



CONTACT INFORMATION

David Keyes, dekeyes@llnl.gov
 Van Henson, vhenson@llnl.gov
 Web site: www.llnl.gov/iscr

Laboratory scientists in technical discussions. We hosted Dr. Anthony Mezzacappa of Oak Ridge National Laboratory, who is the director of DOE's Terascale Supernova Initiative; Professor Marsha Berger of New York University's Courant Institute, who was recently inducted into the National Academy of Engineering for her pioneering work on adaptive mesh refinement algorithms and software; Professor David Ceperley of the University of Illinois, who was recently inducted into the National Academy of Sciences for his work on Quantum Monte Carlo methods; and Dr. David Shaw, the chief scientist of D.E. Shaw Research, who is developing a special-purpose supercomputer to speed up by two or more orders of magnitude the performance of molecular dynamics simulations of importance to pharmaceutical design. Altogether, Computation hosted 226 short-term visits, in addition to the 69 summer students.

Academic collaborations have positioned Computation to be one of the two largest players in the mathematics and computing components of DOE's Scientific Discovery through Advanced Computing (SciDAC) Program.



One way the ISCR facilitates collaborative science is by sponsoring scientific workshops. Here, the participants of the 2006 Workshop on Computational Methods in Transport gather after an off site dinner. The event drew 62 participants, representing 16 universities, 8 government laboratories, and 4 industrial partners, including organizations from Russia, Canada, and the European Union.

Many SciDAC projects are led by academic principal investigators. Computation participates in nine of the 13 projects that were announced in September. There is a strong correlation between pre-existing academic collaborations and successful SciDAC proposals. The research of two Computation postdoctorate researchers who moved on to faculty positions was honored with plenary talks at the 2006 SciDAC Annual Meeting.

Software breakthroughs that are delivered to applications

groups at the Laboratory are often initiated or catalyzed by academic visitors. For example, LLNL's *hypr*, a scalable solver library, now includes the first provably scalable solver for the positive semidefinite form of Maxwell's equations on general unstructured meshes, which is directly traceable to a March seminar by computational mathematician Jinchao Xu of Penn State.

SIGNIFICANCE

By managing joint research, identifying and hosting researchers from the academic community, and operating an active internship program, the ISCR fosters computational science collaborations for Computation and encourages long-term academic research agendas that address Laboratory research priorities. Ideas flow in both directions, and the Laboratory cultivates its future workforce.

5.07 — Techbase Investments Advance Core Competencies



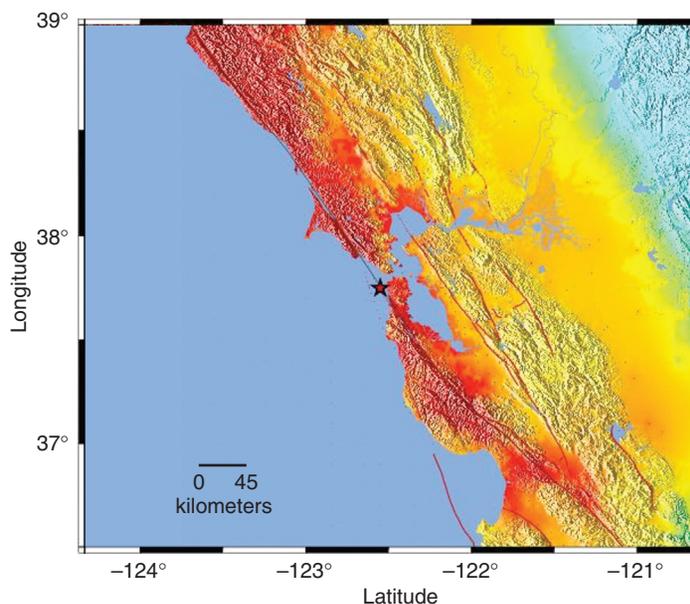
CONTACT INFORMATION

David L. Brown, dlb@lml.gov

OVERVIEW

To support the evolving computational needs of the Laboratory, Computation must continually improve its core mathematical and computer science competencies. The discretionary resources available to accomplish this come from the Laboratory Directed Research and Development program, which funds basic and applied research activities, and from Computation's internal technology base (techbase) program, which funds the reduction-to-practice of existing technologies, technological developments characterized by limited technical risk, or evaluations of promising technologies that may enhance Computation's effectiveness. Techbase projects are initiated either by management direction or through an internal review process where Computation personnel have an opportunity to propose new ideas for consideration. Decisions on what projects to fund are based on the technological needs of the directorate and technical quality of submitted proposals. The goal is to develop new or improved computational capabilities that help us support multiple Laboratory programs.

Computation's techbase program funded enhancements to the Wave Propagation Code (see Section 3.05), including a new materials database interface that enabled U.S. Geological Survey materials data to be used in this recreation of the 1906 San Francisco earthquake.



PROGRESS IN 2006

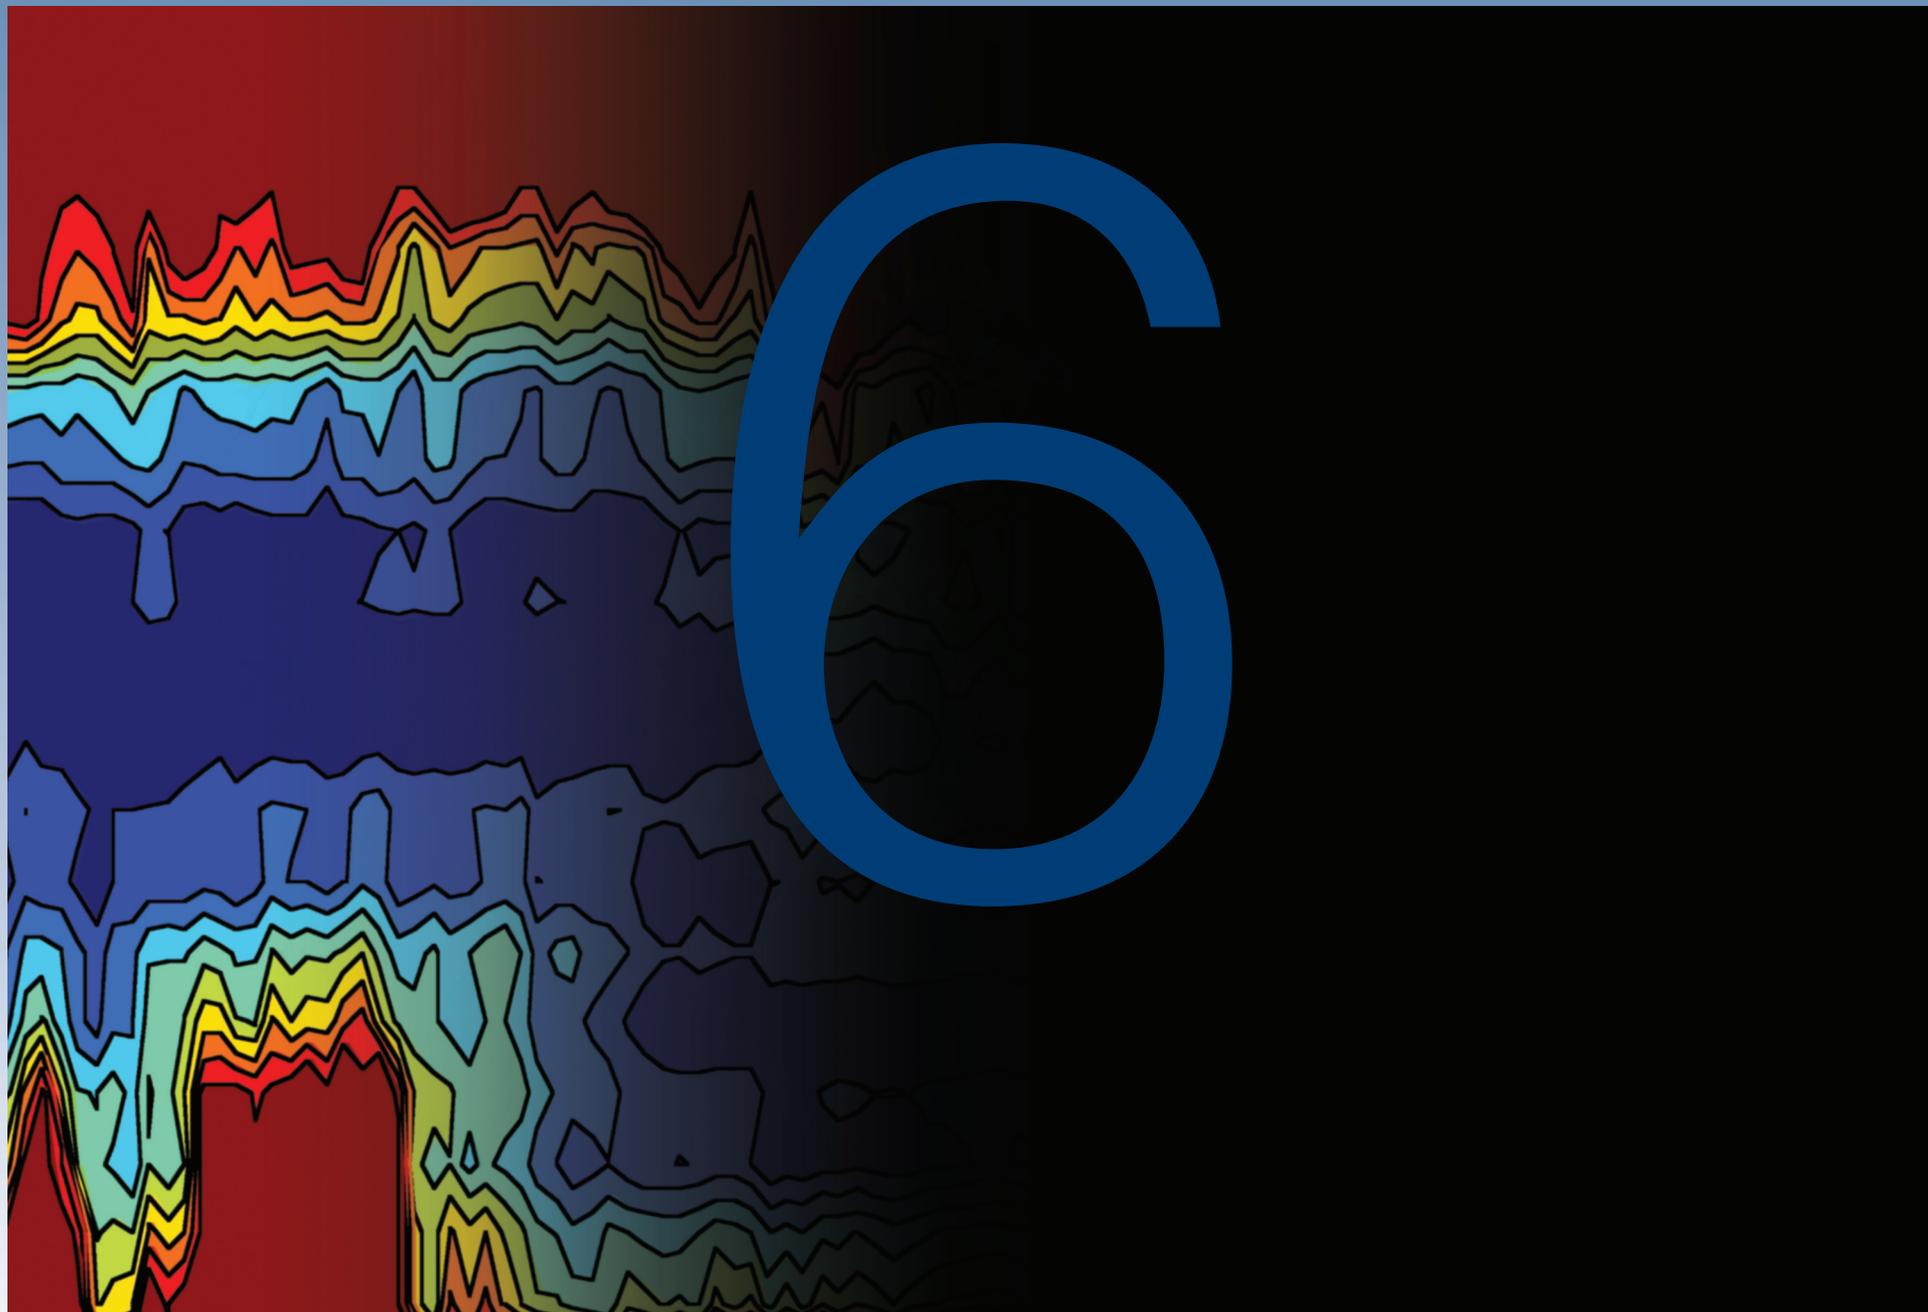
During the past year, our techbase investments were aimed at improving our capabilities in scalable simulation, biodefense computing, data sciences, and data-intensive computing. One particularly effective techbase-funded activity involved making enhancements to the Computation-developed, multiprogrammatic elastic wave simulation code. The new features enhance our ability to support several programs at the Laboratory, including the Enhanced Test Site Readiness program; the Yucca Mountain program; and nondestructive evaluation programs, such as those that study optics imperfections at the National Ignition Facility. Now, we can better support the Laboratory's nonproliferation mission because of enhanced underground explosion modeling and underground facilities characterization capabilities.

Computation's techbase program funded the evaluation of various software tools aimed at improving many of the Laboratory's large-scale simulation codes. In addition, the program supported the development of optimized general-purpose algorithms and

tools for protein structure modeling and enhanced technology for DNA melting temperature prediction, both of which will enhance our ability to support the Laboratory's biodefense programs and the DOE Joint Genome Institute. Techbase funding was also used to evaluate new computing hardware options that might support existing and future Laboratory programs, including the IBM Cell Broadband Engine, the Netezza Performance Server® data warehouse appliance, and various graphics processing units, such as the PhysX™ gaming processor from AGEIA Technologies, Inc.

SIGNIFICANCE

Computation's techbase investments allow us to establish, maintain, and improve core competencies that enhance our ability to provide programs at the Laboratory with state-of-the-art computational technologies. The techbase program allows us to evaluate new software technologies and incorporate them into our software base. We also use techbase to make our staff more effective by providing them with enhanced capabilities and enabling them to educate themselves about state-of-the-art computational tools.



6.01 — Computation Directorate Industrial Collaborators

Company	Topic	LLNL contact
IBM	Scalable systems, multiple areas	Mark Seager and others
IBM	High-performance storage system	Mark Gary
Appro	Scalable capacity clusters	Mark Seager
Appro, Supermicro	IPMI cluster management tools	Albert Chu
Cray	Scalable systems, open source strategy	Mark Seager
Etnus	Debuggers	Dong Ahn
OpenWorks	Valgrind memory tool development	John Gyllenhaal
SGI, Krell Institute	Performance tool PathForward	Martin Schulz
Kitware	Visualization toolkit	Kathleen Bonnell
CEI-EnSight	Visualization package	Eric Brugger
Tungsten Graphics	Chromium render server	Eric Brugger
R3vis Corporation	Chromium render server	Eric Brugger
Hewlett-Packard	Scalable global file system PF	Brent Gorda
Hewlett-Packard, Bull, NUDT	SLURM resource management software	Morris Jette
Cluster File Systems	Lustre file system deployment	Brent Gorda
Red Hat	Operating systems	Mark Seager
OpenMP Consortium	Shared memory programming models	Bronis de Supinski
Schlumberger	Synthetic seismic data set generation	Shawn Larsen
URS Corporation	Earthquake hazard estimation	Shawn Larsen
Society of Exploration Geophysicists	Advanced subsurface modeling	Shawn Larsen
Voltaire, OpenIB Fabrics	InfiniBand Linux kernel software	Matt Leininger
ParaTools, Inc.	Performance tools for HPC systems	Chris Chambreau
IQuum	DNA signature target design	Tom Slezak
Liquid Computing	Hardware alpha testing	Brent Gorda
ClearSpeed	Hardware beta testing	Brent Gorda
DataDirect Networks	Raid 6 R&D for I/O systems	Brent Gorda
Cluster Resources, Inc.	MOAB workload manager	Don Lipari
Limit Point Systems	LibSheaf	Bill Arrighi

6.02 — Academic Outreach

University	Faculty	Activity type	Topic	Funding source	LLNL contact
UC Berkeley	Chris McKee	Joint Research	Computational Astrophysics Consortium	DOE SciDAC	Louis Howell
UC Berkeley	Heidi Kuzma	Subcontract	General-purpose graphics processing unit computing	DOD	David Bremer
UC Berkeley	James Demmel	Joint Research	Terascale-optimal partial-differential-equation solvers	DOE SciDAC	Lori Diachin
UC Berkeley	Jonathan Wurtele	Subcontract	Theoretical and numerical investigations of Raman backscatter	G&A	Richard Berger
UC Davis	Bernd Hamann	Subcontract	Analysis and visualization of scientific data using topology-based methods	LDRD	Valerio Pascucci
UC Davis	Bernd Hamann and Kenneth I. Joy	Joint Research	Visualization and Analytics Center for Enabling Technologies	DOE SciDAC	Valerio Pascucci
UC Davis	Bertram Ludaescher	Joint Research	Scientific Data Management Center for Enabling Technologies	DOE SciDAC	Terence Critchlow
UC Davis	Dave Wittman	Joint Research	Large Synoptic Survey Telescope long-term data analysis	LDRD	Don Dossa
UC Davis	Greg Miller	Subcontract	Algorithm development for interface tracking in a compressible medium	DOE MICS	David Trebotich
UC Davis	Kenneth I. Joy	Subcontract	Real-time three-dimensional viewing of large time-dependent aerial imagery	DOE NA-22	Mark Duchaineau
UC Davis	Kwan-Liu Ma	Subcontract	Interactive visualization of large mixed-resolution volume data	ASC-PPPE	Steve Louis
UC Davis	Kwan-Liu Ma	Joint Research	Large-graph data visualization	LDRD	Valerio Pascucci
UC Davis	Kwan-Liu Ma	Subcontract	Large-graph visualization for homeland security	DOE-OCIO	Tony Bartoletti
UC Los Angeles	Alan Laub	Multilocation Appointment	Computational Science and Engineering Consortium	OPC	Van Emden Henson
UC Los Angeles	Heinz-Otto Kreiss	Consultant	Adaptive methods for partial differential equations	DOE MICS	Lori Diachin
UC Los Angeles	Kendall Houk	Subcontract	Using a combined QM/MM approach with CPMD-rare event methods to study enzymatic reaction mechanisms	G&A	Christopher Mundy
UC Los Angeles	Mark Green	Industrial Mentor	Research in industrial projects for students	DHS	Tina Eliassi-Rad
UC Los Angeles and Weizman Institute (Israel)	Achiezer Brandt	Consultant	Geometric and algebraic multigrid techniques	DOE MICS	Rob Falgout
UC San Diego	Allan Snavely	Joint Research	Performance Engineering Research Institute	DOE SciDAC	Bronis de Supinski
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	Ilkay Altintas	Joint Research	Scientific Data Management Center for Enabling Technologies	DOE SciDAC	Terence Critchlow
UC San Diego	Michael Holst and Dan Reynolds	Subcontract	Implicit solvers and preconditioning techniques for simulations of magnetohydrodynamics and core-Collapse supernovae	DOE MICS	Carol Woodward
UC San Diego	Randy Bank	Consultant	Numerical solutions of partial differential equations, multilevel iterative methods, and adaptive grid generations	ASC ITS	Rob Falgout
UC San Diego	Reagan Moore, Tim Barnett, Michael Norman, and Randy Bank	Subcontract	LLNL-UCSD scientific data management	UCDRD	Van Emden Henson

6.02 — Academic Outreach (cont.)

UC San Diego	Scott Baden	Subcontract	Data-driven execution of latency-tolerant algorithms	G&A	Dan Quinlan
UC San Francisco	David Saloner	Joint Research	Use of computational microfluidics tools to model the hemodynamics of vascular disease	DOE MICS	David Trebotich
UC Santa Barbara	Linda Petzold	Consultant	Numerical methods for sensitivity analysis of DAE systems; algorithmic implementation of methods aimed at large-scale applications	DOE MICS, ASC ITS	Carol Woodward
UC Santa Cruz	Darrell Long	Subcontract	Scalable file systems for high-performance computing	ASC-PSE	Bill Loewe
UC Santa Cruz	Stan Woosley	Joint Research	Computational Astrophysics Consortium	DOE SciDAC	Louis Howell
Brown University	Andries van Dam	Subcontract	Developing immersive VR techniques for interactive scientific visualization	ASC-PPPE	Steve Louis
California Institute of Technology	Daniel I. Meiron	ASC Alliance	Center for Simulating the Dynamic Response of Materials	ASC	Dick Watson
Carnegie Mellon University	Christos Faloutsos	Subcontract	Group discovery in time-evolving complex networks	LDRD	Tina Eliassi-Rad
Carnegie Mellon University	Omar Ghattas	Joint Research	Terascale-optimal partial-differential-equation solvers	DOE SciDAC	Lori Diachin
Columbia University	David Keyes	Consultant	Algorithms for the solution of partial differential equations on massively parallel computers; general high-performance computing	Overhead	Steve Ashby
Columbia University	David Keyes	Joint Research	Terascale-optimal partial-differential-equation solvers	DOE SciDAC	Lori Diachin
Columbia University	David Keyes	Subcontract	Acting Director, Institute for Scientific Computing Research	ASC ITS, Overhead	Steve Ashby
Cornell University	Sally McKee	Sabbatical	PERI, ASC CSSE	ASC Program	Bronis de Supinski
Duke University	Herbert Edelsbrunner	Joint Research	Discrete methods for computing continuous functions	LDRD	Valerio Pascucci
Florida Institute of Technology	Jim Jones	Subcontract	Multigrid approaches for systems of partial differential equations	ASC ITS	Rob Falgout
Georgia Institute of Technology	Jaroslav (Jarek) Rossignac	Subcontract	Advanced prediction for the compression of regular and irregular data sets	LDRD, Overhead	Peter Lindstrom
Krell Institute	Lucille Kilmer	Subcontract	Department of Energy High-Performance Computer Science Fellowship Program—LLNL portion	ASC ITS	John May
McMaster University	Nedialko Nedialkov	Sabbatical	Nonlinear solves and differential equations	ASC ITS	Radu Serban
New York University	Michael Shelley	Joint Research	Numerical methods for fluid dynamics with free boundaries	DOE MICS	Petri Fast
New York University	Olof Widlund	Joint Research	Terascale-optimal partial-differential-equation solvers	DOE SciDAC	Lori Diachin
North Carolina State University	Mladen Vouk	Joint Research	Scientific Data Management Center for Enabling Technologies	DOE SciDAC	Terence Critchlow
Northwestern University	Alok Choudhary	Joint Research	Scientific Data Management Center for Enabling Technologies	DOE SciDAC	Terence Critchlow
Northwestern University	Alok Choudhary	Subcontract	Support for enhanced dyninst testing and initial steps toward open binary editing environment	ASC-PSE	Bill Loewe
Ohio State University	Umit V. Catalyurek	Subcontract	Clustering streaming graph data	DHS	Andy Yoo

Ohio Supercomputer Center	Kevin Wohlever	Joint Research	Data-intensive computing	ASC	Steve Louis
Pacific Northwest National Laboratory	Eric Stephan	Subcontract	DHS IDS Data Integration and Dissemination Workshop	DHS	David Brown
Penn State University	Ludmil Zikatanov	Subcontract	Optimal AMG interpolation and convergence theory	ASC ITS	Rob Falgout
Portland State University	Karen Karavanic	Subcontract	Performance analysis infrastructure for petascale applications	ASC-PSE	John May
Princeton University	Kai Li	Subcontract	R&D scalable display software	ASC-PPPE	Steve Louis
Purdue University	Ziqiang Cai	Summer Faculty	Multigrid methods for the H(div)-type pseudostress system of the Stokes equations	ASC ITS	Barry Lee
Rensselaer Polytechnic Institute	Mark Shephard	Joint Research	Terascale simulation tools and technologies	DOE SciDAC	Lori Diachin
Rensselaer Polytechnic Institute	Don Schwendeman	Subcontract	Development of numerical methods	DOE MICS	Bill Henshaw
Rensselaer Polytechnic Institute	Don Schwendeman	Subcontract	Development of numerical methods for mathematical models of high-speed reactive and nonreactive flow	DOE MICS	Bill Henshaw
Rice University	John Mellor-Crummey	Subcontract	Open-source software technology for transforming scientific problems	DOE MICS	Dan Quinlan
Rice University	John Mellor-Crummey	Joint Research	Performance Engineering Research Institute	DOE SciDAC	Bronis de Supinski
San Diego State University	Calvin Johnson	Subcontract	Frontier computations in the structure of atomic nuclei	G&A	Erich Ormand
Stanford University	Parvis Moin	ASC Alliance	Center for Integrated Turbulence Simulations	ASC	Dick Watson
Stanford University	Pat Hanrahan	Subcontract	General purpose graphics processing units computing	DOD	David Bremer
Stanford University	Pat Hanrahan	Subcontract	PC-based visualization clusters and software	ASC-PPPE	Steve Louis
Stanford University	Pat Hanrahan	Subcontract	Scalable visualization of large time-varying data sets	ASC	Holger Jones
Stanford University	Stuart Marshall	Joint Research	Large Synoptic Survey Telescope camera data output hardware and imaging formats	LDRD	Don Dossa
State University of New York, Stony Brook	Jim Glimm	Joint Research	Terascale simulation tools and technologies	DOE SciDAC	Lori Diachin
State University of New York, Stony Brook	Doug Swesty	Joint Research	Solvers for supernova simulation	DOE SciDAC	Carol Woodward
Technion-Israel Institute of Technology	Irak Yavneh	Summer Faculty	Novel applications for algebraic multigrid	DOE MICS	

6.02 — Academic Outreach (cont.)

Texas A&M University	Joe Pasciak	Subcontract	Algebraic multigrid for problems for h-curl	DOE MICS	Panayot Vassilevski
Texas State University	Anne Ngu	Summer Faculty	Working with local and external collaborators to enhance workflow-based infrastructure	DOE SciDAC	Terence Critchlow
University of Arizona	Adam Burrows	Joint Research	Computational Astrophysics Consortium	DOE SciDAC	Louis Howell
University of Arizona	Tim Axelrod	Joint Research	LSST data management and high-speed data acquisition and analysis	LDRD	Don Dossa
University of Chicago	Don Lamb	ASC Alliance	Center for Astrophysical Thermonuclear Flashes	ASC	Dick Watson
University of Colorado	John Ruge	Subcontract	Algebraic multigrid, spectral AMG, FOSPACK, applications to ALE3D problems	ASC ITS	Rob Falgout
University of Colorado	Steve McCormick	Joint Research	Terascale-optimal partial-differential-equation solvers	DOE SciDAC	Lori Diachin
University of Colorado	Steve McCormick	Subcontract	FOSPACK, AMG, adaptive AMG, compatible relaxation, parallelization, and application codes	ASC ITS, DOE MICS	Rob Falgout
University of Colorado	Tom Manteuffel	Joint Research	Geometric and algebraic multigrid techniques; FOSLS approach to solving partial differential equations	ASC ITS	Rob Falgout
University of Colorado	Tom Manteuffel	Joint Research	Solution methods for transport problems	ASC	Peter Brown
University of Delaware	Richard Braun	Joint Research	Free boundary numerics and thin film fluid flow using moving overset grids	DOE MICS	Petri Fast
University of Illinois, Urbana-Champaign	John C. Hart	Subcontract	Topological analysis for Rayleigh–Taylor instabilities	LDRD	Valerio Pascucci
University of Illinois, Urbana-Champaign	Laxmikant Kale	Joint Research	FastOS	DOE MICS	Terry Jones
University of Illinois, Urbana-Champaign	Michael T. Heath	ASC Alliance	Center for Simulation of Advanced Rockets	ASC	Dick Watson
University of Illinois, Urbana-Champaign	Mike Folk	Subcontract	HDF5 parallel hierarchical data formats	ASC-DVS	Eric Brugger
University of Maryland	Jeff Hollingsworth	Joint Research	Performance Engineering Research Institute	DOE SciDAC	Bronis de Supinski
University of Massachusetts	David Jensen	Subcontract	Relational pathfinding	DHS	Tina Eliassi-Rad
University of Michigan	Peter Honeyman	Subcontract	NFSv4 scalable performance, security, and wide-area optimization	ASC-PSE	Bill Loewe
University of Minnesota	Douglas Arnold	Participating Institution	Institute for Mathematics and Its Applications	ASC ITS	Van Emden Henson
University of New Mexico	Anil Prinja	Subcontract	Validation and verification for radiation and neutron transport	ASC	Peter Brown

University of New Mexico	Thomas Hagstrom	Sabbatical	High-order structure grid methods for wave propagation on complex unbounded domains	DOE MICS	Bill Henshaw
University of North Carolina	Dan Reed and Rob Fowler	Joint Research	Performance Engineering Research Institute	DOE SciDAC	Bronis de Supinski
University of North Carolina	Henry Fuchs	Subcontract	Continued group tele-immersion research	ASC-PPPE	Steve Louis
University of San Francisco	Benjamin (Pete) Wells	Summer Faculty	General issues in computer science involving sequential machines, string algorithms, foundations of programming, etc.	ASC ITS	John May
University of Southern California	Bob Lucan and Mary Hall	Joint Research	Performance Engineering Research Institute	DOE SciDAC	Bronis de Supinski
University of Tennessee	Jack Dongarra	Joint Research	Terascale-optimal partial-differential-equation solvers	DOE SciDAC	Lori Diachin
University of Tennessee	Jack Dongarra and Shirley Moore	Joint Research	Performance Engineering Research Institute	DOE SciDAC	Bronis de Supinski
University of Texas at Austin	Clint Dawson	Subcontract	Discretizations and splitting methods for radiation-diffusion problems	ASC ITS	Carol Woodward
University of Texas at Austin	Kazushige Goto	Subcontract	Optimization and tuning of linear algebra libraries for LLNL IBM systems	ASC-PSE	Mark Seager
University of Texas at San Antonio	Qing Yi	Subcontract	Program analysis and loop optimization	DOE MICS	Dan Quinlan
University of Utah	Charles Hansen	Consultant	Data exploration, multiresolution scientific data visualization, and algorithm design	ASC DVS	Mark Duchaineau
University of Utah	Charles Hansen, Christopher Johnson, Steven G. Parker, and Xavier Tricoche	Joint Research	Visualization and Analytics Center for Enabling Technologies	DOE SciDAC	Valerio Pascucci
University of Utah	Claudio Silva	Subcontract	Studying the topology of point-set surfaces	LDRD	Valerio Pascucci
University of Utah	David W. Pershing	ASC Alliance	Center for the Simulation of Accidental Fires and Explosions	ASC	Dick Watson
University of Utah	Ellen Riloff	Subcontract	Bioforensics text extraction	DHS	Terence Critchlow
University of Utah	Steve Parker	Joint Research	Scientific Data Management Center for Enabling Technologies	DOE SciDAC	Terence Critchlow
University of Waterloo	Hans de Sterck	Subcontract	Improving complexity of parallel AMG	DOE SciDAC	Ulrike Yang
University of Wisconsin	Bart Miller	Subcontract	Support for enhanced dyninst testing and initial steps toward open binary editing environment	ASC-PSE	Martin Schulz
University of Wisconsin	Jason Kraftcheck	Subcontract	Mesquite software development	DOE MICS	Lori Diachin
Worcester Polytechnic Institute	Homer Walker	Consultant	Development and analysis of iterative methods for large-scale linear and nonlinear algebraic systems	ASC ITS	Carol Woodward

6.03 — Publications

BOOK CHAPTERS

- Dahlgren, T. (2006), "Parallel adaptive mesh refinement," in *Parallel Processing for Scientific Computing*, M. Heroux, P. Raghavan, and H. Simon, Eds. (SIAM Publications: Philadelphia, PA), pp. 143–162 (UCRL-BOOK-205864).
- Diachin, L. A. and R. D. Hornung (2006), "Parallel PDE-based simulations using the Common Component Architecture," in *Numerical Solutions of Partial Differential Equations on Parallel Computers (Lecture Notes in Computational Science and Engineering)*, A. M. Bruaset and A. Tveito, Eds. (Springer-Verlag: New York, NY), pp. 327–384 (UCRL-BOOK-208032).
- Falgout, R. D. and U. M. Yang (2006), "A survey of parallelization techniques for multigrid solvers," in *Parallel Processing for Scientific Computing*, M. Heroux, P. Raghavan, and H. Simon, Eds. (SIAM Publications: Philadelphia, PA) (UCRL-BOOK-205864).
- Falgout, R. D., J. E. Jones, and U. M. Yang (2006), "The design and implementation of *hypre*, a library of parallel high-performance preconditioners," in *Numerical Solutions of Partial Differential Equations on Parallel Computers (Lecture Notes in Computational Science and Engineering)*, A. M. Bruaset and A. Tveito, Eds. (Springer-Verlag: New York, NY) pp. 267–294 (UCRL-BOOK-208032).
- Yang, U. M. (2006), "Parallel algebraic multigrid methods—high-performance preconditioners," in *Numerical Solutions of Partial Differential Equations on Parallel Computers (Lecture Notes in Computational Science and Engineering)*, A. M. Bruaset and A. Tveito, Eds. (Springer-Verlag: New York, NY) pp. 209–236 (UCRL-BOOK-208032).

JOURNAL PAPERS

- Allan, B. A., et al. (2006), "A component architecture for high-performance scientific computing," *Int. J. High-Performance Computing Applications* **20** (2), 163–202 (UCRL-JRNL-208635).
- Appelo, D., et al. (2006), "A new absorbing layer for elastic waves," *J. Comput. Phys.* **215** (2), 642–660
- Armstrong, R., et al. (2006), "The CCA component model for high-performance scientific computing," *Concurr. Comp—Pract. E.* **18**, 215–229 (UCRL-JRNL-208616).
- Baker, A. H., S. Crivelli, and E. R. Jessup (2006), "An efficient parallel termination detection algorithm," accepted by *International Journal of Parallel, Emergent, and Distributed Systems* (UCRL-JRNL-204452).
- Baker, A. H., J. M. Dennis, and E. R. Jessup (2006), "On improving linear solver performance: A block variant of GMRES," *SIAM J. Sci. Comput.* **27**, 1608–1626 (UCRL-JRNL-204081).
- Baker, A. H., R. D. Falgout, and U. M. Yang (2006), "An assumed partition algorithm for determining processor intercommunication," *Parallel Comput.* **32** (5–6), 394–414 (UCRL-JRNL-215757).
- Bala, G., et al. (2006), "Biogeophysical effects of CO₂ fertilization on global climate," *Tellus Series B Chem. Phys. Meteorol.* **58** (5), 620–627 (UCRL-JRNL-221030).
- Beller, H. R., et al. (2006), "The genome sequence of the obligately chemolithoautotrophic, facultatively anaerobic bacterium *Thiobacillus denitrificans*," *J. Bacteriol.* **188**, 1473–1488.
- Beller, H. R., et al. (2006), "Whole-genome transcriptional analysis of chemolithoautotrophic thiosulfate oxidation by *Thiobacillus denitrificans* under aerobic versus denitrifying conditions," *J. Bacteriol.* **188** (19), 7005–7015 (UCRL-JRNL-220798).
- Berman, E. S. F., et al. (2006), "Distinguishing monosaccharide stereo- and structural isomers with TOF-SIMS and multivariate statistical analysis," *Anal. Chem.* **78** (18), 6497–6503 (UCRL-JRNL-221140).
- Beuning, P. J., et al. (2006), "A noncleavable UmuD variant that acts as a UmuD' mimic," *J. Biol. Chem.* **281** (14), 9633–9640 (UCRL-JRNL-216587).
- Brezina, M., et al. (2006), "Adaptive algebraic multigrid methods," *SIAM J. Sci. Comput.* **27** (4), 1261–1286 (UCRL-JRNL-203501).
- Brezina, M., C. Tong, and R. Becker (2006), "Parallel algebraic multigrids for structural mechanics," *SIAM J. Sci. Comput.* **27**, 1534–1554 (UCRL-JRNL-204167).
- Bringa, E. M., et al. (2006), "Shock deformation of face-centred-cubic metals on subnanosecond timescales," *Nat. Mater.* **5** (10), 805–809.
- Callahan, S. P., et al. (2006), "Progressive volume rendering of large unstructured grids," *IEEE T. Vis. Comput. Gr.* **12** (5), 1307–1314.
- Chang, B. (2006), "A photon-free method to solve radiation transfer equations," accepted by *J. Comput. Phys.* (UCRL-JRNL-218676).
- De Sterck, H., U. M. Yang, and J. J. Heys (2006), "Reducing complexity in parallel algebraic multigrid preconditioners," *SIAM J. Matrix Analysis and Applications* **27**, 1019–1039 (UCRL-JRNL-206780).
- Dentener, F., et al. (2006), "The global atmospheric environment for the next generation," *Environ. Sci. Technol.* **40**, 3586–3594 (UCRL-JRNL-217619).
- D'haeseleer, P. (2006), "What are DNA sequence motifs?" *Nature Biotechnol.* **24**, 423–425 (UCRL-JRNL-218978).
- Diachin, L. F., et al. (2006), "A comparison of two optimization methods for mesh quality improvement," *Eng. Computation.* **22**, 61–74 (UCRL-JRNL-221001).
- Dobrev, V. A., et al. (2006), "Two-level preconditioning of discontinuous Galerkin approximations of second-order elliptic equations," *Numer. Linear Algebr.* **13** (9), 753–770 (UCRL-JRNL-219141).
- Dong, S., et al. (2006), "Spectral surface quadrangulation," *ACM Trans. Graphics* **25** (3), 1057–1066.
- Falgout, R. D. (2006), "An introduction to algebraic multigrid," *Comput. Sci. Eng.* **8** (6), 24–33 (UCRL-JRNL-220851).
- Falgout, R. D., J. E. Jones, and U. M. Yang (2006), "Conceptual interfaces in *hypre*," *Future Gener. Comp. Sy.* **22**, 239–251 (UCRL-JRNL-148957).
- Fast, P. and M. J. Shelley (2006), "Moore's law and the Saffman-Taylor instability," *J. Comput. Phys.* **212**, 1–5 (UCRL-JRNL-208653).
- Fattebert, J.-L., et al. (2006), "Finite element approach for density functional theory calculations on locally refined meshes," accepted by *J. Comput. Phys.*
- Fattebert, J.-L. and F. Gygi (2006), "Linear-scaling first-principles molecular dynamics with plane-waves accuracy," *Phys. Rev. B* **73**, 115124 (UCRL-JRNL-216052).
- Fernandez, J. C., et al. (2006), "Gas-filled hohlraum experiments at the National Ignition Facility," *Phys. Plasmas* **13**, 056318.
- Fournier, K. B., et al. (2006), "Electron-density scaling of conversion efficiency of laser energy into L-shell x-rays," *J. Quant. Spectrosc. Ra.* **99**, 186–198 (UCRL-JRNL-209373).
- Friedlingstein, P., et al. (2006), "Climate-carbon cycle feedback analysis: Results from the (CMIP)-M-4 model intercomparison," *J. Climate* **19**, 3337–3353.
- Garrido, I., et al. (2006), "Iterative schemes for time parallelization with application to reservoir simulation," *Math. Comput.* **75**, 1403–1428 (UCRL-JRNL-211688).
- Gatti-Bono, C., et al. (2006), "An anelastic allspeed projection method for gravitationally stratified flows," accepted by *J. Comput. Phys.*
- Ghasemi, J., et al. (2006), "A new algorithm for the determination of protolytic constants from spectrophotometric data in multiwavelength mode: Calculations of acidity constants of 4-(2-pyridylazo)resorcinol (PAR) in mixed nonaqueous-water solvents," *Talanta* **68**, 1201.
- Griebel, M., D. Oeltz, and P. S. Vassilevski (2006), "Space-time approximation with sparse grids," *SIAM J. Sci. Comput.* **28** (2), 701–727 (UCRL-JRNL-211686).

- Griffith, B. E., et al. (2006), "An adaptive, formally second-order accurate version of the immersed boundary method," accepted by *J. Comput. Phys.*
- Gunney, B. T. N., A. M. Wissink, and D. A. Hysom (2006), "Parallel clustering algorithms for structured AMR," *J. Parallel Distr. Com.* **66** (11), 1419–1430 (UCRL-JRNL-216580).
- Henshaw, W. D. (2006), "A high-order accurate parallel solver for Maxwell's equations on overlapping grids," *SIAM J. Sci. Comput.* **28** (5), 1730–1765 (UCRL-JRNL-215684).
- Henshaw, W. D. and D. W. Schwendeman (2006), "Moving overlapping grids with adaptive mesh refinement for high-speed reactive and nonreactive flow," *J. Comput. Phys.* **216** (2), 744–779 (UCRL-JRNL-215227).
- Hornung, R. D., A. M. Wissink, and S. R. Kohn (2006), "Managing complex data and geometry in parallel structured AMR applications," accepted by *Eng. Comput.* (UCRL-JRNL-214559).
- Jankun-Kelly, M., et al. (2006), "Vortex visualization for practical engineering applications," *IEEE T. Vis. Comput. Gr.* **12** (5), 957–964.
- Jones, J. and B. Lee (2006), "A multigrid method for variable coefficient Maxwell's equations," *SIAM J. Sci. Comput.* **27** (5), 1689–1708 (UCRL-JRNL-204137).
- Kim, T., et al. (2006), "Mesh-independent convergence of the modified inexact Newton method for a second order non-linear problem," *Numer. Linear Algebr.* **13** (1), 23–47 (UCRL-JRNL-206774).
- Kokko, E. J., et al. (2006), "As-built modeling of objects for performance assessment," *J. Comput. Inf. Sci. Eng.* **6** (4), 405–417 (UCRL-JRNL-215446).
- Kolev, T. V. and P. S. Vassilevski (2006), "AMG by element agglomeration and constrained energy minimization interpolation," *Numer. Linear Algebr.* **13** (9), 771–788 (UCRL-JRNL-219462).
- Kolev, T. Z., J. E. Pasciak, and P. S. Vassilevski (2006), "H(curl) auxiliary mesh preconditioning," accepted by *Numer. Linear Algebr.* (UCRL-JRNL-224227).
- Kreiss, H.-O., and N. A. Petersson (2006), "A second-order accurate embedded boundary method for the wave equation with Dirichlet data," *SIAM J. Sci. Comput.* **27**, 1141–1167 (UCRL-JRNL-202686).
- Kreiss, H.-O., and N. A. Petersson (2006), "An embedded boundary method for the wave equation with discontinuous coefficients," *SIAM J. Sci. Comput.* **28**, 2054–2074 (UCRL-JRNL-215702).
- Kulp, K. S., et al. (2006), "Essiac and Flor-Essence herbal tonics stimulate the in vitro growth of human breast cancer cells," *Breast Cancer Res. Treat* **98**, 249–259 (UCRL-JRNL-216009).
- Laney, D., et al. (2006), "Understanding the structure of the turbulent mixing layer in hydrodynamic instabilities," *IEEE T. Vis. Comput. Gr.* **12** (5), 1053–1060 (UCRL-JRNL-222696).
- Lawler, J. V., et al. (2006), "Cynomolgus macaque as an animal model for severe acute respiratory syndrome," *PLoS Med.* **3**, 677–686.
- Lindstrom, P. and M. Isenburg (2006), "Fast and efficient compression of floating-point data," *IEEE T. Vis. Comput. Gr.* **12** (5), 1245–1250 (UCRL-JRNL-220406).
- Liu, L., et al. (2006), "XWRAPComposer: A multipage data extraction service," *International Journal of Web Services Research* **3** (2), 33–60 (UCRL-JRNL-218270).
- London, R. A., et al. (2006), "Low mode control of cryogenic ICF fuel layers using infrared heating," *Fusion Science and Technology*, **49**, 608–615 (UCRL-JRNL-213460).
- Loots, G. G., et al. (2006), "Array2BIO: A comprehensive suite of utilities for the analysis of microarray data," *BMC Bioinformatics* **7** (307), 1471–1478 (UCRL-JRNL-219041).
- Kubista, M., et al. (2006), "The real-time polymerase chain reaction," *Mol. Aspects Med.* **27**, 95–125.
- Ma, J., et al. (2006), "Multiscale simulation using generalized interpolation material point (GIMP) method and molecular dynamics (MD)," *Computer Modeling in Engineering and Sciences* **14** (2), 101–117.
- Malfatti, M. A., et al. (2006), "The urinary metabolite profile of the dietary carcinogen 2-amino-1-methyl-6-phenylimidazo[4,5-b]pyridine is predictive of colon DNA adducts after a low-dose exposure in humans," *Cancer Res.* **66** (21), 10541–10547.
- Manteuffel, T. A., et al. (2006), "Spatial multigrid for isotropic neutron transport," accepted by *SIAM J. Sci. Comput.*
- Marathe, J., F. M. Mueller, and B. de Supinski (2006), "Analysis of cache coherence bottlenecks with hybrid hardware/software techniques," *ACM Transactions on Architecture and Code Optimization* **3** (4), 390–423 (UCRL-JRNL-222439).
- Max, N. (2006), "Hexahedron projection for curvilinear grids," accepted by *Journal of Graphics Tools* (UCRL-JRNL-222234).
- McDonald, J. W., et al. (2006), "Hard x-ray and hot electron environment in vacuum hohlraums at the National Ignition Facility," *Phys. Plasmas* **13**.
- McEnerney, J., et al. (2006), "Convex polarities over ordered fields," accepted by *J. Pure Appl. Algebra*.
- McEnerney, J., et al. (2006), "Alternative evidence for non-negativity," accepted by *J. Pure Appl. Algebra*.
- Natarajan, V., et al. (2006), "Segmenting molecular surfaces," *Comput. Aided Geom. D.* **23**, 495–509.
- Nevins, W. M., et al. (2006), "Characterizing electron temperature gradient turbulence via numerical simulation," *Phys. Plasmas* **13** (12).
- Nguyen, D. H., and P. D'haeseleer (2006), "Deciphering principles of transcription regulation in eukaryotic genomes," *Molecular Systems Biology* **2**.
- Nishtala, R., et al. (2006), "When cache block sparse matrix vector multiply works and why," accepted by *Appl. Algebr. Eng. Comm.: Special Issue on Computational Linear Algebra and Sparse Matrix Computations* (UCRL-JRNL-217633).
- Pasyanos, M. E., et al. (2006), "Reconciling a geophysical model to data using a Markov chain Monte Carlo algorithm: An application to the Yellow Sea-Korean Peninsula region," *J. Geophys. Res.—Solid* **111**, B03313.
- Pesavento, J. B., et al. (2006), "Identification of a thermo-regulated glutamine-binding protein from *Yersinia pestis*," accepted by *Protein Sci.*
- Petzold, L., et al. (2006), "Sensitivity analysis of differential-algebraic equations and partial differential equations," *Computers and Chemical Engineering* **30** (10–12), 1553–1559.
- Reed, E. J., et al. (2006), "Analysis of simulation technique for steady shock waves in materials with analytical equations of state," *Phys. Rev. E* **74** (5).
- Reynolds, D. R., et al. (2006), "A fully implicit numerical method for single-fluid resistive magnetohydrodynamics," *J. Comput. Phys.* **219** (1), 144–162.
- Rodgers, A., et al. (2006), "Site response in Las Vegas Valley, Nevada, from NTS explosions and earthquake data," *Pure and Applied Geophysics* **163**, 55–80.
- Scherlis, D. A., et al. (2006), "A unified electrostatic and cavitation model for first-principles molecular dynamics in solution," *J. Chem. Phys.* **124**, 74103, (UCRL-JRNL-217141).
- Scherlis, D. A., J.-L. Fattebert, and Nicola Marzari (2006), "Stacking of oligo and polythiophene cations in solution: Surface tension and dielectric saturation," *J. Chem. Phys.* **124**, 194902.
- Shen, Z. Q., K.-L. Ma, and T. Eliassi-Rad (2006), "Visual analysis of large heterogeneous social networks by semantic and structural abstraction," *IEEE T. Vis. Comput. Gr.* **12**, 1427–1439 (UCRL-JRNL-221358).
- Shindell, D. T., et al. (2006), "Multimodel simulations of carbon monoxide: Comparison with observations and projected near-future changes," *J. Geophys. Res.—Atmos.* **111**, D19306.
- Stanfield, R., et al. (2006), "Antibody elbow angles are influenced by their light chain class," *J. Mol. Biol.* **357** (5), 1566–1574 (UCRL-JRNL-218128).
- Stevenson, D. S., et al. (2006), "Multimodel ensemble simulations of present-day and near-future tropospheric ozone," *J. Geophys. Res.—Atmos.* **111**, D08301.
- Trebotich, D., et al. (2006), "A hybrid fluid-particle method to simulate flow and transport of DNA in microdevices," *J. Nanosci. Nanotechnol.* (UCRL-JRNL-225206).
- van Noije, T. P. C., et al. (2006), "Multimodel ensemble simulations of tropospheric NO₂ compared with GOME retrievals for the year 2000," *Atmos. Chem. Phys.* **6**, 2943–2979 (UCRL-JRNL-218209).

6.03 — Publications (cont.)

- Vassilevski, P. S. and L. T. Zikatanov (2006), "Multiple vector preserving interpolation mappings in algebraic multigrid," *SIAM J. Matrix Anal. A.* **27** (4), 1040–1055 (UCRL-JRNL-208036).
- Wilson, B., et al. (2006), "PURGATORIO—A new implementation of the INFERNO algorithm," *J. Quant. Spectrosc. Ra.* **99**, 658–679.
- Xie, S. C., et al. (2006), "Developing large-scale forcing data for single-column and cloud-resolving models from the Mixed-Phase Arctic Cloud Experiment," *J. Geophys. Res.—Atmos.* **111** (D19).
- Yoon, S.-E. and D. Manocha (2006), "Cache-efficient layouts of bounding volume hierarchies," *Comput. Graph. Forum* **25** (3) (UCRL-JRNL-219070).
- Yoon, S.-E. and P. Lindstrom (2006), "Mesh layouts for block-based caches," *IEEE T. Vis. Comput. Gr.* **12** (5), 1213–1220 (UCRL-JRNL-220368).
- Zhou, C. E., et al. (2006), "MvirDBA microbial database of protein toxins, virulence factors and antibiotic resistance genes for biodefence applications," *Nucleic Acids Res.* **35**, 391–394 (UCRL-JRNL-218095).
- Zhou, C. E., et al. (2006), "MannDB—A microbial database of automated protein sequence analyses and evidence integration for protein characterization," *BMC Bioinformatics* **7**, 459 (UCRL-JRNL-221584).
- Ziock, K. P., et al. (2006), "Source-search sensitivity of a large-area, coded-aperture, gamma-ray imager," *IEEE Trans. Nucl. Sci.* **53**, 1614–1621.
- ## CONFERENCE PAPERS
- Bagherjeiran, A. and C. Kamath (2006), "Graph-based methods for orbit classification," in *Proc. Sixth SIAM Int. Conf. Data Mining* (SIAM Publications) (UCRL-CONF-215802).
- Bihari, Barna L. and P. N. Brown (2006), "High order finite volume nonlinear schemes for the Boltzmann transport equation," in *Lecture Notes in Computational Science and Engineering* (Springer-Verlag), **48**, 401–421 (UCRL-PROC-210944).
- Bollen, G., et al. (2006), "Rare isotope accelerator—Conceptual design of target areas," in *Nuclear Instruments and Methods in Physics Research: Section A* (Elsevier Science BV), **562**, 915–920.
- Brannick, J., et al. (2006), "Extending the applicability of multigrid methods," in *J. Phys.: Conference Series* (IOP Publishing, LTD), **46**, 443–452 (UCRL-PROC-224817).
- Bremer, P.-T., et al. (2006), "Understanding the structure of the turbulent mixing layer in hydrodynamic instabilities," in *J. Phys.: Conference Series* (IOP Publishing, LTD), **46**, 556–560.
- Bronevetsky, G. and B. R. de Supinski (2006), "Formal specification of the OpenMP memory model," to appear in *Second Int. Workshop on OpenMP* (Springer-Verlag) (UCRL-CONF-221452).
- Chapman, B. M., et al. (2006), "Toward enhancing OpenMP's work-sharing directives," in *Lecture Notes in Computer Science* (Springer-Verlag), **4128**, 645–654 (UCRL-CONF-221478).
- Chartier, T., et al. (2006), "Spectral element agglomerate AMGe," to appear in *Lecture Notes in Computational Science and Engineering* (Springer-Verlag) (UCRL-PROC-212439).
- Eder, D. C., et al. (2006), "Optimization of experimental designs by incorporating NIF facility impacts," in *J. Phys. IV* (EDP Sciences), **133**, 721–725.
- Elliott, A., et al. (2006), "Enhancing seismic calibration research through software automation and scientific information management," in *28th Seismic Research Review: Ground-Based Nuclear Explosion Monitoring Technologies* (DOE/DOD), **1**, 967–975 (UCRL-CONF-222610).
- Friedrich, S., et al. (2006), "Design of a bolometer for total-energy measurement of the linear coherent light source pulsed x-ray laser," in *Nuclear Instruments and Methods in Physics Research: Section A* (Elsevier Science BV), **559**, 772–774.
- Gallagher, B. (2006), "Matching structure and semantics: A survey on graph-based pattern matching," 2006 AAAI Fall Symposium, in *Capturing and Using Patterns for Evidence Detection: Papers* (AAAI) (UCRL-CONF-220852).
- Gygi, F., et al. (2006), "Large-scale electronic structure calculations of high-Z metals on the BlueGene/L platform," in *Proc. Supercomputing 2006* (ACM/IEEE) (UCRL-PROC-220592).
- Hinkel, D. E., et al. (2006), "Laser coupling to reduced-scale targets at NIF early light," in *J. Phys. IV* (EDP Sciences), **133**, 237–241.
- Hittinger, J. A. F. and M. R. Dorr (2006), "Improving the capabilities of a continuum laser plasma interaction code," in *J. Phys.: Conference Series* (IOP Publishing, LTD), **46**, 422–432 (UCRL-CONF-222175).
- Ipek, E., et al. (2006), "Dynamic program phase detection in distributed shared-memory multiprocessors," to appear in *Workshop on NSF Next Generation Software Program* (IEEE) (UCRL-CONF-219596).
- Ipek, E., et al. (2006), "Efficiently exploring architectural design spaces via predictive modeling," in *ACM Architectural Support for Programming Languages and Operating Systems* (ACM) (UCRL-CONF-223240).
- Iseburg, M., et al. (2006), "Streaming compression of tetrahedral volume meshes," in *Pacific Graphics 2006* (Canadian Information Processing Society), 115–121 (UCRL-CONF-217274).
- Järvi, J., et al. (2006), "Algorithm specialization in generic programming: Challenges of constrained generics in C++," in *Proc. 2006 ACM SIGPLAN Conf. Programming Language Design and Implementation* (ACM), 272–282.
- Kamath, C. (2006), "Mining science data," in *J. Phys.: Conference Series* (IOP Publishing, LTD), **46**, 500–504.
- Keyes, D. E., et al. (2006), "Implicit solvers for large-scale nonlinear problems," in *J. Phys.: Conference Series* (IOP Publishing, LTD), **46**, 433–442.
- Kline, J. L., et al. (2006), "Measurements of gas-filled halfraum energetics at the National Ignition Facility using a single quad," in *J. Phys. IV* (EDP Sciences), **133**, 919–923.
- Koniges, A. E., et al. (2006), "Modeling NIF experimental designs with adaptive mesh refinement and Lagrangian hydrodynamics," in *J. Phys. IV* (EDP Sciences), **133**, 587–593.
- Kumfert, G., et al. (2006), "How the common component architecture advances computational science," in *J. Phys.: Conference Series* (IOP Publishing, LTD), **46**, 479–493 (UCRL-CONF-222279).
- Landen, O. L., et al. (2006), "The first experiments on the National Ignition Facility," in *J. Phys. IV* (EDP Sciences), **133**, 43–45.
- Liang, Y., et al. (2006), "BlueGene/L failure analysis and prediction models," in *Proc. Int. Conf. Dependable Systems and Networks* (IEEE), 425–434.
- Love, N. S. and C. Kamath (2006), "An experimental comparison of block matching techniques for detection of moving objects," in *Proc. Society of Photo-Optical Instrumentation Engineers* (SPIE—International Society of Optical Engineering), **6312**, C3120–C3120 (UCRL-CONF-221486).
- McCandless K. N., et al. (2006), "Making and propagating elastic waves: Overview of the new wave propagation code WPP," in *Proc. 2006 Grace Hopper Celebration of Women in Computing* (ACM online) (UCRL-CONF-221316).
- McDonald, J. W., et al. (2006), "Hard x-ray and hot electron environment in vacuum hohlraums at NIF," *J. Phys. IV* **133**, 313–315 (UCRL-PROC-215740).
- Miller, W. M., et al. (2006), "Exploiting data parallelism in the Image Content Engine," in *Proc. Society of Photo-Optical Instrumentation Engineers* (SPIE—International Society for Optical Engineering), **6203**, Q2030–Q2030 (UCRL-CONF-219867).
- Nagayama, T., et al. (2006), "Multiobjective method for fitting pinhole image intensity profiles of implosion cores driven by a Pareto genetic algorithm," in *Review of Scientific Instruments* (American Institute of Physics), **77**, F525–F525.
- Ollivier-Gooch, C., et al. (2006), "The TSTT mesh interface," in *44th AIAA Aerospace Sciences Meeting and Exhibit Online Proceedings* (AIAA).

- Pasyanos, M. E., et al. (2006), "The use of geophysical models for nuclear-explosion monitoring," in *Nuclear Explosion Monitoring Research and Engineering* (National Nuclear Security Administration), **1**, 176.
- Quinlan, D., et al. (2006), "An extensible open-source compiler infrastructure for testing," in *Lecture Notes in Computer Science* (Springer-Verlag), **3875**, 116–133.
- Quinlan, D., et al. (2006), "Annotating user-defined abstractions for optimization," in *IEEE IPDPS 2006 Proceedings* (IEEE) (UCRL-PROC-217573).
- Quinlan, D., et al. (2006), "Support for whole-program analysis and verification of the One-Definition Rule in C++," in *Proc. Static Analysis Summit* (National Institutes of Standards and Technology), 2–35 (UCRL-CONF-222122).
- Ramirez, A., et al. (2006), "Joint reconstructions of liquid plumes using a Markov Chain Monte Carlo approach," in *Proc. 8th Int. Conf. Greenhouse Gas Control Technologies* (Elsevier) (UCRL-PROC-220448).
- Roberts, D. L. and T. Eliassi-Rad (2006), "A position paper: Value of information for evidence detection," AAAI Fall Symposium, in *Capturing and Using Patterns for Evidence Detection: Papers* (AAAI), 37–44 (UCRL-CONF-223973).
- Saini, S., et al. (2006), "Performance evaluation of supercomputers using HPCC and IMB benchmarks," in *Parallel and Distributed Processing Symposium* (IEEE).
- Schmid, T. E., et al. (2006), "Increased human DNA strand damage is associated with occupational exposure to benzene," in *Environmental and Molecular Mutagenesis* (Wiley-Liss), **47**, 418–418.
- Schmid, T. E., et al. (2006), "Vitamin C intake modifies the effects of aging on DNA strand damage in sperm," in *Environmental and Molecular Mutagenesis* (Wiley-Liss), **47**, 424–424.
- Schneider, M. B., et al. (2006), "X-ray flux and x-ray burnthrough experiments on reduced-scale targets at the NIF and OMEGA lasers," in *J. Phys. IV* (EDP Sciences), **133**, 1205–1208 (UCRL-CONF-214849).
- Schulz, M. and B. R. de Supinski (2006), "A flexible and dynamic infrastructure for MPI tool interoperability," in *Proc. Int. Conf. Parallel Processing* (IEEE), 193–200 (UCRL-CONF-221608).
- Schulz, M., D. Kranzmueller, and B. R. de Supinski (2006), "Exploring unexpected behavior in MPI," in *Lecture Notes in Computer Science* (Springer-Verlag), **4208**, 843–852.
- Shumaker, D. E. and C. S. Woodward (2006), "Implicit solution of nonequilibrium radiation diffusion including reactive heating source in material energy equation," in *Computational Methods in Transport* (Springer-Verlag), 353–370 (UCRL-CONF-212026).
- Sjogreen, B., et al. (2006), "A model equation for peak formation in two-phase flow," in *Hyperbolic Problems: Theory, Numerics, and Applications* (Yokohama Publishers).
- Streitz, F. H., J. Glosli, and M. Patel (2006), "Simulating solidification in metals at high pressure: The drive to petascale computing," in *J. Phys.: Conference Series* (IOP Publishing, LTD), **46**, 254–267 (UCRL-CONF-223218).
- Trebotich, D. (2006), "Modeling complex biological flows in multiscale systems using the APDEC framework," in *J. Phys.: Conference Series* (IOP Publishing, LTD), **46**, 316–321 (UCRL-CONF-222419).
- Trebotich, D., G. H. Miller, and M. D. Bybee (2006), "A hard constraint algorithm to model particle interactions in DNA-laden micro flows," in *ECCOMAS CFD 2006: European Conf. Computational Fluid Dynamics* (TU Delft) (UCRL-PROC-225265).
- Trebotich, D., G. H. Miller, and M. D. Bybee (2006), "Modeling of particle interactions in DNA-laden flows at the microscale," in *Proc. 26th Annual Int. Conf. IEEE Engineering in Medicine and Biology Society* (IEEE), 6670–6673 (UCRL-CONF-224475).
- Trebotich, D., G. H. Miller, and M. D. Bybee (2006), "Particle interactions in DNA-laden flows at the microscale," in *Proc. Second Int. Conf. Transport Phenomena in Micro and Nanodevices* (IEEE) (UCRL-PROC-225266).
- Trinitis, C. and M. Schulz (2006), "Fifth international special session on current trends in numerical simulation for parallel engineering environments: New directions and work-in-progress," ParSim 2006, in *Lecture Notes in Computer Science* (Springer-Verlag), **4192**, 356–357 (UCRL-PROC-222517).
- Vuduc, R., et al. (2006), "Improving distributed memory applications testing by message perturbation," in *Workshop on Parallel and Distributed Systems: Testing and Debugging* (UCRL-PROC-221395).
- White, D. A., B. J. Fasenfest, and M. L. Stowell (2006), "A parallel computer implementation of fast low-rank QR approximation of the Biot-Savart law," in *Progress in Electromagnetics Research Symposium* (Electromagnetics Academy), 420–424 (UCRL-CONF-216989).
- Willcock, J. et al. (2006), "Accelerating sparse matrix computations via data compression," in *ACM Int. Conf. Supercomputing* (ACM), 307–316.
- Wissink, A., et al. (2006), "Adaptive urban dispersion integrated model," in *Sixth Symposium on the Urban Environment* (UCRL-PROC-216813).
- Yee, H. C. and B. Sjogreen (2006), "Adaptive numerical dissipation control in high order schemes for multi-D nonideal MHD," in *Computational Fluid Dynamics* (Springer-Verlag).
- Yee, H. C. and B. Sjogreen (2006), "Performance of adaptive dissipation control in high order methods for the resistive MHD equations," in *Hyperbolic Problems: Theory, Numerics, and Applications* (Yokohama Publishers).

6.04 — Acronyms and Abbreviations

3D — three-dimensional

ASC — Advanced Simulation and Computing

ALE — Arbitrary Lagrangian Eulerian

AMG — algebraic multigrid method

AMR — adaptive mesh refinement

BG/L — BlueGene/L

C&A — certification and accreditation

CCC — Capability Computing Campaigns

CIO — Chief Information Officer

CMDB — Configuration Management Database

CSP — Computer Security Program

DOE — Department of Energy

DNT — Defense and Nuclear Technologies

FN — foreign national

FPMD — first-principles molecular dynamics

GPU — graphics processing unit

HPC — high-performance computing

ISCR — Institute for Scientific Computing Research

ISQA — institutional software quality assurance

ISSBN — Institutional Support Services Backbone Network

IT — information technology

ITPD — Information Technology Protection Division

LC — Livermore Computing

LCLS — Linac Coherent Light Source

LDRD — Laboratory Directed Research and Development

LLNL — Lawrence Livermore National Laboratory

LPOM — Laser Performance Operations Model

LUSciD — Livermore–UC San Diego Scientific Data Study

M&IC — Multiprogrammatic and Institutional Computing

MCR — Multiprogrammatic Capability Resource

NAA — network access architecture

NAPA — Novel Architecture and Progressive Algorithms

NIF — National Ignition Facility

NNSA — National Nuclear Security Administration

SCFN — sensitive country foreign national

SciDAC — Scientific Discovery through Advanced Computing

SPMD — single-program, multiple-data

SQA — software quality assurance

TCO — total cost of ownership

TF — trillion floating-point operations per second

TSF — Terascale Simulation Facility

UQ — uncertainty quantification

WPP — Wave Propagation Project