NA-ASC-500-07 Issue 5

December 2007

The View from HQ

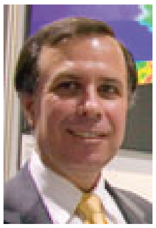

By Dimitri Kusnezov

The View from HQ

It is time for New Year's Resolutions or perhaps New Year's resolve. I know that, in general resolutions are easy to make and easy to let slide. My intent is to resolutely move forward on an initiative in which I believe very strongly.

For some time I have been concerned that if weapons funding continues to decrease, the nation risks losing a critical scientific capability that has been built up through decades of excellence in application to the security and reliability of the nuclear weapons stockpile. Although certain agencies have taken advantage of this strength by leveraging the expertise that exists at the Defense Program laboratories to address technical challenges for which they desperately need help, they have had no incentive to pay for more than just the short-term help they buy. There are some exceptions, but they are few and far between.

NNSA, together with the defense laboratories' directors, has crafted a vision that expands its charge to include threats to national security and response to emergencies beyond the nuclear weapons arena. This vision does not lessen our commitment to fulfill the statutory responsibility for maintaining the nuclear stockpile; it does imply, however, that we recognize the importance of bringing the talents and capabilities of these national laboratories to bear on a broadened spectrum of national security issues, ranging from challenges by proliferant states to achieving energy independence.

We must apply both the expertise and the facilities we have to help defend the nation against a wide range of threats. There are two sides to this coin; one side is our responsibility to use our experience in technical problem solving to serve the nation; the other side is to partner with other federal agencies to be a resource to them and to encourage them to share in the costs of preserving these capabilities—intellectual capital as well as physical infrastructure.

I realize that the defense laboratories are going through difficult times, having to make hard decisions with respect to staffing and supporting the physical infrastructure. As the New Year unfolds, I resolve to work with you to address the damage that resource shortfalls may cause and to expand our future possibilities from within the NNSA framework. I believe it is crucial to ensure the continuing relevance of the work and the vitality of the laboratory workforce by enlarging the national defense mission, and by engaging the next generation of scientists. We are developing a mechanism to do both.

We have an opportunity to be the agency of choice for providing technical responses to the full span of national security challenges. Not everyone supports nuclear weapons, but everyone recognizes the need to apply technology toward threats to the safety of our citizenry and to improve the quality of our lives. Please work with me to enable NNSA to rise to the challenge as I resolve to work with and support you in this endeavor.

Nuclear Forensics Traces Nuclear Explosions

If someone has detonated a nuclear device, would we be able to identify where the material came from? Nuclear Forensics is "post-event" analysis of a nuclear event, and it would help to discover the source of a nuclear device. The ASC integrated codes already have significant capability to simulate nuclear events with a certain degree of confidence, but those are mostly our devices and our underground nuclear tests, under controlled conditions. What happens if, say, a terrorist group steals a device from a foreign country and sets it off in an underground parking garage? What can the ASC codes do to help with understanding the event?

The National Technical Nuclear Forensics (NTNF) effort is a multi-agency, national effort involving personnel who will perform emergency response, consequence management, and modeling with rapid turnaround to aid field work, as well as longer term work for device characterization. ASC has supported a fully integrated, bi-lab (Los Alamos National Laboratory [LANL] and Lawrence Livermore National Laboratory [LLNL]) effort to improve the fundamental nuclear data and code capabilities in the modern ASC production codes in support of NTNF. On Wednesday, November 12, a group of LANL and LLNL personnel, along with representatives from HQ, DTRA, DHS, and other organizations in NNSA, heard a briefing by LANL and LLNL on efforts so far and a path forward. Mark Chadwick (LANL) gave some introductory remarks. Chris Clouse (LLNL) described the code efforts at the two laboratories, and Bob Little (LANL) summarized their nuclear data efforts. This work is leading to a joint LLNL/LANL FY09 Level 1 Milestone for validated, fully operational capability for NTNF.

At the end of the presentation, there were discussions on opportunities and other work. For example, a question was posed whether there is a Nuclear Forensics problem that can utilize the resources of BlueGene/L. What capabilities may need to be developed for understanding device output signatures? And how can ASC codes help the other agencies to focus resources in their work?

Innovative Computing Technique Earns Livermore Team the Gordon Bell Prize

From left to right are the 2007 Gordon Bell winners from Lawrence Livermore National Laboratory: Jim Glosli, Fred Streitz, Robert Rudd, David Richards, and Kyle Caspersen. IBM's John Gunnels is not pictured.

A team of scientists from the Lawrence Livermore National Laboratory Physical Sciences Directorate and IBM won the prestigious Gordon Bell Prize in Peak Performance with a breakthrough physics calculation run on the recently expanded BlueGene/L system. The award was announced at the Supercomputing 2007 conference (SC07) in Reno, Nevada. Computation Associate Director Dona Crawford said, "These scientific computing awards underscore the vital role the ASC Program plays in NNSA's Stockpile Stewardship Program as well as in national security in a global context."

By performing extremely large-scale molecular dynamics simulations using an innovative computational technique, the Livermore Gordon Bell team was able to study, for the first time, how a Kelvin-Helmholtz instability develops from atomic-scale fluctuations into micron-scale vortices. This simulation of unprecedented resolution was made possible by the innovative computational technique used—a technique that could change the way high-performance scientific computing is conducted.

The Kelvin-Helmholtz instability arises at the interface of fluids in shear flow and results in the formation of waves and vortices. Waves formed by Kelvin-Helmholtz instability are found diverse natural phenomena, such as waves on a windblown ocean, sand dunes, and swirling cloud billows. While Kelvin-Helmholtz instability has been thoroughly studied for years and its behavior is well understood at the macro-scale, scientists did not clearly understand how it evolved at the atomic scale until now. Understanding how matter transitions from a continuous medium at macroscopic length scales to a discrete atomistic medium at the nanoscale has important implications for such laboratory research efforts as National Ignition Facility laser fusion experiments and developing applications for nanotube technology.

Sandia and Livermore Collaborate on Radiation-MHD Code Development

Dimitri Kusnezov (NNSA), Joan Woodard (Sandia), and Bruce Goodwin (Lawrence Livermore), recently signed a Memorandum of Agreement "to establish an explicit collaboration between Sandia National

Laboratories (SNL) and Lawrence Livermore National Laboratory (LLNL) for the purposes of developing, verifying, and validating a radiation-MHD simulation capability for high density physics (HEDP) and aboveground experiments (AGEX) within LLNL's Kull framework. A goal for this collaboration is to leverage the complementary strengths of the LLNL and SNL projects (radiation-hydrodynamics and magnetohydrodynamics [MHD], respectively) to create 'best in class' capabilities in a cost-effective fashion."

This development will build upon the successful collaboration in Implicit Monte Carlo radiation transport that has existed for several years between these two laboratories. The memorandum states: "A joint repository for the KULL code and associated third party libraries, material data, and other support software will be established for the combined code team utilizing appropriate network connections and software. The focus of the initial collaboration will be the development of MHD modules, expanding to other areas." The project will draw upon LLNL expertise in multi-physics radiation-hydrodynamics algorithms and the three dimensional (3-D) MHD algorithms developed by Sandia. The 3-D MHD algorithm in the ALEGRA code recently demonstrated quantitative agreement with x-ray powers and energy measurements from wire array z-pinch implosions on the Z accelerator (Sept. 2007 ASCeNews). Inclusion of these MHD algorithms into the Kull and continued development of the radiation-hydrodynamics capabilities by a distributed codevelopment team is expected to deliver an important HEDP simulation capability in support of both Sandia and LLNL HEDP missions.

Micron-Scale Atomistic Studies of Shock Ejecta Production Using BlueGene/L

Los Alamos researchers recently completed a series of simulations of shock ejecta production in copper. This fragmentation and atomization process is difficult to study experimentally, and various theories have been proposed. Atomistic-level simulations will contribute to the development of physics-based models as part of the LANL "Science@Scale" effort (see Sept. 2007 ASCeNews article). In particular, the dependence of ejecta production and transport on shock pressure (e.g., below and above the shock melting transition, to study material strength effects) and background gas (with either a vacuum as in previous simulations or an inert gas atmosphere) have been studied.

The simulation used the entire BlueGene/L machine at Livermore, which was expanded in the fall of 2007 to 212,992 CPUs. The longest simulation ran continuously for 88 hours (400,000 timesteps, or one nanosecond of simulated time) with nearly 800 million atoms generating 101 checkpoint dumps (4.18 TB of data) and 3024 images (14.5 GB). The nearly 20 million CPU-hours of this simulation alone are equivalent to more than 2 CPU-millennia, setting a new standard for HPC stability.

These simulations utilized a quasi-2D geometry, 5.7 microns in length and a 2.23-micron periodic cross-section (but only 1.5 nm in the third direction, also periodic, to preserve 3D equation-of-state and transport properties). Copper atoms were described by an embedded atom method (EAM) potential. The free surface opposite the impact plane is initially machined with a profile matching that measured in recent tin ejecta experiments at LANL,* with approximately 1:40 length scaling so that the ~1 micron experimental amplitude is ~25 nm. Earlier molecular dynamics simulations with a single machining groove demonstrated jet formation, with subsequent necking instabilities leading to jet breakup and droplet formation at later times.

Time sequences from two such simulations are shown in the figures, with the initial surface roughness and shock pressure equal in both cases. Color represents the local density in each case.

*M. B. Zellner et al., "Effects of shock-breakout pressure on ejection of micron-scale material from shocked tin surfaces," J. Appl. Phys. 102, 013522 (2007).

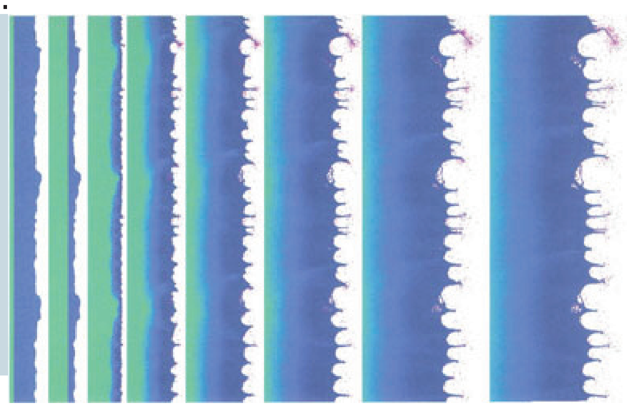

For shock ejection into a vacuum, we observe jet formation from small-scale features at early times: Series of 8 snapshots, one every 10 ps, for shock ejection from a roughened copper surface). These jets, surrounded by a cloud of atomic ejecta, expand and subsequently merge or break up over time.

At later times: snapshot 300 ps after shock encounters the free surface), the longer-wavelength machining defects produce a Richtmyer-Meshkov instability, with three "bubbles" of vacuum pushing into the copper surface, and three large "spikes" of copper protruding out, with smaller-scale jets superimposed throughout. The bubbles and spikes are very asymmetric, as one would expect for Atwood number A=1.

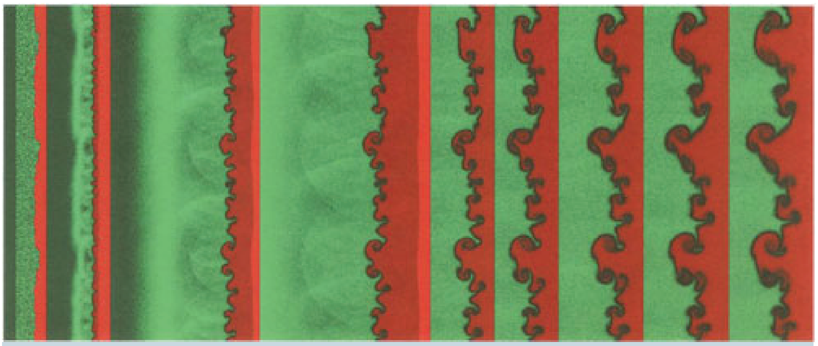

A background gas (A < 1) inhibits jetting and leads to a more symmetric Richtmyer-Meshkov instability, as seen in the second simulation: Series of 9 snapshots, one every 50 ps, for Richtmyer-Meshkov instability development as a shock wave is transmitted from copper to a dense (0.5 g/cc) gas.) In the third and fourth frames, one can clearly see the transmitted gas shock (still nonplanar), as well as a complex interaction of rarefaction fans from the roughened Cu/gas interface.

First Demonstration of Predictive Capability for Weapon Systems in Hostile Nuclear Environments

A Sandia team has completed the first demonstration of a high-fidelity predictive capability for combined system-generated electromagnetic pulse (SGEMP) effects. Without underground testing, a robust SGEMP predictive capability such as that demonstrated by Sandia is needed to design and qualify new systems (e.g., RRW1).

SGEMP effects occur when intense x-rays generate a pulse of electrons in the re-entry vehicle/body (RV/RB) interior that generates strong electromagnetic (EM) fields inside the system. These fields can potentially disrupt electronics. For this milestone, the Sandia team modeled the SGEMP response of a cable, a small cavity, and an enclosed electronic component. For the first time, a prediction of the electrical insult delivered to a subsystem from the combined SGEMP effects was possible. The effects were self-consistently calculated using a lumped-parameter circuit model that integrated all the effects.

The cavity and component SGEMP effects are especially difficult to simulate. For cavity SGEMP, complex physical interactions occur between radiation, EM, and plasma, which are only partly understood. Further development is continuing of models for the physical response of cavity gases as well as component dielectrics and foams in intense radiation environments.

The computational demands for the radiation transport part of the cavity SGEMP analysis were very demanding, requiring up to 700,000 CPU-hours on Purple to resolve the distribution of photo-emitted electrons to the necessary precision. Predicting component SGEMP was also very challenging. For the radiation transport analysis, the electron current distribution across more than 30 million finite elements was resolved to micron resolution at material boundaries. Red Storm was extensively utilized, and each radiation transport calculation required 30,000 CPU-hours. This same mesh was the largest ever used for time-domain EM analysis.

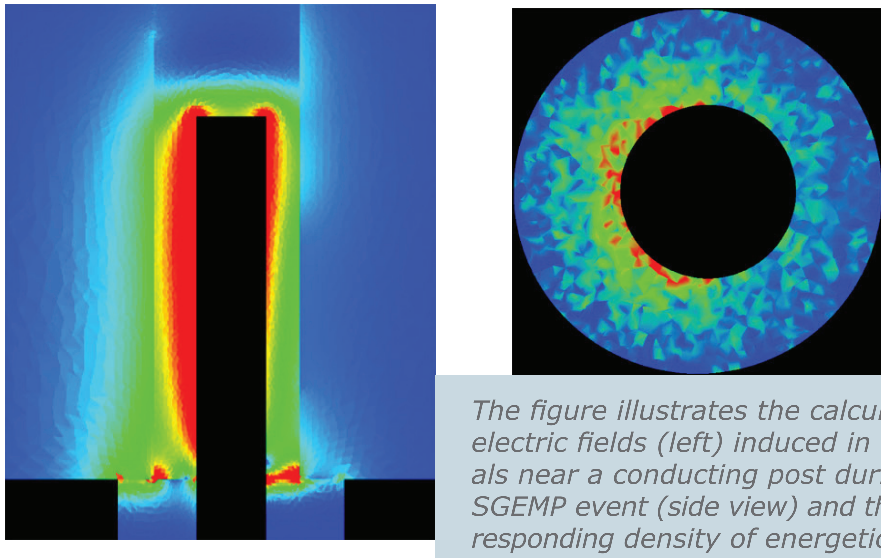

The figure illustrates the calculated electric fields (left) induced in materials near a conducting post during an SGEMP event (side view) and the corresponding density of energetic electrons (right) near the post (top view).

ASC Software Quality Engineering Group Gives Defects Some Static

The ASC Verification and Validation (V&V) software quality engineering (SQE) group at Lawrence Livermore National Laboratory used a new static analysis tool, the automated static analyzer (ASA), to analyze over 2 million lines of ASC source codes and to verify that more than 3,000 defects were repaired in much less time than would be possible using standard debugging techniques. This work was performed in support of an FY07 level 2 ASC milestone.

The ASA is an automated tool that can analyze C, C++, and Java source code for software defects, including security vulnerabilities. Livermore has two different ASA tools with similar capability. One is used by the SQE group, and one is used by other programs at Livermore. Both ASA tools return low false positive rates (20 percent as compared to 50 percent for other static analyzers); however, the ASC tool used by the SQE group provides additional features such as metrics and trending and architectural analysis that was important to the ASC Program.

Livermore software quality engineers are using ASA to help code teams improve code reliability. In the past, finding software defects involved using a run-time debugging tool that would stop the code when a defect was encountered. The developer would trace through the code to identify the cause of the defect and then develop a fix. This could take anywhere from a minute to several days per defect, and some ASC codes had several thousand defects. The ASC ASA tool can analyze as many as 1,600 defect types within an hour or two, thus eliminating months of code developer time. ASA tools will not replace the benefits of peer inspections or functional testing, since the tool is unaware of user requirements.

Plans for FY08 include analyzing the entire suite of ASC codes through the ASA tool and then analyzing the initial reports to pre-filter out false positive defects for developers. This will save considerable time because pre-filtering is the most time-consuming part of the task. Research is also under way to evaluate using a Fortran-to-C converter on legacy Fortran codes and then analyzing the results with the current ASC ASA tool.

Accident Response of SafeGuards Transporters Simulated

For the first time, Sandia analysts simulated the accident response of a full-scale SafeGuards Transporter (SGT), including the weapons cargo. Testing of full-scale SGT accident scenarios is very limited and, as a result, provides little insight into the effects of different impact conditions and does not have the ability to assess the efficacy of design modifications of the transporter. As a result, computational simulation is the only way to address the diverse scenarios that must be explored.

The high-fidelity SGT finite element model, containing 20 million degrees-of-freedom was run with the ASC transient dynamics capability in Sandia's SIERRA Mechanics code suite. When the simulation results

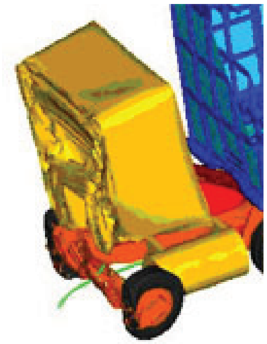

were compared to the existing full-scale test results, it was found that the model faithfully replicated the experimental results. The simulation of the 300-millisecond crash event required 500,000 CPU-hours on Red Storm.

Computer simulations using Sandia's finite element model will be used to identify worst-case scenarios and assess design modifications, as well as quantifying margins and uncertainties associated with accident scenarios. These simulations will exercise the unique capabilities in Sandia's engineering codes and will utilize Red Storm and the next generation of capability machines. Since weapon safety requires the ability to assess weapon component response from insults external to the SGT, there is a need for very high-fidelity models, and some scenarios will require even higher fidelity models than we can currently run.

Pit Lifetime Assessment Team Awarded Defense Programs Award of Excellence

On Wednesday, December 5, 2007, Mr. William Ostendorff, NNSA's Principal Deputy Administrator, presented an Award of Excellence to the Los Alamos National Laboratory (LANL) Pit Lifetime Assessment Team. The 120-member team was responsible for meeting the Laboratory's 2006 Level 1 Milestone for lifetime assessments, which involved substantial ASC resources, including BlueGene/L. The team delivered a science-based assessment of plutonium-aging effects on pit lifetimes and primary performance for the deployed LANL stockpile systems. Their studies included a reanalysis of the underground nuclear test records and an extensive experimental and computational investigation of the mechanical, physical, and chemical property changes caused by plutonium aging.

High-Fidelity B61 Weapon-in-a-Fire Sensitivity Analysis Study Completed

A sensitivity analysis study using massively parallel calculations of a B61 subject to various weapon-in-a-fire scenarios was recently completed on Sandia's Thunderbird computer. These studies employed Sandia's Design Through Analysis Realization Team (DART) tools to construct the computational model of the B61, along with Sandia's Sierra thermal/mechanical simulation software to model the fire environment. A typical run time for one of these Sierra calculations is on the order of 50 cpu-years. In addition, this study included the first-ever coupling of a Sierra fire simulation with Sandia's DAKOTA uncertainty quantification and optimization software toolkit. This was done in an effort to identify the maximum heating conditions for a weapon in a specific fire accident scenario. The maximum heat fluxes predicted by Sierra provide high-fidelity time and space varying boundary conditions that will be used in follow-on Sierra thermal simulations of the B61's internal components. Overall, these combined high-fidelity fire/thermal Sierra calculations will improve the technical basis understanding of B61 system response to abnormal thermal environments.

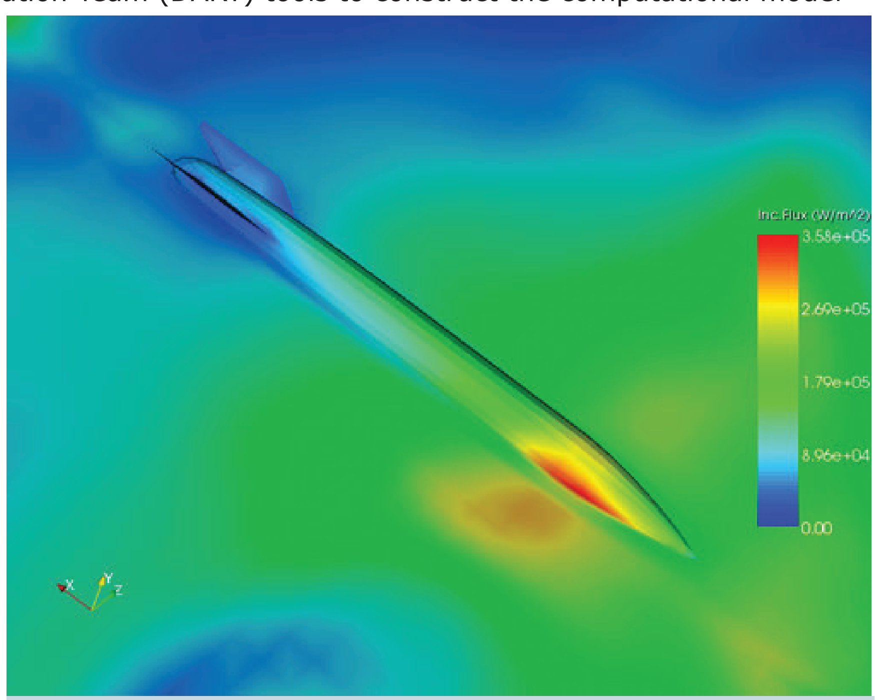

A bomb partially buried in snow colored by the Sandia code Fuego predicted contours of instantaneous incident heat flux from a fire.

Common Tri-Lab Capacity Computing Moves Closer to Goal

The NNSA initiative to establish common computational environments on tri-lab computing systems will take major steps forward in FY08. For several years, a tri-lab effort known as Tripod has been under way to promote cost-effective, portable environments on the Linux-based cluster systems at Livermore, Los Alamos, and Sandia National Laboratories.

In FY07, a tri-lab procurement was led by Livermore for a common Linux hardware, called the Tri-Laboratory Linux Capacity Cluster (TLCC07) systems. NNSA headquarters selected five proposals for FY08 support to establish a common production software stack for the TLCC07 hardware. The largest was submitted by Livermore and Sandia to deploy a common cluster management software stack based on Red Hat Enterprise Linux (RHEL) on all the TLCC07 clusters. This system software, called the Tripod Operating System Software (TOSS), is a tri-lab packaging of the CHAOS/SLURM environment that has been in production on Livermore systems for several years. In addition to the system software stack, four smaller proposals were selected for support, pending funding.

With the introduction of TOSS, many of the software management processes will be broadened and formalized to embrace and support tri-lab computing. Accomplishments thus far include, deploying the TOSS support infrastructure, including a server for the TOSS software repository and for bug reporting by the three laboratories; releasing the Alpha version of TOSS in mid October; each laboratory successfully booting TOSS on at least one Linux node; and successfully running all the synthetic workload applications using LANL's Gazebo test harness on a test system. The initial generally available release of software will be deployed with the first TLCC07 cluster.

Prototype of a new tri-lab Linux capacity cluster. Scalable units (SU) will be aggregated into clusters of two, four, six, or eight SU, with each cluster available for computing across the three defense laboratories.

Expanded BlueGene/L Retains Rank as World's Fastest Supercomputer

BlueGene/L retained its number one ranking on the new Top500 list of the world's fastest supercomputers at the international Supercomputing 2007 conference in Reno, Nevada. BlueGene/L clocked 478.2 teraFLOPS on LINPACK, the industry standard measure of high-performance computing. Recently expanded to accommodate the growing demand for high-performance systems able to run the most complex nuclear weapons science calculations, BlueGene/L now has a peak speed of 596 teraFLOPS.

The BlueGene/L system was formerly housed in 64 cabinets arranged in eight rows. The new system is housed in 104 cabinets in 13 rows with an expanded node count of 106,496 (up from 65,536) and an expanded processor count of 212,992 (up from 131,072). The total memory of the system increased from 32 tebibytes to 68 tebibytes with the addition of nodes with twice the memory.

"Expanding the BlueGene/L system allows us to explore a new class of applications important to our mission and is an important step toward the predictive, fully integrated 3D weapons calculations vital to NNSA's stockpile stewardship mission," said Michel McCoy, head of the ASC Program at Lawrence Livermore.

The NNSA BlueGene/L supercomputer housed at Lawrence Livermore National Laboratory was named the world's fastest computer for a seventh straight time, according to the Top500 list released November 2007.

The BlueGene/L upgrade, most notably the additional memory, allows scientists from the three nuclear weapons laboratories to develop and explore a broader set of applications than the single package weapons science oriented work that has been the mainstay of the machine in the past. For example, BlueGene/L had been used widely for materials science calculations such as assessing materials at extreme temperatures and pressures. Now it will be much easier to run more complex applications related to modeling integrated systems as opposed to focused exploration of one area of physics or chemistry.

Roadrunner Videotape Segments on the Web

A Los Alamos National Laboratory communications team created a video about the Roadrunner Project for an exhibit at SC07. Divided into five segments on the Web site, the video features top-level managers, computer scientists, and plasma physicists discussing the challenges and what we have to look forward to in the future of high-performance computing. The segments are 1) Petascale: The Next Generation of Computing; 2) Hybrid Computing Architecture; 3) Programming for the Future; 4) Science@ Scale; and 5) The Science and Technology Challenges. Go to http://www.lanl.gov/orgs/hpc/roadrunner/video.shtml

Speed of Light Is Too Slow

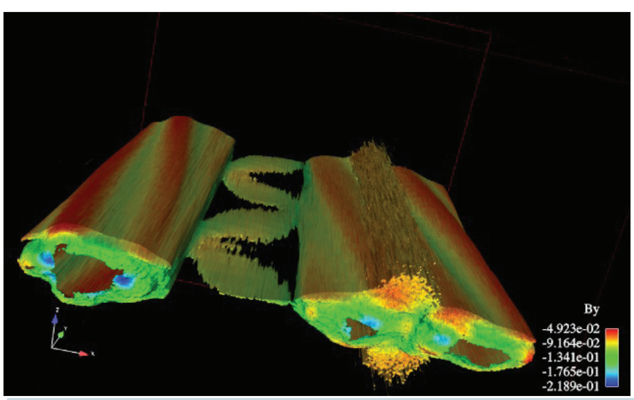

Magnetic islands and drift-kink instability observed in 3D large-scale fully kinetic VPIC simulations run on Roadrunner base system of magnetic reconnection in electron-positron plasma.

The following press release highlights a computer code called VPIC that takes advantage of optimized microprocessors such as the Roadrunner base system. The release was published by the American Physics Society during the annual meeting of the Division of Plasma Physics the week of November 12, 2007.

"Ultrafast Programs Developed for Future Ultrafast Computers (http://www.lanl.gov/news/index.php/fuseaction/nb.story/story_id/11931)," by Kevin N. Roark

Livermore Computer Scientists Receive R&D100 Award for *hypre*

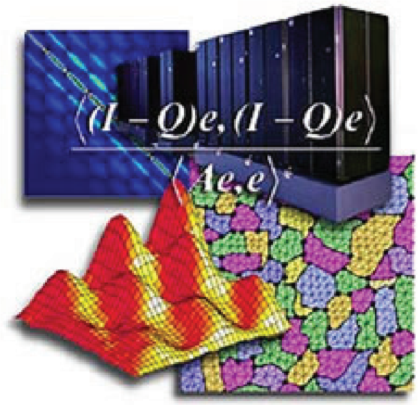

Lawrence Livermore National Laboratory computer scientists have developed a software library called hypre that allows researchers to use supercomputers such as BlueGene/L and ASC Purple more effectively to conduct larger, more detailed simulations faster than before. This cutting-edge technology garnered a 2007 R&D100 award from the trade journal R&D Magazine for developing advances with commercial potential among the top 100 industrial inventions worldwide for 2006.

The ability to run large, detailed simulations in less time is critical to stockpile stewardship and other DOE/NNSA national security programs that use massively parallel machines with as many as tens of thousands of processors. Physical systems are described by complicated sets of mathematical equations that must be solved on computers to simulate reality. Large systems of linear equations are ubiquitous in scientific and engineering simulation codes, and solving these systems of equations is often the most time-consuming function of a code. Consequently, robust and efficient algorithms to solve systems of linear equations are in great demand.

Hypre, which stands for "high performance preconditioners," is a software library unique in its ability to provide solution algorithms that are effective on a wide variety of problems, easily accessible using multiple user interfaces, and effectively exploit the full computational power of today's high performance computers.

The challenge for parallel linear solver algorithms has been scalability. An application code is scalable if it can use additional computational resources effectively. If the size of the problem and the number of processors are increased proportionally, the goal is to keep the computing time approximately the same. In practice, however, as simulations grow to be more detailed and realistic, computing time may increase dramatically even when more processors are added to solve the problem.

Hypre provides linear solver algorithms developed specifically to be scalable on large numbers of processors. As a result, some simulation times may be reduced by orders of magnitude—by as much as a factor of 30. Simulations that previously took days can now be run in hours or less.

Algorithms Decrease Solution Time Ten-Fold

Sandia's ASC Algorithms Milestone team, led by Roscoe Bartlett, developed and demonstrated a full vertical integration of numerical algorithms in the Trilinos Library with the Charon semiconductor device modeling code. Compared to prior approaches, these new algorithms have demonstrated up to a ten-fold reduction in time-to-solution while providing more accurate and robust results.

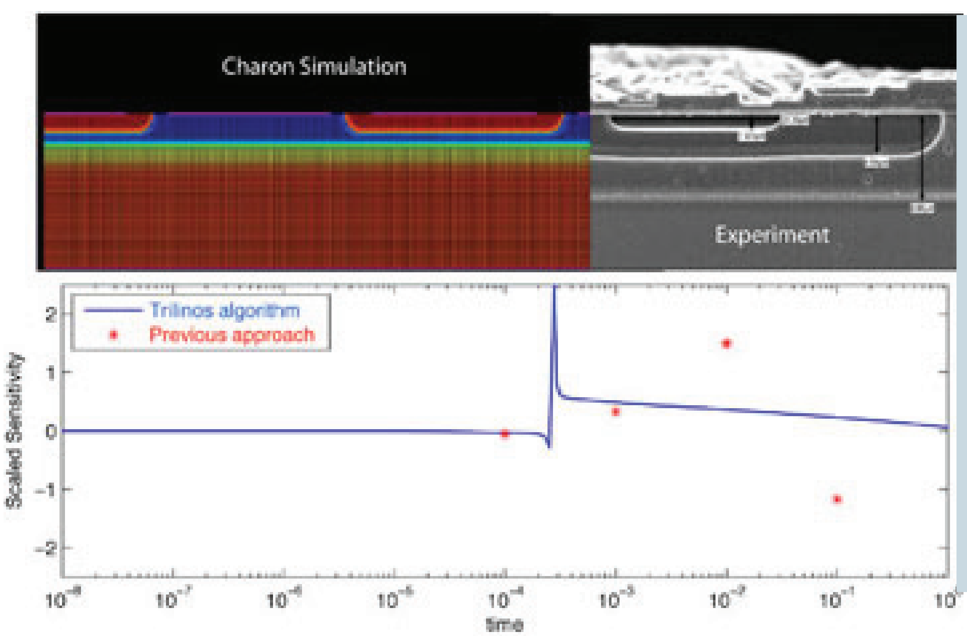

The top images compare typical Charon simulations of a semiconductor device to experimental imagery. The bottom plot shows a time history of the sensitivity of a Charon device model to an input parameter computed using the Trilinos vertically integrated algorithms and the prior point-wise approach. Trilinos algorithms return the entire time history of the sensitivity at 1/10 the computational cost and with greater accuracy.

Algorithms ranging from linear algebra, preconditioners, and iterative linear solvers, all the way up through nonlinear solvers, transient solvers, and optimization were integrated with Charon, in a flexible and modular, yet high-performance and scalable manner. The Qualification Alternatives to SPR (QASPR) program is using these new capabilities in Charon to explore the sensitivity of transistor device models to input parameters and, in the future, to optimize parameters based on experimental data. The versatility of these vertically integrated Trilinos components was also highlighted by their application to the SIERRA engineering simulation suite, with the development of an intrusive optimization algorithm used to solve a prototype MEMS actuator design problem.

These new algorithmic capabilities are available in Trilinos version 8.0, released August 31, 2007, and are being deployed to help bridge the research-development-application spectrum for predictive simulation. You can learn more about Trilinos as well as download the full suite of open-source software at http://trilinos.sandia.gov/.

SC07 Cluster Challenge Committee Proclaims: Perceived Entry Barrier to Supercomputing Has Dropped Significantly

Cluster Challenge being introduced prior to event.

At Supercomputing 2007 (SC07), six teams of undergraduates, supported by their chosen vendor partners, competed in a test of performance on the conference floor. Teams were given a few simple rules (26 amps and no one on the team can have a degree) and the summer to design the system that would perform best on a workload consisting of Linpack benchmarks (the industry standard High Performance Computing [HPC] Challenge) and three open source applications: GAMESS, a molecular chemistry application from AmesLab; POP, an ocean circulation model from Los Alamos; and POVRay, a popular ray tracing application.

Cluster Challenge Committee Chair Brent Gorda, of Lawrence Livermore National Laboratory, observed, "The results showed that if you have a need for simulation computing, it's reasonable to expect that you can use local college or university talent and commonly available software and applications to get started on that work."

Six teams of undergraduates participated from the U.S., Canada, and Taiwan. The highest Linpack score of 420 gigaFLOPS was announced to cheers and applause at the Top500 Birds of a Feather (BOF). That performance would have made the list of the Top500 supercomputers just three years ago. For this to be done by a half-dozen undergraduates shows that there is no significant barrier for entry-level HPC. Conventional wisdom suggested that any system on the Top500 had to be backed by an institution, yet the challenge proved that a team of undergraduates could make the list.

The overall winning team was from the University of Alberta, in Edmonton, Canada, combined with SGI. The other competing teams were Stony Brook University and Dell, National Tsing Hua University (Taiwan) and ASUSTek, University of Colorado and Aspen Systems, Indiana University and Apple, and Purdue University and HP.

Representing ASC at Supercomputing 2007, from the Ground Up

Tri-lab specialists in hardware, networking, electrical, computer support, graphic design, videography, communication, and administration at Livermore, Los Alamos, and Sandia national laboratories combined resources and talents to design, build, and staff the NNSA ASC tri-lab booth at Supercomputing 2007 (SC07). The display by the three laboratories, Y-12, and the five alliances showcased technical breakthroughs achieved on behalf of the nation's Stockpile Stewardship Program (SSP).

The conference was from November 10 to 16, but planning began in April with a kick-off meeting held at Lawrence Livermore National Laboratory. As plans for the booth developed to fit the 40-foot-by-40-foot

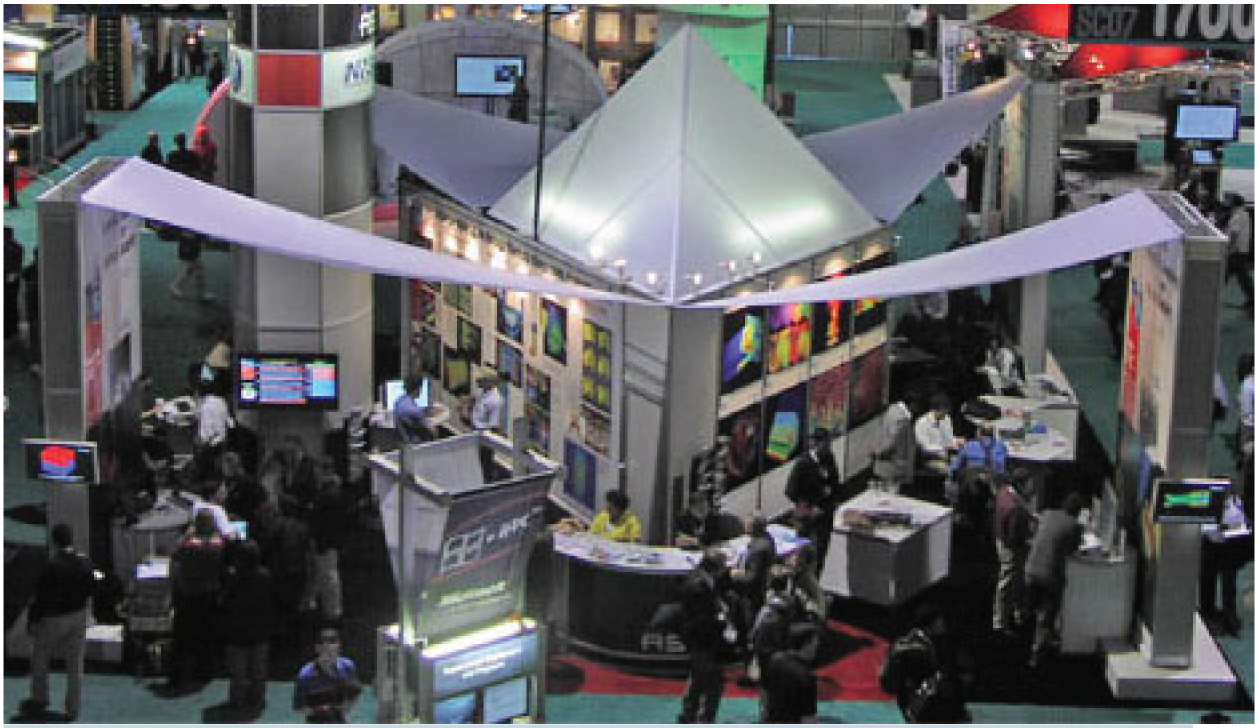

ASC Program booth at Supercomputing 2007, as seen from the balcony above the conference floor. The curved reception desk is in the front center. Note the ASC work breakdown structure wall (left), the ASC simulation gallery (right), and the four corner pillars featuring ASC's largest supercomputers.

space, a theme emerged—Scientific Computing for National Security. This theme was reflected in booth walls, graphics, fliers, posters, and brochures. The booth included many presentations and demonstrations, including video displays and a 3D area with a large powerwall.

A major reason ASC attends Supercomputing is to showcase the science. Top of the list was Gordon Bell prize-winner Fred Streitz (Lawrence Livermore) and team, winning for a first-of-a-kind simulation of Kelvin-Helmholtz instability in molten metals.

DOE-supported researchers across the Complex had a role in:

- 18 of the 54 peer-reviewed technical papers
- 5 of the 7 panel discussions
- 10 of the 25 tutorial sessions
- 6 of the 10 workshops
- 2 of the 15 invited "Masterworks" talks
- 13 of the 45 research posters

For more information on DOE's HPC expertise at SC07, see the following:

- Article in HPC Wire http://www.hpcwire.com/hpc/1880961.html
- Flier for BlueGene/L upgrade https://asc.llnl.gov/publications/sc2007-bgl.pdf
- Flier for new tri-lab Linux capacity clusters https://asc.llnl.gov/publications/sc2007-tlcc.pdf
- Brochure describing tri-lab ASC Program https://asc.llnl.gov/publications/asc07.pdf
- Press release for Gorden Bell winner: Kelvin-Helmholtz instability in molten metals

http://www.llnl.gov/pao/news/news_releases/2007/NR-07-11-04.html

High-Performance Computer Support Team Reinvents Itself

"Better, Faster, Cheaper" is the three-pronged goal of Los Alamos National Laboratory's High-Performance Computer Support Team (HPCST) for in-house computer hardware support. The HPCST has reinvented itself by redefining its long-established mission of monitoring all LANL's supercomputers, clusters, and facilities. In addition to monitoring these systems, the team has taken on the responsibility of providing maintenance and support for all the supercomputer systems currently installed at the laboratory. The HPCST staffs the operations center in the Metropolis Center for Modeling & Simulation 24 hours a day, 7 days a week. It was this availability of personnel, in part, that sparked the innovative initiative to provide in-house computer maintenance and support services.

The benefits realized by the hardware support initiative have gone beyond the reduction of maintenance costs typically paid to a vendor support company. The Mean Time to Repair (MTTR) figures are also significantly reduced across all systems because the response time to hardware failures has been virtually eliminated. The lower MTTR has resulted in higher system availability figures.

The team has been able to accomplish this goal by making some strategic hires and providing computer hardware repair training for all of its members. Every member of the team can quickly diagnose and isolate a failed node in a system and replace it to bring the cluster up to 100% availability. The failed nodes are then repaired and tested offline utilizing a testing workstation that the team also maintains.

The maintenance and repair of supercomputer systems at LANL has provided a greater technical challenge for the HPCST staff. The computer operator's job has evolved from hands-off monitoring to monitoring and a hands-on maintenance and repair effort. Technical challenge, cost savings, and success have resulted in higher job satisfaction for the members of LANL's HPCST.

Weapons Radiochemistry

The radiochemistry team at Los Alamos has compiled the ingoing isotopic ratios of materials for four events of interest to LANL's Applied Physics (X) Division designers for the first quarter of FY08. Radiochemistry will process these data through a code developed by Chemistry Division, producing radiochemical parameters needed by X Division to run their codes.

Researchers who want more information on these radiochemical parameters may find two invited papers to Defense Research Review currently undergoing review. They are entitled: "Fundamental data and determination of fission yield as applied at LANL. Part I. (U)" and "Fundamental data and determination of fission yield as applied at LANL. Part II. Re-evaluation of historic 99-Mo K factors (U)."

The work of the LANL radiochemistry team was recently recognized by LANL in the form of a Large Team Distinguished Performance Award for their efforts on the Combined Nuclear Response Team. The team performed dedicated service to the U.S. government before, during, and after the October 9, 2006, North Korean nuclear test. The team demonstrated the ability to apply unique approaches to the challenging and evolving threat of nuclear proliferation.

The Combined Nuclear Test Response Team at LANL

ASC Salutes

"The reward for me," said **Brian S. Pudliner**, code physicist at Lawrence Livermore National Laboratory, "is to sit in the audience at a high-level technical briefing and hear about an issue of national interest being solved with a tool I helped create."

With a Ph.D. in Physics from the University of Illinois and a Master's degree in Computer Science, Brian began his career at Livermore as a computer scientist specializing in high performance computing (HPC). Continuing on the same project, he moved to the Livermore's Weapons & Complex Integration (WCI) Principal Associate Directorate (formerly, Defense & Nuclear Technologies Directorate) in conjunction with the ASC Program to work as a physicist and help laboratory scientists take advantage of HPC for large simulations of nuclear weapons. Brian now runs a project comprised of 22 code physicists and computer scientists. He and his team continue to push the boundaries of what is possible in weapon physics simulation—more physics at resolutions never reached before. The codes he and his team develop are used to address complex problems in stockpile stewardship by Livermore's weapons designers.

According to Frank Graziani, Livermore Boost Program lead, Brian and his team are a major factor behind the success of ASC as a national program and why code development WCI is of such high quality. "Brian's unique ability to combine rigorous scientific thinking with the pragmatic has resulted in a tool that pushes the envelope on scientific inquiry. What was impossible 5 to 10 years ago, Brian and his team have made possible. The fact that the tool Brian and his team have created is fast, robust, and accurate has resulted in it being the major workhorse for primary design at LLNL."

Brian was part of the original Burn Code Project Team. This team was cited for completion of the first full-system, 3D simulations of a nuclear weapon explosion, using a massively parallel, multi-physics weapons code. Today, Brian, his code team, and the laboratory's weapons designers use these original first-of-a-kind codes in their day-to-day work. What was once cutting edge—only three years ago—is now standard operating procedure.

"People who are curious thrive here," said Brian. There are an enormous number of unsolved problems in the physics and computer science arenas for this type of large-scale simulation. We are never in danger of getting bored. Being part of a multi-disciplinary team, at the edge of HPC, addressing problems of national interest, makes this an amazing environment to work in."

Upcoming Events

http://www.sandia.gov/NNSA/ASC/news/events.html

ASC Web Site

http://www.sandia.gov/NNSA/ASC/

Newsletter Points of Contact

Send submittals to:

Denise Sessions—Los Alamos National Laboratory denise@lanl.gov

Andrea Baron—Lawrence Livermore National Laboratory baron1@llnl.gov

Reeta Garber—Sandia National Laboratories ragarbe@sandia.gov

Who's Who

ASC Program Managers—Headquarters

Director, Office of Advanced Simulation and Computing, NA-114

Dimitri Kusnezov—Dimitri.Kusnezov@nnsa.doe.gov

Bob Meisner—bob.meisner@nnsa.doe.gov
Njema Frazier—njema.frazier@nnsa.doe.gov
Thuc Hoang—thuc.hoang@nnsa.doe.gov
Sander Lee—sander.lee@nnsa.doe.gov
Ed Lewis—edgar.lewis@nnsa.doe.gov
Erich Rummel—erich.rummel@nnsa.doe.gov
April Commodore—april.commodore@nnsa.doe.gov
Karen Pao—karen.pao@nnsa.doe.gov
Watti Hill—watti.hill@nnsa.doe.gov

ASC Program Managers—Labs

Los Alamos National Laboratory

John Hopson—jhopson@lanl.gov

Cheryl Wampler—clw@lanl.gov

Ralph Nelson, Deputy Program Director (acting)—ran@lanl.gov

Lawrence Livermore National Laboratory

Michel McCoy—mccoy2@llnl.gov

Lynn Kissel, Deputy Program Manager—kissel1@llnl.gov

Sandia National Laboratories

Tom Bickel—tbickel@sandia.gov

Paul Yarrington—pyarrin@sandia.gov

Martin Pilch—mpilch@sandia.gov

Art Ratzel—acratze@sandia.gov

James Peery—jspeery@sandia.gov

Developed and maintained by Sandia National Laboratories for NA.114

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company,
for the United States Department of Energy's National Nuclear Security Administration
under contract DE-AC04-94AL85000.