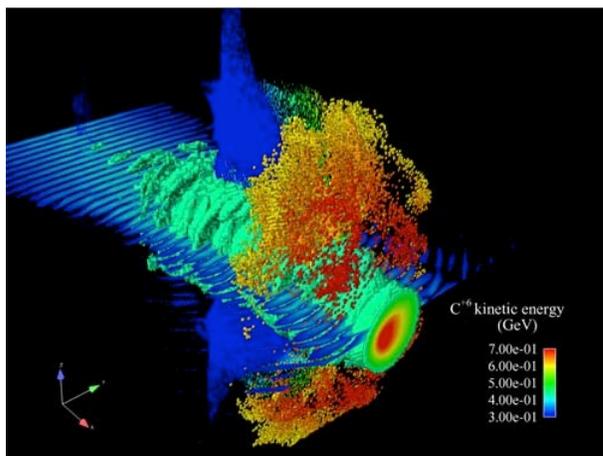


“The Meisner Minute” editorial does not appear in this issue. It will resume in the March issue of this newsletter.—*Editor’s note*

First-Ever 3D Kinetic Simulations of a Novel Laser-Driven Ion Acceleration Mechanism Enabled by Petascale Computing

A recent article in *Physical Review Letters*¹ is the result of extensive analysis of “Science at Scale” calculations during the stabilization and open science phases of Roadrunner. Enabled by petascale computing, scientists at Los Alamos National Laboratory (LANL) discovered a new class of laser-generated ion sources that can be used to resolve fundamental uncertainties in weapons physics codes.

This work explored for the first time the complex three-dimensional nature of a revolutionary new class of laser-ion accelerators, the “Breakout Afterburner” (BOA), that was discovered by the authors. Analysis of these surprisingly rich dynamics led to a result that overturned decades of conventional wisdom about symmetry breaking in the interaction of a high-intensity laser with plasma. The BOA mechanism was discovered in kinetic simulations using the authors’ VPIC code and was experimentally realized in recent experiments at the Trident laser facility at LANL. The BOA leads to laser-ion acceleration with (for the same laser intensity and spot size) order-of-magnitude higher ion energy and laser-to-ion conversion efficiency than conventional laser-accelerators, while at the same time making quasi-mono-energetic beams, as needed for applications such as weapons science experiments, ion-based “fast ignition” inertial fusion energy, and hadrontherapy of tumors. To find the article online, go to www.doi.org and type DOI:10.1103/PhysRevLett.107.045003.



¹ L. Yin, B. Albright, K. Bowers, et al., “Three-dimensional dynamics of breakout afterburner ion acceleration using high-contrast short-pulse laser and nanoscale targets,” *PRL* 107, 045003 (2011).

The above illustration depicts large-scale 3D VPIC simulations at realistic solid target density clearly showing the generation of a beam of GeV carbon ions. Combined with a robust experimental effort, the VPIC kinetic simulations of ion acceleration by short-pulse laser are enabling the development of a new generation of ion beam sources and may enable the success of ion-based fast-ignition inertial confinement fusion.

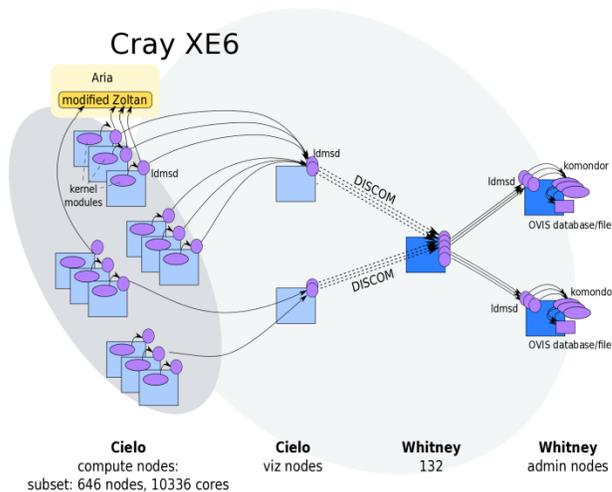
Feedback System Demonstrated for Dynamic Resource-Aware Computing on Cielo

Researchers at Sandia have successfully demonstrated a resource characterization and feedback system for dynamic resource-aware computing. The Sierra application Aria, running on 10,112 processing elements on Cielo, was rebalanced throughout its run time in response to dynamic resource conditions.

Resource characterizations used to drive re-partitioning via the Trilinos Zoltan partitioner were determined using Sandia's OVIS HPC monitoring and analysis system. This work demonstrated the design and development of integrated capabilities needed to enable an application to dynamically rebalance “in the face of changing application needs and platform state,” a key capability need identified by the NNSA Exascale Tools Working Group.

The efficacy of utilizing this type of feedback is tightly coupled to the ability to analyze how resources are being utilized. This, in turn, requires knowledge of what attributes to consider and how to weight the values of those attributes. Real time visualization and post-run analysis of data were used to determine attributes of interest and weighting functions. These analyses were performed remotely by streaming resource data from Cielo in New Mexico over a high-speed wide area network to the Sandia California TLCC cluster, Whitney.

The same wide-area architecture was demonstrated live at SC11 Conference (held in November in Seattle, WA) with data from Sandia's Cielo Del Sur Cray XE6 system streamed live to computers on the show floor. This work has piqued the interest of application developers, partitioner developers, and resource management researchers at Sandia and resulted in a joint FY12 LDRD (Laboratory Directed Research and Development) project to further develop interactions between large-scale applications and these enabling elements.



The illustration above represents the Architecture for Feedback System for Dynamic Resource-Aware Computing. An application, running on Cielo in New Mexico was rebalanced throughout its run time in response to dynamic resource conditions. Determination of rebalancing criteria was made by remote analysis of data streamed to the Whitney cluster in California.

Los Alamos ASC Team Effectively Analyzes Radioactive Debris from Mock Nuclear Attack

The ASC Physics & Engineering Models Threat Reduction (ASC/TR) project supported a recent international National Technical Nuclear Forensics (NTNF) Exercise named “Opal Tiger.” The drill successfully tested the readiness of the US and UK laboratories to perform nuclear forensic analysis on collected material samples from the detonation of an unknown nuclear device. In the Opal Tiger scenario, the detonation occurred in the UK. Assessment of the radiochemical and prompt data provided by the Data Evaluation team at Los Alamos National Laboratory during the exercise encompassed the solution space for both the characteristics of the materials and the weapon type. This accurate weapons-modeling assessment was enabled by new capabilities provided by the ASC Program for both the weapons physics code as well as the Forensics Inversion Tool Suite (FITS). The NNSA (including colleagues from LLNL) joined the FBI and the Defense Department, among others, in the Opal Tiger exercise.

SC11 Booth Showcases ASC Program’s Contributions to Science and Technology

This year’s Supercomputing Conference (SC11), held in Seattle from November 12-18, once again provided the ASC program an opportunity to showcase its contributions to science and technology through high-performance computing (HPC). The ASC booth theme was “Taking on the World’s Complex Challenges.”

“The ASC booth showcases the science and technology employed at the NNSA laboratories for DOE’s ASC program,” said Russ Goebel, this year’s booth lead and Sandia technical staff member. “It also gives us the opportunity to promote exchange and collaboration within the HPC community.”

The booth featured five booth zones: Advancing Science Frontiers, Impacting Global Issues, Technology Provides the Tools, Partnerships Accelerate Innovation, and Solutions through Collaboration. Models and simulations gave visitors a peek at the work being done in each area; short movie clips drew attention to the dawn of exascale computing, and the challenges and rewards it will bring. Scientists and researchers from Sandia, Los Alamos, and Lawrence Livermore National laboratories, along with university and industry partners, were also on hand to give presentations and demonstrations.

As the leading conference on HPC, networking, storage and analysis, SC provides attendees and exhibitors the opportunity to connect with the best and brightest in the HPC world. Nearly 11,000 SC11 attendees listened to technical talks, attended workshops, and visited exhibits on HPC technical advances and resulting modeling and visualization capabilities.



ASC Booth at SC11 in Seattle, WA.

IBM Unveils BlueGene/Q Supercomputer at SC11 Conference



Kim Cupps, Livermore Computing Division Leader at Lawrence Livermore, talks to reporters after the unveiling of BlueGene/Q (the cabinet to her right) at SC11 in Seattle.

The Advanced Simulation and Computing (ASC) BlueGene/Q supercomputing system that will be deployed at Lawrence Livermore National Laboratory (LLNL) as Sequoia was officially unveiled in a brief ceremony at the start of SC11. Jim Herring, IBM director of High Performance Computing (HPC) Offerings, kicked off the event in IBM's SC11 booth before a group of reporters. Kim Cupps, representing LLNL, paid tribute to the long-standing partnership with IBM and the computing breakthroughs that have resulted.

"We're looking forward to running codes on the 20-petaFLOP/s (quadrillion floating operations per second) Sequoia system in 2012," Cupps said. "We expect to achieve many exciting results in areas of national importance, including uncertainty quantification, materials modeling, energy modeling, laser plasma interaction, and climate change."

"This machine is an amazing achievement. We began our partnership with IBM more than 15 years ago with a goal of achieving 100 peak teraFLOP/s in 10 years," she said.

"We achieved our goal twice in 2005: with the 360-teraFLOP/s BlueGene/L machine and the 100-teraFLOP/s Purple machine."

Cupps noted that just six years later "we are standing in front of a machine 55 times more powerful than BlueGene/L and 10 times more energy efficient."

For the second time this year, BlueGene/Q was ranked number 1 on the [Green500 list of the world's most energy efficient computers](#). Energy efficiency remains one of the greatest challenges for next-generation exascale supercomputers.

"Energy efficiency was a critical factor in selecting this machine and will continue to be of paramount importance as we move toward more powerful machines in the future," Cupps said.

The BlueGene/Q Sequoia system was scheduled for delivery to LLNL for NNSA's ASC Program starting in December 2011, with deployment in 2012. When completed, Sequoia is expected to be one of the most powerful supercomputers in the world.

New Lawrence Livermore Supercomputer Tops Graph 500 Benchmark for Data-Intensive Computing



The BlueGene/Q Prototype II, currently located at the IBM T.J. Watson Research Center in New York and soon

to be delivered to Lawrence Livermore National Laboratory (LLNL) as the new Sequoia supercomputer for the Advanced Simulation and Computing (ASC) Program, won first place on the [Graph 500 list](#). Blue Gene/Q was able to traverse more than 254 billion graph edges per second (TEPS), more than two and one-half times more edges per second than the next machine on the list.

The win, announced to a packed room at the SC11 conference in Seattle, corroborates the new machine's data-intensive computing abilities. Accepting the award were Fabrizio Petrini of IBM and Kim Cupps, Livermore Computing Division Leader.

Petrini said of the BlueGene/Q prototype that "we've made tremendous progress," and that the new system is a "convergence of algorithmic and architectural innovation."

LLNL had multiple entries on this year's Graph500 list, including several entries submitted by Maya Gokhale and Roger Pearce that used solid-state drive storage arrays to hold the graphs. At no. 48 on the Graph 500 list, Leviathan (with a single 40-core node, 1 terabyte of memory, and 12 terabytes of Flash storage) was able to process a graph of 1 trillion edges. The Leviathan graph was 64 times larger than the LLNL top-ranked entry. Using Flash storage allowed the single compute node to solve a very large problem, but at a slower speed than the massively parallel supercomputer; thus its lower rank on the list. Gokhale notes that LLNL is able to address both speed (measured by TEPS) and scale (problem size) by using the two differing architectures. Read the press release, [LLNL Leverages ioMemory to Process 68 Billion Node Graph](#), on the Fusion-io blog.

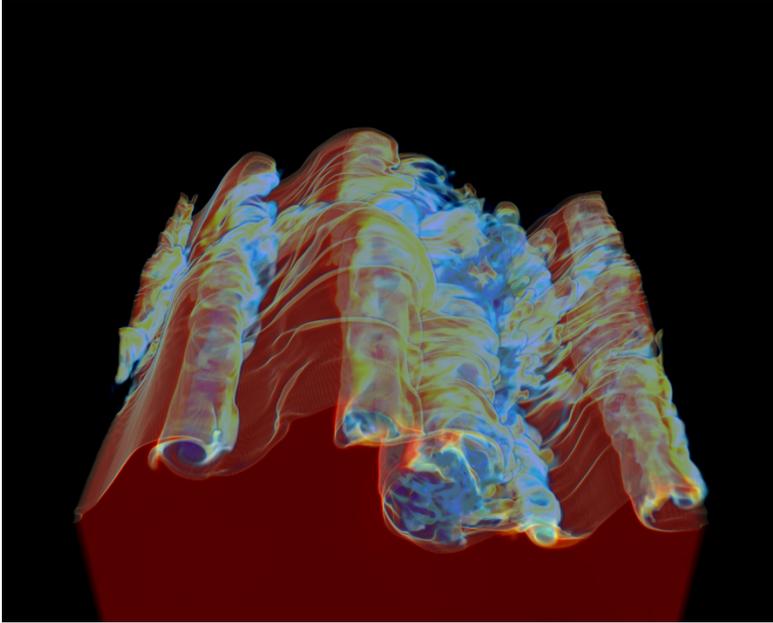
Graph 500 has been published for three years, growing from nine entries the first year to 50 entries this year. The list ranks the world's most powerful computer systems for data-intensive computing. Graph 500 gets its name from graph-type problems—algorithms—that are a core part of many analytics workloads in applications, such as those for cyber security, medical informatics, and data enrichment.

Sandia's NNSA/ASC Intel/Appro Many Integrated Core (MIC) Test Bed Results Showcased

Arthur, Sandia's first-of-a-kind NNSA/ASC MIC architecture cluster, was highlighted at SC11 with early simulation results showcased in the NNSA/ASC research exhibit and overview talks in the Intel and Appro booths. Sandia will use Arthur to support research into advanced computer architectures and looks forward to collaborating with Intel and Appro in investigating Intel MIC architecture. Arthur will support exploration of advanced programming models like OpenMP and Intel Cilk™ Plus that map current MPI applications to Intel® Xeon® processors with MIC architecture co-processors, and research into system software support for advanced data movement capabilities.

Sandia's Appro Xtreme-X™ multiple rack test bed consists of 42 heterogeneous nodes integrated with a QDR IB interconnect. The test bed will evolve in three phases. The baseline phase one test bed with 84 Intel® Xeon® 5600 processors and 84 Intel Knights Ferry co-processors was delivered by Appro in September 2011. Phase two is for the CPUs to be upgraded to Intel® Xeon® E5 processors sometime in early 2012. Phase three will be for the co-processors to be further upgraded with pre-production Intel® Knights Corner co-processors in 2012.

Arthur is integrated into Sandia's external collaboration network to facilitate interactions with collaborators at Intel, Appro, LANL, LLNL, and beyond.



University of Minnesota's turbulent fluid dynamics 3D visualization results showcased at SC11 using eight Knights Corner co-processors of Arthur.

Los Alamos Scientists Discuss Domain-Specific Languages at SC11

A significant topic of discussion over the last year within the ASC community and at SC11 is the impact that multi- and many-core processor architectures will have on mission-critical applications. Many key algorithms, methods, and overall software development techniques will need to be refined, redesigned, and re-implemented with the trends of increasing parallelism, concurrency, and complex memory hierarchies in mind. The next decade of computer systems will force software applications to spend an ever-increasing amount of time programming at the lower levels of rapidly changing hardware designs. The impact will be that less time will be spent focusing on critical, higher level, application requirements and capabilities.

To address this concern, scientists from the Applied Computer Science group at Los Alamos National Laboratory (LANL) propose that we need to design and program at a higher level of abstraction. Additional layers of abstraction will not only shield the applications from the nuances of the underlying architecture, but will also enable developers to write code more closely related to the domain of their application. While there are several traditional approaches to providing abstractions, such as object-oriented programming, LANL has been exploring the use of domain-specific languages (DSLs). The primary goal of this effort is to improve developer productivity and increase portability and performance.

As the name implies, a DSL provides the programmer with a set of domain-centric syntax, semantics and data structures for developing software. With this approach, a DSL compiler is capable of reasoning and optimizing based on knowledge of the domain itself.

In comparison, general-purpose compilers must address codes from an extremely wide range of domains and in the end must always err on the side of correctness when optimizing code. For example, a FORTRAN compiler will be unable to infer that a particular series of array access operations represent the traversal of an adaptive mesh data structure.

This approach has an added advantage of shifting portions of the workload away from the responsibilities of the application codes to the designers and programmers in the Computer Science community. To address the concern of DSLs requiring a full-blown, custom compiler infrastructure, LANL is actively exploring domain-centric constructs embedded within a full C/C++ compiler tool chain. An example of DSL code appears in the box below. This approach not only significantly reduces compiler development costs but also allows applications to be progressively modified to take advantage of DSL constructs, while the host language supports the remainder of the code. Furthermore, to be successful this approach requires enhanced collaborations between the computer and computational science communities.

```
int main(int argc, char *argv[])
{
    const int  NTIME_STEPS    = 5000;
    const int  MESH_WIDTH     = 512;
    const int  MESH_HEIGHT    = 512;
    const float MAX_TEMPERATURE = 100.0f;

    // Define a uniform mesh.
    uniform mesh HeatMeshType {
        cells:
        float t1, t2; // store two temperature fields at cell center.
    };

    // Create a two-dimensional instance of the mesh.
    HeatMeshType heat_mesh[MESH_WIDTH, MESH_HEIGHT];

    // Time steps loop.
    for(unsigned int t = 0; t < NTIME_STEPS; ++t)

        float d2dx2 = cshift(c.t1, 1, 0) - 2.0f * c.t1 + cshift(c.t1, -1, 0);
        d2dx2 /= dx * dx;

        float d2dy2 = cshift(c.t1, 0, 1) - 2.0f * c.t1 + cshift(c.t1, 0, -1);
        d2dy2 /= dy * dy;

        t2 = (alpha * dt * (d2dx2 + d2dy2)) + c.t1;
    }
    . . . .
}
return 0;
}
```

An example of DSL extensions embedded within the C++ programming language. In contrast to general-purpose languages, the compiler is capable of reasoning and optimizing based on the knowledge of the domain itself.

The domain-centric focus of a DSL presents a natural opportunity to support a cross-discipline co-design effort. A DSL must (by definition) be designed to support the requirements of the target domain. In addition, the language can be designed to support the effective, high-performance, utilization of emerging computer architectures. The DSL's compiler infrastructure can further facilitate the low-level mapping of algorithms to multiple low-level architecture design choices with minimal impact on the higher-level, application-aware abstractions.

LLNL Software Quality Engineering Represented at Two Worldwide Conferences

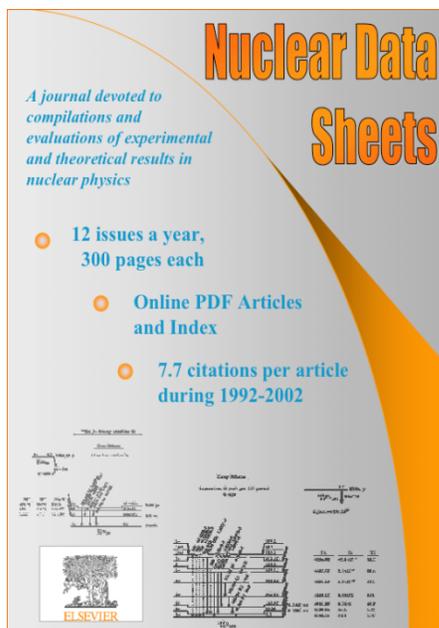
The Better Software Conference East was held November 6–11 in Orlando, Florida. Greg Pope, Software Quality Engineering (SQE) Group Leader and Verification and Validation (V&V) Project Leader for Advanced Scientific Computation (ASC), gave one of the keynote addresses, entitled "No Silver Bullet? Silver Buckshot May Work." Greg's talk reflected on his years in the industry and the many great processes and new tools that have been promoted, and his presentation was Webcast live worldwide.

"It seems that someone is always promising a cure-all—the proverbial 'silver bullet'—for software woes," Greg said. The most common request Greg gets from software developers and managers is to look at their development process and tell them "how to make it better." Greg's goals for his presentation were to promote understanding of what "better" really means, to discuss common problems and potential solutions, and to help conference attendees become empowered to make personal and group practices better. He examined valuable ideas that seem to reincarnate themselves periodically and explored the challenges of today's modern software. "Although there may not be a silver bullet for your software woes," Greg said, "perhaps there is 'silver buckshot'—a collection of techniques and tools to solve common problems—which, when properly aimed by capable professionals, will make your software better."

The Better Software Conference East offers more than 100 learning and networking sessions over six days, tutorials, training classes, keynotes, classes, bonus sessions, and other opportunities. Five multi-day training and certification courses covering testing, Scrum, product owner, agile, and other topics were offered, well as networking opportunities. In the training sessions, Greg and Ellen Hill from LLNL's Scalable Algorithms and Solvers Group presented a white paper entitled, "A Software Quality Engineering Maturity Model."

"Attending the conferences allows us to hear from and network with other industry leaders. It permits us to be aware and take advantage of the latest tools and techniques," Greg said. In October 2011, Bill Oliver, also in the SQE group, gave a talk on "Quantifying the Benefits of Static Analysis As It Relates to Testing" at the Starwest Software Testing, Analysis, and Review Conference in Anaheim that was well received. Bill also chaired two technical tracks. " As a result of these and other outreach efforts, we are getting new SQE business from other national labs and agencies, and we hope to see this new business area keep growing."

Los Alamos ASC Scientists Help Upgrade ENDF/B-VII.1 Nuclear Data Evaluations



In December 2011, the Evaluated Nuclear Data File (ENDF/B) was upgraded to ENDF/B-VII.1. This is the first upgrade since 2006. Researchers from LANL—from a variety of technical organizations—have played integral roles in obtaining new experimental data by developing new and revised neutron cross section and fission yield evaluations, validating the updated data base, and serving as lead authors of peer-reviewed publications to document this work.

Each month, the journal *Nuclear Data Sheets* publishes compilations and evaluations of experimental and theoretical results in nuclear physics. The titles of articles of five Los Alamos scientists published in the December 2011 *Nuclear Data Sheets* (published by Elsevier, Inc.) go to www.sciencedirect.com/science/journal/00903752. They are listed below. The articles document work performed to upgrade the nation's nuclear cross

section data base, the Evaluated Nuclear Data File. This file is maintained by the US National Nuclear Data Center at <http://www.nndc.bnl.gov/>. ENDF is the source for nuclear cross section libraries that scientists use in the nuclear transport packages in the ASC codes.

1. “ENDF/B-VII.1 Nuclear Data for Science and Technology: Cross Sections, Covariances, Fission Product Yields and Decay Data,” M.B. Chadwick et al., *Nuclear Data Sheets* 112 (2011) pp. 2887 – 2996.
2. “ENDF/B-VII.1 Neutron Cross Section Data Testing with Critical Assembly Benchmarks and Reactor Experiments,” A.C. Kahler et al., *Nuclear Data Sheets* 112 (2011) pp. 2997 – 3036.
3. “Quantification of Uncertainties for Neutron-Induced Reactions on Actinides in the Fast Energy Range,” P. Talou et al., *Nuclear Data Sheets* 112 (2011) pp. 3054 – 3074.
4. “Energy Dependence of Plutonium Fission Product Yields,” J.P. Lestone, *Nuclear Data Sheets* 112 (2011) pp. 3120 – 3134.
5. “Fission Product Yields for 14 MeV Neutrons on ^{235}U , ^{238}U and ^{239}Pu ,” M. MacInnes et al., *Nuclear Data Sheets* 112 (2011) pp. 3135 – 3152.

Participants Feel the Heat in LLNL Cyber Security Exercise



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

Out of the FIRE and into the INFERNO was the idea behind a cyber-security exercise hosted by Lawrence Livermore National Laboratory (LLNL) that brought together the NNSA's elite cyber responders. The cyber-security responders from NNSA sites around the country gathered for three days of cyber combat to repel a would-be intruder—the kind of attacks NNSA/DOE institutions see with increasing frequency.

Cyber attacks often target multiple institutions within DOE and NNSA (not to mention other federal agencies), necessitating a coordinated response, according to Neale Pickett, the Los Alamos National Laboratory (LANL) cyber security specialist who organized the exercise hosted by Livermore. "There's increasing recognition among the NNSA sites that we can no longer continue to operate as islands."

Called Tracer INFERNO, the exercise was based on an attack experienced earlier this year by a number of DOE/NNSA labs, and is a play on the Tracer Forensic Individual Response Event (FIRE) Pickett has organized at LANL, a workshop/competition that focuses more on internal team building and training for people new to the field.

Because it was based on a real incident and responders wanted to turn up the heat on the challenge, the exercise at LLNL was called an Intensive Network Forensic Exercise on Real Network Operations, or INFERNO for short. Conceived for NNSA's elite cyber responders, "Tracer Inferno is a more advanced event focused on building relationships across the complex," Pickett said.

"There's a lot of expertise at each of the labs," he said. "What is going to happen as a result of this exercise is that people are going to know each other and have experience working together. We will be better able to respond to future threats."

"These kind of cyber attacks are happening more frequently," said Matt Myrick, a senior cyber security engineer in LLNL's Cyber Security Program. "The more we talk among ourselves, the more we see patterns and the more we can share information across NNSA."

Participating in Tracer Inferno were 26 people from NNSA institutions, including Y12 (Oak Ridge, Tenn.), LANL, Sandia National Laboratories, Savannah River, Pantex, the Kansas City Plant, NNSA's Information Assurance Response Center, and LLNL.

"NNSA recognizes the need to build a complex-wide cyber-defense capability," Pickett said, noting that the exercise is a relatively inexpensive way to build such a response capability. "Everyone thinks it's a good idea," Myrick said. "Through exercises such as this we can develop a multi-institutional team of incident responders and we can have a pool of people to draw from. This is an idea that is now being formalized."

Greg Bronevetsky Receives Presidential Early Career Award

Lawrence Livermore National Laboratory (LLNL) computer scientist Greg Bronevetsky was named a recipient of a Presidential Early Career Award for Scientists and Engineers (PECASE) for helping advance the state-of-the-art in high performance computing (HPC). PECASE is the highest honor bestowed by the U.S. government on science and engineering professionals in the early stages of their independent research careers. Bronevetsky was one of 94 early career scientists and engineers to be recognized this year.



"To receive such recognition at this stage of my career is a great honor," Bronevetsky said. "This award is especially gratifying as it not only recognizes scientific achievement, but also the importance of this research to the nation."

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

"The research Greg is doing is critical to the development of next-generation supercomputers at a time of fierce global competition in HPC," said Dona Crawford, associate director for Computation at LLNL. "Greg's research embodies the spirit of innovation that is a hallmark of this laboratory and underscores Lawrence Livermore's global leadership in HPC."

The awards, established by President Clinton in 1996, are coordinated by the Office of Science and Technology Policy within the Executive Office of the President. Awardees are selected for their pursuit of innovative research at the frontiers of science and technology and their commitment to community service as demonstrated through scientific leadership, public education, or community outreach.

In 2010, Bronevetsky received a U.S. Department of Energy Early Career Award consisting of research funding of \$500,000 a year for five years. After receiving his doctorate in computer science from Cornell University, Ithaca, N.Y. in 2006, Bronevetsky came to Livermore as a Lawrence post-doctoral fellow. He became an employee in 2009.

Laboratory Physicist Recognized with NNSA Award



Steve MacLaren

Steve MacLaren, a physicist with Lawrence Livermore National Laboratory's (LLNL) Weapons and Complex Integration Principal Directorate, has been recognized with the NNSA Defense Programs' Employee of the Quarter Award. Recipients of the awards are recognized for going beyond the call of duty in supporting the mission of NNSA's Defense Programs.

"The Defense Programs Employee of the Quarter awards recognize the commitment of the men and women from throughout the national nuclear security enterprise," said Don Cook, NNSA's deputy administrator for Defense Programs. "The leadership and achievements of the award recipients have contributed directly to the many successes that NNSA Defense Programs has recently enjoyed."

MacLaren received the recognition based on his work as the lead designer in the recently completed High Energy Density (HED) experiments on both the National Ignition Facility (NIF) and the Z machine. These experiments delivered validation data for 3D simulations that enabled LLNL to develop and implement key physics-based models for the Stockpile Stewardship Program.

MacLaren led the effort in using state-of-the-art 3D Advanced Simulation and Computing Program (ASC) simulation capabilities, accounting for the experimental and diagnostics configurations, to perform both pre-shot simulations and post-shot analyses. Data showed excellent agreement with pre-shot simulation and validated LLNL's ability to predict complex radiation hydrodynamics behaviors.

"I feel very fortunate to be in a position to take advantage of the tremendous potential for science presented by the combination of LLNL's unique simulation resources with the emerging experimental capabilities at NIF and the Z machine," he said. "But for me, the real privilege has been the opportunity to work with teams of highly skilled, talented and enthusiastic scientists, whether code developers, designers, experimentalists or target fabrication experts, whose dedicated work is recognized by this award."

Leader in Applied Nuclear Physics Wins Prestigious Award

Mark Chadwick, leader of the Computational Physics (X-CP) Division at Los Alamos National Laboratory (LANL), has been honored with an E.O. Lawrence Award. (LANL's Advanced Simulation and Computing (ASC) Program is part of X-CP Division.)

Chadwick's award is in the national security and nonproliferation category for innovative scientific contributions to advance understanding of fission product yields and other key nuclear reactions resulting in the resolution of a longstanding problem in national security. His work on modeling neutron cross sections on plutonium, americium, uranium and radchem materials, with measurements by collaborators at the Los Alamos Neutron Science Center (LANSCE), has established robust metrics for validating simulations.

Chadwick chairs the national collaboration that oversees the development of Evaluated Neutron Data Files (ENDF), the US's evaluated cross section sets, and is lead author for the ENDF/B-VII database. Chadwick holds bachelors and doctorate degrees from the University of Oxford.

The X-CP Division develops the multi-physics simulation codes at Los Alamos, including the MCNP neutronics code and the Rage and Flag hydrodynamics codes, as well as some of the underlying materials, equation of state, plasma, atomic and nuclear models, databases, and algorithms.



Mark Chadwick from Los Alamos National Laboratory is honored with an E.O. Lawrence Award

ASC Salutes Greg Pope Software Quality Engineering in 3D



“*Making software better* means making it better in three dimensions,” said Greg Pope, Software Quality Engineering (SQE) Group Leader and Verification and Validation (V&V) Project Leader for the Advanced Science and Computing (ASC) program at Lawrence Livermore National Laboratory (LLNL).

“The software itself, the user’s and stakeholder’s experience of the software, and the experience of the scientists building the software all need to be continuously better,” he added. “We have learned from the past that prescriptive and a *one-size-fits-all* software quality approaches fail to achieve wide-spread adoption among researchers. Instead, we’ve found that a risk-based, graded approach works best to balance agility with discipline.”

“Prescribing that weapon scientists use a particular software process or tool is known as a ‘push’ strategy of adoption. Allowing scientists to choose from a number of attractive software development tools and process solutions that are compliant for a given risk level is the ‘pull’ strategy of adoption. The pull strategy of adoption allows the weapon scientists to dynamically choose from among the best contemporary processes and tools that meet their needs. It also institutionalizes the commitment to software quality. If the tool or process is not going to make the scientist’s development experience better, it is probably going to be resisted or bypassed, with good reason.”

This is Greg’s philosophy. With it firmly in place, he and his team work closely with ASC code developers and code end users to provide value-added SQE capabilities, such as compliance traceability, risk grading, assessments, static and dynamic code analysis, and build, test, release automation.

Greg admits that a 3D pull approach to software quality is more challenging to implement than simply beating people over the head with orders (also known as “audit and punish”), but he believes LLNL’s approach improves compliance while giving creative freedom to the ASC program’s scientists. Greg appreciates and credits the strong commitment of LLNL ASC management and hands-on practitioners to the success of software quality assurance. Predictive uses of the codes, high performance computing, and uncertainty quantification are all putting higher demands on software quality. What was acceptable even five years ago as quality software would fall short by contemporary software quality standards.

Greg has more than 40 years of experience developing software in the commercial and government sectors. Prior to joining LLNL in 2001, Greg founded and ran a software testing company, patented automated software testing tools, and held management and technical positions involving mission-critical testing of military systems and development of software code for avionics and aerospace uses including experimental flight test of helicopters. Greg has given industry keynote addresses, written technical papers, taught on software quality internationally, and has been a consultant.

For more on Greg's approach to SQE, see the article: [LLNL Software Quality Engineering Represented at Two Worldwide Conferences](#).

[Link to article earlier in content.]

ASC Relevant Research

Sandia National Laboratories Citations for Publications in 2011

1. Adolf, D. B., Neidigk, M. A., Neilsen, M. K., Chambers, R. S., Spangler, S. W., Austin, K. N. (2011). "Packaging Strategies for Printed Circuit Board Components, Volume I: Materials & Thermal Stresses," Sandia Technical Report. **SAND2011-4751**.
2. Brundage, A. L. (2011). "Modeling Compressive Reaction in Shock-Driven Secondary Granular Explosives," ASME Paper No. AJTEC2011-44130, pp.T20095-T20095-10, ASME/JSME 2011 8th Thermal Engineering Joint Conference (AJTEC2011), Honolulu, HI, USA. DOI: 10.1115/AJTEC2011-44130. **SAND2011-1516 P**.
3. Brundage, A. L., Gump, J. C. (2011). "Modeling Compressive Reaction and Estimating Model Uncertainty in Shock Loaded Porous Samples of Hexanitrostilbene (HNS)," *Bulletin of the American Physical Society*, 17th Biennial International Conference of the APS Topical Group on Shock Compression of Condensed Matter, Vol. 56, No. 6. Abstract ID: BAPS.2011.SHOCK.D2.5 <http://meetings.aps.org/link/BAPS.2011.SHOCK.D2.5>. **SAND2011-5103 C**.
4. Dalbey, K. R., Karystinos, G. N. (2011). "Generating a Maximally Spaced Set of Bins to Fill for High-Dimensional Space-Filling Latin Hypercube Sampling," *International Journal for Uncertainty Quantification*, Vol. 1, Issue 3, pp. 241-255. DOI: 10.1615/Int.J.UncertaintyQuantification.v1.i3.40. **SAND2010-7239 J**.
5. De Chant, L. J. (2010). "Expression for Supersonic Fluctuating Drag Force Magnitude due to Ambient Thermodynamic Disturbances," *AIAA Journal*, Vol. 48, No. 12, pp. 2976-2979. DOI: 10.2514/1.J050715. **SAND 2010-4229 J**.

6. Field Jr., R. V., Edwards, T. S., Rouse, J. W. (2011). "Modeling of Atmospheric Temperature Fluctuations by Translations of Oscillatory Random Processes with Application to Spacecraft Atmospheric Re-Entry," *Probabilistic Engineering Mechanics*, Vol. 26, Issue 2, pp. 231–239. Available online 15 July 2010. DOI: 10.1016/j.proengmech.2010.07.005. **SAND 2010-3069 J.**
7. Field Jr., R. V., Grigoriu, M. A. (2011). A Poisson Random Field Model for Intermittent Phenomena with Application to Laminar-Turbulent Transition and Material Microstructure," *Applied Mathematical Modelling*, Vol. 35, Issue 3, pp. 1142–1156. Available online 11 August 2010. DOI: 10.1016/j.apm.2010.07.059. **SAND 2009-2833 J.**
8. Gray, G. A., Fowler, K. R. (2011). "Traditional and Hybrid Derivative-Free Optimization Approaches for Black-Box Functions," Computational Optimization, Methods and Algorithms, *Studies in Computational Intelligence*, Vol. 356/2011, pp. 125-151. DOI: 10.1007/987-3-642-20859-1_7. **SAND2011-0724 P.**
9. Eldred, M. S., Swiler, L. P., and Tang, G. (2011). "Mixed Aleatory-Epistemic Uncertainty Quantification with Stochastic Expansions and Optimization-Based Interval Estimation," *Reliability Engineering and System Safety (RESS)*, Vol. 96, Issue 9, pp. 1092-1113. Available online 7 April 2011. DOI: 10.1016/j.ress.2010.11.010. **SAND2010-1799 J.**
10. Gray, G. A., Fowler, K. R. (2011). "The Effectiveness of Derivative-Free Hybrid Methods for Black-Box Optimisation," *International Journal of Mathematical Modelling and Numerical Optimisation*, Vol. 2, No. 2, pp. 112-133(22). Publication date 2011-04-01. DOI: <http://dx.doi.org/10.1504/IJMMNO.2011.039424>. **SAND2010-5448 J.**
11. Gullerud, A. S., Emery, J. M., Jamison, J. D. (2010). "Computational Assessment of Brittle Fracture in Glass-to-Metal Seals," Proceedings of the ASME 2010 International Mechanical Engineering Congress & Exposition, Vancouver, BC, Canada. **SAND2010-5605 C.**
12. Jones, R. E., Zimmerman, J. A., Templeton, J. A., Zhou, X. W., Moody, N. R., Reedy Jr., E. D., Kimmer, C. J., Delph, T. J., Oswald, J., Belytschko, T., Lloyd, J. T., McDowell, D. L. (2011). "Atom-to-Continuum Methods for Gaining a Fundamental Understanding of Fracture," Sandia Technical Report. **SAND2011-6031.**
13. Mitchell, J. A. (2011). "A Nonlocal, Ordinary, State-Based Plasticity Model for Peridynamics," Sandia Technical Report. DOI: 10.2172/1018475. **SAND2011-3166.**
14. Pautz, S. D., Pandya, T. M., Adams, M. L. (2011). "Scalable Parallel Prefix Solvers for Discrete Ordinates Transport in Multidimensions," *Nuclear Science and Engineering*, Vol. 169, No. 3, pp. 245-261. Online. **SAND2010-2212 J.**
15. Reedy, E. D., Boyce, B. L., Foulk, J. W., Field, R. V., de Boer, M. P., Hazra, S. S. (2011). "Predicting Fracture in Micrometer-Scale Polycrystalline Silicon MEMS Structures," *Journal of Microelectromechanical Systems (JMEMS)*, Vol. 20, Issue 4, pp. 922-932. DOI: 10.1109/JMEMS.2011.2153824. **SAND2011-8918 J.**

16. Romero, L. A., Torczynski, J. R., Kraynik, A. M. (2011). "A Scaling Law near the Primary Resonance of the Quasiperiodic Mathieu Equation," *Nonlinear Dynamics*, Vol. 64, No. 4, pp. 395-408. DOI: 10.1007/s11071-010-9870-8. **SAND2010-2054 J.**
17. Vaughan, C. T., Rajan, M., Barrett, R. F., Doerfler, D., Pedretti, K. T. (2011). "Investigating the Impact of the Cielo Cray XE6 Architecture on Scientific Application Codes", Workshop Proceedings, Workshop on Large-Scale Parallel Processing (LSPP), 25th IEEE International Symposium on Parallel and Distributed Processing, IPDPS 2011, Anchorage, AK, USA. DOI: 10.1109/IPDPS.2011.342. **SAND2011-3341 C.**

LALP-12-009