Stockpile Stewardship Quarterly

Defense Programs | *Stockpile Stewardship in Action* | Volume 3, Number 2

Message from the Assistant Deputy Administrator for Research, Development, Test Capabilities and Evaluation, Chris Deeney

Up to the moment a computerized numerically controlled machine first cuts metal at Plutonium Facility-4 at Los Alamos National Laboratory, or the first lithographic exposure occurs at the Microsystems and Engineering Sciences Application facility for a Life Extension Program, all our major activities are performed on computers. From design to certification, the Advanced Simulation and Computing (ASC) program provides the codes, platforms, file sharing systems, and even the helpline! The Accelerated Strategic Computing Initiative (ASCI) and its successor, ASC, blazed the path for and continue to drive the post-underground testing (UGT) stockpile stewardship and management program.

The use of high performance computing coupled with detailed verification and validation has transformed our business, and it is now the model for others. How transformative has it been? The answer is well described in a few articles in this issue of the *Stockpile Stewardship Quarterly*. In addition, Figure 1(a) shows a picture of our first 1 teraflop (TF) machine which required nearly 1,600 square feet of space and 850 kilowatts (kW) of power. Raj Hazra of Intel is seen in Figure 1(b) holding a 1 TF chip; it required 0.25 kW. How truly amazing! More sobering, however, is that a 1 exaflop system would need a million TF chips with a total power consumption approaching 1 gigawatt, which is approximately the power consumption of a Caribbean island.

Unfortunately, the scientific demands of stewardship and the future modernization efforts require us to continue to move from reliance solely on the old UGT database to a truly predictive capability. Moreover, the rigorous demands of Quantification of Margins and Uncertainties and modern design efforts call for thousands of calculations and the ability to perform full-system three-dimensional calculations on timescales that impact our design cycles, so we need to climb the next factor of 50 from Sequoia to exascale. It will not be easy as described in the last article, but we don't do easy!

Figure 1. (a) ASCI Red broke the teraflop barrier in December 1996 to become the world's first teraflop supercomputer (a teraflop represents one trillion mathematical operations per second.) (b) Raj Hazra of Intel is holding a 1 TF chip in 2012.

NNSA Bids Dr. Christopher Deeney Farewell

After 7 years with NNSA, Dr. Christopher Deeney is leaving federal service to serve as the Chief Technology Officer for National Security Technologies, LLC. Join the NNSA family in bidding him farewell and wishing him continued success. He will be missed!

Inside this Issue

2 Roadrunner—On the Road to Trinity
3 The Cielo Petascale Capability Supercomputer: Providing Large-Scale Computing for Stockpile Stewardship
7 Sequoia, NNSA's Most Powerful Computer, Prepares for Classified Work
8 Toward Exascale with Co-design and Proxy Applications

Comments
The *Stockpile Stewardship Quaterly* is produced by the NNSA Office of Research, Development, Test Capabilities and Evaluation. Questions or comments regarding this publication should be directed to Terri.Stone@nnsa.doe.gov.
Technical Editor: Dr. Joseph Kindel, Publication Editor: Millicent Mischo

Roadrunner—On the Road to Trinity

by Eileen Patterson, Bill Archer and Cheryl Wampler (Los Alamos National Laboratory)

Roadrunner, the world's first petascale computer, gave weapons scientists a view of the challenging supercomputing future. For NNSA, that future will be one step closer with a machine named Trinity.

Providing the high levels of assurance needed for stockpile assessments and life extensions is a daunting task, so the supercomputers provided for weapons scientists to use tend to push computation's leading edge. Los Alamos National Laboratory's (LANL's) Roadrunner supercomputer won world fame for setting a computing speed record. Roadrunner has completed its life and was decommissioned in March 2013, but NNSA continues to push computing capabilities to new, higher levels to meet the Stockpile Stewardship Program's (SSP's) growing needs. So the complex is anticipating its next big computer, a machine still to be fully conceived: Trinity. Trinity will be designed and procured by a LANL-Sandia National Laboratories partnership—the New Mexico Alliance for Computing at Extreme Scale—and physically sited at Los Alamos.

Roadrunner was a step on the road to Trinity, and it was a huge step—one for the record books.

Built for Speed

Like its namesake, New Mexico's state bird, Roadrunner was certainly speedy. In June 2008, it placed #1 on the Top500, the international list of the world's fastest supercomputers. It was the world's first petascale supercomputer, the first to reach a sustained quadrillion (thousand trillion or 1,000,000,000,000,000) floating-point operations per second: petaflops.

Roadrunner also had a unique, and controversial, design (architecture) for a supercomputer. It combined two different kinds of processors, making it a hybrid. It had 6,563 dual-core general-purpose processors (AMD Opterons), with each core linked to a special graphics processor (PowerXCell 8i) called a "Cell." The Cell was an enhanced version of a specialized processor originally designed for the Sony Playstation 3, adapted specifically to support scientific computing.

Although hybrid computers had existed before, no one had ever tried that approach at supercomputing scale, and many doubted that a hybrid supercomputer would work. So for LANL and IBM, collaborators in designing and building the computer, Roadrunner was a leap of faith—but a leap with purpose.

High-speed calculation was part of the purpose. Although simulation only predicts weapon performance while in contrast nuclear testing demonstrated it, simulation pays a two-part dividend: greater detail about the internal processes feeding into performance and increased fidelity in spatial and temporal resolution. When a computer is fast enough to improve detail and fidelity, with reasonable turnaround time, the resulting simulations deepen scientists' understanding of weapon behavior. As part of its SSP work, Roadrunner took on a troubling, long-standing gap in understanding about energy balance in a weapon and its relation to yield and made a significant contribution to that understanding. In addition to resolving these issues from our test history, Roadrunner also performed full-system simulations that match selected data for certain underground nuclear tests for the first time. Stockpile and weapons science studies that were enabled by Roadrunner include the following projects:

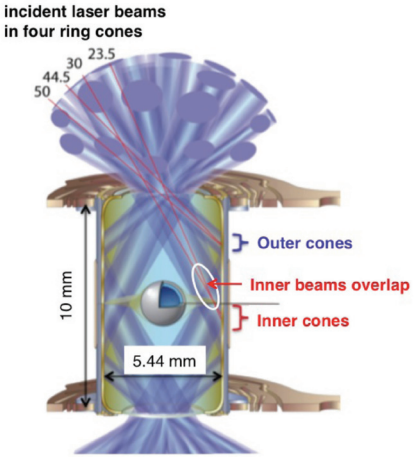

Figure 1. Schematic of a NIF ignition target showing the inner and outer laser beam cones and the simulation region where inner beams overlap. The simulation uses the VPIC Particle in Cell code.

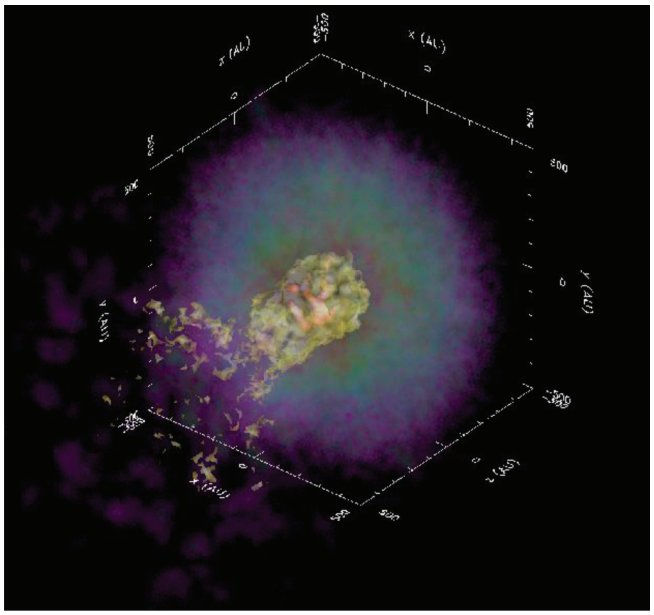

Figure 2. Researchers used Roadrunner to predict light (radiation) output from supernovae. The figure shows a 3D computer radiation-hydrodynamics simulation of a supernova.

- The largest plasma simulation ever done to date to understand how laser energy couples to plasma instabilities and what limits the backscatter, contributing to an understanding of laser-driven inertial confinement fusion ignition (see Figure 1).

- Significant scientific accomplishments on Supernovae Light-Curves: A hybrid rad-hydro code (Cassio) was developed and demonstrated on Roadrunner with the Implicit Monte Carlo radiation package Cell-accelerated and the Eulerian hydrodynamics running unchanged on the Opterons (see Figure 2).
- Critical calculations to complete the energy balance Level-1 milestone.
- LANL simulations for peer review of Lawrence Livermore National Laboratory resolution of a B83 issue.
- Simulations to validate an improved mix model for use in integrated multi-physics codes.
- Predictive Capability Assessment Suite (PCAP) studies to improve our ability to understand and measure progress of a predictive simulation capability for stockpile stewardship by improving models by matching simulations with data from underground nuclear tests.
- Direct Numerical Simulations to improve understanding of turbulence.
- Studies on the effects of nuclear detonation in ionosphere/inner magnetosphere.

However, the need to improve detail and fidelity for full scale weapons calculations continues to drive the need for even faster and bigger computers.

The Cielo Petascale Capability Supercomputer: Providing Large-Scale Computing for Stockpile Stewardship by Douglas W. Doerfler (Sandia National Laboratories) and Manuel B. Vigil (Los Alamos National Laboratory)

The Cielo supercomputer was deployed in February 2011, by the Los Alamos National Laboratory (LANL) and Sandia National Laboratories' (SNL's) Alliance for Computing at the Extreme Scale (ACES) partnership (see sidebar on page 4) to solve the most challenging problems and the largest parallel simulations through the Advanced Simulation and Computing (ASC) Program for the NNSA tri-laboratories—Lawrence Livermore National Laboratory, LANL, and SNL. Now in its fifth Capability Computing Campaign (CCC5), Cielo is working fulltime for stockpile stewardship. It was designed to provide a stable, robust environment for both the weapons designer/analyst communities and the weapons code development community. Designers and analysts run simulations at near-petascale with integrated weapons codes requiring little adaptation. Meanwhile, code developers explore the code modifications necessary for the next generation of computers. The need for three-dimensional (3D), full-scale end-to-end weapons simulations with increased fidelity is one of the major drivers of computing requirements for stockpile stewardship, Cielo has taken a step toward meeting those requirements.

CCC5, which began on August 5, 2013, carried an impressive workload for stockpile stewardship and weapons science. The machine is in such high demand that allocation requests exceeded the available computing resource by more than 5½ times. The number of campaigns that include simulations that each will consume nearly the entire machine is growing. Five campaigns have plans to include jobs each needing more than 128,000 cores: 3D weapons studies, safety studies, and 3D particle transport studies. Co-design scaling studies by SNL will use the full system.

The challenge to stockpile stewardship is to accurately assess the behavior of weapons with the goal of ensuring robust and reliable performance without the benefit of nuclear testing. This drives the need for fine-scale numerical resolution and advanced models for physics and material behavior, for both weapon system simulations and for the sciences supporting the models in those simulations. Cielo is being used for 3D full-system weapon calculations with medium spatial-temporal resolution and more fidelity in computationally complex physics models. Cielo does not have sufficient computational capability for full-system weapons calculations that require 3D and high spatial-

Alliance for Computing at the Extreme Scale

Los Alamos National Laboratory (LANL) and Sandia National Laboratories have collaborated to create a New Mexico center for high performance computing, the Alliance for Computing at the Extreme Scale (ACES). ACES is funded by the DOE/NNSA's Advanced Simulation and Computing program[1] and was formed to enable the solution of critical national security problems through the development and deployment of high performance computing technologies. The first ACES project was developing and deploying a production petascale supercomputer, Cielo.[2,4,5] Cielo was deployed in 2011 and is sited at the Strategic Computing Complex Facility (also known as the Nicholas C. Metropolis Center) at LANL.

temporal resolution; these types of calculations are the design point for the next ASC Advanced Technology System, Trinity, which is targeted for the 30-60 petaflop range, and will be installed at LANL through the ACES partnership in 2015.

Examples of calculations that have run on Cielo include:

- Simulation of long-standing nuclear test mysteries in order to improve and/or validate our scientific understanding of key weapons physics;
- Detailed simulation of primary implosion and burn;
- Simulations for assured safety and security, performance, and survivability of non-nuclear weapons system components;
- Improved National Ignition Facility (NIF) target designs, implosion stability studies, and beam propagation and laser-plasma interactions;
- Warm dense plasma simulations;
- Materials studies, such as strength and phase;
- Validation of turbulence models for Rayleigh-Taylor instabilities; and
- Simulations to support secure transportation.

Cielo is a Cray XE6 architecture supercomputer.[3] It was installed in two phases: the first phase was completed in December 2010, and in the second phase, the system was upgraded from 1.03 to 1.37 PF (107,264 to 143,104 compute cores). In fall of 2011, Cielo successfully passed a major program review—the ASC Level 2 Cielo Production Capability Readiness Milestone–in which the tri-lab user community validated that the machine is successfully providing production capability compute cycles for their simulations.

The following paragraphs discuss simulations run on Cielo.

Asteroid Simulation. Cielo is the first computer large enough in terms of both computing power and memory to run a large-scale asteroid mitigation problem. This research addresses the threat of potentially hazardous objects (PHOs) in space that might collide with Earth. Many methods of mitigating PHOs have been proposed, including nuclear options that would cause an explosive disruption, or stand-off momentum/velocity transfer. Nonnuclear methods include gravity attractors and solar energy absorption. In Figure 1, scientists are calculating a nuclear explosion on the surface of a fictitious asteroid, which is nonspheical like the asteroid Itokawa and composed of a rubble pile rock.[6]

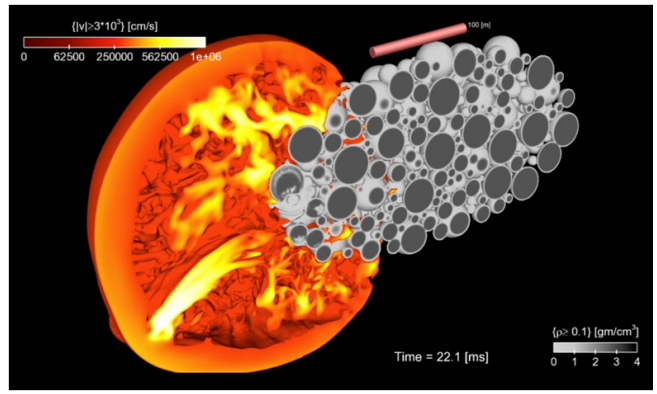

Figure 1. A full-physics, full-geometry, 3D simulation using 32,000 processors on Cielo to model a 500-kiloton nuclear blast to explode a 5-kilometer ficticious asteroid.

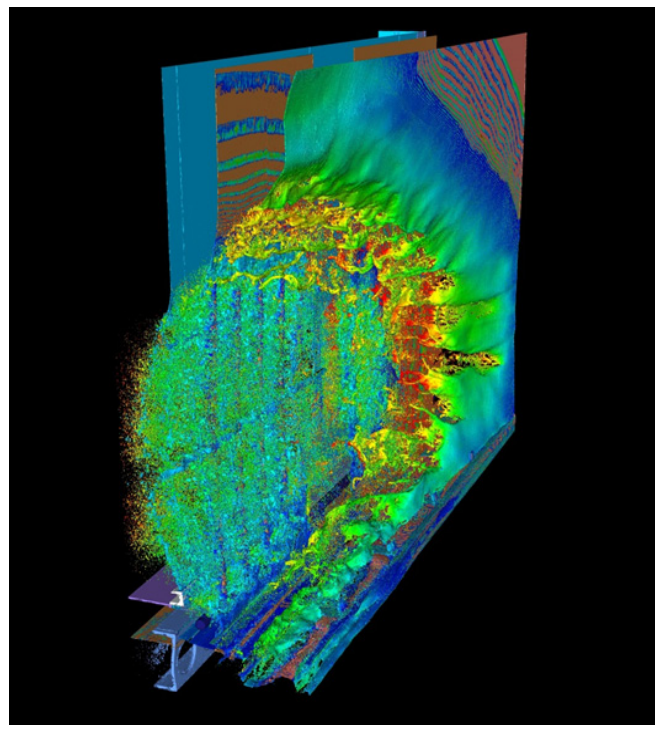

Figure 2. Response of lightweight, blast-resistant structure as simulated by CTH.

Structural Response to Blast. A 32,768-processor simulation using the shock physics code CTH helps engineers understand the response of structures under severe blast

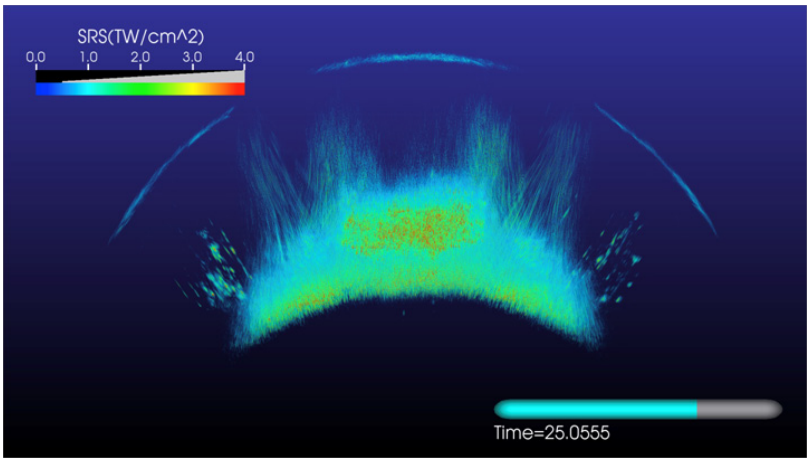

Figure 3. Visualization of the laser light backscattered by Stimulated Raman Scattering (SRS) as calculated by PF3D. The SRS occurs in bursts and is shown near the end of a burst when the scattered light comes mostly from the region where the quads overlap (the arc of light at the bottom).

loading conditions so that the robustness of these structures may be improved. In the simulation in Figure 2, a structure composed of sheet metal is loaded by an explosive blast. The plot shows the velocities of various components of the structure as they are torn apart. Accurately resolving these thin parts requires a very fine mesh resolution, which can presently be achieved using only very large-scale computing platforms such as Cielo.[7]

Integrated Simulation of Laser Plasma Interaction (LPI) in NIF Experiments. This project investigates how to account for backscattered light in a self-consistent manner in simulations of ignition experiments at NIF. The pF3D code is used on Cielo to simulate backscatter in overlapping beams; interesting temporal and spatial structure is seen. Cielo was recently used to simulate three interacting beams, as shown in Figure 3. The spatial variability of the backscattered light shows an enhancement due to the interaction of the beams.[8]

Summary

Cielo has delivered to the Stockpile Stewardship Program and the tri-lab weapons community the ability to perform the most taxing simulation, those that can only be performed at the petascale. Cielo's architecture was chosen to allow migration of the existing integrated weapons codes with minimal effort, thus extending their ability to conduct critical simulations while the code teams work to adapt the codes to the challenging new computing architectures.

References

[1]http://nnsa.energy.gov/asc

[2]http://www.lanl.gov/orgs/hpc/cielo/index.shtml

[3]http://www.cray.com/Products/Computing/XE.aspx

[4]James Ang, D. Doerfler, S. Dosanjh, S. Hemmert, K. Koch, J. Morrison, M. Vigil, "The Alliance for Computing at the Extreme Scale," Cray Users Group 2010 Proceedings, Edinburgh, Scotland, May 24-27, 2010.

[5]Doug Doerfler, M. Vigil, S. Dosanjh, J. Morrison, "The Cielo Petascale Capability Supercomputer," Cray Users Group 2011 Proceedings, Fairbanks, Alaska, May 23-26, 2011.

[6]Method: The RAGE Hydrocode in 3D, with a 1-megaton Energy Source on the Surface of the Object. Researcher: Bob Weaver, Los Alamos National Laboratory.

[7]Method: CTH Finite Difference Code. Researchers: Steve Attaway, Shivonne Haniff, Joel Stevenson, and Jason Wilke; Sandia National Laboratories.

[8]Method: Code pF3D. Researcher: Steve Langer, Lawrence Livermore National Laboratory.

New Centers of Excellence for Academic Computational Science Partnerships

NNSA recently selected six universities as centers of excellence that will focus primarily on predictive science. Listed below, the first three are Multidisciplinary Simulation Centers (MSCs) and the others are Single-Discipline Centers (SDCs):

- Uncertainty Quantification-Predictive Multi-disciplinary Simulation Center for High Efficiency Electric Power Generation with Carbon Capture, University of Utah, Salt Lake City
- Center for Exascale Simulation of Plasma-Coupled Combustion, University of Illinois Urbana-Champaign
- Predictive Simulations of Particle-laden Turbulence in a Radiation Environment, Stanford University, Stanford, California
- Center for Compressible Multiphase Turbulence, University of Florida, Gainesville
- Center for Exascale Radiation Transport, Texas A&M University, College Station
- Center for Shock Wave-processing of Advanced Reactive Materials, University of Notre Dame, South Bend, Indiana

The MSCs will receive $3.2 million and the SDCs will receive $1.6 million each year for five years under NNSA's Predictive Science Academic Alliance Program II (PSAAP II) agreement. The PSAAP II centers will develop the science and engineering models and software for their large-scale simulations utilizing methods of verification and validation and uncertainty quantification, with an additional focus on extreme-scale computing. The goal of these disciplines is to enable scientists to make precise statements about the degree of confidence they have in their simulation-based predictions. For more information about PSAAP II, visit http://nnsa.energy.gov/asc and select the *University Partnerships* link.

Development of X-ray Thomson Scattering Diagnostic for Z Accelerator

X-ray Thomson scattering (XRTS) is an important diagnostic designed to measure temperature and warm dense matter properties on Z experiments. As part of the effort to implement the XRTS diagnostic to measure dynamic materials properties (DMP) on Z, the diagnostic was characterized first in the Z-Beamlet Laser (ZBL) calibration chamber and subsequently on Z. ZBL irradiated a manganese (Mn) foil to generate 6-keV x-rays that scattered from an ambient TPX (CH_2) foam sample mounted onto a Z-DMP load. The x-ray scattering signal was measured with high spatial and spectral resolution, thus validating the experimental design. Figure 1 shows the spectrally- and spatially-resolved data obtained on a Z DMP shot from the Mn source x-rays and the 90° scattered x-rays from the foam. Four shots have been conducted so far on Z to develop the XRTS capability (Shot 2436 in November 2012, Shot 2474 in February 2013, Shot 2490 in March 2013, and Shot 2526 in June 2013). Obtaining the data is challenging: ZBL had never been implemented on Z DMP experiments and the spectrometer is subjected to Z-generated debris and x-ray background. The XRTS technique will be used on Z to measure the temperature in shocked beryllium and lithium deuteride.

—Tommy Ao, Eric Harding, and Jim Bailey of Sandia National Laboratories and the XRTS Team

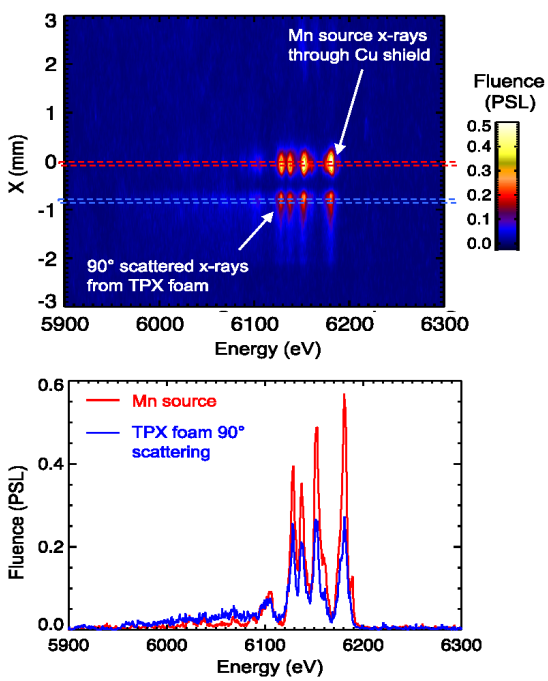

Figure 1. Spatially-resolved spectra on Shot 2474 from manganese source x-rays and x-rays scattered from TPX foam. Fluence is measured in photostimulated luminescence.

Sources Sought Notification

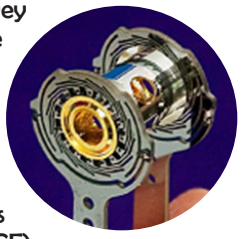

NNSA is conducting a market survey to identify interested parties capable of providing manufacturing services to support the fabrication of targets for use at the National Ignition Facility (NIF). NNSA has an ongoing need for a wide variety of very complex targets for inertial confinement fusion (ICF) and related high energy density physics (HEDP) experiments. ICF/HEDP targets are complex, multi-component systems often employing specialized materials. Many of the targets are experimental and involve tasks that can only be performed at Lawrence Livermore National Laboratory (LLNL). However, NNSA is interested in exploring if some of the standard targets, target subcomponents, or target preparation processes could be performed by other contractors to NNSA or subcontractors to Lawrence Livermore National Security (LLNS). Through this Sources Sought notification, NNSA hopes to determine if sufficient market capability exists outside of LLNS to warrant proceeding with a possible solicitation in this area.

To explain the ICF and high energy density fusion programs, the NIF, and the industrial capabilities required to manufacture targets, NNSA will be hosting a Target Fabrication Industry Day on September 11, 2013. The event will be held at LLNL in Livermore, California. Details about the event and how interested parties may participate are available at www.nnsa.energy.gov/icf/niftargetfabday.

Note: This is not a solicitation announcement for proposals. No contract will be awarded from this announcement.

Target Fabrication Industry Day
September 11, 2013
www.nnsa.energy.gov/icf/niftargetfabday

● ● ● ● ●

Fusion Power Associates (FPA) Awards

David Meyerhofer of the University of Rochester Laboratory for Laser Energetics will receive the 2013 FPA Leadership Award for his many scientific and technical contributions to fusion development and for his leadership. The recently retired David Crandall has been selected for the FPA 2013 Distinguished Career Award for his decades of contributions to fusion research. The awards will be presented at the annual meeting and symposium on December 10-11, 2013 in Washington, DC.

The FPA awards have been given annually since 1980 to individuals who have shown outstanding leadership qualities in accelerating the development of fusion as an energy source. A list of previous recipients and information about the FPA annual meeting and symposium are posted at http://fusionpower.org.

Sequoia, NNSA's Most Powerful Computer Begins Tri-Lab Classified Work by Brian Carnes and Kim Cupps (Lawrence Livermore National Laboratory)

Sequoia (see Figure 1), a world-class IBM BlueGene/Q computer sited at Lawrence Livermore National Laboratory (LLNL) by the DOE/NNSA Advanced Simulation and Computing (ASC) Program, started its first classified Capability Computing Campaign (CCC) on July 1. Work perfomed by researchers from NNSA's three nuclear weapons laboratories supports the Stockpile Stewardship Program—the effort to ensure the safety, security, and reliability of the nation's nuclear deterrent without underground testing. Sequoia's mammoth computational power will be used to assess physical weapons systems and provide a more accurate atomic-level understanding of the behavior of materials in the extreme conditions present in a nuclear weapon. Prior to transitioning to classified work in early April, Sequoia was used to explore a broad range of science to fully develop the capabilities of the system. Tri-lab researchers tested Sequoia's power and versatility by running unclassified science codes relevant to NNSA missions. The early science runs were part of the "shakeout" of the 20-petaFLOP/s peak IBM BlueGene/Q system.

Figure 1. Fully integrated Sequoia system, December 2012, sited at Lawrence Livermore National Laboratory.

Sequoia is the third generation of BlueGene supercomputers developed by the IBM-LLNL partnership during the past 10 years. Its predecessors, BlueGene/L and BlueGene/P, were quite successful and won a number of awards, including the National Medal of Technology and Innovation in 2009, given by President Barack Obama. The BlueGene/Q architecture brings many innovations over the previous generations, including 16 cores per node, multithreaded cores, a five-dimensional torus interconnect, water cooling, and optical fiber links. The system was designed for optimal power efficiency, reliability, and price-performance. Sequoia with its many specialized low-power cores, has a sustained speed of 16 petaflops while consuming only 8 megawatts of electrical power. The 20 petaFLOP/s Sequoia system has a staggering 1.6 million processor cores. Sequoia's architecture introduces many of the challenges associated with future computing architectures, including harnessing massive parallelism effectively and achieving reasonable reliability, availability, and serviceability given millions of components and 11.5 miles of fiber-optic cable.

The integration of Sequoia took place between January (first rack deliveries) and December 2012 (final system acceptance). The integration challenges were varied and daunting. They included issues with water-system commissioning, compute-rack weights, faulty glue on electromagnetic interference gaskets, innovative power supplies not sharing load under extreme conditions, compute-card pins bent during hardware maintenance, link errors, hot optics and, most significantly, an application-specific integrated circuit quality defect discovered relatively late in the acceptance process. Early unclassified work on the machine allows NNSA researchers and IBM computer scientists to work out the bugs and optimize the system before it transitions to classified work and the limited access and security that entails. Los Alamos National Laboratory researchers have run asteroid and turbulence simulations, and Sandia National Laboratories scientists have been scaling the Large-Scale Atomic/Molecular Massively Parallel Simulator (LAMMPS) compute-intensive physics module, named SNAP, to prepare for future simulations exploring the properties of tantalum (see Figure 2).

Initial efforts by LLNL scientists include a QBox first principles molecular dynamics code examination of the electronic structure of heavy metals, research of interest to stockpile stewardship. Qbox was developed at LLNL to perform large-scale simulations of materials directly from first principles, allowing scientists to predict the properties of complex materials without having to carry out experiments.

In addition, LLNL scientists have investigated thermo-nuclear (TN) burn in doped plasmas, exploiting the full capability of Sequoia and the code developed for this

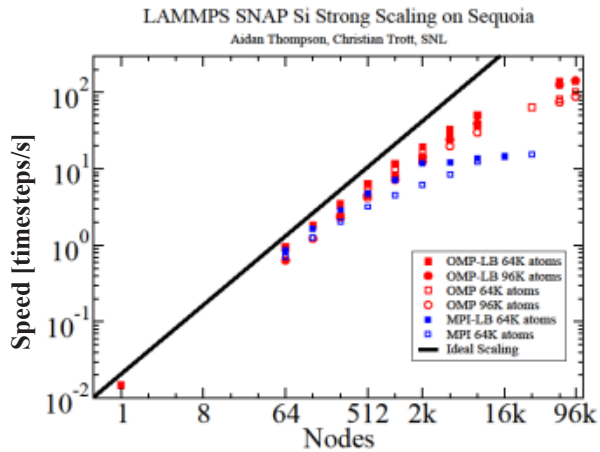

Figure 2. Performance of different SNAP implementations on Sequoia.

purpose. Following a benchmark exploration of the density and temperature dependence of TN burn in undoped hydrogen plasma, researchers began a series of extreme-scale simulations of TN burn in the presence of small fractions of a percent of high-Z dopants. These studies will be used to deepen scientists' understanding of the effect of dopants on TN burn, physics that is vital to capsule design for the National Ignition Facility, LLNL's laser fusion experiment.

Sequoia also has demonstrated its amazing scalability with a three-dimensional simulation of the human heart's electrophysiology. Using Cardioid, a code created in a partnership between LLNL and IBM scientists, researchers modeled the electrical signals moving throughout the heart. Cardioid has the potential to be used to test drugs and medical devices prior to tests on humans.

Sequoia's power enables suites of highly resolved uncertainty quantification (UQ) calculations. UQ is the quantitative characterization and reduction of uncertainty in computer simulations through the running of very large suites of calculations that characterize the effects of minor differences in the systems. Sources of uncertainty are rife in the natural sciences and engineering fields. UQ uses statistical methods to determine likely outcomes.

In the June 2013 Top500 list, Sequoia dropped to number three behind China's Tianhe-2 machine and Oak Ridge National Laboratory's Titan. However, Sequoia retained its number-one ranking on the Graph 500, the ability to solve big data problems—finding the proverbial needle in the haystack.

Toward Exascale with Co-Design and Proxy Applications by Allen McPherson (Los Alamos National Laboratory), Rob Neely (Lawrence Livermore National Laboratory), and Robert Hoekstra (Sandia National Laboratories)

NNSA's Stockpile Stewardship Program (SSP) requires continual increases in computing performance to support demands for increased physics fidelity, more detailed computational models, large-scale three-dimensional (3-D) simulations (for example, Figure 1), and predictive capabilities with narrow uncertainty bounds. The Advanced Simulation and Computing (ASC) program delivers these capabilities. Computer system vendors have responded to this need by developing increasingly powerful computers. In the past, this increased capability was provided by, among other techniques, increasing the clock rate at which processors executed instruction—a faster clock rate directly drove faster execution of codes. Recently, electrical power consumption has become the limiting factor in chip design. If faster cores are not possible, increased system performance can be achieved by using more cores—many more cores. Sequoia, the most recent ASC Advanced Technology System (ATS) housed at Lawrence Livermore National Laboratory, is an example of this massively multicore architecture—an IBM BlueGene/Q platform with over 1.6 million cores.

This quest for power efficiency also drives a move to heterogeneity in architectures through the addition of accelerators and co-processors that specialize in floating

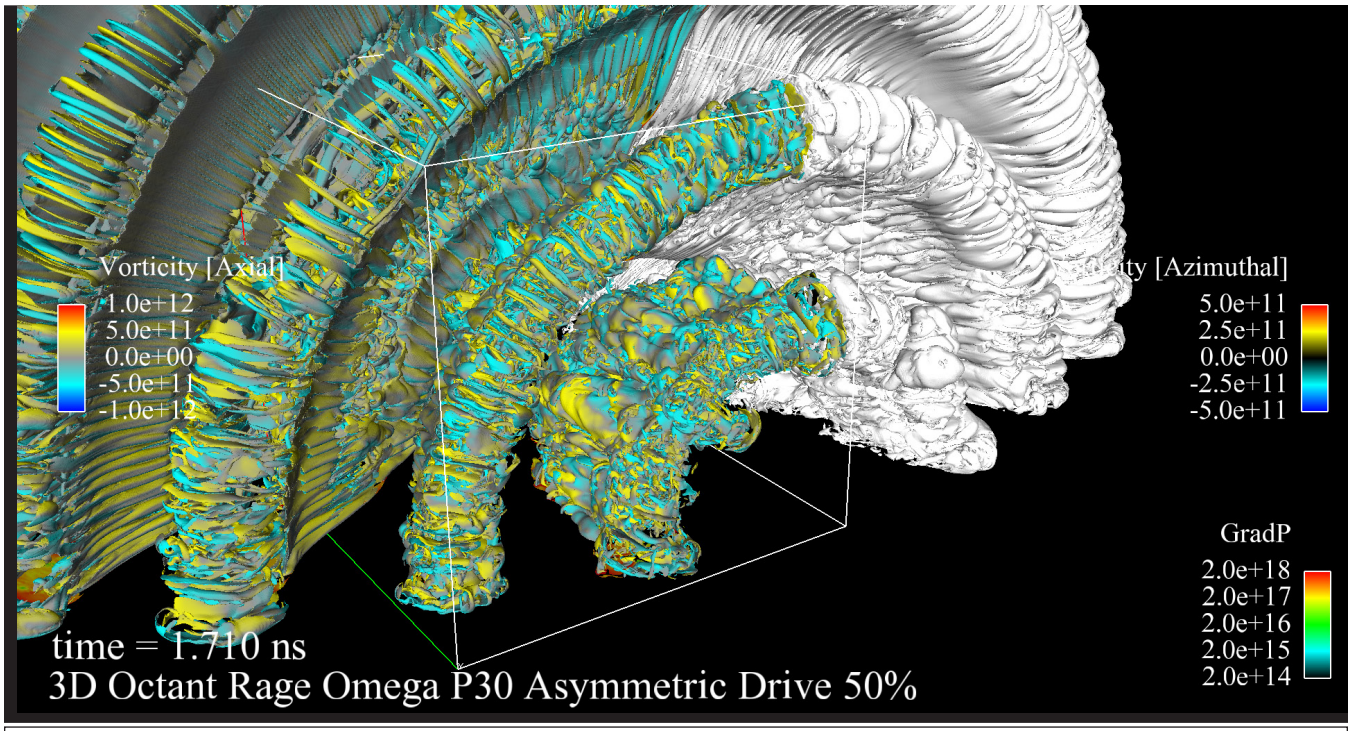

Figure 1. A late time snapshot from the 0.05 micron 3D RAGE simulation of the P30 Omega capsule showing a close-up view of the development of the turbulence in the two bubbles nearest the polar axis. Image taken from LA-UR-11-03527, courtesy of Vincent Thomas and Robert Kares, Los Alamos National Laboratory.

point performance at much lower energy per floating point operation. Los Alamos National Laboratory (LANL) pioneered the use of heterogeneous configurations with the Roadrunner system's hybrid nodes containing AMD Opteron processors and Cell accelerator boards. As another early example, Oak Ridge National Laboratory's (ORNL's) Jaguar machine combined Opteron processors with NVIDIA Fermi graphics processing units (GPUs) as accelerators. Jaguar paved the way for Titan at ORNL—an even larger GPU-based system. Together, Titan and Sequoia now represent the top two fastest machines in the Nation today, making it clear that the trend toward these more complex systems is well underway.

Whether through the use of accelerators such as GPUs, or massive parallelism using huge numbers of computational cores, these architectural trends require significant evolution of ASC's Engineering and Physics Integrated Codes (EPICs) to fully exploit the billion-way parallelism expected to realize the performance benefits of these new architectures. Many of the high performance computing (HPC) system characteristics we've come to rely on for the past 15 to 20 years are being turned on their head, and the way we design our systems and develop our software must evolve with the changes using a strategic and holistic approach.

Co-Design

As we begin to fully grasp the degree of change required for continued accelerated success in HPC, DOE/NNSA has embraced a process of co-design to evolve and transition codes into the next decade's exascale era. Broadly speaking, co-design is a process of end-to-end optimization—from a code's fundamental physics and numerical methods, through its algorithms and data structures, to the hardware and system software that the code will ultimately run on. This process of co-design is implemented through formation and management of collaborative, multi-disciplinary teams that include DOE scientists and representatives from hardware and software vendors. The team conducts its investigation through a tightly coupled cycle of application, algorithm and system software development, performance modeling, hardware simulation, and hardware design. Through the co-design process, optimal design tradeoffs are identified for hardware, the run-time environment, parallel file systems, physics algorithms and packages, and programming models. The interrelationships within the team are fluid and agile, reorganizing and realigning themselves as necessary throughout the lifetime of the effort.

Proxy Applications

A key focus area of co-design is the development and use of proxy applications (or proxy apps for short). Proxy apps are small, tractable vehicles for experimentation and collaboration that represent a simplification of characteristics of real applications (e.g., algorithms, data structures, memory layouts, parallelism, I/O models, programming languages, and coding styles) that are of interest to DOE. To enable rapid exploration and evaluation of these characteristics, proxy applications are much smaller and simpler to understand than the full applications. They are used in the co-design process as concrete examples for component and system designers to understand DOE's software requirements. ASC will make proxy apps generally available to the co-design community, something that has historically been difficult or impossible to do with our full applications. Openly available proxy apps, unlike benchmarks, are intended to facilitate the two-way communication required by co-design to optimally evaluate trade-offs in the system hardware, and inform the developer community of programming best practices for emerging architectures.

The ASC program, in collaboration with the DOE Office of Science Advanced Scientific Computing Research (ASCR) Program co-design center leadership, is standardizing on taxonomy[1] to differentiate different types of proxy apps, in order from simplest to most complex:

- **Kernels** are one or more small code fragments or data layouts that are used extensively by the applications and whose efficient operation are deemed essential to optimal performance on next generation advanced systems.
- **Skeleton apps** reproduce the memory or communication patterns of a physics application or package but make little or no attempt to investigate numerical performance.
- **Mini apps** are the most common form of proxy app and combine some or all of the dominant numerical kernels contained in an actual stand-alone application but models physical phenomena less realistically.

Each of the NNSA tri-labs has at its disposal a number of proxy applications[2,6,7,8] that have either been produced over the years as part of the ASC machine procurement process or are being newly developed as co-design interfaces for hardware and software vendors. NNSA proxy apps will be used for design, evaluation, and acceptance of the upcoming Trinity procurement, the next ATS to be installed at LANL by the LANL/Sandia Alliance for Computing at Extreme Scale (ACES) partnership. Additionally, many of the DOE FastForward vendors are using collections of proxy apps representing each of the six NNSA laboratories and ASCR co-design centers. FastForward is an exploratory effort funding vendors to develop processor, memory, and storage technologies of interest to DOE. The goal of FastForward is to accelerate development of technologies important to exascale computing. These common proxy apps are already being used by vendors, and facilitate apples-to-apples comparison of results and approaches.

Proxy apps have already proven valuable well beyond their originally intended purpose as co-design tools. These "sanctioned" proxy applications facilitate low-overhead research and development in runtime systems, languages, programming models, algorithms, tools, file systems, visualization techniques, etc. Using these proxy applications can help researchers across the HPC community target their work to areas that are known to be of interest to DOE. Similarly, proxy apps facilitate DOE's student programs (such as LANL's Co-Design Summer School[3]) by providing a small, focused problem domain that can produce useful results in a short period of time.

ASC and DOE Advanced Scientific Computing Research Partnership

ASC co-design is progressing in tight collaboration with the DOE ASCR Program, which is funding exascale research including three co-design centers[4] and multiple XStack projects[5]. It is not unusual to find DOE national laboratory researchers working in both ASC and ASCR projects. This deep penetration of co-design, and the use of proxy applications within DOE will facilitate cross-pollination and transition of technologies from research to production.

Conclusion

We are in the midst of a sea change in computer architectures. Complex, accelerated, massively parallel machines are currently being installed for production use. These architectural trends will continue into the exascale era and the foreseeable future. The ASC program is facing the largest overhaul of its billion-dollar code base in its history. The process of co-design, and our internal and external collaborative relationships, will enable us to influence and optimize the design of the hardware systems and software tools and start overhauling our code base much earlier in the procurement process.

Acknowledgements

Michael Heroux (SNL) and Sriram Swaminayran (LANL) contributed to this article.

References

[1]Mike Heroux, Rob Neely, Sriram Swaminarayan, "ASC Co-Design Proxy App Strategy, proxyapps.lanl.gov.

[2]ASC proxy applications are cataloged at proxyapps.lanl.gov.

[3]Los Alamos IS&T Co-Design Summer School, codesign.lanl.gov/summer-school.

[4]The ASCR co-design centers are at science.energy.gov/ascr/research/scidac/co-design/.

[5]The ASCR XStack projects , www.xstack.org.

[6]Los Alamos co-design web site, codesign.lanl.gov.

[7]Lawrence Livermore co-design web site , codesign.llnl.gov.

[8]Sandia Mantevo proxy application site , mantevo.org.

Next-generation 6.5 Megavolt Laser-triggered Gas Switch for Higher Voltage Operation of Z

Researchers at Sandia National Laboratories (SNL) have developed a next-generation laser-triggered gas switch for the Z accelerator (see Figure 1). Z uses 36 such switches, which are the final command-triggered elements of the accelerator and hence determine the performance of each Z shot. The baseline switch presently used on Z includes a plate that provides mechanical support for the switch's cascade section. The new switch uses instead a cantilevered cascade section, which for electrical reasons facilitates the self-break operation of the cascade electrodes. The new design reduces substantially the switch prefire rate, flashover rate, and timing jitter, while simultaneously allowing an increase of the Z Marx charge voltage from 85 kV to 95 kV. At 95 kV, the peak voltage across each switch will be 6.5 megavolts.

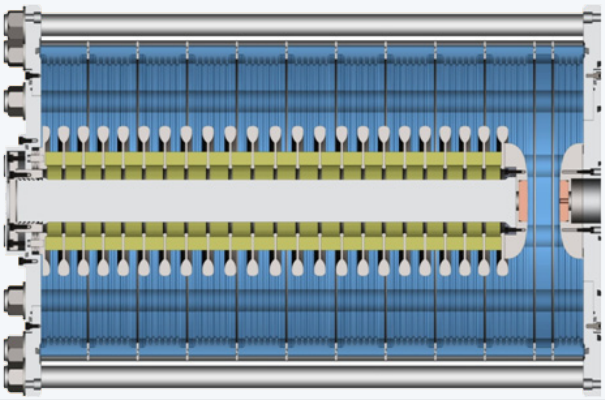

Figure 1. Next-generation 6.5-megavolt laser-triggered gas switch for the Z accelerator.

As of July 12, we have upgraded each of Z's 36 modules to the next-generation laser-triggered gas switch. On Z Shot 2538, which was conducted on July 24, the timing spread of Z's 36 pulse-forming lines was 15 ns, which is the best ever achieved on the refurbished Z, and matches the best performance on the original Z accelerator. To date, we have tested the new switch on 884 switch shots. One of the 36 new switches has survived over 106 shots on Z and is still working flawlessly, which suggests that the new design can support the desired switch lifetime (which is at least 100 shots). Over the course of these shots, we have observed no prefires, no flashovers, and a jitter that is a factor of two less than that of the baseline switch. We have also observed that the new switch performs well over much larger pressure and voltage ranges than the baseline switch, which will improve substantially the accuracy with which we can predict the current pulse generated by a Z shot, especially one that generates a shaped current pulse. The switch is the first 6.5-MV laser-triggered gas switch developed by the international pulsed-power community. The switch upgrade, combined with upgrades to Z's pulse-forming lines, insulator stack, and the magnetically insulated transmission line-convolute system, will allow Z to operate at 95 kV. The higher-voltage operation will increase substantially material pressures, x-ray yields, and fusion-neutron yields achieved on Z.

— R. Focia, T. Avila, M. Jones, P. Jones, and A. Kipp (SNL), K. LeChien (NNSA), M. Lopez, R. McKee, S. Ploor, J. Porter, J. Potter, M. Savage, J. Schwarz, D. Spencer, W. Stygar, and P. Wakeland (SNL), and S. White (LMATA Government Services)