NA-ASC-500-13 Issue 26

ASC eNews Quarterly Newsletter December 2013

In this issue:

The Meisner Minute

Bob Meisner

Greetings once again from our nation's capital. Considering you will be reading this after the holidays, I hope that you and your families had a safe and personally rewarding holiday. A lot has happened since my last *Meisner Minute* when we bid farewell to Reeta Garber, the "mother" of this newsletter, who was dedicated to better communicating the widespread work and accomplishments of the ASC Program. So, to start, I apologize for turning this minute into three minutes.

As the year came to a close, I note that not much has changed administratively for the program. As you know, the President signed the Budget Agreement into law the day after Christmas. This agreement all but ensures that there will not be another government shutdown. But the two-year budget agreement is not an appropriation, so we must await final appropriations for both FY14 and FY15. The President also signed the National Defense Authorization Act into law, which directs NNSA to develop and carry out a plan to incorporate exascale computing in the Stockpile Stewardship Program. This is not an exascale program, but recognition that stockpile stewardship requires exascale computing to be successful, and that it must be explicitly planned for.

Speaking of exascale, Secretary Moniz continues his high-visibility engagement. Last September, ASC and the Department of Energy (DOE) Office of Science labs jointly briefed him on the need for, and status of, the exascale effort. Secretary Moniz has asked his re-established Secretary of Energy Advisory Board (SEAB) to look into the activity and provide him with a recommendation on future steps. Since that time, the SEAB has (1) convened a fact-finding meeting organized around seven panels in early November 2013, including one focused on national security applications; (2) heard classified mission briefings by the defense laboratories; and (3) recommended that the Secretary appoint a task force to delve deeper into the DOE (and national) need for exascale computing. In January, Secretary Moniz acted positively on the recommendation with direction for the task force to complete their work by June of 2014. You can expect that members of the SEAB will engage many of you in additional fact-finding efforts, but know that their mission will carry them broadly into the larger computing community to both those providing and consuming services. They have set for themselves an ambitious task in a short time, and their recommendations will have important consequences for the nation. Please give them your best if called upon.

Any day, you should expect to hear who won the Trinity/NERSC-8 procurement. The New Mexico Alliance for Computing at Extreme Scales (ACES) team, comprised of staff from Los Alamos National Laboratory and Sandia National Laboratories, is doing an excellent job of bringing a world class machine to bear on stockpile stewardship problems. As I have mentioned in other forums, we should not expect this machine to emerge ranked number one on the TOP500 list, but we should expect it to outperform any other supercomputer on the planet for our applications. I will also be interested in Trinity's performance results on the evolving High Performance Conjugate Gradient (HPCG) benchmark.

During the past quarter, we received approval to proceed with preparing and releasing a Request for Proposals for the Advanced Technology System, known as Sierra, replacing Sequoia in 2017. Because of the successes you have achieved in partnering with the DOE Office of Science, we have decided to increase the complexity, and the potential pay-offs, in the Sierra procurement. To wit, we have commissioned a three-way partnership know as CORAL—Collaboration of Oak Ridge National Laboratory, Argonne National Laboratory, and Lawrence Livermore National Laboratory —to deliver three systems from, ideally, two vendor partners.

I hope to report next quarter that personnel churn in program leadership at HQ has settled down. In the meantime, people continue to move. With Dr. Christopher Deeney's acceptance of a position as the Chief Technology Officer at the Nevada National Security Site (NNSS), the duty of acting Assistant Deputy Administrator for Research, Development, Test, and Evaluation (NA-11) is being shared by the NA-11 office directors. LTC Michael Severson, the ASC Verification and Validation (V&V) Subprogram Manager, has moved up to the role of Deputy Assistant Deputy Administrator for NA-11. Until a new federal lead is assigned, Doug Wade and Jay

Edgeworth will jointly oversee the ASC V&V and Integrated Codes Subprograms. Jason Pruet has left the ASC Program, but not the building, to pursue career broadening opportunities in global security. In his place we welcome Adam Boyd who will do "double duty" as a program manager in both the Science Campaigns and the ASC Physics and Engineering Models (PEM) Subprogram. Dan Orlikowski has also taken on additional duties in PEM until his detail comes to an end this spring. Finally, I want to take this opportunity to welcome Keith Matzen as the Sandia ASC Exec, replacing Wendy Cieslak who retired in August 2013.

As 2013 has come to an end, I reminisce and observe that your accomplishments are the pride of the program, a few of which I have enumerated below.

- Eliminated historic discrepancies between simulated and measured yield through high-fidelity simulations of a W78 underground test with modern codes; thus increasing our confidence level for the annual assessment report.
- Advanced a computationally efficient laser weld modeling technique through better simulations of bending and shear loading.
- Revealed underground test metric impacts may be larger than expected by studying neutron down-scattering reactions in stockpile applications.
- Produced a user-guide of reaction history data to facilitate comparisons between DSW simulation results and underground test (UGT) associated data.
- Improved quantification of margins and uncertainties (QMU) process, enabling simulation-based safety assessments with multiple abnormal thermal failure modes; applicable to the B61 Life Extension Program (LEP) and other stockpile systems.
- Analyzed atomic simulations of aging metal tritides, studying mechanisms accommodating helium bubble growth degrading tritium storage materials.
- Created two new capabilities for the B61: simulating structural response of a weapon to acoustic signals and designing acoustic signals to create specific vibration responses.
- Supported 2013 Annual Assessment Review, uncertainty quantification (UQ) showed W78-0 high/low temperature test model is adequate simulator for current activities.
- Updated classified equation of state (EOS) tables for the weapons design community, including major improvements in accuracy, thermodynamic consistency, and range.
- Showed significant improvement in pin fits when phenomena not typically captured with most high explosive burn models was empirically included in B61 models.
- Transitioned Sequoia to the classified environment focused on strengthening the foundations of predictive simulation by running UQ studies and building more accurate physical models by running high-fidelity weapons science calculations.
- Pushed tri-lab Linux capacity clusters (TLCC2) into the "green operating zone" meeting high demand for DSW workload.
- Continued collaboration with DOE Office of Science/NNSA by selecting and funding five DesignForward interconnect projects.

- Completed early acquisition phases of NNSA's first two Advanced Technology Systems (ATS) for delivery in 2015 and 2017. Trinity, the first system, is in vendor negotiations, while Sierra is preparing to release a request for proposals.
- Led community engagement in defining the next generation, international benchmark for high-performance computing with potential to replace LINPACK.

Let me close with two simple notes. First, thanks for distinguishing another year of exceptional support in making the nation's nuclear weapons policies reality. And, last, but certainly not least, I wish you and your families all the best for a healthy and prosperous new year.

Return to top

LANL ASC Program High-Performance Computing Update

Trinity/NERSC-8 Request for Proposal Approved!

The Request for Proposal (RFP) was released to high-performance computing vendors on August 6, 2013, for two advanced technology supercomputers for the US Department of Energy (DOE). The RFP is the formal announcement seeking bidders for the National Nuclear Security Administration Trinity system, to be operated by the New Mexico Alliance for Computing at Extreme Scale (ACES), a collaboration of Sandia and Los Alamos national laboratories, and NERSC-8, a DOE Office of Science supercomputer to be operated by the National Energy Research Scientific Computing (NERSC) Center at Lawrence Berkeley National Laboratory. After evaluating the proposals, LANL will award a subcontract for Trinity by the end of calendar year 2013. It is projected that the Trinity platform will be installed at Los Alamos by September 2015.

This is the first time a joint RFP has been jointly developed under the sponsorship of the DOE Office of Science (NERSC) and the NNSA ASC (ACES-LANL and SNL) Programs. The approval and release of the Trinity/NERSC-8 RFP is an important milestone in the Trinity acquisition process. This comes after more than a year spent by the Trinity and NERSC teams developing and improving the technical requirements, and after a rigorous design review and independent project review of the technical

ASC Strategy Adopts Two Computing Platform Classes

According to the [ASC Program's *Computing Strategy*] document published in May 2013, the new ASC platform acquisition plan includes two computing platform classes:

- Commodity Technology (CT) systems, and
- Advanced Technology (AT) systems

The CT systems provide computing power to a large percentage of the design and analysis community by leveraging predominately commodity hardware and software. In contrast, the AT systems are the vanguards of the HPC platform market and incorporate features that, if successful, will become future commodity technologies.

and project requirements.

Trinity will be the first NNSA advanced technology (AT) system. (See sidebar for details.) Trinity is specified to address weapons science simulations of large-scale Directed Stockpile Work (DSW) calculations in support of the Predictive Capability Framework and mission needs, while at the same time taking a significant step forward in computing platform technology.

Cielo Update: Application and System Improvements Resulting from Cielo's Third Campaign. The Cielo supercomputer is working full time on Stockpile Stewardship Program (SSP) problems from all three labs: LANL, LLNL, and SNL. The problems are typically 3D, and span the entire range of SSP simulations.

Among the Cielo Capability Computing Campaign 3 (CCC3) projects was a Los Alamos Directed Stockpile Work (DSW) calculation that put exceptional stress on the memory and I/O systems of Cielo. This simulation was run at 4096 nodes, or about half of Cielo's 8894 nodes. Each node has 16 cores, so the simulation was using 65,536 cores.

Once runs began, memory exhaustion problems soon followed. Code teams scrambled to conserve memory for particular observed failure cases. These efforts eliminated the first memory exhaustion symptoms, but the simulation continued to hit memory limits. Debugging these issues was difficult, in part because the traditional debugging tools did not scale to the needed size of 4096 nodes.

Simulation Too Large for Cielo. Eventually we realized that the desired DSW simulation was just too large to run on Cielo; it will have to wait for the Trinity system. The code team was then tasked to understand why the simulation was too large for the system. Three efforts began in parallel. The first effort examined the CCC3 usage model, which answered the question about what is the minimum fraction of the machine needed to run this problem. The second effort examined the memory use of the application to get enough head room to run the entire simulation. Finally, we put controls on the compute node write rate to reduce the number of nodes needed for this large simulation.

In summary, this exercise demonstrates that we have DSW problems too big for Cielo, and that it continues to be very difficult to debug large jobs on large machines (for example, we need better tools). We will use the lessons learned from these efforts to help define the next-generation platforms and the tools we need to accomplish running these large DSW simulations.

Return to top

SIERRA Delivers New Integrated Solid Mechanics/Structural Dynamics Capabilities and Prepares for Next Generation Computing

Sandia SIERRA developers successfully completed a FY13 milestone that greatly improved the interoperability between the Solid Mechanics (SM) and Structural Dynamics (SD) modules that included enabling both modules to directly access the same numerical model data. This milestone was driven by a continuing pull to add needs driven functionality to our analysis tools to support nuclear weapon LEPs, ALTs, Annual Assessments and stockpile maintenance issues, and to prepare our codes for next generation computers. In particular, there are many high priority uses that require greatly improved cross functionality between the modules, such as, nonlinear modeling capabilities (e.g., large scale joint slippage) in traditionally linear shock and vibration applications; efficient linear capabilities (e.g., use of super elements for fixed linear components) in traditionally nonlinear solid mechanics applications; enhanced solution transfer capabilities between modules; eigenvalue analysis capabilities on a SM model; and common contact mechanics capabilities for SM and SD models, e.g., the ability to automatically determine and enforce tied interfaces within a multi-part model. All of these foundational cross-functional capabilities were enabled by this milestone activity.

In addition, this milestone positioned both the SM and SD modules to be able to adapt to next generation computer platforms by connecting them to the new SIERRA Toolkit components, such as the mesh and solutions fields data representations. These Toolkit components are being designed to be able to adapt to next generation computers by building in the agility to deal with the now unknown architectures for these platforms. The use of Toolkit components is a major part of the SIERRA system transition strategy for MPI+X parallel computing needs.

One final major accomplishment of this milestone was that of placing the SM and SD modules under a common interface, the Sandia CompSimUI. This user interface provides a unifying and common user experience. The significance of the unification goes beyond just having a common interface, as the CompSimUI plays a key role for the inclusion of "Intrinsic V&V" concepts, including a natural way to define and execute simulations to establish uncertainty bounds on simulation results.

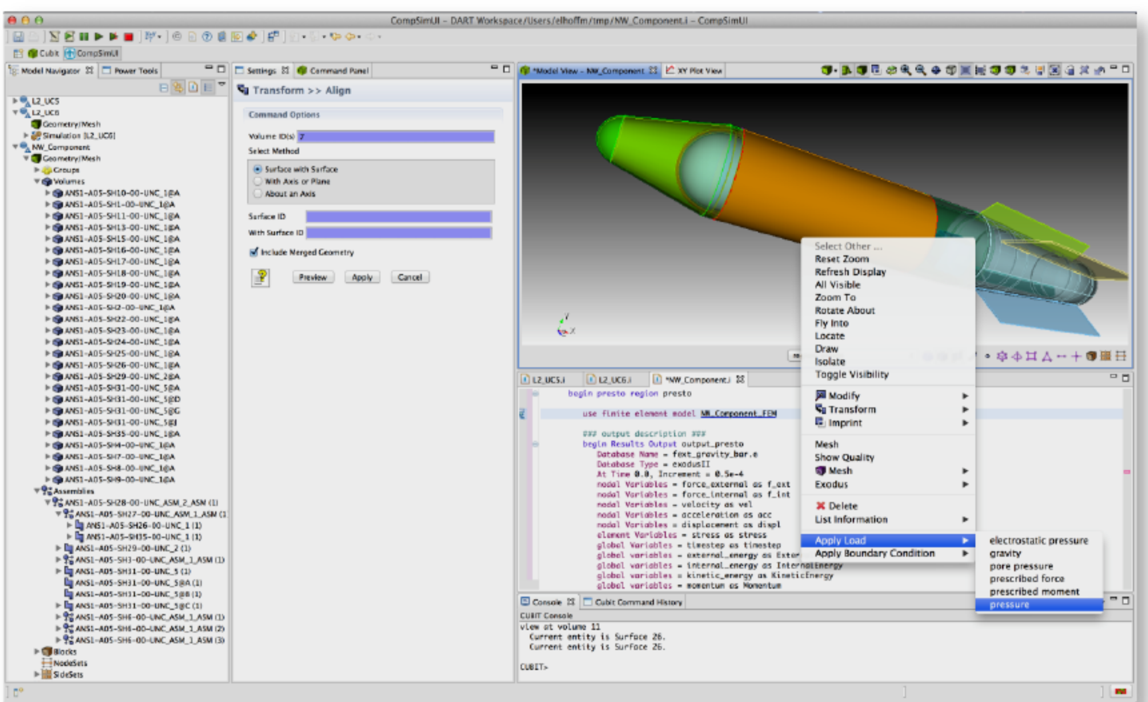

The CompSimUI working with a SM/SD model and SIERRA input file.

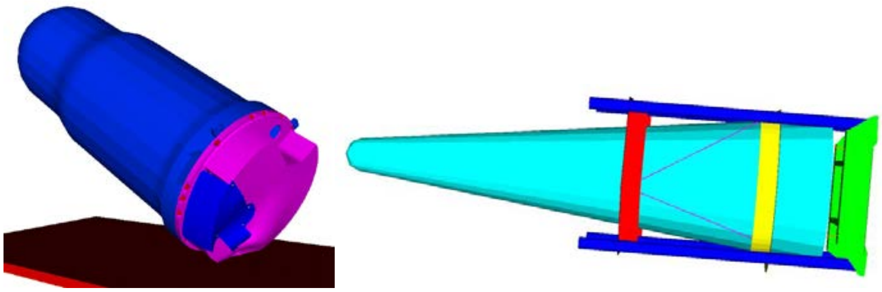

New SIERRA/SM-SD capabilities enable Nonlinear Preloads followed by Shock and Vibration Analyses, and eigenvalue analysis of SM models.

Return to top

Sub-System Level Validation of Detached Eddy Simulation Model Accomplished in SIERRA

The spin-up of stockpile gravity bombs that occurs during the firing of Spin Rocket Motors (SRM) is a critical portion of the flight. During SRM firing, a large region of vortical flow is generated leading to interaction between the vortex and the fin. Simulating the vortex-fin interaction is an important element in predicting spin-up and quantifying performance margins for such systems. Previous work[1] has shown that the standard two-equation Reynolds-Averaged Navier-Stokes (RANS) turbulence models do not possess sufficient fidelity to accurately predict the vortex-fin interaction produced by the SRM. An unsteady turbulence model[2] that contains improved physics has been implemented in Sandia's SIERRA gas dynamics code. This validation effort attempts to quantify the capability of new unsteady turbulence models to predict vortex-fin interaction produced by SRMs using sub-system (jet-in-crossflow).

The configuration studied here corresponds to that studied experimentally by Beresh, et al.[3,4,5]. The setup consists of a supersonic jet issuing from a conical nozzle on the floor of a 12" x 12" test section in a blow-down wind tunnel. The tunnel Mach number is 0.8. The nozzle has an exit diameter of 0.375" and a design Mach number of 3.73. In the experiments, the nozzle orientation is varied from vertical to cant angles up to 45°. Here we only use the 0° cant angle data. PIV measurements of the velocity field (both mean and turbulent fields) are available at a streamwise location that is 33.8 jet diameters downstream of the jet and on the spanwise symmetry plane (for the 0° cant angles). The flow conditions studied here correspond to the test conditions in Beresh, et al.[4] with primary focus restricted to the Mach 0.8 and jet-momentum-ratio of 10.2. Figure 1 illustrates a sample snapshot of a prediction at Mach 0.8 and jet-momentum-ratio of 10.2.

In carrying out the validation, care was taken to ensure that the appropriate boundary conditions were prescribed. This was done by carrying out a precursor calculation of the TWT geometry and by matching the measured boundary layer profile and pressure gradient profile. In addition, accuracy was ensured of the nozzle boundary conditions by comparing the pressure expansion in the nozzle to the measured values.

The meshes used in this work spanned a range from the coarsest mesh of 8.2 million nodes to the finest mesh obtained by uniformly refining the coarse mesh by a factor of 2 in each dimension, resulting in a fine mesh of about 65 million nodes. These meshes were generated using industry standard practices – near wall resolution such that y+=1.0 and jet shear layer resolution such that ~50 points were present across the jet in the coarse mesh. Mesh refinement studies on these meshes were carried out along with time step refinement and CFL refinement. In addition, the flow field was also computed using a different mesh that had a significantly different topology and provided much better refinement of the jet shear layer. In addition calculations were repeated using a different flow solver using the same model and using a different unsteady turbulence model.

The essential conclusions of the effort can be summarized as follows:

1. Mesh refinement study showed that errors on coarse mesh are smaller than errors on fine Mesh (see Figure 2);
2. Overall predicted flowfield characteristics showed strong anomalies;
3. The time step refinement Study showed that errors remain constant with temporal refinement;
4. Detailed analysis of the model showed that the model has a strong dependence on function that determines RANS vs. LES regions of the flow.

Based on these results, it was concluded that the SST-DES model in its existing form could not reliably predict jet-in-crossflow flow field. A strong dependence on the wall distance function needs to be better understood and calibrated or alternate methods to eliminate its unpredictable effects need to be devised. During the validation effort, a substantial of amount of knowledge was gained in understanding the numeric and physical mechanism in the subsystem problem. This knowledge has led to a Physics & Engineering Models (PEM) project to improve model adequacy.

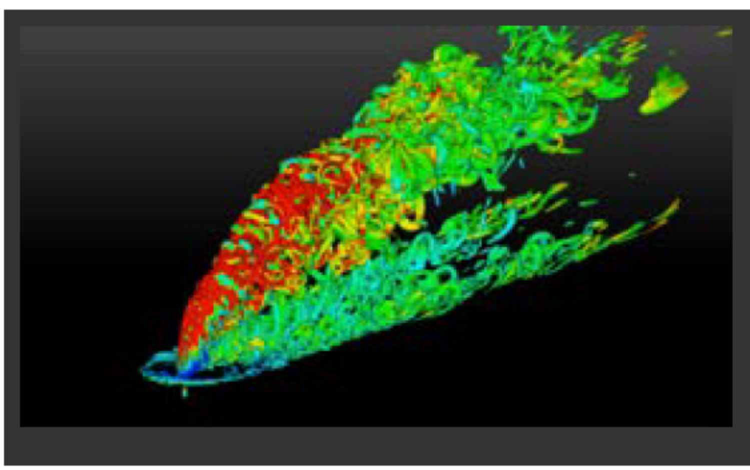

Figure 1: Details of the Turbulent Flow Field Captured Using DES

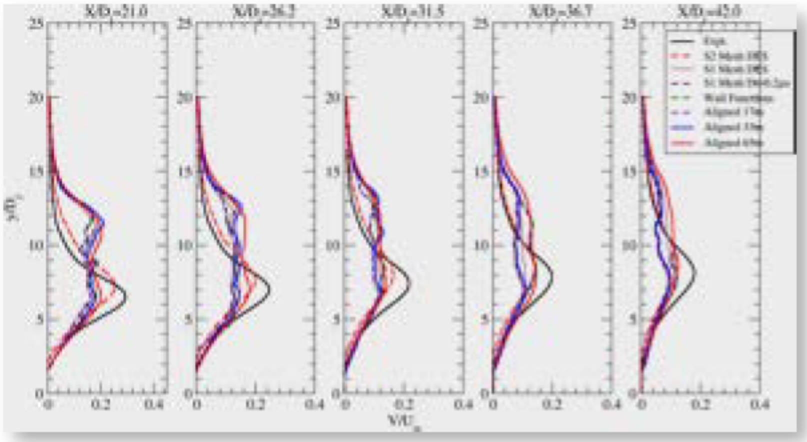

Figure 2: Comparisons of various time steps and grid levels with experimental data.

References

1. Arunajatesan, S., "Evaluation of Two-Equation RANS Models for Simulation of Jet-in-Crossflow Problems", AIAA-2012-1199.
2. Strelets, M., "Detached Eddy Simulation of Massively Separated Flows", AIAA-2001-0879.
3. Beresh, S. J., Henfling, J. F., Erven, R. J., and Spillers, R. W., "Crossplane Velocimetry of a Transverse Supersonic Jet in a Transonic Crossflow," AIAA Journal, Vol. 44, No. 12, 2006, pp. 3051–3061.
4. Beresh, S. J., Henfling, J. F., Erven, R. J., and Spillers, R. W., "Penetration of a Transverse Supersonic Jet into a Subsonic Compressible Crossflow," AIAA Journal, Vol. 43, No. 2, 2005, pp. 379–389.
5. Beresh, S. J., Henfling, J. F., Erven, R. J., and Spillers, R. W., "Vortex Structure Produced by a Laterally Inclined Supersonic Jet in Transonic Crossflow", Journal of Propulsion and Power, Vol. 23, No. 2 2007, pp 353-363.

Return to top

Update on LANL's Visualization and Data Analysis Project

Visualization and data analysis are key capabilities in taming and understanding the ever-increasing large data sets generated from extreme-scale scientific simulations. In the last fiscal year (FY13), LANL's ASC Visualization and Data Analysis Project has: (1) achieved a 7.8-billion-cell simulation on the Cielo supercomputer; (2) successfully completed "In situ Analysis for ASC Simulations," an L-2 milestone; and (3) designed and deployed an upgrade to the Powerwall visualization theater.

7.8-Billion-Cell EnSight Visualization on Cielo. To perform the largest simulation so far on the Cielo supercomputer, researchers at LANL used the visualization nodes on Cielo in combination with the ViewMaster 2 visualization cluster at LANL and the EnSight 10 visualization software to create 3D animations of vortex tube evolution in a 7.8-billion-cell adaptive mesh refinement (AMR) simulation of the transition to turbulence.

Fluid dynamic, large eddy simulations of the transition from coherent vortical structures to fully turbulent flow typically require extremely high spatial resolutions in 3D in order to achieve a sufficiently large, effective Reynold's number to observe the development of the turbulence directly. LANL researchers used 64,000 processors on the ASC Cielo machine to perform the 3D ultra-high resolution AMR simulation of the transition to turbulence. This simulation, shown at the right, reveals in unprecedented detail how fully turbulent flows arise from the instability of coherent vortices in the fluid. The simulation shows a 3D inertial confinement fusion (ICF) implosion done with the RAGE code. The view shows a closeup of the vortex tubes near the polar axis of the capsule as they twist around one another and reconnect to create the turbulence.

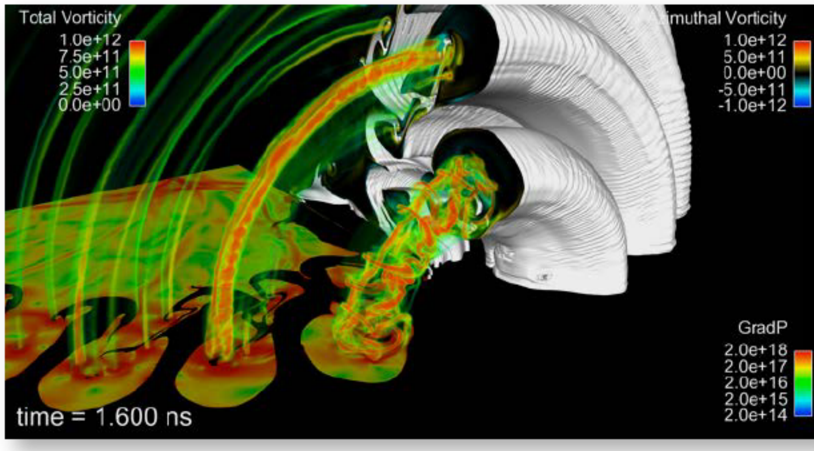

In order to visualize this very large AMR dataset, the full infrastucture of petascale visualization for Cielo was utilized. The EnSight 10 visualization software was run in a distributed mode with an EnSight rendering client running on LANL's ViewMaster2 visualization cluster connected over the internal 10 GigE network to 200 EnSight servers running remotely on 50 viz nodes of the Cielo platform. The 200 EnSight server processes running in parallel read data from the simulation stored on Cielo’s Lustre filesystem and extract visible geometry, which is transmitted over the internal 10 GigE network to the EnSight rendering client on ViewMaster2. The rendering client produces stereo images that are then interactively shipped over the video connections of the LANL Visualization Corridor directly to the researcher's office for 3D viewing. Using this infrastructure, LANL researchers have made 3D stereo animations of the development of the vortex tubes in the transition to turbulence in this revealing 3D simulation.

Rich New Set of Capabilities Offered for Users of ASC Integrated Codes. LANL completed an ASC CSSE Level-2 milestone by demonstrating *in situ* data analysis and visualization in ASC integrated code xRage (see illustration). By transforming simulation data at run time, in situ analysis and visualization provides new options for exploring trade-offs in the design of extreme-scale systems, and simultaneously opens up a rich new set of scripted and automated analytical capabilities for the end users of ASC integrated codes. Possible benefits include savings in disk space, computing time, and analysis time, with options for saving analysis products more densely, e.g., with greater temporal resolution. In situ techniques are expected to play an increasingly important role for future ASC systems, which will be challenged by storage and data movement constraints. The members of the Data Science at Scale team are working with designers in LANL’s Theoretical Design Division to integrate these new capabilities into their modeling work flow.

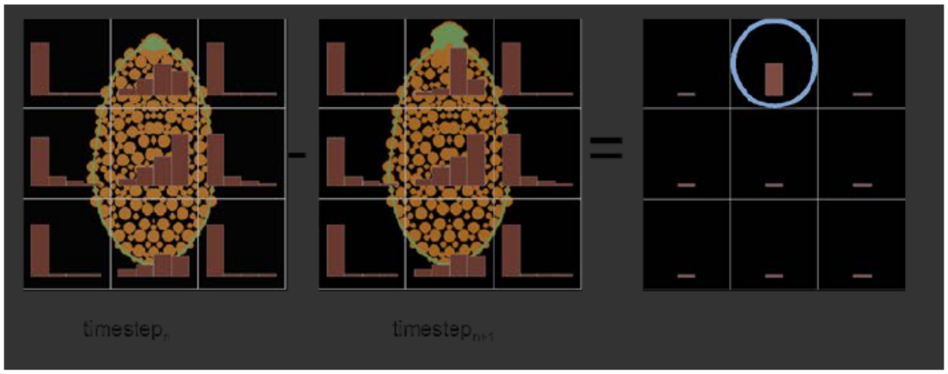

Illustration of in situ analysis with automated camera placement allows a simulation to generate a visualization that tracks the most interesting active areas of physics at run time. The physical variables are spatially decomposed and histogram time differences are used to identify currently active regions.

Upgrade to Powerwall Theater in the Metropolis Center for Modeling & Simulation. The Powerwall Theater (PWT) in the Metropolis Center for Modeling and Simulation at LANL has undergone an extensive upgrade. The PWT is an innovative facility that enables researchers to view models and simulations they have created using some of the world's fastest supercomputers. The new theater features a total of 40 Christie Digital Systems Mirage WU-L 3D LED projectors in a 5-wide by 4-high configuration with double-stacked projectors in each node onto a new ultra-high-fidelity glass by Stewart Filmscreen Corporation. The facility includes automated blending and alignment and continuous color and brightness matching. The technologies combined create seamless, integrated visualization to serve for scientific work, presentations, and meetings.

This system will provide the seamless appearance and the multi-window functionality typical for a single-projection solution (as in movie theaters and single monitors). At the same time, it has the high resolution required for ASC visualization that cannot be reached with a single projector.

To see a press release on the LANL PWT, go to http://www.prweb.com/releases/2013/7/prweb10956400.htm

Return to top

Lawrence Livermore's Vulcan Brings 5 petaFLOP/s Computing Power to Collaborations with Industry and Academia

The Vulcan supercomputer at Lawrence Livermore National Laboratory is now available for collaborative work with industry and research universities to advance science and accelerate the technological innovation at the heart of U.S. economic competitiveness.

A 5-petaFLOP/s (quadrillion floating point operations per second) IBM Blue Gene/Q system, Vulcan will serve Lab-industry projects through Livermore's High Performance Computing (HPC) Innovation Center as well as academic collaborations in support of DOE/National Nuclear Security Administration (NNSA) missions. The availability of Vulcan effectively raises the amount of computing at LLNL available for external collaborations by an order of magnitude.

LLNL's Fred Streitz and Doug East, in front of the Vulcan supercomputer. (Photo by Laura Schulz and Meg Epperly/LLNL.)

"High performance computing is a key to accelerating the technological innovation that underpins U.S. economic vitality and global competitiveness," said Fred Streitz, HPC Innovation Center director. "Vulcan offers a level of computing that is transformational, enabling the design and execution of studies that were previously impossible, opening opportunities for new scientific discoveries and breakthrough results for American industries."

Six recently concluded industrial collaboration projects from Livermore Lab's initiative—called the hpc4energy incubator—illustrate benefits companies have realized through the application of supercomputer technologies and expertise to energy applications. See the Web [http://hpc4energy.org/incubator/] for details. Beyond these examples the availability of Vulcan enables even larger systems to be simulated over longer time periods with greater fidelity and resolution.

In addition to publishable incubator programs, the HPC Innovation Center provides on-demand, proprietary access to Vulcan and connects companies and academic collaborators with LLNL's computational scientists and engineers, and computer scientists to help solve high-impact problems across a broad range of scientific, technological and business fields. As projects are initiated, the HPC Innovation Center rapidly assembles teams of experts and identifies the computer systems needed to develop and deploy transformative solutions for sponsoring companies. The HPCIC draws on LLNL's decades of investment in supercomputers, HPC ecosystems, and expertise, as well as the technology and know-how of HPC industry partners. Companies interested in access to Vulcan, along with other supercomputing services are invited to respond to the posted Notice of Opportunity.

Vulcan also will serve as an HPC resource for LLNL's Grand Challenge program and other collaborations involving the Laboratory's Multi-Programmatic and Institutional Computing effort. Such collaborations with academic and research institutions serve to advance science in fields of interest to DOE/NNSA including security, energy, bioscience, atmospheric science and next generation HPC technology.

On the June 2013 industry-standard Top500 list of the world's fastest supercomputers, Vulcan ranked as the world's eighth-fastest HPC system. This makes it one of the world's most powerful computing resources available for collaborative projects.

During its shakeout period, Vulcan was combined with the larger Sequoia system, producing some breakthrough computations, notably setting a world speed record of 504 billion events per second for a discrete event simulation—a collaboration with the Rensselaer Polytechnic Institute (RPI). This achievement opens the way for the scientific exploration of complex, planetary-sized systems. See the Web [https://www.llnl.gov/news/newsreleases/2013/Apr/NR-13-04-06.html]

Housed in LLNL's high performance computing facility, Vulcan consists of 24 racks, 24,576 compute nodes and 393,216 compute cores. For addition specs, see the Web [https://computing.llnl.gov/?set=resources&page=OCF_resources#vulcan]

Return to top

ZFS: Targeted File System Development for Sequoia and Beyond

Why ZFS for Sequoia? ZFS was chosen as the best solution for an alternate backend for the large Lustre file system on Sequoia because it affordably offered scalability, manageability, data integrity, and other features (for example, the ability to repair and manage the file system without taking it down and the emphasis on data integrity are both winning attributes). Note that according to Wikipedia, ZFS originally stood for zettabtye file system; a zettabyte is equal to 1 billion terabytes.

The Grove high-speed, high-capacity parallel file system at Lawrence Livermore.

"Putting ZFS in allows us to scale to the needs of Sequoia and beyond and puts an emphasis on data integrity and availability," said Mark Gary, deputy division leader for Livermore Computing. "The work done for ZFS allows for the insertion of other file systems under the covers in Lustre. We've defined a canonical layer under which ZFS, the existing ldiskfs, or other file systems can be inserted."

ZFS is a new approach to data management. It uses the concept of storage pools to manage the storage of data so that file systems are no longer constrained to individual devices. Because there was no ZFS implementation for Linux, LLNL began developing the ZFS port for Linux in 2008. The team continues the targeted development necessary for ZFS to perform well in LLNL's HPC environment while leveraging development by the global community; however, not all developments are deliberate. Recently, it was decided to turn on ZFS' compression feature. Surprisingly, 1.7-to-1 compression was observed. Complaints about somewhat slow data transfer rates were silenced when users realized that their files were actually almost twice the size they believed them to be.

So for the uninitiated, what is the significance of ZFS in real-world comparisons?

The ZFS/Lustre file system for Sequoia has a 55-PB capacity—that's 55 million billion characters worth of information. If that much data was written in a book, the book would be more than five times the thickness of the distance from the earth to the moon.

Grove (collectively, Sequoia's 23,040-disk ZFS file system and its Lustre software) has been observed transferring data at 850 GB/s, the equivalent of writing over 100,000 CDs every second. As noted by Mark, "We buy file systems for their speed. A computer is just a big, hot doorstop while it is waiting for data to transfer."

The ZFS team works closely with the global Lustre community, and ZFS development work directly benefits LLNL as well as the global community while also supporting core laboratory missions and delivering dependable computing resources to users.

Return to top

Twenty-five Reasons to Love MCNP6

Los Alamos released the Monte Carlo N-Particle Transport Code, MCNP6, and the ENDF/B-VII.1-based nuclear data cross-section library, for use with MCNP6. It is available to users through the Radiation Safety Information Computational Center (RSICC) at Oak Ridge National Laboratory — one of the world's leading resources for nuclear computational tools and services. MCNP6 is used for a wide variety of applications, including Stockpile Stewardship, threat reduction, criticality safety, nuclear nonproliferation, safeguards, space radiation effects, medical and health physics, radiation shielding and measurement, radiation effects from an urban nuclear weapon detonation, fossil fuel exploration, and more. MCNP6 Version 1.0 was used for detailed analysis of diagnostic ports of the International Thermonuclear Experimental Reactor model.

MCNP is the generally acknowledged "gold-standard," and was requested by 586 people in the first week of its availability — an impressive demand for the code. The code can be used on small desktop PCs to national lab HPCs, using both MPI and shared-memory threading parallelism.

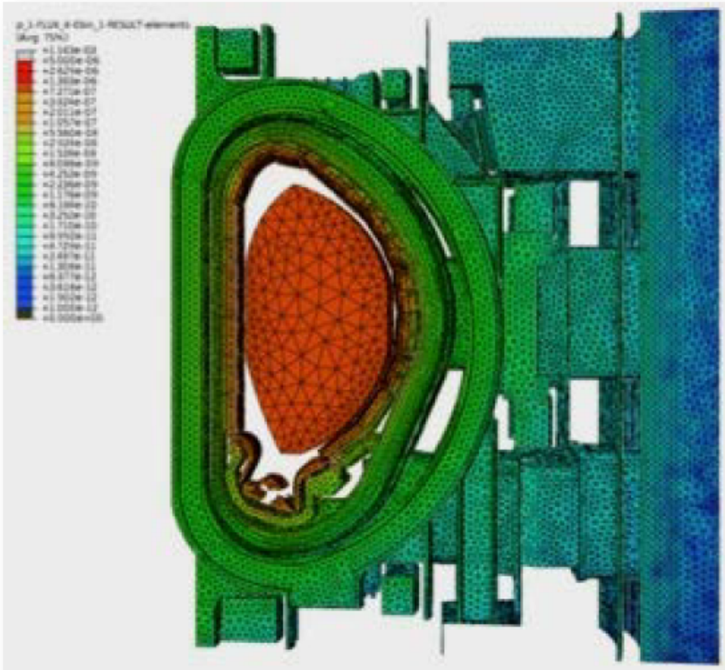

International Thermonuclear Experimental Reactor (ITER) model (20-degree section used for detailed analysis of diagnostic ports) calculation with MCNP6 Version 1.0. The color scale shows neutron flux intensity.

More than the merger of two radiation transport codes, MCNP5 and MCNPX, the MCNP6 release contains more than 25 new features. These features include, but are not limited to:

- Transport calculations on unstructured mesh geometries from and feedback into the engineering program Abaqus;
- Generation of structured mesh geometries from standard MCNP input files, for development of hybrid deterministic and Monte Carlo methods, calculation of sensitivity profiles of Keff with continuous-energy physics for cross sections, fission multiplicities and spectra, for quantifying nuclear data uncertainties on criticality;
- Tracking particles through materials, such as air, when magnetic fields are present;
- New source and tally options (background, beta-decay, cosmic rays, spontaneous fission/decay) for detector modeling and homeland security applications;
- Continuous $S(\alpha, \beta)$ data and methods, for improved modeling of low-energy neutron interactions with different materials;
- Improved delayed particle emission from fission and radioactive decay, and
- Transport of 34 discrete radiation particle types plus arbitrary heavy-ion transport (for example, neutrons, electrons, photons, muons, protons).

Return to top

RAMSES Preparing to Deploy Improved Solver Capabilities for Higher Frequency Electromagnetics

A new algorithmic capability has been implemented and tested in EIGER (an ASC component of the RAMSES code suite). EIGER solves Maxwell's equations using an integral equation formulation that results in having to solve a matrix equation where the matrix is dense. The storage of this matrix causes the method to be memory limited and thereby limits the size of problems (with respect to the frequency of excitation) that can be solved. To help extend the usefulness of the EIGER technique a matrix compression algorithm has been implemented, which alleviates the matrix storage issue. The matrix compression procedure is an algebraic method that computes low-rank approximations of sub-blocks of the original matrix. In addition, it performs this "on the fly" resulting in an efficient method both in storage and numerical complexity.

Two very diverse test problems have been calculated with the new algorithm and compared to solutions based on a standard direct matrix solve. The first problem was a conceptual aircraft known as the VFY 218. It was solved at an excitation frequency of 1. GHz and resulted in a matrix order of approximately 1 million. The matrix was compressed by 97%, using a large compression tolerance. The results compared well with the direct solve solution. The second problem considered an object with a cavity interior that was connected via thin slots. The thin-slot algorithm is a sub-cell algorithm designed specifically for geometries encountered at Sandia. For this problem the matrix was compressed by 39% (set by the compression tolerance) and the results also compared well with the direct solution method. Future improvements will concentrate on implementing improved matrix pre-conditioners to speed up run-time.

As a result of these efforts, the new algorithm implemented in EIGER will allow for the computation of a large variety of complex geometries at frequency ranges higher than previously achieved. These problems are relevant to the stockpile modernization program, in particular with respect to electromagnetic reliability and interference effects.

Return to top

CORAL Collaboration Ushers in the Future of DOE Computing

A memorandum of understanding (MOU) between the Office of Science (SC) and NNSA's Office of Defense Programs coordinating exascale activities in the two organizations focuses on driving U.S. scientific discovery and economic competitiveness by enabling high performance scientific computing on a new generation of computers.

The new collaboration among SC's Oak Ridge National Laboratory (ORNL) and Argonne National Laboratory (ANL) and NNSA's Lawrence Livermore National Laboratory (LLNL), referred to as CORAL, will accelerate key R&D for DOE applications and acquire advanced computing resources for delivery in 2017 to meet the mission needs of SC's Advanced Scientific

Computing Research (ASCR) Program and NNSA's ASC Program. Other initiatives under this MOU include FastForward and DesignForward.

From one request for proposals (RFP) issued by LLNL, CORAL will choose two different computing system architectures and procure a total of three systems, with unique architectures for ANL and ORNL and one of the two architectures for LLNL. The CORAL Request for Information (RFI) was issued on Dec. 13, 2012. The CORAL team has identified several key requirements for systems procured under the CORAL RFP, including a 4*x* to 12*x* improvement on specific benchmarks that represent important DOE applications as compared to Sequoia, Titan (ORNL), or Mira (ANL). The RFP will emphasize memory capacity, memory bandwidth, and operating characteristics such as power draw and resiliency.

Once vendors are selected, the multiyear lab/awardee goal is to co-design the computers and have both R&D contracts jointly managed by the three labs. Each lab will manage and negotiate its own computer procurement contract and may exercise options to meet its specific needs. This methodology was designed to promote a rich HPC ecosystem in which different applications may favor different architectures while also mitigating risk. Mission requirements must be met even in the face of delays or failure of one system or a particular vendor's product.

"CORAL is as much about a long-term partnership with an HPC vendor as it is a procurement for leadership computers," said Bronis de Supinski, LLNL/ICC's chief technology officer and representative on CORAL. "We expect this vendor partnership to be a five-year development and production lifespan. Risk reduction and risk mitigation will be built into the overall procurement strategy from the start, reflecting lessons learned from earlier large procurements as well as understanding new challenges unique to current technologies."

LLNL's machine from the CORAL procurement is currently being referred to as the SIERRA Advanced Technology System. As described in the *SIERRA Advanced Technology System Critical Decision 0* document, "requirements for the SIERRA AT system, as with all previous large ASC platform acquisitions, present significant challenges in terms of functionality, scalability, and performance. Challenges include technical and schedule delivery issues and will require detailed contract and project management for successful execution.

Lessons learned from previous ASC platform acquisitions are factored into the planning for these types of large systems. While careful planning is essential to meeting requirements, unforeseen events and changes are likely to occur, based on previous experiences. These events can only be successfully addressed by a strong partnership that goes beyond ordinary vendor-customer relationships. It must be a partnership where teaming, mutual respect, and an honest desire to achieve success are present on the part of all parties involved.

"We also intend to jointly fund nonrecurring engineering work to incentivize the vendor to accommodate particular and specific needs of the Stockpile Stewardship Program early enough in the design of the solution to impact the outcome significantly," added Bronis. "An example of this would be improvements in compiler technology to support threading or an alternative memory hierarchy. More than two full years of lead-time are necessary for maximum impact."

A second collaboration that includes Los Alamos, Sandia, and Lawrence Berkeley laboratories was also formed. Collaboration groups were based on common acquisition timing to reduce the number of RFPs to which vendors need respond and the number of reviews the labs need to perform. In addition, R&D funds will be pooled.

For more information, see the CORAL Vendor Meeting Web page.

Return to top

PIPER Lays the Groundwork for Performance Tools on Exascale Systems

The Performance Insight for Programmers and Exascale Runtimes (PIPER) project, led by Lawrence Livermore Computer Scientist Martin Schulz, is a recent Office of Science X-Stack program initiative to address key challenges to creating an exascale software ecosystem.

PIPER is a multi-institutional collaboration that includes Peer-Timo Bremer and Todd Gamblin as LLNL co-principal investigators as well as researchers from Pacific Northwest National Laboratory, Rice University, and the Universities of Maryland, Utah, and Wisconsin. Together, the PIPER team will develop new techniques for measuring, analyzing, attributing, and presenting performance data on exascale systems.

Exascale architectures and applications will be much more complex than today's systems, and to achieve high performance, radical changes will be required in high performance computing (HPC) applications and in the system software stack. In such a variable environment, performance tools are essential to enable users to optimize application and system code. Tools must provide online performance feedback to runtime systems to guide online adaptation, and they must output intuitive summaries and visualizations to help developers identify performance problems.

The figure highlights the team’s overall approach and illustrates how performance data will be collected and analyzed at different levels of the exascale software stack. The output of the performance analysis is either displayed intuitively to users using visualization tools, or it is fed back to the exascale runtime for online optimization.

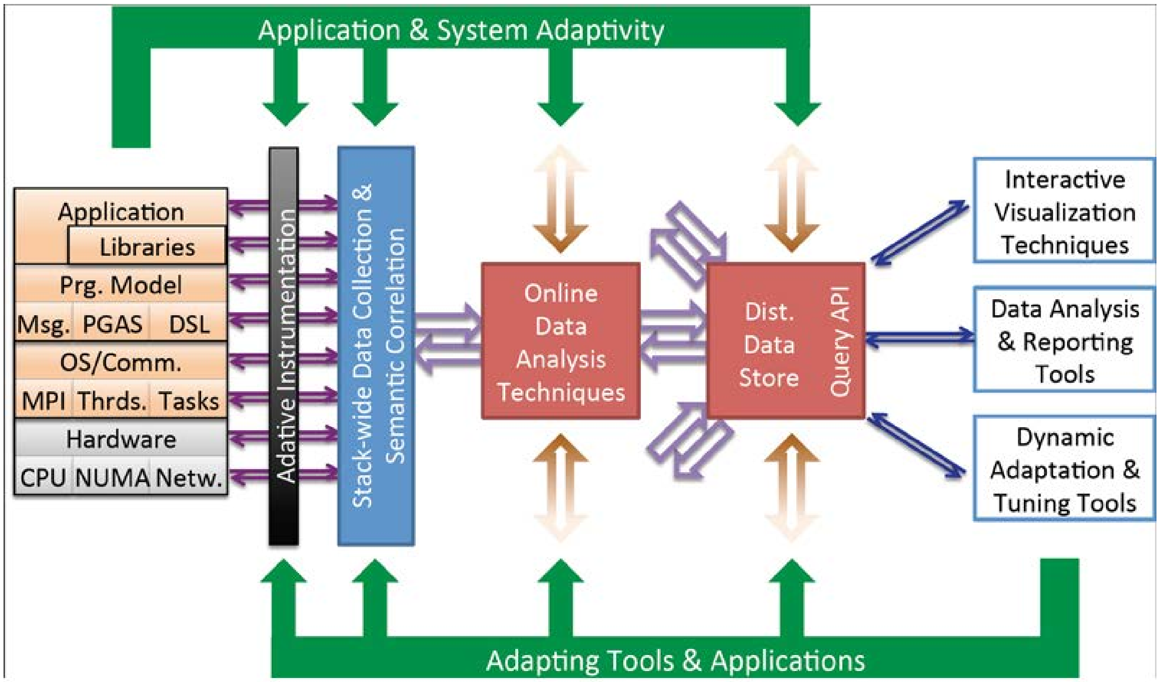

The software architecture of PIPER.

The PIPER team will work closely with other X-Stack, co-design, and exascale operating system/runtime efforts to lay the foundation for online performance measurement and analysis across the exascale software stack. Insights from PIPER, coupled with the adaptive runtime infrastructures built into related projects, will enable users to turn the predicted increases of raw system performance in future machines into tangible results, including mission critical advances and new scientific discoveries.

Return to top

Advances in Next Generation of Portals Networking Software

Scalable Interconnects is an important technology for future HPC platforms that is not being adequately addressed by industry. A team from Sandia National Laboratories has been working on extending the state of the art in high performance networking. Work on version 4 of the Portals networking specification (the previous generation was deployed with the Cray XT SeaStar interconnect) has been advancing along with a software reference implementation. The reference software implementation has been extended to work with multiple communication methodologies including MPI, GasNet (Global address space Network), and SHMEM (Shared MEMory). Improved support includes functionality over traditional Ethernet networks to facilitate development of protocols on top of Portals. This new network transport support also enables more rapid development of next generation operating systems that use Portals networks, such as Kitten, another Sandia-led project. In addition, shared memory support has been added to the baseline InfiniBand-based Portals reference implementation, to better represent how a Portals hardware implementation would behave.

Work continues into the use of operations that can be triggered on the networking hardware by lightweight counting events, and how these can be leveraged for high performance non-blocking collective operations. Non-blocking collective operations are a new feature of the recently approved MPI-3.0 standard, and Portals is well positioned to efficiently support them. Good performance of such operations is key to enabling the ability of ASC to deliver the computing capabilities needed to meet its national security goals in the future.

The advanced systems technology testbed, Teller, is being used for Portals Reference Implementation Development.

Return to top

Mini-Sequoia Code Support Deployed at Sandia

Sandia has obtained a small (half-rack) proxy computer for Sequoia users to enhance the Integrated Codes team efforts of code conversion and optimization. Sandia maintains our code production facilities in the unclassified restricted network (Sandia Restricted Network – SRN), which presents a challenge for the developers in optimizing for the Sequoia ASC supercomputer system since resides in the classified network at Lawrence Livermore National Lab (LLNL). The SIERRA and RAMSES code groups are composed of many rapidly changing elements that must be checked for quality and accuracy on a nightly basis as updates to the codes are checked in daily. Running up to 20,000 nightly regression tests on a remote platform located on a different network would be highly inefficient requiring uploads of many megabytes of code and transfers of files between sites. Therefore, deploying a mini-Sequoia platform at Sandia on the same network as the code repositories was identified as a requirement to ensure code compatibility and adherence to the production quality practices of the Integrated Code

The mini-Sequoia system. (Photo: Courtesy of Steve Monk, Sandia)

Teams.

Initial code porting to Sequoia identified issues with the IBM compiler linker sequence and performance limitations that will be addressed more quickly with a local, unclassified resource available for extensive debugging and optimization activities. Sandia will provide an access path for administrators at LLNL who will ensure the Sequoia software environment is replicated onto the Sandia system in a timely manner. The Intersite HPC network will facilitate this interaction and access while providing appropriate controls for privileged accounts that are required for system updates and maintenance.

The combined efforts of Sandia and LLNL in supporting a consistent up to date software environment for Sandia code development, and the flexibility of a local resource will ensure the highest levels of efficiency for Sandia developers preparing codes to run on Sequoia. This in turn will ensure that Sandia analysts will be able to take full advantage of the Sequoia CCC allocations of time for vital NW and ASC computational simulation needs In FY14 and beyond.

Return to top

News from ISC'13; Sequoia Remains at the Top of the Graph 500 List

Among the news from the 2013 International Supercomputing Conference (ISC'13) in Leipzig, Germany, was that the Sequoia system had retained its No. 1 ranking on the Graph 500 list. Bronis de Supinski, Livermore's chief technology officer, accepted the No. 1 Graph 500 certificate on behalf of the lab.

The Graph 500 benchmark measures the speed with which a system can "connect the dots" within a massive set of data. Sequoia was able to traverse 15,363 giga edges per second on a scale of 40 graph (a graph with 2^40 vertices).

"The Graph 500 provides an additional measure of supercomputing performance, a benchmark of growing importance to the high performance computing community," said Jim Brase, deputy associate director in Computation. "Sequoia's top Graph 500 ranking reflects the IBM Blue Gene/Q system's capabilities. Using this extraordinary platform, Livermore and IBM computer scientists are pushing the boundaries of the data-intensive computing critical to our national security missions."

Bronis de Supinski (center) accepts the No. 1 Graph 500 certificate at the 2013 International Supercomputing Conference in Leipzig, Germany, on June 19.

Extracting essential nuggets of information, sometimes called data mining, and recognizing patterns in large data sets is important to such Livermore missions as nonproliferation, atmospheric monitoring, intelligence, bioinformatics, and energy. For example, computer scientists are developing techniques for organizing data that allow nonproliferation analysts to pinpoint anomalous behavior in a huge quantity of data.

Several Livermore scientists were on hand at ISC’13 to present recent research and participate in international collaborative discussions.

At a session on memory system design, Lawrence Scholar and Texas A&M Ph.D. student Roger Pearce presented innovative graph traversal algorithms they developed.

Bronis de Supinki gave a half-day Advanced OpenMP tutorial that covered many OpenMP 4.0 topics. Bronis is the chair of the OpenMP Language Committee and has led the effort to produce a new specification that will be adopted in July. Bronis also participated in a breakfast panel focused on Lustre.

In addition to a full-day tutorial on advanced parallel programming with MPI, Martin Schulz participated in a half-day tutorial on supporting performance analysis and optimization on extreme-scale computer systems. Martin also co-organized a Birds-of-a-Feather session and gave an invited talk on MPI, during which he discussed the current state of MPI 3.0 as well as expectations for future versions.

Also during the conference, the Department of Energy and the Japanese Ministry of Education, Culture, Sports, Science, and Technology (MEXT) held a joint workshop on a new U.S.–Japan collaboration on Trans-Petascale System Software. Together with Koji Inoue from Kyushu

University, Martin Schulz gave a talk on a collaborative project on power management and optimization.

Additionally, LLNL attendees participated in an invitation-only strategy workshop of the Virtual Institute for High Productivity System (VI-HPS), which Livermore joined last year. VI-HPS is an association of tool developers or tool-developing research groups from around the world who share a goal to promote tools, provide training, and work toward integrating the individual tools.

Return to top

Exploring the "Prius Effect" Energy Performance Dashboard for Supercomputing

A paper co-authored by Anna Maria Bailey, manager of Livermore's high performance computing facilities, and presented at the annual High Performance Power Aware Computing (HPPAC) Workshop highlighted LLNL's leadership role in driving energy efficiency measures for supercomputing.

The paper, "Recommendations for HPC Energy Management Dashboard Displays," explores how data and metrics should be displayed to provide the "Prius effect" for supercomputing. The Prius effect refers to the use of dashboards to provide real-time, actionable visual data. By observing the fuel consumption screen, Prius drivers begin to see the direct connection between the way they drive and the rate of fuel consumption. In turn, they become more aware of their behavior and are motivated to drive in a way that keeps fuel consumption low. Is there an equivalent metric that HPC facility managers and others like Anna Maria should use?

The paper was motivated by a survey taken in early 2012 of DOE HPC centers about their use of dashboards for driving energy efficiency. Only a few DOE labs reported having energy performance dashboards (LLNL, Lawrence Berkeley National Laboratory, and the National Renewable Energy Laboratory), and even those were described as partial, piecemeal, and under construction. Anna Maria's paper makes recommendations for what data should be included in an HPC center energy performance dashboard.

Anna Maria co-authored the paper as a member of the Energy Efficient High Performance Computing Working Group "Dashboard Team." The team suggested that facility managers choose from dashboard metrics like power usage effectiveness (PUE) and cooling efficiency (kW/ton). The team also recommended dashboard metrics for information technology (IT) managers. A suggestion that generated significant interest at the HPPAC workshop was an IT efficiency metric (work output/watt), where work output depends on how each HPC center defines its work output.

"It's important to remember that dashboards do not make decisions—facility managers do," Anna Maria said. "A dashboard is only one tool, or in some cases, an integrated set of tools for helping to measure, analyze, and manage how HPC energy is consumed. We're still far from identifying a Prius effect, but we're making steady progress in optimizing HPC energy consumption. LLNL is recognized as a leader in that effort for the HPC community."

The HPPAC Workshop was held May 20 in Boston as part of the International Distributed and Parallel Processing Symposium and was co-chaired by Computation's Bronis de Supinski and Dong Li of Oak Ridge National Laboratory. Topics ranged from more efficient networking chip design to virtualization techniques for cloud computing.

Return to top

Livermore Scientists Win Best Paper Award at International Conference on Supercomputing

A paper on massively parallel loading earned Livermore Scientists Dong Ahn, Matt LeGendre, Todd Gamblin, Bronis de Supinski, and external collaborators Felix Wolf from the German Research School for Simulation Sciences (part of RWTH Aachen University) and Wolfgang Frings from the Jülich Supercomputing Centre the Best Paper Award at the International Conference on Supercomputing, held in Eugene, Ore. in June.

The winning paper addressed the problem of dynamic loading on massively parallel systems, which—if carried out in an uncoordinated manner—can lead to file-system access storms that manifest themselves much like a denial-of-service attack, seriously degrading performance and scalability. The proposed solution efficiently coordinates the associated parallel file system operations with a scalable network of cache server processes.

Return to top

Los Alamos Celebrates Its 70th Anniversary with the Second Los Alamos Primer

Throughout 2013, LANL will be celebrating its 70th anniversary of service to the nation. In June, as part of the celebration, LANL conducted a lecture series — the Second Los Alamos Primer — to enact a modern-day version of the "Primer Lectures" originally convened in 1943. The first "Los Alamos Primer" lecture series took place in conjunction with the start-up of "Project Y," a part of the Manhattan Project. Project Y would eventually become Los Alamos National Laboratory.

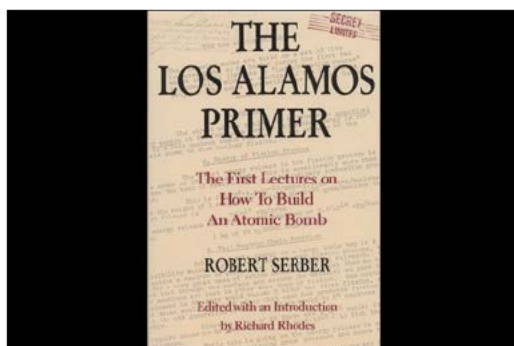

At the startup of Project Y, the United States was faced with the carnage of World War II. The brightest students from the nation's best colleges and universities were recruited for the project and soon became the world's first nuclear weapons scientists and engineers. Director Oppenheimer tasked his Berkeley protégé Robert Serber to begin the Los Alamos lectures by presenting an essential introductory overview of the relevant nuclear physics for making an atomic bomb. Following each day's lecture, Serber's original notes were expanded and annotated, based on the questions and discussions among audience participants. Formulas, graphs, and simple drawings from the blackboard were added. The resulting document, titled the

Los Alamos Primer, was mimeographed and handed out to every newly arriving Project Y scientist.

Lectures and discussions at the Second Primer event were aimed to explore the changing stewardship landscape, seek new ways to meet its challenges, celebrate LANL's successes, and inspire the current and next generation of scientists. The 2-day event featured guests from the national security community and speakers from laboratory staff and managers. The lectures examined the geostrategic, military, academic, diplomatic, and political environment confronting Los Alamos. A few photos from the event are shown here.

(from left) Mary Hockaday, John Sarrao, Alan Bishop, Gary Grider, and Bill Archer discuss "Decadal Challenges for Science and Engineering." [Photo by Richard Robinson, LANL.]

The Weapons Designers Roundtable was an overview of the technical issues confronting experienced designers and the challenges of training emerging designers during the moratorium on underground nuclear testing. From left: Robert Webster, Brian Lansrud-Lopez, Langdon Bennett, John Scott, James Mercer-Smith, John Pedicini, and Gary Wall. [Photo by Richard Robinson, LANL.]

The Laboratory Directors' Roundtable was moderated by Principal Associate Director Bret Knapp (center). From left: Former Laboratory Directors Pete Nanos, Robert Kuckuck, and Michael Anastasio, and current Director Charlie McMillan. [Photo by Richard Robinson, LANL.]

Return to top

R&D 100 Award for Mantevo "Mini-App" Suite: "Mini-Apps" Play a Crucial Role in the Future of High Performance Computing

The Mantevo 1.0 suite of "mini-apps" has received a 2013 R&D 100 award. This suite, which is named for the Greek word *mantevo* meaning predict, has been a pioneering effort to create proxies for full-scale applications, which are often extremely large and complex, making them difficult to deploy to new and exploratory high performance computing architectures. By representing crucial aspects of the larger application in a much smaller, open source "mini-app", researchers and industry can rapidly evaluate these new architectures and gain insight to improve both the hardware and software. With the surging interest in "co-design", tightly integrated design of hardware and software, by the high performance community, these "mini-apps" have become crucial tools for communication between the national laboratories, academia, and the computing industry. They are being rapidly adopted and used by such companies as AMD, IBM, Intel and Nvidia to tune their experimental computing architectures for optimal performance on real world applications. This Sandia National Laboratories-led project has been a joint effort with Lawrence Livermore National Laboratory, Los Alamos National Laboratory, United Kingdom collaborators at the Atomic Weapon Establishment and NVIDIA Corporation.

Mantevo website: http://mantevo.org

Youtube video:
http://www.youtube.com/watch?v=j6INXyg5jK0&list=PLA89E1F3535785901&index=16

Return to top

ASC Salutes Maya Gokhale

Thriving on Change

Computer scientist Maya Gokhale appreciates the unpredictability and rapid pace of change in her chosen field. "You never know where computing is going to go, and that's what's exciting about it," she says. "It's harder to look forward five years in computer science than it is in other technical fields."

Since joining LLNL's Center for Applied Scientific Computing in 2007, Maya's research has centered on developing innovative high-performance computing (HPC) system architectures that effectively store, retrieve, and use the reams of data that modern scientific research produces. For one such effort, Maya and her team of researchers have been working to augment the memory of a computer server or compute node with large, parallel arrays of solid-state storage devices. Data stored in these nonvolatile memory storage arrays are both permanent and close to the compute node, allowing for fast access and manipulation. The viability of this novel architecture was proven in the June 2011 international Graph 500 competition, when Maya and her team achieved rank 7 using a single compute node with nonvolatile memory. A machine near the top of the Graph 500 list can efficiently analyze vast quantities of data to find the hidden gems of useful information. Not only has the nonvolatile memory research supported data-intensive computing in fields as diverse as social network analysis and bioinformatics, it has also boosted exascale computing preparation (exascale supercomputers that will be able to perform at least one quintillion operations a second—at least 100 times what today's machines can do).

While the nonvolatile memory effort involves adding permanent memory to compute nodes, another of Maya's data-science projects is its perfect complement, as it entails adding compute functionality to the memory or storage system. Latency is the delay that occurs as a packet of data travels between two points—for example, from memory to the CPU. By enabling the computer to perform low-level computations within the memory itself, rather than moving the data to the compute node before performing the operation, Maya and her colleagues are alleviating some of the latency problem. "This is an incredible way of increasing memory bandwidth, as the internal bandwidth in the memory is quite large," she says. Maya considers in-memory computing to be one of the most exciting and potentially transformative research innovations underway today in computer architecture.

In 2013, Maya was named a distinguished member of Livermore's technical staff, an honor only bestowed on a tiny fraction of researchers. She has previously worked in industry and academia, but one reason she favors working at a national laboratory is because her colleagues here help her identify and test the real-world applications for her ideas. She also appreciates the chance to make direct national security contributions through her work at LLNL. The daughter of two educators, Maya was the first in her family to pursue a career in science and technology. After having endured years of technical dinner-table conversations with Maya and her computer engineer husband, her son and daughter swore off computer science as a career. But now, both use computers extensively in their scientific disciplines. Says Maya, "It was quite a challenge for my husband and me to juggle careers and kids, and we are grateful to have wonderful children."

I also like doing backbends and headstands in yoga. You can look at the world from a different perspective!

Return to top

ASC Relevant Research (Publications)

Los Alamos National Laboratory — Citations for Publications (not previously listed)

Akkan, H., Lang, M., Liebrock, L. (2013). "Understanding and isolating the noise in the Linux kernel," *International Journal of High Performance Computing Applications,* Vol. 27, No. 2, pp. 136-146. DOI:10.1177/1094342013477892.

Bazilevs, Y., Akkerman, I., Benson, D.J., Scovazzi, G., Shashkov, M.J. (2013). "Isogeometric analysis of Lagrangian hydrodynamics," *Journal of Computational Physics,* Vol. 243, pp. 224-243. DOI:10.1016/j.jcp.2013.02.021.

Bochev, P., Ridzal, D., Shashkov, M. (2013). "Fast optimization-based conservative remap of scalar fields through aggregate mass transfer," *Journal of Computational Physics,* Vol. 246, pp. 37-57. DOI:10.1016/j.jcp.2013.03.040.

Ciappina, M.F., Pindzola, M.S., Colgan, J. (2013). "Fully differential cross section for O8+-impact ionization of Li," *Physical Review A,* Vol. 87, No. 4, pp. 042706. DOI:10.1103/PhysRevA.87.042706.

Colgan, J., Kilcrease, D.P., Magee, N.H., Armstrong, G.S.J., Abdallah, J., Sherrill, M.E., Fontes, C.J., Zhang, H.L., Hakel, P. (2013). "Light element opacities from ATOMIC," *High Energy Density Physics,* Vol. 9, No. 2, pp. 369-374. DOI:10.1016/j.hedp.2013.03.001.

Csanak, G., Fontes, C.J., Inal, M.K., Kilcrease, D.P. (2013). "Creation, destruction, and transfer of atomic multipole moments by electron scattering: Liouville-space formulation," *Journal of*

Physics B-Atomic Molecular and Optical Physics, Vol. 46, No. 8, pp. 085202. DOI:10.1088/0953-4075/46/8/085202.

Ding, X., Lookman, T., Zhao, Z., Saxena, A., Sun, J., Salje, E.K.H. (2013). "Dynamically strained ferroelastics: Statistical behavior in elastic and plastic regimes," *Physical Review B,* Vol. 87, No. 9, pp. 094109. DOI:10.1103/PhysRevB.87.094109.

Ferenbaugh, C.R. (2013). "A comparison of GPU strategies for unstructured mesh physics," *Concurrency and Computation-Practice & Experience,* Vol. 25, No. 11, pp. 1547-1558. DOI:10.1002/cpe.2894.

Ferouani, A.K., Inal, M.K., Csanak, G. (2013). "Ne IX line G-ratio in a non-Maxwellian and anisotropic plasma," *Journal of Physics B-Atomic Molecular and Optical Physics,* Vol. 46, No. 8, pp. 085701. DOI:10.1088/0953-4075/46/8/085701.

Grabowski, P.E., Markmann, A., Morozov, I.V., Valuev, I.A., Fichtl, C.A., Richards, D.F., Batista, V.S., Graziani, F.R., Murillo, M.S. (2013). "Wave packet spreading and localization in electron-nuclear scattering," *Physical Review E,* Vol. 87, No. 6, pp. 063104. DOI:10.1103/PhysRevE.87.063104.

Gumberidze, A., Thorn, D.B., Fontes, C.J., Najjari, B., Zhang, H.L., Surzhykov, A., Voitkiv, A., Fritzsche, S., Banas, D., Beyer, H., Chen, W., DuBois, R.D., Geyer, S., Grisenti, R.E., Hagmann, S., Hegewald, M., Hess, S., Kozhuharov, C., Martin, R., Orban, I., Petridis, N., Reuschl, R., Simon, A., Spillmann, U., Trassinelli, M., Trotsenko, S., Weber, G., Winters, D.F.A., Winters, N., Yu, D., Stohlker, T. (2013). "Electron- and Proton-Impact Excitation of Hydrogenlike Uranium in Relativistic Collisions," *Physical Review Letters,* Vol. 110, No. 21, pp. 213201. DOI:10.1103/PhysRevLett.110.213201.

Haines, B.M., Grinstein, F.F., Schwarzkopf, J.D. (2013). "Reynolds-averaged Navier-Stokes initialization and benchmarking in shock-driven turbulent mixing," *Journal of Turbulence,* Vol. 14, No. 2, pp. 46-70. DOI:10.1080/14685248.2013.779380.

Harribey, T., Breil, J., Maire, P.H., Shashkov, M. (2013). "A swept-intersection-based remapping method in a ReALE framework," *International Journal for Numerical Methods in Fluids,* Vol. 72, No. 6, pp. 697-708. DOI:10.1002/fld.3763.

Hemez, F. (2012). "Scientific computing and model validation at Los Alamos," Proceedings of International Conference on Noise and Vibration Engineering, 2012. pp. 4807-4807.

Hill, R.N., Shashkov, M. (2013). "The Symmetric Moment-of-Fluid interface reconstruction algorithm," *Journal of Computational Physics,* Vol. 249, pp. 180-184. DOI:10.1016/j.jcp.2013.04.037.

Iaccarino, G., Sharp, D., Glimm, J. (2013). "Quantification of margins and uncertainties using multiple gates and conditional probabilities," *Reliability Engineering & System Safety,* Vol. 114, pp. 99-113. DOI:10.1016/j.ress.2012.11.026.

Ichikawa, T., Moller, P., Sierk, A.J. (2013). “Character and prevalence of third minima in actinide fission barriers,” *Physical Review C*, Vol. 87, No. 5, pp. 054326. DOI:10.1103/PhysRevC.87.054326.

Kulkarni, A., Ionkov, L., Lang, M., Lumsdaine, A. (2013). “Optimizing process creation and execution on multi-core architectures,” *International Journal of High Performance Computing Applications*, Vol. 27, No. 2, pp. 147-161. DOI:10.1177/1094342013481483.

Lipnikov, K., Reynolds, J., Nelson, E. (2013). “Mimetic discretization of two-dimensional magnetic diffusion equations,” *Journal of Computational Physics*, Vol. 247, pp. 1-16. DOI:10.1016/j.jcp.2013.03.050.

Ma, P.L., Gattiker, J.R., Liu, X.H., Rasch, P.J. (2013). “A novel approach for determining source-receptor relationships in model simulations: a case study of black carbon transport in northern hemisphere winter,” *Environmental Research Letters*, Vol. 8, No. 2, pp. 024042. DOI:10.1088/1748-9326/8/2/024042.

Mihaila, B., Stan, M., Crapps, J., Yun, D. (2013). “Impact of thermal conductivity models on the coupling of heat transport, oxygen diffusion, and deformation in (U, Pu)O2-x nuclear fuel elements,” *Journal of Nuclear Materials*, Vol. 433, No. 1-3, pp. 132-142. DOI:10.1016/j.jnucmat.2012.09.017.

Neudecker, D., Capote, R., Leeb, H. (2013). “Impact of model defect and experimental uncertainties on evaluated output,” *Nuclear Instruments & Methods in Physics Research Section A-Accelerators Spectrometers Detectors and Associated Equipment*, Vol. 723, pp. 163-172. DOI:10.1016/j.nima.2013.05.005.

Owhadi, H., Scovel, C., Sullivan, T.J., McKerns, M., Ortiz, M. (2013). “Optimal Uncertainty Quantification,” *SIAM Review*, Vol. 55, No. 2, pp. 271-345. DOI:10.1137/10080782x.

Rollin, B., Andrews, M.J. (2013). “On generating initial conditions for turbulence models: the case of Rayleigh-Taylor instability turbulent mixing,” Journal of Turbulence, Vol. 14, No. 3, pp. 77-106. DOI:10.1080/14685248.2013.790549.

Rosa, M., Warsa, J.S., Perks, M. (2013). “A Cellwise Block-Gauss-Seidel Iterative Method for Multigroup S-N Transport on a Hybrid Parallel Computer Architecture,” Nuclear Science and Engineering, Vol. 174, No. 3, pp. 209-226.

Rossi, M.L., Taylor, C.D. (2013). “Equations of state for crystalline zirconium iodide: The role of dispersion,” Journal of Nuclear Materials, Vol. 433, No. 1-3, pp. 30-36. DOI:10.1016/j.jnucmat.2012.08.020.

Rubinstein, R., Clark, T.T. (2013). “Reassessment of the classical closures for scalar turbulence,” Journal of Turbulence, Vol. 14, No. 2, pp. 71-98. DOI:10.1080/14685248.2013.769685.

Sambasivan, S.K., Shashkov, M.J., Burton, D.E. (2013). “A finite volume cell-centered Lagrangian hydrodynamics approach for solids in general unstructured grids,” International

Journal for Numerical Methods in Fluids, Vol. 72, No. 7, pp. 770-810. DOI:10.1002/fld.3770.

Sanati, M., Albers, R.C., Lookman, T., Saxena, A. (2013). “First-order versus second-order phase transformation in AuZn,” Physical Review B, Vol. 88, No. 2, pp. 024110. DOI:10.1103/PhysRevB.88.024110.

Schmidt, W., Collins, D.C., Kritsuk, A.G. (2013). “Local support against gravity in magnetoturbulent fluids,” Monthly Notices of the Royal Astronomical Society, Vol. 431, No. 4, pp. 3196-3215. DOI:10.1093/mnras/stt399.

Schoffler, M.S., Stuck, C., Waitz, M., Trinter, F., Jahnke, T., Lenz, U., Jones, M., Belkacem, A., Landers, A.L., Pindzola, M.S., Cocke, C.L., Colgan, J., Kheifets, A., Bray, I., Schmidt-Bocking, H., Dorner, R., Weber, T. (2013). “Ejection of Quasi-Free-Electron Pairs from the Helium-Atom Ground State by Single-Photon Absorption,” Physical Review Letters, Vol. 111, No. 1, pp. 013003. DOI:10.1103/PhysRevLett.111.013003.

Subramanian, G., Perez, D., Uberuaga, B.P., Tome, C.N., Voter, A.F. (2013). “Method to account for arbitrary strains in kinetic Monte Carlo simulations,” Physical Review B, Vol. 87, No. 14, pp. 144107. DOI:10.1103/PhysRevB.87.144107.

Zhang, H.L., Fontes, C.J. (2013). “Relativistic distorted-wave collision strengths for the 16 Delta n=0 optically allowed transitions with n=2 in the 67 Be-like ions with 26 <= Z <= 92,” Atomic Data and Nuclear Data Tables, Vol. 99, No. 4, pp. 416-430. DOI:10.1016/j.adt.2012.04.004.

Sandia National Laboratories — Citations for FY13, Q4

Key: DOI = Digital Object Identifier; URL prefix of DOI is: http://dx.doi.org/

Brown, A. L. (2013). “Predictive Impulse Dispersal of Liquid Employing a Code Coupling Methodology,” Proceedings of the Seventh International Seminar of Fire and Explosion Hazards, Providence, RI. **SAND2013-0363 C**.

Childs, J., Geveci, B., Schroeder, W., Meredith, J., Moreland, K., Sewell, C., Kuhlen, T., Bethel, E. W. (2013). “Research Challenges for Visualization Software,” Computer, Vol. 46, Issue 5, pp. 34-42. **SAND2012-8316 J**.

Hendry, G. (2013). “The Role of Optical Links in HPC System Interconnects,” 2013 IEEE Optical Interconnects Conference, Santa Fe, NM, pp. 5-6. **SAND2013-0252 C**.

Mandadapu, K. K., Jones, R. E., Zimmerman, J. A. (2013). “On the Microscope Definitions of the Dislocation Density Tensor,” Mathematics and Mechanics of Solids. Available online 13 June 2013. **SAND2012-10363 J**.

Noble, D. R. (2012). "A Conformal Decomposition Finite Element Method for Material Death," Proceedings of ASME 2012 Computational Heat Transfer Summer Conference, Rio Grande, Puerto Rico, USA. ASME Paper No. HT2012-58464, pp. 1059-1963. **SAND2012-0152 C**.

Pedretti, K., Vaughan, C., Barrett, R., Devine, K., Hemmert, K. S. (2013). "Using the Cray Gemini Performance Counters," Proceedings, Cray User Group Meeting, Napa, CA. **SAND2013-3482 C**.

Rajbhandary, A., Bunde, D. P., Leung, V. J. (2013). "Variations of Conservative to Improve Fairness," 17th Workshop on Job Scheduling Strategies for Parallel Processing (JSSPP), in Conjunction with 27th IEEE International Parallel & Distributed Processing Symposium (IPDPS 2013), Boston, MA. Electronic collection of presented papers were distributed at the workshop. **SAND2013-1734 C**.

Schultz, P. A. (2012). "Simple Intrinsic Defects in GaP: Numerical Predictions," Sandia Technical Report **SAND2012-3314**. Unlimited release.

Schultz, P. A. (2012). "Simple Intrinsic Defects in InP: Numerical Predictions," Sandia Technical Report **SAND2012-3313**. Unlimited release.

Schultz, P. A. (2013). "Simple Intrinsic Defects in InAs: Numerical Predictions," Sandia Technical Report **SAND2013-2477**. Unlimited release.

Tang, L., Hu, S., Chen, D. Z., Niemier, M., Barrett, R. F., Hammond, S. D., Hsieh, G. (2013). "GPU Acceleration of Data Assembly in Finite Element Methods and Its Energy Implications," 2013 IEEE 24th International Conference on Application-Specific Systems, Architectures and Processors (ASAP), Washington, DC, pp. 321-328. **SAND2013-3083 C**.

Corrections to Prior Submittals

Submitted FY13 Q1:

Piekos, E. S., Grillet, A. M., Ingersoll, D., Streeter, N. D., Noble, D.R. (2012). "Modeling of Thermal Battery Initiation Using Level Sets," Proceedings of the 45th Power Sources Conference, Las Vegas, NV. **SAND 2012-5184 C**. Available on CD ROM only from National Technical Information Service (NTIS) at http://www.ntis.gov.

Correction:

Piekos, E. S., Grillet, A. M., Ingersoll, D., Streeter, N. D., Noble, D.R. (2012). "Modeling of Thermal Battery Initiation Using Level Sets," Proceedings of the 45th Power Sources Conference, Las Vegas, NV. **SAND 2012-2725 C**. Unclassified title and document . Available on CD ROM only from National Technical Information Service (NTIS) at http://www.ntis.gov.

Return to top