NA-ASC-500-09 Issue 10
March 2009

The Meisner Minute

> … ASC now plays the central integrating role previously performed by nuclear tests, and is the only arena in which all aspects of the program are tested together.—**Defense Science Board Task Force on the National Nuclear Security Administration (NNSA) Strategic Plan for Advanced Computing, January 2, 2009**

Our newsletter this month focuses on the variety of contributions by ASC to national defense. This is a timely theme and it has been echoed in several external studies that have been completed in the past few months (including the one cited above). A recent demonstration of the central role of ASC is that the first refurbished W76 nuclear warhead has been accepted into the U.S. nuclear weapon stockpile by the Navy, culminating a ten-year effort that made extensive use of ASC-provided simulations and computing

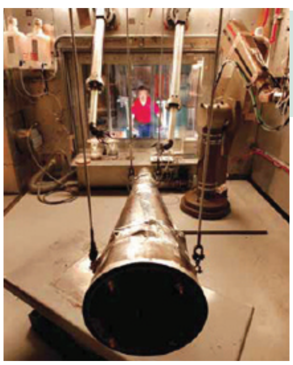

Hostile shock testing of the W76-1 at Sandia National Laboratories

http://nnsa.energy.gov/news/2286.htm

This is a major accomplishment for the weapons program and points to the critical importance of science and technology in national defense.

The ASC codes and capabilities are also used to inform decisions made by the plants on safety and efficiency of their processes. For example, the managers of the Pantex cells, where nuclear weapons are dismantled and examined, use ASC tools to determine safe levels for the material density in their cells as well as to examine safe ways to clamp the weapons for examination.

A strong science technology & engineering (ST&E) capability contributes to saving money in the operation of the plants by optimizing production and manufacturing processes. An example is using ASC simulation tools to reduce the costs of neutron generator tube production and qualification by reducing the number of costly neutron generator (NG) prototypes that are built en route to a final NG design.

Our contributions extend beyond the immediate benefits to the nuclear weapons program as acknowledged in the following excerpt:

The Department of Energy's laboratory system provides invaluable support to the nation … the system is the wellspring of the talent and tools needed to address a multitude of national problems, such as non-proliferation research, nuclear threat reduction, nuclear forensics, bioterrorism defense, missile defense, countering improvised explosive devices, nuclear energy, and alternative energy options.—Interim Report of the Congressional Commission on the Strategic Posture of the United States, December 11, 2008

The ability of the science base of the national laboratories to contribute to national security is a major theme for NNSA for the future and one that we have been developing over the past two years as the Administrator's "Focus Area 4" effort. The acceptance of this broadened mission by the NNSA and the Defense Laboratories has already resulted in an MOU with the Defense Threat Reduction Agency (DTRA), and I expect that as we move forward over the next several years, it will become an increasingly important part of our workload.

I am heartened by the acknowledgement of our value by the external community and appreciate the dedication of those responsible, in both direct and supporting roles, for the technical work lauded in these studies.

DoD/DOE Strategic Alliance Workshop Held on Computational Simulation

On November 20, 2008, the first in a series of DoD/DOE strategic alliance workshops was held on the topic of verification and validation methods for computational simulation. Attendees included representatives from DoD research organizations (Air Force Research Laboratory [AFRL], Army Research Laboratory, Office of Naval Research), along with representatives from the DOE/NNSA weapons labs (SNL, LANL, LLNL) and NNSA HQ.

The TridentII/D5 (Photo from US Navy)

The working group identified the following three priority areas for collaboration: (1) development of methodologies for simulation-based risk-informed certification of weapon subsystems/systems when testing is limited or impossible; (2) development of mathematical/statistical methodologies to validate simulation models using experimental test data; and (3) education of the technical work force on the use of these methods. These three collaboration areas are directly aligned with the interests of the Advanced Simulation and Computing (ASC) program.

This DoD/DOE collaboration will leverage several existing DoD/DOE projects and will also establish new partnerships. One current project is a Sandia/Navy study involving computational simulations of blast loads on ships. This study will allow for computational tool sharing, as well as knowledge sharing on topics such as verification, validation, sensitivity analysis, and uncertainty quantification. Additional DoD/DOE partnerships are being explored on topics such as large-scale blast-structure interactions and directed energy applications. Near-term activities in this DoD/DOE collaboration include DOE participation in a DoD verification and validation workshop in San Diego in March 2009, along with a sensitivity analysis and uncertainty quantification methods workshop at AFRL in May 2009. Future meetings of the DoD/DOE strategic alliance team will focus on topics such as large dataset analysis, renewable fuels, new power sources, and advanced materials.

Application High-End Computing Research Needs Investment in I/O and Storage R&D

A multiagency working group—with representation from ASC—to help manage overall government investments in high-end computing research & development predicts that in the near future, gaps and open problems in the file systems and I/O in high-end computing will be formidable. These challenges present areas in need of new and continued investment in R&D and standardization that the government should pursue.

The High-End Computing Interagency Working Group's members are from Los Alamos, Argonne, Sandia, Pacific Northwest, and Oak Ridge national laboratories, and NASA. They point out in a paper published in the January issue of the Association for the Computing Machinery Special Interest Group on Operating Systems that sites in the near future will routinely deploy supercomputers with hundreds of thousands of processors. They anticipate that storage bandwidth requirements will go from tens of gigabytes per second to terabytes per second.

For efficient complex science, online storage requirements to support workflows will approach the exabyte range. It will require the ability to handle a more varied I/O workload ranging seven orders of magnitude in performance characteristics, extremely high metadata activities, and management of trillions of files. It will be necessary to share data using global or virtual enterprise wide-area networks with flexible and effective security. The number of storage devices needed in a single coordinated operation could be in the tens to hundreds of thousands. It will be increasingly difficult to manage enterprise-class global parallel file/storage systems due to the number of elements involved, which may approach 100,000 spinning disks with widely varying workloads.

The paper acts as a preview to the National Science Foundation's 2009 High End Computing University Research call. To see the paper "Coordinating Government Funding of File system and I/O Research through the High-End Computing University Research Activity" go to http://institute.lanl.gov/hec-fsio/docs/HEC-FSIO-FY08-Gaps_RoadMap.pdf.

Roadrunner to Expand Hybrid Computing Applications

On December 22, 2008, Los Alamos National Laboratory (LANL) officially accepted the Roadrunner Phase 3 system from IBM. Having met this significant contractual milestone, LANL is moving to the next step of system and code stabilization. See the timeline graphic showing Roadrunner's implementation schedule from August 2008 to January 2010. Weapons and Open Science code runs are being used to accomplish the stabilization.

As part of the system and code stabilization effort for the ASC Roadrunner petascale installation at LANL, 10 Open Science projects have been chosen from a field of 29 proposals in a selective process to use Roadrunner this spring. The Open Science runs will increase the number of codes that can take advantage of the Roadrunner hybrid architecture, and will be the driver for many other applications worldwide. Abstracts of the LANL Open Science projects can be found at http://www.lanl.gov/roadrunner/rropenscienceabstracts.shtml Roadrunner has already successfully run codes including VPIC (plasma physics), SPaSM (molecular dynamics), MC transport (IMC), and deterministic transport (Sweep3D).

Roadrunner installed in the Nicholas C. Metropolis Center for Modeling and Simulation at Los Alamos National Laboratory. (Photo by R. Robinson, LANL)

The Weapons and Open Science stabilization is an exciting time for scientists in the high performance computing (HPC) and ASC communities. It is a time of rapid and unprecedented innovation. Roadrunner is large enough to model physical processes that are closer to nature and achieve greater precision in the solutions to enormously complex problems. The fastest computer in the world, Roadrunner is the first to run petascale simulations, and represents a significant step toward meeting ASC's major predictive capability goals.

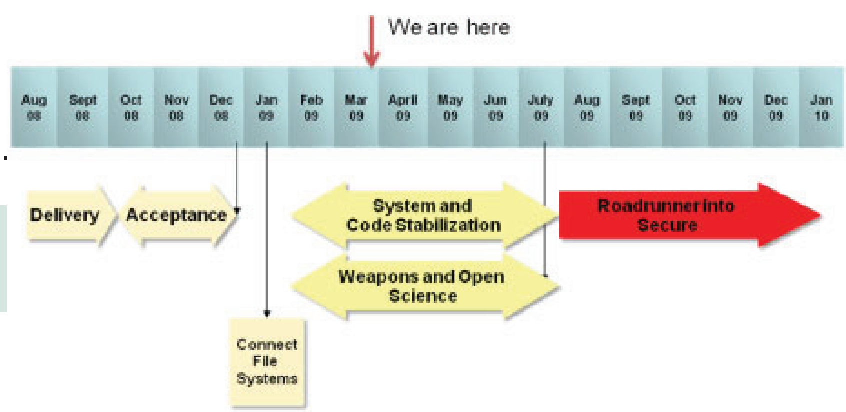

Roadrunner timeline from August 2008 to January 2010.

NNSA Awards IBM Contract to Build Next Generation Supercomputer

Two new IBM supercomputing systems will help continue to ensure the safety and reliability of the nation's aging nuclear deterrent. Sequoia will be a 20 petaFLOPS system based on future BlueGene technology, with delivery in 2011 and deployment in 2012. An initial delivery system, Dawn, a 500 teraFLOPS BlueGene/P system, is being delivered and installed. Dawn will lay the applications foundation for multi-petaFLOPS computing on Sequoia.

Sequoia is expected to be the most powerful supercomputer in the world and will be approximately 10 times faster than today's most powerful system. To put this into perspective, if each of the 6.7 billion people on earth had a hand calculator and worked together on a calculation 24 hours per day, 365 days a year, it would take 320 years to do what Sequoia will do in one hour.

The Sequoia systems will be located at Lawrence Livermore National Laboratory and be focused on strengthening the foundations of predictive simulation through running very large suites of complex simulations called uncertainty quantification (UQ) studies. In addition, the machines will be used for weapons science calculations necessary to build more accurate physical models. This work is a cornerstone of NNSA's Stockpile Stewardship program to ensure the safety, security, and reliability of the U.S. nuclear weapons stockpile today and into the future without underground testing. http://nnsa.energy.gov/defense_programs/index.htm

Sequoia will have 1.6 petabytes of memory, 96 racks, 98,304 compute nodes, and 1.6 million cores. Though orders of magnitude more powerful than such predecessor systems as ASC Purple and BlueGene/L, Sequoia will be 160 times more power efficient than Purple and 17 times more so than BlueGene/L.

"Sequoia represents an extremely bold step in simulation for the ASC Program. We have, in the past, explored, matured, and exploited high-performance computers to build the 3D codes and do our day-to-day work as efficiently as possible," said Michel McCoy, head of LLNL's ASC program. "To tackle the challenges of predictive simulation we face today requires us to build better science models and quantify uncertainty in our integrated calculations. To do this, we need a Sequoia-class system. Never before have we taken an advanced HPC architecture system and proposed to inject it into a mainline production environment at this scale. We will work closely with our customers and collaborators at all three labs to move the critical codes onto this multicore architecture." For more information, see the NNSA press release.

Kansas City Plant Acquires Tri-Lab Capacity Cluster Platform

As reported in the December 2008 issue of this newsletter, Kansas City Plant (KCP) has leveraged the work done by the three NNSA defense laboratories in an effort to bring forward the next generation of capacity platforms to the Nuclear Weapons Complex.

During KCP's initial visit to Sandia on June 19, 2008 to discuss the possibility of KCP getting a small Tri Lab Capacity Cluster (TLCC) system for a production role at their site, Sandia's TLCC personnel gave briefings on platform and deployment issues. With the help of Lawrence Livermore and Appro (the system vendor), KCP ultimately bought a 34 node compute cluster under the LLNL TLCC subcontract.

In September of 2008, KCP's cluster was ready for initial load testing at the vendor facility. Sandia ran the complete Pre-ship TLCC Synthetic Work Load (SWL) against the machine. The system was then shipped to KCP where Appro performed the site integration. Sandia personnel were granted accounts and access to the KCP system and performed the Post-ship SWL to validate the machine. Shortly thereafter, KCP administrators spent a week at Sandia training to "install and set up TOSS" on their machine. KCP left this training session with enough tools to be self-sufficient with the day-to-day operations of their cluster. Finally, Sandia assisted KCP with many tasks (e.g., resource manager setup) to help get their machine into production and provide on-going support as needed.

At the TriPod Workshop on Common Computing Environment held January 13th in Albuquerque, Jim Mahoney, KCP Plant TLCC Development Lead thanked the entire ASC team: "I want to summarize by saying we are very pleased with the TLCC machine and I thank you for all of the support that we have had from the tri-lab."

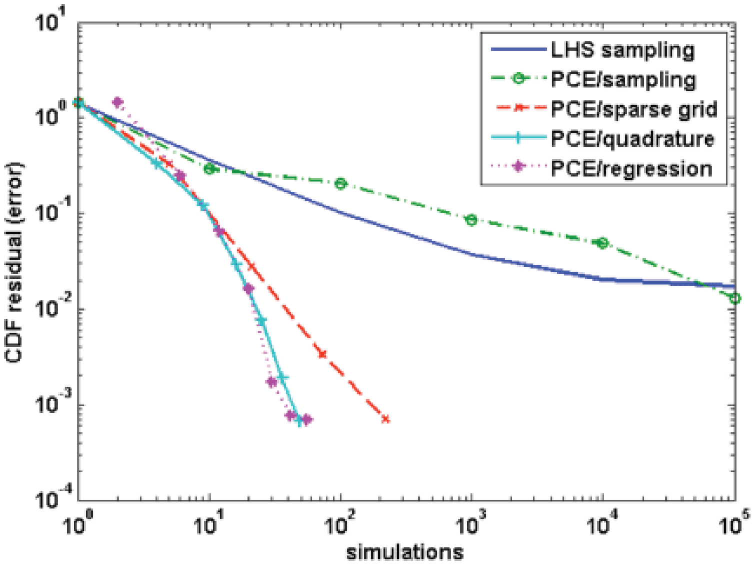

DAKOTA's stochastic expansion methods can resolve response probability distributions considerably more efficiently than traditional sampling-based approaches. For the "log ratio" benchmark problem, the plot depicts exponential convergence for polynomial chaos approximations constructed with advanced multi-dimensional integration techniques (sparse grid, quadrature, and regression), compared with √n convergence typical of Latin hypercube sampling. These advanced methods will be deployed for UQ and QMU with ASC Full-System Models.

SIERRA Successfully Deployed at KCP in Support of Defense Programs Applications

The latest SIERRA engineering simulation framework has been deployed on the KCP Tri-Lab Capacity Cluster (TLCC) platform in support of various Defense Programs (DP) applications. The codes have been very successful, both on implementation and operations. The projects performed since last fall include various W76 applications–supporting current production build schedules. Examples include a Lightning Arrestor Connector abnormal loading review, Trajectory Strong Link feed-thru header sealing evaluations, encapsulation flow of electronics, header welding operations, vendor forging operations, safety reviews of mixing operations, and high-voltage abnormal drop loadings. Advanced development support for internal PDRD projects is also being performed with the tools, focused on fluid/solid interactions and coupling methods. The figure below shows a design change to increase durability by eliminating solder joints and a crystal based on SIERRA simulations.

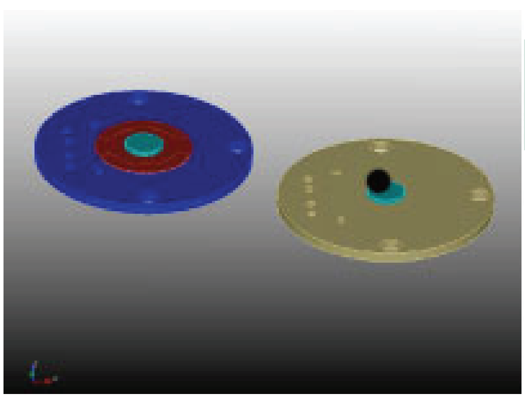

A design change based on Sierra simulations.

The tools, techniques, and fast processing utilizing the SIERRA framework have been critical to meeting schedule targets on the W76-1 arming fuzing & firing (AF&F) builds. Uncertainties in manufacturing, including abnormal environment issues and rate production concerns, are being addressed with the toolset. Specifically, the Aria, Adagio, Presto, and Calore software tools are being used, along with coupling using the Calagio code. It is anticipated that KCP will continue to use the simulation methods to support AF&F production concerns, and use this methodology of evaluations for new production schemes for future LEPs.

KCP would not have been able to advance so quickly with the tools had it not been for the dedicated work done to deliver the codes to a Production site. This includes the efforts on the application builds, the work done for stability of the codes, and the custom training and support given to make this remote effort successful.

Structural Simulation Toolkit Mitigates Risks in Supercomputer Deployments

Taking chances with supercomputing is a very unhealthy business. Costs are in the tens to hundreds of millions. Complexity is high with numerous commodity and customized parts requiring integration. Innovative architectures like multicore come with extreme uncertainty. The simulation codes to be run are diverse. Potential unknowns are everywhere. Sandia has responded to this high risk by developing the Structural Simulation Toolkit (SST) to provide early insight into how future supercomputers will perform.

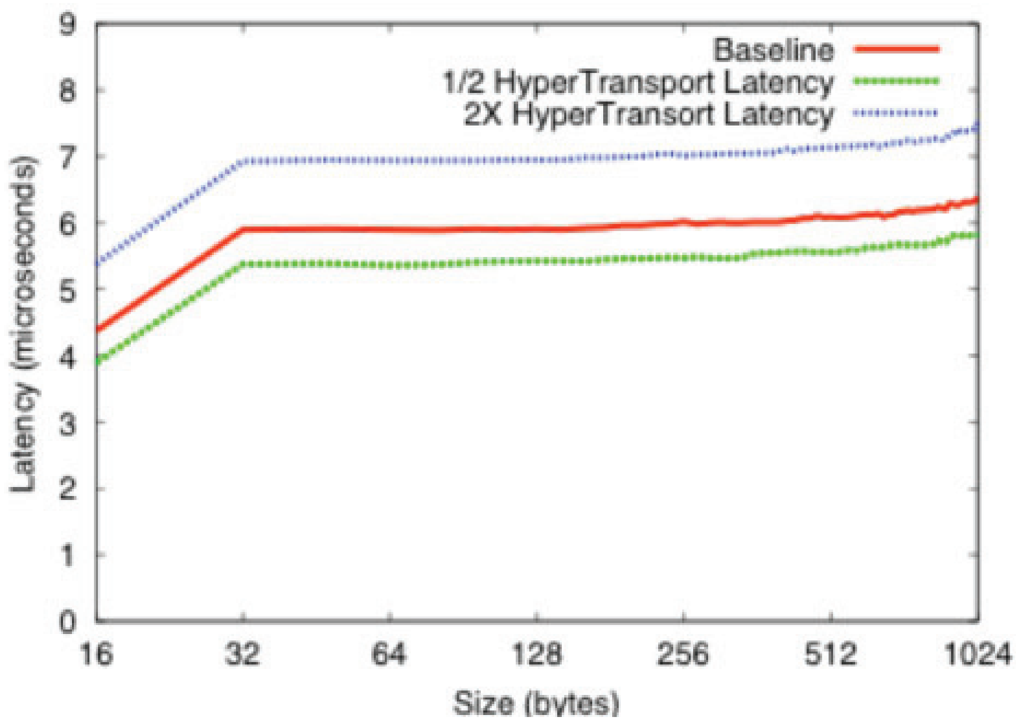

Simulations reveal potential enhancements to the Red Storm network.

The SST enables researchers to explore the performance of supercomputers running complex simulations virtually, all before funds and resources are committed. The SST has already had dramatic impact on the success of supercomputer deployments. It has also identified significant issues in scaling multicore hardware and pinpointed differences between current performance tests and real-world applications.

A specific example is of the use of SST to model the network used on Sandia's Red Storm Supercomputer. These simulations helped identify bottlenecks and also reveal how future enhancements to network interfaces and components could improve performance.

Understanding the performance limitations of future supercomputers is key to supporting a range of simulation-based science activities. The field of computer architecture is currently exploring a range of new hardware and software techniques. The Structural Simulation Toolkit provides a modular framework to explore these novel hardware and software systems.

Study of Turbulence Mixing Showcases Science and Computation

Some of the largest fully resolved simulations of turbulence mixing to date showcase the ASC Program's scientific and computational science capabilities. New and unexpected physics have been revealed in the study of mixing driven by strong pressure gradients. Applications of interest include inertial confinement fusion targets, laser induced launching of a flyer plate, stellar pulsations, and supernova explosions.

The turbulence mixing problem, being a result of the processes of advection, stirring, and molecular mixing occurring at different scales—is unusually complex. In the flow visualization image shown, a large pressure gradient drives the interpenetration of two pure fluids of very different density.

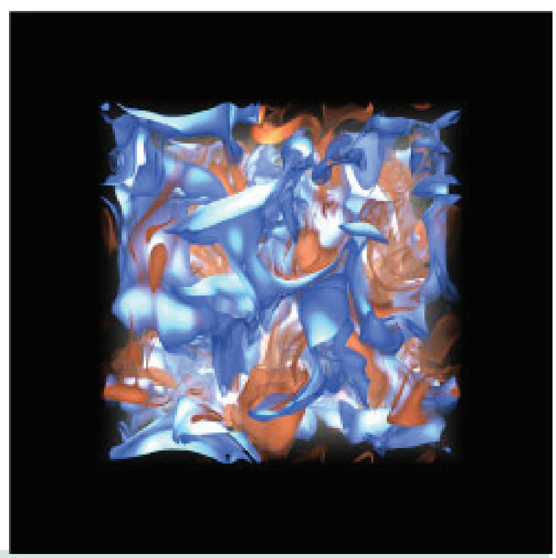

A still image showing the visualization of two different density fluids. The heavy (blue) fluid has density three times larger than the light (red) fluid.

Unlike most mix processes studied, the turbulent mixing process between two very different density fluids is highly skewed: the light fluid mixes much more rapidly than the heavy fluid. Mathematically speaking, this is because the equations in the mixing of large density materials have cubic nonlinearities as well as the usual quadratic nonlinearities. In short, the pure heavy (blue) material lasts longer than pure light (red) material, and this effect increases with density ratio.

Experiments and numerical simulations such as these produce an overwhelming amount of data and complexity. Visualization is one of the best tools to deal with this data and complexity. Scientists Pat McCormick and Steve Martin at Los Alamos have developed a novel visualization code for heterogeneous architectures, such as Roadrunner's new innovative architecture. Using Martin's code, they have produced movies of a simplified model problem that isolates specific mixing physics. In the attached movie, the turbulent mixing problem starts from rest with two pure, different density fluids (blue heavy and red light) filling a closed box.

Los Alamos researchers Daniel Livescu, of the Computer, Computation, and Statistical Sciences Division, and J. Ray Ristorcelli and Rob Gore, of the Applied Physics Division, have published this work in journals,[1] diverse conference proceedings, and book chapters.

[1] *J. Fluid Mech.* 591, p. 43 (2007) and 605 p. 145 (2008), and *J. Turbulence* (to appear)

The Dawn of Petascale Computing at Lawrence Livermore

IBM recently delivered the final seven racks of the Dawn Supercomputer to LLNL. All 36 racks of the system are being installed in the Terascale Simulation Facility. Dawn is a predecessor to the gargantuan 20-quadrillion-operations-per-second Sequoia system, which will be delivered in 2011.

At 500-teraFLOPS (a half petaFLOPS), Dawn (a BlueGene/P system) will serve as a bridge and accelerate the development of the petaFLOPS computing Sequoia will make possible. Once acceptance of the system is complete in late March, Dawn will be used as a code-development and scaling platform for tri-lab weapon science codes targeted for Sequoia.

The timeline for Dawn delivery, integration, testing, and acceptance is aggressive. Delivery began on Jan. 20, and the machine is expected to be available to users at the end of March.

Application-Level Services Enable In Situ Analysis of Shock Physics Codes

The ability to perform meaningful, real-time analysis and visualization of scientific data, either simulated or real, is a significant advance that has the potential to fundamentally change the way computational scientists work. In situ analysis provides immediate feedback that allows the scientist to quickly detect bugs, identify areas of interest, or interactively modify the behavior of a running application. Such a capability will dramatically improve productivity of the application scientist and significantly reduce time-to-solution for time-critical ASC applications.

In a joint effort between the Scalable Data Analysis and Scalable I/O groups of the ASC Computational Systems and Software Environments (CSSE) project, researchers at Sandia National Laboratories are exploring the use of "Application-Level Services" for in situ analysis of the shock physics code CTH. An application-level service is a separate application that executes on either compute nodes or service nodes of an HPC system. In contrast to system-level services (e.g., file services), an application-level service is dedicated to a single application.

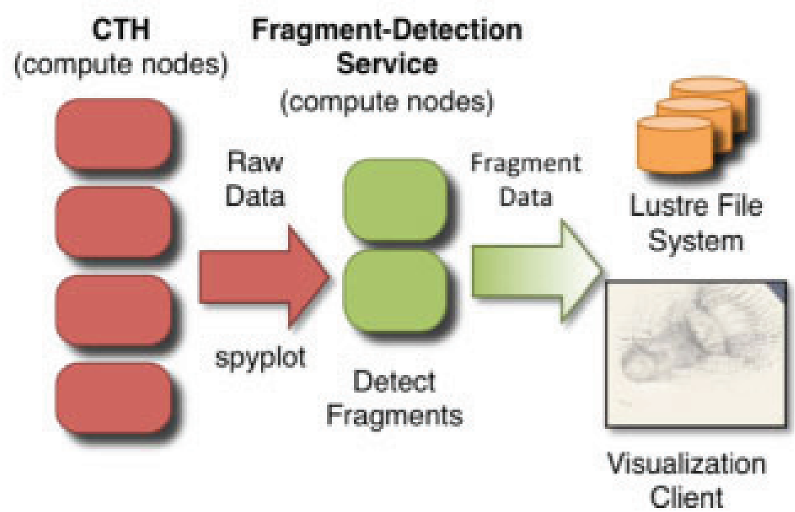

The fragment detection service provides on-the-fly data analysis with no modifications to CTH.

The application-level service for CTH detects and tracks material fragments generated by the simulation of high-speed impacts. Fragment detection and tracking provide valuable insight to the scientist, but present a number of computational challenges. First, fragment tracking is data intensive because it requires data from every time-step calculation. The I/O requirements (both capacity and write bandwidth) of performing fragment tracking as an off-line task make this option impractical. Second, fragment detection is computationally expensive. If integrated with the CTH code, fragment detection adds as much as 30% to each simulated time-step calculation, significantly impacting scalability of the application. In addition, integrating detection into the CTH code adds complexity and significant programming burden to the application developer.

The application-level service to detect and track fragments executes on a separate partition of compute nodes on the HPC system and leverages technologies developed by the Lightweight File Systems project and Paraview, two ASC-funded projects. Offloading fragment detection to a separate partition of compute nodes provides several advantages: it allows concurrent CTH computation and fragment detection; it significantly reduces the I/O to storage by outputting only the fragment information; and, perhaps most importantly, it does not require major modifications to the CTH code. The application-level service operates on requests generated by the spyplot I/O API (the API that is already used by CTH).

The application-level service approach represents a significant advance in the ability to perform in situ analysis of scientific data and will have a broad impact in the greater computational-science community. Our work with CTH demonstrates the first significant use of this technology and will serve as a model for future efforts with other codes.

Report Now Available from ASC Risk Management Workshop for High-Performance Computing Centers

The RMTAP workshop, held Sept. 17 and 18, 2008, in San Francisco, CA, convened to assess current and emerging techniques, practices, and lessons learned for effectively identifying, understanding, managing, and mitigating risks associated with acquiring leading-edge computing systems at high-performance computing centers (HPCC). The report from the workshop detailing discussions and findings is now available online.

Sponsored by the DOE—jointly by the Office of Science and the NNSA—and hosted by LLNL, the workshop was targeted at HPCC managers and key staff who are planning for leading-edge computational systems.

The overall workshop findings were the following:

- Standard risk management techniques and tools are in the aggregate applicable to projects at HPCCs and are commonly employed by the HPC community
- HPC projects have characteristics that necessitate a tailoring of the standard risk management practices
- All HPCC acquisition projects can benefit by employing risk management, but the specific choice of risk management processes and tools is less important to the success of the project
- The special relationship between the HPCCs and HPC vendors must be reflected in the risk management strategy
- Best practices findings include developing a prioritized risk register with special attention to the top risks, establishing a practice of regular meetings and status updates with the platform partner, supporting regular and open reviews that engage the interests and expertise of a wide range of staff and stakeholders, and documenting and sharing the acquisition/build/deployment experience
- Top risk categories include system scaling issues, request for proposal/contract and acceptance testing, and vendor technical or business problems

Seager Honored by *Federal Computer Week*

Lawrence Livermore's Mark Seager was selected by *Federal Computer Week* magazine as one of this year's "Federal 100" top executives from government, industry, and academia who had the greatest impact on government information systems in 2008.

Seager was selected because of "the difference you made in the way agencies, companies and government officials develop, acquire, manage and use information technology." The nomination was submitted by industry collaborators for Seager's leadership of the Hyperion Project, a collaboration with ten industry leaders to advance next-generation Linux 1q performance computing clusters.

"Hyperion represents a new way of doing business. Collectively we are building a system none of us could have built individually," Seager said when the project was announced at SC08 last November. "The project will advance the state-of-the-art in a cost-effective manner, benefitting both end users, such as the national security labs, and the computing industry, which can expand the market with proven, easy to deploy large-and small-scale Linux clusters."

LANL Computing Platforms Available for PSAAP Centers

As of February 2009, Los Alamos National Laboratory has announced a large increase in the amount of computing power available for use at Los Alamos by the Predictive Science Academic Alliance Program (PSAAP) Centers. To read more about the PSAAP Centers, see the announcement in the March 2008 ASC eNews.

For more details and contact information about the resources at LANL, go to the web site http://www.lanl.gov/asc/ and click on the PSAAP link under the heading Documents.

Three new platforms, including the hybrid Roadrunner system Cerrillos, are available. Users are encouraged to take advantage of these powerful resources:

Los Alamos National Laboratory

Lobo

272 compute nodes, 4,352 cpu, 38 Tflop system. Quad-core, quad-socket AMD Opteron w/ Infiniband

https://bear.lanl.gov/drupal/?q=tlcc_home (requires cryptocard to access)

Coyote

1,290 compute nodes, 2,580 cpu, 13 Tflop system. Dual-socket AMD Opteron w/ Infiniband. Also 9 serial nodes with two AMD Opterons per node.

https://bear.lanl.gov/drupal/?q=coyote_home (requires cryptocard to access)

Cerrillos

Hybrid Roadrunner system sized at 1 CU (180 compute nodes). Each node is a "tri-blade" with a dual-socket, dual-core Opteron and four cell processors.

https://bear.lanl.gov/drupal/?q=phase3_home (requires cryptocard to access)

ASC Booth Showcases Numerous Achievements at SC08 in Austin, TX

With a theme of "Leading HPC—Past, Present and Future," the ASC Program once again demonstrated the numerous tri-lab advances in high performance computing at the International Conference on High-Performance Computing, Networking, Data Storage, and Analysis. Sponsored by the Association for Computing Machinery (ACM) and the Institute of Electrical and Electronics Engineers (IEEE) Computer Society, this year's conference was held in Austin, TX, from November 15th to the 21st.

Led by Sandia National Laboratories this year, the booth was once again visited by U.S. and international experts and students in the field of high performance computing. ASC participants from the each laboratory served on conference committees, presented at technical conference sessions, gave demonstrations at the booth, and set up the booth. Booth participants had a full schedule of presentations and demonstrations.ASC (formerly ASCI) has been exhibiting its lab-developed technologies at Supercomputing since 1996. The first tri-lab booth was set up at the conference in Pittsburgh.

ASC Salutes Sue Kelly

For the past seven years, **Sue Kelly** has been a driving technical force behind the acquisition, development, bring-up, and sustainment of the Red Storm massively parallel processing supercomputer located in Albuquerque. Sue began working on the Red Storm project soon after the contract was awarded to Cray, Inc., in 2002. The project involved considerable engineering R&D, some of which was performed by Sandia and was led by Sue. Sandia provided the scalable run time system software, which included the operating system, its libraries, the application launch utility, and the node allocator.

This work on Red Storm was a natural follow-on to Sue's prior assignment on the ASCI Red Supercomputer. When the contract with Intel ended, Sue stepped in to lead the system software effort on ASCI Red. Sue and her team assumed responsibility for the multiple million lines of source code that comprised the system software. They kept the system operational for three more years.

Prior to those HPC projects, Sue worked on various high performance storage systems used by Sandia scientific computing systems. She was an original HPSS developer.

Sue is now a member of the NM Alliance for Computing at Extreme Scales (ACES) team, which is tasked with acquiring ASC's next capability computer system to replace Purple. The ASC Program is delighted to have her 20+ years of system software experience on this important acquisition.

Newsletter Points of Contact

Send submittals to:

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

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

Reeta Garber—Sandia National Laboratories ragarbe@sandia.gov

Who's Who

ASC Program Managers—Headquarters

Director, Office of Advanced Simulation and Computing, NA-121.2

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

ASC Program Managers—Labs

Los Alamos National Laboratory

John Hopson—jhopson@lanl.gov

Cheryl Wampler—clw@lanl.gov

Lawrence Livermore National Laboratory

Michel McCoy—mccoy2@llnl.gov

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

Sandia National Laboratories

James Peery—jspeery@sandia.gov

Art Ratzel—acratze@sandia.gov

Developed and maintained by Sandia National Laboratories for NA.121.2

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