ASC ASCeNews QUARTERLY NEWSLETTER

NA-ASC-500-07 Issue 4

September 2007

The View from HQ

By Dimitri Kusnezov

Perceptions are not always reality

"Why does Dimitri hate our Lab?" This question was posed to a colleague recently—and passed on to me. My first response was surprise, which quickly turned to introspection. I do not think of the labs in such a context; rather, I think of them as dynamic entities. I see the labs through the people, through the trust and professional relations that develop over time and through their ability to work with us and deliver on commitments. While I might claim the question is ill-posed, that does not address the real issue: How does such a perception arise?

The purpose of the labs was forged in a time of extraordinary national need. The labs were formed to support national goals through an often intense scientific competition and mutual review. Their value to the country and the corresponding financial support were a matter of national consensus. In the early days, resources were made available in almost unlimited amounts, and hard choices were rarely necessary.

Today, the need to support such a far-reaching range of scientific and engineering capabilities at the labs is less clear. Budgets are bitterly contested at the congressional, agency, and the laboratory levels. Decisions that ensure a viable future for the labs need to be made in the context of skill sets, competencies, and performance. Often, difficult decisions result in winners and losers and lead to the perception of more or less favored institutions. The implementation of the 2030 complex transformation agenda may appear as though Washington managers are rewarding the plants and the real work at the expense of the more basic science-oriented activities.

Despite this appearance, I am committed to supporting the science and technology on which the success of NNSA has always depended. But, at the same time, it is necessary for me to stretch the available resources to cover a broader spectrum of activities and to insist that the monies we allocate are leveraged to the greatest extent possible. The ASC Program, like the other NNSA programs, requires that the labs evolve toward a more interdependent system of cooperating research-and-development institutions, working together and avoiding unnecessary duplication, and moving away from the old model of completely self-sufficient, competing labs. The long-term future of the labs depends on a stronger partnership with each other and with NNSA and away from the adversarial relationships that are more common than not today.

It is in this context that I must drive future directions, often make unpopular choices, and recognize that I can be seen as playing favorites and selecting winners. I strive for a partnership between Headquarters and the sites and am seriously committed to fairness, based on objective criteria, external reviews and scrutiny, and not likes or dislikes. I must confess that this does not make life easy. If I didn't love the labs and value and respect the science and technology that comes out of these institutions, I couldn't do this job.

ASC Participates in Second Annual Modeling and Simulation Exhibition on Capitol Hill

From left: Karen Pao, April Commodore, Ken Alvin, Thuc Hoang, and Dimitri Kusnezov.

The ASC Program recently participated in the second annual Modeling and Simulation Exhibition on Capitol Hill in Washington, DC. The exhibition was hosted by the Congressional Modeling and Simulation Caucus, which aims to promote the potential impact of modeling and simulation technologies on the nation's military and defense capabilities, as well as on the broader economy. Currently, the Caucus is very focused on the use of simulation in military training and preparedness, and the exhibition strongly reflected these types of applications. The ASC Program was a unique and broadening contributor to the exhibition of simulation capabilities, with an emphasis on predictive science and high-performance computing.

ASC Headquarters, with significant support from Sandia's ASC program office, presented a booth at the exhibition where interested constituents, congressional staffers, and members of Congress could learn more about the predictive capabilities of ASC simulation codes, software environments, and computers. As part of the showcase, ASC displayed an interactive large-scale fire simulation as well as various visuals of simulation codes from Sandia, Lawrence Livermore, and Los Alamos national laboratories.

U.S. Congressman Randy Forbes (R-VA) and Congresswoman Heather Wilson (R-NM) stopped by the ASC booth while promoting H. Res. 487, a recently introduced Congressional resolution, which encourages the ongoing development of computer science and mathematic applications for the general welfare of the United States and its defense capabilities.

World's Fastest Supercomputer Delivers Breakthrough Science Simulations for the NNSA's Nuclear Weapons Program

The BlueGene/L supercomputer at Lawrence Livermore National Laboratory topped the list of the world's fastest computers for a record sixth straight time, according to the new Top500 list released Wednesday, June 27, at the International Supercomputing Conference in Dresden, Germany.

Built by IBM, BlueGene/L (BG/L) clocks in at 280.6 teraFLOPS (trillion floating operations per second) on the LINPACK, the industry standard for supercomputer performance. BGL is a workhorse machine for the U.S. Department of Energy's National Nuclear Security Administration's (NNSA) effort to ensure the safety, security, and reliability of the nation's nuclear deterrent without underground nuclear testing, known as the Stockpile Stewardship Program.

NNSA's Advanced Simulation and Computing (ASC) Program's 100 teraFLOPS Purple system, another IBM machine at Lawrence Livermore National Laboratory, dropped from fourth to sixth on the Top500 list.

"Since BG/L went into production in early 2006, it has performed beyond our expectations and delivered for the ASC Program. BG/L's architecture has proven suitable for a much broader range of applications than originally envisioned," said Dona Crawford, Associate Director for Computation. "Likewise, ASC Purple also has demonstrated the system's ability to deliver weapons simulations of unprecedented spatial resolution for the Stockpile Stewardship Program."

The detailed computer simulations of nuclear weapons performance produced by the ASC Program using BG/L, ASC Purple, and other supercomputers at the three nuclear weapons labs are a cornerstone of stockpile stewardship. ASC is a tri-lab program uniting the high-performance computing expertise of NNSA's Los Alamos, Lawrence Livermore, and Sandia national labs.

BG/L's three year reign as the world's fastest supercomputer has seen significant progress in code development and the achievement of numerous milestones for NNSA's Stockpile Stewardship Program. For example, simulations on BGL helped answer critical questions about plutonium aging—a key to understanding the life expectancy of nuclear weapons systems. Breakthrough calculations/simulations run on the machine have over the last two years garnered three Gordon Bell Prizes, widely regarded in the computing community as the Oscars of high-performance computing.

Simulations on BG/L of high explosives, super-ionic water, and graphite-to-diamond experiments have provided scientific insights and/or confirmed results of earlier physical experiments. Livermore scientists used BG/L to perform the first instability simulation with a Reynolds number large enough to determine the nature of turbulence beyond the mixing transition, a feat that made the cover of *Nature Physics* magazine in August 2006.

In June 2006, BGL set a new world mark for a scientific application with a sustained performance of 207.3 teraFLOPS on the "Qbox" computer code for conducting materials science simulations. This represents a leap forward in scientists' capability to perform predictive simulations of large, complex high-Z metals relevant to stockpile science and was awarded the 2006 Gordon Bell Prize for peak performance.

Meeting Predictivity Milestones with High-Performance Computing

High-performance computing (HPC) is needed to support the Predictivity milestones for certification. In particular, HPC can be used to eliminate knobs by solving the many microphysics issues in relevant regimes and over the relevant range of spatial scales. This is called Science@Scale. This strategy developed at Los Alamos involves multi-scale computing in which they couple multiphysics codes with unit-physics codes to obtain more physics-based design tools. For certification, the multiphysics codes must describe a variety of phenomena and, thus, must employ many approximate sub-grid physics models. In contrast, the unit-physics codes are nearly *ab initio* because they focus on specific issues and utilize few approximations. These include simulations such as molecular dynamics (MD) and particle-in-cell (PIC), which can be used to calculate basic material properties and transport rates under the realistic conditions set by the multiphysics codes. In order to achieve this multi-scale computing, they estimate needing petaFLOPS computing for a NIF-scale object and more for larger applications. This may require

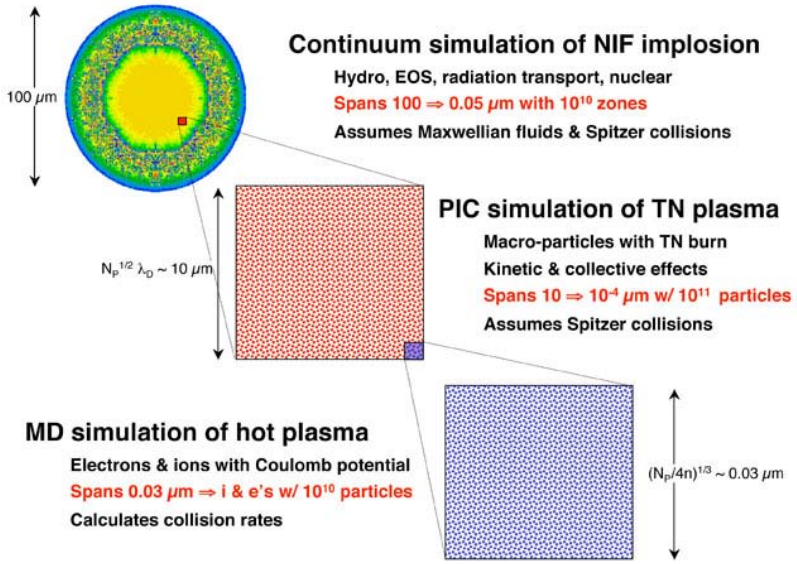

Multi-scale computing will allow near *ab initio* simulation of a NIF-scale implosion.

an extrapolation of HPC to heterogeneous architectures, such as Roadrunner (sited at Los Alamos National Laboratory), that employ a combination of conventional processors (Cores) and single-instruction, vector processors (Cells). Researchers at Los Alamos have successfully converted several unit-physics codes to the Roadrunner architecture, including PIC, MD, and hydrodynamics codes. For more information, contact Guy Dimonte at dimonte@lanl.gov.

V&V Analysis Conducted of Large Simulation Results Using Scalable Python Scripting Capability

Verification of complex, large results from Sandia's simulation codes is crucial to understanding the science behind the results. Sandia's scalable python scripting capability, released in ParaView 3.0, enables flexible, advanced V&V analysis of large data within a familiar python scripting environment. Using this capability, analysts write small modules in python, addressing specific solutions, or codes. The scalable python capability manages the rest (data distribution, load balancing, parallel python execution), making it simpler than ever for analysts to work on large data analysis problems.

The image below shows results computed for the Noh problem, which describes a canonical, spherically symmetric, hydrodynamic implosion and has a known analytic solution. The initial velocity is directed radially inward. As the solution evolves, a strong shock wave propagates outward. Most simulation codes overpredict the temperature at the origin, often referred to as "overheating." In practical applications overheating is often observed when shock waves reflect from walls or material interfaces. With its scripting capability, Paraview computes the exact solution, the volume, and the volume-weighted error for each element of the simulation – all using python code written by the analyst.

ParaView 3.0 is an open source application available for download at http://www.paraview.org.

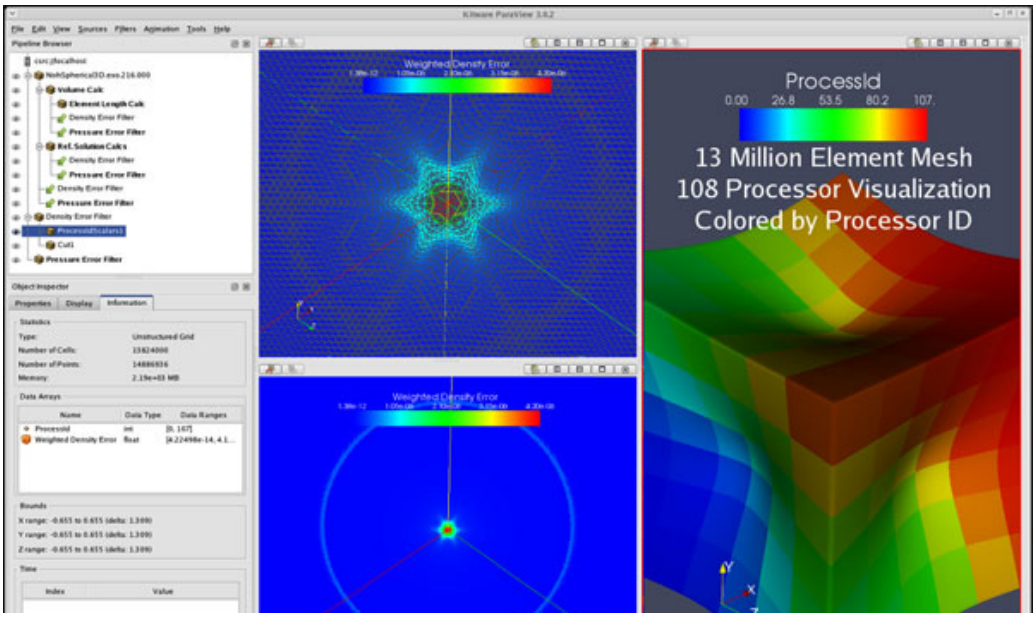

Left: Analysis of simulation results performed and displayed with ParaView 3.0. This image shows the results of analyzing the results of a mesh containing 13 million elements. The analysis and visualization was run on 108 processors and the results were displayed on a PC in the analyst's office. Note the view of the data in the right panel, in which the mesh is colored by processor ID. This shows how the simulation results were distributed for the analysis. The two smaller panels in the middle of the screen shot show the results of comparing the simulation results with an ideal result, computed in python. The weighted error is then shown in detail at the center of the mesh (top), and at a larger scale (bottom), showing the error around the shockwave.

Improving Supercomputer Reliability via Data Mining

Computer logs often provide critical information about malfunction or misuse, but finding and correlating the clues interspersed among millions of lines of time-stamped text messages generated by supercomputers is in itself a challenging data mining task. The Sisyphus toolkit is the result of four years of research and development on how to efficiently find the important nuggets of information in supercomputer logs. Now in production use on Red Storm (sited at Sandia National Laboratories), it has automatically detected—and more importantly isolated—a wide range of problems including failures (disks, I/O controllers, network interfaces, power supplies, and memory), misconfigurations (BIOS, RAID controller, system software, and inconsistent versions), and problematic user behavior (unbalanced RAID stripe usage, inappropriate remote monitoring). This has enabled focused proactive and reactive responses by system administrators, thus increasing system reliability.

Sisyphus is based on the premise that similar computers correctly executing similar workload should produce similar logs—and thus, anomalies warrant investigation. It automatically ranks log files and colorizes words based on information theory, answering the questions, "What is the most unusual logfile?" and "Exactly what makes this logfile unusual?" It works with ASCII text (rather than numerical data as most anomaly detectors do) and is general enough to be used with any computer logs. It provides useful file and word statistics in tabular and plot formats, and includes web and command line interfaces. See http://www.cs.sandia.gov/~jrstear/sisyphus/ for video demos, downloads, and documentation.

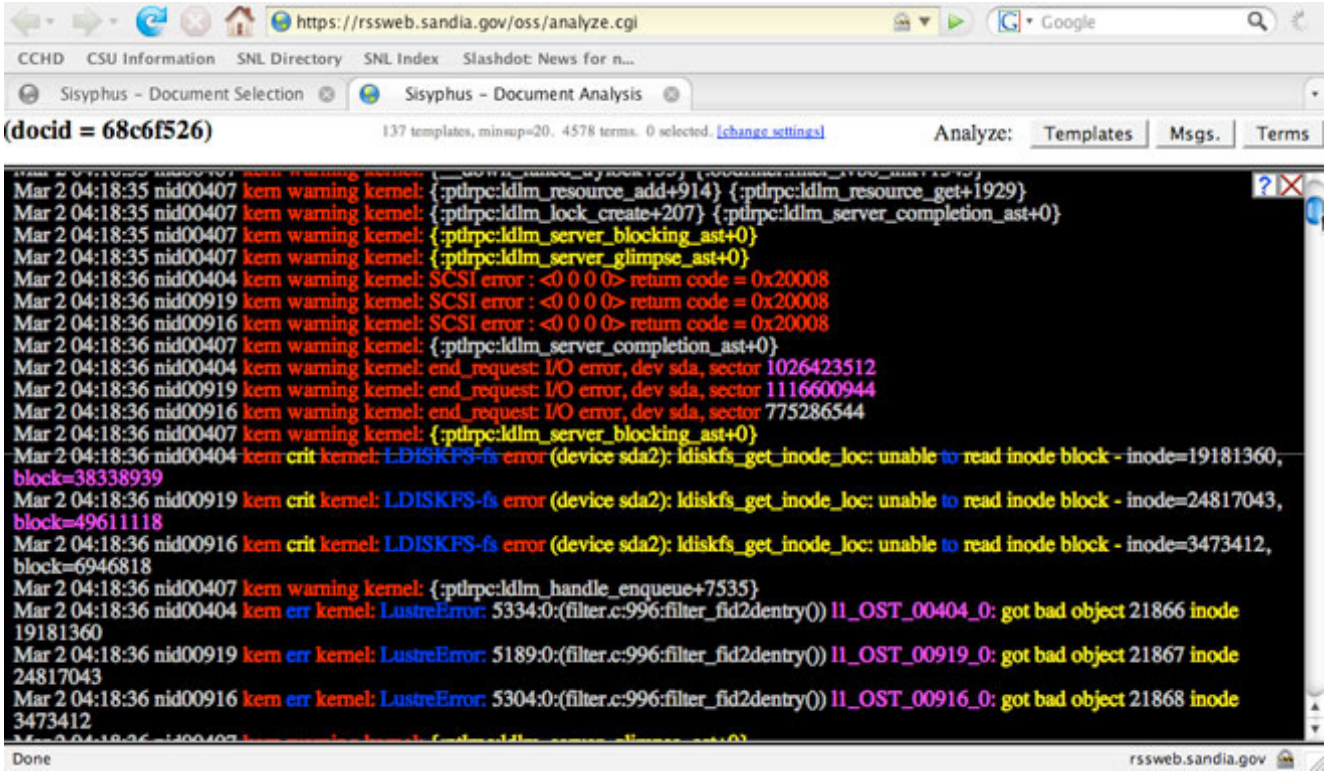

Automatically highlighted symptoms of a common cause failure (underlying RAID controller) on multiple Red Storm I/O nodes.

New Model for Tantalum Developed under ASC Multiscale Modeling Effort

More than 10 years ago, an effort began to develop improved, physically based models of material constitutive behavior for use in integrated multiphysics codes. The challenge has been that the models are needed for programmatic applications in regimes of pressure, strain-rate, and temperature, where relevant experimental data are scarce or nearly impossible to obtain. Building upon the ASCI paradigm at the time, a multiscale modeling strategy was developed to bring together and leverage material modeling activities from a variety of length scales. This technique uses new and improved computational simulations involving detailed physics to enable the development of better macroscopic models for the integrated codes.

Lawrence Livermore National Laboratory has just completed a milestone for the Physics & Engineering Models (PEM) element of the ASC Program that utilized the multiscale modeling methodology to develop a new strength model for tantalum. This effort involved combining theoretical developments with modeling and simulation efforts involving quantum molecular dynamics, classical molecular dynamics, dislocation dynamics, and integrated continuum simulations. Dislocation dynamics is a relatively new simulation capability, using the Livermore parallel code ParaDiS, where the motion, generation, and interaction of dislocations (crystallographic irregularities) are modeled as a way to predict realistic single crystal behavior. The new model has been successfully used to run large-scale, multiphysics programmatic simulations.

The development of the new tantalum strength model using the multiscale modeling methodology has truly been enabled through ASC model developments and modern platforms. Calculations required for this milestone involved utilizing one-third of the BlueGene/L machine to perform dislocation dynamics simulations requiring more than 200 million CPU-hours—equivalent to running a 1985 Cray Y-MP from the great ice age of 20,000 years ago until now. Additional molecular dynamics simulations required more than 10 million CPU-hours.

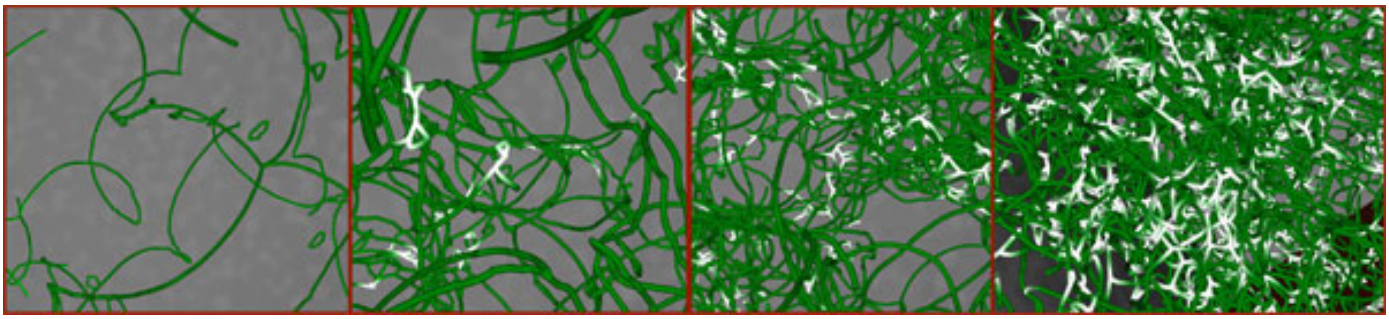

Four snapshots in time of a ParaDiS dislocation dynamics simulation, where complexity of the problem continuously increases as dislocations multiply and interact.

New Interface Reconstruction Method for Multi-Fluid Flows Developed as Part of LANL ASC Strategic Research

Los Alamos researchers Rao Garimella, Sam Schofield, and Marianne Francois, along with Raphael Loubere from Universite Paul-Sabatier, France, have developed a new interface reconstruction method for multi-fluid flows that is completely material-order independent. This is in contrast to traditional interface reconstruction methods, which produce very different results depending on the order in which materials are processed, leading to errors in advection of the materials.

The new method reconstructs the interfaces in three steps. First, the relative locations or approximate centroids of the materials are determined using a linear reconstruction of the volume fraction function over multi-material cells. Next, the approximate centroids of the materials are used to construct a weighted Voronoi or Power diagram subdivision of the cell while exactly matching the input volume fractions. Finally, the interface segments are smoothed using a constrained optimization process in order to minimize the slope discontinuity with interfaces in the neighborhood.

The method shows marked improvement in the reconstruction of multi-material interfaces compared to order-dependent methods. It can reconstruct straight lines exactly, and initial experiments indicate that it is second-order accurate. Finally, it also shows improved accuracy in dynamic advection tests of multi-material shapes. A preliminary implementation of the method is in LANL ASC Code Project A.

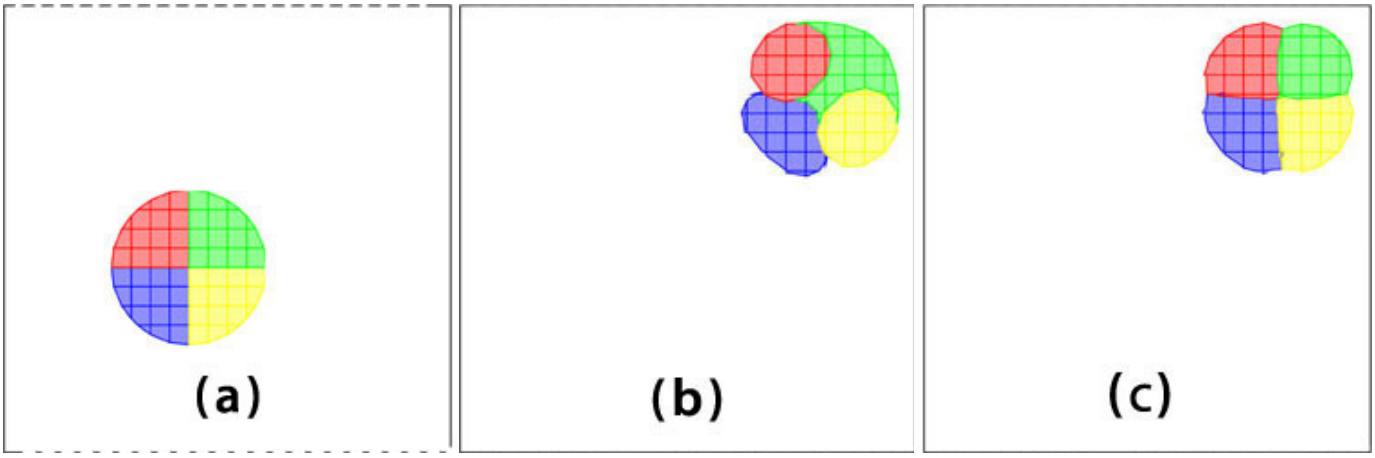

Diagonal advection of a 4-material bubble on a 40x40 grid with a velocity of (1.1,1.1): (a) initial configuration, (b) results of order-dependent method after 100 timesteps, and (c) results of new order-independent method with smoothing.

New Conservative Remap Method in FLAG

Los Alamos mathematical modeling and analysis researchers Sam Schofield and Rao Garimella have implemented a new, prototype conservative remap method in the FLAG Arbitrary-Lagrangian-Eulerian (ALE) code that is based on exact intersections of the old mesh with the relaxed mesh. In ALE calculations, the mesh moves and distorts with time. When the mesh becomes tangled or of poor quality (Fig. 1a), it is smoothed, the solution values are mapped, in a conservative manner, from the old to the smoothed mesh (Fig. 1b), and the simulation continues. In the FLAG code (LANL ASC Code Project B), the existing remap method is based on approximating the changes in the values in the mesh cells between the new and old meshes. Because of accuracy constraints, it is limited to small displacements.

With our new remap method, the smoothed mesh cells are intersected with a linear reconstruction of the field values on the old mesh to determine how much of each conserved quantity is located in the new mesh cell. Materials are remapped by intersecting the new mesh cell with any reconstructed interfaces on the old mesh. There is no loss of accuracy for large displacements between the meshes and monotonicity is guaranteed if the linear reconstruction is monotonic.

Figures 1a and 1b contrast the old vs the new remap. The new intersection-based remap allows the mesh smoother to take larger steps, maintaining a quality mesh.

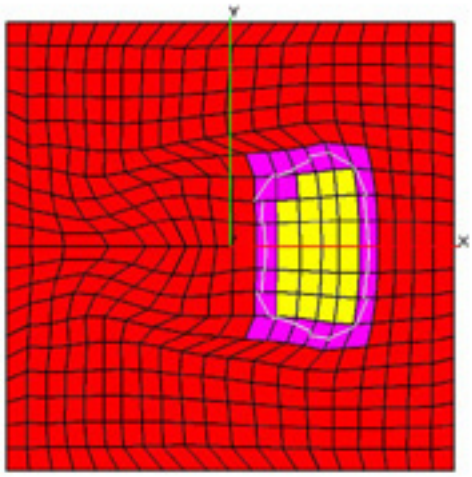

Fig. 1a. Old Remap Method

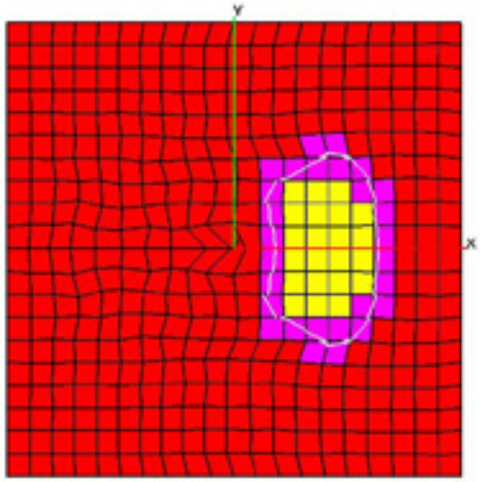

Fig. 1b. Intersection-based Remap

Figures 2a and 2b contrast the old vs the new remap. The corner coupling in the new method better maintains the material shape when moving diagonally across the mesh.

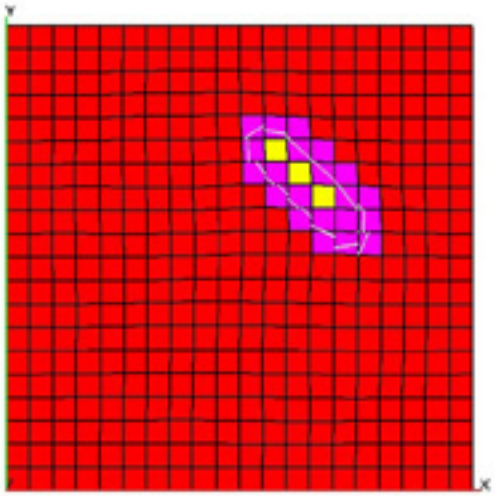

Fig. 2a. Old Remap Method

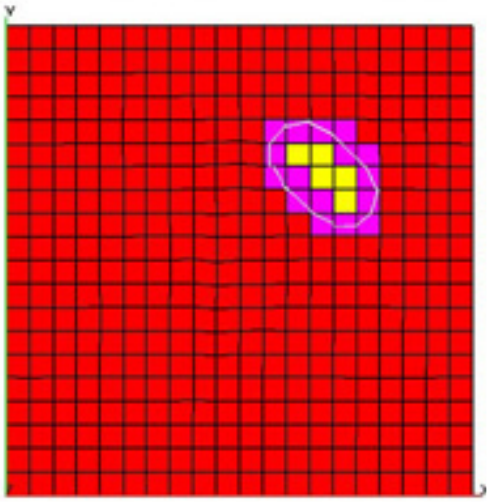

Fig. 2b. Intersection-based Remap

Progress on Simulating Wire-Array Z Pinch Implosions

A Sandia team has implemented a mass inflow boundary condition into the ALEGRA High Energy Density Physics code, which has allowed researchers to make significant progress in the study of Z pinches. A wire-array Z pinch is an annular array of tiny metal wires (~10 micron diameter) which, when pulsed with current, is subject to a magnetic force that implodes the pinch. The resulting stagnation and thermalization of the pinch on the cylindrical axis of symmetry yields a copious amount of x-rays, thus making the Z pinch an attractive radiation source. The first 50 to 80% of the Z pinch lifetime is described by a mass ablation phase, during which stationary wire cores cook off material that is subsequently swept towards axis by the magnetic field, thus resulting in a radial redistribution of mass. The ablation phase is difficult to simulate because it is both fully 3-dimensional (3D) in nature and requires very high resolution to model the small diameter wires.

The new mass inflow boundary condition models the ablation phase as a boundary condition, thereby allowing researchers to focus on the ensuing implosion phase. Z pinch simulations using the mass inflow model have produced the best comparison to date with experimental x-ray powers and radiographic images. Also, significant qualitative understanding has been made regarding the importance of 3D effects on the implosion.

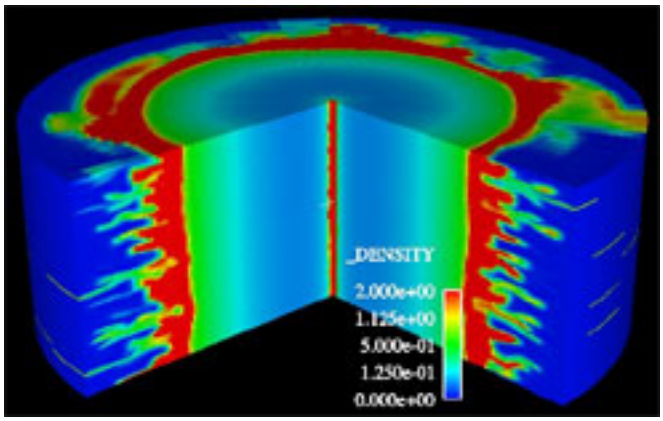

An example simulation of a Z pinch implosion using the mass inflow model, illustrating the highly 3D nature of the pinch.

Moab Rollout at Sandia

The Moab workload manager from Cluster Resources, Inc., was adopted by ASC as a Tri-Lab standard job scheduler last August. The job scheduler manages user's compute jobs, allocates resources, and gathers accounting data. Users would have a common set of job management tools, systems administrators would be able to share expertise, and management would gain a common set of reporting metrics. The initial goal of the Tri-Lab Workload Management Board was to ensure a uniform job scheduling environment across all new ASC platforms. However, with a site-wide license in place, all three labs are moving rapidly to deploy the Moab scheduler on existing systems too.

In August 2007, Sandia completed the rollout of Moab on Red Storm, Red RoSE, Black RoSE, and all of Sandia's HPC capacity clusters, well ahead of schedule. Moab provides new capabilities for Red Storm, including the ability to request specific memory configurations for compute jobs, and the ability to effectively use backfill algorithms to run smaller, shorter jobs when nodes are available. These new capabilities have greatly improved the system management in support of the peak workload associated with the end of the fiscal year.

Tri-Labs Workload Management Board met at Sandia in July. During that meeting, the team agreed on near-term goals for FY08 based on having Moab running on the ASC platforms, including common usage reporting across ASC platforms, common management configurations, and a shared Web site on the classified network that will host Moab documentation.

LANL Roadrunner Base Capacity System Running ASC Codes

The Roadrunner project completed an ASC Level 2 milestone in June 2007, a demonstration of Roadrunner's integration into the LANL classified computing system and success in running ASC codes. These milestone criteria were outlined in the FY07 ASC Implementation Plan. The system had been accredited for classified computing in May 2007.

Following the completion of this milestone, on schedule, the Roadrunner system delivered over 10 million CPU hours in July and August 2007 to the Stockpile Stewardship Program, running five application codes and system checks. LANL classified computer resources, including the Roadrunner base capacity resource, are also computing an important Directed Stockpile Work (DSW) study.

The Roadrunner project will formally transition the base capacity system to LANL HPC operations in September 2007.

LANL Roadrunner Procurement Team Wins Small Team Distinguished Performance Award

LANL's Director Anastasio announced that the Roadrunner procurement team had been honored with a Small Team Distinguished Performance Award. Congratulations to team members Manuel Vigil and Ray Miller, High-Performance Computing-Division Office (HPC-DO); Diana Little, Computing Operations and Support (HPC-2); and Bart Burson, Acquisition Services Management-Purchasing (ASM-PUR). This team provided the leadership for the technical planning, procurement, and project management for the Roadrunner Platform.

Director Anastasio presenting an individual distinguished performance award last year.

Sandia CSRI Workshop on Mathematical Methods for Verification and Validation

The Computer Science Research Institute (CSRI) of Sandia National Laboratories held a very successful workshop on Mathematical Methods for Verification and Validation (V&V), from August 14 to 16, 2007, at the Hyatt Regency Tamaya Resort. The workshop was chaired by Clayton Webster, the FY08 John von Neumann Fellow, and co-organized by fellow Sandians Scott Collis, Tim Trucano, and David Womble, as well as Prof. Max Gunzburger from the School of Computational Science at Florida State University. The conference Web site can be found at http://www.cs.sandia.gov/CSRI/Workshops/2007/MMVV/.

It is well recognized in the both the academic and laboratory communities that V&V must be an essential component of our research efforts if we expect to amplify our future predictive capabilities and therefore, the focus of this meeting was to emphasize the technical content that is required for successful and consequential V&V. To accomplish this goal, a diverse group of 25 university and 30 NNSA laboratory researchers gathered to consider foundational mathematical, statistical, and computational methods for V&V in complex application areas for predictive computational science. Talks surveyed important themes, including: stochastic sampling, stochastic differential equations (sdes), uncertainty characterization, error estimation, design of computer experiments and reliability of computational science, as well as open research problems in the topics discussed. In addition to talks discussing methods for V&V, we also organized talks aimed at some broad areas of applications where V&V methods have had and will have a significant impact. These topics included astrophysics, hydrology, chemical reactions, and the Global Nuclear Energy Partnership (GNEP). In total, there were 10 invited one-hour lectures and 12 short "rapid fire" talks. A copy of the program and a downloaded version of each presentation can be found at http://www.cs.sandia.gov/CSRI/Workshops/2007/MMVV/program.html.

A key feature of the workshop was the structured twice-daily discussion sessions. These sessions facilitated interactions among methods experts and applications scientists and proved very beneficial to all those attending. Future directions that methods development should take to be effective for V&V were debated. A consensus emerged that several developments in Uncertainty Quantfication have proved to be very promising and can be specificially beneficial to V&V. The workshop also held a discussion session on the training of scientists in V&V. Serious training is an important issue, both for those who want

to specialize in the development of specific V&V methods and also for the general computational science community that want to understand and incorporate V&V methodologies in their work. This session addressed issues such as what courses students need to take to be V&V savvy, the creation of new courses to train students in V&V, and practical mechanisms to give students experience in V&V in realistic settings.

An overview of the presentations and a summary of the discussion sessions will be written in a forthcoming workshop summary white paper by Max Gunzburger, Tim Trucano, and Clayton Webster.

Left to right: Pavel Bochev (Sandia National Laboratories), Max Gunzburger (Florida State University), Qiang Du (Penn State University), Clayton Webster (Sandia National Laboratories), John Burkardt (Virginia Tech University), Yanzhao Cao (Florida A&M University)

Recent Publication Highlights Science-Based Prediction at LANL

"Science-Based Prediction at LANL," an article about how LANL scientists have taken on the challenge of building computational tools to examine complex physical processes by using innovations that extend beyond the classic scientific method, was published in the *SciDAC Review*, Issue 4, Summer 2007, p. 33, http://www.scidacreview.org/0702/html/hardware.html.

ASC Salutes

Dr. Michael (Mike) Heroux is a Distinguished Member of the Technical Staff in the Scalable Algorithms Department at Sandia National Laboratories. His research interests include the solution of challenging engineering and science problems, especially the solution of large-scale implicit problems and the development of robust preconditioners. His contributions in the past nine years have led to breakthroughs in large-scale circuit modeling, non-Newtonian fluid simulation, and classical density functional theories. As part of this work, Dr. Heroux has also been the primary developer of numerous linear algebra packages that are used by hundreds of other solver and application developers and are available as part of the Trilinos project. These tools, especially the Epetra package of scalable linear algebra classes, have been used within the ASC Program to develop scalable solvers across every major Sandia ASC-funded application.

Dr. Heroux leads the Trilinos project as principal architect and lead developer. Trilinos is an effort to develop and implement robust enabling software using modern object-oriented software design and software engineering processes and tools, while still leveraging the value of established libraries. Trilinos is an R&D100 award recipient and is increasingly the preferred software framework for hundreds of application and solver developers throughout the world. In the coming year, the scope of Trilinos is expanding to include access to partitioning and load balancing tools, advanced discretization algorithms, and more, with the goal to provide a rich vertical software stack upon which applications can be quickly developed. Mike has led the Trilinos project from the very beginning, expanding it to the present effort involving dozens of computational scientists and applications.

Dr. Heroux is an expert in large-scale scientific and engineering computing. In addition to his interests in algorithm research and development, he plays a key role in understanding and improving application performance on high-performance computing (HPC) systems and is keenly interested in software engineering principles applied to engineering and science applications. He is actively involved in the HPC research community as a member of the Society for Industrial and Applied Mathematics and the Association for Computing Machinery. He has published extensively on HPC topics including the recent book *Parallel Processing for Scientific Computing*.

Prior to joining Sandia in 1998, Dr. Heroux worked for Cray Research, Inc. and Silicon Graphics, Inc. (SGI), developing mathematical libraries and applications for Cray and SGI systems. As the applications representative for new systems design, he was also actively involved in collaboration with systems designers to determine application performance needs on future systems, including the T90, J90, T3E, and SV1 systems.

Upcoming Events

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

ASC Web Site

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

Newsletter Points of Contact

Send submittals to:

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

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

Reeta Garber—Sandia National Laboratories ragarbe@sandia.gov

Who's Who

ASC Program Managers—Headquarters

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

Dimitri Kusnezov—Dimitri.Kusnezov@nnsa.doe.gov
Bob Meisner—bob.meisner@nnsa.doe.gov
Njema Frazier—njema.frazier@nnsa.doe.gov
Thuc Hoang—thuc.hoang@nnsa.doe.gov
Sander Lee—sander.lee@nnsa.doe.gov
Ed Lewis—edgar.lewis@nnsa.doe.gov
Erich Rummel—erich.rummel@nnsa.doe.gov
April Commodore—april.commodore@nnsa.doe.gov
Karen Pao—karen.pao@nnsa.doe.gov
Denise Peoples—denise.peoples@nnsa.doe.gov

ASC Program Managers—Labs

Los Alamos National Laboratory

John Hopson—jhopson@lanl.gov

Cheryl Wampler—clw@lanl.gov

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

Lawrence Livermore National Laboratory

Michel McCoy—mccoy2@llnl.gov

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

Sandia National Laboratories

Tom Bickel—tbickel@sandia.gov

Martin Pilch, QMU and Management Support—mpilch@sandia.gov

Paul Yarrington, Weapon Simulation & Computing—pyarrin@sandia.gov

SAND 2006-6930W

Developed and maintained by Sandia National Laboratories for NA.114

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