UCRL-ID-153222

# Advanced Simulation and Computing Program

Lawrence Livermore National Laboratory

***A Summary Report***
*to the*
***Director's Review Committee***

Michel G. McCoy

Acting ASCI Program Leader
Defense and Nuclear Technologies Directorate
Lawrence Livermore National Laboratory

June 2003

DISCLAIMER



This work was performed under the auspices of the U.S. Department of Energy by University of California Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

# Table of Contents

# Overview of the ASCI Program at LLNL

*Michel G. McCoy*
*Acting ASCI Program Leader at LLNL, Defense and Nuclear Technologies Directorate*

## Introduction

It has now been three years since the Advanced Simulation and Computing Program (ASCI), as managed by Defense and Nuclear Technologies (DNT) Directorate, has been reviewed by this Director's Review Committee (DRC). Since that time, there has been considerable progress for all components of the ASCI Program, and these developments will be highlighted in this document and in the presentations planned for June 9 and 10, 2003. There have also been some name changes. Today, the Program is called "Advanced Simulation and Computing," Although it retains the familiar acronym ASCI, the *initiative* nature of the effort has given way to sustained services as an integral part of the Stockpile Stewardship Program (SSP). All computing efforts at LLNL and the other two Defense Program (DP) laboratories are funded and managed under ASCI. This includes the so-called legacy codes, which remain essential tools in stockpile stewardship.

The contract between the Department of Energy (DOE) and the University of California (UC) specifies an independent appraisal of Directorate technical work and programmatic management. Such represents the work of this DNT Review Committee. Beginning this year, the Laboratory is implementing a new review system. This process was negotiated between UC, the National Nuclear Security Administration (NNSA), and the Laboratory Directors. Central to this approach are eight performance objectives that focus on key programmatic and administrative goals. Associated with each of these objectives are a number of performance measures to more clearly characterize the attainment of the objectives. Each performance measure has a lead directorate and one or more contributing directorates. Each measure has an evaluation plan and has identified expected documentation to be included in the "Assessment File." Assessment files generally contain a variety of reviews, and in particular available DRC input.

Of the 39 performance measures, two are of particular interest for this review.

1. *Measure 2.3—Demonstrate the ASCI simulation and modeling capabilities that support the ongoing needs of stockpile assessment and certification*. Associated with this measure are three metrics.
   a. Quality of improvement in predictive capability via increasingly complex simulations;
   b. Meeting Milestones including 2-D Primary Burn Code, 3-D Safety Simulation explosive initiation and application of advanced code to secondary performance;
   c. Report status of 3-D code development.
2. *Measure 6.3—Develop and implement, with NNSA & other DOE programs, plans to support optimal use of scientific research and test facilities & capabilities (e.g., NIF, DAHRT, Terascale Computing Facilities, LANSCE) at both Laboratories*. Associated with these deliverables are metrics
   a. Optimize the facilities for terascale computing
   b. Availability of White
   c. Progress on the Purple 100 teraOPS/s contract
   d. Progress toward a "desktop-to-teraflop" environment.

In the past, this Committee has focused very heavily on what would be considered covered under Measure 2.3. However, in this new model, the Committee will also be asked to review, along with the Computation Director's Review Committee, progress in the development of terascale facilities.

## The Burn Code Review Committee

Throughout the presentations and the articles in this summary report, reference is made to various *Burn Code Reviews*. The ASCI Burn Code Review Panel (BCRP) is an independent body chartered by LLNL and LANL to review progress on the high level nuclear application milestones. The BCRP first met to review progress toward the CY99 3D Primary Burn Milestone and has continued to meet and report its findings twice a year.

The panel is currently chaired by Professor Kim Molvig, MIT. The BCRP external members are knowledgeable about nuclear weapons technology and high-end computing. In addition, review members from the three national laboratories offer an element of peer review. Several members of the NNSA staff also participate in the review.

For the January 2003 review, there were nine reviewers from LLNL, LANL, and Sandia plus four reviewers from DOE/NNSA. The external members of the January 2003 BCRP were

Marvin L. Adams, Texas A&M
James M. Ferguson, LLNL retiree (Simulation)
Kaye D. Lathrop, LANL retiree (Simulation)
William R. Martin, University of Michigan
Paul C. Messina, Cal Tech (Emeritus)
Kim Molvig, Chair, MIT
D. Reid Worlton, LANL retiree (Primary Design)
Robert F. Lucas, ISI, USC
Jill Dahlburg, General Atomics

## Overview of ASCI Program Components

The ASCI Program is composed of five components. The first is Defense Applications and Modeling (DAM), which develops applications necessary to achieve the ASCI goal of providing the means to assess and certify the safety, performance and reliability of nuclear weapons and their components. The second is Integrated Computer Systems (ICS), which provides the computing platforms and centers necessary to field these. The third is Simulation and Computer Science (S&CS), which provides the infrastructure necessary to connect the applications and platforms into an integrated system. The final two components are not covered in this Review, namely the university partnerships and program management and integration.

The budget is broken down in such a way as to maintain balance between elements. For FY03, the ASCI budget breakdown is shown in the figure below.

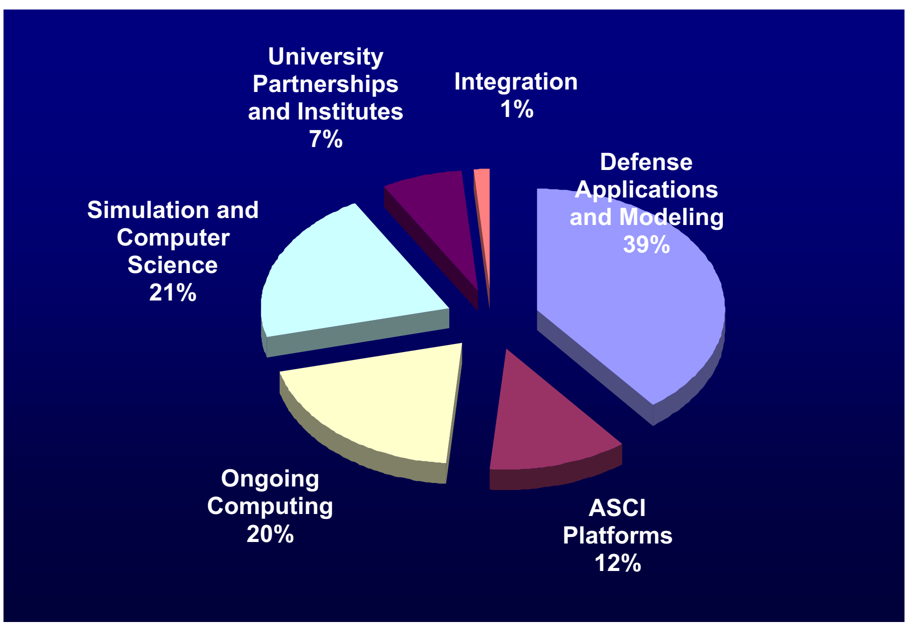


**ASCI's integrated development of high fidelity 3D weapons simulations requires many program elements.**

Here we have divided the ICS budget into two pieces: Platforms and Ongoing Computing. The ASCI budget for FY03, following the long continuing resolution and the rescission was $621M. This represented a cut of $41M from the FY03 Future Years Nuclear Security Program (FYNSP) guidance. The Program mid-year at the laboratories could not absorb such a large cut by cutting staff or local hardware procurements. Los Alamos consequently offered and the ASCI executives and Headquarters agreed, to forsake the delivery of the third sector of the ASCI “Q” machine, leaving the system at 20 teraOPS total. While this took the immediate pressure off of the budgets, it also exacerbated an existing capacity problem for the three laboratories. This issue is being addressed in part by the procurement of some low-cost capacity clusters at Los Alamos, and by LLNL through strategically chosen and well-negotiated early delivery systems within the 100-teraOPS ASCI Purple contract. However, long term the need for cycles, both for capacity and for capability is significant.

It is important to note that, while the Platform slice is markedly reduced by the $41M cut absorbed in FY03, platforms are never the dominant component of the ASCI effort. People are. Typically, 70% of the budget goes to people, 22% to hardware, either in platforms or other support infrastructure, and 8.5% to contracts, including university partnerships.

## ASCI in Transition

For at least the past year, a number of developments have acted to encourage a re-examination of the ASCI strategies and role of ASCI within the Program. Drivers for this include an increasing need for strong coupling between stockpile requirements and the ASCI effort, in particular, the size and number of platforms procured. A second driver comes primarily from outside of NNSA regarding the technological and business strategy selected by ASCI for its Platform acquisitions: Is it the right strategy to continue to focus heavily on riding the commodity curve or should ASCI work towards “building a computer for science?” Both areas are subject to reviews this year, the former by the JASONS and the latter by a panel of the National Academy of Sciences, as requested by Senate language in the FY03 funding legislation.

The former concern, regarding the coupling between requirements and ASCI, is certainly the deeper question. That the question is even being asked has much to do with the nature of ASCI during its initial years—its strong growth period. ASCI was born as an initiative, the intention being to build the new set of 3D, high fidelity applications necessary to manage stockpile stewardship in the absence of nuclear tests in Nevada. In order that these applications would be useful, it was necessary to site a commensurate and balanced computing infrastructure matched to the requirements coming from these codes as they evolved. The intention was that by mid-decade the complex would have the demonstrated the capability to model a full system at fidelity sufficient to resolve the “known unknowns” otherwise known as an entry-level, high fidelity, full system calculation. In particular, such a calculation, running on a 100-teraOPS computer for about two months, would be sufficiently resolved that the errors arising from numerical discretization could be largely distinguished from those emanating from outstanding physics issues. Today, we believe that in 2D such an outcome will certainly be the case. In 3D, with the availability of highly resolved 2D resolution studies for reference, estimates can be made regarding the magnitude of the numerical error. The 3D results will help as well to determine the optimal size for future capability platforms, taking account three issues: (1) any remaining resolution questions, (2) improved physics as needed and (3) time to solution. The latter, for the designer, is an important consideration if a 3D calculation at this scale is to be a routine part of the workload. This emerging 3D capability represents a major accomplishment for the Program as it is organized and managed at LLNL. The work until today has *generated and demonstrated a capability*. The future has to stress *integrating this capability* into the program and continuously improving the models and infrastructure so that the tools are easily accessible and used routinely.

Today, in response to the growing maturity of the applications and to pressing stockpile work, scientists demand to use the computers and the codes to support stockpile deliverables. This is in addition to using resources to meet milestones demonstrating growing functionality in the codes themselves. Tying ASCI level-1 milestones to stockpile work was a good first step at addressing these two competing factors. Much of the work described in this document and in the associated presentations is related to the successful attainment of these difficult milestones. The table shows the existing ASCI level-1 milestones as formulated in FY02.

However, the access issues go well beyond the major 3D calculations. These new codes are being applied aggressively to the W80 stockpile life extension program (SLEP), to campaigns and to significant findings investigations (SFIs) in a capacity mode–that is to say, an individual calculation shares the computer with other calculations running concurrently. As ASCI becomes increasingly integral to the attainment of SSP goals, it is important that the appropriate balance between capacity and capability infrastructure reside on computer room machine floors. Insufficient cycles to meet demand for capacity calculations will result in the capability resource being used exclusively for the smaller calculations, forgoing the insight and improvement in assessment possible from routine use of higher dimensional calculations.

There is yet another important development. The Stockpile Stewardship Program itself is undergoing a planning process to pull together integrated deliverables through 2009. While this evolution is not the subject of this Review, ASCI has participated in the development of a new set of SSP Level-1 milestones critical to the entire Program. The ASCI planning process for future lower-level ASCI milestones and deliverables will feed directly into many of these level-1 milestones. So the process of integration into the SSP, or the evolution from *initiative to program,* is irreversibly under way. The issue for ASCI is to maintain world leadership in computing *while* employing technical strategies that are of maximum value to the broadest possible programmatic goals and *while* planning a path to the future in which use of 3D applications is a routine part of the workload. Some of these strategies will be described in the sections and talks devoted to platforms and terascale infrastructure. Much progress has been made, and the story at LLNL is powerful, representing the confluence of three very successful local efforts: applications development, strategically chosen science, and terascale platforms within a balanced computing infrastructure.

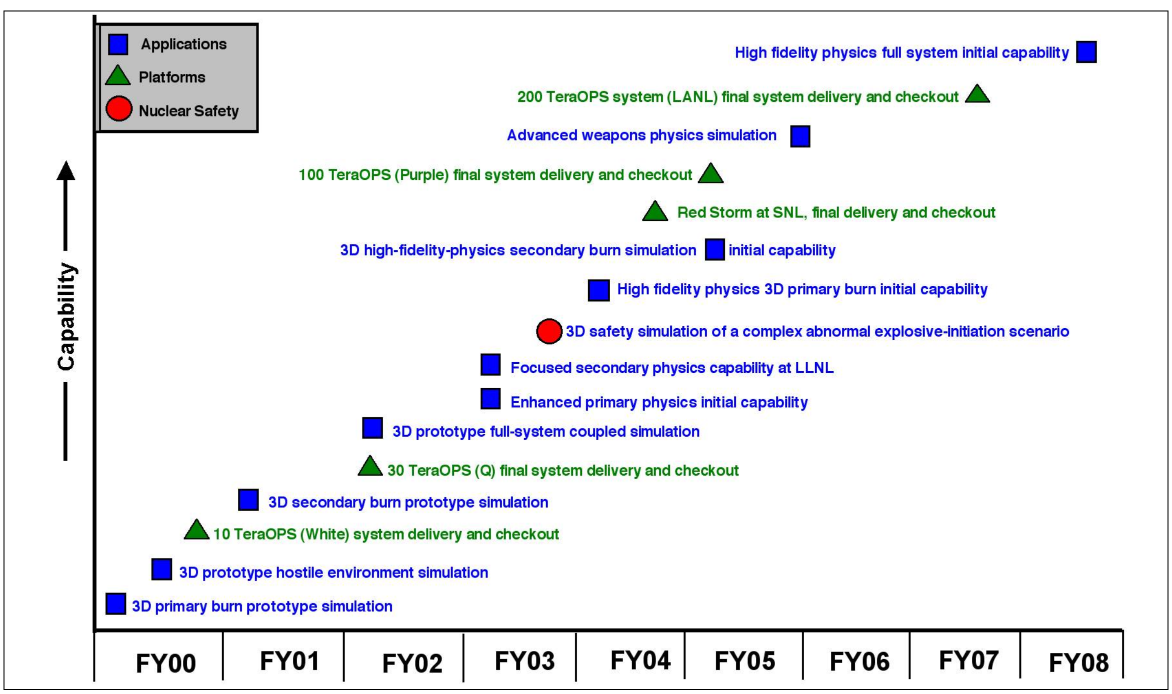


**ASCI is delivering critical capability on schedule.**

## Major Accomplishments in CY02

In this section, we have selected ten major ASCI accomplishments during the reporting period. Most of these will be described in detail in the Review itself, where classified discussion is possible, and also to the extent possible in the succeeding sections of this unclassified document. These accomplishments follow the order of the presentations.

### 1. Terascale Facilities

The LLNL ASCI classified and unclassified computing facility has consistently demonstrated high levels of availability and customer service to the Tri-Lab community. Three major ASCI milestone calculations (one from Sandia and two from LLNL) were completed on the LLNL flagship system—ASCI White-in 2002. The White system continues to be used heavily by all three laboratories. We have received strong positive feedback for our model for Tri-Lab coordination and management that we established for White. LANL is now using this model as a basis for developing its ASCI Q system. White and the other ASCI systems at LLNL continue to show high utilization due to reliable hardware and software and the scheduling that encourages maximum throughput for simulation runs. We have achieved our performance goals for the networking and storage environment for White. In addition, we have helped improve the bandwidth of the classified wide-area network, which transports data among the Tri-Labs, demonstrating that the distance-computing model is delivering for the Program.

### 2. Purple and BlueGene/L (BG/L) Contract

At *Supercomputing 2002* in Baltimore (the premier high performance computing conference in the US), the Secretary of Energy announced the ASCI Purple Contract, a partnership between IBM and NNSA. This represented the culmination of more than two and one-half years of work. The contract provides options (1) to procure either a 100-teraOPS or a 60-teraOPS system for classified weapons work and (2) to procure an unclassified 360-teraOPS computational science research platform (BG/L) both for stockpile stewardship science, and (3) to begin the work of scaling applications to the 30,000 to 50,000 processor level necessary to achieve PF/s (petaFLOPS/s) computing levels by 2010. The systems are scheduled to be sited in December 2004, in the first computer room of the newly constructed Terascale Simulation Facility (TSF). There are a number of associated systems being delivered under this Contract that in themselves total about 30 teraOPS of computational capability. These systems are intentionally dual-use, as development platforms for Purple and to provide much needed capacity relief for the program. LLNL is (for the third time) on track to deliver state of the art platforms to the program.

### 3. CY02 3D Secondary Performance Milestone

A-Program's three-dimensional massively parallel design code was conceived from the beginning not only to provide the required physics algorithms and models but also to explore the potential advantages to our scientific research and development environment of modern computer science techniques. The ASCI Burn Code Review Committee (BCRP) recently validated the results for the latest ASCI milestone. They agreed that the simulations presented met or exceeded all of the objectives laid out for this milestone by demonstrating the code's ability to simulate, in three dimensions, relevant experiments, obtaining results that compare within reasonable bounds to both experimental data and the output from simulations performed with other well-established codes. This ASCI milestone represents a large portion of the deliverables associated with Appendix F Performance Measure 2.3. In addition to meeting ASCI milestones, this code is beginning to provide important programmatic results for our directed stockpile work (DSW).

## 4. CY03 Nuclear Safety Milestone

Progress toward the CY03 Nuclear Safety Milestone, a "Nuclear Safety Simulation of a Complex Abnormal Initiation Scenario," was considered by the BCRP at their January 2003 meeting. They reported: "B Division is also on track to complete the CY03 Nuclear Safety Milestone. The Panel is impressed by the capability demonstrated in this Project to simulate nuclear weapon safety. We believe it is highly likely that this Milestone will be met." Subsequent to that review, the calculations required to accomplish this milestone, a series of 3D simulations of a complex safety-related underground nuclear test, were completed prior to the March 31, 2003, deadline. The Nuclear Safety Milestone calculations were run on the ASCI IBM White machine. Execution times ranged up to 100 hours on as many as 800 processors. These results are to be presented to the BCRP at its next meeting in August 2003. Completion of this milestone is an important step in deploying simulation capabilities for use by weapons scientists and engineers, in conjunction with data from relevant experiments, to assure the continued safety, reliability, and performance of the US nuclear stockpile.

## 5. CY02 Enhanced Primary Physics Milestone

The achievement of the CY02 Enhanced Primary Physics Initial Capability was also affirmed by the ASCI BCRP in its meeting in January 2003. Their report states: "The panel unanimously agreed that the Livermore Code group met and exceeded the CY02 milestone. The team performed high-resolution 2D simulations of several NTS events, and presented detailed comparisons with measurements. These exercises demonstrated the ability to perform simulations far exceeding legacy code capabilities both in resolution and in model content." This milestone consisted of a suite of high-resolution 2D primary burn calculations for several NTS events, run with another code on the ASCI White computer at LLNL. The calculations were compared with data from those tests and related nonnuclear experiments. These calculations contributed to several current stockpile activities. The CY02 Enhanced Primary Physics Milestone also included a formal code release to the B Program design community.

## 6. Parallel Explicit Nonlinear Stress Analysis Simulation

The ParaDyn code team has built a parallel application directly upon the architecture of the serial DYNA3D code, which remains the archetypal example of an explicit nonlinear stress analysis code. ParaDyn is the workhorse stress analysis tool for engineering DSW at LLNL and LANL. This effort has earned the endorsement from peers at LANL in the sincerest manner by their adoption of ParaDyn as their main production code. During the Engineering Directorate's own recent Director's Review Committee visit, John J. Ruminer, LANL's Deputy Division Leader for Engineering Sciences and Applications, spoke extemporaneously as to their appreciation for the code capabilities and user support they receive from LLNL. The major ParaDyn effort in recent years has been the implementation of all contact algorithms previously available in DYNA3D. Analysts at LANL recently benchmarked the performance of ParaDyn on ASCI White for a model of a re-entry vehicle (RV) forward mount. ParaDyn was shown to be extremely efficient in comparison to competing codes. We have taken the same model, executed it on 128 CPUs using the hardware performance monitoring facilities and determined it sustains 12% of the theoretical peak floating-point performance. Simpler models attain higher rates, but this benchmark embodies the characteristics of day-to-day engineering analysis in support of DSW.

## 7. ASCI Verification and Validation Program

The credibility of computational simulations depends on the degree to which the capability has been demonstrated through quantitative Verification and Validation (V&V) evaluations.

These involve comparisons of simulation results to known analytic/semi-analytic solutions (Verification), to other, well-validated codes, and high quality data (Validation). Among the significant recent accomplishments can be included: (1) completion of numerous documented V&V analyses, (2) development and application of weapon relevant semi-analytic verification problems, and (3) first round reliability and confidence quantification (application of V&V methodology) derived from model/data uncertainties for stockpile systems assessments. This method was also applied during W76 1X-109 Peer Review. Development of Uncertainty Quantification (UQ) approaches contributes directly to quantification of the quality of ASCI simulation and modeling capabilities. Verification and validation analyses performed this year have contributed directly to the evaluation of simulation quality and sufficiency when they are applied to important stockpile issues, including for the W76 1X-109 and SFI Peer Reviews, W80 LEP, SFI resolution and weaponization analyses for W87 and B83 issues, W88 and B61 Peer Review activities.

## 8. Equation of State and Dynamics of Metals Research

Materials properties efforts are broadly based and aggressively studied at LLNL given their direct interest to the stockpile. Highlighted here are just two of several major recent advances. A major effort within Material Properties Research is the Equation of State (EOS) project that seeks to use advanced quantum based theoretical and computational methods to develop next-generation multiphase EOS. Among the significant accomplishments is included a new correlated treatment of f-electrons which has been successfully applied to the volume collapse transition in the prototype metal Ce. We have developed and tested a first-generation dynamical mean field theory (DMFT) code and successfully applied this code to the analogue alpha-gamma volume collapse transition in Ce as well as performing initial parallel calculations for plutonium. Second, the dynamics of Metals Program is an aggressive multi-scale modeling program whose goal is to develop predictive models of strength, plasticity and failure in metals. Most significant among the successes this year was (1) the first linking of all the length scales for a Ta simulation, with dislocation mobilities developed on the atomistic scale, (2) crystal yield points developed using the mobilities in dislocation codes in microscale work, and (3) homogeneous polycrystalline yield surfaces developed using a mesoscale simulation of polycrystals and the crystal yield surfaces and elasticity models. The result was a texture-sensitive polycrystalline strength model at the continuum level.

## 9. High Explosive Modeling

Notable progress has been made both in high explosive (HE) macroscopic and grain-scale models. In the macroscopic domain, CHEETAH, a thermochemistry code, has been linked with a chemical kinetic framework, to the ARES hydrodynamic code. New algorithms have been developed that make the computations affordable on current computers. For the first time, individual chemical species can be tracked through the detonation process. This also means that EOS measurements on individual products of detonation, such as water, can be related to the EOS of the detonation product mixture in a scientifically rigorous way. For grain-scale modeling, ASCI-class computational resources have allowed calculation on the size scale of a representative volume element, which corresponds roughly to a single zone in a hydrodynamic calculation. Form factors for reaction growth have been derived and implemented in a successful macroscopic model. An external review of the most recent grain-scale simulations was conducted in January 2003. The review panel report concluded that LLNL has made significant progress in understanding high explosive behavior at the grain scale, and that the team met all milestone requirements associated with grain-scale modeling.

## 10. Turbulence

Turbulence modeling plays a major role in the Program to account for mixing effects. Using the MIRANDA code to model the University of Arizona experiment on impulsively accelerated interfaces, LLNL uncovered a mixing mechanism driven by generation of secondary vorticity from the interaction of the complex density structure and the pressure field in the flow—basically a centrifugal version of the Raleigh-Taylor instability. A video describing this work was honored as a winning entry in the 20th Annual Gallery of Fluid Motion at the *55th Annual Meeting of the APS, Division of Fluid Dynamics*, Nov. 24–26 2002. In addition, a very large (720 × 720 × 1620) Raleigh-Taylor simulation in progress (at this writing) is providing some surprises, particularly by exhibiting a linear growth rate of the mixing layer (as opposed to quadratic, which has been the traditional description). While firm conclusions await the results of follow-on work, it is apparent that issues of initial conditions and assumptions about the scaling of flow need to be critically re-examined, and that this work will result in a significant change in our understanding of the Raleigh-Taylor instability, one way or the other.

# Terascale Facilities: The Livermore Computing Program

*Terri Quinn*
*Deputy Department Head, Integrated Computing and Communications Department*

## Introduction and Background

The LLNL ASCI high-performance computing enterprise—Livermore Computing (LC)—provides the computational capabilities for the program with computing services, computer science research, and computational science research. The LC program funds ~170 FTEs and stewards two large computing centers (one classified and one unclassified), and the Institute for Terascale Science, which is a computer and computational science research center. This research center will not be a focus of this talk.

During the seven-year life of ASCI, the LC has become world-class, and, in some interesting sub-areas, one could argue that the Laboratory is in a class by itself. We have deployed some of the largest computers in the world and shaped them into valuable tools for science. Our computers have high-availability with a full-featured software environment. The storage, visualization, and network architecture is as balanced as promising technology, decent budgets, and careful planning can ensure, and the user environment is recognized as being greatly enhanced by a robust customer service investment and culture. The LC is devoted to meeting the requirements of Tri-Lab stockpile stewardship programs at LLNL, Sandia, and Los Alamos National Laboratory (LANL). This enterprise has been similarly successful in meeting the multivariate science and technology requirements coming from programs at LLNL by leveraging ASCI developments.

The LC computing resources consist of 19 teraflop/s of classified computing and 24 teraflop/s of unclassified computing. Classified systems meet a variety of needs and include the ASCI large-scale parallel White and Blue–Pacific sustained stewardship teraflop (SST) systems. There are Linux-Intel clusters that are sized to accommodate the large number of small- to mid-sized (16- to 128-node) parallel jobs, which are now common as part of normal classified production workload—an extremely cost-effective system. We continue to re-balance these computing resources to meet the demands of stockpile stewardship work.

ASCI unclassified systems include the ASCI Linux Cluster (ALC), Frost, and Blue–Pacific. The ALC is an Intel–Linux system delivered on December 30, 2002, as part of the ASCI Purple contract. Initial use of ALC is as a development platform for the Open Source Lustre parallel file system that is a key strategic technology needed for LC. Later, ALC will be targeted for use by ASCI University Alliance users allowing the retirement of the aging Blue–Pacific system. Today, Frost is heavily used by ASCI University Alliances, as well as by ASCI scientists for unclassified programmatic work.

The support infrastructure for computing at LLNL includes networks, archival storage, shared file services, and visualization platforms. Our goal is to provide a balanced framework for simulation science, allowing users to generate, transport, visualize, understand, and safely protect their data. The ever-increasing capabilities of the computing platforms require that the infrastructure be scalable and state-of-the-art. Time spent generating data that cannot be navigated, visualized, or protected is wasted time. A balanced environment is thus an essential enabler for computationally based science at LLNL. We achieve this balance through a significant amount of planning based on the requirements for bandwidths and capacity of storage devices associated with the computing platforms, the archival storage systems, shared file systems, and the visualization platforms. These requirements are based on user input and historical data.

## Advances in LC's Balanced Computing Environment

A focus of this year has been to maintain the stability and availability of the ASCI platforms. Utilization of the ASCI systems including Frost, Blue, White, and SKY is extremely high, with each system being an average of 80 percent utilized (see chart below). Hardware and software reliability is high. With a mixture of technical and hand scheduling we have been able to achieve high throughput for the simulation runs. Three major ASCI application milestones were successfully completed as planned on White in CY02, one by Sandia and two by LLNL.

To completely integrate ASCI White and make it fully productive, we had to re-engineer the computing infrastructure, storage, networks, and visualization. Last year's planning called for tripling the bandwidth from the ASCI White platform to the archival storage system, which we achieved delivering over 800 MB per second to archival storage from White. This increase in bandwidth led to a capacity growth in the archive of twice the amount previously stored in all of LC's history, which, although anticipated, was still an impressive achievement. Storage demands continue to grow, particularly in the unclassified environment, which saw a near doubling of data. After an astonishing six-fold (6×) increase in classified bytes stored between 2000 and 2001, classified storage demand has leveled off in 2002. Since the 6× increase was clearly attributable to the availability of White, a major focus for storage in the upcoming year will be the preparations for the 100-TF ASCI Purple system in 2004.

Livermore Computing strives to provide a stable, usable, leading-edge parallel application development environment for our computers. A full-featured environment that respects standards and is essentially the same on each computer can significantly increase the productivity of applications developers. We call this the Livermore Model, and it has been a successful strategy in support of our application development teams for over 10 years. One goal of this strategy is to reduce the time needed to port applications to new computers. Through collaborations with vendors and other third-party software developers, we have ensured a complete environment, while steering their application development to exploit emerging technologies.

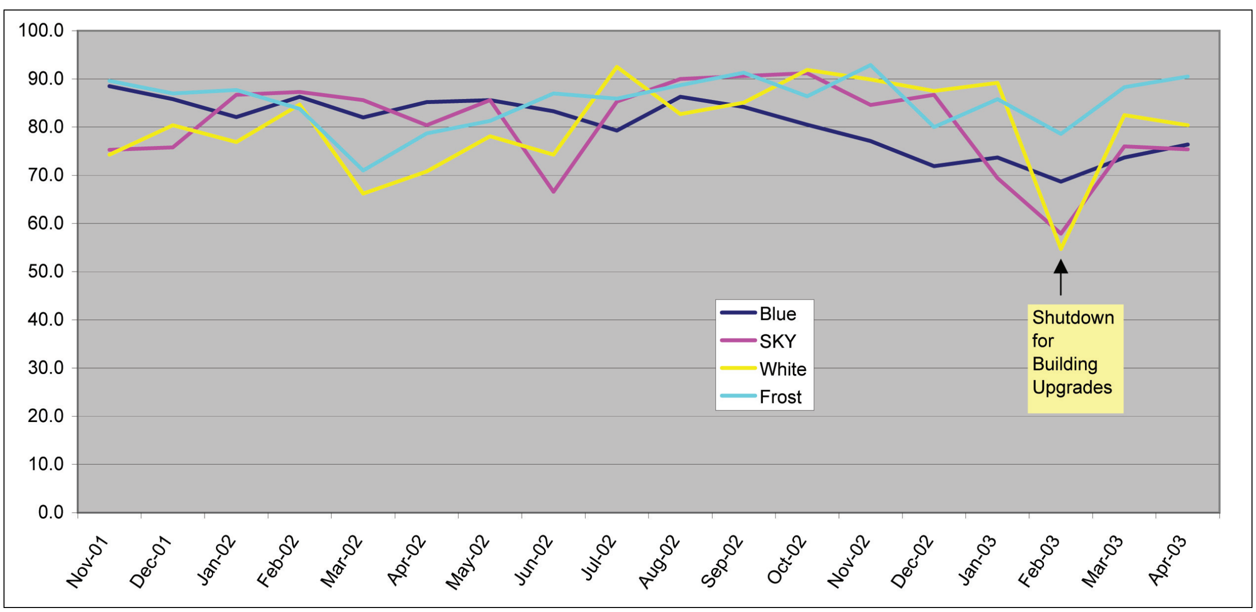


**Utilization of ASCI Platforms from November 2001 to April 2003.**

Livermore Computing supports three large visualization facilities called PowerWalls (two classified and one unclassified), which are used for formal reviews and presentations, and data analysis working groups. In addition, LC supports 22 weapon-designer and code-developer offices with direct video connections to the classified SGI visualization servers. To deliver visualizations to these offices and the facilities, we maintain a separate video network. This includes many fiber optic paths, electronic modems, and a variety of video signal switching systems. This video switching system and the SGIs themselves are aging and have not been as useful as we had planned. Weapons scientists want the ability to explore their datasets quickly with high-quality and high-resolution visualizations and have not met to the extent desired these requirements with the current technology. We lead an effort in ASCI to develop a distributed, parallel visualization (rendering) software solution. This multi-year Open-Source effort had its first release last fiscal year. This collection of software systems forms the backbone of the commodity cluster-based scalable rendering system and opened the door for us to deploy such a system in production. We have deployed our first Linux-based visualization cluster last year. The parallel visualization cluster (PVC) was delivered in September 2002, and represents a strategic platform for the continued research and development of our Linux-based visualization efforts.

The ASCI wide area network (WAN) is one part of the Tri-Lab computing infrastructure that was designed, deployed, and managed as a single entity by all three national laboratories. The WAN is a 2.4-gigabit classified network (250-MB payload) that interconnects the ASCI resources at the three laboratories. When initially put into operation in March 2001, the data movement applications achieved 100-megabyte throughput. In the past year, LLNL led the effort to optimize the WAN edge router, doubling the throughput to 200 MB. In addition, the network teams at the three laboratories increased reliability of the WAN and reduced the maintenance overhead by installing the highest-speed encryptors (OC-48 ATM NSA Type-1), which recently became available. We now routinely route 20 TB of classified data among the three laboratories each month, which only a few years ago was the size of the entire storage archive at LC.

## Tri-Lab ASCI Computing (White and Q)

Equitably and effectively sharing the ASCI platforms between the three laboratories is a basic premise of the ASCI program. Successfully achieving this goal required the laboratories to establish a Tri-Lab infrastructure. The Tri-Lab infrastructure at LLNL is representative of that at the other laboratories. It includes the platform machine, application environment, other local ASCI resources (e.g., storages, networks, and visualization servers), customer support services (e.g., hotline), and the operational support staff. The Tri-Lab infrastructure at LLNL is thoroughly integrated into the existing organization.

The assignment of the ASCI White computer as a 3-laboratory resource led to the formation of a new Tri-Lab coordination and management process (and culture), which is now in place, and which ensured successful completion of all three laboratories' major code team milestones. User requirements are gathered, and the corresponding effort needed to meet these requirements is documented in a process termed the "Usage Model." At LLNL our planning process ensures that these efforts and resultant capabilities are integrated into our production environment. The first Usage Model was developed by LLNL for White; LANL is now leading this effort since it is deploying the Q platform.

As part of the Tri-Lab coordination process, policies were defined for White resource allocations for each laboratory and scheduling policies. A set of well-defined job priority tiers was created to assist with enhanced priority run (EPR) scheduling decisions. Weekly EPR meetings have continued (via video teleconference) with representatives from all three laboratories. These meetings were run by LLNL and proved to be fundamental to the successful use of White and to completion of each calendar year's milestone calculations; they are also used to schedule time on LANL's ASCI Q machine.

Weekday coordination on White was assisted through the LC Hotline (our customer service entity). Weekend (and weeknight) coordination by LC's operations staff and was assisted through e-mail reflectors, with each laboratory specifying its jobs requirements for the weekend in tier-order, with complete contact information so that LC operations could talk to users about job status as needed. Tri-Lab management of White through these mechanisms worked so well that, with very few exceptions, it was unnecessary for the Tri-Lab White Access Policy Committee to convene to set new policy.

The ASCI Q machine, sited at LANL, is scheduled for general availability to the Tri-Lab user community shortly after meeting an ASCI Simulation and Computer Science (S&CS) milestone in October 2003. LLNL, Sandia, and LANL are jointly working towards this milestone that will demonstrate that the needed tools and user environment are available on Q. Preparations for running applications on the Q machine have been underway in the Tri-Lab community for over a year.

## Performance Measures of Quality

The presentation "Terascale Facilities: the Livermore Computing Program" covers many of the performance measures cited in section 6.3 of Appendix F, but not all. One not covered is the deployment of a unified trouble ticket system for the LLNL Stockpile Stewardship Program. We are presently deploying such a system to improve customer service. We have a due date of September 2003 to have a unified trouble ticket system in place for the DNT directorate and LC hotline. We have deployed this system for the NAI directorate and are nearing completion with a deployment for the NIF directorate. Although we are currently behind with our planned deployments, our re-worked plan still has us meeting our September 2003 deadline.

## Future Work

We are challenged to meet the continuing demands coming from scientists to access the datasets they have generated either computationally or experimentally and to have these sets close and understandable. While research in LC on managing large datasets is important, and progress is being made, it is also important to begin viewing the scientist's work area as an extension of the computer center itself—the "desktop-to-teraflop" environment. While the term "desktop-to-teraflop" rolls easily off the lips, its definition does not. What is certain is that achieving such a result means solving a myriad of seemingly minor, but devilishly confounding technical and social challenges, from the perspectives of security, data-transfer rates, connectivity, and customer service, including end-to-end diagnostic

We are in the second year of planning for the final Purple deliveries late 2005—the 100 teraflop/s system and BlueGene/L. This represents a 10-fold increase in compute power at the LC requiring a commensurate increase in the infrastructure. To meet this challenge, we are promoting careful planning, just-in-time procurements, exploitation of technology advances, collaborations with universities and industry, and phased integration schedules.

## ASCI at Livermore Platform Strategy

*Mark Seager*
*ASCI Platforms Principal Investigator*

ASCI Purple is the fifth-generation main line ASCI procurement. Predecessor procurements were ASCIs Red, Blue, White, and Q. The Purple procurement has resulted in a contract with IBM for the delivery of multiple platforms. Primary deliveries under this contract are *Purple C*, the 100-teraFLOP/s delivery based on IBM pSeries technology (Power5 microprocessors) and BlueGene/L, the 180/360 teraFLOP/s computational research engine based on IBM PowerPC system-on-a-chip technology. The Purple C is the fulfillment of the original ASCI platform roadmap to get to 100 teraFLOP/s by 2004. As such, Purple C will impact the Stockpile Stewardship Program (SSP) as the capability platform that can, for the first time, simulate both primary and secondary at high resolution.

Going from 20 teraFLOP/s with ASCI Q at LANL or 12.3 teraFLOP/s with ASCI White at LLNL to Purple C at 100 teraFLOP/s represents the largest jump in capability the program has attempted. Thus, several risk reduction strategies must be employed to increase the probability of success. Lessons learned from previous generation ASCI procurements indicate that introducing multiple technologies (e.g., node, interconnect, and software as was done with White) with the final delivery can lead to multiple layers of problems that lengthen the amount of time required to stabilize the system. Accordingly, the Purple contract is structured to include the Early Demonstration of Technology Vehicle (EDTV).

Purple's major risk reduction strategy, EDTV, is currently envisioned to deploy pSeries technology in the pEDTV machine with Federation interconnect and the first delivery of software targeted at Purple C. In addition, the SSP is currently experiencing an extreme capacity crunch. In order to respond to this pressing requirement, the second EDTV machine based on xSeries (IA-32 or IA-64 or Opteron microprocessors and Quadrics Elan3 or Elan4 interconnect) will be delivered as xEDTV. Taken together EDVT = pEDTV + xEDTV represent more than 17 teraFLOP/s of additional capacity to the program. This is more than ASCI White and only slightly less than ASCI Q. In addition, the ASCI Linux Cluster (ALC) machine was delivered in December 2002 as the file system development platform for BlueGene/L. At 9.2 teraFLOP/s, it represents a significant capability machine as well and is currently being utilized by ASCI science runs as a method for shaking down the Lustre Open Source global file system. See the table titled "Purple Platforms at a Glance," which shows the key characteristics of the platforms delivered under the ASCI Purple budget line.

The Purple C system is targeted for delivery in December 2004. The Purple contract with IBM is structured so that the final decision on the size of Purple must be made by the ASCI program in November 2003. The Purple contract has negotiated both the Purple C (100 teraFLOP/s) and Purple (60 teraFLOP/s) configurations, payment and delivery schedules. No contract modification is necessary to execute either Purple C or Purple.

BlueGene/L (BG/L) represents a new technology line for IBM. It is the result of collaboration between IBM Research and the ASCI Advanced Architectures Program, with LLNL as the lead laboratory. The results of the ASCI research and development (R&D) contract feed into the BlueGene/L build contract with IBM. Purple C and BlueGene/L are components of the same fixed-price contract with IBM, but have separate statements of work and payment schedules.

| Purple Platforms at a Glance | | | | | |
|---|---|---|---|---|---|
| **ALC** | **xEDTV** | **pEDTV** | **Purple C** | **BG/L** | **Attribute** |
| 9.2 | 5.2 | 12.3 | 100 | 360 | **Peak computational rate (TF/s)** |
| 3.84 | 1.54 | 4.0 | 50.0 | 16.512 | **Aggregate memory (TiB)** |
| 176 | 90 | 140 | 2,000 | 400 | **Aggregate global disk (TB)** |
| 3.072 | 1.64 | 5.12 | 133 | 360.5 | **Aggregate memory bandwidth (TB/s)** |
| 6.0 | 2.0 | 6.0 | 108 | 40 | **Delivered global I/O bandwidth to applications (GB/s)** |
| 0.680 | 0.768 | 2.05 | 12.6 | 137.6 | **Aggregate intra-SMP link bandwidth (TB/s)** |
| 115.2 | 69.12 | 37.27 | 223/443 | 0 | **Aggregate local disk (mirrored) capacity (TB)** |
| 38.4 | 19.20 | 25.6 | 47.28 | 0 | **Delivered aggregate local I/O bandwidth to applications (GB/s)** |
| 76x1Gb/s | 24x1 Gb/s | 24x1 Gb/s | 32x10 Gb/s | 1,024x1 Gb/s | **External Ethernet networking** |
| 1,920 | 768 | 2,048 | 12,608 | 131,072 | **Number of processors** |
| Pentium 4 Xeon | Pentium 4 Xeon | Power4+ | Power5 | PowerPC 440 SOC | **Microprocessor technology** |
| 0.357 | 0.143 | 0.591 | 7.5 | 3.09 | **Power required for computer and cooling** |
| Linux 2.4 | Linux 2.4 | AIX 5.2 | AIX 5.3 | Linux/HPK | **Operating System** |
| Intel, PGI | Intel, PGI | IBM XL | IBM XL | IBM XL | **Compilers F90, C, C++** |
| SLURM/DPCS | LURM/DPCS | oadLeveler/ DPCS | LoadLeveler/ DPCS | SLURM/DPCS | **Resource management** |
| Etnus TotalView | Etnus TotalView | Etnus TotalView | Etnus TotalView | Etnus TotalView | **Debugger** |
| December 2002 | December 2003 | December 2003 | Dec 2004 | Dec 2004 | **Delivery date** |

Because substantial results from the R&D contract (e.g., the BlueGene/L prototype) will occur after the Purple contract was executed, it was necessary to couch the requirements of the BlueGene/L statement of work as "targets," and provide both ASCI and IBM with the opportunity to review progress towards these targets and exit the contract without penalty if necessary. Review of the targets is based on a holistic approach to system functionality, rather than the approach of traditional statements of work.

The BlueGene/L payment schedule was structured to minimize ASCI financial exposure prior to these decision points. Two major decision points have been identified: (1) BlueGene/L GO/NOGO in late FY2003, and (2) BlueGene/L parts order in the first quarter of CY04. The GO/NOGO decision point will review the results of the R&D contract with emphasis on the results obtained from the BlueGene/L prototype. The BlueGene/L GO/NOGO will include an external review committee recommendation. Current thinking on the GO/NOGO decision point is that it may involve two steps, review of the R&D contract results and review of the prototype results. The parts-build decision point will be based in the price and availably of the required DDR-SDRAM chips and the yield, operating frequency and functionality of the BlueGene/L compute ASCI and IBM's ability to build the entire system and deliver a system that will be useful for computational science (not just computer science).

Culminating in a design of a computing system with a target peak of 180/360 teraFLOP/s[1], the development effort will build a 512-node prototype system in 2003 to verify critical technology components. The successful conclusion of this development effort will

[1] BlueGene/L is built of two-processor nodes; the processors are fully symmetric. In communications coprocessor mode, one processor is used for computation and one processor is dedicated to communications, with a target peak processing power of 183.5 teraFLOP/s. In virtual node mode, each processor is available for both communications and computation, giving a target peak processing capability of 367 teraFLOP/s.

enable the build of a full $2^{16}$ (65,536) processors system in late 2004 Purple contract. BlueGene/L has important design characteristics that differ from many current-generation systems. The architecture utilizes multiple complementary interconnects with a high-bandwidth, low-latency 3D-torus, combining-tree, barrier, and interrupt networks. Gigabit Ethernet from 1,024 I/O nodes (expandable by a factor of 8) provide access to external disks. Two double-precision floating-point units per processor (4 FPU/node) enhance both the computational and communications capability of the teraFLOP/s. LLNL is working successfully to develop a broad community of support and interest for BlueGene/L within application code groups at the Laboratory, within the Tri-Lab community and other organizations. At the present time, it is anticipated that BlueGene/L will have a huge impact on the SSP in multiple scientific areas.

The compiled list of applications of programmatic interest expected to effectively utilize BlueGene/L is currently over 40 applications and is growing. These are tri-Laboratory applications from LANL, LLNL, Sandia and ASCI Academic Alliances (see "Applications" table below). Early adopter applications include:

- First-principles molecular dynamics {GP, Francois Gygi}
- Atomistic materials models {LAMPSS, Jack Reaugh}
- Dislocation dynamics {DD3D, Wai Cai}
- Continuum modeling {ALE3D}.

BlueGene/L will perform materials simulations at time and length scales that allow overlap calculations of models from different scales. In addition, BlueGene/L will perform molecular dynamic materials properties simulations at a scale allowing direct comparison with experiment. For example, simulations of up to 1 $\mu m^3$ needed for National Ignition Facility (NIF) experiments are possible with BlueGene/L.

| Applications | Scientific Importance | Mapping to BlueGene/L Architecture |
|---|---|---|
| Ab initio molecular dynamics (multiple impacts in materials and computational biology) | A | A |
| Three-Dimensional Dislocation Dynamics | A | A |
| Multi-scale materials modeling | A | A |
| Atomistic dynamic simulations extended to macroscopic timescales (GP, MD3D, DDCMD, GFMD) | A | ~~B~~ A |
| Hydrodynamic instability | A | A |
| Turbulence: Rayleigh-Taylor instability | A | A |
| Shock turbulence | A | A |
| Turbulence and instability modeling (sPPM) | B | A |
| Computational gene discovery (mpiBLAST) | A | ~~B~~ A |

## Accomplishments

**Purple**. Calendar 2002 activities for Purple focused on completing the procurement. On February 21, 2002, the Purple request for proposals (RFP) was officially released from an LLNL website. Multiple, highly competitive responses were obtained from the industry by April 29, 2002. The Tri-Lab review team evaluated the responses during the week of

May 6, 2002. LLNL Director and Weapons Program executives were briefed on May 13, 2002. Tri-Lab ASCI executives were briefed at LLNL on May 15, 2002, and the Laboratory and Tri-Lab ASCI executives all concurred with the selection and high-level negotiations strategy. IBM was officially selected for negotiations on May 31, 2002.

Negotiations with IBM began on June 19, 2002 and concluded almost four months later on October 17, 2002. Purple Contract Review Board was convened on November 1, 2002. Less than two weeks later, DOE/NSSA contract approval was obtained on November 13, 2002. This paved the way for DOE Secretary Abraham with IBM Senior VP for Systems and Technology, Nick Donofrio to announce the Purple Contract at the *Supercomputing 2002* conference in Baltimore, Maryland, a week later. With a budget of $290M, this contract with IBM is the largest contract ever let by the University of California at the LLNL. For a procurement and resulting contract of this magnitude and complexity, the speed with which this process executed from receipt of bids to announcement was truly astounding. The entire procurement team, from technical reviewers to DOE/NNSA contract review responded to the need for quick response on this contract execution.

**Linux Strategy**. Our Linux strategy is focused on utilizing the cost/performance curve of commodity hardware and Open Source software (Linux and clustering tools) for large-scale Production scientific computing at LLNL by extending the Livermore model into this Beowulf (commodity clusters) environment. Our Linux strategy for providing a Production Linux cluster is based on the following concepts

- Provide the same Production software environment and hardware architecture (“Livermore Model”) on multiple platforms. Exceptions for proprietary products will be based on “best of breed” distinction.
- Examples of the latter are the Intel/KAI compilers and Etnus TotalView debugger.
- Transition from proprietary solutions to Linux commodity clusters by starting simply and small and incrementally increasing functionality and scaling up the clusters.
- Work with multiple Open Source consortia to close the gaps between Open Source Linux clustering state of the art and our Production software environment requirements.
- Grow a set of vendor partners in the development consortia that develop products on the Open Source cluster offering (and compete on quality, service and name brand).
- Manage the consortia with a high performance storage system- (HPSS-) like management model through an executive team and technical teams working on specific technical features/areas.
- Comparing the required list of features for the Livermore model on Linux clusters and a summary of the state of the art in Linux clusters quickly leads to the areas for where LLNL should deploy local development efforts, (1) cluster tools; (2) cluster file systems; and (3) resource management. In addition, Linux kernel compilers and tools development and support are critical issues.

Our Linux deployment strategy is to start on small clusters with a small group of people (both computer scientists and computational scientists) interested in this new technology. The second phase of this strategy is to deploy slightly larger clusters into production to learn the issues associated with production usage of Linux clusters and to test the idea that 64-128-way dual CPU clusters are usable and highly robust. The third phase is to then scale these “Compute Node Scalable Units” in large aggregations to build world-class clusters and put these into production. Beyond that, we can deploy Linux clusters routinely at small and medium scale for capacity production computing and with some effort at very large scale as capability production platforms. The objective at every step is to deliver something into production or to lay the foundation for later production usage. This tight coupling between the advanced technology office and production computing portions of the computer center maximizes the benefit to DNT of emerging technologies and minimizes the time for deployments. It also keeps both teams focused on bringing in state-of-the-art technology that actually works.

We executed this strategy for three years. The first year was a start-up exercise in which we gained experience with Linux clusters in partnership with Compaq (now HP) on the four-way SMP Alpha-based system. The second year, the cost performance advantages of dual Xeon (Pentium 4) commodity nodes and the outstanding delivered performance of these nodes on LLNL scientific applications immediately drove us to move to small production clusters with this technology. This was the genesis of the 1.5 teraFLOP/s parallel capacity resource (PCR) clusters for the secure computing facility (SCF). These three clusters (one 128-way dual Xeon with Quadrics, one 88-way dual Xeon with Quadrics and one development cluster with 24 dual Xeon nodes) have been in production for over eighteen months. Usually we deploy technology in the open computing facility (OCF) first and then the SCF. With Linux clusters, having development resources on the open network is sufficient, as the Integrated Computing and communications Department (ICCD) takes responsibility for most problems with the clusters and can debug both hardware and software problems as effectively on the SCF and OCF.

Based on this experience, we executed the next step in the plan: a world-class cluster for capability computing. This step took the form of a capability Linux cluster for multiprogrammatic and institutional computing (M&IC) (MCR) on the OCF. The ICCD executed an extremely aggressive procurement for a 1,152 node, 11.2 teraFLOP/s in record time. This was quickly followed by the ASCI Linux Cluster (ALC), a 960 node, 9.2 teraFLOP/s Linux cluster for BlueGene/L file system development activities. The extremely complex MCR procurement activity (over six contracts) overlapped with the ASCI Purple procurement. From start to finish, however, the MCR was delivered and started science runs within one month of the original target. MCR was announced at SC2002 as the fifth fastest computer according to the LINPACK TOP500 list. After SC2002, MCR performance was improved to 7.654 teraFLOP/s and would place fourth on the TOP500 list, just ahead of ASCI White and just behind ASCI Q. MCR is currently producing stunning computational science results at scale.

Another component of our Linux strategy includes forming of long-term working relationships with industrial partners. To this end, we included funding in the multiprogrammatic capability cluster (MCR) contract for one full-time employee (FTE) at Linux NetworX to work on the Simple Linux Utility for Resource Management (SLURM) project. Others have hailed this strategy within the Tri-Lab community as groundbreaking because it acknowledges that Open Software is not free. Second, we have a support contract with Red Hat (Linux) for on-site kernel support. We are expanding this contract for FY2003 to include two FTEs working on clustering and code development tools (e.g., compilers and hardware performance monitors). In addition, we have been instrumental in getting Red Hat to focus on high-performance computing style Linux clusters.

## Advanced File and Storage Systems Partnerships

A critical component for high-performance cluster computing is a global parallel cluster file system. We are actively engaged in development and deployment of this technology with the Tri-Lab community through several ASCI PathForward activities. LLNL is the lead laboratory for the Lustre development contract with HP [with Intel and Cluster File Systems (CFS) as subcontractors]. Through this “Hendrix Project,” ASCI is working to develop scalable meta-data clusters for Lustre as well as performance, scalability, and robustness enhancements. In particular, we are actively engaged with HP and CFS in Lustre testing on multiple cluster platforms. In addition, we have direct contracts with Cluster File Systems for at least a two-FTE level of effort for stabilization of Lustre Lite (the first major Lustre delivery) on MCR and ALC (see Linux strategy above). As such, LLNL is the first institution fielding Lustre on medium and large clusters. We are actively working with HP and CFS to do testing and develop testing methodology and testing tools.

## A Program ASCI Overview

*James A. Rathkopf*
*Associate A Program Leader*

Simulation codes represent an essential component of A-Program's support of the National Stockpile Stewardship Program. A-Program weapon designers depend heavily on the use of numerical simulations to arrive at key decisions regarding the assessment of the safety, security and reliability of today's United States nuclear weapon stockpile. Within A-Program, simulation code development and use (baselining), analysis of past underground tests (UGT), and design and analysis of new relevant above ground experiments (AGEX) represent a balanced approach to supporting the current stockpile and positioning ourselves to respond to future stockpile needs and requests.

ASCI support all of the scientists, engineers, and computer scientists developing and maintaining A-Program's simulation codes. The codes range from our newest three-dimensional massively parallel design code to a 40-year-old legacy code. These codes support a wide variety of Laboratory program and missions, for example weapon design, output and effects, inertial confinement fusion capsule design for the National Ignition Facility, design of beam interaction targets for the Dual Axis Radiographic Hydrodynamic Test facility (DARHT) at Los Alamos National Laboratory (LANL), and the simulation of laser interactions with tissue for medical and dental technology.

At the present time A-Program's Modern Production Code (MPC) continues to be the main design tool used by the Program. As part of the ASCI Program, this code has undergone significant algorithmic and computational science improvements vital to our effort. Essential physics algorithm and models have been added, and it has been extended to allow for its effective use on our newest computer architectures (also provided by ASCI) to a level of parallelism appropriate for the physics modeled. Due to the extensive validation of the code, it represents our current "physics gold standard" against which other codes are compared. This code also serves the important function of being the "test bed" for new algorithms and models destined for our newest three-dimensional code. The members of our more experienced MPC development team by providing this "test bed" capability also serve the vital role of mentors for the newer and less experienced members of our three-dimensional code development team.

A-Program's three-dimensional massively parallel design code was conceived from the beginning to not only provide the required physics algorithms and models but also to explore the potential advantages to our scientific research and development environment of modern computer science techniques. These computer science decisions coupled with new hardware platforms have resulted in a steep learning curve.

Despite this challenge, the code has met its past two ASCI high-level milestones. The ASCI Burn Code Review committee recently validated the results for the latest ASCI milestone. They agreed that the simulations presented met or exceeded all of the objectives laid out for this milestone by demonstrating the code's ability to simulate—in three dimensions—relevant experiments, obtaining results that compare within reasonable bounds to both experimental data and the output from simulations performed with other well-established codes. These ASCI milestones represent a large portion of the Appendix F Performance Measure 2.3. In addition to meeting our ASCI milestones, this code is beginning to provide important programmatic results for our directed stockpile work (DSW).

The focus of our massively parallel three-dimensional code development is to provide programmatic deliverables for current stockpile work and meet future ASCI high-level application coded milestones. The future milestones are focused on relevant stockpile, work while providing improved physics capabilities that run fully in three-dimensions. The A-Program verification and validation (V&V) effort works hand-in-hand with the code development team to aid in the V&V of the algorithms and physics models being implemented in the code.

Integral to the success of the use of our three-dimensional code for DSW work is the ability to create meshes that faithfully represent the geometries, masses, and materials required for the simulation. A-Program's mesh generation tools provide this capability and have already successfully generated accurate meshes containing millions of zones. Future development of our mesh generation tools is being focused on handling more complicated geometries and providing a more robust user interface so that designers can more easily generate the meshes required.

In addition to our main design codes A-Program is also developing a number of specialized codes necessary to support programmatic needs. These codes include output and effects codes, high numerical order hydrodynamic codes to explore hydrodynamic instabilities and turbulence, and diagnostic replication codes required to fully understand experimental results obtained on existing AGEX facilities.

A-Program's ASCI-funded effort into the study of hydrodynamic instabilities is focused on understanding the linear through highly nonlinear development of the Richtmyer-Meshkov and Rayleigh-Taylor instabilities. The goal of this effort is to develop high-order numerical simulation codes to provide direct numerical simulations of these fluid instabilities in order to understand actual experiments and to obtain "numerical experimental results." These tools and results are in turn being used in the development of large eddy simulation (LES) models as well as other subgrid models.

In addition to aiding in the verification and validation of our design codes, the V&V effort within A-Program is also exploring methods of quantifying the uncertainties associated with numerical simulations, particularly when simulating conditions for which no direct experimental data exists. This effort is extremely important for the development and implementation of the new assessment methodology recently adopted by LANL and LLNL. This assessment methodology is being referred to as Quantification of Margins and Uncertainties (QMU). To address the determination of the uncertainties A-Program's V&V team has partnered with a Sandia National Laboratories code development effort. Sandia has created a tool that encompasses a significant array of uncertainty quantification algorithms as well as optimization routines. A-Program is one of the first users to attempt to use the tool on this complicated problem. An extremely good working relationship has developed between A-Program and the Sandia development team. Initially, this uncertainty quantification development effort is being limited to one- and two-dimensional simulations. However the focus of the effort will be to assure that the tools and techniques can be used for three-dimensional simulations as well.

ASCI supported code development and platform acquisition has been and will continue to be vital to the ultimate success of A-Program's stockpile stewardship mission. While designers ultimately make the final decisions regarding stockpile design and assessment issues, their reliance on simulation tools and the platforms they run on is significant and growing.

## B Program ASCI Overview

*Thomas F. Adams*
*Associate B Division Leader*

B Program simulation and modeling capabilities developed under ASCI provide essential tools for stockpile assessment and certification. ASCI efforts in B Program include two major weapons code development projects, continued development and support for legacy weapons codes, research on advanced material models for metals and high explosives, and a broad range of code and model verification and validation activities. The focus of this discussion is on the development and deployment of the two "ASCI" codes and the legacy codes. These codes support virtually all B Program activities, including

- Analysis of stockpile issues, such as significant findings investigations (SFIs)
- Support for the W80 Life Extension Program
- Fulfillment of peer review and independent assessment responsibilities
- Design and analysis of hydrotest, subcritical, laser, and other nonnuclear experiments
- Baseline studies of stockpile systems
- Work supporting the science campaigns and related efforts, such as reanalysis of archived nuclear test results, to improve our understanding of essential weapons physics phenomena.

The most important overall accomplishment of B Program simulation and modeling is the deployment and active use of the new codes by designers and weapons researchers to meet programmatic objectives. This accomplishment is apparent in almost all B Program stockpile stewardship activities. Specific indicators (performance measures of quality, Appendix F, Sec. 2.3) of the successful development of these codes include the achievement of "ASCI Level-1 Milestones." These level-1 milestones, originally designed as technology capability demonstrations, have been adapted to be more closely integrated with program activities. The FY03 level-1 milestones were the *CY02 Enhanced Primary Physics Milestone*, the *CY03 Nuclear Safety Milestone*, and a *Grain-Scale Explosive Dynamics Milestone*. The Grain-Scale Milestone, which was listed as an ASCI milestone in previous years, and was redefined as an internal LLNL milestone, will be discussed elsewhere in this document (see Fried, "High Explosive Modeling"). The achievement of these milestones, coupled with the increased use of the new codes, marks a transition from the initial development phase of ASCI to a deployment phase, where the codes are used to meet stockpile objectives.

The achievement of the CY02 Enhanced Primary Physics Initial Capability was affirmed by the ASCI Burn Code Review Panel (BCRP) in its meeting in January 2003. Their report states: "The panel unanimously agreed that the Livermore Code group met and exceeded the CY02 milestone. The team performed high-resolution 2D simulations of several NTS events, and presented detailed comparisons with measurements. These exercises demonstrated the ability to perform simulations far exceeding legacy code capabilities both in resolution and in model content."

This milestone consisted of a suite of high-resolution 2D primary burn calculations for several NTS events, run on the ASCI White computer at LLNL. The calculations were compared with data from those tests and related nonnuclear experiments. While the code development team successfully built the code capabilities to support the milestone calculations, much of the execution of the milestone was done by designers. This is another example of the deployment of the code to the user community. These calculations contributed to several current stockpile activities. The BCRP observed: "Much of the reported work was performed and presented by designers, and many of the calculations were, in fact, substantial contributions to various DSW programs, reinforcing the ASCI strategic principle of coherence between milestones and essential programmatic work."

The CY02 Enhanced Primary Physics Milestone also included a formal code release to the B Program design community. The BCRP noted: "As part of the milestone, the B Division Program formally released to the design community a frozen version of the code in concert with the publication of a users' manual. This is a significant accomplishment and represents a major step forward in the integration of ASCI products into the stockpile stewardship program."

Progress toward the CY03 Nuclear Safety Milestone, a "Nuclear Safety Simulation of a Complex Abnormal Initiation Scenario," was also considered by the BCRP at their January, 2003, meeting. They reported: "B Division is also on track to complete the CY03 Nuclear Safety Milestone. The Panel is impressed by the capability demonstrated in this Project to simulate nuclear weapon safety. We believe it is highly likely that this Milestone will be met." Subsequent to that review, the calculations required to accomplish this milestone, a series of 3D simulations of a complex safety-related underground nuclear test, were completed prior to the March 31, 2003, deadline.

The Nuclear Safety Milestone calculations were run with the second B Program "ASCI" code on the ASCI IBM White machine. Execution times ranged up to 100 hours on as many as 800 processors. The multiple calculations were designed to demonstrate mesh convergence and to explore the effects of various physics model assumptions. Comparisons against data from the nuclear test and from a related nonnuclear experiment confirm the quality of the calculations. These results are to be presented to the BCRP at its next meeting in August 2003. Completion of this milestone is an important step in deploying simulation capabilities for use by weapons scientists and engineers, in conjunction with data from relevant experiments, to assure the continued safety, reliability, and performance of the US nuclear stockpile.

The original ASCI milestone sequence includes a CY03 High Fidelity Physics 3D Primary Burn Initial Capability Milestone, which was to be completed by December 31, 2003. As with other high-level ASCI milestones, plans to accomplish this milestone have been adapted to align it more closely with overall program priorities. This milestone builds on 2D calculations that were completed for the CY02 Enhanced Primary Physics Milestone. Work associated with the CY03 milestone is focused on a 3D primary explosion simulation, along with a suite of supporting 2D and 3D simulations, to be run on the ASCI Q machine at Los Alamos National Laboratory (LANL), in support of a current stockpile activity. The Q machine will be used because of its higher performance characteristics, even though it is a new machine with a relatively unstable computing environment. To date, the relevant codes and libraries have been installed and made operational. A significant number of large preliminary calculations have been completed. The prospects for completing the calculations for the milestone are very good.

The accomplishment of ASCI milestones and growing use of the new codes by designers is a reflection of the continuing increase in the capabilities of the codes. This results from the implementation of new or improved models and from the increased robustness and stability of the codes as they are exercised on new applications. The codes are also starting to be used for mesh convergence studies. The ability to run in parallel on finely resolved 2D meshes and on relatively fine 3D meshes has allowed the measurement of convergence characteristics for some aspects of calculations involving hydrodynamics and transport.

One issue that has an impact on both the development and use of the codes has been the growing competition for computing cycles. Even with the availability of the ASCI SKY and White machines at LLNL, and the two sectors of the Q machine at LANL, developers and users are competing for nodes on the machines, especially for large runs. This growing shortage of computing resources decreases the effectiveness of developers, who need quick access for test runs, debugging, and user support. It also decreases the productivity of the users because of delays in their runs or because they scale back the runs they submit. For capacity runs, the codes have been ported to new Linux clusters, where they are running with good efficiency for problems appropriate to the size of those machines. The computing resources limitations are also affecting the ability to run large material properties calculations, such as dislocation dynamics calculations for metals. The material properties calculations, generally running at "standby" priority, do receive sufficient computer time to support reasonable progress. Computer resources similar to those that Blue Gene/L could provide would represent a significant step toward meeting these needs.

# ASCI Advanced Application Development for W Program

Robert M. Ferencz, Leader, Methods Development Group
Daniel C. Badders, Lead Analyst, W80 LEP Group

## Introduction

W Program is responsible for the totality of weaponization engineering activities in support of LLNL's enduring stockpile assets. These activities include system Life Extension Programs (LEP), annual certification, surveillance and Significant Finding Investigations (SFI), process development and liaison for manufacturing and assembly/ disassembly operations. W Program also provides engineering support for physics DSW (e.g., hydro shots), and science campaign activities (e.g., the fielding of sub-critical experiments). Computational modeling, with ever-improving tools, is a major and irreplaceable component of these activities. Our largest computational efforts, in terms of people and machine resources, are allocated to modeling the response of weapon systems to their defined Stockpile-to-Target Sequence (STS) environments. This includes simulation to help design and interpret ground tests intended to replicate these environments. Sophisticated modeling is also brought to bear on the design of unique experimental facilities where the consequences of failure are high, e.g., containment vessels and structures for sub-critical experiments.

The focus of our simulations is mechanical engineering, including stress analysis, heat transfer, and system dynamics. Thus, our code investments have concentrated in these areas. Finite element methodologies, with Lagrangian mesh descriptions, have served extremely well for our basic discretizations. The balance laws of nonlinear solid mechanics are straightforward: the richness in our application space is represented by the variety of material behaviors/models and contact algorithms. A fundamental distinction is drawn between modeling fast, high-frequency transient response and slower, low-frequency or quasi-static response. In this context "fast" or "slow" is a relative characterization: The issue is the time scale of the desired response quantity in relation to the stability limit for explicit time integration on a given mesh. An explicit method uses much less memory per node or element and much less CPU per time step. Thus if the response of interest can be captured in tens or hundreds of thousands of time steps, the explicit method has a natural advantage and is preferred. An implicit method is justified when the response is sufficiently long that relatively large time steps are needed, and thus the memory and CPU costs of solving coupled linear systems of equations is justified. The strategies of explicit and implicit time integration are embodied in two different code projects.

The ParaDyn code team, led by PI Carol Hoover, built directly upon the architecture of the serial DYNA3D code, which remains the archetypal example of an explicit nonlinear structural dynamics code. The ParaDyn team extended the DYNA3D source to incorporate a distributed memory, message-passing parallel programming model, which allows ParaDyn to effectively utilize ASCI hardware resources for engineering analysis. ParaDyn is the workhorse stress analysis tool for engineering DSW at LLNL and at LANL. For example, in March 2002 LANL's Engineering Sciences and Applications–Weapons Response group (ESA-WR) estimated that 85 percent of their total CPU consumption was utilized for ParaDyn simulations, primarily in support of the W76-1 LEP. Likewise, the W80 LEP project team here in W Program makes daily use of ParaDyn.

The Diablo code project is tasked with the creation of an implicit capability for engineering simulation on ASCI platforms to complement the capabilities of ParaDyn. This project is implementing a new architecture with several goals: (a) use of object-based data structures to facilitate future enhancements and maintenance; (b) enable simulations incorporating multiple field theories, e.g., nonlinear stress analysis, heat transfer and fluid flow; and (c) leverage past investments in discretization technology by reuse of software from our legacy implicit codes. The multi-field capability of Diablo is also an enabling architecture for a subsidiary project on Re-entry Vehicle Dynamics.

## Accomplishments

The major ParaDyn effort in recent years has been the implementation of all contact algorithms previously available in DYNA3D. These algorithms enforce constraints to model the interaction of material interface, e.g., unbonded surfaces where the patches of active interaction evolve with time as the system responds to excitation. This capability is essential for STS modeling. Most recently the Lagrange multiplier contact option was parallelized. Its increased accuracy proved useful for a weapon sub-component analysis and contributed toward understanding of the system's complex assembly preload path. Contact algorithms are a substantial computational cost and care must be taken in their parallel implementation. Analysts at LANL recently benchmarked the performance of ParaDyn on ASCI White for a model of an RV forward mount. This detailed model has 21 contact surfaces to simulate the behavior of the threaded surfaces in the joint. In comparing to Sandia National Laboratory's Presto code for executions using 150 CPUs, ParaDyn was found to be 17 times faster. We have taken the same model, executed it on 128 CPUs using the hardware performance monitoring facilities and determined it sustains 12 percent of the theoretical peak floating-point performance. Simpler models attain higher rates, but this benchmark embodies the characteristics of day-to-day engineering analysis in support of DSW.

The Diablo project has established an initial architecture and implemented sufficient functionality to demonstrate parallel implicit thermo-mechanics. This includes an initial penalized contact capability that enforces both mechanical and thermal constraints. One contributor to relatively early parallel execution was the adoption of the finite element interface (FEI) application program interface (API) for linear algebra developed by Sandia. The FEI encapsulated the communication required for management of the distributed left-hand-side matrix. We did port an IBM proprietary direct solver under the FEI to provide a robust capability. Such robustness is paramount for our eventual users, as the mechanical models of weapon systems typically include materials and contact constraints that are highly ill-conditioned and lead to unacceptable performance, or outright non-convergence, of iterative equation solvers. However, the FEI does provide us with a standard interface, which should permit us to evaluate solver packages coming from other ASCI efforts as our own code project matures. As noted, Diablo incorporates a long-term, subsidiary project to establish a capability for high-fidelity modeling of the thermo-mechanics of RV reentry dynamics. One component of that effort has been to leverage the planetary probe community's knowledge of ablation and pyrolysis. This has led us to prototype a capability for surface chemistry coupled with heat transfer and this formulation has been reviewed through publication in a peer-reviewed journal article. This capability has proved to have an interesting, unexpected relevance to current DSW work: this same framework for surface chemistry simulation is now being applied to the study of metal hydriding.

## Appendix F Performance Measure 2.3

W Program Advanced Applications represent a modest portion of the overall ASCI effort at LLNL. Thus while not broad, it is tightly focused on the needs of mechanical engineering analysis for DSW and delivers sustained and ever-improving capabilities to that program. We are not externally reviewed in a formal sense, yet have earned the endorsement of our peers at LANL in the sincerest manner by their adoption of ParaDyn as their main production code, too. During the Engineering Directorate's own recent Director's Review Committee visit, John J. Ruminer, LANL's Deputy Division Leader for Engineering Sciences and Applications, spoke extemporaneously as to their appreciation for the code capabilities and user support they receive from us. It should also be noted that the United Kingdom's Atomic Weapons Establishment (AWE) also utilizes our simulation codes.

## Conclusions and Future Work

The mechanical simulation codes developed within W Program's Methods Development Group are utilized on a daily basis by dozens of engineering analysts to support the Stockpile Stewardship Programs at LLNL and LANL. Our small team, supported by ASCI, TechBase, and LDRD funding, is able to extend and maintain the codes, support our analysts, and research innovative numerical methodologies. Newly launched initiatives will focus on two topics. Adaptive mesh refinement can relieve some of the model development burden of the analyst, while increasing the robustness of numerical solutions and providing some quantitative feedback as to the accuracy/ reliability of the solutions. This should contribute toward characterization of uncertainty as the DSW program deploys QMU methodologies. The second topic, generically known as meshless methods, will explore spatial discretization techniques other than Lagrangian finite elements. Such methods will be more deformation tolerant and can be an important contributor to better modeling of penetration dynamics in support of the recently initiated Robust Nuclear Earth Penetrator (RNEP) phase 6.2 study.

## ASCI Support to the W80 Life Extension Program

*P. Derek Wapman*
*W80 Project Manager*

The W80 nuclear explosive package (NEP) was developed by Los Alamos National Laboratory (LANL) in the 1970s and deployed in the Air Force's Air Launched Cruise Missile (ALCM) and Advanced Cruise Missile (ACM), and the Navy's Tomahawk Land Attack Missile (TLAM) in the early 1980's. The W80 was originally designed to be in the stockpile for 20 years, after which it would be retired and replaced with a newer, more modern design. With the suspension of underground nuclear testing in the early 1990s, the ability to design new weapons and certify their performance using underground nuclear tests ended. The focus of the nuclear weapons program shifted away from developing new designs to extending the lives of the systems that remained in the enduring stockpile. To accomplish this mission the Stockpile Stewardship Program was created to extend the lives of the enduring stockpile weapons and to develop the simulation and testing capabilities needed to certify the performance of the systems and determine their margins and uncertainties. With the loss of underground nuclear tests to assess the effect of changes and the increasing cost and complexity of tests, the need for high fidelity simulations to understand weapon performance increased in importance and played a more central role in certifying the systems.

In October 1998, the Nuclear Weapons Council (NWC) authorized the development of options (Phase 6.2/6.2A) to extend the life of the W80 for another lifetime. Subsequently, LANL completed the Phase 6.2/6.2A process and approval to proceed into development (Phase 6.3) was granted by the NWC in February 2001. Early in Phase 6.2/6.2A, the NNSA determined that Lawrence Livermore National Laboratory (LLNL) would be responsible for the W80 Life Extension Program (LEP) for the planned modifications as well as potential future modifications. LANL continued to maintain responsibility for the existing W80 systems (W80 Mod 0 and Mod 1). Responsibility for the W80 LEP was transferred from LANL to LLNL to accomplish a more balanced workload at the nuclear laboratories and tend to the current needs of the national stockpile.

The process LLNL is following to extend the life of the W80 is fundamentally the same as was used to design a new weapon system. A multidisciplinary project team of physicists, engineers, chemists, and material scientists is working to (1) develop a baseline understanding of the existing system, (2) determine the system margins and uncertainties, (3) modify the system to extend its life, if necessary, and (4) certify that the system continues to meet the requirements specified in the Military Characteristics (MCs), Stockpile-to-Target Sequence (STS), and the Interface Control Documents (ICDs). These responsibilities are being accomplished through a detailed examination of the system to determine problems that could affect the system performance, by modifying the system to address identified problems, and by using an integrated simulation and test program to confirm the new design meets the requirements.

An extensive simulation program is being performed by physics and engineering to design components and tests, predict system performance, and certify the system. The ability to perform these simulations is dependent on the improved code and computer capabilities being provided by ASCI. Without these improvements many of the simulations being used to certify the W80 could not be performed and it would be impossible to analytically assess the design changes without more extensive testing and perhaps underground nuclear tests.

### Accomplishments

The design of the W80 and the changes being implemented by Sandia National Laboratories

(SNL) as part of the W80 LEP have driven the primary design physicists to use the ASCI primary design codes to assess the performance of the LEP primary against the existing system. The ASCI codes were used to design several hydrodynamic tests and predict the test results in advance of the tests. Excellent correlation was seen between the pre-test simulations using the ASCI codes and the test results.

The ASCI secondary design codes are being used to assess the effect of the SNL changes on secondary performance and answer questions raised by several of the key peer review issues. Certification of secondary assembly performance is particularly challenging given the 3D nature of many of the issues that must be resolved.

The engineering ASCI code ParaDyn is being used to model the response of the NEP to the STS environments and also help design and interpret ground tests intended to replicate these environments. The ParaDyn and Diablo codes are also being used to help resolve many of the key peer review issues and for providing more accurate models of the NEP's physical configuration to physics.

## Appendix F Performance Measure 2.3

Measures of the quality of the simulations made with the physics and engineering ASCI codes are routinely performed by the W80 LEP. The W80 was used as the test case for the 2D primary burn milepost that was completed in December 2002 and will be used for the 3D primary burn milepost to be completed in December 2003. The ASCI primary design codes are being validated by the W80 hydrodynamic tests that have been fired to establish a baseline of the system. The secondary design ASCI codes are being validated using small scale and High Energy Density target tests that have been performed to better understand some of the fundamentals of secondary performance. The engineering ASCI ParaDyn code is being validated using data from engineering ground tests that have recently been completed.

## Conclusions and Future Work

The W80 LEP is becoming increasingly reliant on the ASCI codes to provide a detailed understanding of the current system performance and the effect the LEP changes will have on performance. Without the current ASCI codes and the new capabilities planned in the next few years, it will be difficult to certify the W80 LEP design.

# ASCI Verification and Validation Program at LLNL

*Cynthia Nitta*
*LLNL V&V Program Manager, Design Physics Group Leader*

## Introduction

Computational simulation and analyses play a critical role in understanding and coming to technically based decisions on nearly every stockpile issue that we encounter in the Weapons Program. At LLNL, the role of simulation has been central to our contributions in advanced weapon development in both the conventional and nuclear weapons areas throughout the history of the Laboratory. At this critical juncture in the Stockpile Stewardship Program, when the program status and direction is being re-evaluated during the August 2003 Weapons Conference, it is important to appreciate the value of a credible computational simulation capability, especially compared to alternative approaches being discussed such as a return to nuclear testing.

The credibility of our computational simulations depends on the degree to which the capability has been demonstrated through quantitative Verification and Validation (V&V) evaluations. These involve comparisons of simulation results to

- Known analytic/semi-analytic solutions (Verification)
- Other, well-validated codes
- High quality data (Validation).

Because V&V evaluations are predicated upon the use of modern software quality processes (including tailored software process improvement), software quality engineering (SQE) is the fourth cornerstone of the ASCI V&V Program at LLNL. This V&V approach is applied across the primary, secondary, weaponization, and physical data areas.

A major goal of the V&V Program is to deliver quantification of uncertainties contributed by simulations to stockpile system uncertainties. Quantification requires identification of the sources of uncertainties, including from the numerical algorithms, mesh resolution, user code feature choices, physics/material models, physical data input parameters, and validation data. Since the ASCI Program is in its simulation capability delivery phase, most current V&V activities center upon evaluations of the highest priority simulation capabilities available in the codes. The prioritization of V&V analyses depends upon (1) which capabilities are most relevant to current stockpile system need and (2) which capabilities show the highest sensitivities to variations for particular system performance or safety metrics.

## V&V Focus Areas

Current V&V activity focus areas include the following:

- Validation evaluation of physics models in system simulations compared to relevant nuclear device data.
- Establishment of a user-friendly summary of the current V&V status of high priority physics and material models and system simulations, based upon criteria established in each weapon physics area (primary, secondary, weaponization, and physical data) and in SQE.

Another major focus area involves the development of uncertainty quantification methodologies that quantify simulation uncertainties, and eventually establish the linkage between simulation and data uncertainties to stockpile system certification methodologies. Brief descriptions of these focus areas follow.

In FY04, an initial, focused, major validation evaluation will be performed and documented for a stockpile system simulation capability. It involves a 3D secondary performance code resulting in a detailed, quantitative comparison of the simulation with key metrics from experimental data and a determination of uncertainty quantification (UQ) of the accuracy of the simulation. The validation analysis will quantitatively compare the experimental radiographic data from four distinct regions of the experiment with simulated radiographs from the same four regions. These regions will be chosen to provide key information for various aspects of a particular physics capability of the code. A detailed UQ analysis of the accuracy of the code will be performed, taking into account the initial uncertainty of input models such as equation of state (EOS) and the effect of this input uncertainty on output. Many 2D simulations will be performed and sampled by Latin Hypercube sampling techniques to guide the selection of a smaller subset of 3D calculations with the performance code to study the range of uncertainty underlying the calculations. Newly developed techniques to obtain global metrics of the experiment will be used to compare experimental data with code simulations. The outcome of this milestone will be a validation of the secondary performance code's ability to do fundamental focused physics in 3D that pertains to key aspects of a stockpile system.

For FY05, in the primary application area, another major validation evaluation will be performed and documented for a 3D stockpile system simulation capability. A detailed set of verification and validation analyses will be made in 2D and 3D for performance metrics relevant to one or more stockpile weapon systems. Analytic and semi-analytic solutions developed specifically for their relevance to real systems will be applied, and selected. High quality weapon data will be used for comparisons. The outcome of this work will be an initial validation of the current 3D code simulation capability for primary performance through to yield production.

A major focus area is the development and implementation of new methodologies for UQ.

One effort will map out the full phase space of uncertainty in several key output quantities of a stockpile simulation due to uncertainties in a physics model in the secondary simulation code. Uncertainties in physics models that have multiple parametric settings can be represented as uncertainty in an N-dimensional space. The morphology of this N-dimensional space of uncertainty will be quantitatively determined by sophisticated sampling and optimization techniques that constrain the simulations by uncertainties from the detailed measurements of a stockpile event. The phase space of uncertainty in the output quantities of interest for the "as tested" regime of the event will then be extrapolated into regimes that have not undergone testing. We will use new approaches in UQ to determine the uncertainty in our simulations in these new, untested regimes. The outcome of this milestone will be the introduction of a new methodology for UQ and the application of this new methodology to a real stockpile event of interest to the ASCI V&V and Directed Stockpile Work programs.

Another effort is to develop a method that evaluates uncertainties from simulation and experimental data and links them to assessed *Margins and Uncertainties,* then links these quantitatively to reliability and confidence values for system analyses. The quantification of reliability and confidence enables a benefit/cost evaluation to be made for a particular assessed issue. Consideration of several issues then enables a relative value to be placed on the each of the issues, making possible a systematic prioritization of various actions that might be taken. This methodology enables, for example, an evaluation of the relative strategies to reduce uncertainties (e.g., through model improvement or gathering additional experimental data) from a relative value and cost perspective.

An ASCI V&V SQE team at LLNL has been established to provide technical assistance to the ASCI application code teams in building quality into software products by applying the appropriate SQE principles, practices, and processes. These are applied commensurate with program requirements and project resources. There is currently one SQE-knowledgeable, V&V-program-sponsored computer scientist supporting each ASCI code development team in the primary, secondary, and weaponization areas in their software process improvement activities. These individuals will be involved in supporting the code teams as they prepare for a DOE-sponsored Software Quality Audit in FY05.

In FY05, results of V&V evaluations will be rolled-up into a user-friendly summary for physics models and stockpile system simulations across the V&V subprojects relevant to high priority capabilities for the weapons program. Results will include the following:

- Model and system simulation V&V status
- References supporting the completed V&V evaluations
- Appropriate mesh resolution and physical parameter ranges for model use
- Identification of appropriate verification test problems and validation data
- Sensitivities of some performance metrics to variations in model input parameters and example input decks.

V&V standards for high priority computational capability evaluations will be developed for capabilities in each focus area to enable systematic, quantified interpretation of summary results. Each guidance document will include V&V criteria as well as software process improvement guidance for these evaluations. This milestone represents a roll-up of all current V&V evaluations performed to date; the summary will be updated annually to reflect additional V&V quantification progress.

## FY03 V&V Accomplishments to Date

Our V&V accomplishments for FY03 include the following:

- Numerous V&V analyses performed and documented
- Development and application of weapon relevant semi-analytic verification problems
- First round reliability and confidence quantification (application of V&V methodology) derived from model/data uncertainties for stockpile systems assessments (method also applied during W76 1X-109 Peer Review)
- Completion of preliminary work to establish quantification of image comparisons
- Release of unclassified demonstration version of Materials Database graphical user interface (GUI)
- SQE personnel added to support all major ASCI code teams in risk-based, tailored software process improvement; software process improvement plans developed
- Draft Tri-Lab V&V Strategy and Planning documents completed
- Second (annual) Tri-Lab Verification Test Problem Suite established.

In addition, Tri-Lab Software Quality Best Practices Workshop planned the following:

- Tri-Lab Verification Test Suite Workshop
- Tri-Lab ASCI Software Quality Engineering Goals, Principles, and Guidelines document development
- LLNL Site-Specific ASCI Software Quality Engineering Recommended Practices document development
- Numerous invited papers to ASME, AIAA
- DOE Software Quality Forum
- Participation in Spring 2002 Software Quality Assessment,
- Successful completion of an ASCI Level-1 Milestone for V&V Methodology.

## Relationship of Accomplishments to Performance Measures

Development of UQ approaches contributes directly to quantification of the quality of ASCI simulation and modeling capabilities. Verification and validation analyses performed this year have contributed directly to the evaluation of simulation quality and sufficiency when they are applied to important stockpile issues, including for the W76 1X-109 and SFI Peer Reviews, W80 LEP, SFI resolution and weaponization analyses for W87 and B83 issues, W88 and B61 Peer Review activities. In addition, V&V has helped develop analytic tools for evaluation of the quality of the physical and material data that are inputs to the system simulations, through, for example, the development of the materials database GUI.

Finally, V&V has also directly contributed to the favorable reports from the ASCI Burn Code Review Panel, specifically for development of semi-analytic verification test problems and for milestone work in the weapon physics areas.

## Current Issues

Tri-Lab coordination has suffered this year through Continuing Resolution budget constraints placed on the programs until midyear, delaying many planned activities. Budget constraints have also delayed desired collaborations for the development of UQ techniques with external organizations. Machine time constraints have slowed V&V analyses in 2D and 3D, as have personnel resource constraints.

## Materials and Physics Models Development for the LLNL ASCI Program

*Elaine A. Chandler*
*Associate B Program Leader*

Materials and Physics Models Development (M&PM) at LLNL is comprised of two subgroups: Materials Properties and the Subgrid Models. The Materials Properties subgroup is made up of the program elements that develop the materials properties tables and analytical models used for simulations. The program elements include the following:

- Equation of State: Actinides, Hydrogen, and Tables
- Dynamics of Metals
- Opacity
- High Explosive (HE) Detonation Performance Initiative
- Nuclear Properties
- Physics Codes
- Pu Aging due to Self-Irradiation

The Subgrid Models subgroup includes research in Turbulence and in Ejecta Modeling. In this review, we will highlight Dynamics of Metals, Equation of State, High Explosives, and Turbulence work due to time constraints. Details of High Explosives and Turbulence will be described in separate articles. However, we will present a summary of the other activities.

### Materials Properties

The *Equation of State* project seeks to use advanced quantum-based theoretical and computational methods to develop next-generation multiphase equations-of-state (EOS) for plutonium, uranium, hydrogen, and other stockpile materials. Techniques are tested on other materials as well, to understand the range of their applicability and to access data that may be more accessible than weapons materials data. This research thus includes unclassified methodology development and application to transition, rare-earth metals, and actinide metals, as well as close linkage to the static and dynamic experiments in Campaign 2. This project also seeks to develop advanced tabular EOS representations for use in ASCI hydrodynamic simulation codes. There are five major-effort research areas: (1) f-electron correlation and the expansion EOS; (2) high-pressure phase diagram and multiphase EOS in high-*Z* metals; (3) quantum simulations of high-Z metals; (4) hydrogen and low-*Z* EOS; and (5) EOS tables.

Major accomplishments include a new correlated treatment of f-electrons, which has been successfully applied to the volume-collapse transition in the prototype metal Ce. We have developed and tested a first-generation dynamical mean field theory (DMFT) code and successfully applied this code to the analog $\alpha$-$\gamma$ volume-collapse transition in Ce as well as performing initial parallel calculations for plutonium. We expect to complete a DMFT-based expansion EOS for Pu by the end of FY03.

Additionally, A first-generation capability for performing quantum molecular dynamics (QMD) simulations on transition and actinide metals has been developed. The development and testing of a robust parallel code for performing pseudopotential-based QMD simulations of transition and actinide metals at high pressure has been completed, allowing first-principles studies of high-temperature properties with fully coupled electronic and ionic degrees of freedom. Using ASCI White at LLNL, we have carried out QMD simulations of the structural and thermodynamic properties of Mo in both the high temperature solid and liquid phases. Using the ASCI Q machine at Los Alamos National Laboratory (LANL), we have performed the first ever QMD simulations on an f-electron actinide metal, U.

The figure below shows electron-density contours in liquid uranium at a temperature of

2200 K and a pressure of 15 GPa taken from a snapshot of a snapshot of a quantum molecular dynamics (QMD) simulation performed on the ASCI Q machine at LANL.

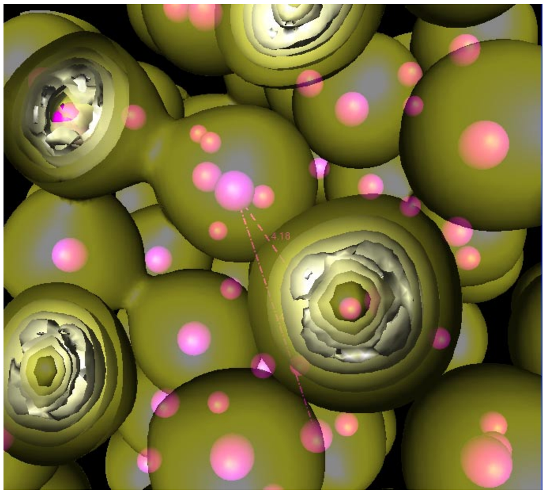

**Electron-density contours in liquid uranium at a temperature of 2200 K and a pressure of 15 GPa.**

Quantum simulation methods for hydrogen and other low-*Z* materials have been improved**.** Innovative improvements in path-integral Monte Carlo (PIMC), variational Monte Carlo (VMC), and density-function-theory (DFT) quantum simulation methods have been achieved and successfully applied to the dense-plasma regime in hydrogen, deuterium, and liquid oxygen. As an outgrowth of this work, a study of dense plasma effects on nuclear reaction rates has also been initiated. New insight has been obtained concerning the enhancement due to ion and electron screening and a first-principles derivation of reaction rates in equilibrium systems has been obtained.

Additionally, Livermore equation-of-state (LEOS) tables software has been improved, and work on a new physics code for EOS generation initiated**.** The LEOS library software has been improved with a new high-quality monotone bicubic interpolator and a linear extrapolation option. The older LLNL EOP and LANL SESAME libraries have been reworked in LEOS format so that these EOS tables are accessible to LLNL hydrocodes without having to write separate interfaces

All ASCI level-2 and level-3 milestones during the past year were met. A total of 21 research papers were published or submitted for publication in the open literature during the same time period, including four *Physical Review Letters*. Approximately 20 invited talks at national and international conferences and workshops have resulted from this research. In addition, three papers have been published in classified literature.

The *Dynamics of Metals* program is a highly visible multi-scale modeling program whose goal is to develop predictive models of strength, plasticity, and failure in metals. It comprises work on the atomistic, micro, meso, and continuum length scales. The accomplishments during the past year were many, and only a few are mentioned here.

Most significant among the successes this year was the first linking of all the length scales for a Ta simulation, with dislocation mobilities developed on the atomistic scale, crystal yield stresses developed using the mobilities in dislocation codes in microscale work, and homogenous polycrystalline yield surfaces developed using a mesoscale simulation of polycrystals and the crystal yield surfaces and elasticity models. The result was a texture-sensitive polycrystalline strength model at the continuum level.

A new highly parallel dislocation dynamics code has been developed and is scaling well with a thousand processors. This has allowed us to attain nearly 1% of strain to date, and to calculate yield point and strain hardening from first principles (see the "dislocation simulation" on page 3). This is a FY03 milestone but, more important, it is also a long-standing dream of materials science. During these simulations the dislocation density reaches $4 \times 10^{12}$ line segments/m$^2$, more than an order of magnitude increase over the starting point. This code has been chosen to run on BlueGene/L during its startup phase.

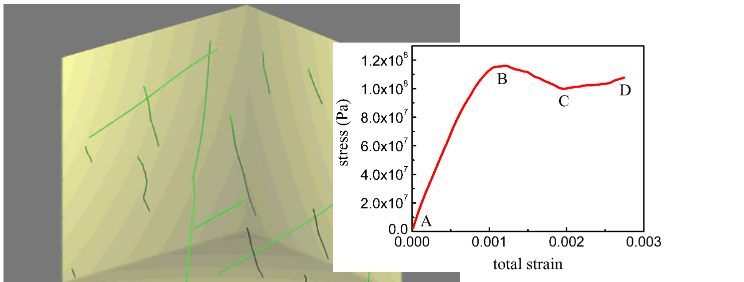


The dislocation simulation box at the beginning of a run with a low density of lines. As the lines increase in number and move under applied stress, the stress-strain curve is mapped out.

**A new highly parallel dislocation dynamics code has allowed us to attain nearly 1% of strain to date, and to calculate yield point and strain hardening from first principles.**

An important issue at the atomic length scale is the crystal structure and the nature and time scale of phase changes. During this year, the evolving atomic structure and kinetics during rapid resolidification in Ta have been simulated in the presence of competing solid phases. Using structurally relevant local-order parameters as diagnostics in a realistic molecular dynamics simulation with quantum-based potentials, we have shown for the first time that a metastable A15 structure appears during solidification of molten Ta, and that crystallization into the equilibrium bcc structure occurs only after this partially ordered phase is lost. The A15 structure shown in figure below is seen to persist for an order of magnitude longer in time than the effectively under-cooled liquid would normally last before crystallization occurs, and thus it has an important impact on the kinetics of resolidification.

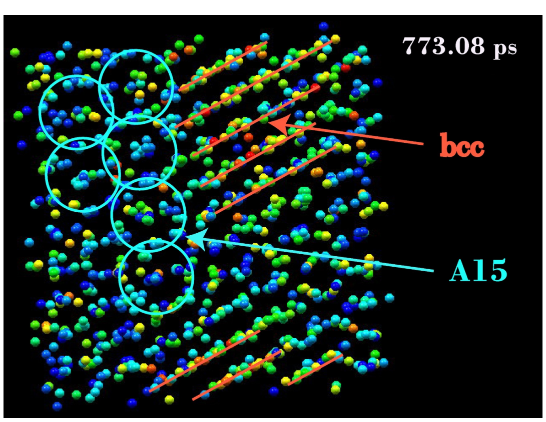


**The A15 phase, which persists for an order of magnitude longer in time than an effectively under-cooled liquid would normally last before crystallization occurs, has an important impact on the kinetics of resolidification.**

At the continuum level, the texture modeling capabilities instituted in previous years in ALE3D has been refined so that we can model the rotation of shape charges built with texture. It has long been a challenge to do this, and Dynamics of Metals Program is the first to demonstrate this capability. In addition, the continuum level researchers have developed a new graphical user interface (GUI), which allows users to graphically develop optimal material parameters for an application by graphically showing strength model predictions and comparing to data sets from the associated data bank. This GUI is now in beta testing.

All ASCI level-2 and level-3 milestones during the past year were met. Approximately 25 papers were published in refereed journals, including four *Physical Review Letters* and two *Philosophical Magazine Letters*. Four review articles have been submitted and accepted by our researchers highlighting the advances in multi-scale modeling associated with our program. They will comprise a special issue of the *Journal of Computer-Aided Materials Design* (JCAD) to be published this summer. Many invited talks were presented worldwide. The opening speech by LLNL's Tomas Diaz de la Rubia at the Spring 2003 Materials Research Society meeting was devoted to a description of the Dynamics of Metals Program. One of the researchers conducted a course in multi-scale modeling at a NATO Advanced Study Summer School. Our researchers also mentored five students at the 2002 Computational Materials Science & Chemistry Summer Institute held at the LLNL Materials Research Institute.

The program element in *Opacity* includes activities in Advanced Opacity, Opacity Databases, and Opacity Verification and Validation (V & V). The goal is to develop and implement improved models of the interaction of radiation with material. They are focusing on developing an 'exact' detailed term accounting (DTA) treatment of local thermal equilibrium (LTE) opacities. The improved accuracy can be seen in the figure at the right, including experiment, DTA, and previous averaged techniques.

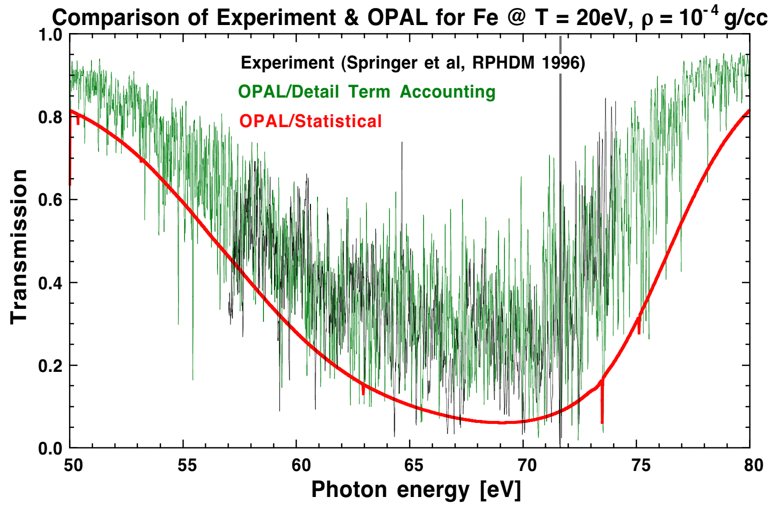


**The goal of Opacity activities is to develop and implement improved models of the interaction of radiation and material.**

Program scientists have built a relativistic, parallelized code called TOPAZ that can compute the detailed opacities for all elements (and relevant temperatures and densities). The initial version of the code has been carefully checked against other known methods as well as in comparison with multi-group transmission experiments for mid-*Z* elements. Having passed these tests, the code will soon be applied to programmatic, above-ground experiments (AGEX), e.g., proposing and guiding opacity experiments for the National Ignition Facility, etc.) and astrophysical questions. This code has been selected as one of eight 'flagship' projects to be given first access to the BlueGene/L machine with its 128,000 processors in the fall of 2004. Additionally, as a result of detailed code inter-comparisons, there were significant improvements in algorithms for computing important cross sections (the improvements were in both accuracy and speed), implementation of new physical models that eliminated serious ad hoc assumptions, and finally, the production of new tables for our five leading opacity codes. Moreover, research and publications have proceeded in high-precision atomic structure, a Laboratory 'signature capability.' In the past year, the group has produced publications in professional journals and delivered talks at conferences, which highlighted the scientific issues in these codes, the research areas, and the applications in the area of helioseismology.

The focus of *Nuclear Properties* and *Physics Codes* is to improve the nuclear databases and particle transport methods in areas that (1) better constrain weapons physics phenomenology or (2) improve the fidelity of simulations of weapon performance and safety. Major objectives include better cross section data for interpreting radiochemical debris signatures and improved transport methods to increase the fidelity of simulations. The program includes elements of code development, database evaluation and management, and experimental and theoretical work to improve nuclear reaction data. Code development highlights include (1) improvements to the Monte Carlo particle transport code library (libmcapm.a) to support laboratory angle biased sampling and thermal neutron scattering and (2) R&D on a new formulation of the radiation transport equation. Database management activities have focused on providing a data management API enabling sensitivity studies to support quantification of margin uncertainties. This has also included new web-based tools to view the data. Theoretical and experimental work has focused on improving cross section data for interpreting radiochemical debris signatures (see McNally graphs at the top of page 5). The theoretical effort has provided improved reaction models and improved calculations of cross sections that are not accessible experimentally. The associated experimental program has focused on high-leverage cross section measurements where the impact on ASCI simulations is so significant that relying on nuclear models is inappropriate.

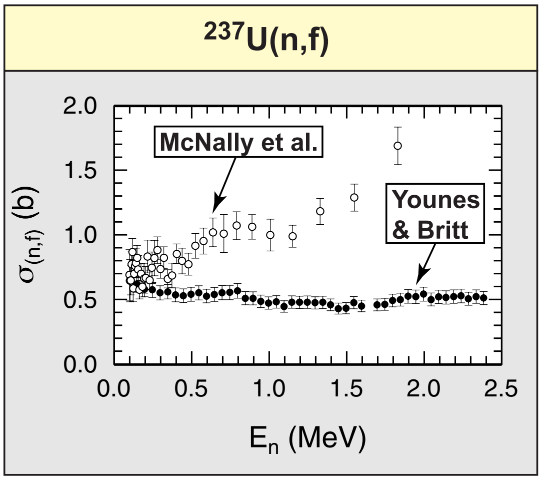


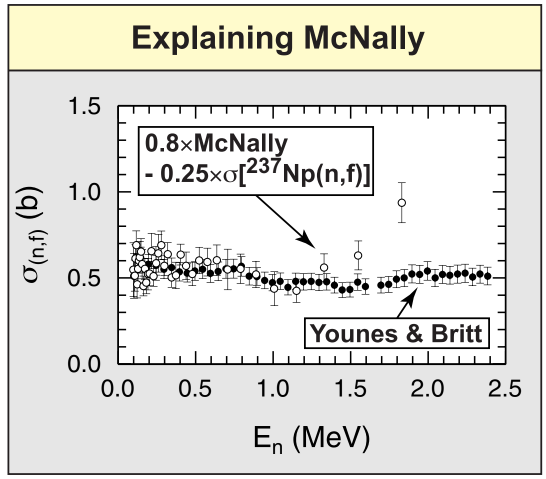


**Previous data (McNally) on the $^{237}$U(n,fission) cross section was inconsistent with criticality assembly integral data. Younes and Britt have re-analyzed particle transfer data to extract this cross section. The new results are consistent with the integral data. The McNally data can be explained assuming a $^{237}$Np contaminant.**

In the last year, this project area has published in journals including *Physical Review Letters* and *Physical Review C Issues,* and in the *Frontiers of Nuclear Research*.

The objective of the project *Pu Aging due to Self-Irradiation* is to develop computational models and codes to predict the changes in properties of Pu and its alloys due to aging. Mechanisms of aging considered are accumulation of actinide decay products, of helium, and radiation-produced defects. In addition, computational methods requiring high-performance computers are being developed for analysis of critical experiments, such as positron annihilation spectroscopy, photoelectron emission spectroscopy, and phase transformations. The computational codes for aging will be validated and calibrated with data from the accelerated aging program, and then applied to forecast aging in the stockpile.

Two codes have been developed for helium bubble formation and for the nucleation of radiation-induced voids (see figure at right). Modeling and computer simulations of bubble formation suggest that radiation damage plays a dual role: Self-interstitial capture keeps helium highly compressed while nearby Pu decays eject helium from the bubbles. These codes have been validated with data from stainless steels irradiated in fission reactors. A suite of first-principle calculations has been performed on the $\alpha$ and $\delta$ phases of Pu-Ga alloys. These calculations have for the first time determined the heat of transformation for the low-temperature martensitic transformation, explained and predicted the density changes produced by Ga, and created the capability to

*He bubbles are observed in actual aged Pu and density is compared with models*

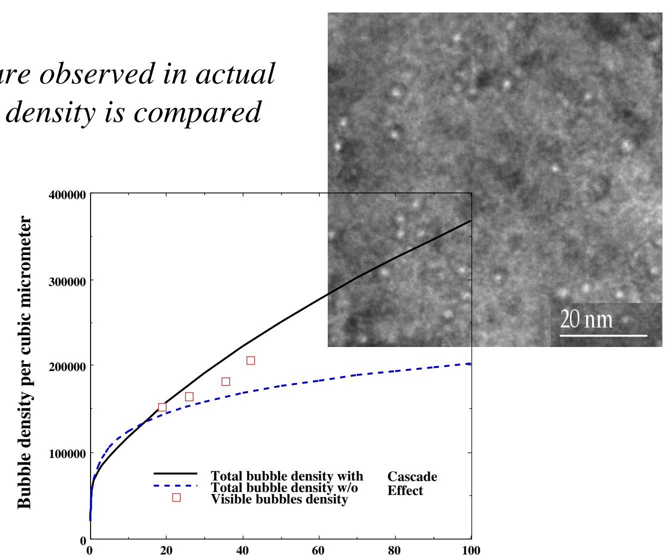

predict long-term stability of Ga-stabilized δ-Pu used in the stockpile.

All milestones during the past year have been met. Four papers were submitted for publication in refereed journals on the subjects of Helium Bubble Formation and Growth and on Void Nucleation. Three papers were submitted on electronic structure calculations for α-Pu, δ-Pu, and Pu-Ga. One paper was published on isentropic compression of irradiated steels, which described and analyzed experimental work funded by Campaign 8 to validate models developed in this project.

## Subgrid Models

The *Ejecta Program* studies the characteristics and transport of the damaged material cast off from a metal surface during shock passage. Ejecta formation work in the last year has focused on the influence of mesoscale effects in alloys. Areas of investigation include schematic effects of grain boundaries, inclusions, and concentrations of eutectic materials. Early direct numerical simulations in 2D indicate that these effects are small and do not affect the overall nature of the formation process elucidated in previous investigations. Validation experiments were developed and fielded with Campaign 2 funding. In these experiments, shocks generated on the LLNL 40-mm gas gun traversed tin targets into one atmosphere of helium. Particles are expected to (1) slow down due to aerodynamic drag based on Reynold's number and particle size and (2) fragment into smaller particles due to pressure exerted on the particle by the gas. For this experiment, only the former effect was expected. The figure at the right shows the comparison with experiment. The black line is the simulated density distribution of tin ejecta, while the light blue and red lines are densities determined from radiography of the experiment (film packs placed in different locations). The green and dark blue lines are similar data taken from a companion experiment fired into vacuum to determine the source mass and velocity distribution. The degree of agreement is quite good, although follow-on experiments using different lots of tin showed markedly different experimental results (actually more like the simulated curve). This result points to the importance of material characterization when analyzing these experiments because it is likely material characteristics played a role in the ejecta formation process. Surface preparation was not a factor as the surfaces were measured to be within a few percent roughnesses of each other. Eight conference papers were presented on this subject during the past year.

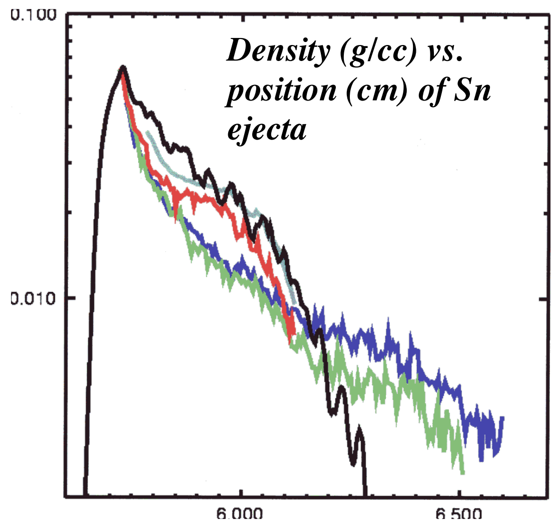


**Simulations (red and black lines) of ejecta into He and vacuum, respectively, show a different power law than expermental data (gree and blue).**

The M&PM Program at LLNL is a vibrant and productive program, developing and applying basic science to improve the modeling capabilities of the ASC codes. It is also a showcase for the use of large-scale computing in basic science, and its success serves an important role in validating the decision of investing in large-scale platforms. It attracts new computer-facile scientists to the lab, and is a big participant in conferences and university collaborations.

## High-Explosive Modeling

*Laurence E. Fried*
*Scientific Capability Leader for Extreme Chemistry, Chemistry, and Chemical Engineering Division*

The purpose of this project is to produce advanced material models of detonating high explosives (HE) for use in ASCI codes and simulations. The accuracy of ASCI code simulations is strongly dependent on the quality of material models used. Historical models miss important aspects of high explosive behavior, such as the detonation velocity size effect (see the figure below). High fidelity simulations require the development of improved material models. Detonation is a truly multiscale phenomenon that couples atomic, grain scale, and macroscopic effects. Accordingly, the scope of the subproject consists of three major efforts: *HE Macroscopic Models, HE Grain-Scale Models,* and *HE Molecular Models*.

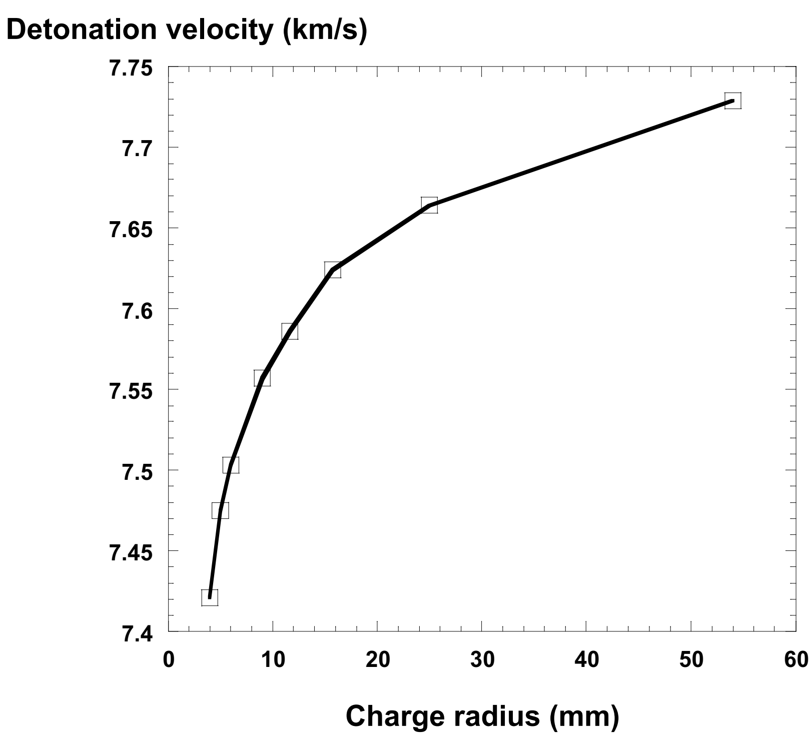


**The detonation velocity of a cylindrical charge of an insensitive high explosive (PBX9502) is a function of charge diameter. The historical beta/program burn model predicts a detonation velocity that is independent of charge diameter.**

### HE Macroscopic Models

An enormous number of engineering-level detonation experiments have been performed in the past. Typically, these experiments have not focused on the chemistry and physics of the HE detonation itself. The goal of this effort is to develop and implement advanced models of HE detonation in ASCI-sponsored hydrodynamic codes. The achievement of this goal will require a number of scientific and algorithmic advances, as well as careful validation against experimental data.

We are currently developing a new approach to reactive flow modeling. We are linking CHEETAH, a thermochemistry code with a chemical kinetic framework, to the ARES hydrodynamic code. New algorithms have been developed that make the computations affordable on current computers. For the first time, we can track individual chemical species through the detonation process. This also means that equation-of-state measurements on individual products of detonation, such as water, can be related to the equation of state of the detonation product mixture in a scientifically rigorous way. Current calculations are performed with between 20–30 chemical species, with 1–2 species out of equilibrium. In the future, calculations with more complicated reaction schemes will be possible. CHEETAH linked to ARES is developing into an attractive alternative to traditional high explosive models in a wide range of applications. The model is particularly suited to aged materials where subtle chemical or morphological changes may occur.

## HE Grain-Scale Models

The objective of the HE grain-scale modeling effort is to integrate detailed simulations of explosive response and derive macroscopic models based on and consistent with the heterogeneous structure of the explosive. The detailed simulations include descriptions of the structure that include explosive crystal grain-size distribution, binder, and inter- and intra-crystalline voids. Micro-structural simulations of the decomposition rates and thermochemical simulations of the product species will be used to improve the thermo-kinetic descriptions of reacting explosives in these simulations.

Simulations have been performed to model a collection of binder-covered explosive grains with specified grain-size distribution and void volume. A shock wave moves through a microscopic explosive domain and chemical reactions are observed. These simulations will provide chemical and rate models of the coarser grain-scale simulations. ASCI-class computational resources have allowed us to perform a calculation on the size scale of a representative volume element, which corresponds roughly to a single zone in a hydrodynamic calculation. We have been able to derive form factors for reaction growth and implement them in a successful macroscopic model.

## Performance Measures of Quality

Macroscopic high explosive models created in this project are validated against both small scale and large-scale experimental data. The models are available for use in ASCI code simulations. The work in this project has led to numerous publications in peer-reviewed journals and conference proceedings. An external review of the most recent grain-scale simulations was conducted in January 2003. The review panel report concluded that we have made significant achievements in understanding high explosive behavior at the grain scale, and that we have met all milestone requirements associated with grain-scale modeling.

## Conclusion

We are making strides in developing science-based models of high explosive detonation. Our approach couples microscopic, grain-scale, and macroscopic phenomena. The resulting models are validated against a wide range of experiments. ASCI-class computational resources have made realistic grain-scale simulations of plastic-bonded high explosives possible. Further enhancements in the accuracy of microscopic, grain-scale, and macroscopic models through more advanced physics and chemistry will yield the first fully quantitative and scientific models of the behavior of insensitive high explosives.

# Direct Numerical Simulation and Large-Eddy Simulation of Mixing Instabilities

*Paul L. Miller*
*Physicist — AX Division, A Program*

## Direct Numerical Simulation (DNS) of an Impulsively Accelerated Interface

### *Introduction*

This presentation describes two separate investigations of mixing instabilities. The first builds upon a long-term collaboration with researchers at the University of Arizona, where Professor Jeff Jacobs and his students developed an LLNL-sponsored experiment on impulsively accelerated interfaces. This work explores a fundamental mixing instability of accelerated density interfaces and serves as validation of the simulation code for follow-on work on Rayleigh-Taylor simulations described below.

The Arizona experiments consisted of a small tank, mounted on vertical rails, containing a heavy fluid and a lighter (miscible) fluid. The tank was oscillated horizontally, creating standing waves on the heavy-to-light fluid interface. The tank was then dropped, bounced on a spring, and allowed to coast up and back down the rails until it came to a rest. Video images documented the development of the impulsively accelerated interface (see Figure 1). We made use of a spectral/compact-difference DNS code named MIRANDA to explore the dynamics and mixing behavior of the flow, using the code in a 2D manner. (DNS or *direct numerical simulation* means that there was no modeling of the small scales; viscosity and diffusion are explicitly calculated.) The flow's viscous scales are fully resolved in the calculations, and the species diffusion scales are captured at a Schmidt number of 100. After simulating the initial conditions and the pre-bounce development, we calculate the post-bounce development and validate the calculations with data from Niederhaus (2000). We then proceed to examine the behavior of a secondary instability that develops in the core of the primary vortex.

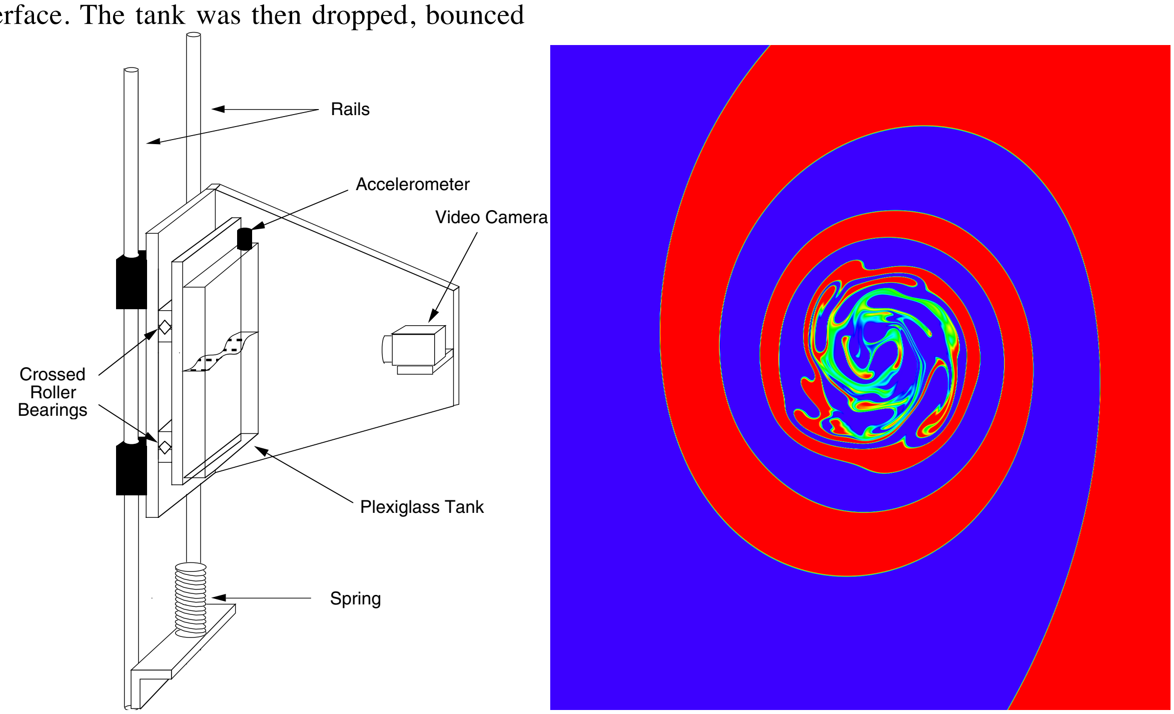


**Figure 1. (left) a schematic of a portion of the Arizona drop-tank apparatus; (right) pseudo-color image of the simulated density field of a roll-up.**

### Results

We found that the secondary instability is driven by the centrifugal pressure gradient of the large-scale vortex interacting with the very steep density gradients of the high-Schmidt-number fluid spiraling into the vortex (Fig. 1). Vorticity is produced, as expected from the baroclinic term in the vorticity equation, at density gradients along the sides of the pressure well. Some interface orientations are Rayleigh-Taylor unstable, and those orientations exhibit additional perturbation growth. Small-length-scale features are generated in the process of the baroclinic instability by the nature of the (cross-product-driven) term, the generation of opposite-sign vorticity, and the subsequent instabilities of the additional vorticity generated. An important confirmation of the physical process was provided by visualizations of the vorticity field and vorticity production term, permitting us to look for (lack of) correlation between the two. Viewing two scalar fields was achieved with a colored height map of the vorticity field (Fig. 2). A movie was produced to allow examination of the developing flow. With two spatial dimensions, two scalar fields, and time, it represents a five-dimensional visualization. What the movie made clear is that new vorticity was being produced, creating additional small-scale structure in the flow, and that it was not simply a redistribution of previously existing vorticity. It is interesting to note that, while the large-Schmidt-number fluids examined here may accentuate the effects, the secondary baroclinic vorticity deposition and subsequent instability described here is not fundamentally limited to such cases. In particular, it is clear that longer time scales, which would permit more time for the instability to develop, and/or stronger pressure gradients, have the potential to compensate for the lack of strong density gradients, in certain flows. We believe that this instability plays an important role in a wide variety of mixing flows.

### Performance Measures of Quality

The MIRANDA code has undergone a variety of verification activities throughout its development and implementation, and has been validated against the University of Arizona experiments in the course of this work. A video describing this work was honored as a winning entry in the 20th Annual Gallery of Fluid Motion, at the 55th Annual Meeting of the American Physical Society, Division of Fluid Dynamics, held in Dallas, Texas on November 24–26, 2002.

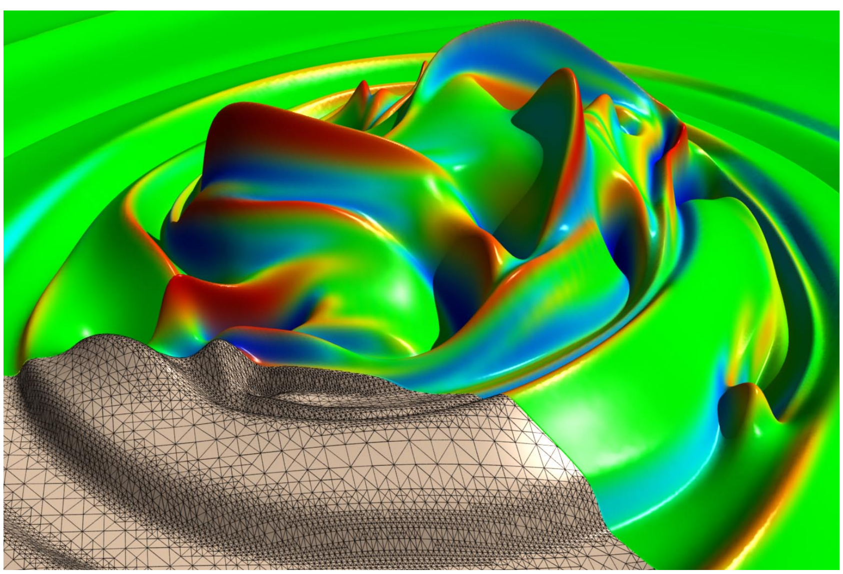

**Figure 2. A height-map rendition of the vorticity field, with vorticity production added as a pseudo-color on the surface. A portion of the visualization mesh has been made visible.**

## Large-Eddy Simulation of a Rayleigh-Taylor Instability

### Introduction

A second investigation, in progress at the time of this writing, is exploring the behavior of the Rayleigh-Taylor instability. It also uses the MIRANDA code employed in the work described above, but in fully 3D mode and as a large-eddy simulation (LES). In short, the LES version makes more efficient use of the range of scales available for the problem by reducing the high-wavenumber requirements by almost an order of magnitude. This permits simulation of much larger effective Reynolds numbers for comparable computational effort.

The Rayleigh-Taylor instability plays an important role in processes such as ICF implosions, and we are pursuing the development of models that will be incorporated into ASCI codes to account for its mixing effects. These simulations are part of our effort to provide simulation data—for the model development and validation—that is not available from experiments.

## Results

The problem was set up in a box 720 × 720 × 1620 mesh points, with the top half of the box filled with density 3.0 fluid and the bottom half with density 1.0 fluid. A multimode initial perturbation was prescribed at the interface between the two densities, and then gravity was imposed in the downward direction. As the problem progresses, the high density fluid sends "spikes" of material downward while the low-density fluid forms rising "bubbles" (see Fig. 3). Black represents the densest fluid, white the lowest density, and the intermediate grays represent the range of densities in between that result from mixing.

A prime objective at the onset of the simulation was to investigate mixing behavior in the developing layer as a function of time. As the layer grows in width, its Reynolds number increases. Theoretical considerations suggest that other flows experience a mixing transition above a threshold Reynolds number, above which there is a marked increase in the proportion of mixed fluid. We are tracking a "mixedness" parameter, as well as exploring the mixing behavior by means of visualizations, fluid mixture histograms, and density power spectra. Density and vertical velocity component power spectra demonstrate that the flow becomes very well developed, exhibiting a decade or more of power-law behavior.

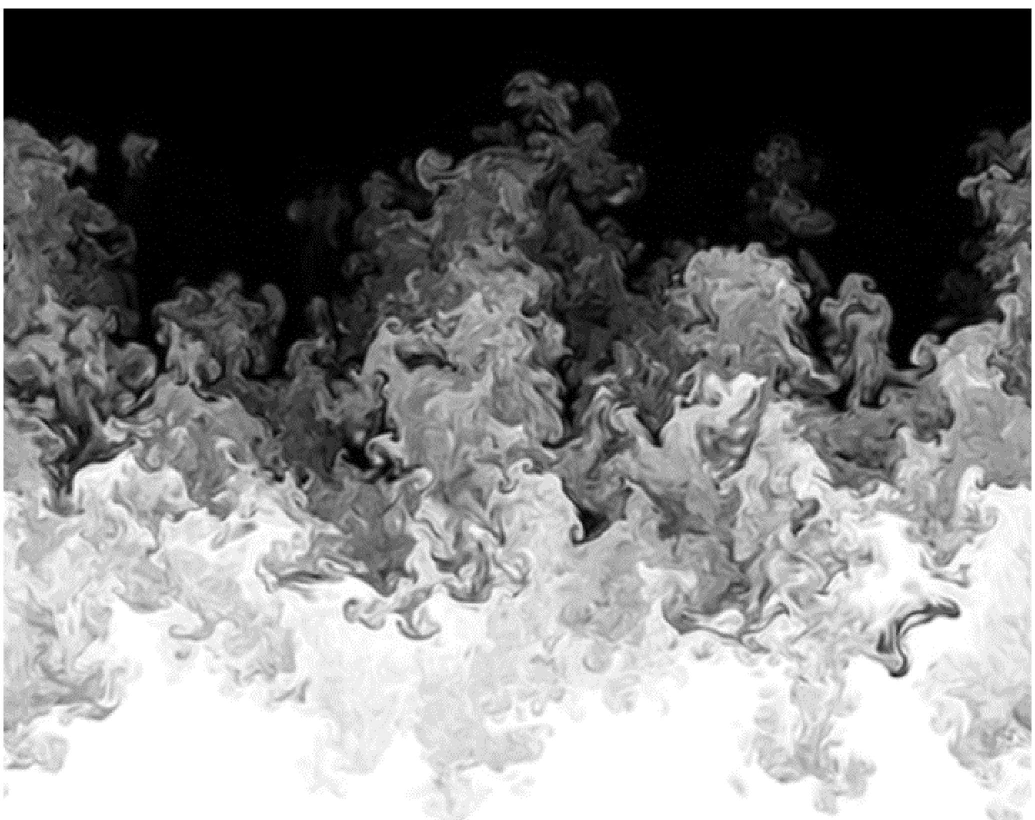

**Figure 3. A grayscale image of the density field from a vertical section through the center of the Rayleigh-Taylor simulation, at non-dimensional time = 100 (near the midpoint of the simulation).**

However, as the simulation progresses, another item of interest has shifted the current attention away from the mixing behavior. The layer width clearly is growing linearly with time. This behavior, observed in some earlier DNS simulations as well, is quite significant because nonlinear Rayleigh-Taylor growth has been described traditionally as quadratic in time. The finding is quite substantial, and we are expending considerable effort to put it in perspective. The linear behavior has stimulated a flurry of theoretical activity to explain it, and a series of moderate sized simulations are underway to explore the parameters and sensitivities of this behavior. At this time it is too early to make definitive statements, except that very interesting things are being uncovered by this run, and that regardless of the final details, it will change our view of the Rayleigh-Taylor instability.

### *Performance Measures of Quality*

These high-resolution LES results have been benchmarked against the best other Rayleigh-Taylor simulations available at this time, including ones by David Youngs at AWE (Atomic Weapons Research Establishment, UK). The MIRANDA code has also, as described above, been validated against the University of Arizona data. Since the simulation is still in progress, peer review has been limited to internal (A Program) personnel so far.

## Conclusion

In our simulation of the University of Arizona experiment, we uncovered a mixing mechanism driven by generation of secondary vorticity from the interaction of the complex density structure and the pressure field in the flow. This mechanism is basically a centrifugal version of the Rayleigh-Taylor instability, and it plays a role in developing atomic-scale mixing in a variety of flows of interest to us. Along the way, we also were able to validate the MIRANDA code against the experimental results. This provided us with some assurance as we moved on to more idealized Rayleigh-Taylor simulations.

Our large Rayleigh-Taylor simulation that is in progress (at this writing) is providing some surprises, particularly by exhibiting a linear growth rate. As described above, this has already generated a lot of activity. While firm conclusions await the results of follow-on work, it is apparent that issues of initial conditions and assumptions about the scaling of the flow need to be critically reexamined, and that this work will result in a significant change in our understanding of the Rayleigh-Taylor instability, one way or another.

# Glossary of Acronyms, Abbreviations, and ASCI Terms

Below is a list of acronyms and abbreviations used in the ASCI Director's Review presentations and summary articles.

| | |
|---|---|
| 2D | two dimensional |
| 3D | three dimensional |
| ACM | Advanced Cruise Missile |
| AGEX | above-ground experiment |
| ALC | ASCI Linux Cluster |
| ALCM | air-launched cruise missile |
| ALE3D | arbitrary Lagrange-Eulerian grid technology code |
| AMR | adaptive mesh refinement |
| API | application program interface |
| ARES | a hydrodynamic code |
| ASCI Purple C | a 100-teraOPS computer (LLNL) |
| ASCI Purple | a 60-teraOPS computer (LLNL) |
| ASCI Q | a 20-teraOPS computer (LANL) |
| ASCI SST | ASCI Blue Pacific Stockpile Sustained Teraop (system) |
| ASCI White | a 12+ teraOPS computer at LLNL |
| ASCI | Advanced Simulation and Computing Program (formerly known as the Accelerated Strategic Computing Initiative) |
| AWE | Atomic Weapons Research Establishment, UK. (formerly AWRE) |
| bcc | body-centered-cubic |
| BCRP | ASCI Burn Code Review Panel |
| Beowulf | commodity clusters |
| BG/L | see BlueGene/L |
| BlueGene/L | a scalable, unclassified 360 teraOPS computational science research platform |
| C&MS | Chemistry and Materials Science (LLNL directorate) |
| CASC | Center for Applied Scientific Computing (program at LLNL) |
| CFS | Computer File Systems, Inc. |
| CHEETAH | a thermochemical code with a chemical kinetic framework |
| CR | continuing resolution (budgets constraints) |
| CTBT | Comprehensive Test Ban Treaty |
| CY | calendar year |
| DAM | Defense Applications and Modeling |
| DARHT | Dual Axis Radiographic Hydrodynamic Test (facility at LANL) |
| DARPA | Defense Advanced Research Projects Agency (DoD) |
| DFT | density-function-theory |
| DisCom | distance and distributed computing and communication |
| DMFT | dynamical mean field theory |
| DNS | direct numerical simulation |
| DNT | Defense and Nuclear Technologies directorate |
| DO | Director's Office (LLNL) |
| DoD | Department of Defense |
| DOE | Department of Energy |
| DP | defense programs |
| DSW | directed stockpile work |
| DT | deuterium-tritium |
| DTA | detailed term accounting |

| | |
|---|---|
| DYNA3D | an explicit nonlinear stress analysis code |
| EDTV | Early Demonstration of Technology Vehicle (ASCI Purple) |
| EOS | equation(s) of state |
| EPR | enhanced priority run |
| ESA-WR | LANL's Engineering Sciences and Applications–Weapons Response group |
| FEI | finite element interface |
| FY | fiscal year |
| FYNSP | Future Years Nuclear Security Program |
| GUI | graphical user interface |
| HE | high explosive |
| HEAF | High Explosives Applications Facility |
| HPSS | high performance storage system |
| I/O | input/output (usually a computer interface) |
| ICBM | intercontinental ballistic missile |
| ICCD | Integrated Computing & Communications Department |
| ICD | Interface Control Document |
| ICF | inertial confinement fusion |
| ICS | Integrated Computer Systems |
| JCAD | *Journal of Computer-Aided Materials Design* |
| KULL | ASCI code, object-oriented framework |
| LANL | Los Alamos National Laboratory |
| LANSCE | Los Alamos Neutron Science Center |
| LASNEX | computer code |
| LC | Livermore Computing |
| LDRD | Laboratory Directed Research and Development (special LLNL R&D funding) |
| LEOS | Livermore equation-of-state |
| LEP | life extension program |
| LES | large-eddy simulation |
| Level-1 milestone | formerly called an ASCI "milepost" |
| LLNL | Lawrence Livermore National Laboratory |
| LTE | local thermal equilibrium |
| M&IC | multiprogrammatic and institutional computing |
| M&PM | Materials and Physics Models |
| MCR | Multiprogrammatic Capability Cluster |
| metadata | data about data; descriptions and other information about other data |
| MIRANDA | a spectral/compact-difference direct numerical simulation (DNS) code |
| MPC | Modern Production Code (A Program) |
| MPP | massively parallel processor |
| NEP | nuclear explosive package |
| NIF | National Ignition Facility (LLNL) |
| NNSA | National Nuclear Security Administration |
| NTS | Nevada Test Site |
| NWC | Nuclear Weapons Council |
| OCF | Open Computing Facility |
| OVERTURE | ASCI code: object-oriented, terascale simulation framework |
| ParaDyn | an explicit nonlinear structural dynamics code |
| PCR | parallel capacity resource |
| petaflops | $10^{15}$ floating-point operations per second |
| PF/s | petaflops per second |

| | |
|---|---|
| PIMC | path-integral Monte Carlo |
| PowerWall | large visualization facility, which is used for formal reviews and presentations, and data analysis working groups |
| PSE | problem-solving environment |
| Pu | plutonium |
| Purple C | a 100-teraOPS system (see “ASCI at Livermore Platform Strategy”) |
| PVC | parallel visualization cluster |
| QMD | quantum molecular dynamics |
| QMU | quantification of margins and uncertainties |
| R&D | research and development |
| RFP | request for proposals |
| RNEP | Robust Nuclear Earth Penetrator |
| RT | Rayleigh-Taylor |
| RV | re-entry vehicle |
| S&CS | Simulation and Computer Science |
| SCF | Secure Computing Facility |
| SFI | Significant Findings Investigations |
| SLBM | submarine-launched ballistic missile |
| SLCM | submarine-launched cruise missile |
| SLEP | Stockpile Life Extension Program |
| SNL | Sandia National Laboratories |
| SQE | software quality engineering |
| SRD | Secret Restricted Data |
| SSP | Stockpile Stewardship Program |
| SST | sustained stewardship teraflop |
| STS | Stockpile-to-Target Sequence |
| teraOPS | trillions of operations per second (also teraFLOPS, teraFLOP/s and TF) |
| TF | teraFLOPS (also teraFLOP/s, teraOPS) or trillions of operations per second |
| TLAM | Tomahawk Land Attack Missile |
| TOPAZ | a relativistic, parallelized code |
| Tri-Lab | Livermore, Los Alamos, and Sandia national laboratories |
| TSF | Terascale Simulation Facility (site of BG/L and ASCI Purple) |
| UC | University of California |
| UGT | underground tests |
| UQ | uncertainty quantification |
| V&V | verification and validation |
| VIEWS | visual interactive environment for weapons simulation |
| VMC | variational Monte Carlo |
| W Program | weapons engineering |
| WAN | wide area network |
| $Z$ | atomic number |