

# Nuclear Weapons Highlights 2007

Associate Directorate for  
Theory, Simulation, and Computation (ADTSC)

**Alan R. Bishop**  
*Associate Director*

**Paul J. Dotson**  
*Deputy Associate Director  
Nuclear Weapons*

**Andrew White**  
*Deputy Associate Director  
High-Performance Computing*

**Stephen R. Lee**  
*CCS Division Leader*

**Alyn Ford**  
*CTN Division Leader, Acting*

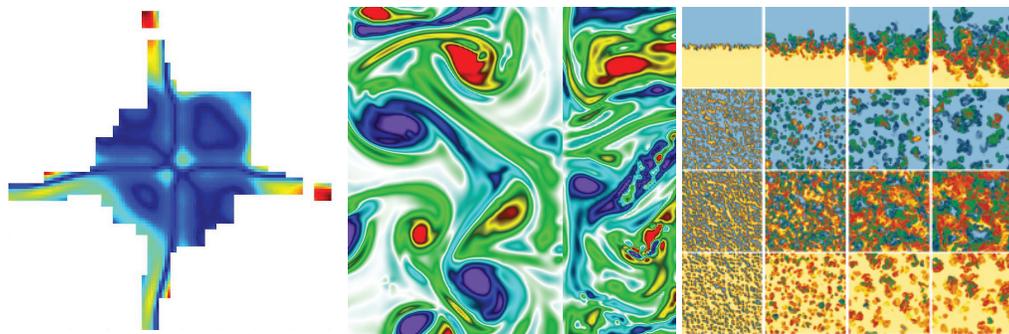
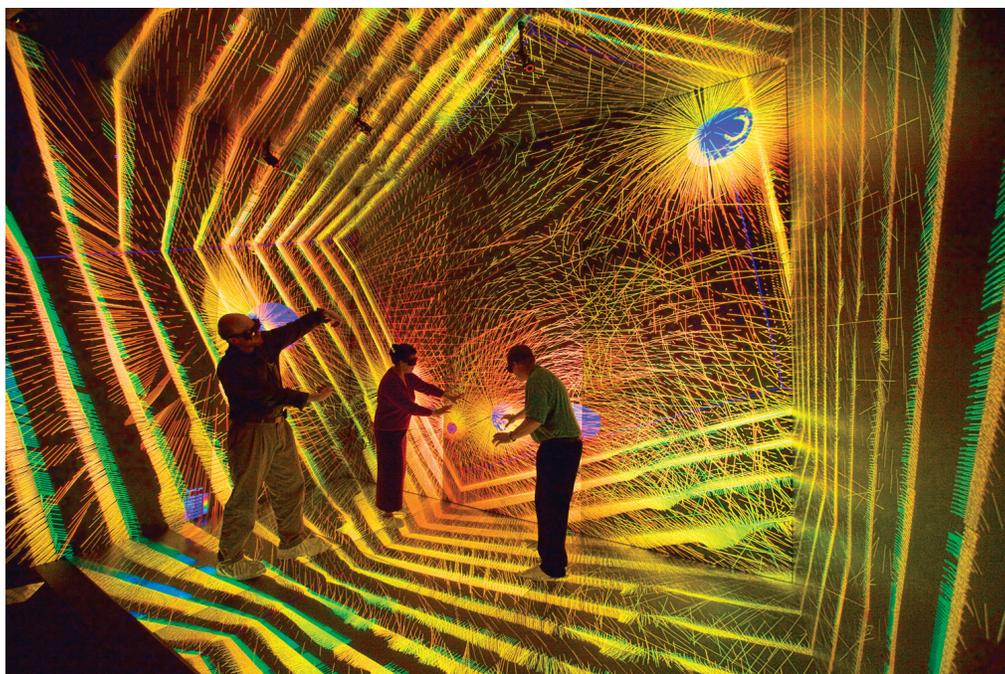
**John F. Morrison**  
*HPC Division Leader*

**Antonio Redondo**  
*T Division Leader, Acting*

**Denise Sessions,  
Susanne King**  
*Editors*

**Sharon Mikkelson**  
*Publication Designer*

**Printing Coordinator**  
*Guadalupe D. Archuleta*



ADTSC web site: <http://www.lanl.gov/orgs/adtsc>

Nuclear Weapons Highlights 2007

LALP-07-041





# Nuclear Weapons Highlights 2007

## Associate Directorate for Theory, Simulation, and Computation (ADTSC)

### Preface

The Theory, Simulation, and Computation Directorate (ADTSC) plays a pivotal role for the Nuclear Weapons (NW) Program, serving as the principal steward for the theory and modeling, simulation, high-performance computing, and computer science capabilities for the Los Alamos National Laboratory. These capability areas serve the NW program in broad areas of physics and engineering for weapons physics, fluid mechanics, and computational methods, statistical sciences and applied mathematics, and high-end computation and computer science, and are also applied to other programs at the Laboratory.

The NW program has formalized its partnership with the Laboratory as a whole by identifying ten broad areas of capability central to the NW mission, then giving division leaders the responsibility, and the resources, for managing the health and development of these capability areas. Of the ten capability areas, division leaders from the ADTSC directorate are managing four of them. These capabilities are Computational Physics and Mathematics (led by Stephen Lee), Computer Sciences (Stephen Lee), High-Performance Computing (John Morrison), and Theory and Modeling (Tony Redondo). These capability leaders are given negotiated financial authority to manage the capability areas on behalf of the NW program.

The future challenge is the more general goal of true science-based prediction, where the natural scientific partnerships among theoreticians and modelers, experimentalists, and the high-performance computing and simulation community are purposefully focused to produce validated simulation tools for a broad array of national-level security problems, on scales befitting a national laboratory.

*Paul J. Dotson*  
*Deputy Associate Director, ADTSC*



# Table of Contents

Preface .....	iii
Organizational Abbreviation Legend .....	xi
Computer Science .....	xii
<b>Computational Architecture Exploration</b>	
Open MPI: A High-Performance, Fault-Tolerant Message-Passing Interface.....	2
<i>Brian Barrett, Ralph Castain, Galen Shipman, CCS-1</i>	
Hiding Communication Overheads in High-Performance Computing.....	4
<i>Jose Carlos Sancho, Kevin J. Barker, Darren J. Kerbyson, Kei Davis, CCS-1</i>	
A Novel Single Projection Method to Calculate View Factors for Radiosity, and a GPU-Based Implementation in a Large Multiphysics Code .....	6
<i>Sriram Swaminarayan, John Turner, CCS-2</i>	
Sparse Matrix-Vector Multiplication and Conjugate Gradient Algorithm on Hybrid Computing Platforms .....	8
<i>David DuBois, Andrew DuBois, Carolyn Connor Davenport, Stephen Poole, HPC-5</i>	
Exploring Solutions to ASC Simulation Code Input/Output Issues .....	10
<i>Steve Hodson, James Nunez, HPC-5</i>	
New Approaches to Fault-Tolerant Computing .....	12
<i>Anders Hansson, Gabriel Istrate, Christian Reidys, Allon G. Percus, CCS-3; Stefan Boettcher, Bruno T. Goncalves, Emory University; Shiva P. Kasiviswanathan, Carnegie Mellon University</i>	
<b>Libraries and Tools</b>	
Scientific Application Development using Eclipse and the Parallel Tools Platform .....	14
<i>Greg Watson, Craig E. Rasmussen, CCS-1</i>	
<b>Software Design and Verification</b>	
A Software Quality Engineering Case Study: The Jayenne IMC Project Gets a New Mesh Type .....	16
<i>Todd Urbatsch, Mike Buksas, Tom Evans, Aimee Hungerford, Scott Mosher, Chris Fryer, Jeff Densmore, Tim Kelley, CCS-2</i>	
<b>Theoretical and Applied Foundations of Data and Management</b>	
Blind Inversion in Speech and Weapons Image Processing.....	18
<i>John Hogden, CCS-3</i>	
<b>Advanced Visualization and Data Management</b>	
Using High-End Collaborative Technologies at Los Alamos National Laboratory .....	20
<i>Cindy Sievers, CCS-1</i>	
Feature Painting from 2-D Feature Grids and 3-D OSO Models onto 2-D Cartesian and Cylindrical Grids .....	22
<i>Martin Staley, T-7</i>	
PixelVizion: An NPU-based Compositor for the Visualization of Large Datasets .....	24
<i>Carolyn Connor Davenport, Andrew DuBois, David DuBois, Stephen Poole, HPC-5; Laura Monroe, David Pugmire, HPC-4</i>	

Pick-n-Place: A Virtual Reality Assembly Tool.....	26
<i>David Pugmire, Laura Monroe, David Modl, HPC-4</i>	
Linux Clusters Now Drive Metropolis Center Visualization.....	28
<i>David G. Modl, HPC-4; Andrew Shewmaker, HPC-5</i>	
<b>System Software</b>	
Dynamic Kernel Instrumentation on Clusters .....	30
<i>Nathan DeBardeleben, William Rust, HPC-4</i>	
<b>Hydrodynamics Methods and Algorithms .....</b>	<b>33</b>
Effects of Pressure-Temperature Equilibrium on Richtmyer-Meshkov Instability.....	34
<i>John W. Grove, Thomas Masser, CCS-2</i>	
Numerical Simulations of Raleigh-Taylor Turbulence with a Complex Acceleration History.....	36
<i>Praveen Ramaprabhu, Malcolm J. Andrews, CCS-2; Guy Dimonte, X-1</i>	
The Physics of Radiative Shock Waves .....	38
<i>Robert B. Lowrie, CCS-2; Rick M. Rauenzahn, T-3; Jarrod Edwards, Texas A&amp;M University</i>	
Calculations of the Single-Mode Rayleigh-Taylor Instability: RAGE and FLAG Hydrocode Comparisons.....	40
<i>Marianne M. Francois, Robert B. Lowrie, CCS-2; Jimmy Fung, Mark A. Kenamond, and Edward Dendy, X-3</i>	
Multimaterial Interface Reconstruction Using Particles and Power Diagram.....	42
<i>Rao V. Garimella, Samuel P. Schofield, Raphael Loubere, T-7; Marianne M. Francois, CCS-2</i>	
Simulation of the Rayleigh-Taylor Instability Using Atomistic Methods .....	44
<i>John L. Barber, Kai Kadau, Peter S. Lomdahl, T-14; Timothy C. Germann, Brad L. Holian, X-1-SMMP; Berni Alder, Lawrence Livermore National Laboratory</i>	
<b>Radiation Transport Methods and Algorithms .....</b>	<b>47</b>
Discrete Diffusion Monte Carlo for Radiation-Transport Simulations on Adaptive-Refinement Meshes.....	48
<i>Jeffery D. Densmore, Thomas M. Evans, Michael W. Buksas, CCS-2</i>	
Residual Monte Carlo Methods.....	50
<i>Thomas M. Evans, Scott Mosher, CCS-2</i>	
Implicit Monte Carlo Methods for Three-Temperature Physics .....	52
<i>Thomas M. Evans, Jeffery Densmore, CCS-2</i>	
Computational Cost Analysis of the Milagro Implicit Monte Carlo Code .....	54
<i>Scott W. Mosher, Timothy M. Kelley, CCS-2</i>	
Analysis of the Stability, Monotonicity, and Time Discretization Error of the Implicit Monte Carlo Method .....	56
<i>Scott W. Mosher, Jeffery D. Densmore, CCS-2</i>	
Verification and Validation of Radiation Transport Schemes .....	58
<i>Chris Fryer, Aimee Hungerford, Tom Evans, Mike Buksas, Todd Urbatsch, CCS-2; Darrell Peterson, Randy Kanzleiter, X-2</i>	
Student Projects in Radiation Transport Methods, Summer 2006.....	60
<i>Todd Urbatsch, Ed., CCS-2</i>	
Fourier Analysis of Parallel Block-Jacobi Splitting with Transport Synthetic Acceleration in Two-Dimensional Geometry.....	62
<i>Massimiliano Rosa, James S. Warsa, and Jae H. Chang, CCS-2</i>	

<b>Coupled Physics and Multiscale Methods and Algorithms.....</b>	<b>65</b>
Buoyancy-Driven, Variable Density Turbulence .....	66
<i>Daniel Livescu, CCS-2; J.R. Ristorcelli, Jr., X-3</i>	
Detailed Experimental Measurements of Turbulence Model Constants .....	68
<i>Arindam Banerjee, Malcolm J. Andrews, CCS-2</i>	
How Much Laser Power Can Propagate through Fusion Plasma? .....	70
<i>Pavel M. Lushnikov, Landau Institute and University of New Mexico; Harvey A. Rose, T-13</i>	
Self-Similarity, Scale-Separation, and Universality in Rayleigh-Taylor, Boussinesq Turbulence .....	72
<i>Natalia Vladimirova, University of Chicago; Michael Chertkov, T-13</i>	
Replicate Variation.....	74
<i>Michael L. Fugate, CCS-6; Timothy Wallstrom, T-13</i>	
Modeling Interfacial Surface Tension in Fluid Flow .....	76
<i>Marianne M. Francois, James M. Sicilian, CCS-2; Douglas B. Kothe, Oak Ridge National Laboratory</i>	
<b>Sensitivity Analysis, Error Propagation, Uncertainty Quantification, and Verification and Validation .....</b>	<b>79</b>
Statistical Models for Stockpile Health Assessment .....	80
<i>Alyson Wilson, Mike Hamada, Christine Anderson-Cook, Todd Graves, Scott Vander Weil, Benjamin Sims, Greg Wilson, CCS-6</i>	
Uncertainty Quantification for Simulation-Based Predictions .....	82
<i>Dave Higdon, Brian Williams, Jim Gattiker, Mike Fugate, CCS-6</i>	
Code Assessment for Adaptive Mesh Refinement Simulations.....	84
<i>Shengtai Li, T-7; William J. Rider, X-1</i>	
<b>Manufacturing Simulation and Computation.....</b>	<b>87</b>
Visualization of Complex Mechanical Models.....	88
<i>David Pugmire, HPC-4</i>	
Truchas: A Simulation Tool for Manufacturing.....	90
<i>ASC Telluride Project Team</i>	
<b>Applied Mathematical Sciences.....</b>	<b>93</b>
Multimaterial Interface Reconstruction from the Moment Data.....	94
<i>Vadim Dyadechko, Mikhail Shashkov, T-7</i>	
A Comparison of Multilevel Preconditioners for Solving Multimaterial Equilibrium Radiation Diffusion Problems on Locally Refined Grids .....	96
<i>Bobby Philip, T-7; Michael Permice, Idaho National Laboratory; Wayne Joubert, Xylon Software; Bryan Lally, CCS-2</i>	
An Efficient, Numerically Stable, and Scalable Parallel Tridiagonal Solver.....	98
<i>Christian Ketelsen, University of Colorado; Markus Berndt, J. David Moulton, T-7</i>	
Benchmarking Diffusion Discretizations on Polyhedral Meshes.....	100
<i>J. David Moulton, Konstantin Lipnikov, T-7; Jimmy Fung, Scott Runnels, X-3</i>	
Direct Search Methods with Equality Constraints .....	102
<i>David W. Dreisigmeyer, HPC-4</i>	

Numerical Schemes for Hydrodynamics Based on Multidimensional Riemann Solvers .....	104
<i>William W. Dai, HPC-4; Paul R. Woodward, B. Kevin Edgar, University of Minnesota</i>	
Tensor Multiplication on Parallel Computers .....	106
<i>Bryan Rasmussen, HPC-4</i>	
<b>Materials Science and Equation of State.....</b>	<b>109</b>
Evolution of the Equation of State Production Code OpenSesame .....	110
<i>Nicolas Bock, Eric Chisolm, Travis Peery, Scott Crockett, Carl Greeff, T-1</i>	
Lattice Vibrations of Tin Under Pressure .....	112
<i>Sven P. Rudin, T-1</i>	
Nanostructured Actinide Materials: Lead-Plutonium Superlattices .....	114
<i>Sven P. Rudin, T-1</i>	
Tabular Multiphase Equation of State for Plutonium .....	116
<i>Eric Chisolm, Carl Greeff, Shao-Ping Chen, Scott Crockett, J. D. Johnson, Sven Rudin, Duane Wallace, and John Wills, T-1; Jonathan Boettger, Kevin Honnell, X-1-SMMP; David Pimentel, HPC-1</i>	
Equation of State Developments in T-1 .....	118
<i>The T-1 EOS Team</i>	
Progress on the Campaign 2, Level-2 Milestone: Kel-F 800 Experimental Characterization and Model Development .....	120
<i>Brad T. Clements, T-1</i>	
Constitutive Modeling of Army Face Shields .....	122
<i>Brad T. Clements, T-1</i>	
Thermal Excitations and the Equation of State of Tantalum .....	124
<i>Carl W. Greeff, Sven P. Rudin, Scott D. Crockett, John M. Wills, T-1</i>	
A Sample of Some New Equations of State .....	126
<i>Scott Crockett, Leonid Burakovsky, Sven Rudin, T-1</i>	
Zero-Temperature Magnetic Transition in an Easy-Axis Kondo Lattice Model .....	128
<i>Jian-Xin Zhu, T-11; S. Kirchner, Rice University; R. Bulla, Universitat Augsburg; Qimiao Si, Rice University</i>	
Dynamically Driven Phase Transformation in Damaged Composite Materials: Micromechanics.....	130
<i>JeeYeon N. Plohr, Bradford E. Clements, T-1; Frank L. Addessio, T-3</i>	
<b>Atomic and Molecular Theory and Modeling .....</b>	<b>133</b>
Aging of the Binder in Plastic-Bonded Explosive PBX-9501 and Free-Radical Oxidation .....	134
<i>Joel D. Kress, T-12</i>	
Multi-Length Scale Modeling of Highly Filled Polymer Composites.....	136
<i>David E. Hanson, Cynthia Reichhardt, T-12; Axinte Ionita, T-1</i>	
Time-Reversible <i>ab initio</i> Molecular Dynamics .....	138
<i>Anders M. N. Niklasson, T-1; C. J. Tymczak, Texas Southern University; Matt Challacombe, T-12</i>	
Recent Developments in Modeling Atomic Processes in Dense Plasmas .....	140
<i>James Colgan, Joseph Abdallah, Jr., T-4; Ben Streufert, University of Virginia; Christopher J. Fontes, X-1-NAD</i>	
Quantum Molecular Dynamics Calculations of Opacities of Mixtures .....	142
<i>Dan Horner, Lee Collins, T-4; Joel Kress, T-12</i>	

<b>Nuclear Physics .....</b>	<b>145</b>
Iridium Cross-Section Evaluations.....	146
<i>Patrick Talou, Toshihiko Kawano, Shannon Cowell, T-16; Ronald O. Nelson, Matthew Devlin, Nik Fotiades, LANSCE-NS</i>	
High-Spin Isomeric States of Iridium and Gold .....	148
<i>Patrick Talou, Toshihiko Kawano, Shannon Cowell, T-16; Ronald O. Nelson, Matthew Devlin, Nik Fotiades, LANSCE-NS</i>	
Measurement and Analysis of the ${}^6\text{Li}(n,t){}^4\text{He}$ Cross Section .....	150
<i>Gerald Hale, T-16; Matthew Devlin, LANSCE-NS</i>	
Calculation of Charged-Particle Transport for NIF Mix Diagnostics.....	152
<i>Anna C. Hayes, T-16; Gerard Jungman, T-6; Paul A. Bradley, X-2</i>	
Calculations of the Proliferation-Resistance of Recycled Fuel.....	154
<i>Anna Hayes, Bill Wilson, T-16; Holly Trelue, X-1</i>	
Large-Scale Calculation of Fission Barrier Parameters for $171 \leq A \leq 330$ .....	156
<i>Peter Möller, Arnold J. Sierk, T-16</i>	
Fission Fragment Excitation of Nuclear Isomers in Hot Dense Environments.....	158
<i>Anna Hayes, Jim Friar, T-16; Gerard Jungman, T-6</i>	
Improved Neutron Fluence Spectra for Simulations of the Role of the Isomer of U-235.....	160
<i>Anna Hayes, T-16; Gerard Jungman, T-6; Merri Wood-Schultz, X-2</i>	
<b>High Explosives and Organic Materials.....</b>	<b>163</b>
Microstructure Effects on PBX Mechanical Response .....	164
<i>Scott G. Bardenhagen, T-14</i>	
Atomistic Simulations of Shock Waves in Polycrystalline Iron.....	166
<i>Kai Kadau, T-14; Timothy C. Germann X-1-SMMP; Peter S. Lomdahl, T-14; Robert C. Albers, T-11; J.S. Wark, A. Higginbotham, University of Oxford; Brad L. Holian, X-1-SMMP</i>	
Molecular Dynamics Comes of Age: 320-Billion-Atom Simulation on BlueGene/L .....	168
<i>Kai Kadau, Peter S. Lomdahl, T-14; Timothy C. Germann, X-1-SMMP</i>	
Shock-Front Broadening in Polycrystalline Materials .....	170
<i>John L. Barber, Kai Kadau, T-14</i>	
New Shock-Induced Phase Transformation in RDX.....	172
<i>Marc J. Cawkwell, Thomas D. Sewell, T-14</i>	
The Core Structure of the [001] Screw Dislocation in alpha-RDX .....	174
<i>Marc J. Cawkwell, Thomas D. Sewell, T-14</i>	
The Mechanics of LANL Foam Pads .....	176
<i>Scott G. Bardenhagen, T-14; Brian M. Patterson, MST-7; Carl W. Cady, MST-8; Mathew W. Lewis, WT-1; Dana M. Dattelbaum, DE-9</i>	
Elastic Stress and Strain Distributions in Polymeric Materials.....	178
<i>Panagiotis Maniadis, T-CNLS; Turab Lookman, Kim Ø. Rasmussen, T-11; Edward M. Kober, T-14</i>	
Gravitational Effects in Reactive Flows: From Physics to Large Eddy Simulation Modeling.....	180
<i>Sergei Chumakov, T-CNLS; Natalia Vladimirova, University of Chicago</i>	

<b>Production Desktop, Supercomputing, and Support .....</b>	<b>183</b>
Managing Software Tools in a High-Performance Computing Environment.....	184
<i>David R. Montoya, Georgia A. Pedicini, David O. Gunter, HPC-4</i>	
Tripod Project as a Process.....	186
<i>Richard L. Graham, CCS-1; Josip Loncaric, HPC-5; Steven R. Shaw, HPC-3</i>	
Quantifying Checkpoint Efficiency .....	188
<i>John T. Daly, HPC-4</i>	
Unified Data Model: A Library for Parallel I/O and Data Management.....	190
<i>William W. Dai, Robbie Aulwes, HPC-4</i>	
Applications Using the Unified Data Model .....	192
<i>Robbie Aulwes, William Dai, HPC-4</i>	
Creating Future Network Architectures .....	194
<i>Robert Martinez, Tim Merrigan, Denny Rice, Todd Bowman, CTN-5</i>	
The ASC Roadrunner System at LANL: Near-Term Capacity with a Future Hybrid Accelerated Option .....	196
<i>Manuel Vigil, HPC-DO</i>	
<b>High-Performance Computing Systems Operations .....</b>	<b>199</b>
Concurrent Application Testing Based on a Throughput Metric.....	200
<i>John T. Daly, HPC-4</i>	
Facilitating High-Throughput ASC Calculations.....	202
<i>John T. Daly, HPC-4</i>	
Seventeen Years of the New Mexico Supercomputing Challenge .....	204
<i>David Kratzer, HPC-3</i>	
Request Tracker (RT): A Tool for Software Change Control .....	206
<i>Georgia A. Pedicini, David R. Montoya, HPC-4</i>	
Monitoring Software Health in a Heterogeneous HPC Environment.....	208
<i>David R. Montoya, Georgia A. Pedicini, David O. Gunter, Jennis Pruett, Jennifer M. Frye, HPC-4</i>	
<b>Author Index .....</b>	<b>211</b>

## Organizational Abbreviation Legend

### **Computer, Computational, and Statistical Sciences Division**

CCS-1 Computer Science for High-Performance Computing  
CCS-2 Computational Physics and Methods  
CCS-3 Information Sciences  
CCS-6 Statistical Sciences  
CCS-DO Division Office

### **Computing, Telecommunications, and Networking Division**

CTN-DC Departmental Computing (CTN-1, -2, and -3)  
CTN-1 Departmental Computing  
CTN-2 Departmental Computing  
CTN-3 Departmental Computing  
CTN-4 Telecommunications  
CTN-5 Network Engineering  
CTN-DO Division Office

### **High-Performance Computing**

HPC-1 Scientific Software Solutions  
HPC-2 Computing Operations and Support  
HPC-3 High-Performance Computing Systems  
HPC-4 High-Performance Computing Environments  
HPC-5 High-Performance Systems Integration

### **Theoretical Division**

T-1 Equation of State and Mechanics of Materials  
T-3 Fluid Dynamics  
T-4 Atomic and Optical Theory  
T-6 Theoretical Astrophysics  
T-7 Mathematical Modeling and Analysis  
T-8 Elementary Particles and Field Theory  
T-10 Theoretical Biology and Biophysics  
T-11 Condensed Matter and Statistical Physics  
T-12 Theoretical Chemistry and Molecular Physics  
T-13 Complex Systems  
T-14 Explosives and Organic Materials  
T-15 Plasma Theory  
T-16 Nuclear Physics  
CNLS Center for Nonlinear Studies  
T-DO Division Office

# Computer Science

Computer science is a broad field that encompasses theory, mathematical activities such as design and analysis of algorithms, performance studies of systems and their components, and the estimation of reliability and availability of systems by probabilistic techniques. The field encompasses the study of computers, their design, and their uses for computation, data processing, and systems control including design and development of computer hardware and software, and programming.

In the past, the development and advance of computational engines (and the development of computer science as a field) was largely driven by the needs of users in the scientific arena. Not surprisingly, the early development of supercomputers was largely driven by the nuclear weapons program. However, in the past decade or so, the advance of computing has been driven by different demographics in the financial and gaming industries.

There are essentially three major elements within computer science: (1) architecture, which includes all levels of hardware design and the integration of hardware and software components to create a complete computer system; (2) the software, or computer programs that cause the computer to carry out tasks, which broadly includes software design and engineering, programming language, operating systems, information systems and databases, artificial intelligence, computer graphics, and visualization; (3) theory, which includes computational methods and numerical analysis of both data structures and algorithms.

Edsger Dijkstra, a Dutch computer scientist who won the Turing Award in 1972, perhaps described the field best by saying, "Computer science is no more about computers than astronomy is about telescopes." Computer science is still very much an evolving field, and the subsections of this document speak to the rich technical diversity in the field today. These brief articles outline the state-of-the-art in computer science research and development at Los Alamos National Laboratory at the present time.

The section on “computational architecture exploration” is focused on techniques for enhancing application throughput on the latest computing systems as well as the exploration of new and emerging technologies that will comprise tomorrow’s supercomputing systems for applications of interest. Not surprisingly, this involves hardware and computing technology and software.

The section on “libraries and tools” is centered on software tools for the application developer, and software engineering techniques that enhance the agility of computational physics software—a necessity with today’s rapid changes in computing technology.

The section on “theoretical and applied foundations of data and management” speaks to theory and algorithmic advances that couple two seemingly disparate applications—speech recognition and image processing for the nuclear weapons program. The article scratches the surface of information sciences and is an indicator of the potential energy behind the emerging information sciences thrust at Los Alamos.

The “advanced visualization and data management” section is also strongly related to information sciences, and is centered on the advanced development, tools, and infrastructure geared toward the extraction of knowledge—not simply information—from vast arrays of data produced by simulation tools, which generate volumes of data akin to the entire Library of Congress in a single simulation.

Finally, the last section titled “system software” is about the software at the level of the “metal” in a computer—the operating system and kernel that lie just above the computer hardware and below all of the applications that run on it.

# Open MPI: A High-Performance, Fault-Tolerant Message-Passing Interface

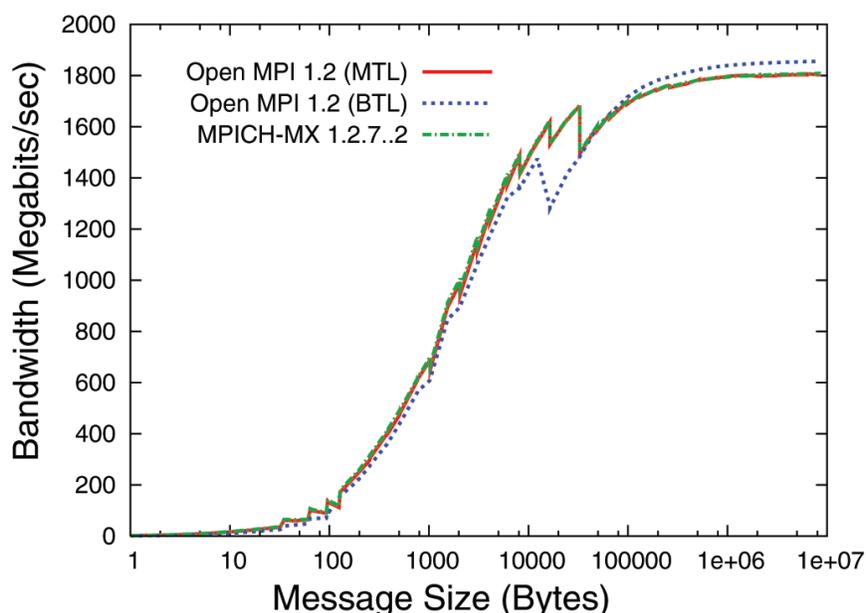
Brian Barrett, Ralph Castain, Galen Shipman, CCS-1

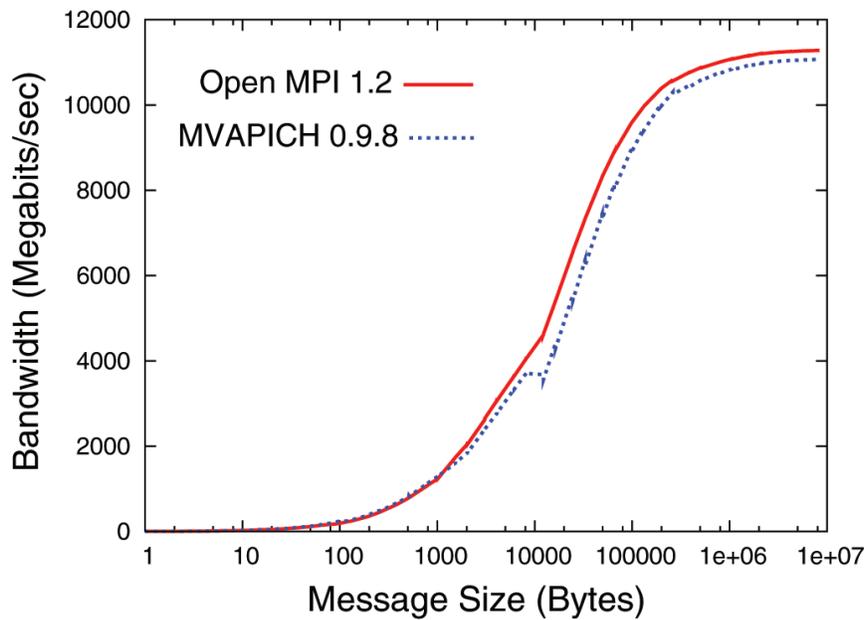
**O**pen MPI is a mature message-passing interface (MPI) implementation developed by Los Alamos National Laboratory (LANL) in collaboration with a number of academic, industry, and national-laboratory partners. Open MPI is designed to run at the large scales currently found within the Advanced Simulation and Computing program. Recently, Open MPI successfully demonstrated its performance at scale at Sandia National Laboratories, where it was an integral part in boosting the high-performance linpack (HPL) performance to 53.00 teraflops, 84.7% of peak performance up from 38.27 teraflops and 71.4% of peak a year earlier. Open MPI is the production MPI at LANL supporting several clusters including the Roadrunner base system.

The Open MPI project was originally started by LANL; Indiana University, Bloomington; and the University of Tennessee, Knoxville. Involvement in the project has since grown to include other national laboratories, universities, and commercial vendors. Support for a number of interconnects, including InfiniBand (MVAPI and Open Fabrics), Myrinet (GM and MX), Cray Portals, TCP, and shared memory, is provided in a single unified release. The Open Run-Time Environment, developed in unison with the Open MPI project, provides a scalable run-time environment with support for process launching utilizing rsh/ssh, PBS/Torque's TM interface, BProc, and SLURM.

The performance of Open MPI is competitive with other high-performance MPI implementations such as MPICH/MX and MVAPICH. Figure 1 illustrates

**Fig. 1.**  
The performance of Open MPI over Myrinet MX, which is comparable to MPICH/MX.





**Fig. 2.**  
The performance of Open MPI over Mellanox InfiniBand, which is comparable to MVAPICH.

the performance of Open MPI over Myrinet MX, which is comparable to MPICH/MX. Figure 2 illustrates the performance of Open MPI over Mellanox InfiniBand, which is comparable to MVAPICH. Open MPI provides 1-byte latency of  $3.11 \mu\text{s}$  using Mellanox InfiniBand vs  $3.15 \mu\text{s}$  for MVAPICH, and  $3.15 \mu\text{s}$  using Myrinet MX vs  $2.97 \mu\text{s}$  for MPICH/MX. These results indicate that a single MPI implementation can provide good performance across multiple architectures while providing the scalability and stability of a production-grade MPI implementation.

The Open MPI team at Los Alamos is currently researching a number of issues related to running in large-scale InfiniBand environments. These include adaptive resource allocation, protocol selection and recovery from spurious host channel adapter (HCA) resets that are seen in large-scale InfiniBand clusters. In addition to generic large-scale InfiniBand clusters, the Los Alamos team is also researching issues related to the full-scale Roadrunner configuration. These include

optimizations for heterogeneous data transfer, run-time support for dynamic processes across heterogeneous processors, and efficient usage of InfiniBand in nonuniform memory architectures.

*For more information contact Galen Shipman at [gshipman@lanl.gov](mailto:gshipman@lanl.gov).*

#### **Funding Acknowledgements**

NNSA's Advanced Simulation and Computing (ASC), System Software and Support, and Tools and Capabilities programs.

# Hiding Communication Overheads in High-Performance Computing

Jose Carlos Sancho, Kevin J. Barker, Darren J. Kerbyson, Kei Davis, CCS-1

Communication delays are becoming a critical factor in the performance of large-scale applications that can utilize many thousands of processors at a time. Performing communication at the same time as computation, commonly known as overlapping, can partially or even fully hide these communication delays leading to an immediate increase in achievable performance.

In recent work in the Performance and Architecture Lab (PAL), a part of the Computer Science for High-Performance Computing group at Los Alamos National Laboratory, we have developed methods to quantify the potential performance improvement from overlapping communication with computation. Insights from this work have shown that overlapping is possible and that networks in high-performance computing (HPC) could be more cost effective.

The methods are predictive and thus designed to examine the potential performance improvements prior to any application recoding. Measurements are made of the computation components that are required prior to a communication and those performed after a

communication. The potential overlap is calculated by comparing the computational time available with the communication delays that will be experienced on the target network. A simple example is shown in Fig. 1 where the potential overlap is the difference in time between the point at which data is produced on the sender and the point at which it is required on the receiver.

Initial work considered four large-scale applications: POP (Ocean Modeling), SAGE with and without adaptive mesh refinement (shock-wave hydrodynamics), Sweep3D (deterministic transport), and HYCOM (Ocean Modeling). A medium-sized system containing 1024 processing cores with an InfiniBand network was used in the analysis.

The predictive nature of the method can be seen in Fig. 2 in which networks of varying bandwidths are considered (from 1 MB/s up to 5 GB/s). The performance is normalized to an ideal case in which communication has zero performance penalties. When overlapping is exploited the potential performance quickly attains near ideal performance even for networks with low bandwidths (in this case at ~400 MB/s). When overlapping is not exploited then the achievable performance gradually nears that of the overlapping case as the network bandwidth increases. This example is for SAGE in which message sizes can be between 10 and 100s of KB.

This result is important in that it shows that networks with lower bandwidths, and hence lower costs, could be utilized

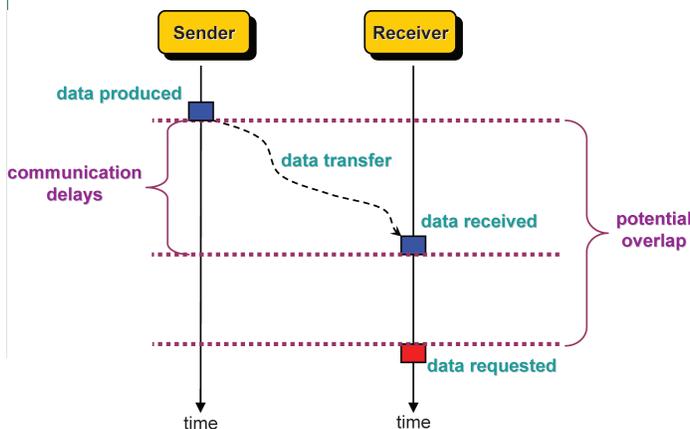
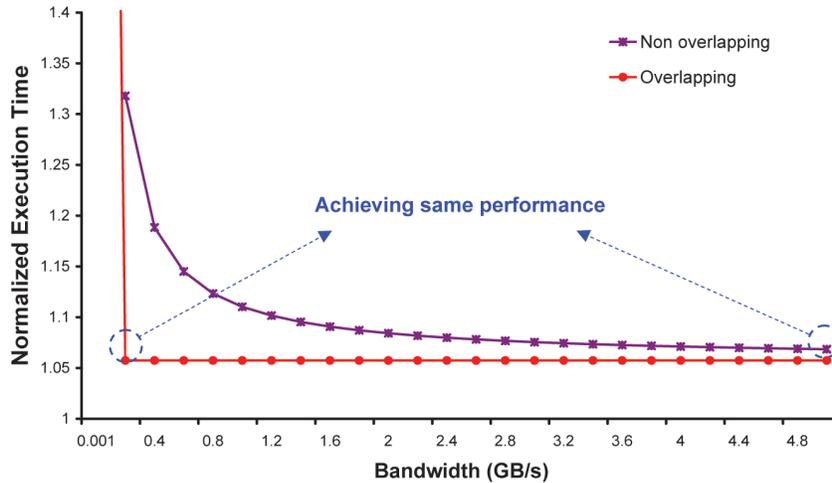
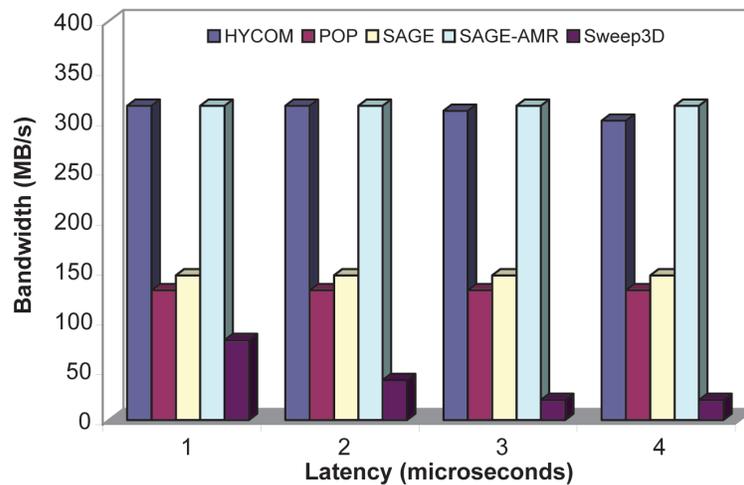


Fig. 1.  
Example showing  
potential overlap.



**Fig. 2.** Potential of overlap in SAGE for networks of various bandwidths.



**Fig. 3.** Bandwidth network characteristics when overlapping is exploited deliver the same performance as the case of nonoverlapping on a 5-GB/s network.

to achieve the same performance as more expensive networks that have higher bandwidths; that is, if overlapping can be exploited by the application. Potential overlap does exist in the four applications analyzed and shows that a reduction in network bandwidth and/or increase in latency can be tolerated. Figure 3 summarizes these results and shows the network characteristics (in terms of bandwidth and latency) that are required for each application when using overlap, compared to the case with no overlap and a network with 5 GB/s bandwidth, and 4  $\mu$ s latency. This work has demonstrated that future HPC systems could benefit from the use of cost-effective networks that support overlapping of communication with computation without negatively impacting application performance.

This work was presented by PAL in the technical program at the Supercomputing Conference (SC06) in November 2006 [1].

*For more information contact Jose Carlos Sancho at [jcsancho@lanl.gov](mailto:jcsancho@lanl.gov).*

[1] J.C. Sancho, et al., "Quantifying the Potential Benefit of Overlapping Communication and Computation in Large-Scale Scientific Applications," Los Alamos National Laboratory report LA-UR-06-3109 (May 2006).

**Funding Acknowledgements**  
NNSA's Advanced Simulation and Computing (ASC), and Defense Advanced Research Projects Agency (DARPA), High Productivity Computing Systems Program.

# A Novel Single Projection Method to Calculate View Factors for Radiosity, and a GPU-Based Implementation in a Large Multiphysics Code

Sriram Swaminarayan, John Turner, CCS-2

**W**e describe a new method for calculating view factors for radiosity using a single Cartesian projection, and its implementation on graphics hardware to accelerate heat transfer in Truchas, a large multiphysics flow and solidification code developed at Los Alamos National Laboratory. By utilizing the hardware accelerated z-buffer of the Graphical Processing Unit (GPU) we are able to achieve speedups of the order of 30 times in the view factor calculation.

## Algorithm

Our single projection method is similar to the Hemicube method [1] in that we do Cartesian projections, but is more efficient since instead of five projections, we only project onto a single plane. This requires modification of the original hemicube equations to explicitly include the distance to the imaging plane as follows:

$$dF = \frac{z_o^2 \Delta^2}{\pi (r^2 + z_o^2)^2} \quad (1)$$

where  $dF$  is the contribution to the view factor matrix by any pixel on the imaging plane,  $z_o$  is the height of the imaging plane from the center of the face of interest,  $\Delta$  is the size of the pixel, and  $r$  is the radial distance to the pixel from the  $z$ -axis.

As is evident from Eq. 1, the pixels closest to the  $z$ -axis will contribute the most to the view factor matrix (inverse dependence on  $r^4$ ). Consequently, we use a finer mesh closer to the center

of the imaging plane and a coarser mesh further away. On the GPU this is achieved by simply changing the field of view while keeping the window size fixed. This allows us to achieve high accuracy in the view factor calculation where it is needed most without sacrificing speed. Figure 1 demonstrates this for a geometry of concentric spheres where three levels of refinement have been used to calculate view factors. The faces are colored by their IDs and the contents of this buffer are read back onto the CPU to combine contributions into the overall view factor matrix. A significant advantage of our method is that we can selectively refine the view factor matrix along directions that need the high resolution due to increased detail in that direction, while at the same time keeping the mesh coarse in other areas.

## Implementation

To interface the GPU code into Truchas, which is written in modern Fortran (F90), required a platform-independent library for generating windows and drawing to them using OpenGL. We decided against using OpenGL Utility Toolkit (GLUT) for this since it does not give the application control over the event loop, and it was not easy to recast the F90 code in such a way that it could be called from the display routine of GLUT. We instead wrote a simple library that opens a window on the client machine, and draws to it with the geometry given. This way, the physics code itself has no graphics code embedded in it and we can easily eliminate the GPU-based extensions

when building on machines without powerful GPUs.

### Results

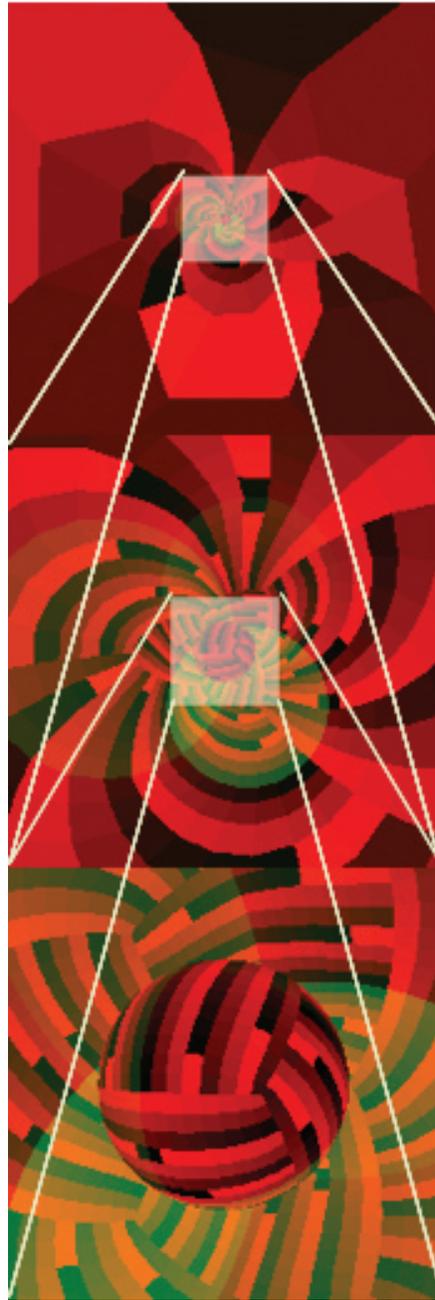
Figure 2 shows the times for calculating the view factor using the Hemicube method on a 3.4 GHz 64-bit Xeon and using the plane projection method on an NVIDIA Quadro FX1400 GPU for a problem containing two concentric spheres. For any given problem size, the GPU implementation is at least 30 times faster than the CPU for a comparable resolution.

*For more information contact Sriram Swaminarayan at [sriram@lanl.gov](mailto:sriram@lanl.gov).*

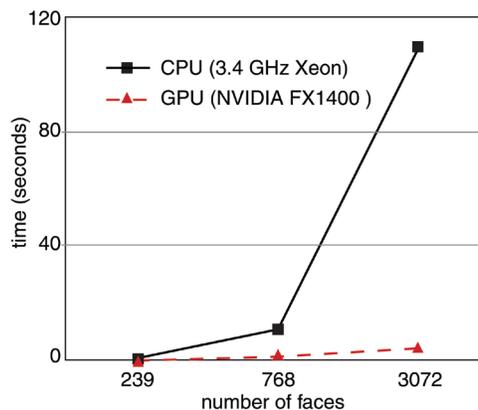
[1] M.F. Cohen, D.P. Greenberg, "Hemi-Cube: A Radiosity Solution For Complex Environments," *Comp. Graphics* **19**, 3 (1985).

### Funding Acknowledgements

NNSA's Advanced Simulation and Computing (ASC), Telluride Project; and Laboratory Directed Research and Development.



**Fig. 1.** Snapshot of one face of the view factor calculation. The faces are colored by their IDs. The frame buffer is read back and the sum of the contributions of each face is used to calculate the view factors for that face. The three panes in the image represent three fields of view used in obtaining this image.



**Fig. 2.** Time taken for the Hemicube on the CPU and the plane projection method on the GPU. The GPU computes the view factors at least 30 times faster than the CPU.

# Sparse Matrix-Vector Multiplication and Conjugate Gradient Algorithm on Hybrid Computing Platforms

David DuBois, Andrew DuBois, Carolyn Connor Davenport, Stephen Poole, HPC-5

**S**parse matrices, derived from systems of partial differential equations (PDEs), occur in physics, mechanical engineering and other fields where a physical phenomenon needs to be mathematically described. These PDEs are used to understand phenomena such as fluid flow, the growth of crystals, gravitation, diffusion, and the behavior of electromagnetic fields.

The solution to a nonsingular linear system  $Ax=b$  lies in a Krylov space whose dimension is the degree of the minimal polynomial of  $A$  (where  $A$  is a matrix,  $x$  &  $b$  are vectors). If this minimal polynomial of  $A$  has a low degree, a Krylov method has the potential of rapid convergence [1]. When solving a system of linear equations  $Ax=b$ , if the coefficient matrix  $A$  is large and sparse, the time required to solve the system by direct methods is too high and requires too much storage.

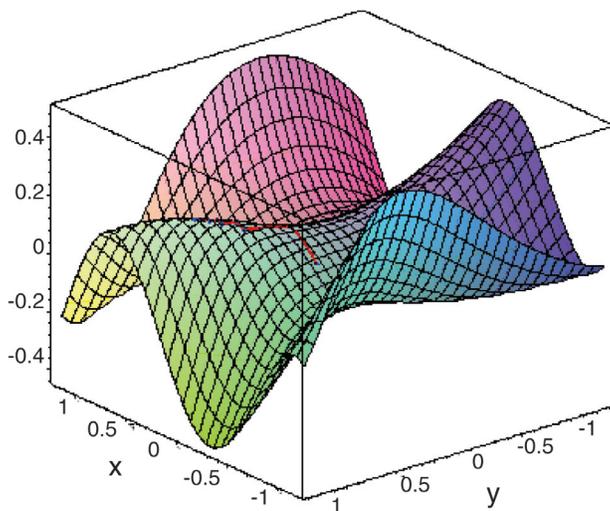
Krylov methods, or solvers like Conjugate Gradient (CG) [2], are particularly well suited for use on large-scale scientific simulation codes that are defined by sparse linear systems. These codes call solvers such as CG, which in turn repeatedly perform

sparse matrix-vector multiplication (SMVM) operations to converge on a result for each time step [3]. To facilitate convergence, CG uses the gradient descent method to minimize a residual vector (Fig. 1) [4].

Double-precision floating point SMVM is the time dominant computational kernel used in iterative solvers like CG. It is imperative that the SMVM operations be computed efficiently, yet the poor data locality exhibited by sparse matrices along with the high memory bandwidth requirements of SMVM result in poor cache utilization in general-purpose processors. Field programmable gate arrays (FPGAs) and the heterogeneous multicore architecture of the Cell processor offer possible alternatives. The Cell architecture is a hybrid computing innovation used in Roadrunner, the IBM supercomputer now being installed at Los Alamos National Laboratory for the Advanced Simulation and Computing (ASC) program.

We have developed a FPGA-based implementation of a double-precision, non-preconditioned, conjugate gradient solver for three-dimensional finite-element or finite-difference methods [5]. Our work uses the SRC Computers, Inc., MAPStation hardware platform (Fig. 2) and the "Carte" software programming environment. We have demonstrated that an FPGA-based system can perform on par with today's processors while running over 30 times slower (i.e., 100 MHz vs 3.4 GHz), which makes the FPGA-based design much more power efficient [6]. This is possible because an FPGA-based system can be designed to more optimally match the computational units to available memory bandwidth, providing a more balanced system. It still suffers from the basic physical constraints of limited I/O and limited memory bandwidth

**Fig. 1.** The "gradient" descent method applied to a function. This image shows the surface of the function (3-D interpretation) [4].



for this and other memory bandwidth-intensive classes of problems. To efficiently use the peak computational capability of FPGA, hybrid (e.g., Cell), or CPU-based systems for this class of problems requires tremendous amounts of memory bandwidth.

**Conclusion**

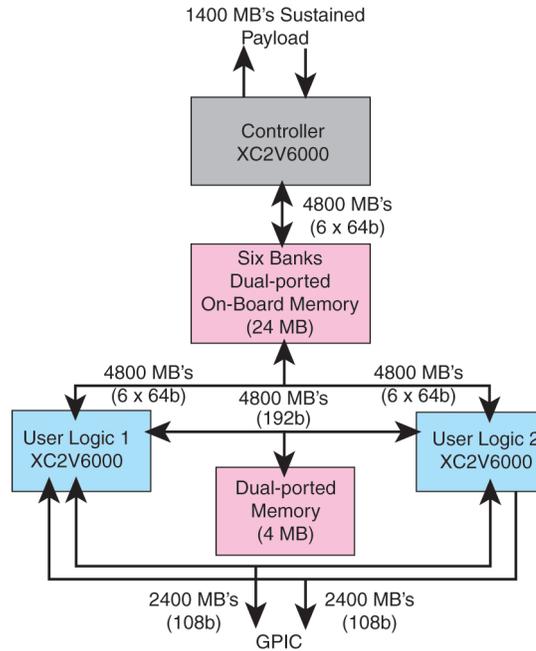
The FPGA results we have presented are deterministic (i.e., not variable) and will scale directly with any improvements in the system user logic frequency, memory bandwidth, and/or memory depth. We are working to port a preconditioned CG algorithm onto the Cell processor. Our goal is to exploit the improved memory bandwidth available on the Cell processor, giving us increased performance over both the traditional, cache-based implementations and our FPGA result.

*For more information contact Andrew DuBois at [ajd@lanl.gov](mailto:ajd@lanl.gov).*

[1] I. Ipsen and C. F. Meyer, "The Idea Behind Krylov Methods," *Am. Math. Mon.* **105**, 10, 889 (1998).  
 [2] J. Shewchuk, "An Introduction to the Conjugate Gradient Method Without the Agonizing Pain," School of Computer Science, Carnegie Mellon University (1994).  
 [3] J. Mellor-Crummey and J. Garvin, *Int J High Perform C*, **18**, 2, 225 (Summer 2004).  
 [4] [http://en.wikipedia.org/wiki/Gradient\\_descent](http://en.wikipedia.org/wiki/Gradient_descent) (definition)  
 [5] D. DuBois, et al., "An Implementation of the Conjugate Gradient Algorithm on FPGAs," to be published.  
 [6] D. DuBois, et al., "Sparse Matrix-Vector Multiplication on a Reconfigurable Supercomputer," submitted for publication to IEEE Transactions on Parallel and Distributed Systems, LA-UR-06-5312.

**Funding Acknowledgements**

This project was supported through the LANL ASC Weapons-Supported Research project.



**Fig. 2.** SRC Computers, Inc., MAPStation's MAP Processor block diagram [6].

## Exploring Solutions to ASC Simulation Code Input/Output Issues

Steve Hodson, James Nunez, HPC-5

**T**here is a conflict between the Los Alamos National Laboratory Advanced Simulation and Computing (ASC) simulation codes and the supercomputers on which they run. The codes can run anywhere from weeks to months to years, yet the computer systems rarely stay up for this amount of time. One method for the simulations running to make progress is for the codes to periodically write files to storage that have all the information in them to restart the run. The stored files are called restart, or checkpoint, files. Normally, each process writes all characteristics for all elements for which it is responsible to a single file. Thus, if a machine goes down, the code can restart close to the place where it stopped with the use of the information in the checkpoint file. In this manner, days, weeks, or months of work are not lost. Since the checkpoint files can be very large, on the order of gigabytes to terabytes, the time to write a checkpoint file is anywhere from seconds to tens of minutes, and the user must balance this cost against the mean time to interrupt (MTTI) of the machine to figure out how often to write a checkpoint file. In addition, the MTTI will vary by machine, maturity of machine, by the size of the job, and the checkpoint time will depend on the application.

Unfortunately, the manner in which many codes write checkpoint files is with very small writes, on the order of tens to a few hundred kilobytes, and the writes are normally not aligned on page boundaries. Not only do they write in this small and unaligned manner, but all processes write to one portion of a single file at one time, i.e., the writes are strided. The codes have a variety of good reasons to write in this manner; it being the simplest way to restart a job from N processors to M processors, and a friendly format for visualization.

Since the simulations cannot afford to lose data, most file systems employ a reliability scheme called RAID 5 that can rebuild data in the event of a failed disk drive by keeping the parity of data stored across all drives. Due to this reliability and the code's I/O (input/output) patterns, every write to disk is possibly a fetch-on-write or read the blocks on disk, update these blocks with the new data, and write to disk and thus is more time consuming. For this reason the HPC-5 File System team explored the benefits and risks of a nonparity-based reliability scheme.

With real application in hand, the file system team, including the authors, tested a variety of I/O techniques to speed up the write time for real applications. Steve Hodson brought a deep understanding of the underlying physics of the applications to the team, and James Nunez brought an understanding of I/O middleware and file systems. The first step in eroding the decade-long problem of slow write bandwidth due to small unaligned I/O, was to analyze the I/O patterns of the application. Analyzing the application's I/O patterns entailed instrumenting the I/O routines and coming up with a profile for a typical run. With that analysis, potential areas that would impact the application's time for writing checkpoint files were found. The initial problem the authors were given was a code that was getting 1 MB/s of write bandwidth, which is intolerably slow. The analysis showed that all processes were writing to a small region of a single file with very small blocks of data. The team did a study of the impact of data aggregation at the middleware (MPI-IO) level; i.e., sending data to a small number of processes and having only that subset write to the file. This study confirmed previous results that having many writers can actually slow I/O

rates. Reducing the number of writers led to a 20-time improvement in bandwidth; to 20 MB/s. Due to rigorous testing and parameter studies, the data aggregation parameters were made the default in Los Alamos Message Passing Interface (LA-MPI) on all production machines.

Even this improved data rate was unacceptable. So Steve continued to search for ways to improve the application's I/O rates. In conjunction with testing James was doing, they realized that their bandwidth results for very similar tests were not matching. An analysis of the differences in their testing showed that they were using different versions of MPI, MPICH vs LA-MPI, which led to the discovery that LA-MPI had a very inefficient implementation of a subroutine that was called hundreds of times by the I/O routines. Once the LA-MPI team optimized the I/O subroutine, the application's write rate increased to approximately 50 MB/s.

Next, Steve undertook the job of activating asynchronous I/O in MPI and a study of how this would affect the application's I/O rates. Although the work with asynchronous I/O helped the bandwidth, it was clear that there were other factors impeding the I/O performance.

With help from the file system vendor Panasas, the team tested a nonparity-based file layout called RAID0. The team conducted a study of collective, independent synchronous, and independent asynchronous I/O against the alternate file layouts. The study concluded that the RAID0 layout alone helps the application, but combined with

asynchronous independent I/O (Fig. 1), the application can achieve over 450 MB/s!

The next step for this team is to work with Panasas to bring a nonparity-based reliable file layout into its product and into the Laboratory. This work has already begun, and testing on a RAID 10, mirror and stripe data across multiple disks scheme, already looks promising.

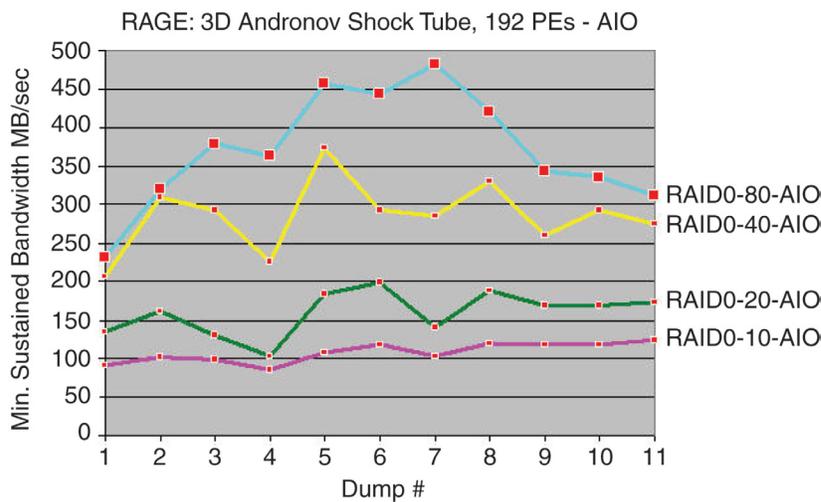
**Conclusion**

Through persistence, experimentation, and methodical testing, significant progress has been made toward understanding and solving the small and unaligned I/O problem that plagues many of ASC major codes. The File System team was able to demonstrate an average of 350 times improvement in write bandwidth for a major code team.

*For more information contact Stephen Hodson at [swh@lanl.gov](mailto:swh@lanl.gov).*

**Funding Acknowledgements**

This research was supported by the NNSA tri-Lab Advanced Simulation and Computing Program.



**Fig. 1.** The top turquoise line shows the dramatic improvement in speed of the application write time on a super-computer when RAID0 is used in combination with a specific I/O method.

## New Approaches to Fault-Tolerant Computing

Anders Hansson, Gabriel Istrate, Christian Reidys, Allon G. Percus, CCS-3; Stefan Boettcher, Bruno T. Goncalves, Emory University; Shiva P. Kasiviswanathan, Carnegie Mellon University

**F**ault tolerance in a parallel and distributed computing architecture relies on rapidly matching processes with resources across a network. This problem, called load balancing, is largely unsolved at the scale needed to support routine petaflop computing. Load balancing is closely related to the discrete optimization problem of graph partitioning. Imagine that each process is represented by a node in a network and that communication between processes is represented by a link. Mathematically, the communication patterns of the processes form a directed acyclic graph, shown in Fig. 1. The goal in graph partitioning is to partition the graph into a certain number of subcomponents, with as few links as possible connecting them. When one imposes additional global constraints, such as the “balanced partition” requirement that all subcomponents be the same size, the problem rapidly becomes computationally intractable as the graph size increases.

Load-balancing problems may be studied under the following generative model: create a dependency graph by placing, independently and with some probability  $p$ , a link between each pair of nodes. Reaching a detailed understanding of the solutions to the balanced partitioning problem, even on simplified random instances of this sort, is a prerequisite to modeling appropriately structured graphs. We have studied characteristics of how these solutions “cluster” in solution space, for the case of partitioning the graph into two balanced

subcomponents. It follows from standard random graph arguments that given  $n$  nodes, if  $pn < 2 \log 2$ , processes can almost always be partitioned perfectly without any links connecting the subcomponents. We have shown that in this case, given a straightforward definition of distance in solution space, all solutions lie within a single “cluster” where each solution is close to others [1].

This has significant implications for the analysis of graph partitioning using methods from statistical physics. The existence of a single cluster enables a key tool for algorithmic analysis, known as *replica symmetry*, to be used for this problem. Our study has resulted in a new algorithm, known as core peeling [2], that reliably finds near-optimal solutions in time scaling quadratically in  $n$  (Fig. 2). Thus, it makes the problem computationally tractable. We expect that a theoretical analysis of the core peeling algorithm could improve what is currently the best theoretical bound on the optimal solution, dating from the 1980s.

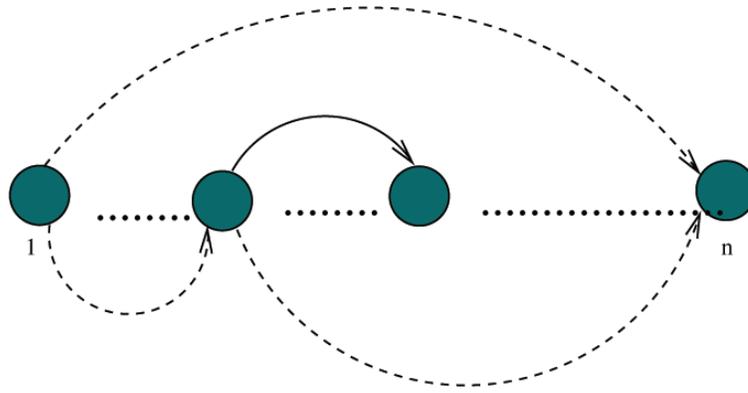
Additionally, we have developed a load-balancing algorithm for computer networks that is based on a local load-redistribution protocol, redistributing a given processor’s data packets according to the value of a “load function” at neighboring nodes [3]. Our analysis shows that under certain system assumptions, the algorithm functions at close to optimal network capacity. Since the approach is robust to a large class of update schedules, it can

provide a methodology for near-optimal dynamic redistribution of computational tasks—an instrumental tool for fault-tolerant network computing.

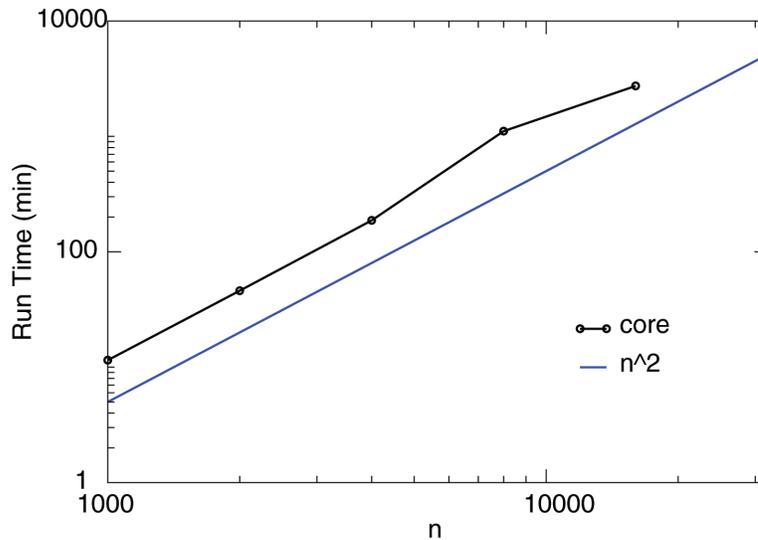
*For more information contact Allon G. Percus at [percus@lanl.gov](mailto:percus@lanl.gov).*

[1] G. Istrate, et al., "The Cluster Structure of Minimum Bisections of Sparse Random Graphs," Los Alamos National Laboratory report LA-UR-06-6566 (2006).  
 [2] B.T. Goncalves, et al., "The Core Peeling Algorithm for Graph Bisection: An Experimental Evaluation," Los Alamos National Laboratory report LA-UR-06-6863 (2006).  
 [3] A. Hansson and C. Reidys, "A Class of Load Balancing Network Routing Protocols," in *Proceedings of European Modeling and Simulation Symposium*; part of International Mediterranean Modeling Multiconference, Barcelona, Spain, Oct. 2006, pp. 37–42. Invited Paper. Also available as Los Alamos National Laboratory report LA-UR-06-5281 (July 2006).

**Funding Acknowledgements**  
 Weapons Supported Research,  
 Computer Science Research Foundation.



**Fig. 1.** Communication patterns on  $n$  processes represented as directed acyclic graph.



**Fig. 2.** Scaling of running time for core peeling algorithm, compared with  $n^2$ .

## Scientific Application Development using Eclipse and the Parallel Tools Platform

Greg Watson, Craig E. Rasmussen, CCS-1

**S**cientific application developers face many problems in today's parallel computing environments. Despite the perception that there are few tools supporting high-performance computing, the opposite is in fact true. Unfortunately, the tools are often difficult to use, support only a few platforms, and in many cases, are incompatible with each other.

An integrated development environment (IDE) is generally considered best practice in the majority of the (nonscientific) software development industry. Eclipse is an open-source IDE that provides a portable, robust, commercial quality environment for a wide range of software development activities. It supports multiple languages, including Java, C, C++, and Fortran, provides a syntax-aware editor, code refactoring, incremental code compilation, source-level debugging, and integrated support for source control systems such as CVS and Subversion.

To assist scientific application developers to improve their productivity, we created the Eclipse Parallel Tools Platform (PTP) project (<http://eclipse.org/ptp>). The aim of this project is to establish a common and portable IDE across a wide range of parallel computing platforms. An important aspect of PTP is that it remains agnostic to the tools actually deployed on the machines, such as compilers, linkers, job schedulers, runtime systems, and performance analysis tools. This ensures portability of the platform, and also provides a tool integration framework that allows the individual tools to share data and functionality.

In addition to the standard features supported by the Eclipse platform, PTP adds a range of functionality that enhances the ability of software engineers to develop codes for parallel machines. These features include:

### **Tools to aid MPI and OpenMP programmers.**

PTP includes tools that simplify the development of MPI and OpenMP programs by providing special content assistance and context-sensitive help. In addition, the tools provide static analysis features that provide advanced error checking and analysis of parallel programming constructs.

### **The ability to monitor and control parallel jobs.**

PTP adds support for monitoring the status of a parallel computer system (such as which nodes are available, etc.) and launch parallel jobs onto the system. Figure 1 shows the Eclipse view of a parallel machine. The status and output from processes of the job can be monitored, and the job can be terminated if necessary. Recent work has added support for submitting jobs through job-scheduling systems such as LSF and MOAB.

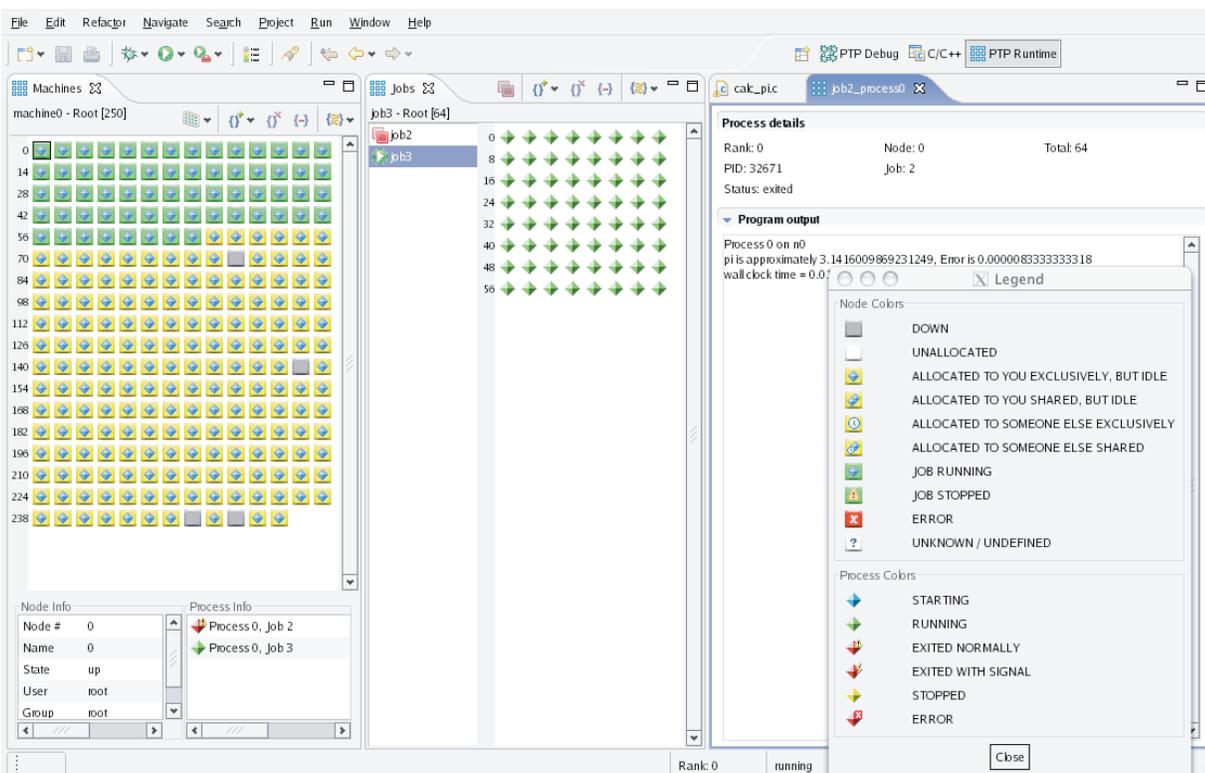
### **The ability to debug parallel programs.**

PTP adds an integrated, scalable, parallel debugger to the IDE. The debugger is launched with the click of a button, and provides the ability to control many processes simultaneously. Figure 2 shows a typical debug session. It supports traditional debugging commands, such as breakpoints, viewing variables, etc., as well as some innovative functions for dealing with multiple processes. The debugger is designed to scale to jobs with many thousands of processes.

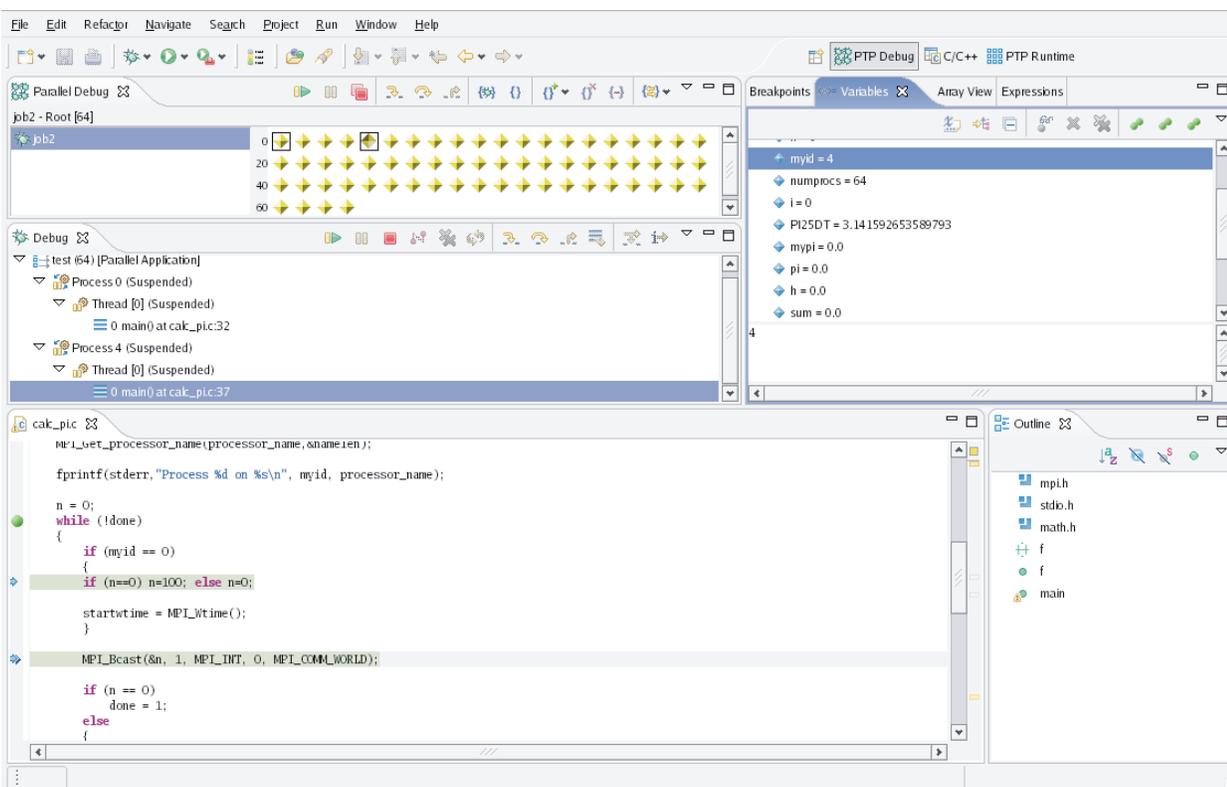
*For more information contact Craig E. Rasmussen at [rasmussen@lanl.gov](mailto:rasmussen@lanl.gov).*

### **Funding Acknowledgements**

NNSA's Advanced Simulation and Computing (ASC).



**Fig. 1.** PTP runtime views showing the node status of a parallel machine (left side) and the process status of a running parallel job (middle). The output from a process, along with a legend showing the different node and process states is also displayed (right side.)



**Fig. 2.** A typical parallel debug session showing the parallel job and associated process status (top left). Directly underneath is a more detailed view of stack frame information from some of the processes being debugged. The program source code and markers indicating breakpoints and current execution locations are displayed (bottom left). Information about the values of variables in one of the processes is also shown (top right.)

## A Software Quality Engineering Case Study: The Jayenne IMC Project Gets a New Mesh Type

Todd Urbatsch, Mike Buksas, Tom Evans, Aimee Hungerford, Scott Mosher,  
Chris Fryer, Jeff Densmore, Tim Kelley, CCS-2

**S**oftware Quality Engineering (SQE) is a support activity for scientific methods and software development. End-user scientists sometimes understandably marginalize SQE, especially when SQE becomes a field of science or philosophy unto itself. However, the importance of software quality cannot be dismissed, especially when there are examples of experts' intuition being irreparably misshapen by bug-ridden software, or when, after decades and 10s of millions of dollars worth of development, a large software program is found to have a bug.

The Jayenne Implicit Monte Carlo (IMC) Project [1] is a collection of software in the Computer, Computational, and Statistical Sciences (CCS) Division that solves the thermal x-ray transport equations using the Fleck and Cummings IMC algorithm [2]. An application of the Jayenne Project software is shown in Fig. 1, where it is coupled to hydrodynamics to model the impact of a hot comet into a granite planet. It is the design of this software that makes testing and many SQE practices possible. The single most important design element is leveled design in which higher-level objects are built upon lower-level objects and there are no cyclic or same-level dependencies. Thus objects can be unit-tested and built upon with some confidence. At the top level of the Jayenne Project software, Milagro is the radiation-only code, and Wedgehog is the interface that can be utilized by application codes. Another design element is representing the independent variables of the mathematical equations as template parameters. For example, a Mesh Type (MT) template parameter represents the spatial variable, and any number of MTs can be built as long as they satisfy the interface requirements. One MT that the Jayenne Project lacked was a 3-D continuous

adaptive mesh refinement (AMR) MT. Here, we describe the testing that went with this new MT.

One tenet of ours is that testing must be repeatable and invocable automatically and on demand. Thus, unit tests are written and stored alongside the actual software. Unit tests verify that the software does what it is required to do. Within each object, or unit, we make use of Design-by-Contract (DBC) assertions that test data coming in, being used, and going out. These DBC statements, which will stop the code if testing fails, can be turned on or off at compile time for either debugging or performance, and they serve to document the requirements of the software. Higher-level unit tests verify ensembles of objects and interfaces. Integral tests, or highest-level unit tests, allow verification against analytic mathematical solutions. "Shunt" tests help verify package interfaces before they are integrated into application codes. Regression tests are simply all these tests monitored over time. The addition of the new 3-D AMR MT in the Jayenne Project software had over 300 unit tests and 30 integral tests at the Milagro level. The impact of the new MT on Wedgehog is shown in the unified modeling language (UML) diagram in Fig. 2, where affected components are shown in red.

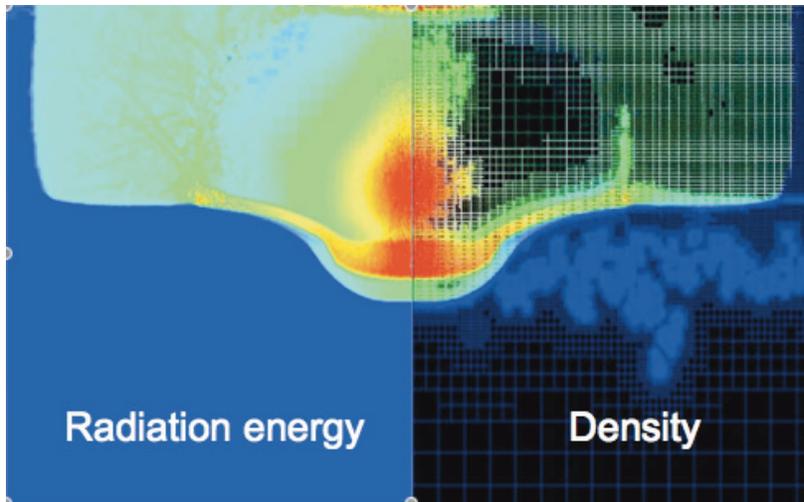
A "Buggy Pageant" is a stunt we perform where we have someone plant bugs in our software and then we find the bugs in front of a live audience. One Buggy Pageant in 2006 pitted us against the most malicious bugs that Russian scientists from our sister lab, VNIIEF, could muster. The average time to find a planted bug is 13.3 minutes. We have since realized that Buggy Pageants would be just as useful as routine team activities whenever we add new capabilities. We performed two Buggy Pageants for the new 3-D AMR MT and

helped improve the testing of the MT and, in the process, educated the entire team in the new software.

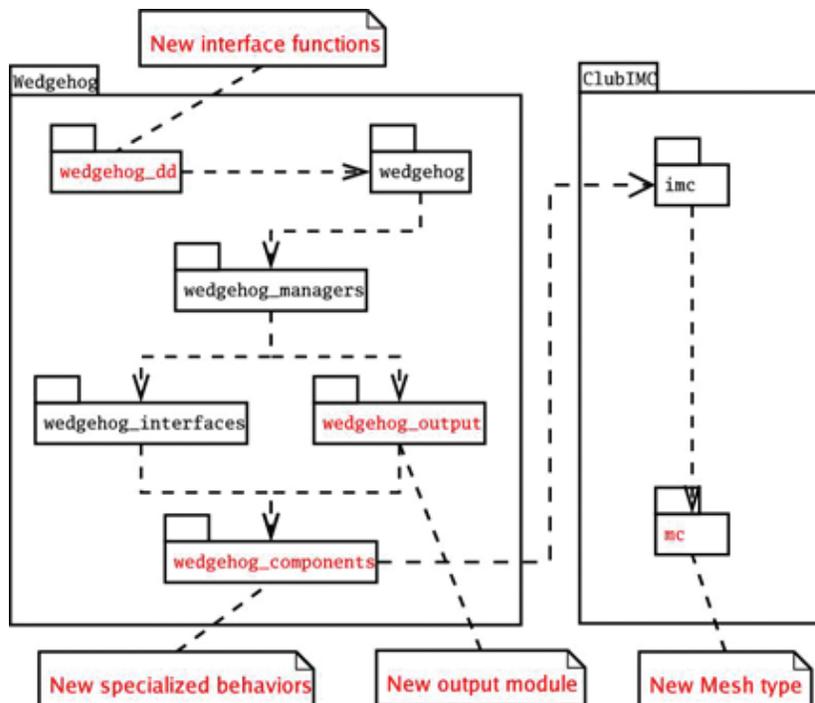
*For more information contact Todd Urbatsch at [tmonster@lanl.gov](mailto:tmonster@lanl.gov).*

[1] T.J. Urbatsch, T.M. Evans, "Milagro Version 2: An Implicit Monte Carlo Code for Thermal Radiative Transfer: Capabilities and Usage," Los Alamos National Laboratory report LA-14195-MS (February 2006).  
 [2] J.A. Fleck and J.D. Cummings, *J. Comp. Phys.* **8**, 313–342 (1971).

**Funding Acknowledgements**  
 NNSA's Advanced Simulation and Computing (ASC), Integrated Codes Program Element, Transport Project.



**Fig. 1.** An application of the Jayenne Project software where it is coupled to hydrodynamics to model the impact of a hot comet into a granite planet.



**Fig. 2.** The impact of the new MT on Wedgehog in the UML diagram, where affected components are shown in red.

## Blind Inversion in Speech and Weapons Image Processing

John Hogden, CCS-3

**T**he speech inversion problem—the problem of inferring the positions of the tongue, jaw, lips, and other *speech articulators* from a recording of a sound pressure wave over time—provides a nice example of a synergy between apparently disparate fields. Not only has speech inversion research benefited from using physical models of wave propagation, but processing techniques developed for speech inversion are starting to be applied to images used for hydrocode validation. We discuss an approach to the inversion problem, emphasizing aspects that generalize to other problems.

Research related to speech inversion includes work by former Los Alamos affiliates such as Kac (founding chairman of the Center for Nonlinear Studies External Advisory Committee) who asked whether it is possible to hear the shape of a drum [1], and Papcun et al. [2] who used a supervised algorithm to learn the mapping between simultaneously collected speech acoustics and articulator positions.

Our work extends the previous work by invoking a very general theorem to invert the mapping from articulator positions to speech acoustics. Since our method does not require simultaneous measurement of acoustics and articulation to learn the relationship, it is a type of blind inversion process. Ignoring caveats that are discussed in depth elsewhere [3], if we have a sufficiently long bandlimited signal,  $\mathbf{x} = [x(1), x(2), \dots, x(T)]$ , that is difficult to

observe (like articulator motions), and an observable signal (like speech acoustics),  $\mathbf{y} = [y(1), y(2), \dots, y(T)]$ , that is a memoryless function of the unobservable signal, i.e.,  $y(t) = f(x(t))$ , then  $\mathbf{y}$  can only have the same bandlimit as  $\mathbf{x}$  if  $f()$  is an affine function, i.e.,  $y(t) = mx(t) + b$ .

The theory stated above tells us that if we transform  $\mathbf{y}$  by any  $g()$  such that  $\hat{x}(t) = h(x(t)) \equiv g(f(x(t))) = g(y(t))$ , and  $\hat{\mathbf{x}} = [\hat{x}(1), \hat{x}(2), \dots, \hat{x}(T)]$  has the same bandlimit as  $\mathbf{x}$ , then  $h()$  will be affine. Thus, if we know the bandlimit on  $\mathbf{x}$ , we can recognize when we have found a  $g()$  that inverts  $f()$  to within an affine transform, and thereby get information about the unobservable signal.

In addition to proving the theorem stated above, we have developed an algorithm for finding  $g()$ . We have successfully tested our function inversion approach using speech data and computer generated, two-dimensional, bandlimited signals [4]. The left side of Fig. 1 shows periodic samples of the computer-generated signals. The right side of Fig. 1 shows the samples transformed by one of the nonlinearities we studied. Using our approach to find  $g()$  we were able to invert the nonlinear transformation. In this case, inverting the function requires reducing the dimensionality of the data, which is an important problem in its own right. We have also successfully applied the technique to synthetic data transformed by an  $f()$  with many-to-one mappings.

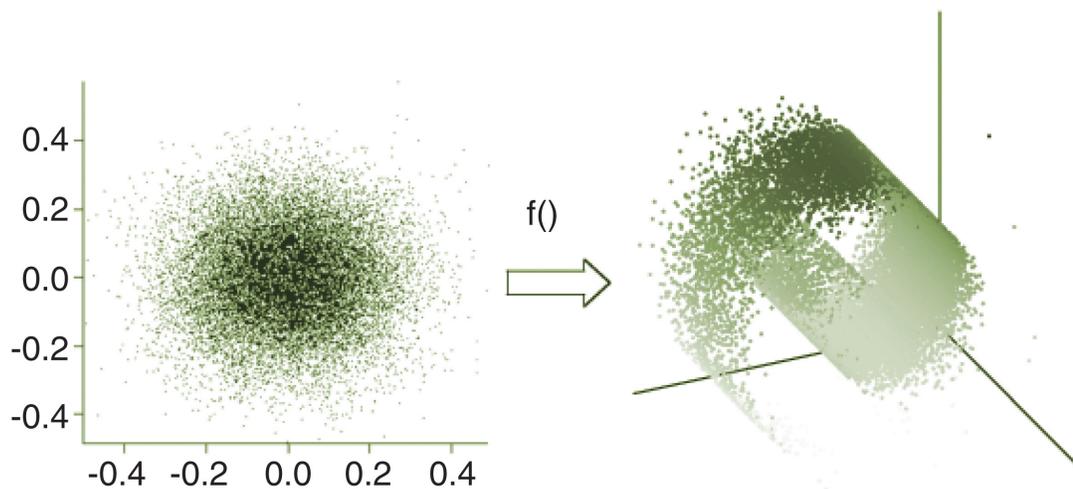
The theory does not require  $x$  and  $y$  to be related to speech. We can also think of  $y(t)$  as being the observable brightness of a pixel in an image and  $x(t)$  as being an unobservable value to be inferred from the image. By using appropriate physical constraints on  $x$ , our theory and closely related ideas have been used to help infer possible mix distributions in both real and simulated PINEX images [5].

*For more information contact John Hogden at hogden@lanl.gov.*

[1] M. Kac, *Am. Math. Monthly* **73** (4), Pt. 2, 1–23 (1966).  
 [2] G. Papcun, et al., *J. Acoustical Soc. Am.* **92** (2), Pt. 1, pp. 688–700 (1992).  
 [3] J. Hogden, et al., “Inverting Mappings from Smooth Paths through  $R^n$  to Paths through  $R^m$ : A Technique Applied to Recovering Articulation from Acoustics,” *Speech Communication, Special Issue on Bridging the Gap Between Human and Automatic Speech Recognition*, Eds. K. Kirchhof & L. ten Bosch (Elsevier, 2007, <http://dx.doi.org/doi:10.1016/j.specom.2007.02.008>).  
 [4] J. Hogden, et al., “Blind Inversion of Multidimensional Functions for Speech Enhancement,” in *Proceedings of the Eurospeech Conference* (2003), pp. 1409–1412.  
 [5] J. Hogden and R. Brewer, “Improvements on Blind Inversion of PINEX Data (U),” JOWOG-32M, Los Alamos, NM, September 26–30, 2005, Los Alamos National Laboratory report LA-CP-05-0990.

#### Funding Acknowledgements

NNSA’s Advanced Simulation and Computing (ASC), Advanced Applications Program.



**Fig. 1.** We did simulation work to see if we could blindly invert the  $f()$  shown above— a Cartesian-to-polar transformation followed by a swiss-roll transformation. We demonstrated that this and other complicated functions can be inverted by our technique.

## Using High-End Collaborative Technologies at Los Alamos National Laboratory

Cindy Sievers, CCS-1

**T**he premier international conference on high-performance computing, networking, and storage, Supercomputing '06 (SC06), convened in November 2006 in Tampa, Florida. As the SCGlobal Showcase '06 General Producer, I worked with SCGlobal to play a major part in SC06 by linking geographically diverse communities from around the world and allowing them to showcase high-end collaborative technologies.

With attendance in excess of 7,000 people, SC06 is an excellent forum for researchers to explore ideas and build collaborations. The latest in high-end collaborative technologies varied in type from simple video conferencing and multipoint information dissemination to highly sophisticated arts productions. Wide and diverse groups of people from business and educational researchers to corporate executives are touched by these technologies enabling those with scarce resources, specialized expertise or equipment, and education resources to collaborate in a timely and cost-effective manner. The SCGlobal presentation had two parts: SCGlobal Showcase and SCDesktop.

### SCGlobal Showcase

Utilizing Access Grid Technologies, the SCGlobal Showcase presentations were made with speakers at both local and remote locations around the world. (See Fig. 1). The Access Grid<sup>®</sup> (AG) is an ensemble of resources including multimedia large-format displays, presentation and interactive environments, and interfaces to Grid

middleware and to visualization environments. These resources are used to support group-to-group interactions across the Grid.

### SCDesktop

From its successful debut last year, SCDesktop returned to offer the supercomputing community the option of participating in SC06 remotely as a "virtual attendee." A virtual attendee accessed sessions, such as the keynote address, plenary sessions, exhibitor forums, and poster sessions, via collaborative technologies that provide two-way audio and video connections to the conference. As part of the registration, virtual attendees received a limited-term license for the collaboration software. Testing and training were provided so that attendees were confident of successful participation.

An added feature of SCDesktop this year was the introduction of Time Delayed Broadcasting. All programs that were offered to virtual attendees were rebroadcast 12 hours later. This time delay allowed European and Asian audiences to enjoy the programs at a better time locally.

Almost 100 sites worldwide participated in SCGlobal, with six continents represented. Remote locations included sites from Canada, United States, United Kingdom, Australia, South America, Mexico, Korea, Alaska, and Hawaii. Site participation ranged from room-based nodes (with more than 10 participants at each room) to individual desktop nodes. We also had up to 40 people in the

SCGlobal room at the SC06 conference center attending the sessions in person. Of special interest were the training sessions, which offered local and remote participants an opportunity to get some hands-on training of new application software developed in the open source AG community.

#### **Access Grid at LANL**

Access Grid technologies are being used at Los Alamos National Laboratory (LANL). We currently have three room-based nodes and several desktop nodes. The main AG node is primarily used for group-to-group technical meetings and seminars. Another AG node, located at the Los Alamos Research Park, is used for the LANL/University of California, Santa Barbara (UCSB) Institute distance learning classes. A third AG node is being built at the LANL Research Library and will host scientific discussions and technical meetings. Numerous examples of how the LANL AG has enabled the Laboratory's programmatic and scientific efforts include 1) collaborations between the LANL Climate, Ocean and Sea Ice Modeling Project within the Climate

Change Prediction Program and its sponsors and others; 2) the partnership between UCSB and LANL to create the Institute for Multiscale Materials Studies; and 3) scientific collaborations with DARPA's High-Productivity Computing Systems program.

*For more information contact Cindy Sievers at [sievers@lanl.gov](mailto:sievers@lanl.gov).*

#### **Funding Acknowledgements**

Los Alamos National Laboratory infrastructure funding.



**Fig. 1.** The SCGlobal room at the Supercomputing '06 conference in Tampa, Florida. The photo shows local presenters as well as remote participants, as seen on the video screen.

# Feature Painting from 2-D Feature Grids and 3-D OSO Models onto 2-D Cartesian and Cylindrical Grids

Martin Staley, T-7

**F**eature painting allows mesh designers to add detail to primary, base meshes by overlaying them with separately meshed shapes, or features. This ability to overlay meshes saves designers from the considerable effort of creating complete meshes directly, beginning again from scratch if a feature’s position, orientation, or size is changed. We have developed two feature-painting libraries, *intergrid* and *interpoint*.

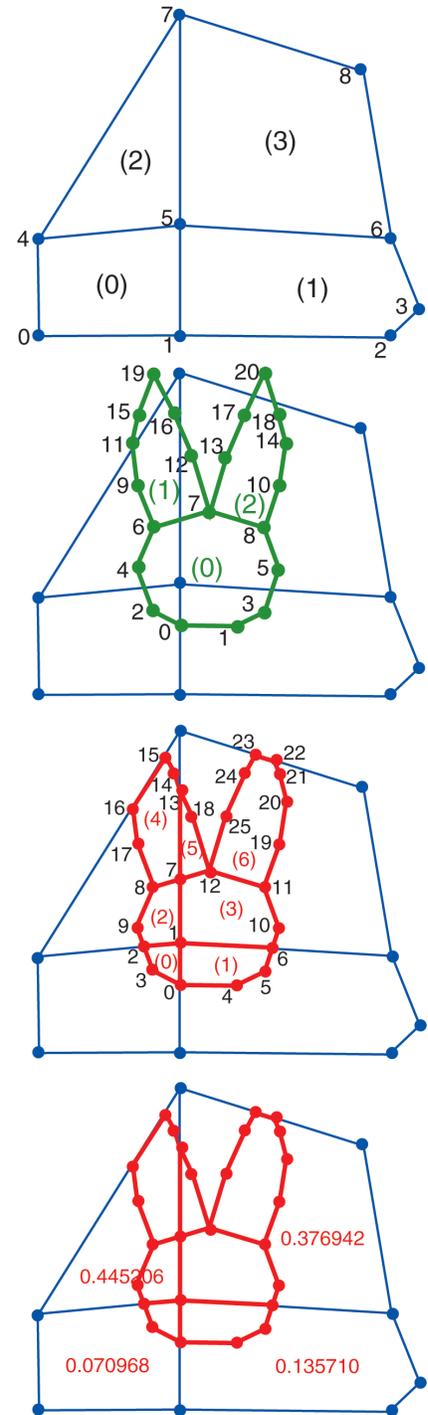
A particular application of feature painting consists of beginning with a base mesh whose cells are of one material, overlaying a feature mesh whose cells are of a different material, and then computing for each base cell the volume fraction of new, feature material. For example, if a particular base cell is four-tenths covered by the feature grid, then that base cell should receive a volume fraction of 0.4.

Precise per-cell volume fractions can be used by codes that allow for multimaterial cells. Other codes can arrange for “stair-stepping:” volume fractions less than 50%, for example, can be clipped to 0, and those 50% or more clipped to 1.

### **intergrid**

To use *intergrid* (“intersect grid”), users supply general, unstructured 2-D grids with convex cells for base and feature. Material information may also be provided. For each base cell of a selected material, *intergrid* computes all of the cell’s intersections with any feature cells, and computes the volume

**Fig. 1.**  
*a) Sample base grid; b) feature grid; c) intersection of base and feature, as computed by intergrid; and d) computed volume fractions.*



fraction of feature material appearing in that cell. The grids may be regarded as Cartesian or cylindrical; *intergrid* uses the appropriate volume weighting in each case.

The trickiest part of this computation is the determination of cell intersections, a seemingly simple action whose robust and efficient implementation in fact requires considerable care.

An important aspect of our implementation was to design an efficient computation of all cell intersections. If the base and feature grids each have  $10^7$  cells, for example, we certainly want to attempt far fewer than  $10^{14}$  intersections.

Figure 1 shows a sample base grid, with nodes enumerated in dark blue and cells parenthesized in black. We then overlay a rabbit feature, enumerated similarly. The third frame shows the intersection grid as computed by *intergrid*, and the fourth frame shows the computed volume fractions, one per base cell.

### **interpoint**

This similarly named extension to *intergrid* accepts an unstructured 2-D base grid as before, now allowing for non-convex cells, plus a feature given in terms of an arbitrary slice through an OSO model. For a selected base material and region of interest in the OSO model, *interpoint* uses point queries (hence the name *interpoint*) to compute volume fractions. We don't use cell intersections, as before, because features now come from constructive solid geometry models, not from grids. Like *intergrid*, this library supports both Cartesian and cylindrical volume weightings.

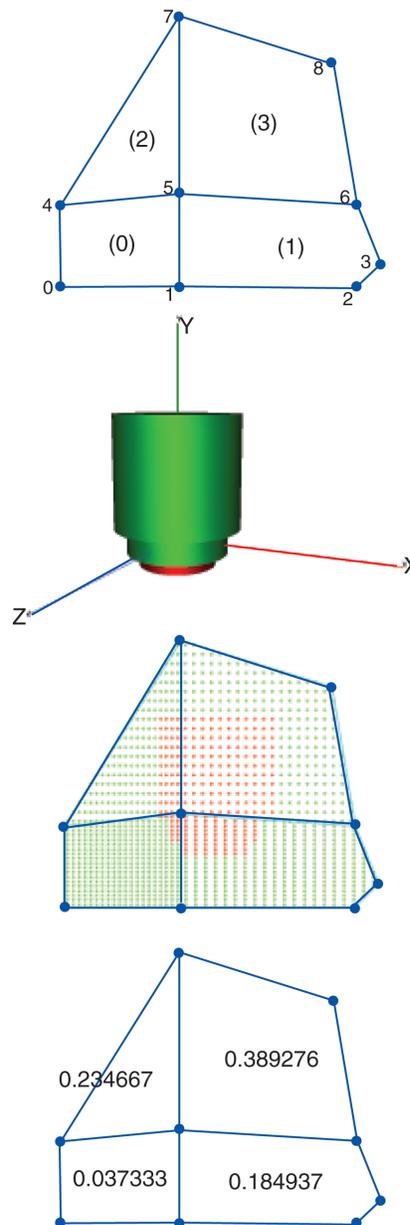
Figure 2 again shows our sample base grid, followed by an OSO feature model. The third frame shows the locations (in 2-D) of point queries, with red indicating values inside the OSO model. Finally, the fourth frame shows the computed volume fractions.

Both libraries are written in ANSI standard C++, and include C-language interfaces. Each library requires the simple inclusion of a C++ header file, without complex makefiles, configure scripts, or library linking. We also have augmented each library with a set of MATLAB<sup>®</sup> codes for visualization, and with example codes and a test suite.

*For more information contact Martin Staley at [staley@lanl.gov](mailto:staley@lanl.gov).*

### **Funding Acknowledgements**

NNSA's Advanced Simulation and Computing (ASC).



**Fig. 2.**  
*a) Sample base grid;  
 b) OSO feature model;  
 c) point-query locations in 2-D; red indicates intersection with OSO feature when transformed to 3-D; and d) computed volume fractions.*

## PixelVizion: An NPU-based Compositor for the Visualization of Large Datasets

Carolyn Connor Davenport, Andrew DuBois, David DuBois, Stephen Poole, HPC-5;  
Laura Monroe, David Pugmire, HPC-4



PixelVizion received a 2006 R&D 100 award.

Advances in high-performance computing have enabled scientists from a variety of disciplines to perform calculations that generate more data than is stored in the entire Library of Congress print collection (17-20 terabytes). A large 3-D problem to visualize a dynamic process may have more than 2 billion cells, with tens of variables per cell, and hundreds of timesteps producing terabytes to petabytes of simulation data to visualize and analyze. Just the sheer volume of data presents a challenge to even the highest-end graphics hardware available today.

Although many advances in speed, memory, and performance have been made in graphics hardware in recent years, the improvement in such hardware is driven by market forces, which are not always aligned with high-performance scientific computing and visualization needs. Much research has been carried out to leverage these advances so that they may be used in visualizing large datasets.

Visualization clusters composed of commodity nodes are often used to perform distributed rendering of large datasets. This type of large-scale data visualization usually involves workload distribution across a cluster of computational nodes, since the datasets are too large to fit on a

single graphics processor. This rendering distribution necessitates a final compositing step to blend the rendered images into one image.

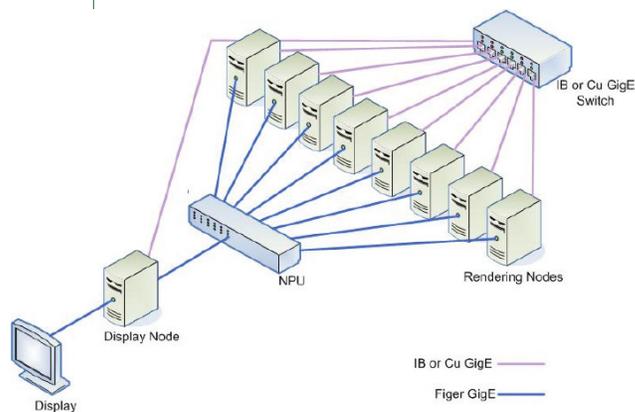
This compositing step is a well-known bottleneck that places an upper bound on performance. Thus, although distributed commercial off-the-shelf (COTS) systems do offer increased graphics horsepower and a better cost/benefit ratio, the compositing step must be improved in order to truly realize the benefits of this type of rendering system.

We have recently developed a novel use of Network Processing Units (NPUs) to perform the compositing step on a COTS cluster [1]. The NPU uses a heterogeneous multicore compute architecture that efficiently maps onto the streaming nature of image composition [2]. We streamed images generated on current graphics cards into the NPU via a simple application-programming interface (API) and allowed the NPU to composite the images.

The NPU-enhanced visualization system makes use of a completely COTS NPU rather than the interconnect to perform high-speed hardware-assisted image composition. The NPU augments a visualization system composed of a traditional COTS rendering cluster equipped with COTS graphics cards and an InfiniBand 4X network connection, with GigE connecting the NPU compositing system to the main cluster (Fig. 1).

We attained compositing speeds beyond that of the comparison software-only composition even when the software-only method had the advantage of a faster high-performance network. Unlike traditional software-based methods that do not scale with cluster size [3], the NPU compositor is able to maintain a nearly constant frame rate

**Fig. 1.**  
An NPU-enhanced visualization cluster.



for arbitrarily large clusters. The theoretical frame rate for such an NPU system built using a GigE network is 28.68 frames per second, and this frame rate is attained by our system.

At the 2005 International Conference on High-Performance Computing, Networking, Storage, and Analysis, we demonstrated 1024x1024 image composition speeds nearly 4X faster than the standard fast binary-swap software-based composition (Fig. 2). This was done using just the compositor with prerendered images.

We also incorporated NPU-based composition into VisIt, the Lawrence Livermore National Laboratory (LLNL) visualization tool [4], to demonstrate the integration of this technology to an existing, full-featured visualization system (Fig. 3), and applied it to real Los Alamos National Laboratory simulation datasets (Fig. 4) [5]. Integrated NPU-accelerated composition achieved a 2X increase over the performance of VisIt using the software binary-swap composition that was implemented.

The difference in speedup between the movieplayer case and the VisIt full-package case above (4X vs 2X) is due to (1) the nonoptimized application pipeline, and (2) the inherent system overhead penalty associated with the use of any full-featured visualization package.

We have since integrated stereo into the VisIt/NPU system. We have also achieved a preliminary integration of another scientific visualization software package, ParaView [6] into the NPU compositor system.

- [1] D. Pugmire, et al., "NPU-Based Image Compositing in a Distributed Visualization System," to appear in IEEE Transactions on Visualization and Computer Graphics 2007.
- [2] B. Bouzas, et al., "MultiCore Framework: An API for Programming Heterogeneous Multicore Processors," Information Sciences Institute, University of Southern California (March 2006).
- [3] A. D. Brydon, et al., "Simulation of

the Densification of Real Open-celled Foam Microstructure," *J. Mech. Phys. Solid* **53**, 12, 2638 (Dec. 2005).

- [4] M. Ogata, et al., "The Design and Evaluation of Pipelined Image Compositing Device for Massively Parallel Volume Rendering," Proceedings of the 2003 Eurographics/IEEE TVCG Workshop on Volume Graphics, 61 (2003).
- [5] VisIt home page, <http://www.llnl.gov/visit/>
- [6] ParaView home page, <http://www.ParaView.org/>

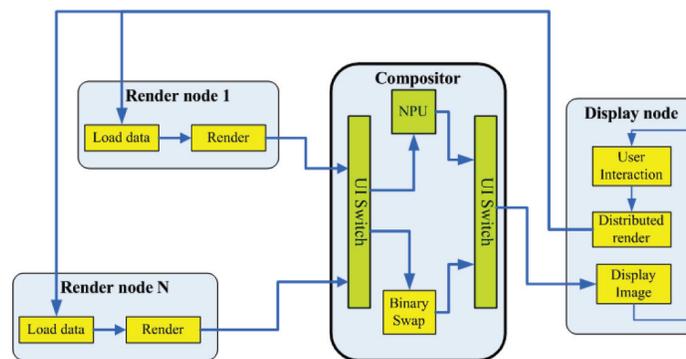
*For more information contact Carolyn Connor Davenport at [cmcd@lanl.gov](mailto:cmcd@lanl.gov).*

### Funding Acknowledgements

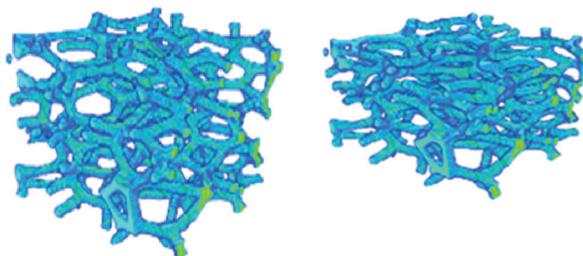
This project was supported through the LANL ASC Simulation and Computer Science Production Visualization project and the LANL ASC Weapons-Supported Research project.

Image	Nodes	NPU fps	BSC fps	Speedup
512 <sup>2</sup>	8	89	26	3.4X
1024 <sup>2</sup>	2	28		
1024 <sup>2</sup>	3	28		
1024 <sup>2</sup>	4	28		
1024 <sup>2</sup>	5	28		
1024 <sup>2</sup>	6	28		
1024 <sup>2</sup>	7	25		
1024 <sup>2</sup>	8	22	6	3.7X
2048 <sup>2</sup>	8	5.5	1.5	3.7X

**Fig. 2.** Comparison of binary swap compositing (BSC) to single NPU compositing on a variety of render node configurations.



**Fig. 3.** The VisIt distributed rendering pipeline. The green blocks indicate modules added to the VisIt pipeline to incorporate NPU-enhanced rendering: specifically, UI switches select between the binary swap or NPU compositor.



**Fig. 4.** Low-density polymeric foam calculation. On the left is the initial configuration; on the right is a deformed intermediate state.

## Pick-n-Place: A Virtual Reality Assembly Tool

David Pugmire, Laura Monroe, David Modl, HPC-4

**V**isualization has long been an important tool for the understanding of complex data. The creation of a 3-D visual representation enables the highly developed human visual system to detect, among other things, trends, correlations, anomalies, and unexpected events in the data. Over the past several years, Los Alamos National Laboratory has invested in providing state-of-the-art facilities and software tools to help scientists in the Advanced Simulation and Computing (ASC) program understand and explore their data. One such facility, the LANL CAVE, driven by a Linux cluster and equipped with an optical motion tracking system, is shown in Fig. 1.

An unexplored application of these capabilities has been the design, manufacture, and assembly of complex mechanical models. These models are developed and studied by a large

number of people, often in different locations. Allowing these parties to interact with an unbuilt, virtual model has the potential to tremendously transform the way the Department of Energy (DOE) complex operates. Further, performing virtual analysis on complex models before they are built, where conflicts, alternative approaches, and optimizations can be identified early on, can yield enormous savings in both cost and time. Our recent work with the Reliable Replacement Warhead (RRW) project provided an opportunity to develop and test these techniques.

As shown in Fig. 2, Pick-n-Place is a virtual reality assembly and disassembly application developed by David Pugmire using CAVELib [1], an industry-standard virtual reality platform. The application is able to read and display CAD models as well as texture-mapped polygonal models to provide realistic environments, such as

**Fig. 1.** The CAVE is a five-surface (left, front, right walls with floor and ceiling) display with 33 projectors lighting 43 million pixels in a 15 x 12 x 10-ft display area, and researchers can work in a 3-D environment.



an assembly bay. Assembly constraints can also be given to specify how parts are able to be mated. Collision detection [3] is used to determine interpenetration of parts, and the stereo sound system is used to play a thud sound when parts collide. The application communicates with the motion tracking system through the Trackd [2] interface. Entering the CAVE, the user can visually inspect the components of the assembly, grab them, and try to assemble them. If parts do not fit together or are obstructing one another, both visual and auditory cues are used.

This tool proved valuable while exploring various design alternatives with personnel at DOE plants and LANL designers to integrate safety studies with plans for their assembly and disassembly procedures.

[1] CAVELib, [www.vrco.com/CAVELib/OverviewCAVELib.html](http://www.vrco.com/CAVELib/OverviewCAVELib.html)

[2] Trackd, <http://www.vrco.com/trackd/Overviewtrackd.html>

[3] T. Hudson, et al., "V-COLLIDE: Accelerated Collision Detection for VRML," Proceedings of VRML 97 (1997).

*For more information contact David Pugmire at [pugmire@lanl.gov](mailto:pugmire@lanl.gov).*

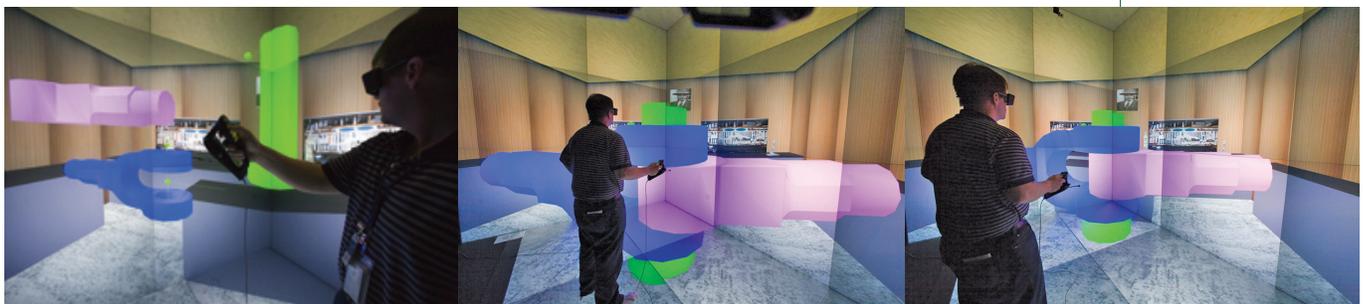
#### **Funding Acknowledgements**

Supported through the LANL ASC Simulation and Computer Science Production Visualization project.

**Fig. 2.**

*A sequence of photos in the CAVE demonstrates how the Pick-n-Place virtual reality tool can be used to assemble and analyze mechanical components.*

*Photos by L. Sanchez*



## Linux Clusters Now Drive Metropolis Center Visualization

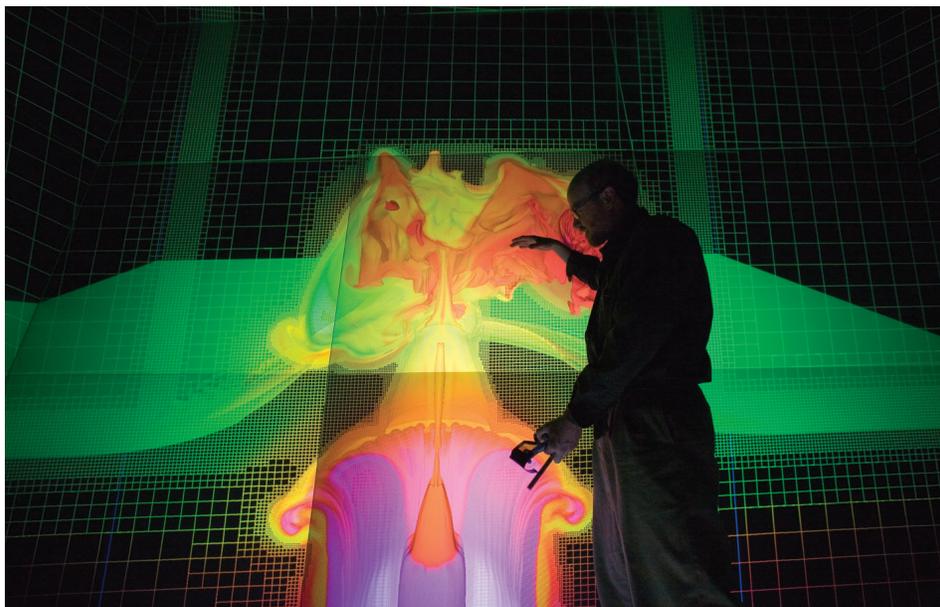
David G. Modl, HPC-4; Andrew G. Shewmaker, HPC-5

**L**os Alamos National Laboratory's Nicholas C. Metropolis Center for Modeling and Simulation houses some of the world's most advanced supercomputing systems and hosts several world-class data visualization display facilities. These facilities include an 85-seat PowerWall theater with 24 Christie Digital Mirage Model 2000 projectors producing 31 million pixels spread across a 22 × 12-ft display screen. There is also a five-surface (left, front, right walls with floor and ceiling) display called La Cueva Grande (the big cave), with 33 Mirage projectors lighting 43 million pixels in a 15 × 12 × 10-ft display area. These facilities and direct KVM (keyboard-video-mouse) connections to the scientists' offices easily enable researchers to compute and analyze their data and share results with colleagues and sponsors. Until recently, these facilities were driven by large SGI (Silicon Graphics, Inc.)

supercomputer systems designed and installed in the late 1990s as part of the DOE Accelerated Strategic Computing Initiative (ASCI) 1.608 teraFLOPS Blue Mountain system.

We have replaced three SGI Origin 2000 supercomputers and one (newer) SGI Onyx 4 3900 with a 264-node HP (Hewlett Packard) Linux cluster system we call ViewMaster. This system is based on the HP xw8200 workstation with the high-end Nvidia Quadro FX 4500 graphics cards that are more than 100 times faster than the SGI system's graphics hardware. We have dedicated 25 nodes in this cluster to drive the PowerWall theater and another 34 nodes are assigned to drive La Cueva Grande. This new technology allows us to visualize much larger data sets much more rapidly. We have waited a long time for cluster technology to be available that would allow us to assemble this system.

**Fig. 1.**  
*HPC-4 staff in an im-  
mersive CAVE  
visualization.*



State-of-the-art computer graphics visualization and analysis requires enormous computing capability, including large amounts of memory per CPU, a high-speed network interconnect between the nodes in the cluster, and a high-end graphics card that is capable of 3-D stereo imaging with Frame Lock technology. Stereo imaging is a technology that allows us to display images in so-called “stereo pairs” (one image is rendered for the left eye and another image is rendered for the right eye of the observer). We use StereoGraphics CrystalEyes 3-D glasses and a synchronization signal from the graphics system to synchronize the glasses with the display screens.

Frame Lock is a technology that allows us to configure computer graphics hardware where images on different systems are synchronized so that each rendered frame spread across the display surface is locked and the images update together. Without this technology, 3-D stereo multipanel displays would be “out of sync” with the other panels, thus the whole displayed image would seem “jumpy.” The SGI equipment from the 1990s provided us with Frame Lock 3-D stereo capability. Cluster technology, just recently available, allows us to update

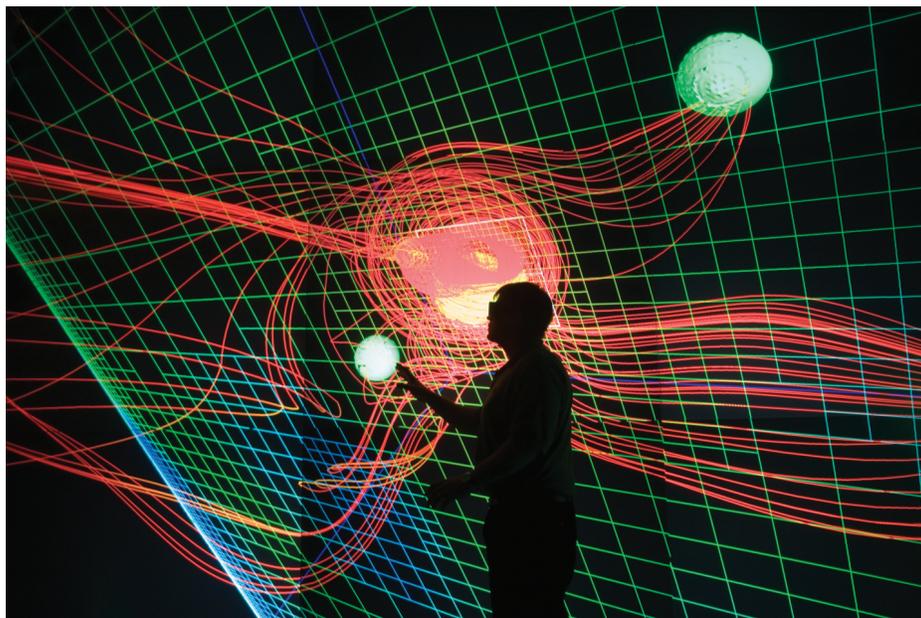
the systems we use for 3-D stereo multipanel display visualization work.

In December 2006 we completed the transition from the SGI systems to the HP/Nvidia-based Linux cluster.

*For more information please contact David Modl at [digem@lanl.gov](mailto:digem@lanl.gov).*

#### **Funding Acknowledgements**

This research was supported by the NNSA tri-Lab Advanced Simulation and Computing Program.



**Fig. 2.**  
*A demonstration simulation in the CAVE.*

*Photos by L. Sanchez*

## Dynamic Kernel Instrumentation on Clusters

Nathan DeBardleben, William Rust, HPC-4

**H**igh-performance computing clusters are increasingly complex systems that bring together disparate subsystems of CPUs, memory, I/O, networks, and file systems. All of these subsystems can affect the performance and stability of the others, leading to a complex balance that must be maintained. As we move to new processors such as the IBM Cell Broadband Engine™ (Cell BE) (Fig. 1) in the Los Alamos National Laboratory Roadrunner supercomputer, we can expect these effects to get worse, not better. The emphasis on correctly tuning applications, system software, and operating system kernels will grow to be even more important for overall system performance.

We have been studying the technology of dynamic kernel instrumentation[2, 3] using SystemTap [1] on a small testbed cluster designed to mimic Lightning, a large machine used at LANL for the Advanced Simulation and Computing program. We have already used this technology to instrument the virtual file system layer to look at potential parallel file system problems. Additionally we have written probes to trace kernel calls, monitor specific files, and examine network activity and assorted other hot spots in the kernel.

With the increasing emphasis on multicore systems and, in particular, the Cell BE processor for Roadrunner, we extended this dynamic probing technology to the Cell BE kernel. The kernel modifications for the Cell BE are

vast and rapidly changing. As with any new technology put into the kernel, there are problems that take time to work out. Furthermore, the complexity of the Cell processor is revealing many nuances about how to correctly program applications for it when one looks at how they affect the kernel. By probing a running Cell kernel, we are able to debug and study the complex interaction between the main processor and the Synergistic Processing Units. With this technique we can look at how a system is behaving without the need to build a custom kernel with debugging features always enabled. This ability to attach to a kernel and then detach at a later point is important in production machines because we do not want to be running with debug kernels waiting for a problem to arise. Furthermore, with the complexities of these systems growing, it is often impossible to predict what a system analyst will need to examine. Therefore, having the ability to apply a dynamic kernel modification is important for looking at a problem as it arises.

### Conclusion

As we experiment with this developing kernel, we are finding many interesting places to glean detailed information about a running system. We believe carefully placed kernel probes can be used to allow the kernel to self-diagnose problems and report them to a system analyst. It may even be possible to build an expert system that attempts to tune itself when it recognizes problems.

For more information contact Nathan DeBardeleben at [ndebar@lanl.gov](mailto:ndebar@lanl.gov).

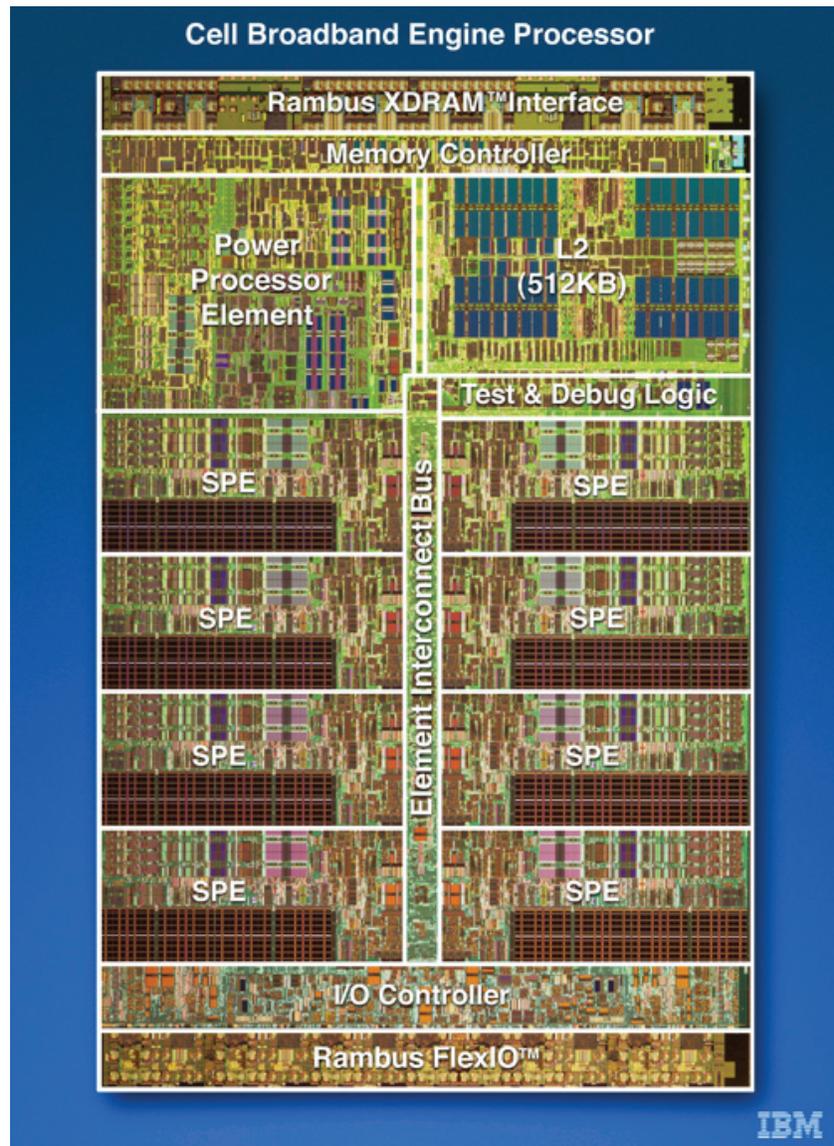
[1] Redhat Linux, *Systemtap Open-Source Project*, <http://sources.redhat.com/systemtap/>

[2] V. Prasad, et al., "Locating System Problems Using Dynamic Instrumentation," *Proceedings of the Linux Symposium 2*, 49 (July 2005).

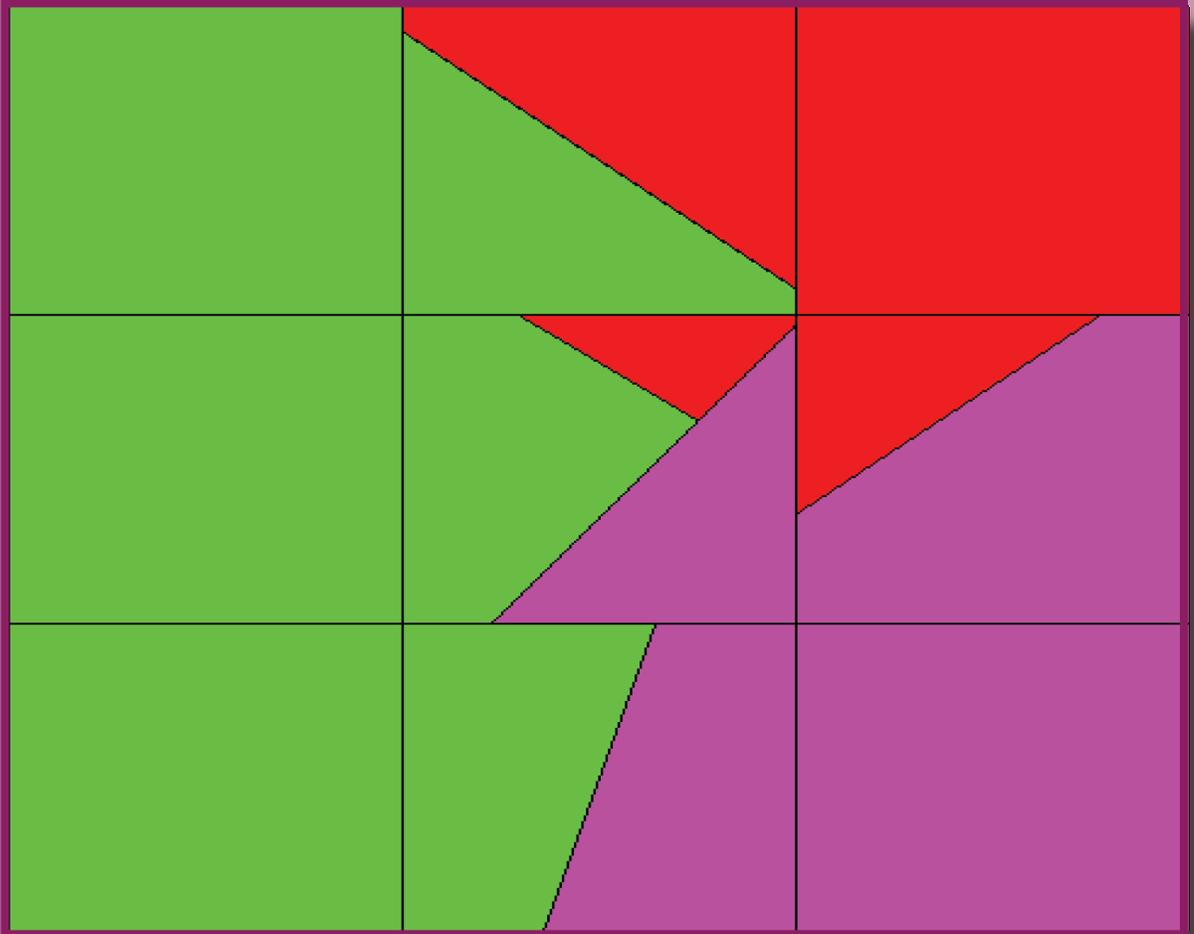
[3] A. Tamches and B. P. Miller, "Using Dynamic Kernel Instrumentation for Kernel and Application Tuning," *Intl. J. High Perf. Comput. Appl.* **13**, 263 (1999).

### Funding Acknowledgements

This research was supported by the NNSA tri-Lab Advanced Simulation and Computing Program.



**Fig. 1.** Developing dynamic probing technology for the complex Cell Broadband Engine™ processor is a goal of this research. Image courtesy of IBM.



# Hydrodynamics Methods and Algorithms

The dynamics of fluid flow is one of the fundamental physical processes in the evolution of the explosion of a nuclear weapon. The articles in this section speak to the current state-of-the-art in the simulation and understanding of some of the underlying physics in hydrodynamics—from the effect of approximations and numerical instability, to the fundamental evolution of turbulence, to the simulation of shock dynamics.

The six articles in this section summarize current developments in:

- the effect of pressure-temperature equilibrium on an important multimaterial fluid instability, the evolution;
- the evolution of turbulence in the interface between fluids of different densities using numerical simulation and experiment;
- the structure of radiative shock waves and the efficacy of their numerical prediction;
- the evolution of Rayleigh-Taylor instabilities and numerical comparisons with theory;
- a new method for the reconstruction of multimaterial interfaces; and
- new atomistic methods for simulating Rayleigh-Taylor instabilities.

These research developments are fundamental to improving the nature of predictive science in the nuclear weapons program.

## Effects of Pressure-Temperature Equilibrium on Richtmyer-Meshkov Instability

John W. Grove, Thomas Masser, CCS-2

The most common numerical method for treating multiple material component mixtures is to assume that the mixture consists of separated components that are in pressure-temperature (PT) equilibrium. For inviscid flows, the resulting partial differential equation model is the Euler equations expressing conservation of component mass, total momentum, and total energy. Such a model is relatively easy to implement in a standard compressible hydrodynamics code, but its appropriateness as a flow model is highly problem-dependent since it assumes that the time scales of interest are long relative to the thermal conduction rates of the material components. For high-speed flows with shocks, the component interactions are too fast for the materials to come into temperature equilibrium, and the use of such a model can lead to substantial errors in the computation.

We studied effects of PT equilibrium on flow hydrodynamics for Richtmyer-Meshkov (shock-driven) instability by comparing two hydrodynamic codes, the Advanced Simulation and Computing (ASC) code RAGE/SAGE, which uses PT equilibrium for mixed cells, and the University at Stony Brook's hydro-code FronTier, which maintains pure material cells separated by tracked material interfaces. FronTier also supports interfaces with shear, while RAGE/SAGE assumes a single velocity amongst mixture components.

We found that PT equilibrium causes several pernicious effects for a flow calculation. Figure 1 shows spurious initialization transients generated at temperature discontinuities. The flow is

simple periodic translation of a slab of hot dense material at constant pressure. The exact solution is computed by FronTier, but even at early time nearly 10% oscillations in the density field are produced near interfaces when PT equilibrium is imposed. At late time these transients in fact dominate the flow. The graph compares three versions of RAGE/SAGE calculations, one with the default setting for the hydrodynamics, one with interface preserving artificial compression, and the other with volume of fluid interface reconstruction (VOF). The behavior of all three schemes is similar.

A second main effect is over-cooling of compressible regions. Figure 2 shows a late-time comparison for a cylindrical implosion of a tin layer into an air cavity. Two major differences are immediately apparent. The multiple shock refractions driving the instability have produced pockets of superheated air that are over an order of magnitude hotter when computed by explicit interfaces in FronTier as compared to the highly mixed PT equilibrium cells in RAGE. The vorticity field shows an even more marked difference, with XSAGE (a cousin of SAGE) yielding a much more diffuse vorticity field than the corresponding FronTier simulation.

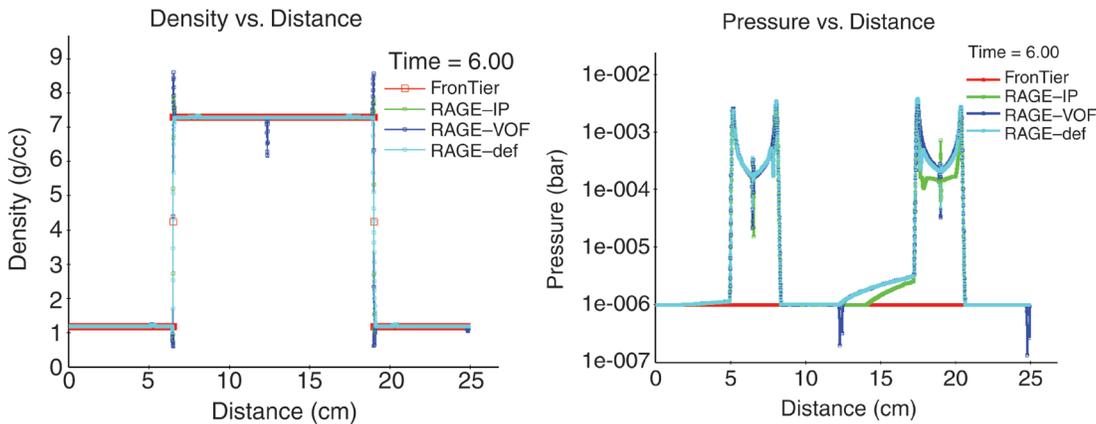
*For more information contact Thomas Masser at [tmasser@lanl.gov](mailto:tmasser@lanl.gov).*

[1] R.D. Richtmyer, *Comm. Pure Appl. Math.* **13**, 297 (1960).

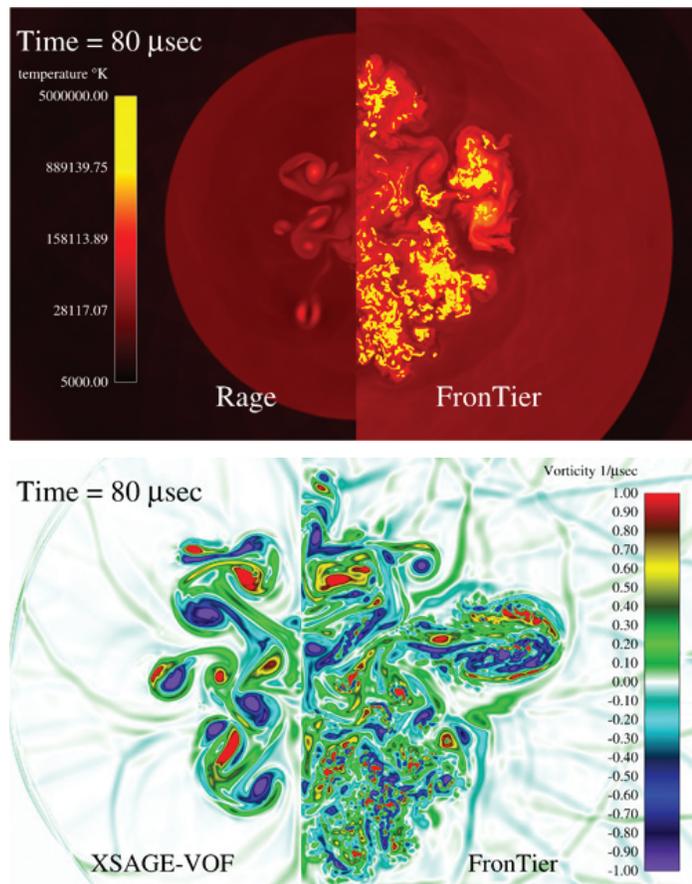
[2] E.E. Meshkov, *NASA Tech. Trans.* **F-13**, 074 (1970).

[3] Michael Gittings, et al., "The RAGE Radiation-Hydrodynamic Code," Los Alamos National Laboratory report LA-UR-06-0027 (2006).  
 [4] J.W. Grove, *J. Appl. Num. Math.* **14**, 213 (1994).

**Funding Acknowledgements**  
 NNSA's Advanced Simulation and Computing (ASC), Computational Physics and Mathematics Program.



**Fig. 1.** Pressure transients created by the imposition of temperature equilibrium at a density-temperature interface.



**Fig. 2.** A late-time comparison of cylindrical Richtmyer-Meshkov instability. FronTier produces substantially hotter pockets of air with a more complex vorticity field than the RAGE/XSAGE calculation.

# Numerical Simulations of Rayleigh-Taylor Turbulence with a Complex Acceleration History

Praveen Ramaprabhu, CCS-2; Guy Dimonte, X-1; Malcolm J. Andrews, CCS-2

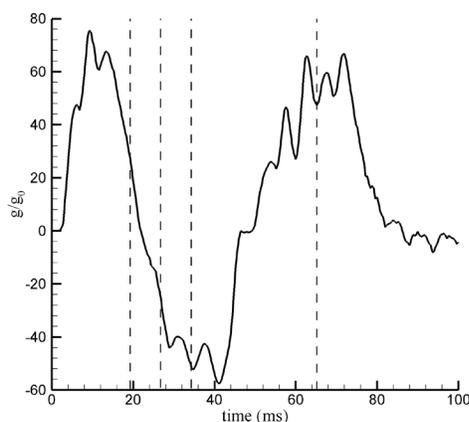
The interface separating two fluids of different densities is unstable if an acceleration is applied from the light fluid to the heavy. The resulting instability is the Rayleigh-Taylor (RT) instability and is a dominant phenomenon in the implosion phase of Inertial Confinement Fusion, where the subsequent turbulent mixing dilutes fuel with pusher material, thus reducing yield. Rayleigh-Taylor-driven turbulence is also an important effect in explaining supernova detonations, and other high-energy density applications. Reliable numerical simulations (NS) of the turbulent phase of RT are thus important to the Validation & Verification (V&V) efforts of Advanced Simulation and Computing (ASC) codes and to the stockpile stewardship program in general.

Hydrodynamic simulations of such high energy density phenomena deal with multiple physics and as a result represent the turbulence using low-order mix models. In this article, we describe NS of RT that will be useful when validating mix models in the demanding setting of a complex acceleration history.

The simulations employ an incompressible MILES [1] code and reproduce experimental runs on the Linear Electric Motor (LEM) [2,3], which used a three-stage acceleration profile—an initial acceleration stage, followed by a sudden deceleration, and a final re-acceleration stage (Fig. 1). It is expected that not all models will completely describe the separation of phases during deceleration. This is a challenging problem for NS because during deceleration, large bubbles reverse direction and are shredded by smaller bubbles in their way, thus generating small-scales that can only be resolved at large-grid resolutions. Furthermore, the calculations of this incompressible flow would have proved challenging to compressible codes due to their propensity to generate pressure waves during sudden changes in acceleration.

Figure 2 (a–d) shows images of the turbulent density field from experiments and simulations, realized at times indicated by the vertical lines in Fig. 1. A broadband spectrum of density perturbations was used to initialize the simulations, since the experiments are susceptible to ambient vibrations initially, which are expected to have a similar spectral structure. A second set of calculations with energy confined to a narrow band of wavelengths was performed, giving inferior agreement with the experimental results compared with the broadband cases [3]. The subsequent evolution of the bubble penetration depth is shown in Fig. 3 and was found to be sensitive to the initial spectral structure and its r.m.s. amplitude. The calculations accurately describe the initial exponential growth

**Fig. 1.** Nondimensional acceleration history ( $g/g_0$ ) used in LEM experiments and NS.  $g_0$  is the earth's gravity.

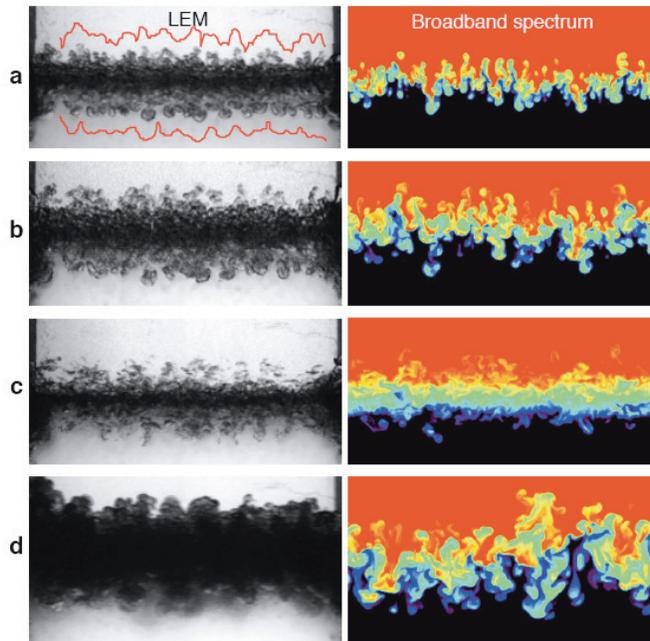


to nonlinearity, the collapse of coherent structures during the deceleration phase, the subsequent recovery during re-acceleration, and the late-time evolution in to self-similarity. It is noteworthy that our high-resolution NS accurately capture the shredding of bubbles into small scales during deceleration, and the associated molecular mixing [3]. We conclude that initial conditions play a critical role in describing RT-driven mixing especially under nonequilibrium conditions.

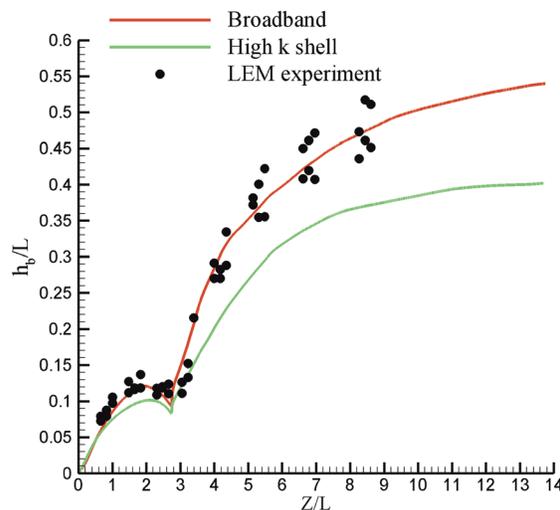
*For more information contact Praveen K. Ramaprabhu at praveen@lanl.gov.*

[1] Guy Dimonte, et al., *Phys. Fluids* **16**, 1668 (2004).  
 [2] Guy Dimonte and M. Schneider, *Phys. Fluids* **12** (2), 304 (2000).  
 [3] Guy Dimonte, et al., "Rayleigh-Taylor Instability with Complex Acceleration History" (submitted to *Phys. Rev. Lett.*).

**Funding Acknowledgements**  
 NNSA's Advanced Simulation and Computing (ASC), Physics and Engineering Models Program.



**Fig. 2.** Turbulent density field images from experiments (left) and NS (right) at (a)  $t=19.6$  ms, (b)  $t=26.8$  ms, (c)  $t=34.4$  ms, and (d)  $t=65.2$  ms.



**Fig. 3.** Bubble penetration depths  $h_b$ -scaled to the cell size  $L$  used in LEM runs, from NS and experiments vs  $Z/L$  ( $Z = \int dt' \int dt'' g(t'')$  cm).

## The Physics of Radiative Shock Waves

Robert B. Lowrie, CCS-2; Rick M. Rauenzahn, T-3; Jarrod Edwards, Texas A&M University

**T**his research investigates the structure of radiative shock waves and the ability of our simulation codes to accurately predict such waves. A radiative shock is a wave in which both hydrodynamic and radiation transport physics play a significant role in the shock's propagation and structure. The addition of thermal radiation broadens the temperature profile of a hydrodynamic shock, in much the same manner as thermal conduction. At very high temperatures, when the material and radiation energies are comparable in magnitude, radiation plays an additional role in determining the shock jump relations.

One goal of this research is to generate semianalytic radiative shock solutions with sufficient accuracy that they may be used to quantify errors in simulation codes. There are very few existing test problems with such nonlinear, coupled physics. The radiative shock problems will also help answer whether temporally unsplit numerical algorithms are significantly more accurate than operator-split algorithms and to help develop better mesh refinement criteria.

The detail that the semianalytic solution provides has also led to the discovery of temperature profiles that have a more intricate structure than has previously been documented [3].

Thus far, we have developed solutions for the Euler equations of gas dynamics, coupled with two different gray radiation models:

- **Equilibrium Diffusion:** The radiation and material temperatures are assumed to be the same. For this simple model, the radiation adds nonlinear thermal conduction and effectively modifies the equation of state.
- **Nonequilibrium Diffusion:** The radiation and material temperatures may be different, and the radiation energy is assumed to propagate via a Fick's law. When the opacity is large, this model reduces to the equilibrium diffusion model.

Figure 1 compares our solutions [3] with results from the RAGE code [1]. RAGE solves a discretization of the nonequilibrium diffusion model. The RAGE results compare well with the semianalytic solution. Such results give us confidence that RAGE can accurately compute highly nonlinear, coupled radiation-hydrodynamics.

Figure 2 shows the errors in the shock location as a function of mesh resolution, for two different simulation codes. For this problem, the opacity is large enough that equilibrium diffusion holds, regardless of the radiation model. Each code uses a different radiation model, yet in this large-opacity limit, both predict the equilibrium diffusion answer with reasonable accuracy. The results also show that there is some benefit in using a second-order time integration scheme. See Ref. [2] for more discussion.

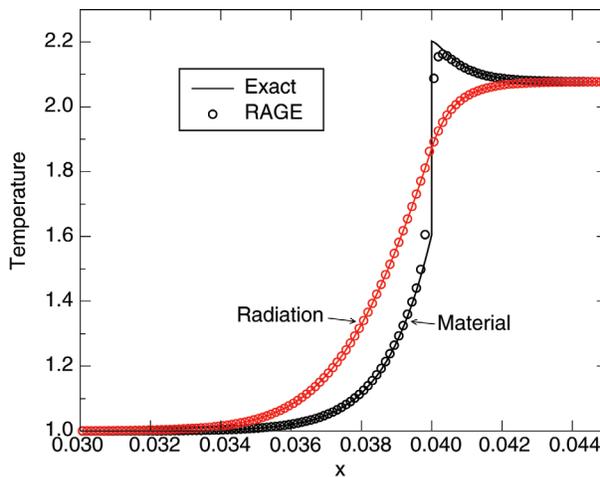
Future work will include additional physics, more advanced radiation models, and using these solutions to verify simulation codes and improve their algorithms.

*For more information contact Robert Lowrie at [lowrie@lanl.gov](mailto:lowrie@lanl.gov).*

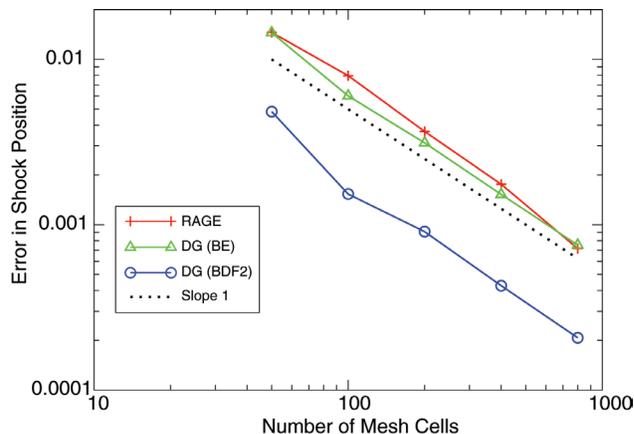
- [1] M. Gittings, et al., "The RAGE Radiation-Hydrodynamics Code," Los Alamos National Laboratory report LA-UR-06-0027 (2006).
- [2] R.B. Lowrie and R.M. Rauenzahn, "Radiative Shock Solutions in the Equilibrium Diffusion Limit," Los Alamos National Laboratory report LA-UR-06-8283 (2006); to appear in *Shock Waves*.
- [3] R.B. Lowrie and J.D. Edwards, "Shock Wave Solutions for Radiation Hydrodynamics," Proceedings of the American Nuclear Society M&C+SNA Meeting, Monterey, 2007.

### Funding Acknowledgements

NNSA's Advanced Simulation and Computing (ASC), Computational Physics and Mathematics Program.



**Fig. 1.** Shock temperature profile for a Mach 2 shock, nonequilibrium diffusion. The shock is propagating right to left. RAGE values are shown at each mesh cell-center. Note that there is an embedded hydrodynamic shock at  $x = 0.04$ .



**Fig. 2.** Errors in computing the shock position as a function of mesh size, Mach 10 shock, in the equilibrium diffusion limit. RAGE and DG(BE) use a first-order time integration method, while DG(BDF2) uses a second-order method. All the methods are second-order in space.

## Calculations of the Single-Mode Rayleigh-Taylor Instability: RAGE and FLAG Hydrocode Comparisons

Marianne M. Francois, Robert B. Lowrie, CCS-2; Jimmy Fung, Mark Kenamond, Edward Dendy, X-3

**T**he Rayleigh-Taylor instability (RTI) is a fluid instability that is often encountered in astrophysics, inertial confinement fusion, and in mixing applications. This instability occurs whenever a light fluid is accelerated into a heavy fluid. This fluid configuration is barotropically unstable; that is, perturbations along the interface tend to grow as pressure and density gradients interact to produce vorticity. The instability is named after Lord J. Rayleigh (1883) [1] and G.I. Taylor (1950) [2].

Because of the RTI importance in many applications and because it has been widely studied, the RTI serves as a good test case for comparing the RAGE and FLAG hydrocodes. The goals of this comparative study are to assess and document the codes' ability and identify areas for code improvements. Although the two codes are designed under different philosophies, their designed intended applications overlap enough to allow for designing cross-code comparisons.

Both the RAGE and FLAG codes solve the compressible Euler equations for multi-materials. However, the codes are different in their approach to solve the equations. RAGE is an Eulerian code with adaptive mesh refinement and employs a high-resolution Godunov scheme [3]. FLAG is an arbitrary Lagrangian-Eulerian (ALE) code with an Eulerian mesh remap capability [4].

Several computations of the single-mode RTI of two ideal gases with density ratio 2:1 have been performed for a variety of numerical parameters.

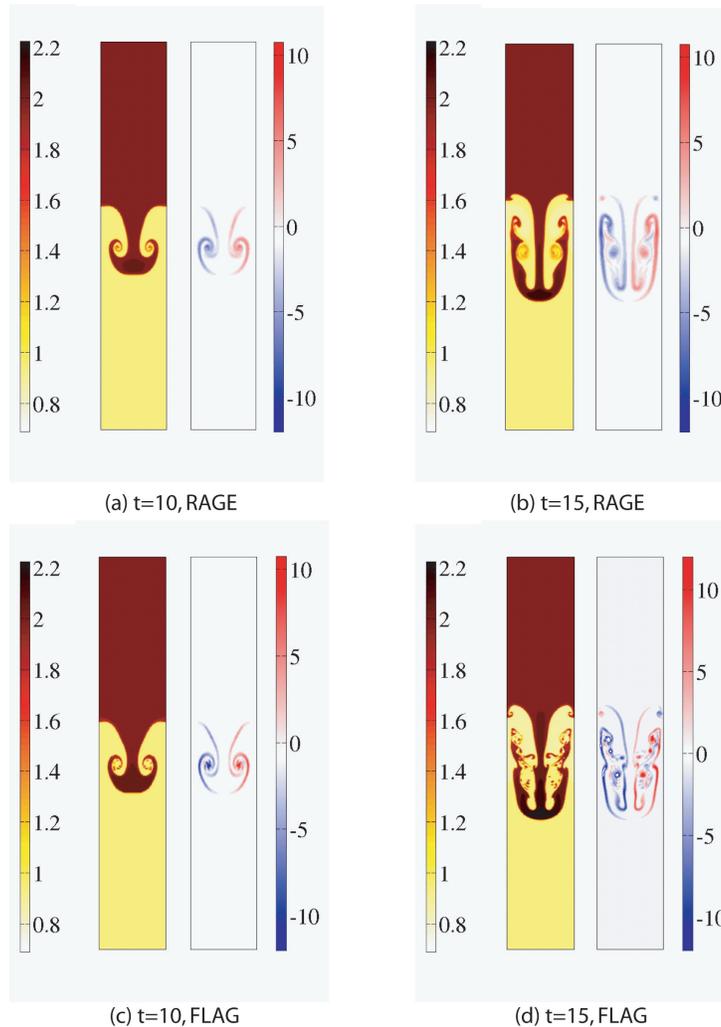
These parameters have included advection schemes for RAGE, mesh strategies for FLAG, and interface treatments for both RAGE and FLAG. This investigation study reported in [5] has shown (1) that both RAGE and FLAG calculations of the RTI are similar in mixing width (interface growth) development, and (2) that the use of high-resolution advection schemes and sharper interface treatments in both hydrocodes produce smaller structures in the flow. Out of all the calculations, the RAGE results using Van Leer as the slope limiter and the interface-capturing method with interface preserver have been found to most resemble the FLAG results when run on a mesh aligned with the initial perturbation and using second-order piecewise linear interface reconstruction LVIRA, as shown in the plots of the density contours and vorticity in Fig. 1. The interface growth metric plotted in Fig. 2 at early times compares well with linear theory. Other metrics such as the surface density and circulation have been considered in [5] to better reflect differences in interface area and symmetry issues, respectively.

This study has identified areas where code improvements are to be considered in the future. In particular, symmetry issues have been observed at later times with both codes using interface reconstruction [5]. In future work, other metrics and image analysis may be developed to better discern the flow structure differences and estimate the manifestation of numerical surface tension and other numerical phenomena that are generated during interface reconstruction.

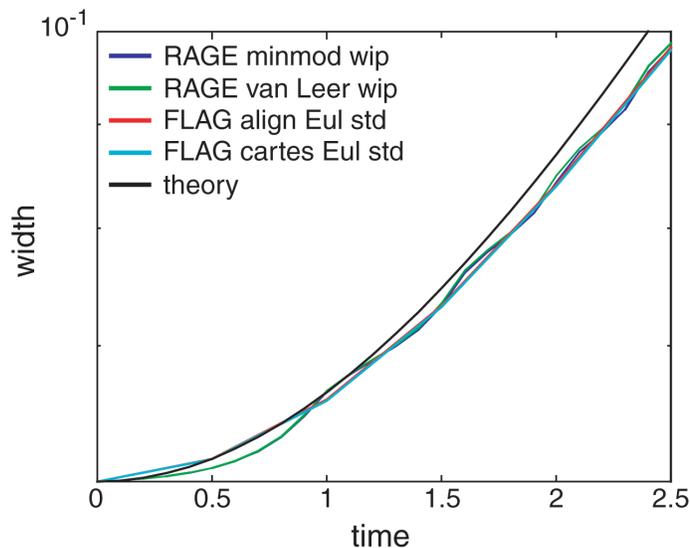
For more information contact Marianne Francois at [mmfran@lanl.gov](mailto:mmfran@lanl.gov).

[1] Lord J. Rayleigh, *Proc. Lond. Math. Soc.* **14**, 170–177 (1883).  
 [2] G.I. Taylor, *Proc. Royal. Soc. Lond. A* **201**, 192–196 (1950).  
 [3] M. Gittings, et al., “The RAGE Radiation-Hydrodynamic Code,” Los Alamos National Laboratory report LA-UR-06-0027 (2006).  
 [4] D.E. Burton, Lawrence Livermore National Laboratory report UCRL-JC-118306 (1994).  
 [5] J. Fung, et al., “Calculations of the Rayleigh-Taylor Instability: RAGE and FLAG Hydrocode Comparisons,” in *Proceedings of Nuclear Explosives Code Developers Conference 2006* (Los Alamos, New Mexico, October 23–27, 2006).

**Funding Acknowledgments**  
 NNSA’s Advanced Simulation and Computing (ASC), Integrated Codes Program.



**Fig. 1.** Density contours (left) and vorticity contours (right) computed with RAGE and FLAG at times  $t = 10$  and  $15$ .



**Fig. 2.** Interface growth vs time at early times computed with RAGE and FLAG and compared to linear theory. RAGE computations are performed using two slope limiters (minmod and Van Leer) with the interface preserver option. FLAG computations are performed on a Cartesian mesh and on mesh aligned with the initial interface perturbation in an Eulerian mode with interface reconstruction.

# Multimaterial Interface Reconstruction Using Particles and Power Diagrams

Rao V. Garimella, Samuel P. Schofield, Raphael Loubere, T-7; Marianne M. Francois, CCS-2

**W**e have developed a new method for volume-conservative reconstruction of multimaterial interfaces in flow simulations. The method reconstructs piecewise-linear interfaces using only the material volume fractions in cells of a computational mesh, and therefore, belongs to the Piecewise Linear Interface Construction (PLIC) class of interface reconstruction methods [1].

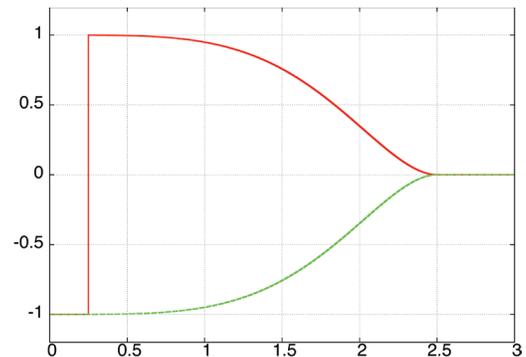
**Fig. 1.** Velocities of two particles of the same material (red) and different materials (green) as a function of particle separation.

The distinguishing feature of our method is that the reconstructed interface in cells with three or more materials has the same topology, independent of the order in which materials are specified. This is different from other PLIC methods, which generate substantially different reconstructions for different material orderings.

Order-independent reconstruction of the interface with the right topology reduces errors such as bits of material breaking off from the bulk of the material and helps to capture the complex physics at multimaterial junctions more accurately.

The technique we have developed reconstructs interfaces in multimaterial cells using particle attraction and repulsion to locate the materials in the cells, and a weighted Voronoi diagram called a Power Diagram [2] to partition the cell into subcells of pure materials.

Initially, particles of different material types are placed randomly in mixed cells containing three or more materials and all neighboring cells (pure or mixed). An attraction-repulsion model that defines particle velocities as a function of particle separation is initiated between the particles. (See Fig. 1.) In the model, particles of the same materials attract each other, in general, and particles of different types repel. When two particles of the same type get too close



to each other, however, they repel each other. When any two particles move too far apart, there is no interaction between them. Using this model, particles are then allowed to migrate to their final positions. The system is stabilized using a simulated annealing type rescaling of particle velocities. The centroids of the clustered particles in their final position determine the location of the different materials in the cell.

The partitioning of the cell into subcells of different materials is done by a Power Diagram of the centroids of the materials. Power diagrams are a partitioning of space that bear a close resemblance to simple Voronoi diagrams. In a power diagram, the *power* of a point,  $x$ , with respect to a site,  $X_i$ , with weight  $w_i$ , is defined as  $(\|x - X_i\|^2 - w_i^2)$ . The *power region* of a site  $X_i$  is the locus of points whose power with respect to  $X_i$  is lesser than their power with respect to any other site  $X_j$ .

The weights of the sites in each cell control the relative areas of their associated polygons. Therefore, if each site in the cell represents the centroid of a material, then the material volume fractions in the cell can be controlled by adjusting the weights. The weights required to exactly match the material volume fraction in any cell are computed by solving the following system of equations:

$$F(w_1, w_2, \dots, w_n) - V = 0$$

A scalar equation in the system represents the discrepancy between the actual volume fraction and the specified volume fraction of a particular material. There are as many equations in the system as there are clusters of particles. Note that there may be more equations than the number of materials since particles of a material may cluster into two or more groups.

We have implemented this method for interface reconstruction in structured and unstructured meshes, and have seen that the method works quite well in capturing the topology of such interfaces. We show some results here to illustrate the capabilities of this method.

In Fig. 2, we illustrate the performance of the particle model. On the left, the initial configuration of particles in a  $4 \times 4$  structured grid is shown. The exact interface in this mesh is shown as a dark diagonal line. The final position of the particles is shown on the right. It can be seen from this figure that the method moves the particles to the positions we expect them to be in.

Figure 3 compares the performance of a traditional Volume of Fluid (VoF) method and the current method in the reconstruction of a three-material T-junction interface. The reconstruction with the PLIC method is perfect when the correct order of materials is chosen (top left), but not very good when a different order is used (top right). On the other hand, our method creates the right interface topology whatever the material ordering (bottom).

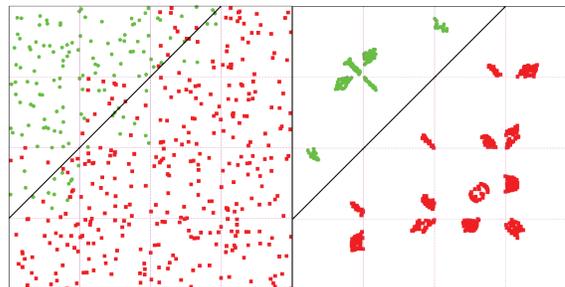
Future work in this research includes smoothing of the interfaces, better particle models to capture details such as filaments, and incorporation of clustering algorithms to detect multiple groups of particles of the same type in a cell.

For more information contact Rao V. Garimella at [rao@lanl.gov](mailto:rao@lanl.gov).

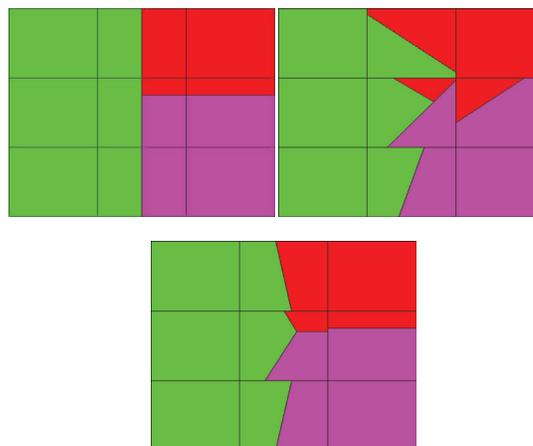
- [1] W.J. Rider and D.B. Kothe, *J. Comp. Phys.* **141**, 112–152 (1998).
- [2] F. Aurenhammer, *SIAM J. Comput.* **16** (1), 78–96 (1987).

**Funding Acknowledgements**

NNSA’s Advanced Simulation and Computing (ASC) Program.



**Fig. 2.** Left—initial distribution of particles in a mesh; Right—final distribution of particles after application of the particle model. The dark line shows the exact interface from which the input volume fractions were computed.



**Fig. 3.** Three-material reconstruction by a traditional PLIC method with the right ordering of materials (top left), PLIC reconstruction with a different ordering of materials (top right), and reconstruction using our order-independent algorithm (bottom).

## Simulation of the Rayleigh-Taylor Instability Using Atomistic Methods

John L. Barber, Kai Kadau, Peter S. Lomdahl, T-14; Timothy C. Germann, Brad L. Holian, X-1-SMMP; Berni Alder, Lawrence Livermore National Laboratory

**T**he Rayleigh-Taylor instability (RTI) occurs when a dense fluid lies on top of a light fluid in the presence of a vertical gravitational field.

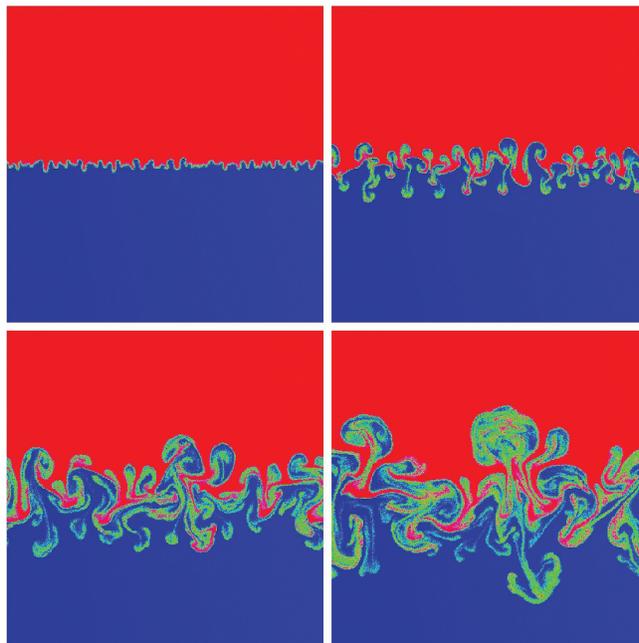
This arrangement is an unstable equilibrium, and the subsequent behavior of the two fluids is an archetypical example of turbulent fluid mixing. The RTI has applications in many diverse fields, such as astronomy, oceanography, and inertial confinement fusion, and it has been studied extensively via both simulation and experiment.

Most previous simulated studies of the RTI have involved the use of continuum techniques based on the Navier-Stokes (NS) or Euler equations. In this work [1,2], we instead make use of two atomistic methods: molecular dynamics (MD) and direct simulation Monte Carlo (DSMC) [3]. In MD, Newton's equations of motion for the positions and momenta of a large number of interacting particles are solved numerically. DSMC is similar in many respects to MD, except for the fact that the interparticle interactions are modeled stochastically,

which results in a speed up factor of approximately 30–50 compared with MD.

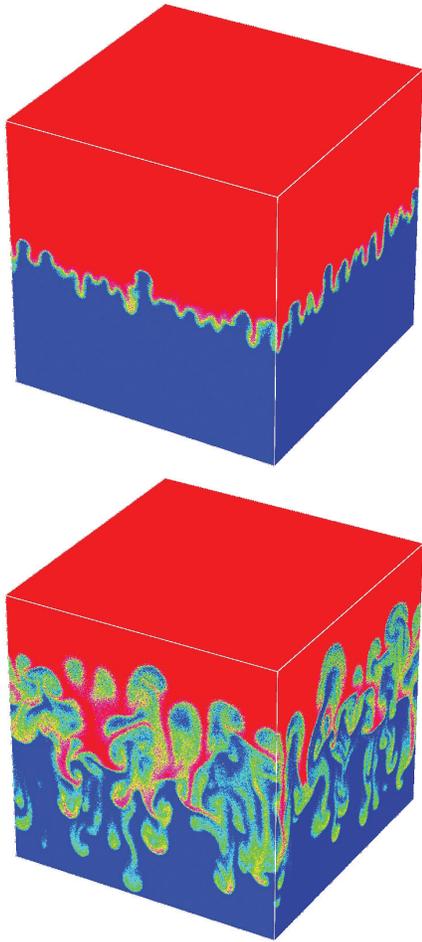
There are several advantages to the use of atomistic methods to fluid turbulence. Continuum models, such as the NS equations, are known to fail in several regimes, such as in extremely rarified (high Knudsen number) flows, and flows at very small length scales. More fundamental atomistic methods are applicable in a much wider set of regimes. Furthermore, for applications involving interface dynamics, such as the RTI, MD, and DSMC have the advantage of allowing an initially flat interface, as opposed to the somewhat artificial initial perturbation spectrum that must be applied in a continuum simulation. In addition, it is much easier to control the degree of miscibility in atomistic methods (particularly MD). The main disadvantages of atomistic methods are the small length and time scales to which they are confined by the limits of current computational capacity. These limits are continually being pushed back, however, as evidenced by the simulations performed in this work.

**Fig. 1.**  
Four snapshots in time from a 2-D, 500-million particle simulation of the RTI using DSMC.



Several large-scale DSMC simulations of the RTI were performed using the Scalable Parallel Short-range MD code [4] on Livermore's BlueGene/L machine. Figure 1 shows the results of a 2-D run involving 500 million particles, corresponding to length and time scales of  $5 \mu\text{m}$  and 140 ns, respectively. Figure 2 shows some results from a 3-D run of 7 billion particles (a new world record), corresponding to  $1 \mu\text{m}$  and 20 ns. The bubble-formation and merger process characteristic of the RTI is evident in both.

During the initial stages of the RTI, linear stability analysis predicts



For more information contact John L. Barber at [jlbarber@lanl.gov](mailto:jlbarber@lanl.gov).

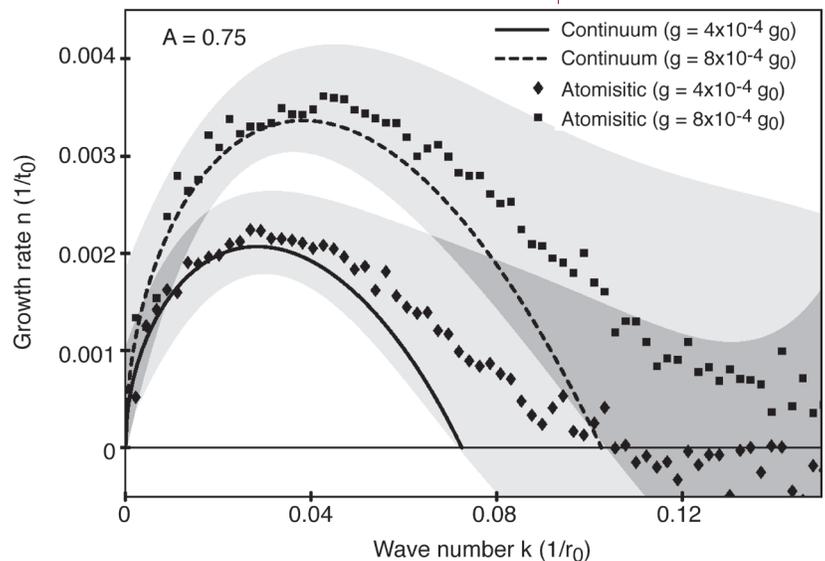
- [1] K. Kadau, et al., *PNAS* **101**, 5851 (2004).
- [2] J.L. Barber, et al., *J. Phys.: Conference Series* **46**, 58–62, (2006).
- [3] G.A. Bird, *Molecular Gas Dynamics and the Direct Simulation of Gas Flows* (Oxford University Press, Oxford, England, 1994).
- [4] P.S. Lomdahl, et al., *Proceedings of Supercomputing 93*, G.S. Ansell, Ed. (IEEE Computer Society Press, Los Alamitos, CA, 1993), p. 520.
- [5] S. Chandrasekhar, *Hydrodynamic and Hydromagnetic Stability* (Dover, New York, 1961).

#### Funding Acknowledgements

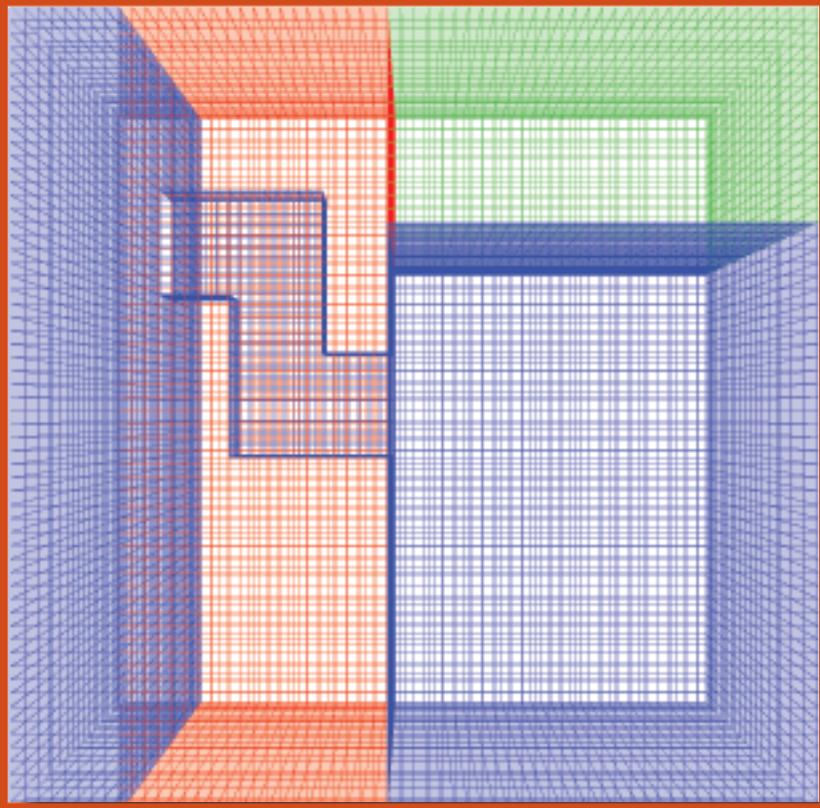
Laboratory Directed Research and Development Program.

**Fig. 2.** Two snapshots in time from a 3-D, 7-billion particle simulation of the RTI using DSMC.

that the various modes present on the interface will grow independently and exponentially with a growth rate  $n(k)$ , where  $k$  is the wavenumber [5]. As a test of how well atomistic simulations agree with continuum predictions, we have used MD to calculate the growth rate for two different gravities. The results are shown in Fig. 3. Note that the shaded envelopes are not error bars, but are instead an estimate of the intrinsic physical spread in the growth rate, which occurs due to the fact that all macroscopic fields (such as density or fluid velocity) become random variables subject to microscale fluctuations at these small length scales. There is good agreement with the continuum prediction at small wave numbers, but increasing discrepancies at larger  $k$ , due to the effects of capillary waves, compressibility, and fluctuations.



**Fig. 3.** The initial growth rate spectrum as calculated via MD for two different gravities. The lines represent the theoretical continuum predictions, and the shaded envelopes reflect an estimate of the physical spread in growth rate at these length scales. (Results in simulation units.)



# Radiation Transport Methods and Algorithms

Like its hydrodynamic cousin, radiation transport is another fundamental physical process in the evolution of the explosion of a nuclear weapon. The articles in this section are centered on the latest developments in this field in critical areas such as new Monte Carlo techniques for efficient and accurate simulations on complex mesh types, important new developments for the implicit Monte Carlo method, verification and validation results for key radiation transport schemes, a valuable new methods result which may lead to a more precise and efficient convergence of transport simulations, and a collage of student projects in radiation transport methods representing a key talent pipeline of critical importance to Los Alamos National Laboratory.

# Discrete Diffusion Monte Carlo for Radiation-Transport Simulations on Adaptive-Refinement Meshes

Jeffery D. Densmore, Thomas M. Evans, Michael W. Buksas, CCS-2

**D**iscrete Diffusion Monte Carlo (DDMC) is a technique for increasing the efficiency of Monte Carlo simulations in diffusive media. Monte Carlo simulation is a well-established method for modeling energy transfer by x-ray radiation transport in applications such as astrophysics and inertial confinement fusion, but can be computationally inefficient in problems with diffusive regions. In this regime, particle histories will consist of many steps due to the small distance between collisions and lack of absorption, a situation that results in an inefficient radiation-transport calculation.

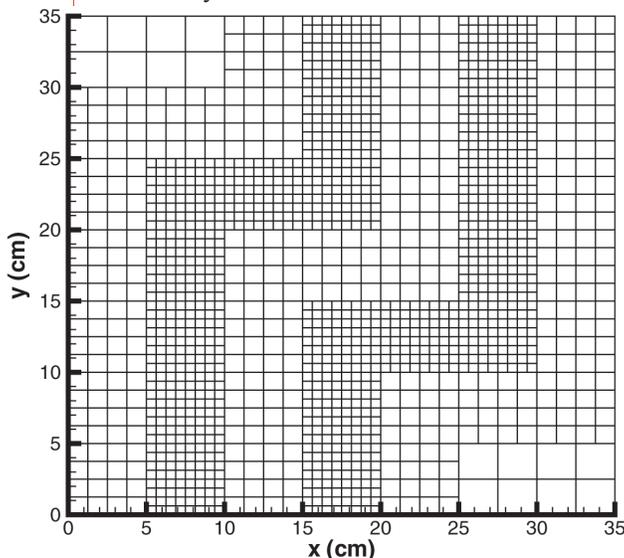
In DDMC, particles take discrete steps between spatial cells according to a spatially discretized diffusion equation. Since each discrete step replaces many smaller Monte Carlo steps, DDMC is more efficient than standard Monte Carlo in diffusive media. In practice, DDMC is combined with standard Monte Carlo in nondiffusive regions to form a hybrid method.

In previous work, we have developed techniques for interfacing DDMC and Monte Carlo simulations [1,2] and presented a hybrid Monte Carlo-DDMC method for nonlinear, time-dependent problems [3]. Recently, we have developed a DDMC technique for adaptive-refinement meshes [4], a type of mesh in which spatial cells may have multiple neighboring cells across each face. Adaptive-refinement meshes are commonly used in Eulerian hydrodynamics calculations, so extending our DDMC methodology to these meshes is important for coupled radiation-hydrodynamics simulations.

We compare our DDMC technique for adaptive-refinement meshes to both standard Monte Carlo and Random Walk (RW), another hybrid Monte Carlo-diffusion method in which particles take macrosteps over spheres according to an analytic diffusion solution [5]. We consider a test problem consisting of a purely scattering region with an embedded duct. The computational mesh for this problem is given in Fig. 1. We form a hybrid method by employing DDMC in the purely scattering region and standard Monte Carlo in the duct. In addition, we use two values of the minimum sphere radius,  $\lambda$ , with RW.

A contour plot of the scalar flux for this problem is given in Fig. 2. In Fig. 3, we plot the scalar flux resulting from these calculations as a function of  $x$  for fixed  $y$ . From this figure with see that the hybrid and  $\lambda = 20$  solutions agree well with standard Monte Carlo, while the  $\lambda = 5$  scalar flux is systematically too low.

Fig. 1. Computational mesh.



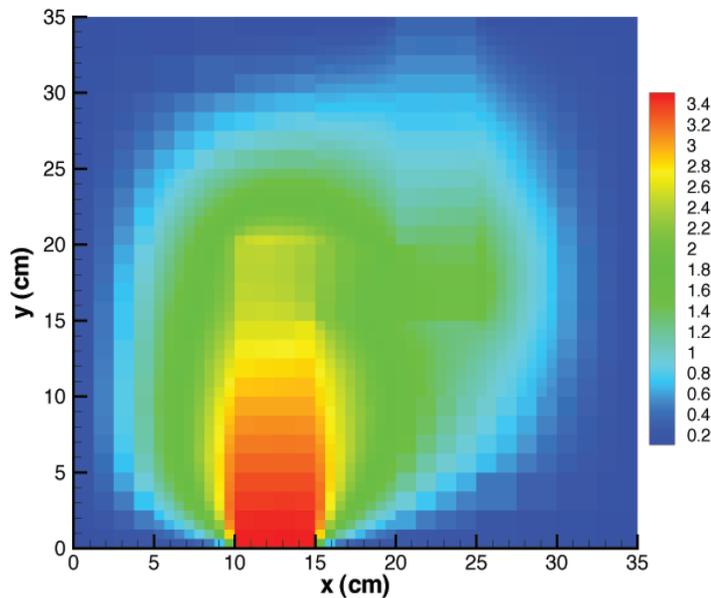
In addition, the hybrid simulation was almost 500 times faster than standard Monte Carlo and over 300 times faster than the most efficient (but inaccurate) RW calculation. Thus, our new DDMC method for adaptive-refinement meshes is more computationally efficient than standard Monte Carlo and RW while yielding accurate solutions.

For more information contact Jeffery Densmore at [jdd@lanl.gov](mailto:jdd@lanl.gov).

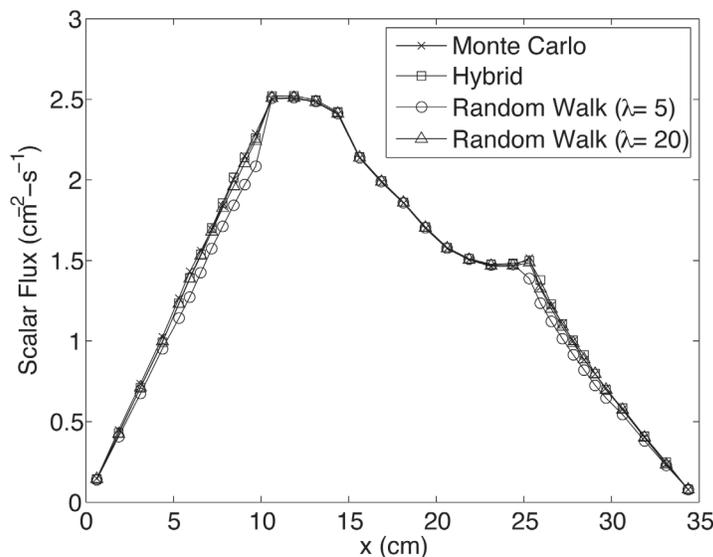
[1] J.D. Densmore, *Ann. Nucl. Energy* **33**, 343 (2006).  
 [2] G. Davidson, et al., *Trans. Am. Nucl. Soc.* **94**, 517 (2006).  
 [3] J.D. Densmore, et al., *J. Comput. Phys.* **222**, 485 (2007).  
 [4] J.D. Densmore, et al., *Trans. Am. Nucl. Soc.* **95**, 541 (2006).  
 [5] J.A. Fleck, Jr., and E.H. Canfield, *J. Comput. Phys.* **54**, 508 (1984).

**Funding Acknowledgements**

NNSA's Advanced Simulation and Computing (ASC), Computational Physics and Methods Strategic Capability.



**Fig. 2.** Contour plot of scalar flux ( $\text{cm}^{-2} - \text{s}^{-1}$ ).



**Fig. 3.** Scalar flux at  $y = 15.3125 \text{ cm}$ . This value of  $y$  corresponds to the row of cells just after the duct makes a turn to the right.

## Residual Monte Carlo Methods

Thomas Evans, Scott Mosher, CCS-2

In Ref. [1], we developed a residual Monte Carlo method that was an application of Halton's Sequential Monte Carlo method [2] to the (1-D) equilibrium (1T) thermal radiation diffusion equation. At that time, Halton's method was not widely known among numerical practitioners. While extending our method to 3-D, we have discovered that the Sequential Monte Carlo method is actually a variant of iterative refinement cast as a residual method. Based more rigorously on iterative refinement, our new Monte Carlo method surpasses our previous 1-D method and allows for extension to more general solution techniques.

The goal of recent work is to develop efficient Monte Carlo solvers for discrete systems. We have developed fast linear solvers using Monte Carlo techniques with two parallel approaches in this effort: one approach for general linear systems and the other for nonsymmetric, multidimension, banded systems that result from discretized partial differential equations.

The focus of the second approach is to solve 3-D diffusion equations for nonlinear thermal radiative transfer. Along these lines, we have extended our previous work by developing an Iterative-Refinement Monte Carlo (IRMC) method for solving sparse matrix systems. We applied this solver to the 1T thermal radiation diffusion equation in 3-D Cartesian geometry.

Consider the following preconditioned, discrete linear system,

$$\mathbf{M}^{-1}\mathbf{D}\phi = \mathbf{M}^{-1}\mathbf{q}. \quad (1)$$

The only requirement on the preconditioned system is that the spectral radius of  $\mathbf{M}^{-1}\mathbf{D}$  must be less than 1. The IRMC method that solves (1) is defined

$$\phi^{l+1/2} = (\mathbf{I} - \mathbf{M}^{-1}\mathbf{D})\phi^l + \mathbf{M}^{-1}\mathbf{q},$$

(fixed-point iteration)

$$\mathbf{r}^{l+1/2} = \mathbf{M}^{-1}\mathbf{q} - \mathbf{M}^{-1}\mathbf{D}\phi^{l+1/2},$$

(estimate residual)

$$\mathbf{M}^{-1}\mathbf{D}\delta\phi^{l+1/2} = \mathbf{r}^{l+1/2},$$

(solve using Monte Carlo)

$$\phi^{l+1} = \phi^{l+1/2} + \delta\phi^{l+1/2}$$

(update  $\phi$ )

The update of the correction term is performed using a Monte Carlo transport calculation [3].

We have applied the IRMC method to the 3-D, 1T, nonlinear radiation diffusion equation. A multimaterial duct problem is shown in Fig. 1. This problem features a 0.5 keV blackbody flux on the low- $x$  side. Radiation is propagated through a dog-legged duct bounded by an opaque wall. An optically thick foil is placed on the high- $y$  side of the outlet. A contour plot of the solution is shown in Fig. 2, and the time-evolution of the temperature at four edit points is shown in Fig. 3. The Monte Carlo solution can be run to arbitrary precision because the convergence of the IRMC method is not bound by the Central Limit Theorem.

We have compared the IRMC method on this problem with standard solution techniques. The IRMC method compared with preconditioned Conjugate Gradient (CG) and Richardson iteration are shown in Table 1. The IRMC method compares very favorably to preconditioned Richardson iteration and is marginally faster than preconditioned CG.

Table 1: Comparison of solution methods for the multimaterial problem. The problem was run to an elapsed time of 10 sh requiring 10157 cycles.

Method	Max Iterations	CPU Time (s)
Preconditioned CG	18	12794.3
Preconditioned Richardson	38	17677.6
IRMC	20	12367.4

We have presented a new Monte Carlo solution method for solving the discrete, time-dependent diffusion equation in 3-D. The IRMC method has been shown to match results using standard solution techniques to arbitrary precision. Also, the new method is faster than preconditioned CG and Richardson iteration.

While we have demonstrated marginal improvements over standard solution schemes in this study, significant improvements could be realized in fully nonlinear-consistent solutions. In these cases, the IRMC method competes with GMRES, which is more costly in memory and time than CG. Another area where the IRMC scheme may have advantages over traditional solution methods is on dentritic, or adaptive, meshes. These meshes yield matrices with poor condition numbers because of the changing cell volumes at different refinement levels. A smart transport algorithm could be developed that more efficiently solves the residual on these types of meshes. These topics will be the focus of future study.

For more information contact Thomas M. Evans at [tme@lanl.gov](mailto:tme@lanl.gov).

[1] T. Evans, et al., *J. Comp. Phys.* **189**, pp. 539–556, 2003.  
 [2] J. Halton, *J. Sci. Computing* **9**(2), pp. 213–257, 1994.  
 [3] J. Hammersly and D. Handscomb, *Monte Carlo Methods* (Spottiswoode, Ballantyne, and Co., London, 1964).

#### Funding Acknowledgements

NNSA's Advanced Simulation and Computing (ASC) Computational Physics and Methods Strategic Capability program, Advanced Discrete Monte Carlo Project.

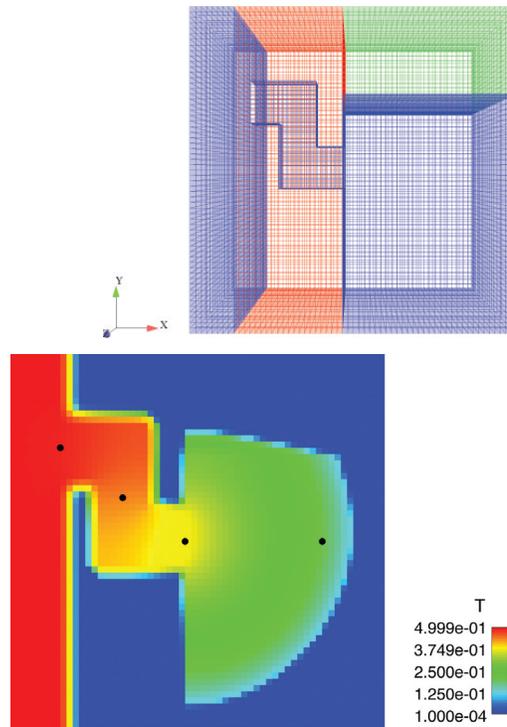


Fig. 1. Multimaterial problem mesh (60x60x60): blue shows the duct region, red shows the wall, and green shows the foil.

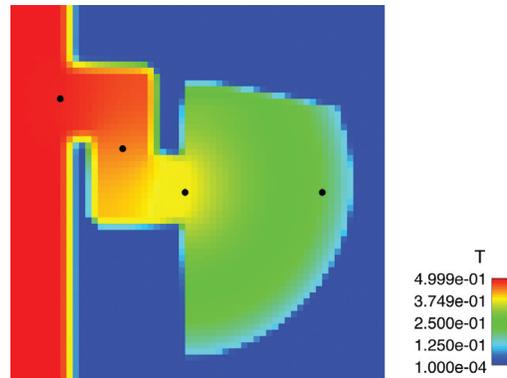


Fig. 2. Contour plot of the temperature at 100 sh on a cut-plane positioned at the midpoint of the z-axis.

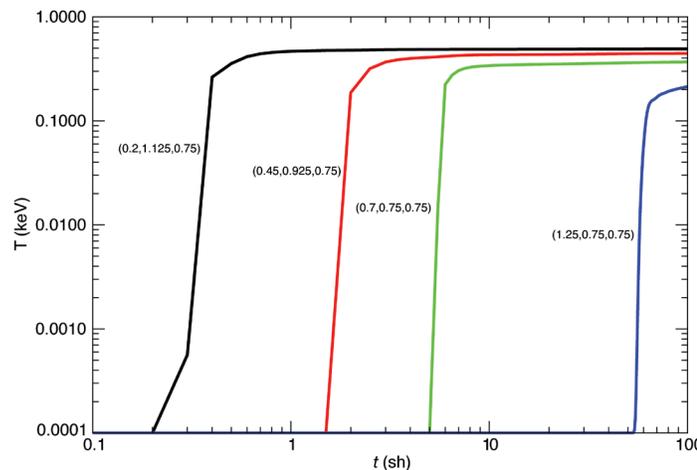


Fig. 3. Time evolution of the temperature at the edit points shown in Fig. 2.

## Implicit Monte Carlo Methods for Three-Temperature Physics

Thomas Evans, Jeffery Densmore, CCS-2

The simulation of thermal radiation propagation ranks among the most difficult class of transport problems. These problems are highly nonlinear, and the fundamental unknown (the radiation intensity) can be a function of seven independent variables [in three dimensions (3-D)]. One of the most successful and widely used methods in thermal radiation transport is the Implicit Monte Carlo (IMC) method [1]. This method is a two-temperature (2T) scheme that includes radiation and material coupling where the matter is represented by a single temperature.

A more accurate description of the radiation and material coupling represents the ions and electrons by distinct, separate temperatures [2]. The resulting three-temperature (3T) equations for the time evolution of the radiation, electron, and ion energies include terms representing electron-ion coupling and conduction [3]. Conventionally, this system of equations is solved using radiation diffusion with operator-split conduction and coupling [4].

The objective of this study is to extend the standard IMC method to include 3T physics. Descriptions of matter that include separate energies for electrons and ions are important in high energy density physics applications and astrophysics [5]. We derive three methods for solving the 3T equations using IMC. The first method (FSIMC) uses the standard IMC technique to simulate radiation transport. The conduction and coupling terms are linearized and split into separate equations that are solved independently. A second method (SCIMC) uses a more robust splitting scheme in which half of the coupling is treated during the transport simulation. The conduction is split from the ion and electron equations and is solved subsequently. Afterwards, the second half of the coupling is solved.

The third method (ECIMC) treats the conduction explicitly and includes all of the conduction and coupling in the linearization of the transport equation. This is a good approximation when the conduction timescales are much longer than the radiation-transport timescales. We expect this to be the case for most problems because the conduction timescales are related to the electron thermal velocity whereas the radiation moves at the speed of light. The resulting system has three equations: a Monte Carlo transport equation and two decoupled energy equations for electrons and ions.

We use Modified Equation Analysis (MEA) [6] to estimate the errors that result from operator splits and linearization. In short, MEA uses Taylor series expansions to convert from a discrete to continuous system of equations. The resulting set of *modified* equations consists of the original continuous system plus error terms. Figure 1 shows the results of each IMC method on a source-driven 3T test problem. Using MEA, we have plotted the magnitude of the error terms for this problem in Fig. 2. The error terms A and B dominate, and because these terms are only present in the SCIMC and FSIMC methods, they account for the improved accuracy of the ECIMC method illustrated in Fig. 1.

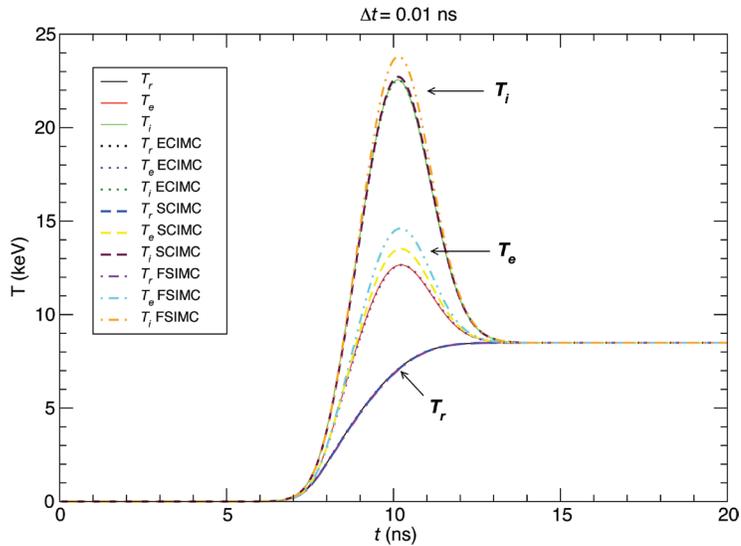
In this study we have presented three IMC methods for performing 3T radiation-transport calculations. We have shown that the ECIMC method is considerably more accurate than the SCIMC and FSIMC methods in problems where the ions and electrons are decoupled. The ECIMC method will only be viable in problems where the electron and ion conduction timescale is slow compared to the radiation timescale. In cases where the conduction

timescale is of the same order as the radiation timescale, the SCIMC method is most optimal. Future work includes extending these methods to fully analyze the effects of radiation conduction and extension to multifrequency.

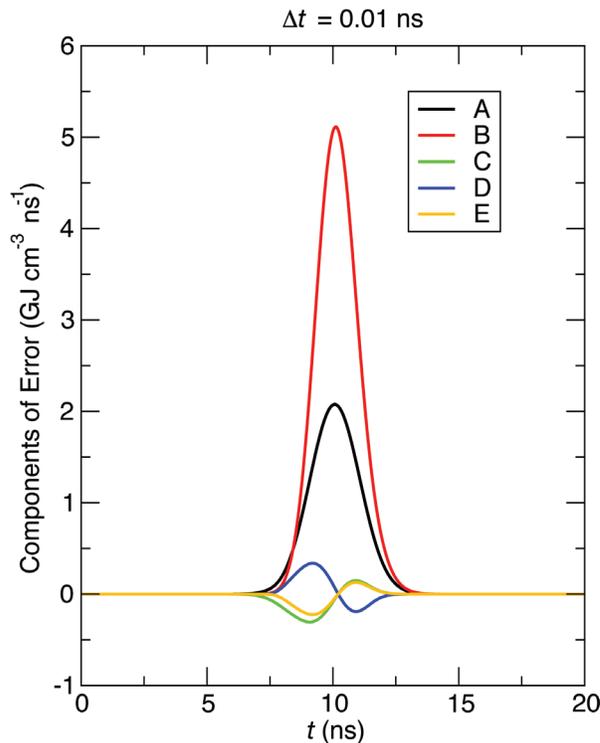
For more information contact Jeffery Densmore at [jdd@lanl.gov](mailto:jdd@lanl.gov)

[1] J.A. Fleck and J.D. Cummings, *J. Comp. Phys.* **8**(3), pp. 313–342 (1971).  
 [2] M. Sincell, et al., *Shock Waves* **9**, pp. 403–411 (1999).  
 [3] J. Huba, *NRL Plasma Formulary*, Naval Research Laboratory, Washington, D.C., 2004.  
 [4] G. Olson and J.E. Morel, "Solution of the Radiation Diffusion Equation on an AMR Eulerian Mesh with Material Interfaces," Los Alamos National Laboratory report LA-UR-99-2949 (June 1999).  
 [5] T.D. Matteo, et al., *Mon. Not. R. Astron. Soc.* **291**, pp. L23–L27 (1997).  
 [6] D. Knoll, et al., *J. Comp. Phys.* **185**, pp. 583–611 (2003).

**Funding Acknowledgements**  
 NNSA's Advanced Simulation and Computing (ASC), Computational Physics and Methods Mission Capability Program, Transport Project.



**Fig. 1.** Comparison of three IMC methods on a Gaussian ion source problem. The ECIMC method is the most accurate.



**Fig. 2.** Magnitude of the error terms that result from the linearization and splitting strategies employed in the three IMC methods. The C, D, and E terms are present in all three methods. The A and B terms are only present in the SCIMC and FSIMC methods. The B term is twice as large in the FSIMC method (not plotted).

## Computational Cost Analysis of the Milagro Implicit Monte Carlo Code

Scott W. Mosher, Timothy M. Kelley, CCS-2

Developers of production-level transport codes are motivated (by their user community if not by their own curiosity) to analyze and understand the computational cost of their methods and algorithms. Ideally, one would like to know an algorithm's execution time as a function of various input parameters. The information is useful to users for predicting run times, to developers for evaluating the software implementation, and to both groups for understanding the limits of the algorithm's practical application. To this end, we have completed a systematic survey [1] of the serial (i.e., single processor) performance of the Milagro code [2] for nonlinear, time-dependent, radiative transfer simulations based on the Implicit Monte Carlo method [3]. More recently, we have developed a mathematical model to predict the execution time of Milagro transport simulations.

The computational cost of an implicit Monte Carlo calculation depends on several input variables, such as the number of geometric cells, cell size, material opacity, number of frequency groups, etc. In the survey described in Ref. [1], we measured the computational cost of several sections of the Milagro code over a wide range of input parameters. We analyzed the results to compare the relative cost of different sections, as well as the scaling of the costs with respect to changes in the input parameters. That analysis confirmed that the relative costs and cost scaling of nearly all code sections were as expected. The analysis also

identified something unexpected—that a major consumer of CPU time (up to 50% in some cases) was the system-level memory allocation function. Further analysis found that the allocation and deallocation patterns in the Milagro code caused the allocation routine to spend a tremendous amount of time sorting freed blocks of memory. The bottleneck was eliminated by switching to an alternative allocation routine.

The data generated during the survey were catalogued and comprise a performance baseline with which to compare timing results from future versions of the code. A similar exercise was completed as part of this project. That is, a comparison was made between the then-current version of the code and version 4\_1\_0 from 2004. It was found that the computational efficiency of the code has increased substantially as a result of several improvements. For example, the total computational time as a function of the number of photon frequency groups for both versions running a sample problem is plotted in Fig. 1. It can be seen that the overall cost of the simulation was reduced by about 25 to 50%.

In recent work, we have been developing a mathematical model to predict the computational cost of the transport operations of Milagro. In a Monte Carlo transport algorithm, the execution time per particle scales with the number of events (e.g., scattering collisions, interface crossings, etc.) that occur in the simulation. The number of events per unit simulation time depends on the input parameters in

a complicated way. For an idealized problem, the mean number of events (of various types) can be predicted analytically. From these analytic expressions, a cost model has been constructed to predict the average computational time consumed per transport cycle. The model contains parameters that depend on the implementation of the transport algorithm and the computer system on which the simulation is conducted. For a particular implementation running on a particular system, the parameters can be estimated by numerically fitting the parametric model to actual timing data.

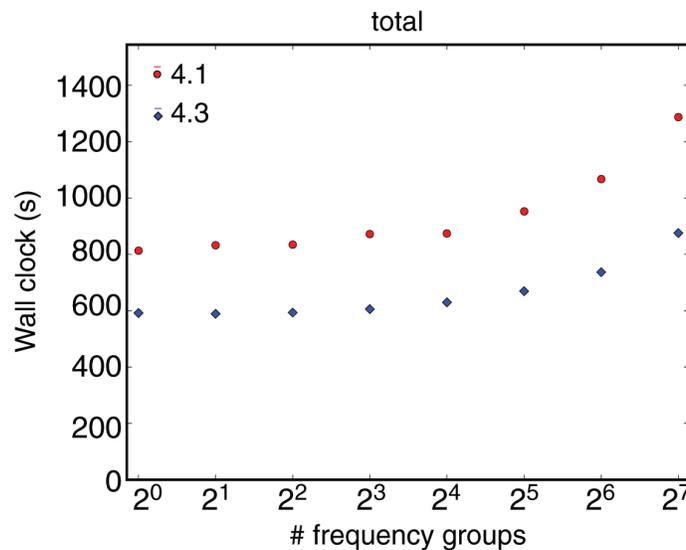
Timing data were collected for a series of Milagro runs in which the independent variables were varied over reasonable sets of values. This led to a total of 144 individual problems (i.e., distinct combinations of the independent variables). The cost model was fit to the Milagro timing data and compared to the actual times. For the full set of data, the average absolute relative error in the predicted transport times was 1.97%. The maximum absolute relative error was 33.02%, which was much higher than expected. There were seven problems for which the predicted times were outside of three estimated standard deviations from the mean. In each of these cases, the model under-predicted the actual time. The actual number of events occurring in these simulations was tallied, and it was found that the model predicted the number of events to within 0.1%. Hence, it currently appears that the additional time was spent performing operations

not incorporated into the model. This issue is currently being investigated. Excluding the seven outliers, the average absolute relative error in mean was 0.97% and the maximum was 4.31%.

**For more information contact Scott Mosher at [swmosher@lanl.gov](mailto:swmosher@lanl.gov).**

[1] T.M. Kelley and S.W. Mosher, "Characterization of Implicit Monte Carlo Projects, Phase 1: Serial Performance of Milagro," Los Alamos National Laboratory memorandum CCS-4:06-03(U) (2006).  
 [2] T.J. Urbatsch and T.M. Evans, "Milagro Version 2 – An Implicit Monte Carlo Code for Thermal Radiative Transfer: Capabilities, Development, and Usage," Los Alamos National Laboratory report LA-14195-MS (2006).  
 [3] J.A. Fleck, Jr. and J.D. Cummings, *J. Comp. Phys.* **8**, 313 (1971).

**Funding Acknowledgements**  
 NNSA's Advanced Simulation and Computing (ASC), Computational Physics and Methods Mission Capability Program, Transport Project.



**Fig. 1.** Log-log plot of total wall-clock time vs number of groups. Red circles: head version 4\_3\_0+ (Dec. 2, 2005), blue diamonds: Milagro-4\_1\_0 (2004).

# Analysis of the Stability, Monotonicity, and Time Discretization Error of the Implicit Monte Carlo Method

Scott W. Mosher, Jeffery D. Densmore, CCS-2

The Implicit Monte Carlo (IMC) method for nonlinear, time-dependent, radiative transfer calculations employs a semi-implicit approximation to the photon absorption-emission process within a time step [1]. This approximation originates from a backward-Euler time discretization of the exact equations followed by a linearization of the radiation-material coupling. An IMC calculation proceeds in a series of time steps, each of which consist of a linear radiation transport calculation followed by a nonlinear material energy update. Although the semi-implicit time discretization allows the use of larger time steps than a purely explicit technique, the method can produce nonphysical solutions if the time step is too large [2]. We are motivated to understand the limitations of this approximation in order to ultimately develop an *a priori* time-step control algorithm that avoids nonphysical behavior.

To this end we have studied gray (frequency-independent), 0-D (space-independent) problems, in which the radiation and material energy densities initially coexist in a nonequilibrium state. As energy is exchanged between the material medium and radiation field (via photon absorption and emission) the system approaches thermal equilibrium asymptotically in time. The time-dependent radiation and material energy densities are described by a differential radiation transport equation that is coupled to a differential material energy density equation. The equations can be either linearly or nonlinearly coupled, depending on the form of the material heat capacity. For the linear case, we have developed conditions that guarantee that the exact IMC solution (neglecting statistical error) is stable (i.e., progresses toward the analytic equilibrium solution) and monotone

(i.e., does so without nonphysical oscillations) [3]. For the nonlinear case, we have obtained an exact solution and have used it to numerically estimate the time order of accuracy of the IMC method [4].

For the linear case, the IMC method generates end of time-step energy densities through a linear transformation of the beginning of time step values

$$\mathbf{x}_{n+1} = \mathbf{A}\mathbf{x}_n \quad (1)$$

where  $\mathbf{x}$  denotes the radiation and material energies. If the time step is constant, then  $\mathbf{A}$  is a constant  $2 \times 2$  matrix of coefficients for the problem at hand. The time stability and monotonicity of the IMC approximation can be investigated by expanding the solution of the IMC equations in the eigenvectors of the  $\mathbf{A}$  matrix

$$\mathbf{x}_n = a\lambda_1\varphi_1 + b\lambda_2^n\varphi_2. \quad (2)$$

where  $\lambda$  and  $\varphi$  denote the eigenvalues and eigenvectors of  $\mathbf{A}$ , respectively, and  $a$  and  $b$  are constants. The first term on the right-hand side is the exact equilibrium state of the system. The second term represents the transient behavior between the initial condition and thermal equilibrium. Since  $\lambda_1 = 1$ , the method is stable whenever  $|\lambda_2| < 1$  and monotonic whenever  $0 \leq \lambda_2 < 1$ . From these conditions, time-step limits have been derived that guarantee stability and monotonicity of the method. These limits have been verified numerically. In Fig. 1, the IMC energy densities are plotted (in arbitrary energy and time units) for a problem in which the time step was chosen to be half of the stability limit. It is clear that the solution is stable but not monotonic and hence nonphysical. In Fig. 2, the time step was chosen to be half of the monotonicity limit. The exact solution is also plotted, and it can be seen that the IMC method yields a physically valid result.

For the nonlinear case, stability and monotonicity limits are much more difficult to obtain since the  $\mathbf{A}$  matrix is no longer constant in time. However, a solution to the exact radiative transfer equations has been obtained (whereas most existing benchmarks have been developed for linear problems). The solution has been used to benchmark the IMC approximation and determine its order of accuracy with respect to the time variable. For that purpose, we examine the integrated error in the IMC-predicted time-dependent material temperature

$$\text{RMS Error} = \sqrt{\frac{1}{N} \sum_{n=1}^N [T_{n+1} - T(t_{n+1})]^2}, \quad (3)$$

where  $T_{n+1}$  and  $T(t_{n+1})$  denote the IMC and exact end of time-step material temperatures, respectively. In Fig. 3, the RMS error is plotted vs time-step size for  $\alpha = 0, 0.5,$  and  $1$ . Here,  $\alpha$  (where  $0 \leq \alpha \leq 1$ ) is a user-selected parameter that controls the time-centering of the absorption-reemission approximation. As expected, the numerical results indicate that the method is first-order accurate in time-step size for  $\alpha = 0$  and  $1$ , and second-order accurate for  $\alpha = 0.5$ . It should be noted that IMC absorption-reemission approximation is most stable when  $\alpha = 1$ , and often this is the best choice.

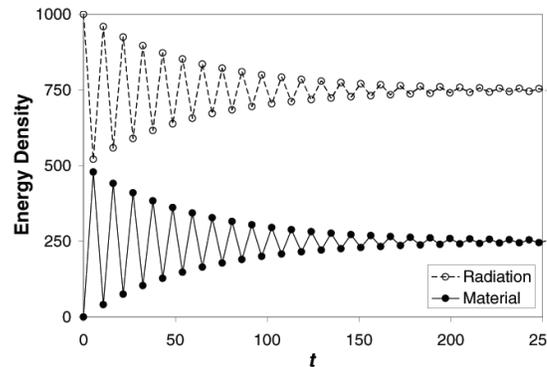
Through the analysis of idealized radiative transfer problems, we have learned much about the stability, monotonicity, and accuracy of the IMC method. In future work, we intend to use these results as a guide to develop *a priori* time step controls for more complex radiative transfer problems.

*For more information contact Scott Mosher at [swmosher@lanl.gov](mailto:swmosher@lanl.gov).*

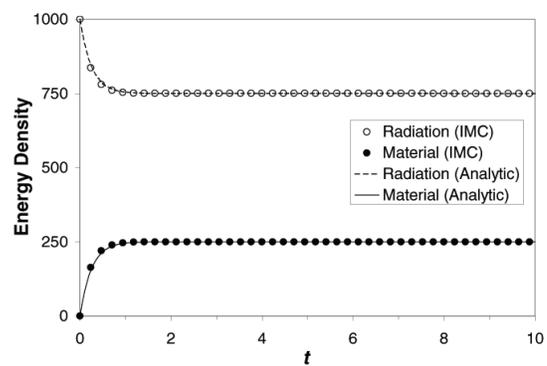
- [1] J.A. Fleck, Jr. and J.D. Cummings, *J. Comp. Phys.* **8**, 313 (1971).
- [2] W.R. Martin and F.B. Brown, *Trans. Am. Nucl. Soc.* **85**, 329 (2002).
- [3] S.W. Mosher and J.D. Densmore, *Trans. Am. Nucl. Soc.* **93**, 520 (2005).
- [4] S.W. Mosher, *Trans. Am. Nucl. Soc.* **95**, 744 (2006).

### Funding Acknowledgements

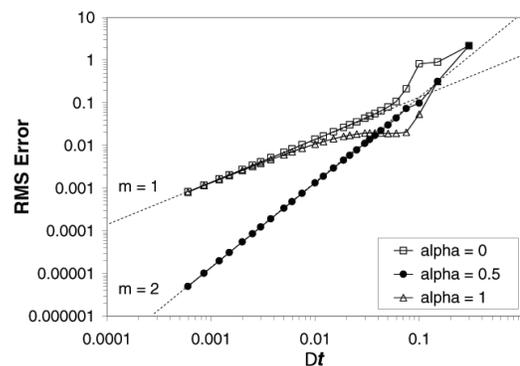
NNSA's Advanced Simulation and Computing (ASC), Computational Physics and Methods Mission Capability Program, Transport Project.



**Fig. 1.** IMC energy densities for a time step at half of the stability limit.



**Fig. 2.** IMC energy densities for a time step at half of the monotonicity limit.



**Fig. 3.** IMC material temperature RMS error.

## Verification and Validation of Radiation Transport Schemes

Chris Fryer, Aimee Hungerford, Tom Evans, Mike Buksas, Todd Urbatsch, CCS-2;  
Darrell Peterson, Randy Kanzleiter, X-2

**A**s part of the Advanced Simulation and Computing (ASC) Verification and Validation program, we have conducted a number of tests on radiation transport schemes in ASC codes. These tests include a range of problems and a variety of algorithms. In this brief summary, we discuss the methodology behind these studies using one specific code test as an example.

Verification of a code can be done in a number of ways: convergence studies (in space, in time, in tolerance settings, etc.) and comparisons of results (e.g., with an experiment or an analytic result or between two different codes or algorithms). Note that comparison with an experiment may also serve to “validate” a code, but not all code-to-experiment comparisons are validation tests [1,2]. We utilize many of these techniques in our tests.

The example in this paper illustrates three verification techniques: code/algorithm comparison, code/experiment comparison, and a convergence study. The experimental setup takes the radiation released by a Sandia Z dynamic Hohlraum source and studies the propagation of the radiation through a target. The radiation flows down a taper into an aerogel foam. The resultant shock is then backlit to provide an image. The first suite of these experiments has already been performed, providing an ideal set of results for comparison.

Figure 1 shows the density contour from our calculations for this experiment at a set time after the injection of the

radiation. The top panel shows a comparison between a certain set of input parameters (right half) versus the latest and greatest set of input parameters (left half) for our code. In a sense, we have just done a code-to-code comparison in this top panel even though the actual code used was the same, but the exact numerical implementation within this code changed. We also compared these results to the experiment and found agreement at the 10–20% level, within the error of the experimental setup. The bottom panel of Fig. 1 shows a comparison of the calculation results for simulations using two different resolutions, one in 2-D, and the other in 3-D. Again, the good agreement implies both that we are converged and that 3-D effects are not too critical for this particular problem setup.

Our simulations are part of a recently completed Level 2 ASC milestone, and our future simulations will continue to play a role in major verification and validation milestones.

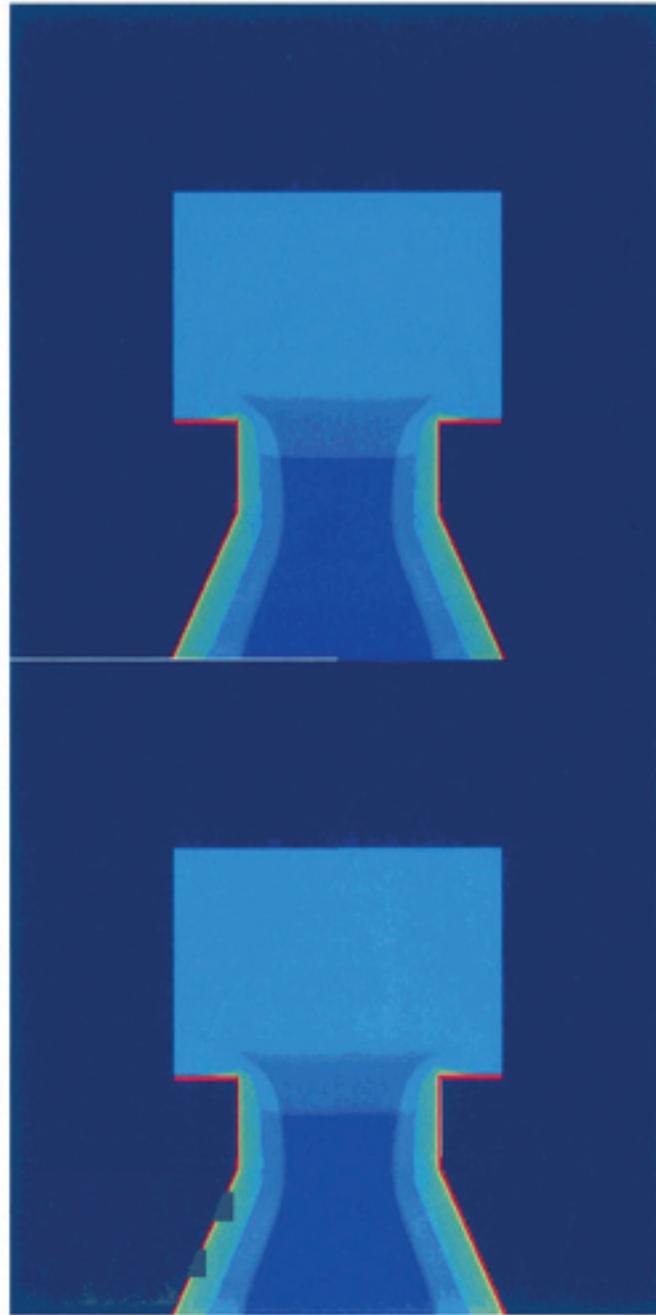
*For more information contact Chris Fryer at [fryer@lanl.gov](mailto:fryer@lanl.gov).*

[1] C.L. Fryer, et al., *Ap. J.* **643**, 292 (2006).

[2] C.L. Fryer, et al., “Supernova Explosions: Understanding Mixing” (submitted to *Int. J. Mod. Phys. D*).

### Funding Acknowledgements

NNSA’s Advanced Simulation and Computing (ASC), Secondary Verification and Validation and Focused Research on Capabilities.



**Fig. 1.**  
*Plots of density for four separate calculations of our experimental setup. Top left: parameters from latest release from milestone project. Top right: old simulation. Bottom left: normal resolution. Bottom right: low resolution.*

## Student Projects in Radiation Transport Methods, Summer 2006

Todd Urbatsch, Ed., CCS-2

**T**he Transport Methods Group, CCS-4, has had a strong summer student program, typically hosting about one student per two to three technical staff members. The CCS-4 mentors have fostered productive and enlightening summers with well-defined work plans. These work plans are designed to engage the student in novel research that is publishable and will positively impact the science behind our programmatic efforts. Prior to the formation of the Computer, Computational, and Statistical Sciences (CCS) Division in October 2000, the Transport Method Group was X-6, which was formed from an X-6/T-1 merger in the early 1980s. In January 2007, CCS-4 merged with CCS-2 (Continuum Dynamics) to form the Computational Physics and Methods Group. To commemorate the final Transport Methods Group summer student class, we briefly summarize the projects from the summer of 2006.

**“Neutrino Streaming Under Weak Gravitational Fields,”** Alex Maslowski, Graduate Research Assistant (GRA), Texas A&M University, DOE Nuclear Engineering Health Physics Fellow  
*Mentor: Kent Budge*

Neutrino particles act as the main energy transport agents in supernovae. The influence of gravitational fields on neutrino transport is nontrivial, and yet its effect remains uncertain. Thus, we derived a cylindrical coordinate streaming operator that accounts for the effects of a weak gravitational field. A second-order (post-Newtonian) approximation to the General Relativity induced curvature was applied. Based on the post-Newtonian metric, the Christoffel symbols were computed and the corresponding geodesic particle trajectories were solved. Using the Lindquist approach for nonorthonormal coordinate systems, we derived a post-Newtonian-order Boltzmann operator for massless particles in cylindrical

coordinates. We also derived the respective energy gravitational shift. The energy gravitational shift was implemented into the Capsaicin Project’s Macho neutrino transport software package.

**“Implementing and Testing the Piecewise Linear Finite Element Method in Capsaicin,”** Teresa Bailey, GRA, Texas A&M University, DOE Computational Science Fellow

*Mentor: Kelly Thompson (currently with Transpire, Inc.)*

We have derived Stone and Adams’ Piecewise Linear Discontinuous Finite Element Method (PWL) for the 2-D RZ transport equation. We implemented and tested it in the Capsaicin Project software. The PWL method is uniquely suited to discretize differential equations on arbitrary polyhedral meshes, whereas methods such as the bilinear discontinuous finite element method (BLD) has been applied only to triangular and quadrilateral grids. We used an asymptotic analysis to determine successful PWL lumping schemes. The convergence rate of PWL on quadrilateral grids for both the transport and diffusive regimes is similar to BLD.

**“Krylov Acceleration for Transport in Binary Statistical Media,”** Erin D. Fichtl, GRA, University of New Mexico

*Mentor: Jim Warsa*

Radiation transport methods for stochastic mixtures using the Levermore, Pomraning, Vanderhaegen coupled equations are slow to converge as one or both materials approach the diffusion and/or atomic mix limits. We propose a three-part acceleration scheme to accelerate convergence. Firstly, we divide the iteration into a series of “inner” material and species iterations to attenuate the atomic mix and diffusion error modes separately. Secondly, we apply atomic mix synthetic acceleration to the inner material iteration and  $S_2$  synthetic

acceleration to the inner species iteration. Finally, we wrap a Krylov iterative solver around each inner and outer iteration. Utilizing a spectral analysis, we compare the effectiveness and efficiency of our new two-step scheme to the basic one-step scheme.

**“Interface Reconstruction in PARTISN,”**  
Michael Reed, GRA, Texas A&M University  
*Mentor: Randy Baker*

Orthogonal meshes greatly simplify numerical discretizations, but they often introduce significant meshing errors into the problem. Traditionally, the mesh is refined until such meshing errors are acceptable, but this can be prohibitively expensive in 2-D and 3-D. Alternatively, we reconstruct material interfaces at the subcell level and consider such interfaces in the governing equations. We have derived and implemented a diamond difference interface reconstruction technique for the discrete ordinates code PARTISN in multidimensional Cartesian and curvilinear geometries. The reduction in error has been shown to be equivalent to half an order of mesh refinement.

**“Fourier Analysis of Parallel, Inexact Block-Jacobi Splitting in 2-D Geometry,”**  
Massimiliano Rosa, GRA, Pennsylvania State University  
*Mentor: Jae Chang*

A Fourier analysis is proposed for the steady state one-group transport problem solved with Richardson Iteration (Source Iteration) using the parallel block-Jacobi (PBJ) algorithm, Bilinear Discontinuous Finite Element Method (BLDFEM) discretization in 2-D XY geometry, and assuming isotropic scattering. Our novel Fourier analysis considers four computational cells and the interfaces of four adjacent computational subdomains and has proven to be effective in capturing the peculiar features of the PBJ algorithm. The results agree with observed behavior in the Capsaicin Project’s 2-D code, Anaheim. Convergence of the PBJ algorithm can degrade and lead to stagnation of GMRES for thin problems. However, our analysis shows that Modified Transport Synthetic Acceleration is an effective preconditioner in the optically thin case.

**“Cell Global-Local Discontinuous Finite Element Spatial Discretizations for Radiation Transport,”** Rick Gleicher, GRA, University of New Mexico  
*Mentor: Jim Warsa*

We propose the Cell Global-Local Discontinuous Finite Element (CGOLD) method for radiation transport on highly unstructured meshes. During a standard discrete ordinates transport sweep across spatial cells, each cell is divided into subelements. Each subelement interpolates the initial scalar flux from the cell and then the cell’s subelements are swept in the current discrete ordinate direction. The subelement’s angular fluxes are obtained and projected back onto the cell. The transport sweep continues and new cell scalar fluxes, obtained from the accumulation of projected angular fluxes, are iterated to convergence. In 1-D, the CGOLD method is shown to be third-order accurate and to have the diffusion limit.

**“Extending Quadruple Range Quadrature,”** Eric Baker, Undergraduate Student (UGS), Oregon State University  
*Mentor: Jon Dahl*

The quadruple-range and octant-range quadratures developed by I.K. Abu-Shumays, require fewer angular directions to achieve the same degree of accuracy as other quadratures. However, their maximum attainable accuracy has been limited by the extent of the angular discretization with which they have been computed. We have developed the capability to produce the quadruple-range and octant-range quadratures and, in so doing, we have verified the previously published values.

*For more information contact Todd Urbatsch at [tmonster@lanl.gov](mailto:tmonster@lanl.gov).*

**Funding Acknowledgements**  
The Advanced Simulation and Computing (ASC), Integrated Codes Program, Transport Project; and DOE Fellowships.

# Fourier Analysis of Parallel Block-Jacobi Splitting with Transport Synthetic Acceleration in Two-Dimensional Geometry

Massimiliano Rosa, James S. Warsa, Jae H. Chang, CCS-2

**F**ourier analysis is traditionally used to study transport iteration schemes in a homogeneous infinite medium. In fact, it represents a valuable tool to understand the behavior of the iteration error modes of various acceleration techniques, either in their continuous or spatially discretized forms. A Fourier analysis has been recently performed in slab geometry [1] for the discrete-ordinates ( $S_N$ ) approximation of the steady-state one-group transport problem solved with Richardson iteration and preconditioned Richardson iteration, using the Parallel Block-Jacobi (PBJ) algorithm. Two types of Transport Synthetic Acceleration (TSA) have been considered as a preconditioner, the traditional "Beta" TSA (TTSA) and a modified TSA (MTSA).

The Fourier analysis for PBJ and for PBJ with TTSA and MTSA is extended to 2-D Cartesian geometry. The spatial discretization considered is a Bilinear Discontinuous Finite Element Method (BLDFEM) and the scattering is assumed to be isotropic. The analysis is verified with results from a 2-D transport code. As stated in [1], the results of the Fourier analysis reveal the effectiveness of iteration algorithms in terms of spectral radius  $\rho$  and minimum eigenvalue  $\Lambda$  of the symmetric part of the iteration matrix. The latter are functions of the quadrature order (number of angles), total macroscopic cross section  $\sigma$  and scattering ratio  $c$ .

First we look at the results of the Fourier analysis for PBJ without any preconditioning using  $S_4$  level-symmetric quadrature,  $\sigma = 1$  and  $c = 0.5$ . Results show that convergence of PBJ algorithm can degrade for problems containing optically thin subdomains, even for values of the scattering ratio  $c$  less than unity. In fact the spectral radius, as the cell widths are decreased, is independent from the value of  $c$ . Furthermore, the minimum eigenvalue

$\Lambda$  of the symmetric part of the iteration matrix may become negative indicating that restarted GMRES, GMRES(m), may stagnate for optically thin problems.

The predictions from the Fourier analysis for PBJ have been compared with the numerical results obtained from a 2-D transport code. To reproduce the conditions of the Fourier analysis, the 2-D transport code ran on four processors with four computational cells per processor ( $N_x = N_y = 2$ ) and with reflective boundary conditions on all four boundaries, in order to eliminate the effect of leakage. The results for  $\rho$  and  $\Lambda$  obtained for a level-symmetric  $S_4$  quadrature with equal weights, a unitary total macroscopic cross-section, and a scattering ratio of 0.5 are compared with the Fourier analysis. In both cases the numerical results are very close to the predicted theoretical values.

Next we examine the number of iterations and the residual for numerical calculation using GMRES(m) with different values of the restart parameter. The maximum number of iterations is set to 1000 and the error tolerance is  $10^{-5}$ . The results show that GMRES(m) may stagnate for a value of the restart parameter equal to the square root of the number of processors, which is number of steps needed to propagate the information across the full problem domain. This finding is in agreement with negative minimum eigenvalues obtained in the Fourier analysis.

The results above indicate that preconditioning of the PBJ algorithm is needed to improve its spectral properties, especially for optically thin problems. Therefore we conducted a Fourier analysis of PBJ accelerated using both traditional TSA (TTSA) and modified TSA (MTSA) as preconditioners in 2-D Cartesian geometry. Since the effective scattering ratio  $c'$  is a measure of the actual computational effort required by each acceleration method, the

comparison of the two methods is carried out for equal values of this parameter. In particular, the results presented refer to the case  $c' = 0$  ( $\beta = 1$ ), which is an upper bound for the spectral radius.

The surface plots are obtained by plotting the spectral radius  $\rho$  from the Fourier analysis as a function of the cell widths, Fig. 1 for TTSA and Fig. 2 for MTSA, respectively. TTSA appears extremely effective for thick problems, for which the spectral radius tends to zero. However, the spectral properties degrade for thin problems. Also, for values of the scattering ratio  $c$  greater than 0.5, the spectral radius may become greater than unity leading to instability, unless the  $\beta$  parameter is properly selected. MTSA is slightly less effective than TTSA for thick problems, where the spectral radius is slightly greater than zero. The spectral properties are better for thin problems. As a matter of fact, the spectral radius is always bounded from above by the value of the scattering ratio  $c$ , for  $\beta = 1$ . MTSA is therefore unconditionally stable.

A Fourier analysis has been implemented for the Parallel Block-Jacobi (PBJ) algorithm and for PBJ with both traditional TSA (TTSA) and modified TSA (MTSA). The results for the unaccelerated algorithm show that convergence of PBJ can degrade and lead to stagnation of GMRES(m) in problems with optically thin subdomains. These predictions have in turn been successfully verified against a 2-D transport code.

Fourier analysis for the accelerated PBJ algorithm indicates that PBJ with TTSA can be effective provided the  $\beta$  parameter is properly tuned for a given scattering ratio  $c$ , but is potentially unstable. Compared to TTSA, MTSA is less sensitive to the choice of  $\beta$ , more effective for the same computational effort ( $c'$ ), and it is unconditionally stable.

For more information contact Jae H. Chang at [jhchang@lanl.gov](mailto:jhchang@lanl.gov).

[1] M. Rosa, et al., "Fourier Analysis of Parallel Inexact Block-Jacobi Splitting with Transport Synthetic Acceleration in Slab Geometry," in *Proceedings of the American Nuclear Society's Topical Meeting on Reactor Physics: PHYSOR-2006*, Vancouver, BC, Canada, September 10-14, 2006 (American Nuclear Society, 2006).

### Funding Acknowledgements

NNSA's Advanced Simulation and Computing (ASC), Computational Physics and Methods Strategic Capability, Advanced Transport Discretizations, and Iterative Methods Project.

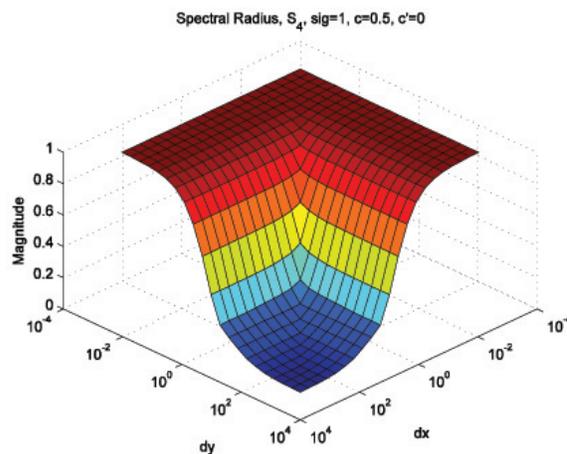


Fig. 1. Fourier analysis for PBJ with TTSA:  $\rho$ .

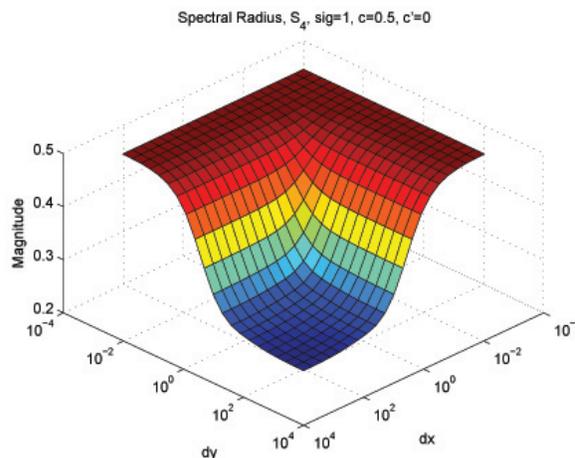
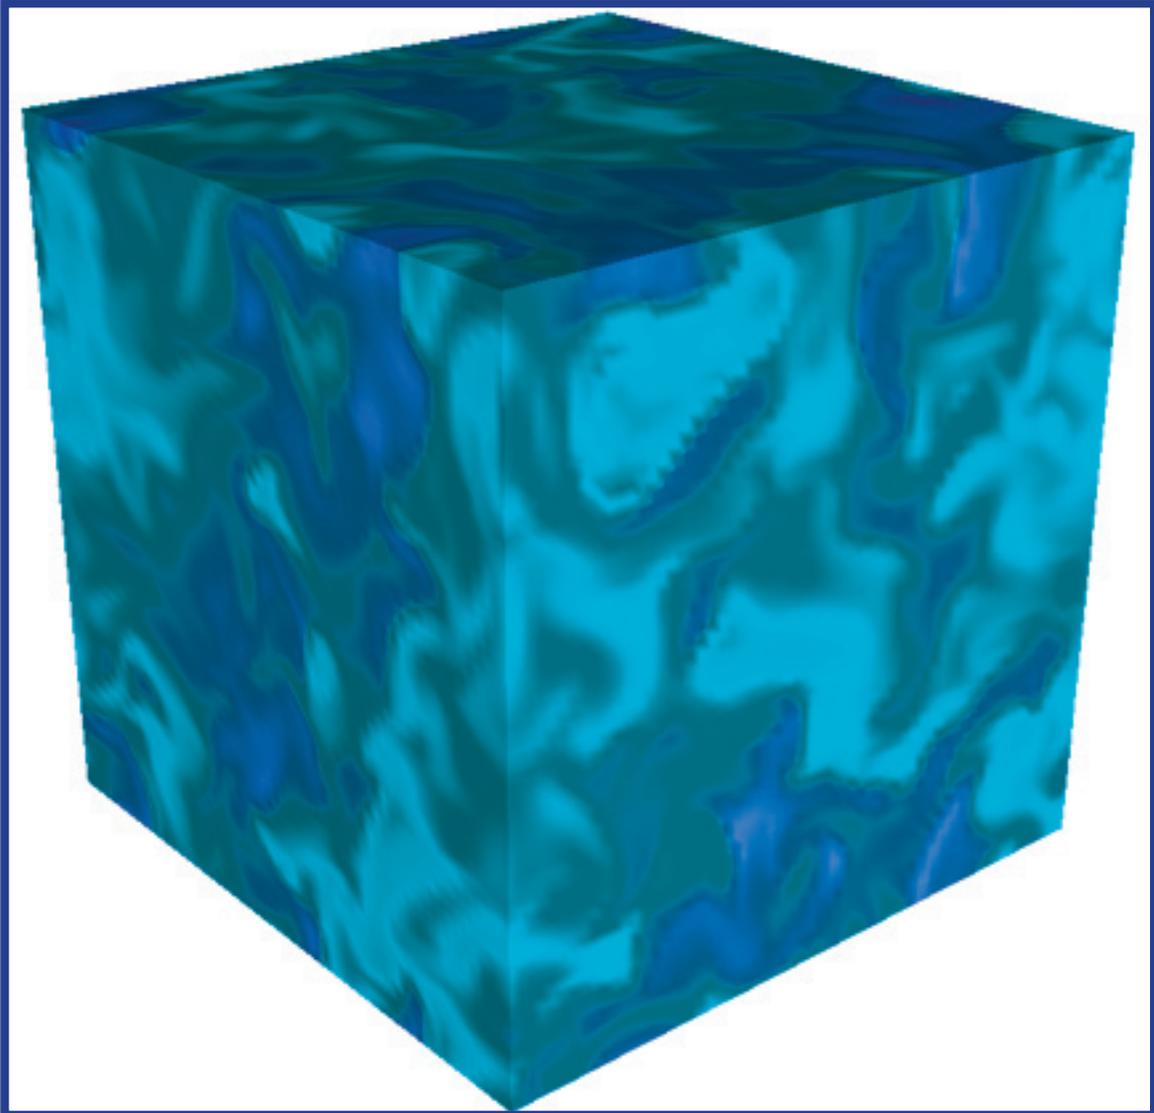


Fig. 2. Fourier analysis for PBJ with MTSA:  $\rho$ .



# Coupled Physics and Multiscale Methods and Algorithms

The fundamental processes in the evolution of the explosion of a nuclear weapon (of which hydrodynamics and radiation transport are two) cannot evolve independently. They are complex, nonlinear, coupled process that must be studied and understood holistically. The articles in this section address the effects and importance of this coupling in the areas of turbulence, the study of internal confinement fusion processes, the effect of multiple scales in the simulation of coupled phenomena, the examination of unmodelled differences in identical experiments, and the effect of the coupling of surface tension with other material properties in modeling low-speed flows.

## Buoyancy-Driven, Variable Density Turbulence

*Daniel Livescu, CCS-2; J. R. Ristorcelli, X-3*

**M**ixing to molecular scale in the presence of turbulence-induced stirring is an important process in many practical applications. In general, the fluids participating in the mixing have different densities and we refer to such flows as variable density (VD) flows. In these flows, the specific volume changes in both time and space depending on the amount of each fluid in the mixture and the resulting velocity field is not divergence-free even for constant density fluids. Variable density mixing is encountered in atmospheric and ocean flows, combustion and many flows of chemical engineering interest, astrophysical flows, etc.

Here we consider a simple form of VD multi-material mixing which involves two miscible fluids with different microscopic densities [1,2], in the presence of a constant acceleration, as occurs in the Rayleigh-Taylor (RT) instability. The current investigation focuses on the nonlinear dynamics and statistics of buoyantly driven turbulence in the statistically homogeneous configuration. As such the new nonlinearities due to large density variations in the advective terms of the Navier-Stokes equations become important. The problem is an extension of the buoyantly generated turbulence in a Boussinesq fluid studied in [3] and in a VD fluid examined in [4]. From the modeling viewpoint, this is a benchmark problem that any turbulence model for VD-RT should handily predict.

In VD turbulence with arbitrary boundary conditions, the two first-order moments, the mean pressure gradient, and the mean specific volume are dynamical variables evolving as the mixing proceeds. For periodic boundary conditions though, the mean pressure gradient can be determined up to a constant gradient, which is a free parameter. This is chosen such that the energy conversion of potential to kinetic energy is maximized (see [2]), ensuring a benchmark flow that is the maximally unstable flow in this configuration.

The flow starts with zero solenoidal velocity in a nonpremixed state and turbulence is generated due to the baroclinic production of vorticity and eventually dies as the two fluids become molecularly mixed. The problem involves both the transition to turbulence and the decay of turbulence as the friction forces overcome buoyancy generation. Results from Direct Numerical Simulations are used to follow the turbulence birth-life-death process and examine the influence of various parameters, Atwood, Reynolds, and Schmidt numbers, and initial length scale of the density field on the mixing. The cases considered cover the Atwood number range  $At = 0.05-0.75$ , in order to examine small departures from the Boussinesq approximation as well as large Atwood number effects. Simulations with resolutions up to  $1024^3$  are performed. The numerical methodology employs a new pressure projection algorithm that treats the pressure step exactly and represents

an improvement over the methods presently used for variable density flows [4,5].

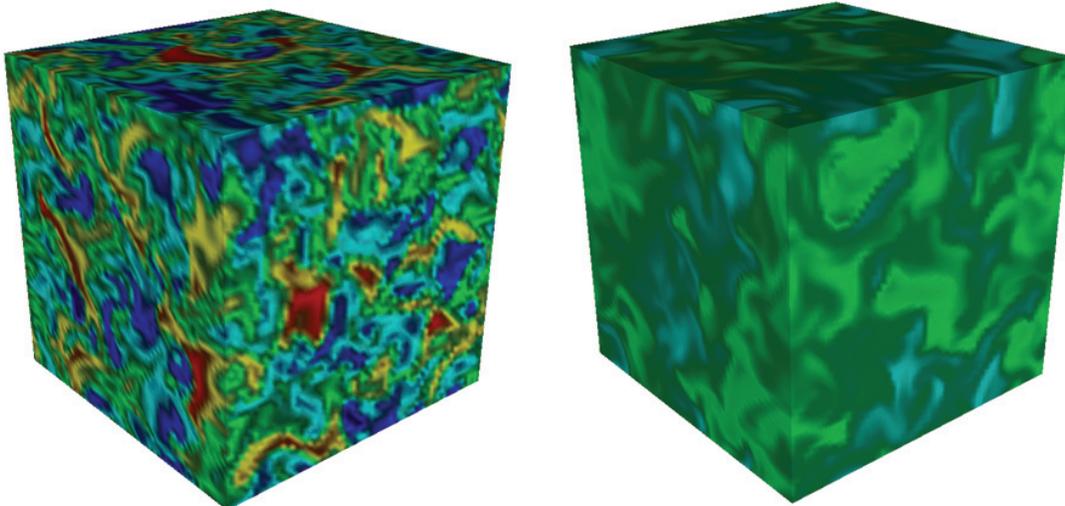
The study shows that the rate of conversion of potential energy into kinetic energy as well as between Favre mean and turbulent kinetic energies are mediated by the mass flux so that the mass flux is likely the most important quantity to predict in lower dimensional models [2,6]. Numerical and analytical approaches are also used to examine the morphology of the scalar structures and the influence of various parameters on the mixing progress. In particular, it is shown that the specific volume density covariance is a better measure of the flow evolution than the density variance for variable density flows, as it directly appears in the dynamical equations. Nevertheless, the normalized density variance can be related to the variance of the excess reactant in an infinitely fast hypothetical chemical reaction. Under special circumstances this represents a metric for the progress of the fully mixed fluid. Neither measure though can express the amount of pure fluid in the flow, which is related to the tails of the density probability density functions (PDF) [7,8]. A model for the density PDF transport equation is proposed to predict the evolution of the pure and mixed fluids.

**For more information contact Daniel Livescu at [livescu@lanl.gov](mailto:livescu@lanl.gov).**

- [1] D. Besnard, et al., "Turbulence Transport Equations for Variable-Density Turbulence and Their Relationships to Two Field Models," Los Alamos National Laboratory report, LA-12303-MS (1992).
- [2] D. Livescu and J.R. Ristorcelli, "Buoyancy-Driven Variable Density Turbulence," submitted to *J. Fluid Mech.*
- [3] G.C. Batchelor, et al., *J. Fluid Mech.* **235**, 349–378 (1992).
- [4] S. Sandoval, Ph.D. thesis, University of Washington (1995).
- [5] A.W. Cook and P.E. Dimotakis, *J. Fluid Mech.* **343**, 69–99 (2001).
- [6] D. Livescu and J.R. Ristorcelli, "Characteristics of Buoyancy-Driven Variable Density Turbulence," in Proceedings of the 10th International Workshop on the Physics of Compressible Turbulent Mixing.
- [7] D. Livescu and J.R. Ristorcelli, "Mixing Characteristics in Buoyancy-Driven Variable Density Turbulence," to be presented at the 11th European Turbulence Conference, Porto, Portugal, June 25–28, 2007.
- [8] D. Livescu and J.R. Ristorcelli, "Mixing Characteristics in Buoyancy-Driven Variable Density Turbulence," submitted to *J. Fluid Mech.*

#### **Funding Acknowledgements**

NNSA's Advanced Simulation and Computing (ASC), Instabilities and Turbulent Mixing Project; Campaign 4, Secondary Assessment Technology.



**Fig. 1.** Typical density field at maximum density variance (left) and late time (right).

## Detailed Experimental Measurements of Turbulence Model Constants

Arindam Banerjee, Malcolm J. Andrews, CCS-2

**E**xperimental data at small Atwood numbers ( $A_t = 0.04$ ) from the Rayleigh-Taylor (RT) mix experiment has been used to perform model constant evaluations. As a key hydrodynamic process during inertial confinement fusion (ICF) implosion, our studies of RT mixing directly impact fundamental understanding of the flow physics, and aid in validation of predictive engineering models for ICF target design and energy deposition. We consider the primary momentum transport flux  $\langle \rho'v' \rangle$ , for which a model representation is to use the gradient diffusion hypothesis [1] as:

$$\langle \rho'v' \rangle = -\frac{\nu_t}{\sigma_t} \frac{\partial \rho}{\partial y} = -\frac{C_\mu k^{1/2} \ell_m}{\sigma_t} \frac{\partial \rho}{\partial y} \quad (1)$$

where  $\nu_t$  is the turbulent viscosity and the active scalar turbulent Prandtl number,  $\sigma_\mu$  takes a value of 0.7 [2]. The primary transport in small  $A_t$  experiments is associated with the largest structures and these are of size  $h_b (= h_s$  at small  $A_t$ ), thus we take the integral length scale  $\ell_m$  as  $h_b$ . Figure 1 shows that the measured density profile is reasonably linear across the mix and thus, the density gradient  $\partial \rho / \partial y = \Delta \rho / 2h_b$  (the factor 2 arises because  $h_b$  is the half mixing width). Our velocity measurements [3] at the centerline indicate that  $u'^2 = w'^2 = 0.3v'^2$  so the turbulence kinetic energy ( $k = 0.5 \langle u'^2 + v'^2 + w'^2 \rangle$ ) is given by  $k = 0.8v'^2$ . Figure 2 plots the collapse of the centerline RMS vertical velocity ( $v'$ )

to obtain a measured centerline  $\alpha_{CL}$  of 0.07. Since  $v'$  can be related to the mix width by [4]:

$$v' = \frac{dh_b}{dt} = 2\alpha_{CL} A_t g t = 2\alpha_{CL} A_t g \frac{x}{U_m} \quad (2)$$

the centerline turbulence kinetic energy is given by  $k = 0.8(2\alpha_{CL} A_t g x / U)^2$ . Substitution into (1) gives the following expression for the measurement of  $C_\mu$ :

$$C_\mu = -\frac{\langle \rho'v' \rangle \sigma_t}{\Delta \rho \sqrt{0.8} \left( A_t g \frac{x}{U} \right) \alpha_{CL}} \quad (3)$$

Figure 3 plots measured values of  $\rho'v' / (\Delta \rho (A_t g x / U))$  that at late time

$$\tau = \frac{x}{U_m} \left( \frac{A_t g}{H} \right)^{1/2} = 1.986 \text{ reach a value of}$$

-0.024, giving a corresponding value for  $C_\mu$  of 0.288. A typical value

quoted in the literature for  $C_\mu$  is 0.09 [5]; however, this lower value

corresponds to turbulent shear flows and flow situations where the rate of

turbulence kinetic energy production and dissipation are balanced [1]. The

present buoyancy-driven flow has a global  $PE_{released} / Dissipation$  of  $\sim 2.0$

which is far from a balanced case.

Measured  $C_\mu$  at the centerline for the different downstream locations are

given in Table 1. It is seen that the value of  $C_\mu$  remains approximately constant

at various downstream locations and is independent of the Reynolds number of

the developing flow.

**TABLE 1.** Turbulence modeling constant ( $C_\mu$ ) at different times ( $\tau$ ).

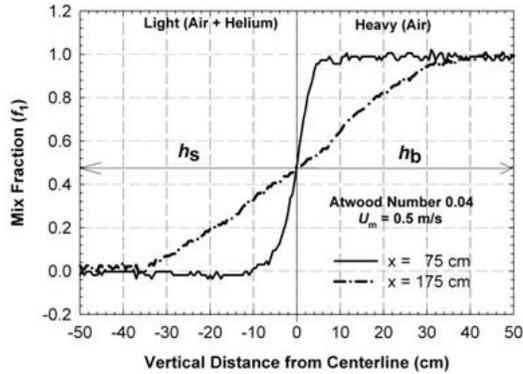
$\tau$	$\frac{\langle \rho'v' \rangle}{\Delta\rho(Agt)}$	$\alpha_{CL}$	$C_\mu$
0.567	-0.0482	0.12169	0.3098
0.851	-0.03812	0.09268	0.3219
1.135	0.02878	0.07709	0.2922
1.418	-0.02634	0.06743	0.3057
1.702	-0.02444	0.06741	0.2838
1.986	-0.02397	0.06520	0.2877
2.213	-0.02363	0.06583	0.2809

For more information contact Malcolm Andrews at [mandrews@lanl.gov](mailto:mandrews@lanl.gov).

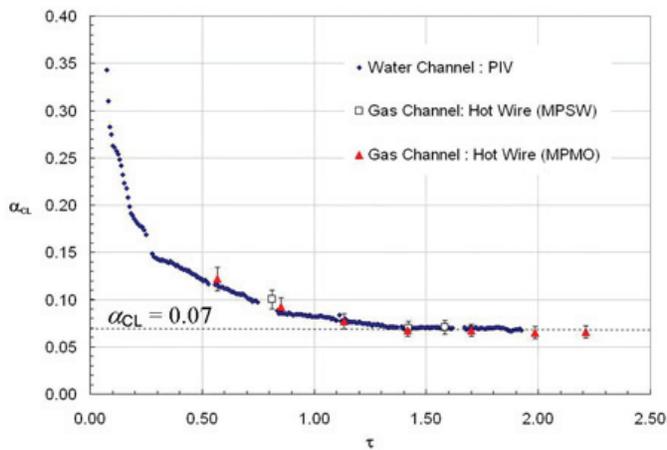
- [1] S.B. Pope, *Turbulent Flows* (Cambridge University Press, 2000).
- [2] D.M. Snider, M.J. Andrews, *Phys. Fluids* **6**, 3324–3334 (1994).
- [3] A. Banerjee, M.J. Andrews, *Phys. Fluids* **18**, 035107 (2006).
- [4] P. Ramaprabhu, M.J. Andrews, *J. Fluid Mech.* **502**, 233–271 (2004).
- [5] B.E. Launder, D.B. Spalding, *Comp. Meth. in App. Mech. & Eng.* **3**, 269–289 (1974).

**Funding**

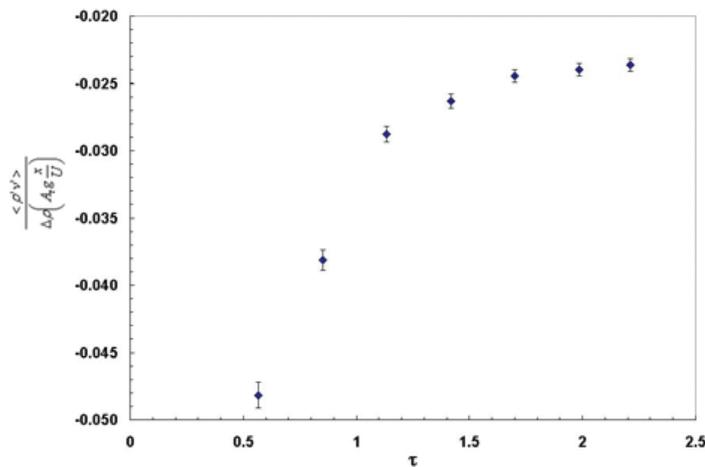
**Acknowledgements**  
 NNSA's Campaign 2,  
 Dynamic Materials  
 Properties.



**Fig. 1.** Mixture fraction profiles from the experimental run at small  $A_t$  ( $= 0.04$ ).



**Fig. 2.** Centerline measurements of velocity RMS at small  $A_t$  ( $= 0.04$ ).



**Fig. 3.** Evolution of  $\langle \rho'v' \rangle$  (primary transport term) for  $A_t = 0.04$ .

## How Much Laser Power Can Propagate through Fusion Plasma?

Pavel M. Lushnikov, *Landau Institute and University of New Mexico*; Harvey A. Rose, *T-13*

**P**ropagation of intense laser beams in plasma raises outstanding technological and scientific issues. These issues are closely tied with inertial confinement fusion (ICF), which requires precise beam control in order to maintain symmetry of spherical target implosion, and so achieve the compression and heating necessary to ignite the fusion reaction. The National Ignition Facility (NIF), where ICF will be attempted, is now under construction. While most engineering features of NIF are now fixed, there are still crucial choices to be made in target designs. Control of intense beam propagation may be ruined by laser beam self-focusing, when a beam digs a cavity in plasma, trapping itself, leading to higher beam intensity, a deeper cavity, and so on.

In laser fusion, the intensity of laser beams is so large that self-focusing can cause disintegration of a laser beam into many small beams, leading to rapid change in beam angular divergence, called beam spray. (See Fig. 1.) Significant beam spray is absolutely unacceptable for attaining fusion since it requires precise laser beam control. It was commonly assumed that the main source of beam spray in fusion plasma is the self-focusing of local maxima of laser intensity (hot spots) randomly distributed throughout the plasma. Although hot-spot self-focusing can be controlled by

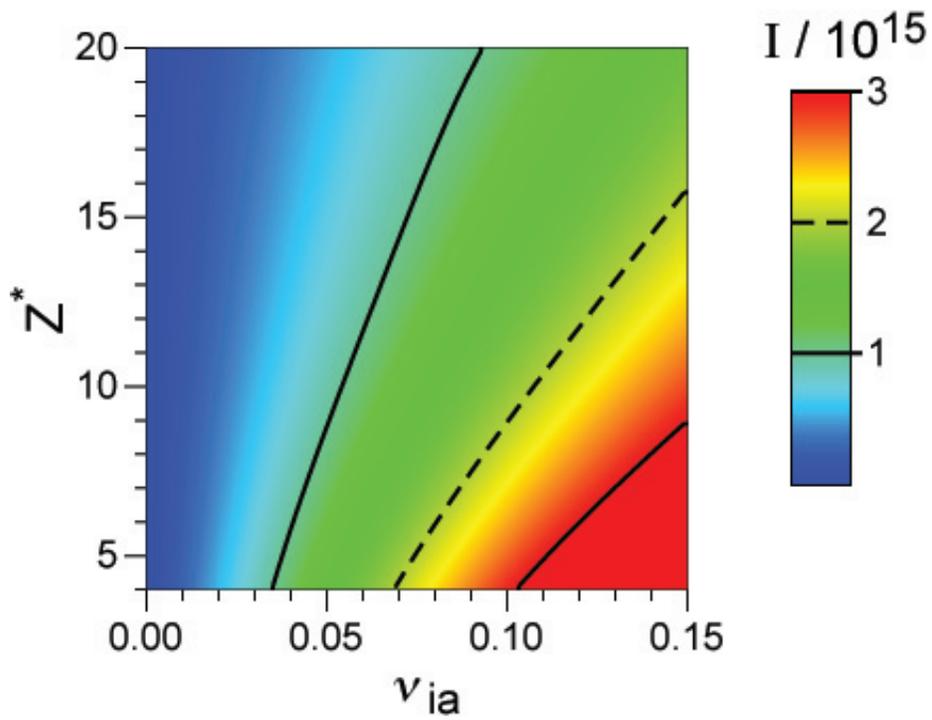
reducing beam correlation time, we find that the main limitation of maximum beam power, which can propagate in plasma without significant beam spray, is determined by collective instability that couples the beam to an ion acoustic wave. We call this instability collective-forward-stimulated Brillouin scatter (CFSBS) because it does not depend on the dynamics of isolated hot spots, but rather the intensity fluctuations as temporally smoothed (averaged) by ion inertia. This collective instability is consistent with the first experimental observation of the beam spray onset while hot-spot self-focusing is not. Our theory [1] suggests how to increase that maximum by appropriate choice of plasma composition that affects damping and thermal transport, with implication for NIF designs.

*For more information contact Harvey Rose at [har@lanl.gov](mailto:har@lanl.gov).*

[1] Pavel M. Lushnikov and Harvey A. Rose, *Plasma Phys. Control. Fusion* **48**, 1501 (2006).

### Funding Acknowledgements

NNSA's Advanced Simulation and Computing (ASC), Laser Matter Interaction.



**Fig. 1.** Predicted laser intensity,  $I$ -in units of  $W/cm^2$ , at onset of beam spray regime for NIF parameters, as a function of dimensionless acoustic damping coefficient,  $\nu_{ia}$  and effective plasma ionization,  $Z^*$ .

# Self-Similarity, Scale Separation, and Universality in Rayleigh-Taylor, Boussinesq Turbulence

Natalia Vladimirova, University of Chicago; Michael Chertkov, T-13

The Rayleigh-Taylor (RT) instability occurs when a heavy fluid is pushed by a light fluid. Two plane-parallel layers of fluid, the heavier on top, are in equilibrium; with the slightest perturbation the heavier fluid moves down under the gravitational field, and the lighter fluid rises. The development of the instability leads to enhanced mixing and growth of the mixing zone. Dimensional arguments, supported by large-scale modeling [1], suggest that the half-width of the mixing zone,  $h$ , grows quadratically at late time,  $h \propto \alpha A g t^2$ , where  $A$  is the Atwood number characterizing the initial density contrast and  $g$  is the gravitational acceleration;  $\alpha$  is a dimensionless coefficient, studies of which have dominated the RT turbulence literature for the last 50 years [2].

We focus here on analysis of the mixing zone internal structure rather than  $\alpha$ . Available phenomenology [3] can be summarized in the following

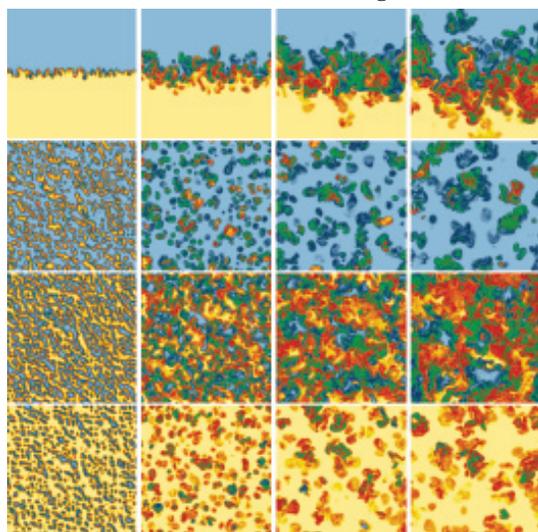
statements: (i) the small-scale turbulence in the mixing zone is adiabatically adjusted to the large-scale buoyancy-controlled dynamics; (ii) in three dimensions, velocity fluctuations at smaller scales are asymptotically decoupled from weaker buoyancy effects; and (iii) the spatial scalings of the velocity and density fluctuations are similar to the scalings in homogeneous Kolmogorov turbulence. Also notice a couple of important deficiencies of the phenomenological approach: (A) it treats all  $z$ -slices within the mixing zone equally; (B) it does not differentiate between the mixing zone width,  $h$ , and the energy-containing scale,  $R_0$ , for the turbulent fluctuations.

Experimentally [4] and numerically [5,6] obtained energy spectra, averaged over a snapshot of the mixing zone, are all consistent with predictions of Ref. [3]. One important consequence of the phenomenology is a decrease of the viscous and dissipative scales,  $\eta$  and  $r_d$  respectively, also predicted in [7] and numerically observed in [6,7].

Here we address questions raised by the phenomenology through simulations from the spectral element code developed by Tufo and Fischer [8]. We consider 3-D incompressible RT flow in the Boussinesq regime with finite viscosity and dissipation (Fig. 1). The two fluids are allowed to mix through diffusion.

- 1) Can the relative dependence of scales be a more reliable indicator of universal behavior than the time-dependence of each individual scale?

**Fig. 1.** Slices of temperature at times  $t = 32, 64, 96, 128$  (left to right). From top to bottom, the images correspond to vertical slices at  $y = 480$  and horizontal slices at  $z = +0.75h, 0, -0.75h$ .



We found that  $R_0$ ,  $\eta$ , and  $r_d$  exhibit monotonic evolution with  $z/h$  (Fig. 2). This scaling behavior is consistent with predictions of Ref. [3].

- 2) How does the energy containing scale,  $R_0$ , compare with  $h$ ? We found that at late time, the ratio of  $R_0$  to  $h$  taken at the center of the mixing zone is  $\approx 1:20$  (Fig. 3).
- 3) Do the turbulent spectra vary as a function of vertical position in the mixing zone? We did not observe any qualitatively new behavior in spectra at any off-centered slices in comparison with the central slice (Fig. 4).
- 4) How different are the scales and spectra corresponding to qualitatively different initial perturbations? We found that the dependence of  $R_0$ ,  $\eta$ , and  $r_d$  vs  $z/h$  on the initial perturbations is weak.

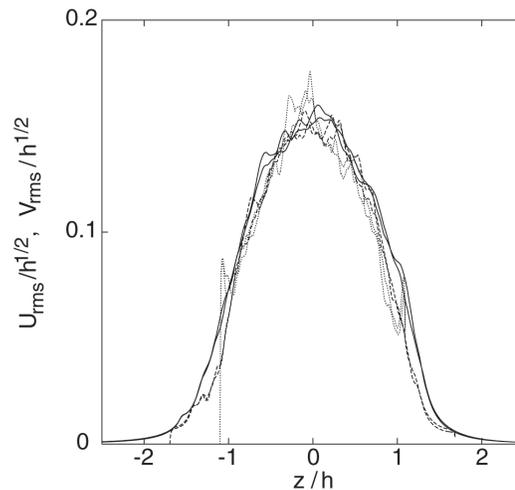
We plan to extend this analysis to account for effects of chemical reactions on the RT turbulence.

*For more information contact Michael Chertkov at [chertkov@lanl.gov](mailto:chertkov@lanl.gov).*

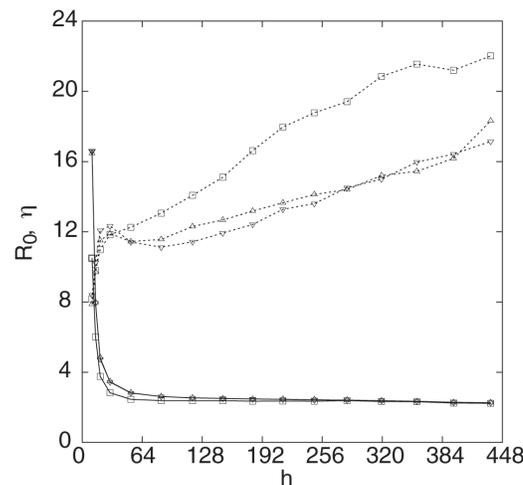
- [1] D.H. Sharp, *Physica D* **12**, 3 (1984).
- [2] G. Dimonte, et al., *Phys. Fluids* **16**, 1668 (2004).
- [3] M. Chertkov, *Phys. Rev. Lett.* **91**, 115001-4 (2003).
- [4] P.N. Wilson, M.J. Andrews, *Phys. Fluids* **14**, 938 (2002).
- [5] Y.N. Young, et al., *J. Fluid Mech.* **447**, 377 (2001).
- [6] W.H. Cabot, A.W. Cook, *Nature Physics* **2**, 562 (2006).
- [7] R. Ristorcelli, T.T. Clark, *J. Fluid Mech.* **507**, 213 (2004).
- [8] H.M. Tufo and P.F. Fischer, *J. Parallel Distr. Comput.* **61**, 151 (2001).

### Funding Acknowledgements

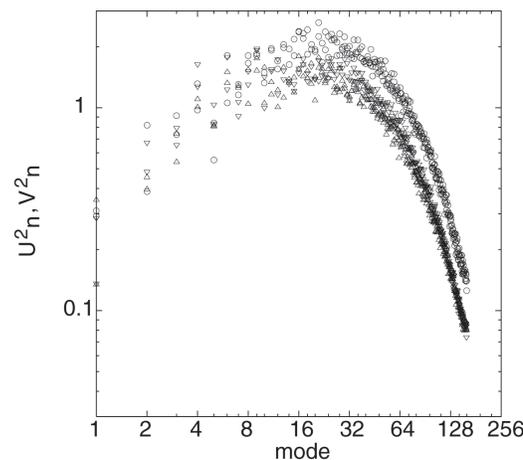
Los Alamos National Laboratory's Weapons Supported Research Program and Laboratory Directed Research and Development Program; and University of Chicago Center for Astrophysical Thermonuclear Flashes.



**Fig. 2.** Horizontal velocities vs height at times  $t = 64$  (solid line),  $t = 96$  (dashed line), and  $t = 128$  (dotted line). Curves taken at different values of  $h$  are almost indistinguishable from each other.



**Fig. 3.** The energy-containing scale (dashed lines) and the viscous scale (solid lines) in the middle of mixing layer. Squares correspond to vertical component of velocity, while triangles correspond to horizontal components of velocity.



**Fig. 4.** Energy spectra at time  $t = 64$  in horizontal planes  $z = 0$  (circles) and  $z \pm 0.75h$  (triangles).

## Replicate Variation

Michael L. Fugate, CCS-6; Timothy C. Wallstrom, T-13

Over the past year, we have been working on the problem of estimating replicate variation, which is the variance in some parameter arising from unmodelled differences in nominally identical experiments, or replicates. Such estimates are essential for many purposes, such as comparing computational models with experimental data.

If all of the experiments are of the same type, the problem is very simple: we simply use the standard variance estimate. Frequently, however, our data come from many different types of experiments, and the variance may differ from one group of experiments to the next.

To understand the nature of the problem, consider two different models we might use to solve this problem. In the first model, we assume that the variance of the groups is completely independent, so we estimate the variance separately for each group. However, there may be few datapoints for each group, so the uncertainty in each of the variance estimates may be so large as to render them almost worthless. If the groups are related, then data from one may give information about others, and we are throwing out a lot of information in each of the individual estimates. In the second model, we assume that the variances of the groups are identical, and we pool all of the data. We now have many datapoints, and the statistical

uncertainty is greatly reduced. However, we may be inappropriately assigning the same variance to groups whose variances are actually quite different. Some middle ground is needed.

The middle ground is provided through the use of Bayesian methods, using an approach known as hierarchical modeling. We introduce different variance parameters for each group. However, we do not allow these estimates to vary freely; they must come from a “population distribution” that is determined, in typical Bayesian fashion, from a combination of “prior information” and data, where the data are from all the groups combined. In this way, the data itself are able to influence whether the estimates for different groups should be constrained to agree closely, or allowed to be independent.

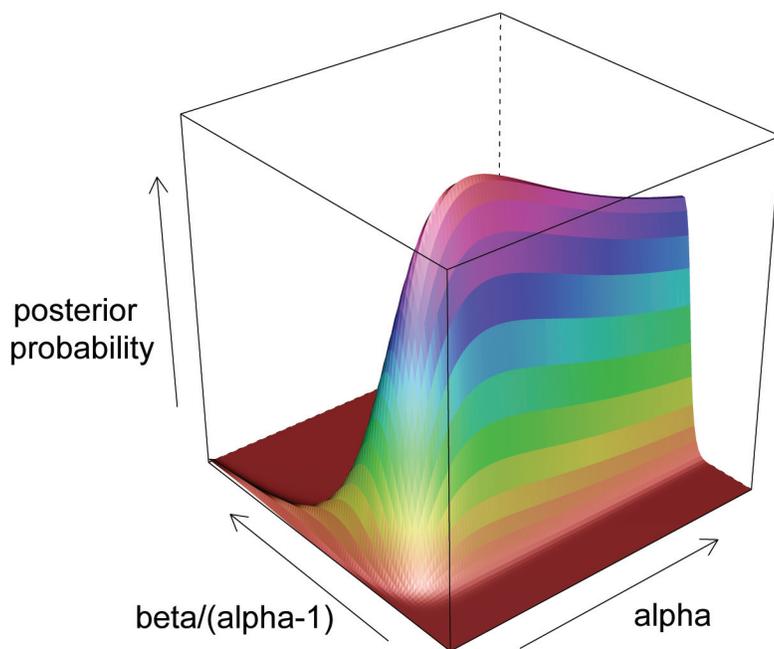
In Fig. 1, we show how the data from several groups of replicates can be used to improve the estimate of the variance in any particular group. The replicate variation is assumed to be described by an “inverse gamma” distribution, parameterized by unknown  $\alpha$  and  $\beta$ . The key parameter is  $\alpha$ : the larger the value of  $\alpha$ , the better we know  $\sigma$ . The figure shows the posterior probability of  $\alpha$  and  $\beta$ , given 20 groups of 6 replicates each, where the prior probability was assumed constant. The peak occurs at about  $\alpha = 10$ , which is close to the value of 15 which was used to generate the data. We can use this information to obtain better estimates of the replicate

variance for any particular group, even though we did not assume that the variance was the same for every group.

*For more information contact Timothy Wallstrom at [tcw@lanl.gov](mailto:tcw@lanl.gov).*

#### **Funding Acknowledgements**

NNSA's Advanced Simulation and Computing (ASC), Verification and Validation Program.



**Fig. 1.**  
*Posterior probability  
of  $\alpha$  and  $\beta$ .*

## Modeling Interfacial Surface Tension in Fluid Flow

Marianne M. Francois, James M. Sicilian, CCS-2; Douglas B. Kothe, Oak Ridge National Laboratory

**I**nterfacial flows with surface tension are frequently encountered in nature and in industrial applications, an example being in material processing. Accurately modeling such flows is challenging because of the discontinuity in material properties (such as density) and because of the interfacial boundary condition (surface tension).

As part of the Telluride project, we have developed a balanced-force fluid flow algorithm to accurately model interfacial surface tension effects. The Telluride project is developing the Truchas code, a 3-D finite-volume multi-physics package that models material processes such as gravity casting. Fluid flow is modeled by solving the Navier-Stokes equations on a fixed mesh. The locations of different materials are tracked using the volume tracking method (also known as the volume-of-fluid method), which assumes a single fluid formulation with averaged material properties in a multi-material cell.

Our surface tension model is based on the continuum surface force (CSF) approach of Brackbill et al. [1] in which the interfacial surface force is transformed to a volume force in the region near the interface via a delta function:

$$\vec{F} = (\sigma\kappa\hat{n} + \nabla_s\sigma)\delta$$

where  $\sigma$  is the surface tension coefficient,  $\kappa$  is the mean interfacial curvature,  $\nabla_s$  is the surface gradient and  $\delta$  is the delta function. The first term of

the equation is acting in the normal direction and the second term in the tangential direction. If the surface tension coefficient is a constant, the tangential component is zero.

The CSF method has the propensity to generate unphysical flow (“spurious currents”) near the interface when surface tension forces are dominant. These spurious currents are best illustrated in the limiting case of an inviscid static drop in equilibrium where Laplace’s formula applies. The causes of these spurious currents are an imbalance of the surface tension and pressure gradient forces and errors in curvature estimations. Therefore, we have devised a balanced-force algorithm [2] such that the normal component of the surface tension force is able to exactly balance the pressure gradient. Because curvatures are proportional to second derivatives of the volume fraction function, it is necessary to smooth the volume fraction to obtain accurate curvatures using a convolution technique or use a more geometrical approach by estimating a height function and discretizing it [2]. Our balanced-force algorithm combined with the convolution curvature method nearly eliminates the spurious currents for the case of the static drop as illustrated in Fig. 1.

We have extended our balanced-force algorithm to model thermocapillary forces that are due to the variation of surface tension coefficient arising from temperature gradients along the interface [3]. Thermocapillary forces are transmitted from the interface to

the bulk of the fluid by viscous forces, inducing Marangoni convection in the flow that causes the interface to deform. In our algorithm, the tangential component of the surface tension forces representing the thermocapillary effect is precisely balanced by the viscous stress.

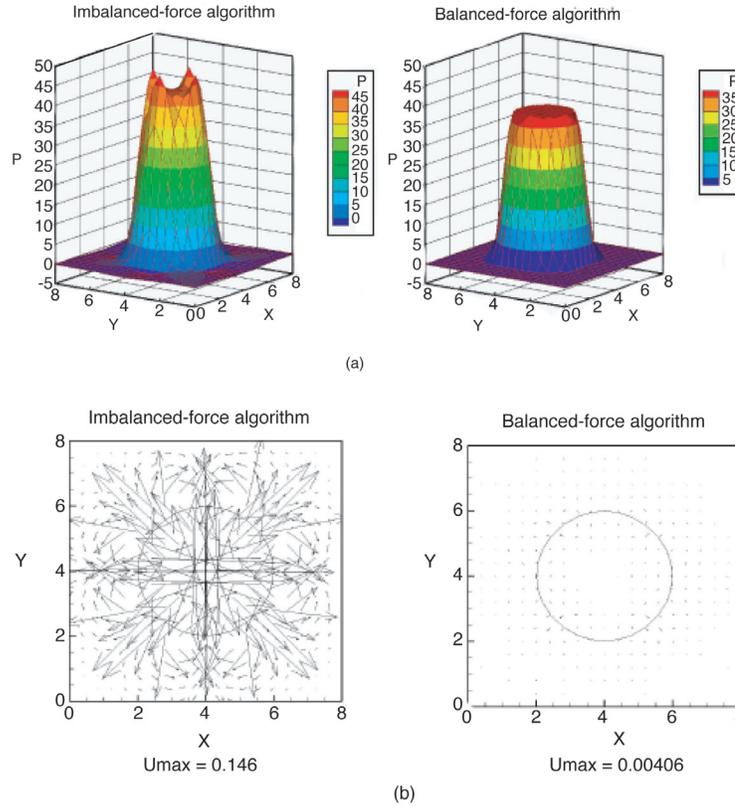
Finally, at rigid boundaries (or wall) we have modified the evaluation of the interfacial normal to represent the contact angle boundary condition as in [1]. Verification of the implementation in Truchas has been achieved by comparison with the analytical results of Sen and Davis [4] for thermocapillary flows in cavities. The steady state results of the simulations are shown in Fig. 2 for a variety of wall adhesion properties (contact angles).

*For more information contact Marianne Francois at mmfran@lanl.gov.*

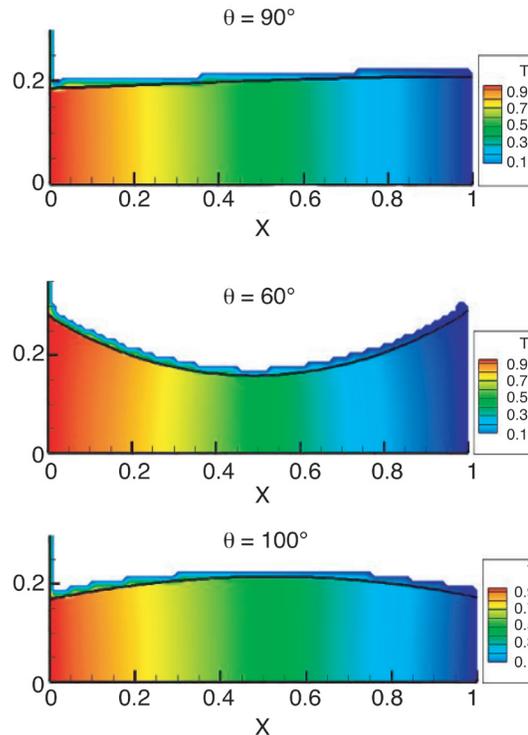
- [1] J.U. Brackbill, et al., *J. Comp. Phys.* **100**, 335–354 (1992).
- [2] M.M. Francois, et al., *J. Comp. Phys.* **213**, 141–73 (2006).
- [3] M.M. Francois, “Modeling of Thermocapillary Forces Within a Volume Tracking Algorithm,” in *Modeling of Casting, Welding and Advanced Solidification Processes—XI*, C.A. Gandin, M. Bellet, Ed. (The Minerals, Metals & Materials Society, 2006) pp. 935–942.
- [4] A.K. Sen and S.H. Davis, *J. Fluid Mech.* **121**, 163–186 (1982).

**Funding Acknowledgments**

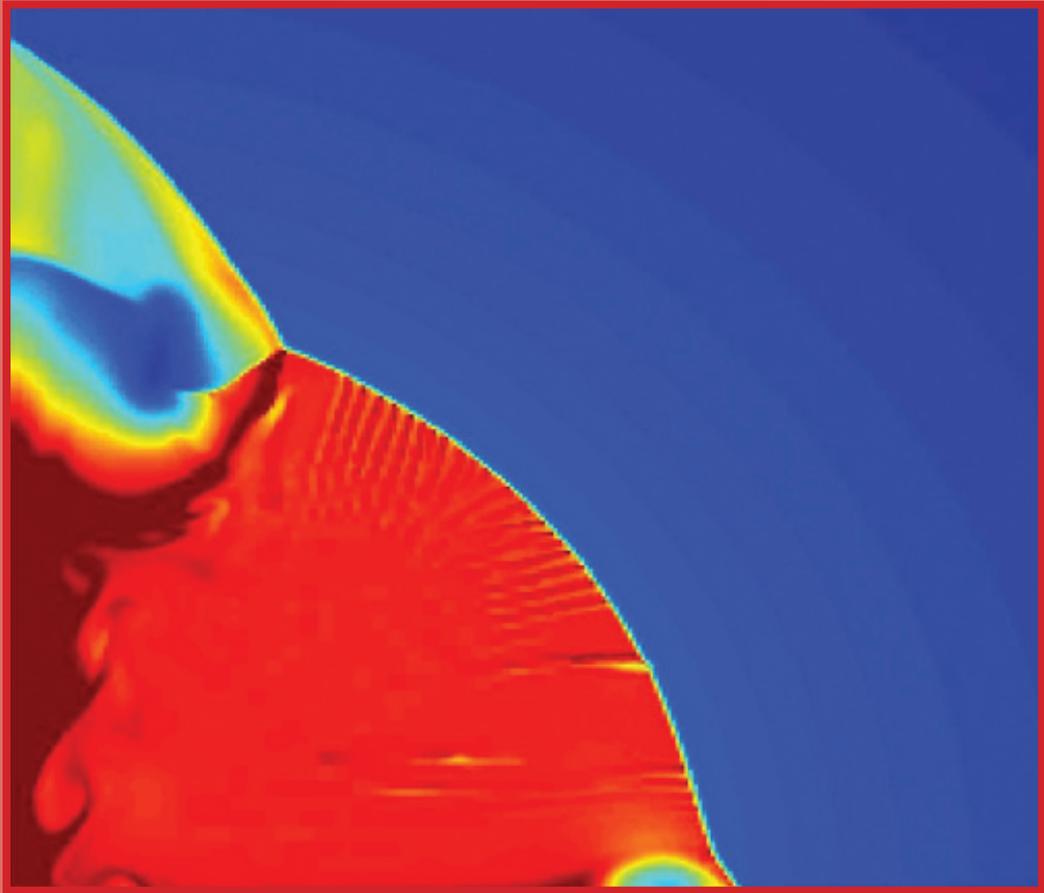
NNSA’s Advanced Simulation and Computing (ASC), Integrated Codes Program.



**Fig. 1.** *Imbalanced vs balanced-force algorithm for the case of an inviscid static drop of radius  $R = 2$  in an  $8 \times 8$  domain with mesh size  $20 \times 20$ ; (a) pressure distribution and (b) velocity field.*



**Fig. 2.** *Temperature contours and interface shape at steady state for thermocapillary flows in cavities for different contact angle  $\theta$ . Initially the interface is flat and the cavity aspect ratio is 0.2. The mesh spacing is uniform of 0.02.*



# **Sensitivity Analysis, Error Propagation, Uncertainty Quantification, and Verification and Validation**

Of critical importance is the study of the effectiveness of our computational physics enterprise. This section deals with various techniques to study, and improve, the accuracy of our simulation tools ultimately driving us closer to true scientific predictive capability. Two articles are presented that utilize statistical sciences to help in this regard. The first is about statistical models to tame the complexities of assessing the reliability of the enduring nuclear stockpile from a myriad of data and information sources. The second presents a discussion of the use of a statistical approach in the quantification of uncertainties in coupled physics simulations. The third article in this section describes a methodological approach to quantifying the accuracy and convergence of a particular numerical hydrodynamic algorithm.

## Statistical Models for Stockpile Health Assessment

*Alyson Wilson, Mike Hamada, Christine Anderson-Cook, Todd Graves, Scott Vander Weil, Benjamin Sims, Greg Wilson, CCS-6*

**T**he Los Alamos National Laboratory (LANL) Statistical Sciences Group (CCS-6) has an ongoing applied research effort to develop methods and software tools to help assess the health of the enduring U.S. nuclear stockpile. Given the complex scientific and data issues associated with understanding the stockpile, CCS-6 works within LANL to apply the best data analysis methods to support decision making. CCS-6 also collaborates with the Department of Defense and other agencies on problems that improve the understanding of weapon stockpile assessment and complex system reliability. (See Figs. 1 and 2).

We have developed statistical methodology for a number of problems arising in reliability assessment, including appropriately handling multilevel data using reliability block diagram and fault tree representations and faithfully dealing with the less than ideal data, which includes biased and incomplete data. Other methodology we have developed is applicable to support quantified margins and uncertainties (QMU) assessments. We have also developed and applied methodology to address a number of questions facing core surveillance and those arising in significant finding investigations (SFIs). Finally, much of the implementation has been done using YADAS, software we have developed to address the statistical challenges often presented by weapons data (see [1] and [yadas.lanl.gov](http://yadas.lanl.gov)).

Our recent successes in supporting the weapons program have included efforts to:

- Combine diverse data types (e.g., pass/fail, accelerated, degradation, lifetime and specification data, and expert judgment) at multiple levels (e.g., system, subsystem, component) to evaluate system reliability. [2,3,4,5].
- Take indirect measurements to make inference about characteristics based on direct measurements that were not taken [6].
- Estimate the proportion of a population with an attribute from data that were purposely biased to contain units with that attribute [7,8].
- Analyze degradation with a model that implies a Weibull lifetime distribution and assess reliability [9].
- Support QMU assessments with statistical models and analyses [10].
- Understand how reduced data collection impacts reliability.
- Understand how to collect degradation data over time for applications like shelf-life programs [11].
- Track and trend surveillance streams from LANL stockpile weapons to help assess stockpile health (see Fig. 3)[12].

*For more information contact Alyson Wilson at [agw@lanl.gov](mailto:agw@lanl.gov).*

[1] T.L. Graves, "Design Ideas for Markov Chain Monte Carlo Software" (accepted by *J. Comput. Graphical Stat.*, Sept. 2006).

[2] T.L. Graves, M. Hamada, "Bayesian Methods for Assessing System Reliability: Models and Computation," in *Modern Statistical and Mathematical Methods in Reliability*, Eds. A. Wilson, N. Limnios, S. Keller-McNulty, Y. Armijo, (World Scientific Publishing Company, October 2005) pp. 41–54.

[3] A.G. Wilson, et al., "Advances in System Reliability Assessment, Los Alamos National Laboratory report LA-UR-05-8755 (2005) (accepted by *Stat. Sci.*, Nov. 2005).

[4] C. Anderson-Cook, et al., "Bayesian Stockpile Reliability Methodology for Complex Systems with Application to a Munitions Stockpile," Los Alamos National Laboratory report LA-UR-05-1959 (March 2005) (accepted by *Military Ops. Res. J.*, Aug. 2006).

[5] C.S. Reese, et al., "A Hierarchical Model for the Reliability of an Anti-aircraft Missile System," Los Alamos National Laboratory report LA-UR-05-9281 (2005) (under revision for *J. Am. Stat. Assoc.*, Dec. 2005).

[6] T.L. Graves, M. Hamada, *Qual. Eng.* **17**, 555–559 (Oct. 2005).

[7] T.L. Graves, M. Hamada, *Qual. Reliability Eng. Intl.* **22**, 385–392 (July 2006).

[8] T. Graves, et al., "Estimating a Proportion Using Stratified Data Arising from Both Convenience and Random Samples," Los Alamos National Laboratory report LA-UR-03-8396 (2003) (resubmitted to *Technometrics*, April 2006).

[9] M. Hamada, *Qual. Eng.* **17**, 615–620 (Oct. 2005).

[10] J. Booker, et al., *Military Ops. Res. J.* **11**, 27–44 (Aug. 2006).

[11] E. Kelly, S. Vander Wiel, "Planning and Evaluating Sampling Designs to Determine Aging Effects," Los Alamos National Laboratory report LA-UR-06-7244 (2006).

[12] C. Faust, et al., "Analysis of Certification and Surveillance Data for Pits: Report 1 (U)," Los Alamos National Laboratory report LA-CP-07-0090 (2007).

### Funding Acknowledgements

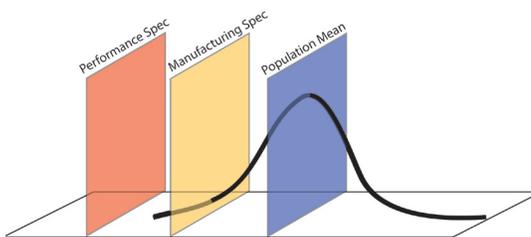
NNSA's Campaign 8, Enhanced Surveillance; and Joint DoD/DOE Munitions Technology Development Program.



**Fig. 1.** Andrew Wiedlea and Nick Hengartner examine a test set used to collect data on the health and functionality of missile components from weapons like the MK 67 submarine-launched mine in the picture.



**Fig. 2.** CCS-6 is collaborating with the Naval Surface Warfare Center to assess the reliability of the RAM missile system.



**Fig. 3.** CCS-6 is working with the LANL Core and Enhanced Surveillance programs to track and trend weapon surveillance data streams. The goal is to assess key weapon characteristics against known specifications and estimate the distribution of these characteristics in the stockpile now and in the future.

## Uncertainty Quantification for Simulation-Based Predictions

Dave Higdon, Brian Williams, Jim Gattiker, Mike Fugate, CCS-6

**A**ssessing the performance, safety, and reliability of the nuclear weapon stockpile is complicated by the Comprehensive Test Ban Treaty that forbids underground nuclear testing. Hence, assessing the current nuclear stockpile relies on the use of detailed, physics-based computer codes. These simulation models can be refined by utilizing information from historical nuclear tests, as well as from ongoing experimental campaigns, to predict weapon behavior. For more information see Refs. [1,2].

Currently substantial statistical research is focused on the development of methodology for utilizing detailed simulation codes to carry out inference for physical systems. Issues such as calibration of simulator input parameters, generation of predictions, and characterization of prediction uncertainty are of particular interest. Simulation of well-understood physical processes is typically based on fundamental physical principles. In such problems, the actual amount of observed field data from this process can often be very limited. It is the simulator code that contains the structure of the actual process being modeled. Because of this, useful inference is possible even with only minimal amounts of observed data on the actual physical system. Figure 1 describes a simple application where a mathematical model enhances our ability to predict how long it will take an object to reach the ground when dropped from a tower.

The statistical methodology we have developed in support of nuclear weapon certification requires a limited amount of field data, which may be from experiments or observations on the physical system, along with an ensemble of computer-model simulations. Gaussian processes are used to model the simulator output at untried input settings. This model for the simulator is then embedded in a larger statistical framework so that parameter estimation (i.e., calibration) and prediction can be carried out.

Although originally developed for weapons-related applications, the statistical methodologies we have developed have potential use in any investigation where a computer model is available to model a given system. To date we have applied these methodologies to a wide range of application areas: engineering, particle physics, manufacturing, climate and cosmology (see Fig. 2).

*For more information contact Dave Higdon at [dhigdon@lanl.gov](mailto:dhigdon@lanl.gov).*

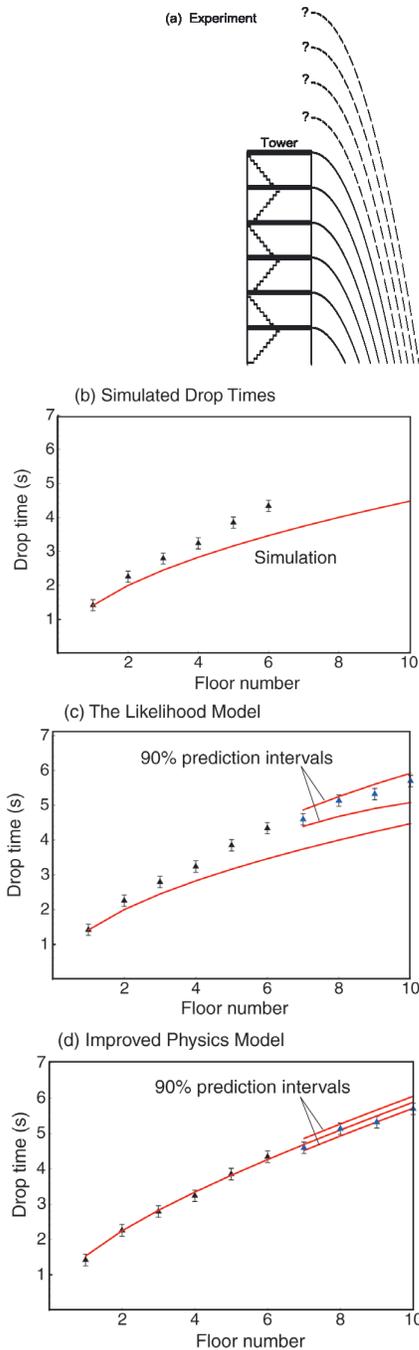
[1] D. Higdon, et al., *SIAM J. Sci. Comput.* **26**, 448–466 (2004).

[2] B. Williams, et al., *Bayesian Anal.* **1**, 765–792 (2006).

[3] K. Heitmann, et al., *Astrophys. J. Lett.* **646**, L1, (2006).

### Funding Acknowledgements

NNSA's Advanced Simulation and Computing (ASC), Verification & Validation Program; and Campaign 1, Primary Assessment Technology.



**Fig. 1.**

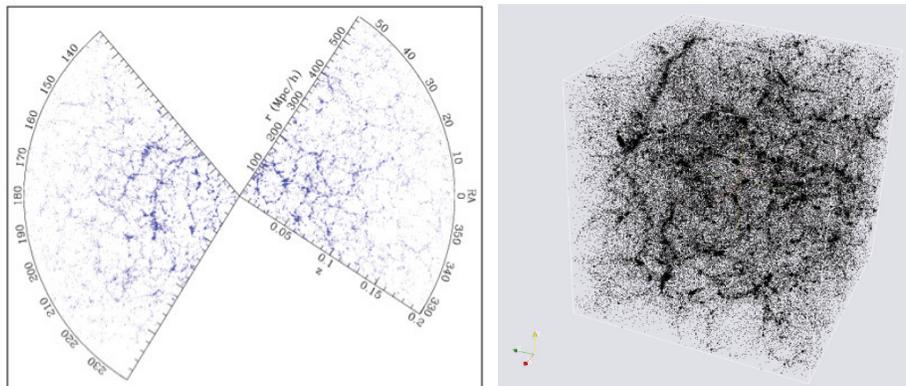
**Dropping an Object from a Tower**

(a) The time it takes an object to drop from each of six floors of a tower is recorded. There is an uncertainty in the measured drop times of about  $\pm 0.2$  s. Predictions for times are desired for drops from floors 7–10, but they do not yet exist.

(b) A mathematical model is developed to predict the drop times as a function of drop height. The simulated drop times (red line) are systematically too low when compared with the experimental data (triangles). The error bars around the observed drop times show the observation uncertainty.

(c) This systematic deviation between the mathematical model and the experimental data is accounted for in the likelihood model. A fitted correction term adjusts the model-based predictions to better match the data. The resulting 90% prediction intervals for floors 7–10 are shown in this figure. Note that the prediction intervals become wider as the drop level moves away from the floors with experimental data. The cyan triangles corresponding to floors 7–10 show experimental observations taken later only for validation of the predictions.

(d) An improved simulation model was constructed that accounts for air resistance. A parameter controlling the strength of the resistance must be estimated from the data, resulting in some prediction uncertainty (90% prediction intervals are shown for floors 7–10). The improved model captures more of the physics, giving reduced prediction uncertainty.



**Fig. 2.**

Methodology originally developed for nuclear weapon certification has been applied to investigations in cosmology. Here observations from the Sloan Digital Sky Survey (left) and large-scale gravity simulations (right) are used to constrain cosmological parameters that describe the evolution of the universe. See Ref. [3] for details.

# Code Assessment for Adaptive Mesh Refinement Simulations

Shengtai Li, T-7; William J. Rider, X-1

**W**e investigated and analyzed the grid convergence issues for adaptive mesh refinement (AMR) code. These issues are related to the RAGE code and have been raised in a previous report [1].

Code verification is extremely important for science-based prediction and simulations. Previous verification focused on the convergence behavior of uniform grid. Using AMR, we can obtain more accurate results with substantially less computational cost. Ideally, AMR should achieve the same accuracy in refinement region as the corresponding fine uniform grid. We expect the results of an AMR grid at least should be better than the results of the coarse uniform grid without local refinement.

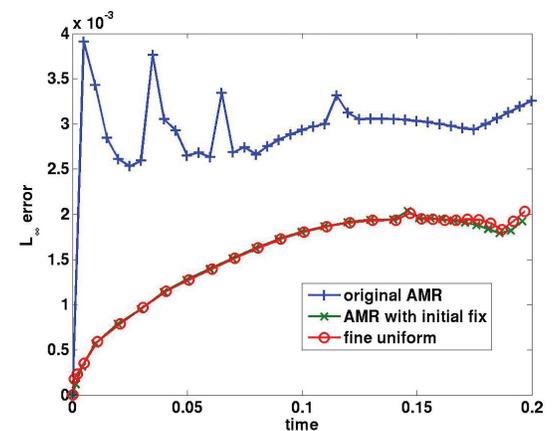
Using two different AMR packages [1], we have investigated three model problems that have exact solutions and represent a variety of problems. Several issues with respect to the RAGE AMR have been found: (a) it has large initialization errors, (b) the numerical error with AMR is larger than without AMR for a high-resolution grid, (c) AMR with more than one-level refinement has larger numerical errors than with only one-level refinement, and many other issues.

We have investigated these issues in more detail and proposed several methods to solve them [2]. In particular, we tested a new mesh initialization for AMR solutions and several refinement criteria to achieve the expected accuracy and convergence rate for AMR simulations.

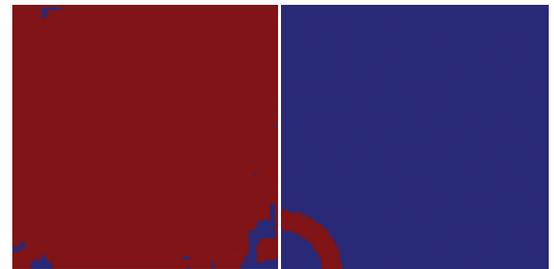
## AMR Initialization Issue

After careful examination of the initial errors and RAGE AMR implementation, we have found that the RAGE AMR never generates the initial AMR mesh based on a user-input routine for exact initial conditions. The initial

error is purely an interpolation error from the coarse grid to the AMR fine grid. We proposed a new mesh generation procedure for RAGE AMR to use the user-input routine. This new procedure does not require too much extra coding and eliminates the source of the initial errors (see Fig. 1).



We also proposed a new AMR flagging approach to turn off the activity test so that AMR does not refine everywhere for Noh's problem (see Fig. 2). This approach also fixes the shock instability for Noh's 3-D spherical problem (see Fig. 3).

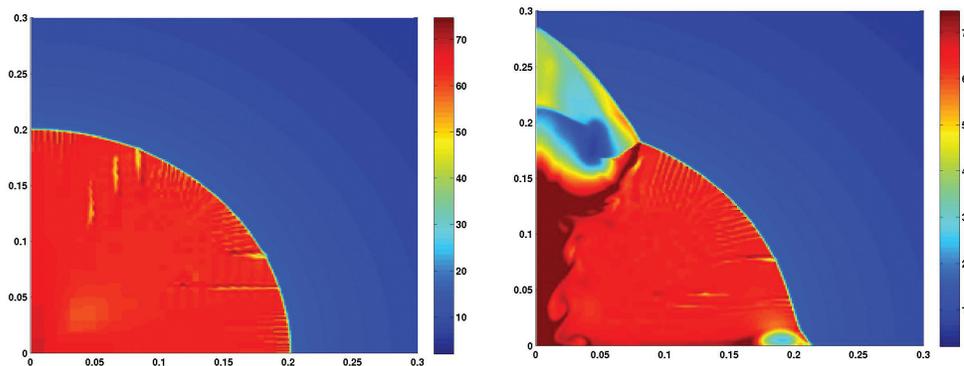


## AMR Refinement Criteria

After detailed analysis, we have found that the numerical solution at the coarse-fine interface between different levels of grid converges only in the first-order accuracy. There-

**Fig. 1.** Numerical error ( $L_\infty$ ) before and after AMR initialization fix for wave's problem.

**Fig. 2.** Mesh refinement before (left) and after (right) AMR fix for Noh's problem. The red color denotes the second level.

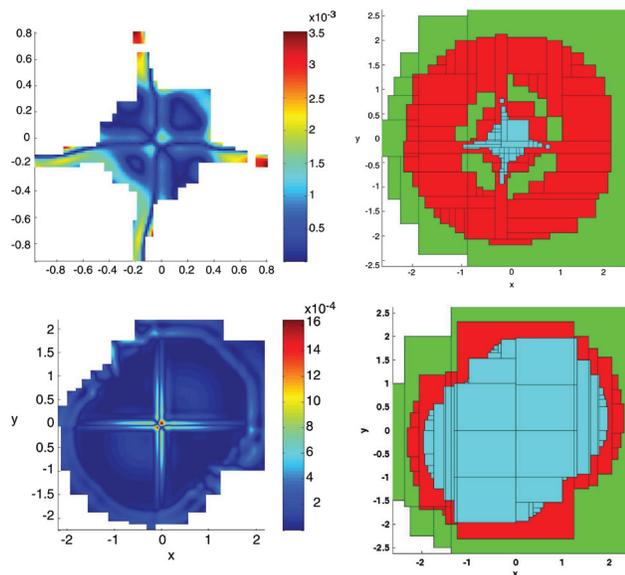


**Fig. 3.** Fixing AMR refinement also removes the shock instability for Noh's 3-D problem. Right: before fix; Left: after fix.

fore, the error near the coarse-fine interface can quickly dominate the error in other regions if the coarse-fine interface is active and not covered by the fine grid.

We implemented and compared several refinement criteria (RC): solution gradient-based, solution curvature-based, RC of FLASH code, Richardson extrapolation type, operator recovery error source detector (ORES) of Lapenta [3], etc. Some of them can catch the large-error region near the coarse-fine interface and refine them with the fine grid. We found the Richardson extrapolation approach and modified ORESD approach performed better than others. Several issues in RAGE AMR are solved by using the new refinement criteria.

The numerical results, Figs. 4 and 5, show that the refinement criteria play an important role in convergence behavior of AMR solutions.



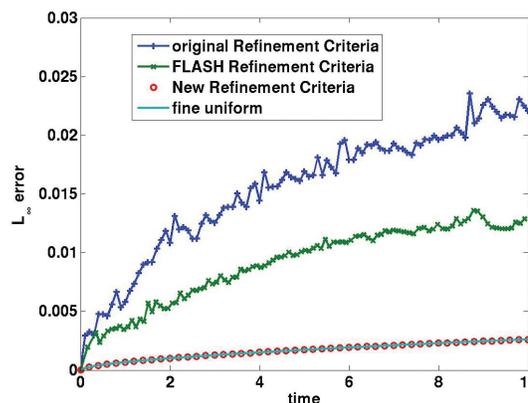
**Fig. 4.** Numerical error for different refinement criteria.

*For more information contact Shengtai Li at [sli@lanl.gov](mailto:sli@lanl.gov).*

- [1] S. Li, et al., "Two-Dimensional Convergence Study for Problems with Exact Solution: Uniform and Adaptive Grids," Los Alamos National Laboratory report LA-UR-05-7985 (October 2005).
- [2] S. Li and W.J. Rider, "Code Verification and Assessment for Adaptive Mesh Refinement Simulations," Los Alamos National Laboratory report LA-UR-06-4595 (2006).
- [3] G. Lapenta, *Intl. J. Numer. Meth. Engng.* 59 (15), pp. 2065–2087 (2004).

#### Funding Acknowledgements

NNSA's Advanced Simulation and Computing (ASC), Verification and Validation.



**Fig. 5.** Numerical error and mesh refinement for AMR with 3-level refinement. The top two is for FLASH's refinement criteria and the bottom two is for ORESD.



# Manufacturing Simulation and Computation

The advent of very powerful and inexpensive computing has made its way into multiple endeavors at deep and fundamental levels. Computational manufacturing is now widely seen as a bridge between design and production, improving the design of the manufacturing process and the quality of the produced product. In this section, the state-of-the-art in the active visualization of the virtual design, assembly, and even testing of complex mechanical parts, paving the way for the transformation of the DOE complex in this regard, is described in the first article. The second article is centered on a very successful integrated coupled physics tool for the simulation of manufacturing processes involving casting. Here we see a tangible example of how simulation and modeling can speed the delivery and quality of manufactured components while reducing production costs and waste.

## Visualization of Complex Mechanical Models

David Pugmire, HPC-4

Visualization has long been an important tool to understand complex scientific data. An unexplored application of 3-D visualization at Los Alamos National Laboratory (LANL) has been the design, manufacture, and assembly of complex mechanical models. These models are typically designed, manufactured, assembled, and tested by a large number of people, often in different locations. Allowing these groups to interact with an unbuilt, virtual model has the potential to transform the way the DOE complex operates.

Pro-Engineer (Pro-E), the DOE's Nuclear Weapons Complex production mechanical design package, and EnSight, the LANL production visualization tool, do not share a common data format, so translation is required to produce visualizations. Typically, this is done by hand in Pro-E for each part; a prohibitively time-consuming and error-prone process for models with hundreds of parts. Because the models were undergoing continual revision, and

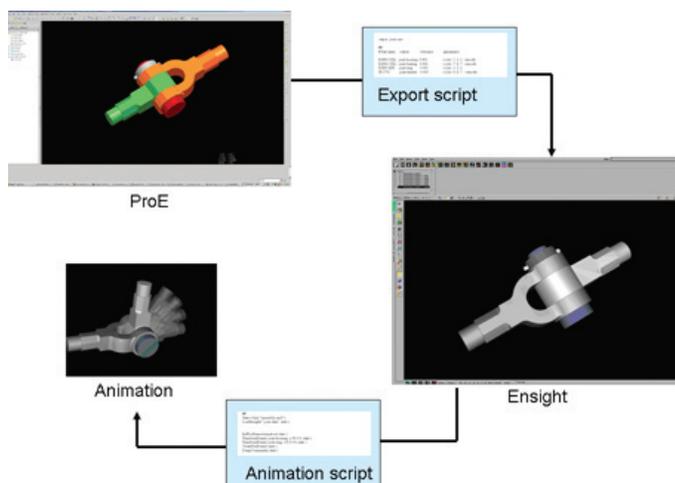
frequent retranslation was required, an automated method for translation was critical.

We accomplished this by using built-in hooks in the Pro-E user interface that allow the integration of user-written code. As shown in Fig. 1, the input to the tool is a text file that specifies the output EnSight file, the parts to be translated, translation tolerance, and display attributes. Once the model is loaded into Pro-E, the translation code is executed, producing files that can be read in to EnSight for visualization. Using this automated sequence, we reduced the time to prepare a large model for visualization from 7-8 hours to about 5 minutes.

Another key component to this project was the creation of animations that illustrated most significant aspects of the models, including assembly, disassembly, and the functioning of subcomponents.

Creating an animation that effectively conveys the intended information can be tedious and difficult with currently available tools. We developed a set of utilities with the Python scripting language that facilitates experimentation and exploration of the animation space. As illustrated in Fig. 1, once a suitable animation is completed, the utilities are capable of generating movies for a variety of output formats

**Fig. 1.** Data translation from the modeling tool Pro-E to the visualization tool EnSight is accomplished using an export specification script that defines the parts to be exported and their appearance. The animation script defines how the parts are moved to create the desired animation.



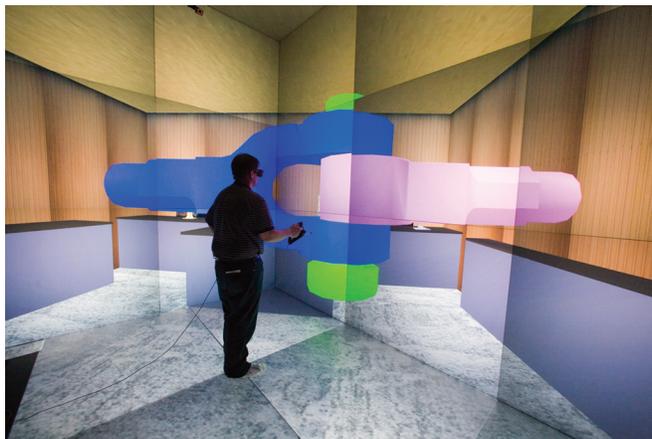
from audio-visual interface format (avi) to high definition, stereo, PowerWall theater, and immersive CAVE movies.

We used these tools in numerous presentations related to the Reliable Replacement Warhead (RRW) project to help illustrate key aspects of complex models, including at several meetings with assembly plant engineers where 1:1 scale models were explored and design alternatives discussed to facilitate manufacturing and assembly processes. Additionally, we analyzed assembly and disassembly animations using the actual plant tooling components inside the immersive CAVE (see Fig. 2). These tools proved to be a valuable component for the RRW project in the exploration and validation of complex models.

*For more information contact David Pugmire at [pugmire@lanl.gov](mailto:pugmire@lanl.gov).*

#### **Funding Acknowledgements**

NNSA's Advanced Simulation and Computing (ASC), Production Visualization project.



**Fig. 2.**

*A user interactively assembling a mechanical joint inside a virtual environment in the CAVE. The glasses and hand-held wand are motion tracked to allow free movement of objects within the virtual world. This capability allows designers to test and validate design options, create and verify safety plans and provide training experience for workers.*

*Photo by L. Sanchez*

## Truchas: A Simulation Tool for Manufacturing

*Advanced Simulation and Computing (ASC) Telluride Project Team*

The Advanced Simulation and Computing (ASC) Telluride project is developing a highly parallel, coupled, continuum scale, multiphysics simulation tool called Truchas to aid nuclear weapon component manufacturing. A large team of researchers has contributed to the development of Truchas for more than a decade. Recent successes demonstrate that Truchas is making significant contributions to responsive and flexible manufacturing processes at Los Alamos National Laboratory and other National Nuclear Security Administration sites.

For example, the LANL Reliable Replacement Warhead (RRW) design relied heavily on casting technology and on modeling for a variety of components within pits, secondaries, and radiation cases. To meet the recent RRW hydro schedule, process modeling with the Truchas code was successfully utilized to develop casting parameters for the radiation midcase and resulted in an acceptable part 19 working days from receipt of the part definition. The short turnaround time hinged on the use of simulation tools including Truchas, and demonstrated a responsive capability for manufacturing that has never been seen in the nuclear weapons complex.

Examples of the types of Truchas simulations that were used to develop the casting processes for RRW are shown in the figures. The induction heating capability allows an engineer to evaluate the heating response of new mold designs. Figure 1 shows the temperatures developed in a hemispherical casting assembly after induction heating for 40 min.

The cooling rate and solidification time interval of the material in a casting are important indicators of casting quality. Figure 2 shows an example solidification time plot for a simple mold geometry.

Truchas is a unique tool because it incorporates all of the continuum physical models necessary to model metal casting processes.<sup>1</sup> Commercial products are also useful for process development, but none incorporate the breadth of models in the current release of Truchas. This set of features enabled the Telluride project to successfully complete a FY05 ASC Level 2 milestone to simulate casting of a plutonium pit from beginning to end. [1]

The following list shows the physical models required for such a simulation.

- Periodic behavior of electromagnetic waves and the resultant heat deposited in the components of the manufacturing system.
- Transfer of heat within system components and between the components and their environment by conduction, convection, and radiation (with view factors).
- Solidification, melting, and allotropic (crystalline) phase changes. These must incorporate the influence of alloy concentrations that can vary within the domain of the alloy (and may result in macro-segregation of alloy concentration).
- Flow of molten alloy. The flow algorithms must deal with dynamic interfaces and must permit evaluation of the redistribution of energy and alloy material due to fluid motion.
- The distortion and dislocation of system components (including mold and the metal product) due to stresses imposed both by temperature changes and by volumetric changes associated with phase change. This model must accommodate both elastic and plastic response to such stress.

Because of the complex geometry of real manufacturing systems, Truchas uses an unstructured computing grid. Truchas implements the numerical algorithms for these physical models as a sequence of separated

solutions to the continuum partial differential equations that describe each phenomenon over a short period of time. The solution is advanced in time by a sequence of such steps.

Manufacturing processes are inevitably transient in nature because they aspire to change the form of materials over a period of time. Truchas is fully transient throughout its modeling and algorithms. It is also fully parallel in its operation, which permits the code to utilize the large computer clusters installed at LANL by ASC. Parallelism is achieved by subdividing the physical domain of the problem and assigning each section to a computing processor.

Each physics model employs a discretization method that is unique to the needs of the model. However, the resulting discretized equations are solved using a limited set of linear and nonlinear solvers. The combination of physical models and numerical algorithms makes Truchas a robust and efficient parallel program for simulating metal casting processes.

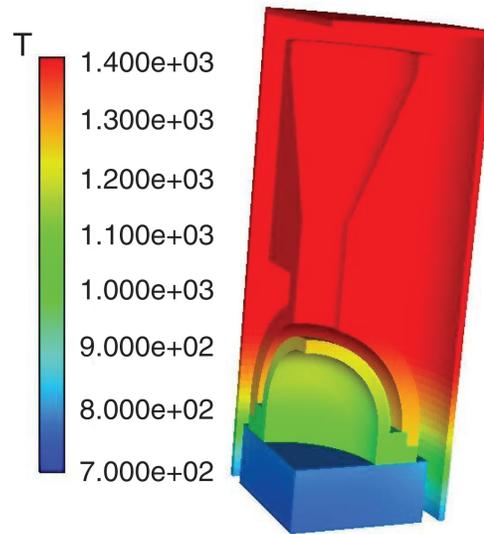
<sup>1</sup>The Telluride project is also working toward simulation of other manufacturing processes, including foam curing and welding.

*For more information contact Neil Carlson at [nnc@lanl.gov](mailto:nnc@lanl.gov).*

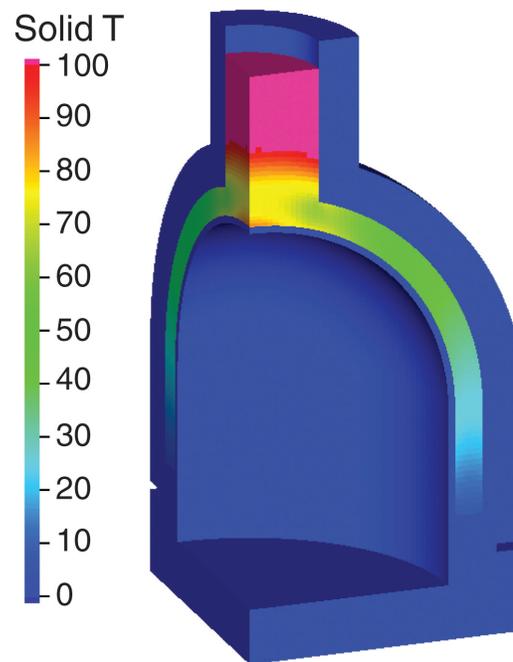
[1] D. Korzekwa, et al., "Truchas Simulation of Casting of the Qual Type 126 Pit (U)," Los Alamos National Laboratory report LA-CP-05-1224 (October 2005).

#### Funding Acknowledgements

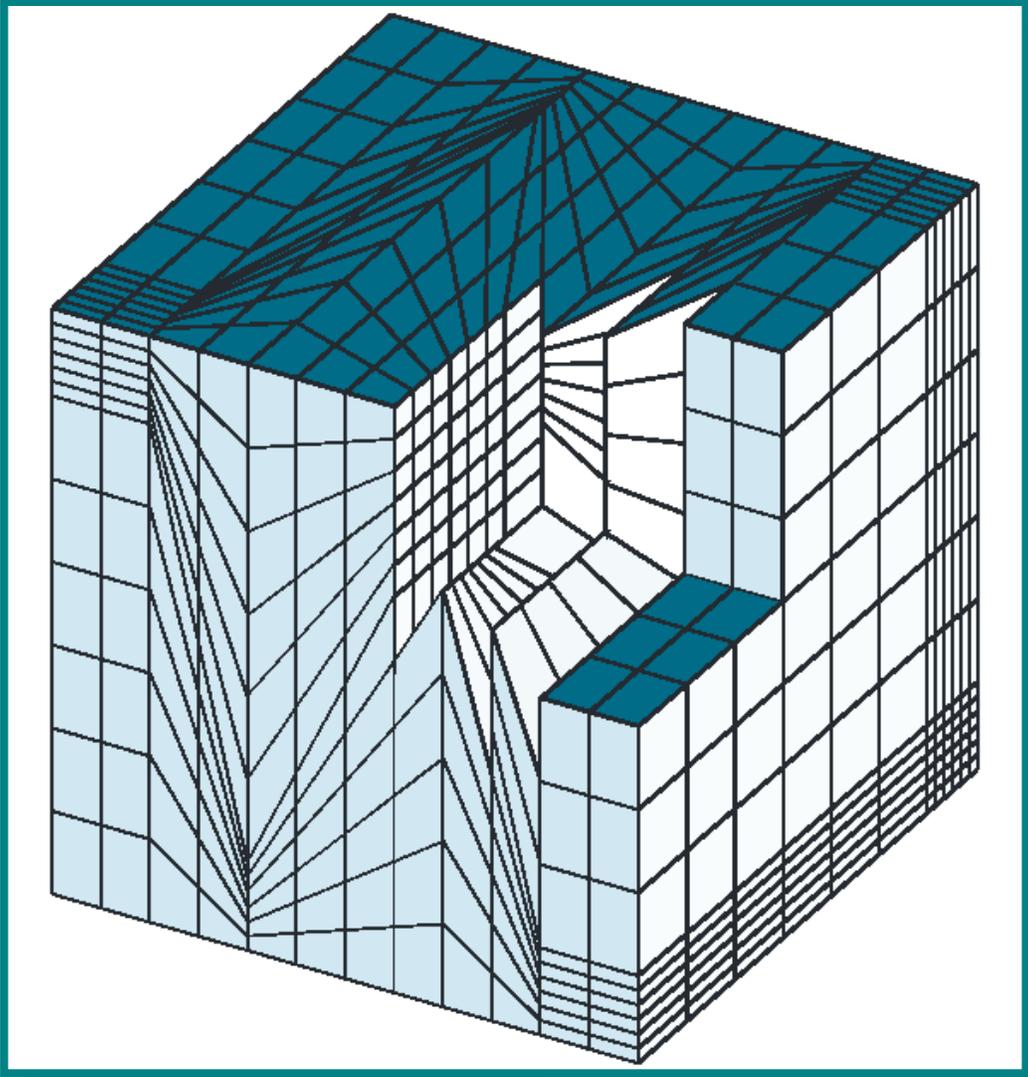
NNSA's Advanced Simulation and Computing (ASC), Integrated Codes.



**Fig. 1.**  
*Temperature Distribution after 40 min of induction heating.*



**Fig. 2.**  
*Solidification time interval prediction.*



# Applied Mathematical Sciences

This section is centered on the contributions of basic and applied mathematics to a myriad of important elements in the nuclear weapons enterprise. The first article describes an important new mathematical technique for accurate interface reconstruction. Other articles discuss efficient and accurate solvers and solver preconditioners of critical importance to multiple applications at Los Alamos National Laboratory. An article on examining discretizations on polyhedral methods also appears, as well as articles on new search and tensor multiplication methods. Like the field of computer science, applied mathematics is a complex and foundational field, with important contributions spanning the gamut of the nuclear weapons program. While these seven articles describe some key research and development activities in this field germane to the nuclear weapons program, one can find the applied mathematics at the foundation of all of the articles in this entire document. No other field of endeavor has both such a broad and deeply penetrating impact on science.

# Multimaterial Interface Reconstruction from the Moment Data

Vadim Dyadechko, Mikhail Shashkov, T-7

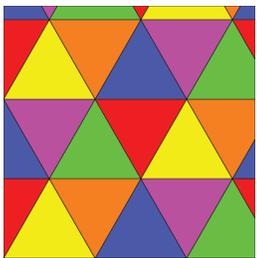
**V**olume-of-Fluid (VoF) methods [1] are widely used to approximate material interfaces in Eulerian fluid flow simulations. Instead of direct interface tracking, the VoF methods calculate the location of the interface at each time step from the solution data, namely from the cell-wise material volumes. This strategy faces no problem changing the topology of the interface dynamically; the choice of the cell-wise material volumes as an input for the interface reconstruction algorithm allows to preserve the volume of each material. Most VoF methods [2] use a single linear interface to divide two materials in a mixed cell; if the cell contains more than two materials, the two-material interface reconstruction algorithm is used repeatedly for extracting the materials from the mixture one by one. The VoF approach has apparent drawbacks: the resolution of the two-material interface reconstruction algorithm is 2 to 3 times lower than the resolution of the grid. There is no way to guess the order in which the materials should be separated from the multimaterial mixed cell, and, even if such an order is known *a priori* (is fixed), due to the limited resolution of the two-material algorithm, the multimaterial VoF reconstruction can be higher-than-first-order accurate only for the

layered material structure, i.e., only if the true interfaces form no junction.

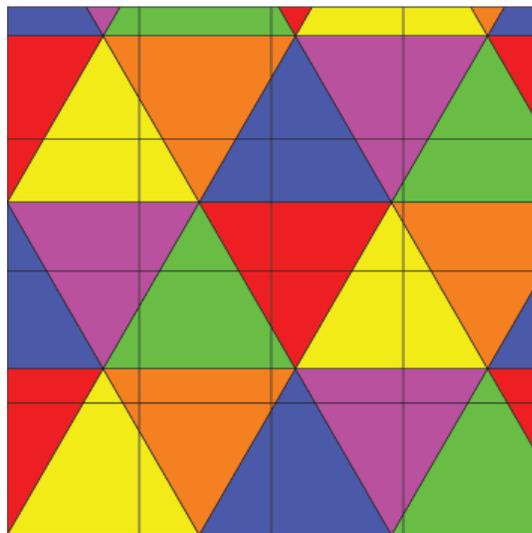
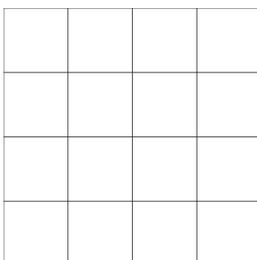
The new Moment-of-Fluid (MoF) technique effectively overcomes the limitations of the VoF approach by utilizing more data: in addition to the volumes, the cell-wise material centroids (the first moments) are used. In [3] we presented the two-material MoF algorithm, which locates the linear interface in a mixed cell by minimizing the defect of the first moment over all the cell partitions that preserve the material volumes. Now [4] we use the same governing principle to perform the polygonal partitioning of a multimaterial mixed cell.

A proper partitioning of the mixed cell with  $M \geq 3$  materials is a challenging problem. We explore two different partitioning schemes. The basic one follows the multimaterial VoF strategy and separates the materials from a mixed cell one by one. There is an essential difference though: the MoF interface reconstruction does not require the user to specify the material order explicitly. The right order is determined automatically by trying all  $M!$  possible material orders and finding the one that results in the minimal defect of the first moment. Another major improvement over the VoF, is that the MoF algorithm does not require the true interfaces to be nonintersecting to guarantee the second-order accurate approximation.

The search for the best partition above has combinatorial complexity in the number of materials: to get the answer, one has to try all  $M!$  material orders. On the other hand, it is reasonable to expect only a limited number of the mixed cells in the whole computational grid to contain 3+ materials. Therefore, for a moderate  $M$ , the computational overhead, associated with the optimal order search, is not likely to be significant. Also, various material orders can be effectively tried in parallel.



+ =



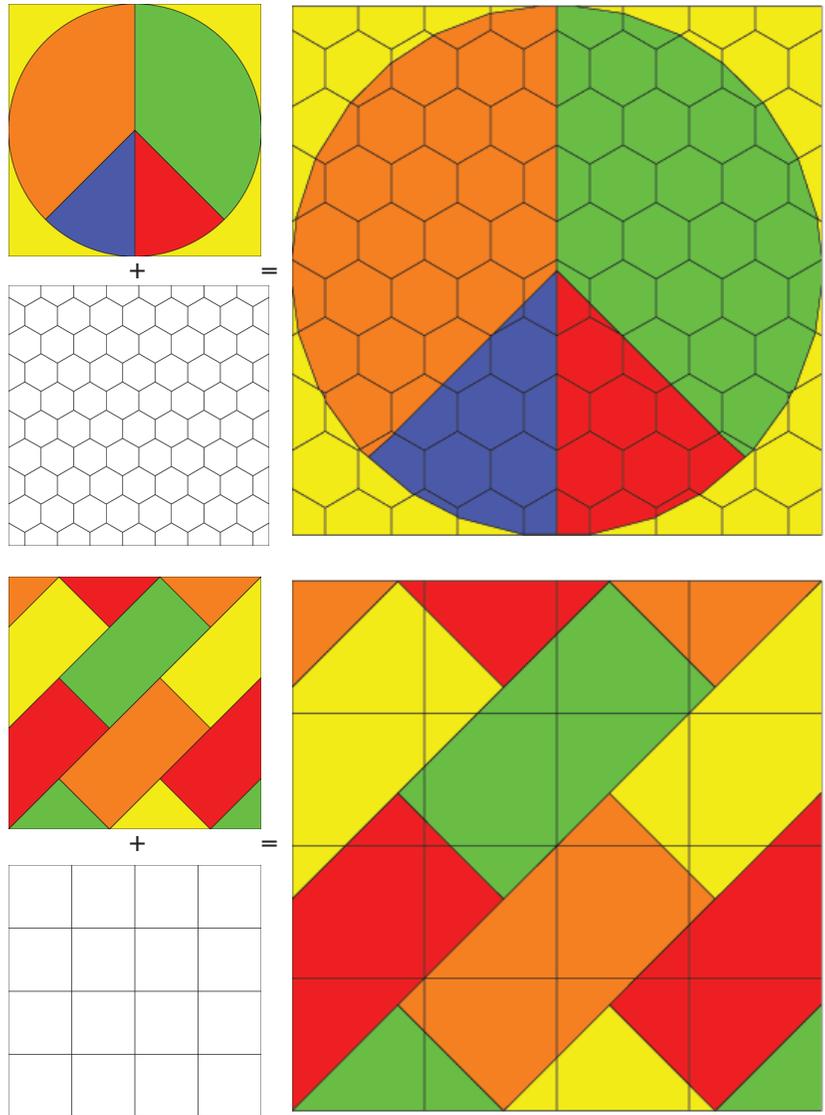
The governing principle of the MoF reconstruction (minimization of the first moment defect) does not limit the choice of the partitioning scheme in any way; in order to achieve a lower defect of the first moment, one can expand the family of trial partitions at will. Thus, along with extracting materials from the mixture in series, we propose to use a more sophisticated partitioning scheme that separates the materials according to the “divide-and-conquer” principle: choose an arbitrary  $m < M$ , separate the mixture of materials  $\bar{1}, m$  from  $m + 1, M$ , and then recursively apply this algorithm to each submixture containing 2+ materials. This procedure allows to generate  $M!(M-1)!$  trial *B-tree partitions* to choose from, which significantly increases the chances of finding an approximate partition that best fits the moment data. Such a partitioning scheme yields the MoF reconstruction of any B-tree partition with sufficiently smooth interfaces to be second-order accurate.

Although we explicitly address only the 2-D case, it is clear that all the partitioning and ordering strategies described are dimension-independent and are applicable in 3-D.

Unlike the VoF competitors, the Moment-of-Fluid interface reconstruction algorithm can partition multimaterial mixed cells in truly automatic manner; it is also capable of reconstructing complex interface junctions with second-order accuracy, which can hardly be achieved with the VoF methods.

For more information on MoF technique go to <http://math.lanl.gov/vdyadechko/research>, or contact Vadim Dyadechko at [vdyadechko@lanl.gov](mailto:vdyadechko@lanl.gov).

[1] C.W. Hirt and B.D. Nichols, *J. Comp. Phys.* **39** (1), 201–225 (Jan. 1981).  
 [2] David J. Benson, *Appl. Mech. Rev.* **55** (2), 151–165 (Mar. 2002).  
 [3] V. Dyadechko and M. Shashkov, “Moment-of-Fluid Interface Reconstruction,” Los Alamos National Laboratory report LA-UR-05-7571 (Oct. 2005).  
 [4] V. Dyadechko and M. Shashkov, “Multi-material Interface Reconstruction from the Moment Data,” Los Alamos National Laboratory report LA-UR-06-5846 (Aug. 2006).



**Fig. 1.**

The figures show three examples of the multimaterial MoF reconstruction. For each case we show the true distribution of the material in the domain (different colors represent different materials) and the computational mesh. The input data for the MoF algorithm (cell-wise material moments) were calculated by intersecting the mesh cells with the true material shapes. We would like to emphasize the exceptional resolution of the method: even though the size of the material “tiles” in the first and the last configurations is comparable to the size of the mesh cells, the MoF algorithm can reconstruct them exactly.

#### Funding Acknowledgements

NNSA’s Advanced Simulation and Computing (ASC) Program.

# A Comparison of Multilevel Preconditioners for Solving Multimaterial Equilibrium Radiation Diffusion Problems on Locally Refined Grids

Bobby Philip, T-7; Michael Pernice, Idaho National Laboratory; Wayne Joubert, Xiylon Software; Bryan Lally, CCS-2

The purpose of this article is to compare geometric multilevel preconditioners tailored for adaptive mesh refinement (AMR) grids (FAC, AFAC, AFACx) with algebraic multigrid (AMG) methods on structured AMR (SAMR) grids.

In order to obtain locally refined grids that are representative of what is encountered in practice, the preconditioners are evaluated within the context of a fully dynamic adaptive simulation [1] of the propagation of radiation or thermal energy using a diffusion approximation.

Each time-advanced solution for the radiation diffusion problem is found using the Jacobian-free Newton-Krylov method in PETSc [2]. Each time step thus requires a sequence of linear system solves. Preconditioning these linear systems requires a diffusion solve [1] for which we compare FAC, AFAC, AFACx, and AMG. SAMRAI [3] is used to manage the complexity of dynamic locally refined grids.

Table 1 shows the iteration counts obtained for FAC. Comparable iteration count results were also obtained for AFACx and LAMG (see [4]).

**Table 1:** The average number of linear iterations per time step is shown. The first column gives the size of the base grid; J is the number of refinement levels. Performance at a fixed finest resolution is obtained by reading diagonally from lower left to upper right. Grid configurations that were not run are denoted by --.

**Table 1: Summary of FAC iterations.**

J	1	2	3	4	5
32 x 32	--	9.7	9.8	10.3	11.6
64 x 64	--	15.1	14.8	14.1	--
128 x 128	7.3	9.0	10.4	--	--
256 x 256	7.5	10.0		--	--
512 x 512	8.3	--	--	--	--

Figure 1 plots the wall-clock time spent in the FAC preconditioner for different base grid configurations. Similar results

were obtained for AFACx (see [4]). The smaller problem sizes due to AMR result in fairly flat graphs as the number of processors is increased beyond 8. Communication costs do not appear to dominate the computations as the number of processors is increased.

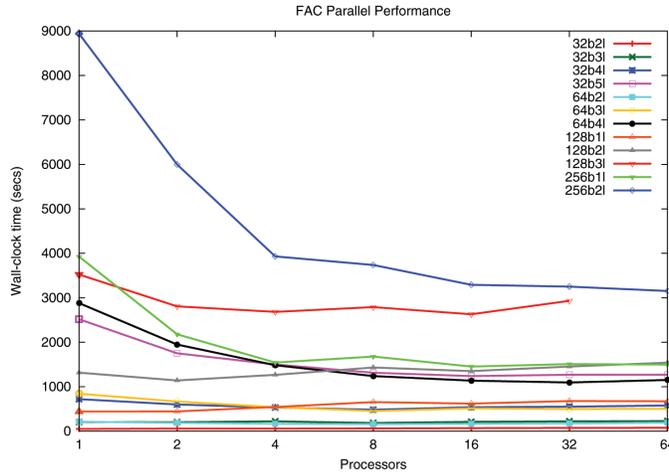
Figure 2 compares the parallel performance of FAC and AFACx for a fixed AMR configuration. Similar plots are obtained for other configurations also (see [4]). AFACx was not optimized or load balanced in a manner that would exploit the asynchronous nature of the algorithm, explaining the similar performance obtained. A possible optimization in the future would be to redistribute the load within an asynchronous preconditioning step.

Figure 3 compares the parallel performance of AMG, FAC, and AFACx on two cases for which we have complete data. In general, AMG requires approximately twice the execution time as FAC, and AFACx, for the model problems tested. We note that the setup phase for AMG, which requires remapping and ordering the grid hierarchy introduces an overhead of roughly 10%. None of the approaches is especially scalable in these evaluations; this can be ascribed in part to the small problem sizes that are made possible through the use of adaptive mesh refinement.

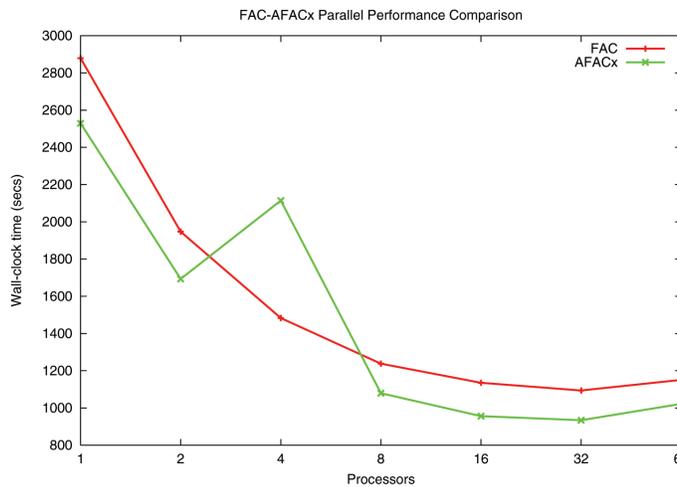
For more information contact  
**Bobby Philip** at  
**bphilip@lanl.gov**.

[1] M. Pernice, B. Philip, *SIAM J. Sci. Comput.* 27 (5), pp. 1709–1726 (2006).  
 [2] S. Balay, et al., "PETSc Users Manual," ANL-95/11, Argonne National Laboratory.  
 [3] R. Hornung, S. Kohn, *Concurrency Comput.: Pract. Exp.* 14, pp. 347–368 (2002).  
 [4] M. Pernice, et al., "A Comparison of Multilevel Preconditioners for Solving Multimaterial Equilibrium Radiation Diffusion Problems on Locally Refined Grids," Los Alamos National Laboratory report LA-UR-06-7231.

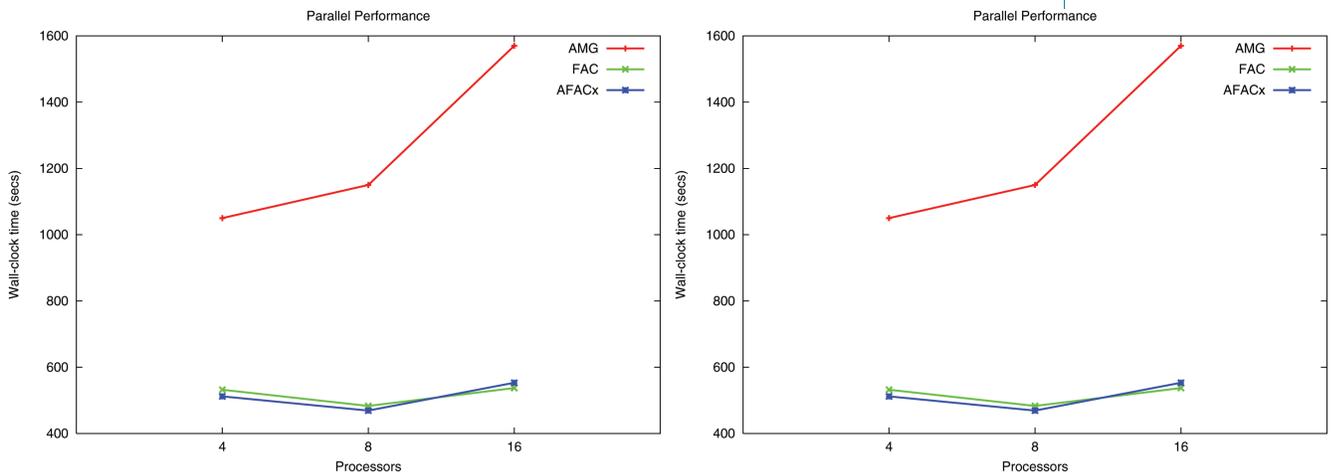
**Funding Acknowledgements**  
 NNSA's Advanced Simulation and Computing (ASC), and Weapons Supported Research (WSR) Program.



**Fig. 1.** Parallel performance of FAC with varying base grid resolution and number of refinement levels.



**Fig. 2.** Comparison of parallel performance for FAC and AFACx for an AMR configuration with a 64 x 64 base grid and four levels of refinement.



**Fig. 3.** Comparison of parallel performance of the various multilevel solvers. The plot on the left is for the case with a 32 x 32 base grid and two additional refinement levels; the plot on the right has the same base grid and three refinement levels.

# An Efficient, Numerically Stable, and Scalable Parallel Tridiagonal Solver

Christian Ketelsen, University of Colorado; Markus Berndt, J. David Moulton, T-7

**L**arge tridiagonal systems of linear equations appear in many numerical analysis applications. In our work, they arise in line relaxations needed by robust multigrid methods, such as the parallel BoxMG code [1], for structured grid problems. We present a new numerically stable and scalable algorithm for the solution of diagonally dominant tridiagonal linear systems of equations that scales well on distributed memory parallel computers. Its multilevel design makes it well suited for distributed memory parallel computers with very large numbers of processors.

## Background

On a serial computer, a diagonally dominant tridiagonal system of linear equations can be solved in “order  $N$ ” steps, by using Gaussian elimination and taking advantage of the special sparsity pattern of the linear system. This resulting algorithm is commonly referred to as the Thomas algorithm.

The first parallel algorithm for the solution of tridiagonal systems was developed by Hockney and Golub. It is now usually referred to as cyclic reduction. Stone introduced his recursive doubling algorithm in 1973. Both cyclic reduction and recursive doubling are designed for fine-grained parallelism, where each processor owns exactly one row of the tridiagonal matrix. In 1981, Wang proposed a partitioning algorithm aimed at more coarse-grained parallel computation typical of shared memory clusters, where the number of processors is much smaller than the number of unknowns. There has also been attention directed toward a

parallel partitioning of the standard LU decomposition algorithm. In 1986, Sun et al. introduced the parallel partitioning LU algorithm that is very similar to Bondeli’s divide and conquer algorithm. These algorithms, while well suited for problems distributed across a moderately large number of processors, do not scale well to very large numbers of processors.

## Algorithm

The tridiagonal linear system is assumed to be distributed across a large number of processors, such that each processor owns a contiguous number of rows (see the top matrix schematic of the figure). Each processor generates from that two interface equations that link neighboring processors (indicated by the red rows in the bottom matrix schematic of the figure). This is done efficiently by applying a Cholesky decomposition to the part of the tridiagonal system that is interior to each processor (indicated by the blue matrix entries of the bottom matrix schematic in the figure). The decomposition can be used to generate the upper, as well as the lower interface equation. The new system of interface equations has properties that are similar to the original tridiagonal system. These interface equations are now assigned in large enough groups to a subset of processors, and the resulting tridiagonal system of linear equations is now solved recursively by computing new lower level interface equations. We proceed with further recursion if the subset of processors is sufficiently large, or solve the new interface system directly on one of the processors in the subset. Once this lowest level set of interface equations is solved, its solution is communicated to the processors

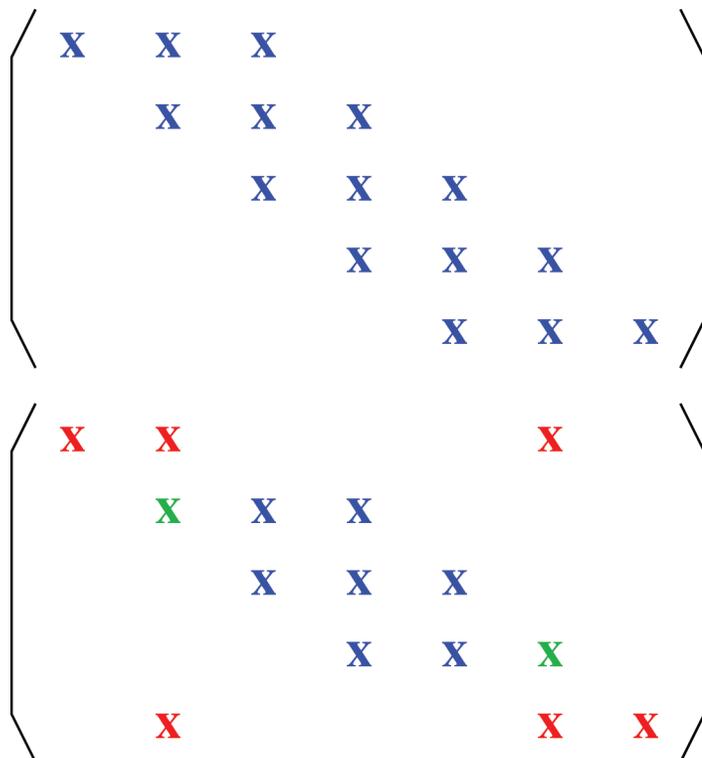
that own the next level up, which then can solve their higher level interface equations, and then on to the highest level, which is the original tridiagonal system of equations. Interestingly, the Cholesky decompositions that were computed along the way to generate each of the interface equations can be re-used now to solve them.

The strength of our new algorithm is that standard numerical components, such as Cholesky decompositions, are used for solving reduced systems of linear equations on each processor. This ensures that the resulting algorithm is numerically stable and easily maintainable. The resulting algorithm can be shown to be numerically stable and to scale well in parallel across very large numbers of processors.

*For more information contact Markus Berndt at [berndt@lanl.gov](mailto:berndt@lanl.gov).*

[1] T.M. Austin, et al., "Parallel, Scalable, and Robust Multigrid on Structured Grids," Los Alamos National Laboratory report LA-UR-03-9167 (2003).

**Funding Acknowledgements**  
 NNSA's Advanced Simulation and Computing (ASC), Advanced Applications Program.



**Fig. 1.** Schematic of a tridiagonal matrix owned by one processor before (top) and after the generation of interface equations (bottom).

# Benchmarking Diffusion Discretizations on Polyhedral Meshes

*J. David Moulton, Konstantin Lipnikov, T-7; Jimmy Fung, Scott Runnels, X-3*

**T**he flexibility of polyhedral meshes, in conjunction with recent advances in robust diffusion discretizations, has created significant interest from both the computational science community and ASC-related applications. Specifically, polyhedral meshes simplify grid generation for complex geometries, readily treat cell- or patch-based adaptively refined meshes (AMR) as degenerate polyhedra, and simplify mesh reconnection algorithms for moving meshes. Simulations of fluid flow in complex geometries indicate that a polyhedral mesh may provide a more accurate solution than a tetrahedral mesh with the same number of elements. Recent progress in Mimetic Finite Difference (MFD) methods has produced a number of accurate and robust discretizations of the time-dependent diffusion equations on polyhedral meshes [1].

Although, these new discretizations lead to linear systems of equations that may be treated by state-of-the-art multilevel linear solvers, issues of efficiency and robustness must still be characterized and addressed. The objective of this work is to engineer an integrated, performance-oriented computational method for the diffusion equation on polyhedral meshes through the development of an automated benchmarking tool.

## Discretization on Polyhedral Meshes

We consider the steady-state diffusion equation written in a conservative form:

$$\mathbf{u} = -D\nabla p \quad \text{and} \quad \text{div } \mathbf{u} = b,$$

where  $D$  is the symmetric diffusion tensor, and  $b$  is the source term. We shall refer to  $p$  as the pressure and  $\mathbf{u}$  as the velocity. We consider three MFD methods that use the same degrees of freedom: two unknowns per mesh face (average normal velocity and average pressure) and one unknown per mesh zone (average pressure). The KLS [2] and BLS [3] methods result in a symmetric indefinite linear system. In contrast, the Morel method [4] results in a nonsymmetric indefinite linear system. In [1], we explain how the elimination of velocity and zone-based pressure unknowns lead to a smaller system involving face-based pressures. We use a Krylov iteration preconditioned with the Los Alamos algebraic multigrid (LAMG) to solve these systems.

## Numerical Results

Polyhedral mesh generation, discretization methods, and multilevel iterative solvers are all under active development, while the underlying implementations may require additional tuning on new architectures. Consequently, we have developed a flexible testing framework in Python (a modern, easy-to-read, object-oriented scripting language), to perform the setup, execution, and report creation for these diffusion problems.

In this study we focus on bottom-line performance, with points appearing closer to the origin corresponding to lower error and less overall computation time. The first case uses a severely distorted structured mesh, dubbed the Kershaw mesh (Fig. 1). This highlights

the higher construction cost of the KLS method, as well as its lower accuracy, with an error more than 6 times greater than BLS and Morel. The second case (Fig. 2) uses a spherical mesh with AMR style refinement leading to many degenerate polyhedral zones. Here, the setup cost plays a significant role in KLS trailing as its accuracy at a given resolution is between Morel and BLS. Overall the BLS discretization appears to provide the best balance of discretization error and solution time for the meshes and solvers considered thus far.

*For more information contact J. David Moulton at [moulton@lanl.gov](mailto:moulton@lanl.gov).*

[1] J.D. Moulton, et al., "Discretization Schemes on Polygonal and Polyhedral Grids for Diffusion Problems." Los Alamos National Laboratory report LA-UR-07-1588.

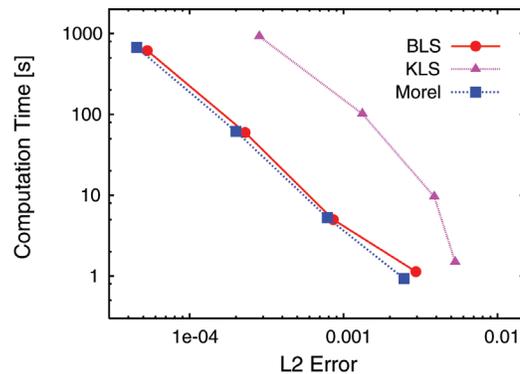
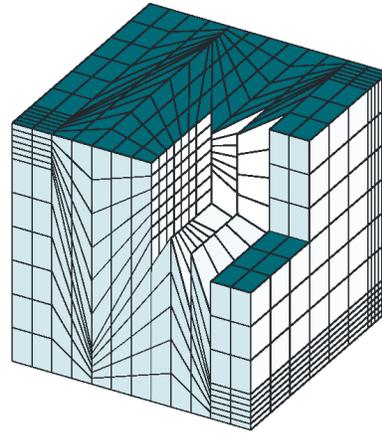
[2] Y. Kuznetsov, et al., *Comput. Geosci.* 8 (4) 301–324 (2004).

[3] F. Brezzi, et al., *Math. Mod. Meth. Appl. S.* 15 (10) 1533–1551 (October 2005).

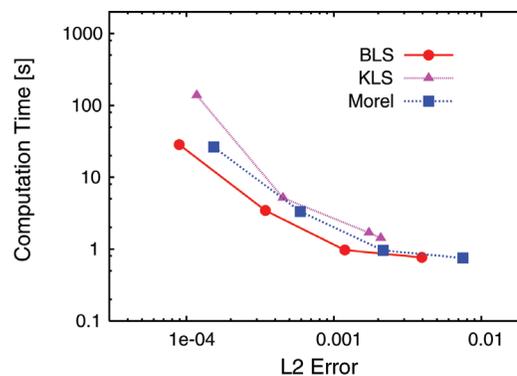
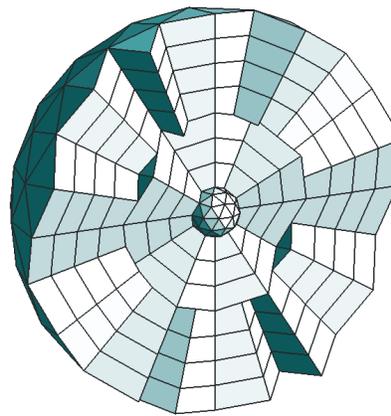
[4] K. Lipnikov, et al., "A New Discretization Scheme on Polyhedral Grids for Diffusion Problems," Los Alamos National Laboratory report LA-UR-06-3727 (2006) (submitted to *SIAM J. Numer. Anal.*).

### Funding Acknowledgments

NNSA's Advanced Simulation and Computing (ASC), Advanced Applications Program.



**Fig. 1.** A series of  $N \times N \times N$  logically structured meshes ( $N = 10, 20, 40, 80$ ), with Kershaw style perturbations were used in this study. A  $12 \times 12 \times 12$  Kershaw mesh with a cutaway of zones lying near the front-right-top corner is shown at the top. The achieved accuracy, measured in the  $L_2$  norm of the error in zone-based pressures, vs total CPU time is shown in the bottom plot. The BLS and Morel discretizations are significantly more efficient than the KLS discretization.



**Fig. 2.** A series of locally refined meshes consisting of prismatic zones with five faces and refined zones with eight faces were used in this study. The zones have approximately an even volume distribution. A cutaway view of a sample mesh is shown at the top. The achieved accuracy in the zone-based pressure, measured in the  $L_2$  norm, vs total CPU time is also shown in the bottom plot.

# Direct Search Methods with Equality Constraints

David W. Dreisigmeyer, HPC-4

**D**irect search methods are derivative-free algorithms for function minimization. They are popular optimization methods for nonsmooth, noisy, and discontinuous. Additionally, a user may not have access to or trust the available derivative information. One major attraction of derivative-free algorithms is the ease with which they can be coded. A typical optimization problem for a direct search method is (P1):

$$\begin{aligned} \text{minimize :} & \quad f(x) \\ \text{subject to :} & \quad g(x) \leq 0. \end{aligned}$$

We'll take the inequality constraints in (P1) as defining some full dimensional region  $\Omega$ . The way that direct search methods typically proceed is to lay out some mesh around the current iterate in  $\Omega$ . The function  $f(x)$  is then evaluated at the points on the current mesh.

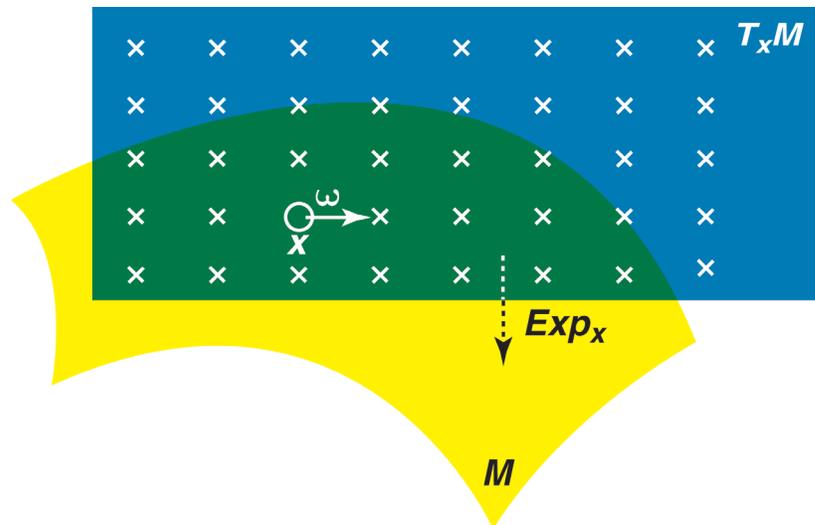
A major difficulty with direct search methods is their ability to handle problems of the form (P2):

$$\begin{aligned} \text{minimize :} & \quad f(x) \\ \text{subject to :} & \quad g(x) \leq 0 \\ & \quad h(x) = 0. \end{aligned}$$

The reason is that the equality constraints define a subspace  $M$  of  $\Omega$  that is vanishingly small. The probability of a mesh point actually lying on  $M$  and satisfying the equality constraints is 0. The way around this is to treat  $M$  as a Riemannian manifold, at least locally [1]. Then, if we start with a point on  $M$  we can efficiently stay on  $M$  as the direct search method proceeds.

What is actually done is to restate (P2) as a problem in the tangent space of some point on  $M$ . Each vector  $\omega$  in the tangent plane will correspond to a unique point  $y$  on  $M$ . We assign to  $\omega$  the function values  $f(y)$  and  $g(y)$ . That is, we pullback the objective function and inequality constraints from  $M$  to the tangent plane. But the tangent plane is just a copy of some Euclidean space, so we can now employ standard direct search methods in the tangent plane. The situation is illustrated in Fig. 1.

**Fig. 1.** General setup for performing a direct search over a manifold. The  $Exp_x$  function is a mapping from the tangent space  $T_x M$  to the manifold  $M$ .



Our technique is especially useful if  $M$  is a Lie group or some other well-understood manifold [2]. In this case, closed form solutions are available for the mapping from  $\omega$  to  $y$ . However, even when  $M$  is defined as a level set, finding the mapping from the tangent space to  $M$  has proven feasible [1, 3].

[1] D. W. Dreisigmeyer, "Equality Constraints, Riemannian Manifolds and Direct Search Methods," submitted to *SIAM J. Optimiz. (SIOPT)*.

[2] D. W. Dreisigmeyer, "Direct Search Algorithms over Riemannian Manifolds," submitted to *SIAM J. Optimiz. (SIOPT)*.

[3] D. W. Dreisigmeyer, "Direct Search Methods over Lipschitz Manifolds," submitted to *SIAM J. Optimiz. (SIOPT)*.

*For more information, contact David Dreisigmeyer at [dreisigm@lanl.gov](mailto:dreisigm@lanl.gov) or visit Data-Driven Modeling and Analysis (DDMA) Team Web site, <http://ddma.lanl.gov/>*

#### **Funding Acknowledgements**

This research was supported by the NNSA tri-Lab Advanced Simulation and Computing Program.

# Numerical Schemes for Hydrodynamics Based on Multidimensional Riemann Solvers

William W. Dai, HPC-4; Paul R. Woodward, B. Kevin Edgar, University of Minnesota

**N**umerical schemes based on Riemann solvers have been widely used in applications of hydrodynamics. Typically these schemes are using 1-D Riemann solvers at interfaces between numerical cells. In Lagrangian or arbitrary Lagrangian Eulerian (ALE) calculations, it is always desired to obtain the flow velocities at grid points. Truly multidimensional high-order shock-capture schemes could play an important role in Advanced Simulation and Computing (ASC) code projects. Approximate multidimensional Riemann solvers, which solve Riemann problems at grid points instead of interfaces between cells, naturally serve the purpose.

Multidimensional Riemann problems have not been mathematically solved yet. Therefore, developing numerical schemes based on multidimensional Riemann solvers is still a challenge. But the hope is that all the 1-D Riemann solvers actually used in numerical schemes are approximate, but not exact. We expect approximate multidimensional Riemann solvers to play an important role in future numerical simulations for hydrodynamics.

In this paper, we briefly outline a numerical scheme based on an approximate multidimensional Riemann solver. In the scheme, we use an approximate solver for 2-D and 3-D Riemann problems to calculate time-averaged values of flow variables at each grid point. Time-averaged fluxes at interfaces between numerical cells are calculated from the time-averaged values of the variables. The scheme is

truly multidimensional, is second-order accurate in both space and time, and satisfies conservation laws for mass, momentum, and total energy exactly.

## Numerical Schemes

The set of 2-D Euler equations is

$$\rho \frac{dU}{dt} = \frac{\partial F_x}{\partial x} + \frac{\partial F_y}{\partial y}. \quad (1)$$

Here,  $U \equiv (v, u_x, u_y, e)^T$ , superscript  $T$  stands for transpose,  $\rho$  is mass density,  $v \equiv 1/\rho$ ,  $u_x$  and  $u_y$  are the components of flow velocity,  $e$  is the specific total energy, and  $F_x$  and  $F_y$  are fluxes in the  $x$ - and  $y$ - directions respectively. A 2-D Riemann problem is Eq.(1) with a set of constant states for each region surrounding a point; for example, four constant states in the four quadrants in a structured mesh. At each time step within a simulation, what are known are mass, momentum, and energy in each numerical cell. The initial condition surrounding each grid point naturally corresponds to a 2-D Riemann problem.

If Eq.(1) is integrated over a cell and one time step,  $0 < t < \Delta t$ , the following equation will be obtained.

$$\langle U \rangle = \langle U \rangle_0 + \frac{\Delta t}{\Delta m} \{ \overline{\oint F_x dy} + \overline{\oint F_y dx} \}. \quad (2)$$

Here,  $\langle U \rangle$  is a space-averaged value of  $U$  over the cell at  $t = \Delta t$ ,  $\langle U \rangle_0$  is its initial value,  $\oint$  is the integral counterclockwise along the perimeter of the cell, and the bar over the integrals stands for the time average during the time step. In our scheme, the time-averaged integral is approximately calculated through the time-averaged values obtained from an approximate multidimensional Riemann solver. To get the second-order accuracy of the scheme, the states surrounding a

grid point will not be the states of the cells, but are the states on domains of dependence.

### Results and Discussions

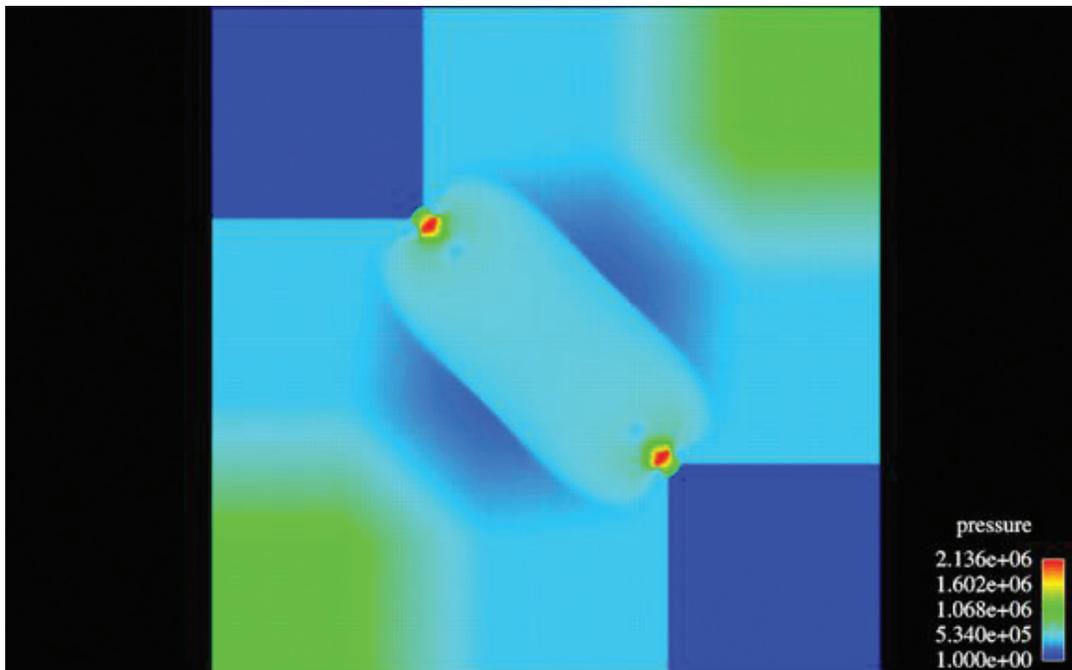
The scheme has been tested in an image processing program, ALE, calculations (Fig. 1). Currently, artificial viscosity has not been introduced into the scheme, but some form of artificial viscosity may be necessary for real equation of state when very strong shocks are present. In the future, we plan to extend the scheme for unstructured meshes.

Figure 1 shows the pressure in an ALE calculation at  $t = 0.0002$  for a 2-D Riemann problem. The initial pressures on the four quadrants are  $10^6$ , 1.0,  $10^6$ , 10 respectively. The initial velocity is zero, and initial density is unity everywhere.

*For more information contact William W. Dai at [dai@lanl.gov](mailto:dai@lanl.gov).*

### Funding Acknowledgements

This research was supported by the NNSA tri-Lab Advanced Simulation and Computing Program.



**Fig. 1.** ALE calculation at  $t = 0.0002$  for a 2-D Riemann problem. The initial pressures on the four quadrants are  $10^6$ , 1.0,  $10^6$ , 10 respectively. The initial velocity is zero, and initial density is unity everywhere.

# Tensor Multiplication on Parallel Computers

Bryan Rasmussen, HPC-4

**T**ensor multiplication is a ubiquitous task in many scientific and engineering applications. While there has been a great deal of research on efficient numerical methods for matrix multiplication on parallel machines, the general multiplication of arbitrary-rank tensors with an arbitrary number of contractions has seen less progress.

We consider a new set of C++ classes that compute, store, and multiply tensors on both serial and parallel platforms. The original, motivating application is in computational chemistry, for example [1], although the work should apply to many disparate areas such as fluid mechanics, general relativity, and quantum mechanics [2].

For the purposes of this project, a “tensor” is simply a multidimensional box of numbers, stored in a logical order. A rank-0 tensor is a scalar, a rank-1 tensor is a vector, etc. The code performs tensor multiplications of the form

$$w_{a_1 a_2 \dots a_m b_1 b_2 \dots b_n} = u_{a_1 a_2 \dots a_m c_1 c_2 \dots c_p} v_{b_1 b_2 \dots b_n c_1 c_2 \dots c_p}, \quad (1)$$

where  $m$ ,  $n$ ,  $p$  and their associated dimensions are completely arbitrary. The multiplication in Equation (1) assumes the Einstein summation convention which is to say that we sum over repeated indices. (For example, the inner product of two vectors is  $u_i v_i$ .)

We also assume that every tensor has a very compact storage method called the  $k$ -index transformation, also known as the Tucker transformation. This transformation says that a rank- $k$  tensor,  $v$ , has a “core tensor,”  $A$ , and characteristic matrices  $z^1, z^2, \dots, z^k$  such that any element of the  $v$  has the form

$$v_{i_1 i_2 \dots i_k} = z^1_{i_1 j_1} z^2_{i_2 j_2} \dots z^k_{i_k j_k} A_{j_1 j_2 \dots j_k}. \quad (2)$$

This transformation becomes very important for minimizing storage and communication costs in the parallel algorithm. The main savings comes from the fact that the core tensor,  $A$ , is usually known *a priori* from straightforward analytic expressions.

To create a parallel algorithm, we assign each processor a different piece of the three tensors in Equation (1). We assume that each processor has enough memory to store one row of  $u$ ,  $v$ , and  $w$  simultaneously, where a row is the piece of the tensor corresponding to a single, fixed index in the first position.

Each processor then creates only its portions of  $u$  and  $v$  one row at a time and multiplies them together to form rows of  $w$ . This approach has several advantages, namely that it requires very little message passing, and it minimizes the number of redundant formations and calculations.

One disadvantage of our approach is that it requires each processor to work on a large piece of the problem for a long time, thus increasing the probability that a single processor failure will sabotage the computation. Another limitation is that we must increase the number of processors in potentially large step-increments in order to take advantage of larger clusters. (The ideal number of processors is an integer multiple or divisor of the number of rows in  $u$ .)

A typical solution to these difficulties is to subdivide the problem and then use a scheduler to balance the workload. Unfortunately, small divisions of labor necessitate redundant calculations, so there is a limit to how we can chop up the problem. At present, the code does include a scheduler with simple rules on how to divvy the pieces of the tensors, but that

capability remains untested. Despite its simplicity (or perhaps because of it), the all-at-once method seems to be very robust, efficient, and scalable.

Consider the following results on the DataStar IBM Power4 computer at the San Diego Supercomputing Center. We experiment with the multiplication of two rank-4 tensors of equal dimensions. The tensors descend from  $k$ -index transformations where the core tensors consist of all ones, and the characteristic matrices all have five columns of uniform random numbers. The theoretical problem size therefore increases with the square of the number of rows,  $N_u$ . Figure 1 shows computation time as a function of  $N_u$ . Each line represents a constant number of processors.

The tensor classes are still under active development, particularly with regard to the  $k$ -index transformation. Through judicious storage of partial calculations, it is possible to reduce the complexity of Equation (2) to  $N^{k+1}$ , where  $N$  is the largest dimension of any of the three tensors. The naïve way of doing this often leads to temporary storage requirements that exceed the memory available to any current processor, and indeed to any processor likely to be produced in the next 20 years.

In theory, we can circumvent this restriction by fusing nests of loops—indeed, there is significant progress in this area—but optimal loop fusion is a complicated process for arbitrary-rank tensors. Other research groups have developed algorithms for generating code automatically to suit specific requirements and platforms [3]. We aim instead for a more general approach that may not be optimal but still allows for increased productivity for practical problems.

Our near-term goal is an efficient, parallel, arbitrary-dimensional algorithm that integrates the  $k$ -index transformation with tensor multiplication to enable a wide variety of new calculations with only a

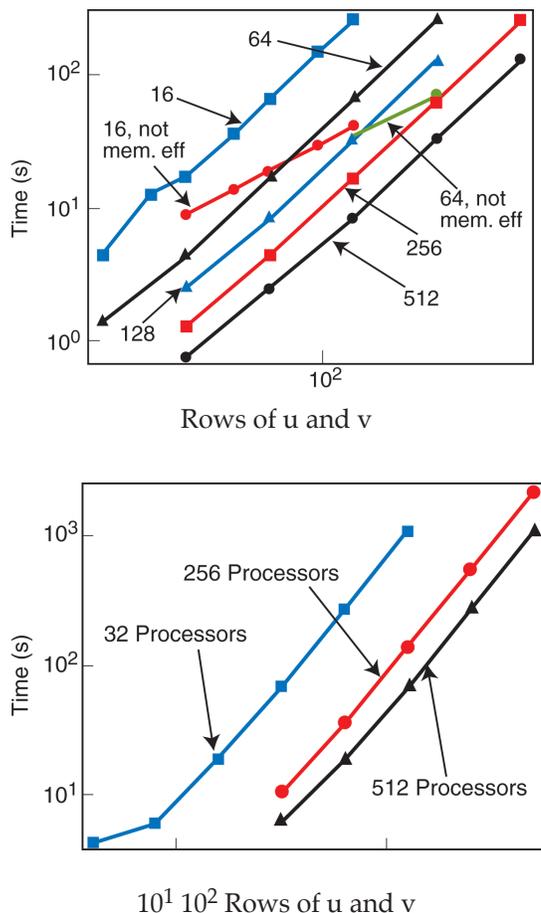
simple set of C++ classes. Preliminary results indicate that this is possible in the near future.

**For more information contact Bryan Rasmussen at [bryanras@lanl.gov](mailto:bryanras@lanl.gov).**

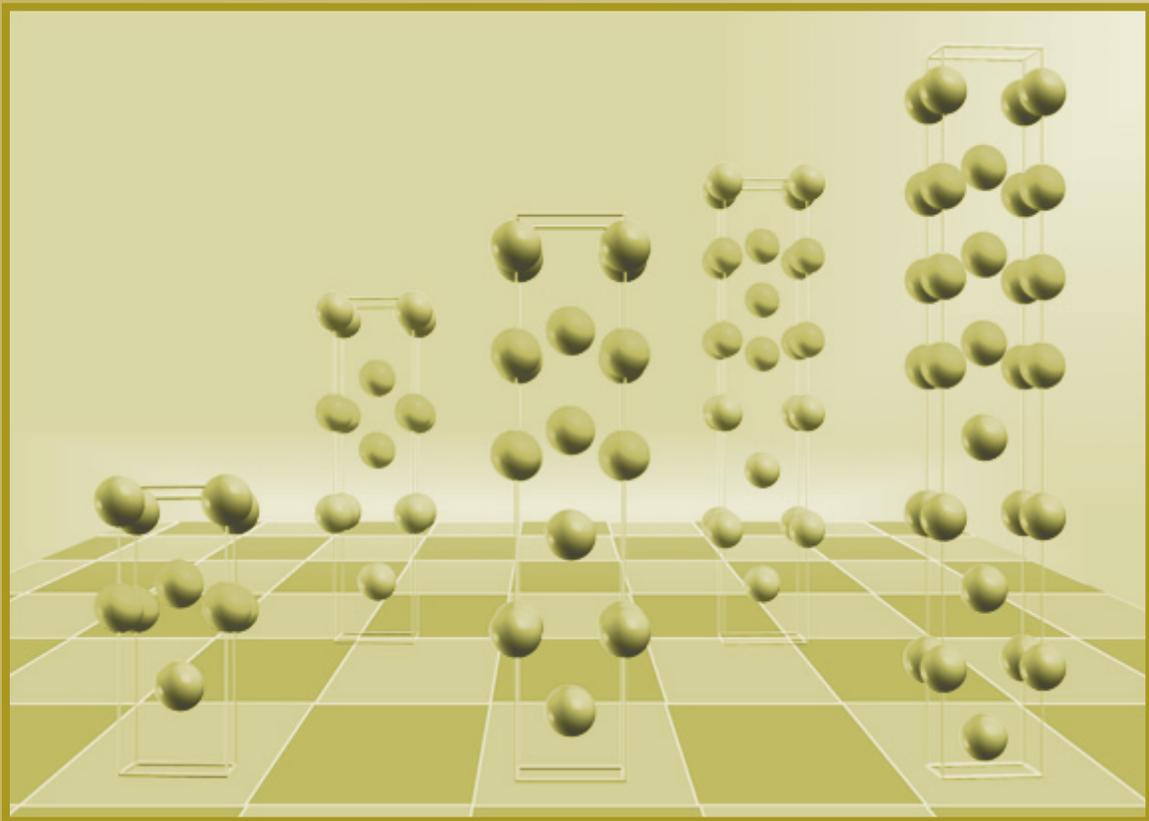
[1] A. P. Rendell, T. J. Lee, and R. Lindh, "Quantum chemistry on parallel computer architectures: Coupled-cluster theory applied to the bending potential of fulminic acid," *Chem. Phys. Lett.*, **194**, 84 (1992).  
 [2] A. J. McConnell, *Applications of Tensor Analysis*, (Dover, New York, 1957).  
 [3] G. Baumgartner, et al., "Synthesis of high-performance parallel programs for a class of *Ab Initio* quantum chemistry models," *Proceedings of the IEEE*, **93**(2), 276 (February 2005).

**Funding Acknowledgements**

This research was supported by the NNSA tri-Lab Advanced Simulation and Computing Program.



**Fig. 1.** Computation time vs  $N_u$  for  $N_u \times 16 \times 16 \times 16$  (top) and  $N_u \times 32 \times 32 \times 32$  (bottom). Scaling appears to be nearly perfect over the ranges studied. The lines labeled "not mem. eff" refer to a capability of the code to trade computation time for memory. This mode will not be practical for most applications due to memory constraints.



# Materials Science and Equation of State

This section includes discussions of current work on the equation of state of plutonium, tantalum, gallium, boron, boron carbide, vanadium, and lithium fluoride. In addition, some of the papers discuss some of the methodology involved in the generation and tabulation of equations of state for condensed materials. Other topics reported in this section deal with lattice vibrations, superlattices, property and constitutive models for binders for high explosives, polymeric materials, phase transformations in composite materials as well as magnetic transitions in the Kondo lattice.

## Evolution of the Equation of State Production Code OpenSesame

*Nicolas Bock, Eric Chisolm, Travis Peery, Scott Crockett, Carl Greeff, T-1*

OpenSesame is the next-generation code for production and management of equation of state (EOS) libraries that has been under development in T-1 for the last few years. This code incorporates the capabilities of its predecessor, GRIZZLY, which has been used with great success in producing EOS for the SESAME database, and it can also calculate and plot information using existing EOS. The code is based on a modular Fortran90 structure that should be easily extensible and maintainable. In the last year we have taken substantial steps toward the release of a reliable and stable Version 1.0.

For users, the most visible code management change is the switch to an updated OpenSesame webpage that can be found at <https://keiki.lanl.gov/opensesame/>. We will periodically release official versions of the code, which can be downloaded from this webpage. We have also simplified the OpenSesame build process by using the standard GNU autoconf/automake process. We support a variety of platform/compiler combinations and intend to add more as the need arises. This year we will implement and release a new graphical user interface (GUI), replacing the old Tcl/Tk-based GUI. The new OpenSesame GUI will be written in Java, making it platform-independent. This choice of language also enables us to release precompiled byte-code, which will make the installation process for users of OpenSesame much simpler.

We have also moved the source code to a subversion server, which will allow users to choose their level of involvement in the code development process. At the

simplest level, users can download the official stable source code from our webpage. At the next level, users with read access to the subversion server can check out the most recent revision directly. Finally, users who want to participate in the code development process can get write access to the subversion server and submit their own changes to the OpenSesame code. In conjunction with the subversion server, we have implemented a ticketing system (trac) that helps organize code development. It can be found at <https://keiki.lanl.gov/cgi-bin/trac.cgi>.

The ticketing system serves three purposes: 1) it allows users to report issues they find when running OpenSesame and track the progress on their resolution; 2) it allows for a more precise organization and handling of quality control issues; and 3) it allows for the explicit use of “milestones,” which define well-focused goals, can be easily tracked by users, and enhances the quality and efficiency of our development efforts.

As an additional step toward quality assurance, we will set up a build farm that will automatically build and verify the most recent revision of the OpenSesame source code. This allows us to quickly and reproducibly identify code regressions. An extensive suite of test inputs to OpenSesame has already been developed to verify the code’s functions; we will expand this suite to cover all major functions and incorporate checking against the suite into the build farm.

We hope that these steps will enhance the quality and productivity of Open-Sesame by allowing quicker and more efficient code development and distribution.

*For more information contact Nicolas Bock at [nbock@lanl.gov](mailto:nbock@lanl.gov) or Eric Chisolm at [echisolm@lanl.gov](mailto:echisolm@lanl.gov).*

**Funding Acknowledgements**

NNSA's Advanced Simulation and Computing (ASC), Materials and Physics Program.

## Lattice Vibrations of Tin Under Pressure

Sven P. Rudin, T-1

**L**attice vibrations provide an important contribution to equations of state (EOS) of materials important to the mission of Los Alamos National Laboratory. Lattice vibrations can in general be measured experimentally, but in practice this can be expensive or in some cases even dangerous.

Density Functional Theory (DFT) calculations of lattice vibrations present a path to improve EOS of materials under many conditions. A typical calculation costs significantly less than an experimental measurement, and materials that pose a risk in measurements can be treated on a computer just like harmless substances. Systematic calculations as a function of, for example, volume, become easily implemented.

In particular for systems at ambient pressure, experience shows that the DFT calculations of lattice vibrations agree well with experimental values. A collaboration with researchers of the High Pressure Science and Engineering Center (HiPSEC) at the University of Nevada, Las Vegas, indicates that this agreement holds for tin at higher pressures.

Figure 1 shows the lattice vibration dispersion calculated for tin in the body-centered cubic crystal structure with a volume corresponding to that found in experiments at a pressure of 64 GPa. The frequencies are plotted for sequences of wave vectors along high symmetry directions. Experiments can certainly measure such dispersions, i.e., the frequencies of particular wave vectors, but under pressure these measurements become much more difficult.

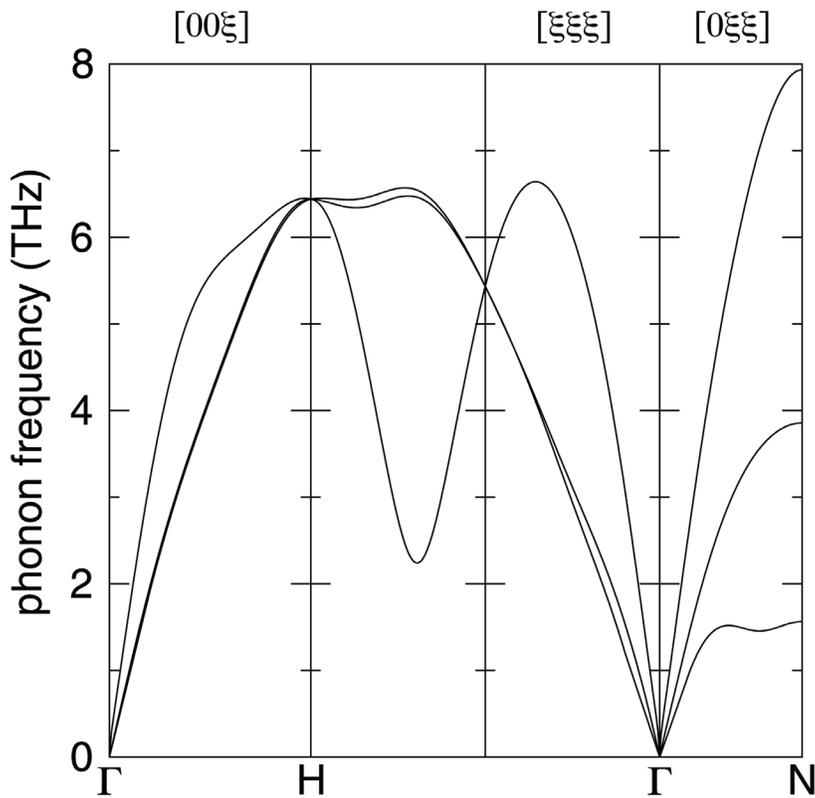
While still not a trivial task, the density of states (DOS) of the lattice vibrations of a crystal under pressure can be measured with less difficulty. Figure 2 shows the lattice vibrations DOS measured by HiPSEC researchers in excellent agreement with the calculated DOS. This agreement adds confidence to equations of state constructed using calculated lattice vibrations.

The collaboration with HiPSEC researchers continues with tin at other pressures where more complex crystal structures appear.

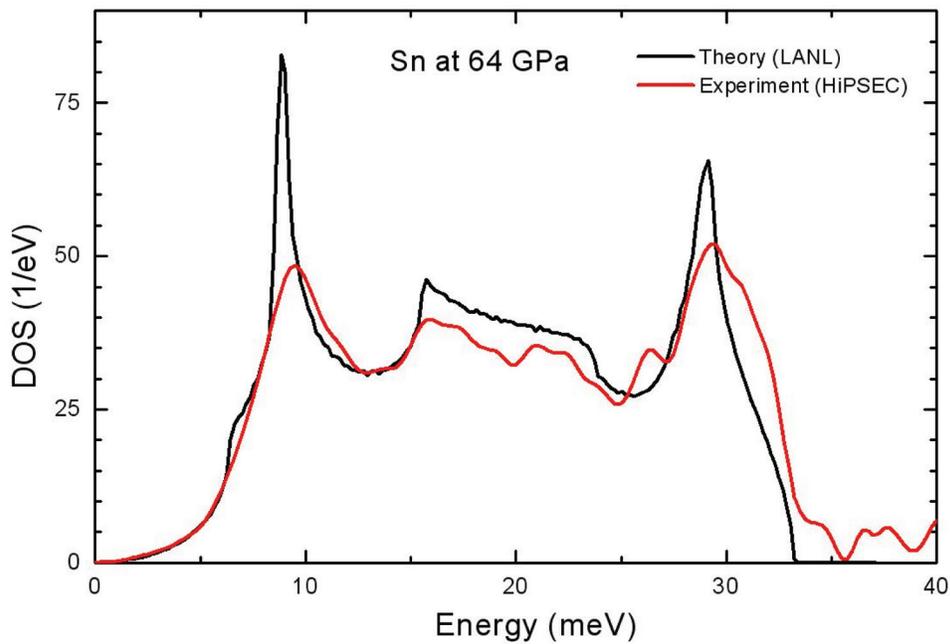
*For more information contact Sven Rudin at [srudin@lanl.gov](mailto:srudin@lanl.gov).*

### **Funding Acknowledgements**

NNSA's Advanced Simulation and Computing (ASC), Materials and Physics Program.



**Fig. 1.** Calculated lattice vibration dispersion of tin in the bcc crystal structure with volume given by experiment at a pressure of 64 GPa.



**Fig. 2.** Lattice vibrations density of states (DOS) of tin at a pressure of 64 GPa.

## Nanostructured Actinide Materials: Lead-Plutonium Superlattices

Sven P. Rudin, T-1

The patterning of materials with nanometer-scale structures can strongly affect physical properties and in some systems even induce them. The effects change as the structural parameters vary, making it possible to fine-tune desired (or undesired) attributes. Such a fine-tuning could yield improved scientific understanding and greater technological control of actinide material characteristics. These characteristics frequently relate to the nature of the 5f electronic states as they vary from itinerant to localized. Designing nanostructures to tailor the 5f electron localization would allow fine-tuning of the associated properties.

Density Functional Theory (DFT) calculations on lead-plutonium (Pb-Pu) superlattices indicate that these form nanostructures with two competing phases strongly related to the alpha and face-centered cubic (fcc) phases of bulk plutonium. Unlike the bulk phases the superlattice phases transform into each other by an easily described distortion, the pairing of Pu planes (see Fig. 1). This pairing occurs with a volume collapse and a delocalization of the 5f electrons along the lines of the Mott transition associated with the bulk phases. Furthermore, the bond lengths between Pu atoms in paired and unpaired planes take on values very close to the short and long bonds of alpha-Pu, respectively.

Figure 1 shows unit cells of Pb-Pu superlattice geometries with itinerant 5f electrons. The layers of Pb atoms remain highly symmetric with equal distances between planes. The layers of Pu atoms

distort with neighboring planes pairing up especially for structures with even numbers of monolayers.

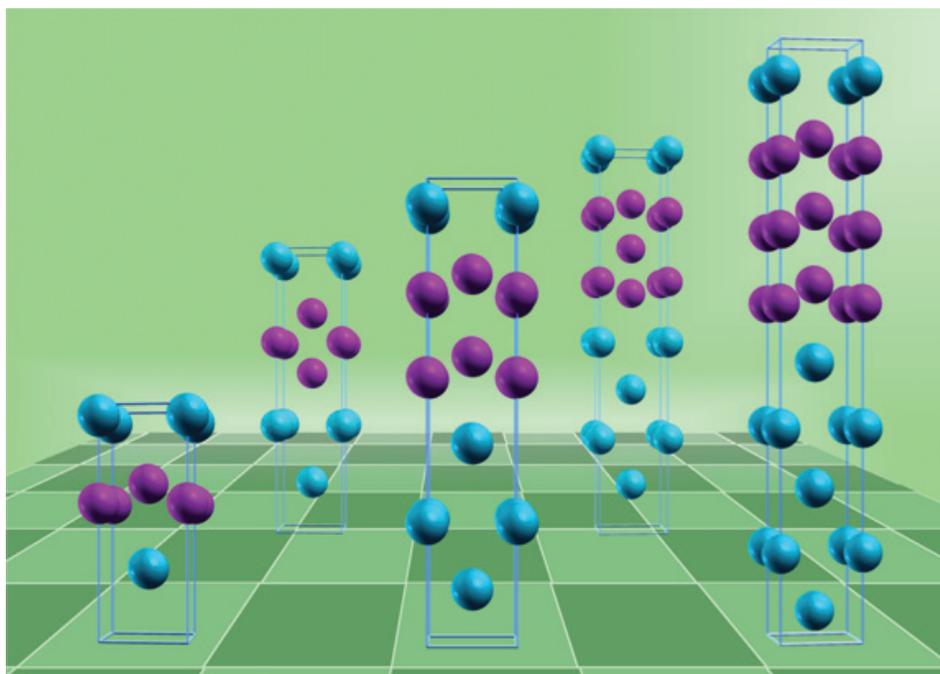
Figure 2 shows that the two superlattice phases appear equally favorable in energy for even numbers of monolayers, where for itinerant 5f electrons the Pu planes all pair up successfully. Calculations allowing localization of the 5f electrons let the system lower its energy without the distortions present in the itinerant case, resulting in a much less dramatic dependence of the energy on layer thickness. With increasing layer thickness the energy of the systems with localized 5f electrons tend towards the value obtained from averaging bulk fcc Pb and bulk fcc Pu (with localized 5f electrons). Similarly the volume shows little dependence on layer thickness.

The toxicity of both elements, the radioactivity of Pu, and shielding properties of Pb make Pb-Pu superlattices unlikely candidates for technological applications, but this system illustrates the rich potential in imparting actinide materials with nanostructures. The strong similarities between the two phases that appear in Pb-Pu superlattices and the alpha and fcc bulk phases of plutonium point to underlying principles that determine Pu's behavior as a solid. In particular the emergence of bond lengths between Pu atoms in paired planes that mirror the two ranges of the bond lengths found in alpha Pu is remarkable and evidence of one such fundamental principle.

For more information contact Sven Rudin at [srudin@lanl.gov](mailto:srudin@lanl.gov).

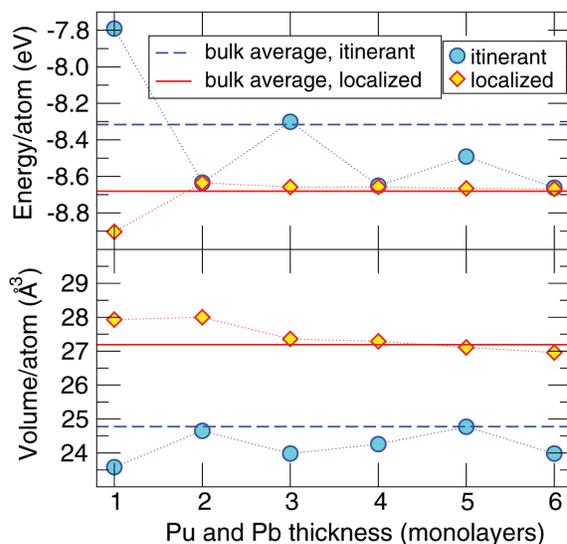
### Funding Acknowledgements

NNSA's Advanced Simulation and Computing (ASC), Enhanced Surveillance Program.



**Fig. 1.**

Unit cells of Pb-Pu superlattice geometries with itinerant 5f electrons. Pb atoms are colored teal, while the Pu atoms are colored magenta. The foreground shows structures with even numbers of monolayers. For odd numbers of monolayers, not all planes can pair up and these frustrated systems cannot lower their energy enough to be favored.



**Fig. 2.**

Energies and volumes for Pb-Pu superlattices as a function of layer thickness and dependent on whether the 5f electrons are treated as localized or itinerant. The average energies and volumes stem from separate pure fcc Pu and pure fcc Pb calculations.

## Tabular Multiphase Equation of State for Plutonium

Eric Chisolm, Carl Greeff, Shao-Ping Chen, Scott Crockett, J. D. Johnson, Sven Rudin, Duane Wallace, John Wills, T-1; Kevin Honnell, Jonathan Boettger, X-1-SMMP; David Pimentel, HPC-1

The dynamical behavior of a material can be significantly impacted by the effects of transformations from one phase to another. Not only do such transitions introduce latent heats and volume changes that can be sizable, but aspects of material behavior such as strength can be phase-dependent as well. The equation of state (EOS) of a material should indicate the equilibrium phase in any situation, so the production of multiphase EOS which provide phase boundaries and phase-specific thermodynamic information is of some interest.

The importance of including these effects is shown in Fig. 1, which includes various thermodynamic curves for tin as functions of pressure and temperature. The light blue dots indicate the phase boundaries as determined by SESAME 2161, a multiphase EOS for tin constructed recently [1]. The three solid lines are the principal Hugoniot (black) and two release isentropes (blue and red) determined by the same EOS; the three dashed lines are the same three curves as determined by SESAME 2160, an older tin EOS that does not treat the solid-solid transition. The blue isentropes terminate in the solid at zero pressure, while the red isentropes terminate in the liquid at that same pressure. Notice that the two principal Hugoniots are rather similar; in effect, the old EOS averages over the shifts in the Hugoniot that the new EOS attributes to the crossings of the phase boundaries. The isentropes, on the other hand, are quite different, and as a result

the two EOS give rather different predictions for the phase the system will relax into after a single shock. The old EOS predicts that the system will relax to pure solid after shocks below approximately 20 GPa and to pure liquid for shocks above about 27 GPa; after intermediate shocks the system will relax into a mixed solid-liquid state. The new EOS also places the lower limit at about 20 GPa, but the upper limit is much higher, at about 36 GPa. It is the inclusion of effects such as latent heat that significantly changes the isentropes, as the figure clearly shows, and thus the predictions for shock release.

We are currently using the multiphase techniques previously applied to tin to create a new EOS for plutonium. We have obtained model parameters for different phases of plutonium from experimental data and first principles calculations, and we have compared the results with additional experiments, leading to new predictions concerning the high-pressure phase diagram. Anomalous properties of the liquid have proved difficult to model consistently, and improving the liquid EOS will be a high priority for subsequent work. Details may be found in [2].

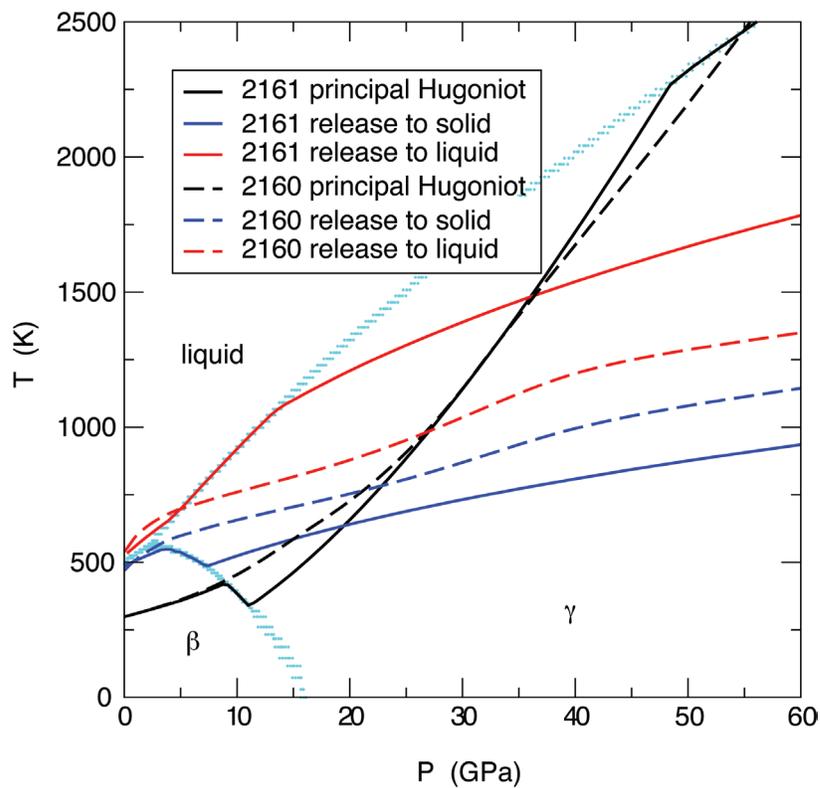
*For more information contact Eric Chisolm at [echisolm@lanl.gov](mailto:echisolm@lanl.gov).*

[1] Greeff, C.W., et al., "SESAME 2161: An explicit multiphase equation of state for tin," Los Alamos National Laboratory report LA-UR-05-9414 (December 2005).

[2] Chisolm, E.D., et al., "Tabular Multiphase Equation of State for Plutonium (U)," Los Alamos National Laboratory report LA-CP-06-1365 (December 2006), to be published in *Proceedings of the Nuclear Explosives Code Developers Conference* (Los Alamos, New Mexico, October 23–27, 2006).

### Funding Acknowledgements

NNSA's Advanced Simulation and Computing (ASC), Materials and Physics Program.



**Fig. 1.** Predictions for shock release to pure solid and pure liquid for SESAME 2161 compared to SESAME 2160. Please see the text for detailed explanation.

## Equation of State Developments in T-1

*T-1 Equation of State (EOS) Team*

**T**he Equation of State (EOS) and Mechanics of Materials Group (T-1) at Los Alamos National Laboratory conducts research in EOS to support the nuclear and conventional defense communities, civilian research communities, and commercial applications. The group maintains the SESAME EOS and Materials Properties Library.

### SESAME Library Update

In February we updated the unclassified SESAME library to include approximately eight new EOS, while the classified library will include five new EOS. The update will also include additional cold + nuclear sub-tables where possible and a large number of edits to the comment tables. We have also produced more than ten additional EOS for special applications.

### Multiphase and Multicomponent Equations of State

We have completed an ASC Level-2 milestone by delivering an improved capability for modeling multiphase materials, and we have also improved our models for plutonium.

- (a) New techniques have been developed for using Density Functional Theory (DFT) to determine EOS properties, and these techniques have been applied to phases of plutonium.
- (b) New models and algorithms for handling multiphase materials have been incorporated into the EOS-production code OpenSesame.
- (c) A new hydrocode interface capability giving users access to phase data has been put into the code EOSPAC.

Combining theory with experimental data, parameters have been determined for eight phases of plutonium, leading to new predictions concerning the high-pressure phase diagram. Anomalous properties of liquid plutonium have proved difficult to model consistently, and improving the liquid EOS will be a high priority in the coming

months, along with verifying the hydrocode implementation of the models.

We continue to develop a new EOS for carbon dioxide as part of the expansion of EOS capabilities within T-1 to multiphase and multicomponent systems. The proximity of carbon to oxygen in the periodic table allows us to use a simple average atom approach. The final EOS will be finished soon.

Beryllium is a very important technological material, specifically in aerospace and nuclear engineering, because of its light weight, high melting point, and large bulk and shear moduli. However, its phase diagram has remained virtually unknown over decades. In particular, its melting curve has never been measured beyond 6 GPa, and numerous empirical models have not converged on beryllium melting temperatures at high pressures. Using the Vienna *Ab Initio* Simulation Package (VASP), we calculated the melt curve to 10 Mbar using *ab initio* MD solid-liquid coexistence simulations. We studied three (bcc, hcp, and fcc) solid phases of beryllium using computational cells of different sizes. The three melting curves, each for a given solid phase, are essentially identical and form a melting curve of beryllium, which is in excellent agreement with the one predicted by the most recent SESAME Be EOS—SESAME 2024.

### Density Functional Theory

To complement experimental data and modeling in boosting the reliability of EOS, our DFT calculations continue to provide data, in particular cold curves and thermal contributions from the electrons and phonons. The reliability of these results is supported by collaborations with researchers of the High Pressure Science and Engineering Center at the University of Nevada, Las Vegas, whose measured phonon densities of states of tin under pressure agree well with results from our calculations.

### Vibration Transit Theory

More progress has been made on the Vibration-Transit (V-T) theory of liquid dynamics, both in the development of the theory and its application to actual systems. Through extensive quenching studies in simulated sodium, we have found further support for the main postulate of the theory, the equivalence of the atomic configurations (called random valleys) in a monatomic liquid. Explicit expressions within the theory have been derived for the vibrational contributions to the density and current autocorrelation functions. These quantities are directly related to the spectra measured by x-ray or neutron scattering. We have gained insight into the physics of these processes by studying the multimode expansions of the autocorrelation functions, and we have also studied the microscopic nature of low-temperature transits. In collaboration with the University of Trento, V-T theory is being applied to water.

We are developing a set of new tools based upon the tessellation of real space into Voronoi polyhedra (Wigner-Seitz cell-like constructions) as part of the ongoing elucidation of V-T theory. Within this theory, liquid state properties are derived from the random valleys and transitions between them. An efficient, general characterization of these random valleys is crucial for both the verification and use of V-T theory. By applying several localized Shannon entropies as well as other statistical tools to the geometry and topology of each atom's Voronoi polyhedron, we are able to effectively probe the local structure of the system to detect various types of localized order or symmetry in real space. Our suite of Voronoi-based tools will thus aid in the identification and classification of random valleys, simplifying the general application of V-T theory.

We have succeeded in combining V-T theory with DFT methods to accurately predict liquid state properties of sodium at ambient pressure, such as the density of the liquid at melt, its bulk modulus, entropy, and internal energy. This approach is as accurate but computationally significantly cheaper than a Carr-Parinello or Born-Oppenheimer

molecular dynamics (MD) simulation. We will submit these results for publication.

### OpenSesame

We have improved the capabilities of the EOS-production code OpenSesame in the following ways:

- (a) We have greatly simplified the build process by using standard GNU utilities and support a variety of platform/compiler combinations, expanding the list as users require.
- (b) We have implemented an issue-tracking system that will allow users to report issues and follow the progress of their resolution.
- (c) We have moved the source code to a subversion server, which allows users to participate in the code development process to whatever degree is desired.
- (d) We have updated the OpenSesame home page, <https://keiki.lanl.gov/opensesame>, to allow users to download source code directly.

Our goal this year is to release a stable version 1.0 and then begin work on two longer-term projects:

- (a) producing a new Java-based graphical user interface, which will ultimately absorb the entire code, to provide superior portability and even further simplify code installation, and
- (b) setting up a build farm that will build and verify the most recent version of the code, incorporating the previously developed validation suite.

*For more information contact Eric Chisolm at [echisolm@lanl.gov](mailto:echisolm@lanl.gov).*

### Funding Acknowledgements

NNSA's Advanced Simulation and Computing (ASC), Materials and Physics Program.

## Progress on the Campaign 2, Level-2 Milestone: “Kel-F 800 Experimental Characterization and Model Development”

Brad Clements, T-1

**K**el-F 800 is the binding polymer for the explosive PBX-9502. The project “Polymer Behavior under Dynamic Loading,” led by Project Leader Brad Clements under Program Manager S. Bingert, Los Alamos National Laboratory (LANL) Associate Directorate for Weapons Physics, of the Joint DoD/DOE Munitions Technology Development Program, is tasked with a Level-2 milestone deliverable for Campaign 2. The project members consist of a team of polymer and materials experts in LANL’s Theoretical (T), Materials Science and Technology (MST), Weapons Engineering Technology (WT), and Dynamic and Energetic Materials (DE) divisions. The deliverable date is October 2007. The deliverables include:

- I. Experimental data including stress-strain measurements over a wide range of temperatures and rates (compressive and tensile), Split Hopkinson Pressure Bar measurements, plate-impact measurements, Differential Scanning Calorimetry measurements, Dynamic Mechanical Analysis, specific volume measurements, heat capacity measurements, and fracture studies and analysis.
- II. A physics-based continuum-level constitutive model that adequately predicts the strain rate and temperature dependence of Kel-F 800 (as determined by experiments listed above).

- III. An assessment of the characterization properties of FK 800. (FK 800 is the Kel-F 800 replacement being manufactured by 3M Corporation.)

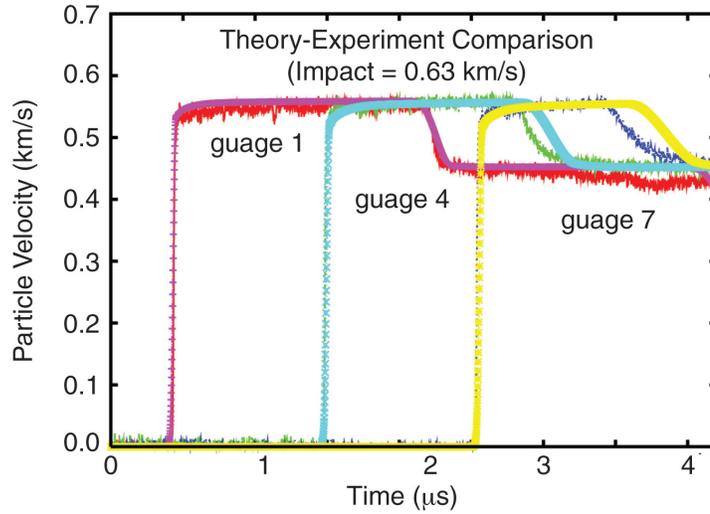
The model must adequately capture the correct equation of state and deviatoric response of Kel-F 800, will be implemented into the finite element code EPIC, and validation tests will be done. Kel-F 800 is a semi-crystalline polymer. We will report crystal volume fractions between 5% and 15%.

Good progress has been made in FY06 towards the completion of the milestone. In Figs. 1–3, experimental and theoretical information on high rate deformation (Fig. 1), the viscoelastic properties (Fig. 2), and uniaxial compression stress-strain measurements (Fig. 3) summarize some of the recent progress made by the team.

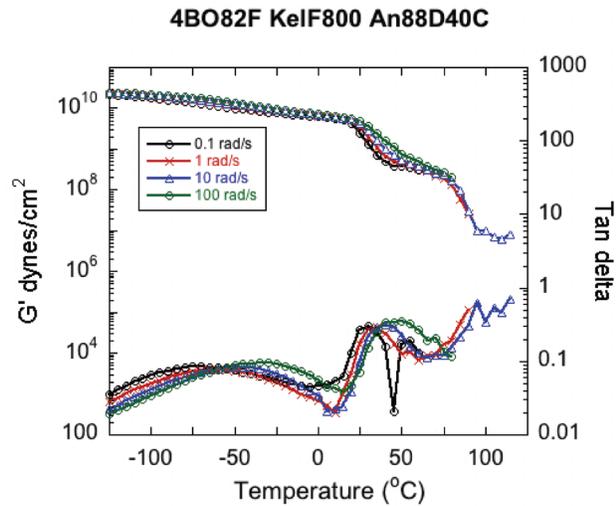
*For more information contact Brad Clements at [bclements@lanl.gov](mailto:bclements@lanl.gov).*

### **Funding Acknowledgements**

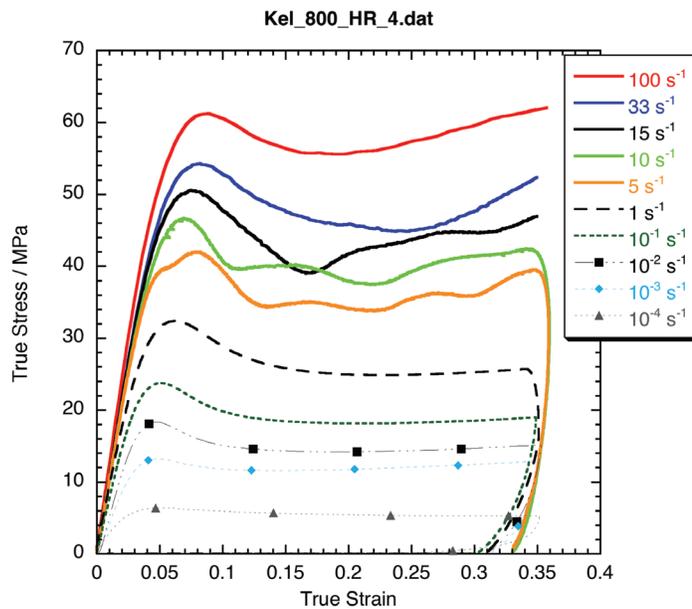
The Joint DoD/DOE MOU Munitions Technology Development Program.



**Fig. 1.** Experimental and theoretical particle velocity profiles for shocked Kel-F 800. The data are from DE Division (Dattelbaum, et al.).



**Fig. 2.** Storage modulus measurements from B. Orlor from MST-7. These measurements give information on the viscoelastic response of Kel-F 800.



**Fig. 3.** Stress-strain measurements from P. Rae and E. Brown of MST-8.

## Constitutive Modeling of Army Face Shields

Brad Clements, T-1

**P**olycarbonate (PC) is used for face-shield applications because of its good optical properties and its high impact-resistance characteristics. The US Army is seeking to better understand the thermo-mechanical properties of PC for its use in transparent armor applications. In collaboration with Army Research Laboratory (ARL) researchers, we have started developing a PC constitutive model. Our goal is to use ARL-measured experimental properties supplemented by published experimental results to construct an improved constitutive model for PC and filled-PC composites.

An amorphous polymer, PC is glassy below temperatures of approximately 423 K. Moy et al. [1] have determined the PC yield stress and obtained stress-strain curves for rates between  $10^{-5}$  to  $10^3$   $s^{-1}$ . The fracture properties and thermal heating of PC have been studied by Bjerke et al. [2]. Polycarbonate [2] was investigated in the glassy phase, and consequently strength loss occurred primarily by brittle failure. We will rely on the work of Ref. 1 and 2 in our model development. Molecular cooperativity theory is used to construct a rate- and pressure-dependent flow stress model. We incorporated viscoelasticity into our model by using ARL Dynamic Mechanical Analysis (DMA) data. This data can be used to determine the rate- and temperature-dependent shear relaxation function.

We developed an equation-of-state (EOS) model to handle the bulk response. To that end, we constructed the Gibbs free energy  $G(P,T)$  as a function of pressure and temperature. The associated derivatives of  $G(P,T)$  with respect to  $P$  and  $T$  then determine the complete thermodynamic properties of PC. We used the specific volume measurements of Zoller [3] and the heat capacity measurements of Cheng and Wunderlich [4] to construct  $G(P,T)$ .

Our EOS captures the behavior of specific volume (Fig. 1) and heat capacity (Fig. 2). (Note that only the zero pressure heat capacity data are measured). In Fig. 3, our predicted bulk modulus is shown. Figure 4 shows the stress-strain data of [1] and our theoretical stress-strain curves; the comparison is excellent. In future work we will attempt to implement glassy polymer damage behavior into our model. In this way we can begin to model the PC face shields as they are subjected to projectile impacts of various speeds.

*For more information contact Brad Clements at [bclements@lanl.gov](mailto:bclements@lanl.gov).*

[1] P. Moy, et al., "Strain Rate Response of a Polycarbonate Under Uniaxial Compression," Session-Impact Behavior of Engineering Materials, Proceedings of SEM Annual Conference on Experimental Mechanics, Charlotte, North Carolina, June 2-4, 2003.

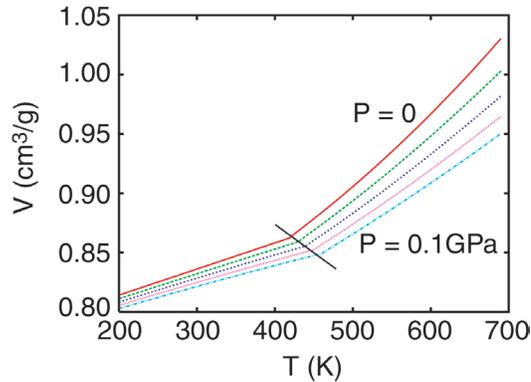
[2] T. Bjerke, et al., *Int. J. Plasticity* **18**, 549-567 (2002).

[3] P. Zoller, *J. Poly. Sci.: Part B: Polymer Phys. Eds.* **20**, 1453-1464 (1982).

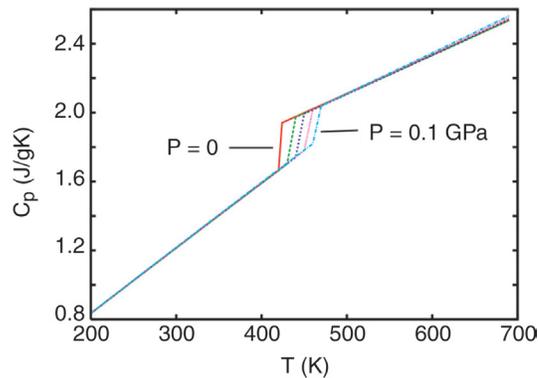
[4] S.Z. Cheng and B. Wunderlich, *J. Poly. Sci.: Part B: Polymer Phys. Eds.* **24**, 1755-1765 (1986).

### Funding Acknowledgements

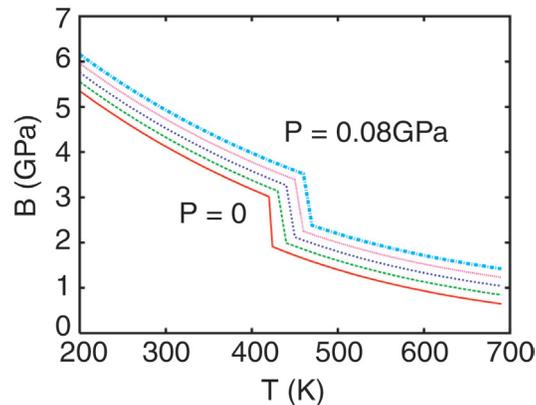
The Joint DoD/DOE Munitions Technology Development Program.



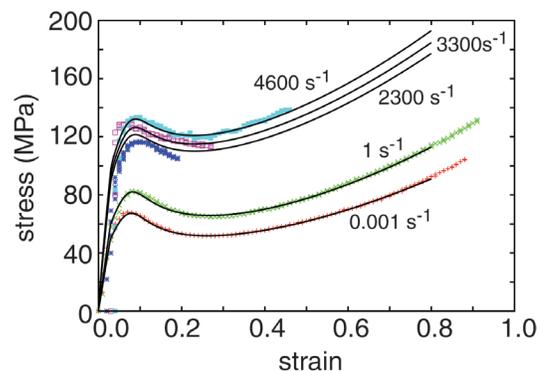
**Fig. 1.** Specific volume for PC [3]. The solid line is the glass transition boundary. Pressures correspond to 0.0, 0.2, 0.4, 0.6, and 0.8 GPa.



**Fig. 2.** Constant pressure heat capacity for PC.



**Fig. 3.** Isothermal bulk modulus for PC.



**Fig. 4.** Experimental [1] and theoretical stress-strain behavior for PC.

## Thermal Excitations and the Equation of State of Tantalum

Carl W. Greeff, Sven P. Rudin, Scott D. Crockett, John M. Wills, T-1

An important part of the mission of the Equation of State and Mechanics of Materials Group (T-1) is the development of highly accurate equations of state (EOS) that cover wide ranges of densities and temperatures. Typically, empirical EOS have been based on experimental data on the shock Hugoniot and extrapolated to other states through the use of models for the temperature dependence of the pressure and internal energy [1]. This kind of extrapolation has formed the basis for the pressure standards used in static compression experiments, and is needed to predict the results of dynamic compression experiments when the loading is more complex than a single shock. The thermal dependence of the EOS is determined by the excitations of the condensed matter system. We aim to exploit the capability of modern electronic structure theory to predict the spectrum of phonon and electron excitations in solids to create EOS of the highest possible accuracy. This work describes application of these ideas to the EOS of tantalum, which is used as a standard in both dynamic and static high-pressure experiments.

Recently, a new experimental technique has been developed [2] that allows for dynamic compression of samples to high pressures with smooth, magnetically driven waves. In principle, this smooth loading maintains nearly constant entropy, in contrast to dissipative shock loading. It is expected that it will soon be possible to carry out isentropic compression experiments (ICE) to pressures of 400 GPa (higher than the earth's center) in Ta. In this range, the

ICE loading should result in much lower temperatures than single shock loading, allowing the investigation of the EOS on a very different path. The design and interpretation of these experiments is significantly more involved than that of standard shock Hugoniot measurements. Supporting the design of ICE and predicting their results is an important motivation for this work.

We formulate the EOS by writing the total Helmholtz free energy as a sum of terms  $F(V,T) = \Phi_0(V) + F_{vib}(V,T) + F_{ei}(V,T)$ , where  $\Phi_0$  is the energy of a perfect lattice, and  $F_{vib}$  and  $F_{ei}$  are lattice vibration and electronic excitation free energies, respectively. To evaluate  $F_{vib}$  and  $F_{ei}$ , we require the densities of states for the phonons and electrons as functions of density. These have been obtained from electronic structure theory. The capability to predict the phonon excitations this way is illustrated in Fig. 1, where the calculated phonon dispersions curves are compared with experiments done at ambient pressure. The calculations use no adjustable parameters, and so the accuracy of the results gives us confidence in our ability to predict the excitation spectrum at high pressure, where it is not directly measured. The static lattice energy  $\Phi_0$  is determined partially empirically, using the accurately measured lattice spacing and bulk modulus, while requiring that  $\Phi_0$  be consistent with theory at high compression.

Having constructed the EOS, we have an accurate basis for predicting properties that are not measured. This is illustrated in Fig. 2, which shows the temperatures along two thermodynamic paths, the

Hugoniot and the principal isentrope. Hugoniot states are those reached in single shock compression, while the isentrope is approached in smooth compression. We see from the figure that at 400 GPa, the pressure of the proposed ICE, the Hugoniot is in the liquid, at a temperature of over 1 eV, while the isentrope is at a modest 750 K. At this low temperature, the isentrope pressure is completely dominated by  $\Phi_0$ , while the Hugoniot pressure has a 100 GPa contribution from thermal excitations.

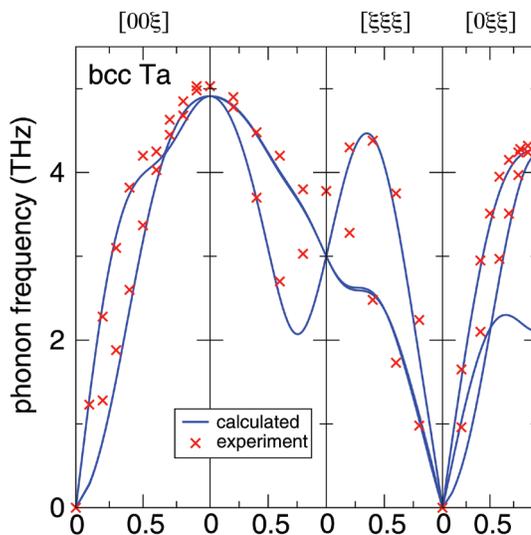
An important aspect of the design of ICE is to control the shape of the drive pulse, so that a shock is not formed as the wave propagates through the sample. With our EOS, we can predict the required pulse shape for a given peak pressure and sample thickness. Results of such a calculation are shown in Fig. 3. Calculations like these are necessary to design a successful experiment, especially an ambitious one like 400 GPa ICE. We are also using our models and codes to investigate the importance of non-ideal effects, like dissipation of plastic work, on ICE. These calculations on Ta illustrate how we are using modern theory and computational techniques to improve the accuracy of our EOS, and support important new experimental capabilities.

*For more information contact Carl W. Greeff at [greeff@lanl.gov](mailto:greeff@lanl.gov).*

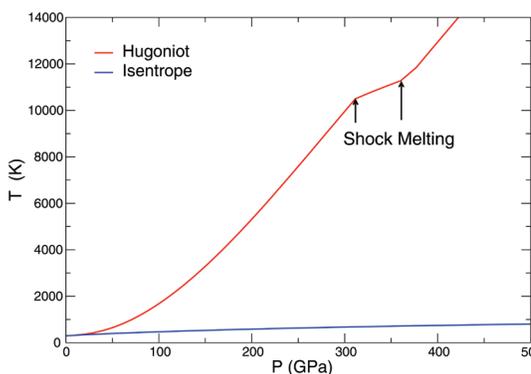
- [1] R.G. McQueen, et al., in *High Velocity Impact Phenomena*, R. Kinslow, Ed. (Academic Press, New York, 1970).
- [2] C.A. Hall, *Phys. Plasmas* 7, 2069 (2000).

### Funding Acknowledgements

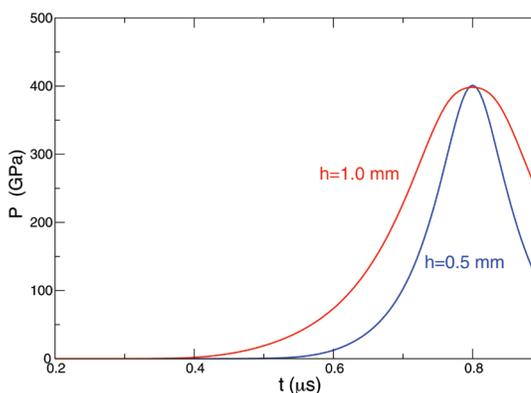
NNSA's Advanced Simulation and Computing (ASC), Materials and Physics Program.



**Fig. 1.** Phonon dispersion curves for Ta at ambient density. Solid blue curves are calculated with density functional theory, and red symbols are experimental data.



**Fig. 2.** Hugoniot and principal isentrope of Ta. The kink in the Hugoniot at 310 GPa is due to the melting transition.



**Fig. 3.** Calculated drive pressures needed to achieve shockless compression of Ta to 400 GPa. Red curve is for sample thickness of 1 mm, and blue curve is for thickness of 0.5 mm.

## A Sample of Some New Equations of State

Scott Crockett, Leonid Burakovsky, Sven Rudin, T-1

**R**ecently, we have been exploring several different categories of materials: metals, glasses, liquids, and ceramics.

By using electronic structure calculations and the best experimental data available to guide our standard Equation of State (EOS) models, we have created new baseline EOS and made improvements on existing EOS. Some of the materials include baseline EOS such as gallium and boron, for which no SESAME table existed, and improved EOS such as boron carbide, lithium fluoride, and vanadium.

Gallium melts slightly above room temperature and has three known solid phases. A new EOS was created by using theoretical calculations and published parameters, which reproduce two of the three solid phases and the liquid phase of the EOS. Figure 1 illustrates how well the new EOS predicts the shock response of the liquid. It is also in agreement with preliminary quantum molecular dynamics (QMD) simulations of Stephane Mazevet, formerly of T-4.

The boron and boron carbide EOS were created in parallel. Both materials exhibit high strength along with other similar characteristics. These EOS were created uniquely to compensate for the strength effect in the shock data. Figures 2 and 3 illustrate the shock response of the new EOS compared with the data.

The new vanadium EOS is similar to the older SESAME EOS at moderate compression, but the Hugoniot for the new EOS agrees much better with the shock data at higher compression than its predecessors (Fig. 4). This new EOS also

takes into account melting, for which it correctly reproduces the available melt data.

The final example, lithium fluoride, is used as a window for VISAR measurements in shock experiments. The new EOS was created using a similar process as the gallium EOS. This new EOS Hugoniot matches new shock data from Sandia National Laboratories, which ranges from two to ten Mbars in pressure. Figure 5 compares this new EOS with the older shock data, QMD calculations, and a previous SESAME table.

In summary, with every new theoretical development and new data from EOS experiments we continually improve SESAME EOS database.

*For more information contact Scott Crockett at [crockett@lanl.gov](mailto:crockett@lanl.gov).*

[1] S. Crockett and L. Burakovsky, "Gallium Equation of State (SESAME 2370)," Los Alamos National Laboratory report, LA-UR-06-8390 (2006).

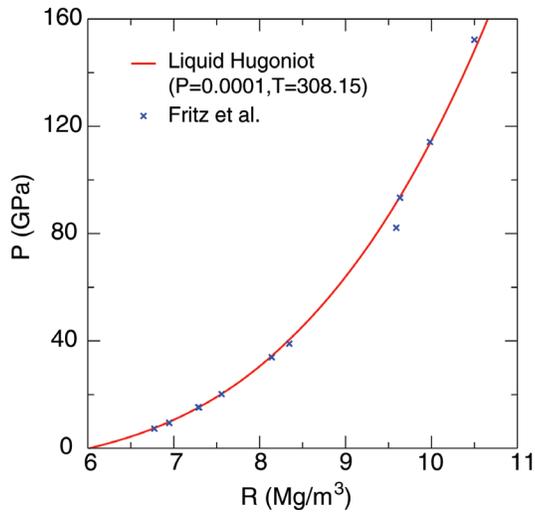
[2] S. Crockett, "Creating New Equations of State for High Strength Materials (Boron and Boron Carbide)," Los Alamos National Laboratory report, LA-UR-06-8402 (2006).

[3] S. Crockett and S. Rudin, "Vanadium Equation of State (SESAME 2552): A Comparison of Old and New SESAME Tables," Los Alamos National Laboratory report, LA-UR-06-8389 (2006).

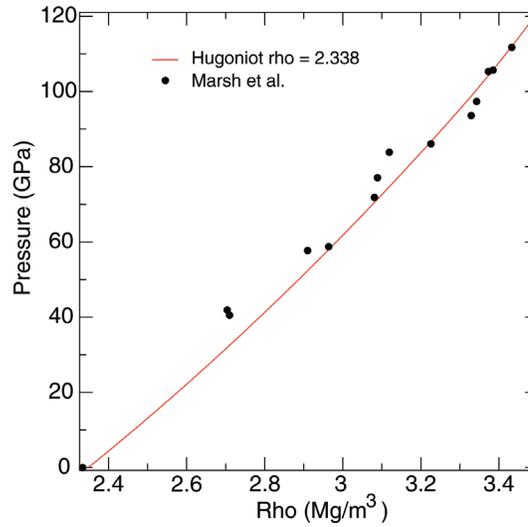
[4] S. Crockett and S. Rudin, "Lithium Fluoride Equation of State (SESAME 7271)," Los Alamos National Laboratory report, LA-UR-06-8401 (2006).

### Funding Acknowledgements

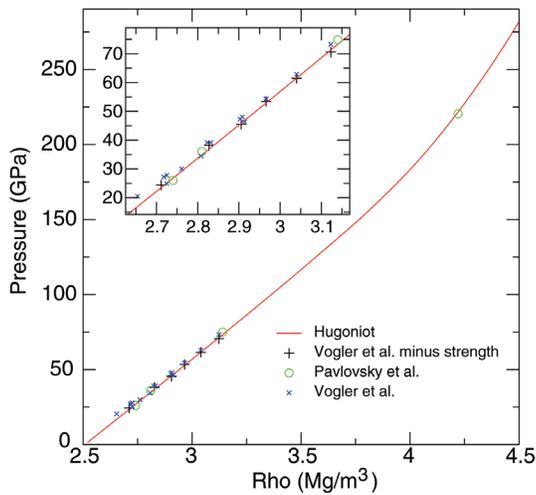
NNSA's Advanced Simulation and Computing (ASC), Materials and Physics Program.



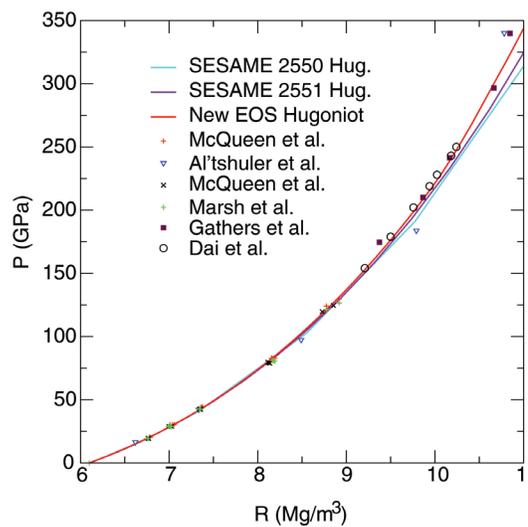
**Fig. 1. Left**  
A gallium pressure Hugoniot compared with shock experiment [1].



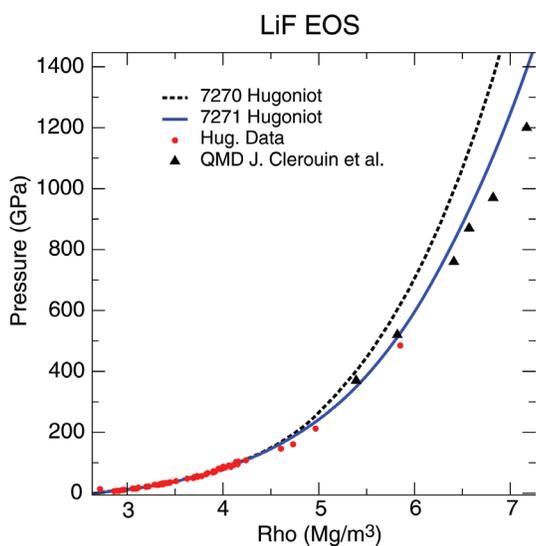
**Fig. 2. Right**  
A boron pressure Hugoniot compared with shock experiment [2].



**Fig. 3. Left**  
A boron carbide Hugoniot compared with shock experiments [2].



**Fig. 4. Right**  
A vanadium Hugoniot compared with older EOS and shock experiments [3].



**Fig. 5.**  
A lithium fluoride Hugoniot compared with QMD simulations, an older EOS, and older shock experiments [4].

## Zero-Temperature Magnetic Transition in an Easy-Axis Kondo Lattice Model

Jian-Xin Zhu, T-11; S. Kirchner, Rice University; R. Bulla, Universitat Augsburg; Qimiao Si, Rice University

Over the past decade or so, a sizable number of (nearly) stoichiometric heavy fermions have been discovered in which the antiferromagnetic transition temperature can be continuously suppressed to zero. These quantum critical materials have not only elucidated the heavy fermion physics but also provided a concrete setting to address the larger question on the nature of quantum criticality. The application of the Landau-Ginzburg-Wilson paradigm considers the fluctuations of the magnetic order parameter as the primary critical modes. The resulting  $T=0$  K spin-density-wave (SDW) quantum critical point (QCP) is Gaussian. However, a host of dynamical, transport, and thermodynamic data suggest that the observed QCPs are non-Gaussian, indicating the existence of additional quantum critical modes. Since there is not yet a universal prescription for the identification of such emergent critical modes, microscopic considerations have been playing an important role.

One theoretical idea invokes the breakdown of the Kondo screening effect at the magnetic QCP to characterize the new critical modes. In the form of local quantum criticality, the destruction of the Kondo effect arises through the decoherence by the magnetic order parameter fluctuations. Microscopically, this picture has been studied through the extended dynamical mean field theory (EDMFT) approach to the Kondo lattice systems.

This approach describes the Kondo lattice in terms of a Bose-Fermi Kondo (BFK) model, with the spectra of its fermionic and bosonic baths self-consistently determined. The EDMFT approach goes beyond the conventional methods. It treats such competition dynamically. An important question is whether the actual zero temperature transition is second order.

We have used the EDMFT to address this issue in a spin-1/2 antiferromagnetic Kondo lattice model with an easy-axis anisotropy [1,2]. We have derived results in real frequency using the bosonic numerical renormalization group (bNRG) method and compare them with Quantum Monte Carlo results in Matsubara frequency. The bNRG results show a logarithmic divergence in the critical local spin susceptibility, signaling a destruction of Kondo screening. The  $T=0$  K transition is nearly second order, with any jump in the magnetic order parameter not exceeding a few percent of the full moment. Our results are important for experiments, not only because the numerical studies play an important role in the understanding of the unusual magnetic dynamics (which itself was the primary initial experimental indication for the non-SDW nature of the QCP), but also because the theoretical picture has crucial predictions for other experiments (such as a jump of Fermi volume and fractional exponent in the Gruneisen ratio) that are actively being examined by ongoing experiments.

More generally, whether unconventional QCPs would be stable and relevant to realistic models/materials or tend to be preempted by first-order transitions are broadly important and also arises in, e.g., the case of deconfined quantum criticality in spin/boson lattice systems.

*For more information contact Jian-Xin Zhu at [jxzhu@lanl.gov](mailto:jxzhu@lanl.gov).*

[1] Jian-Xin Zhu, et al., <http://arxiv.org/abs/cond-mat/0607567>.

[2] Jian-Xin Zhu, et al., *Phys. Rev. Lett.* **91**, 156404 (2003).

### **Funding Acknowledgements**

Weapon Supported Research Program (No. JA2W/JHU1), and NNSA's Advanced Simulation and Computing (ASC), Materials and Physics Program.

# Dynamically Driven Phase Transformation in Damaged Composite Materials: Micromechanics

JeeYeon N. Plohr, Bradford E. Clements T-1; Frank L. Addessio, T-3

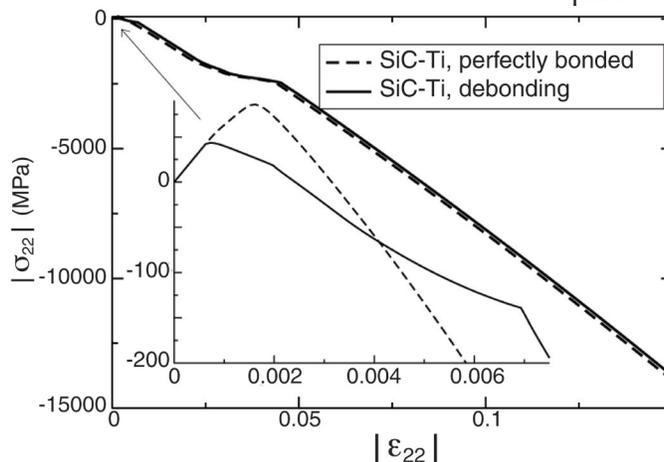
**W**e have developed a theoretical framework to model composite materials when the constituents exhibit viscoplasticity, phase transformation, cracking, and debonding [1,2]. The particular systems we studied are SiC-Ti composites, where Ti undergoes solid-solid phase transformation, SiC cracks, and the interface between these materials debonds. The interplay among these phenomena is complicated and often leads to unexpected materials behavior. In understanding these mechanisms in composite materials, it is crucial to examine them at the micromechanics level. This allows us not only to see the distribution of stress, cracking, and debonding inside the composites, but also to identify the areas where competition among these different physical phenomena occurs. This capability to analyze the micromechanics of the constituents is also useful to validate numerical simulation codes. In this article, we report the following three points that we have learned via micromechanics analysis.

1. Even when the stress in a certain direction is tensile, if the net stress (trace of the stress tensor) is compressive, debonding is inactive. We consider the case of compressive net stress because the phase transformation in Ti is pressure-induced. Still, the contribution from the deviatoric part of the stress can overcome compressive pressure, and the overall stress can be tensile. However, viscoplasticity limits the growth of the deviatoric stress, and all components of the stress eventually become compressive. (See Fig. 1.)

2. Different Ti subcells experience phase transformation at different times and different rates because of the anisotropy caused by SiC subcells. SiC has higher elastic moduli, and the Ti subcells close to SiC subcells are subjected to higher load as a consequence of normal traction continuity. When SiC develops cracks, its effective moduli becomes lower, and the anisotropy in stress distribution among the Ti subcells become weaker. As shown in Fig. 2, Ti subcells show a homogenized phase transformation pattern.

3. The dependence of stress-strain relationship on the rate of loading is less prominent in damaged composites than in composites without damage. We conclude that this rate desensitization results from the effective homogenization of moduli as the SiC develops cracks. (See Fig. 3.)

**Fig. 1.**  
The effect of debonding on the material stress-strain behavior. In this simulation,  $\dot{\epsilon}_{11} = -5.0 \text{ s}^{-1}$ ,  $\dot{\epsilon}_{22} = 3.0 \text{ s}^{-1}$ ,  $\dot{\epsilon}_{33} = 0.0 \text{ s}^{-1}$ .

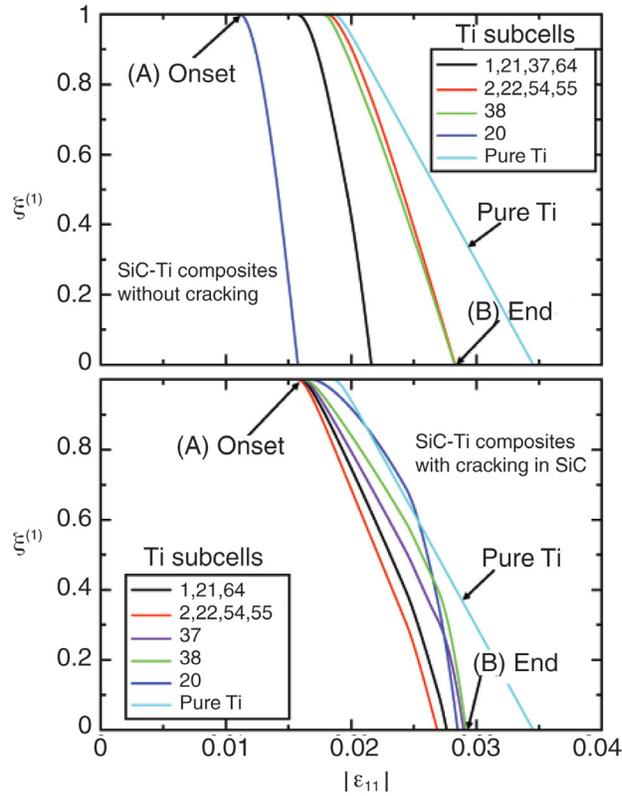


For more information contact  
JeeYeon Plohr at [jplohr@lanl.gov](mailto:jplohr@lanl.gov).

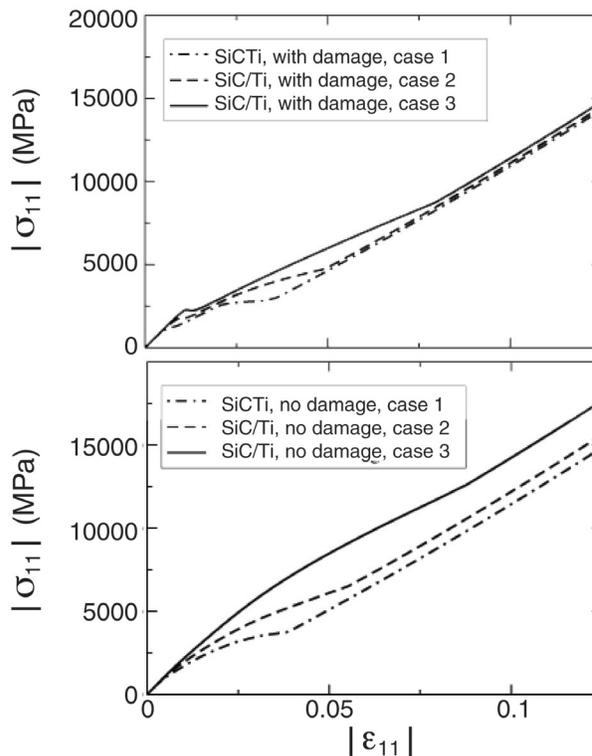
[1] B. Clements, et al., *J. Appl. Phys.* **100** (12), 123520 (2006).

[2] J.N. Plohr, et al., *J. Appl. Phys.* **100** (12), 123521 (2006).

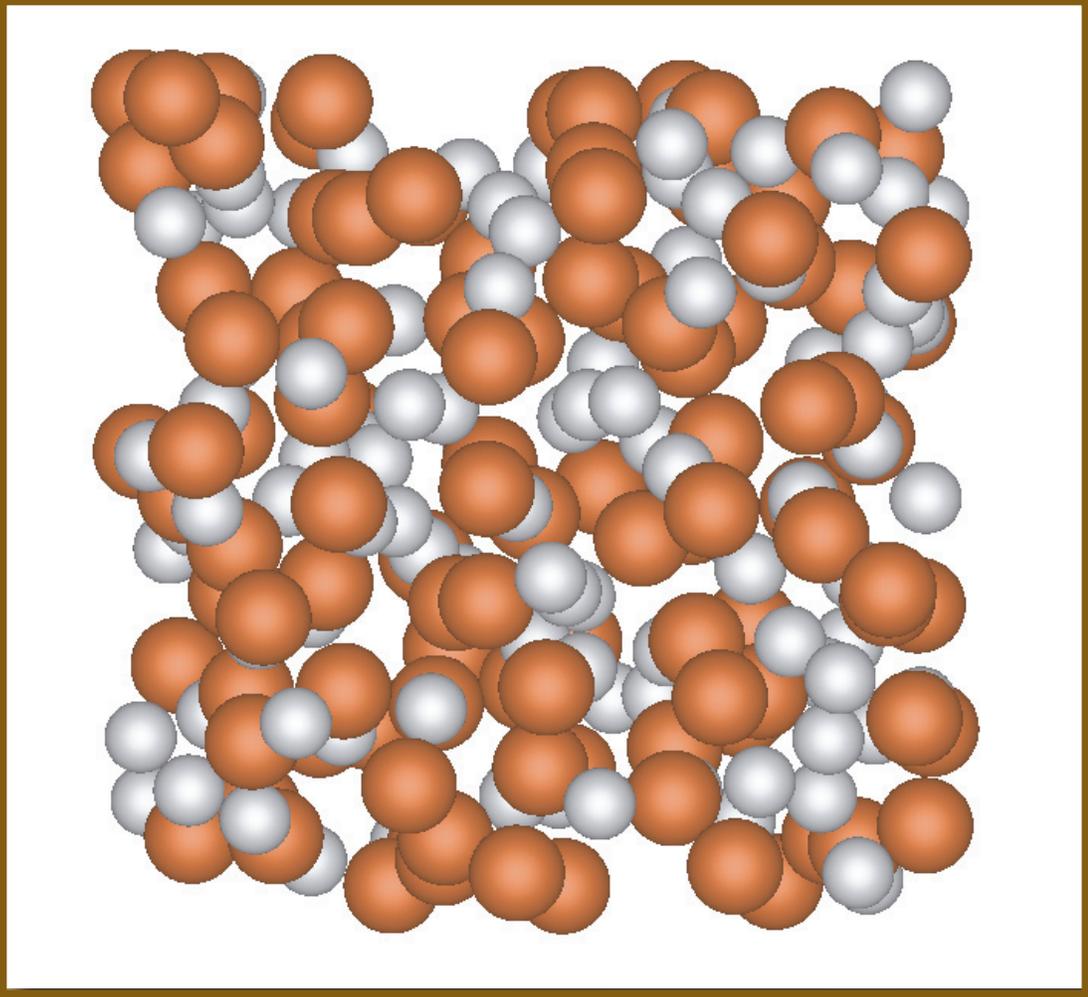
**Funding Acknowledgements**  
NNSA's Advanced Simulation and Computing (ASC), Materials and Physics Program; and Joint DoD/DOE MOU Program.



**Fig. 2.**  
The Ti  $\alpha$ -phase mass fraction  $\xi^{(1)}$  as a function of the macro-strain  $\epsilon_{11}$ , for a uniaxial strain loading rate of  $\dot{\epsilon}_{11} = -1.0 \text{ s}^{-1}$ . Crack growth damage is suppressed (top) and allowed to occur (bottom). Points (A) and (B) indicate the onset and conclusion for the phase transformation in the composite.



**Fig. 3.**  
The effect of strain rate on the (macro) stress-strain behavior of the undamaged (bottom) and damaged (top) composite. In case 1,  $\dot{\epsilon}_{11} = -100 \text{ s}^{-1}$ ,  $\dot{\epsilon}_{22} = \dot{\epsilon}_{33} = 10 \text{ s}^{-1}$ . In case 2,  $\dot{\epsilon}_{11} = -1000 \text{ s}^{-1}$ ,  $\dot{\epsilon}_{22} = \dot{\epsilon}_{33} = 100 \text{ s}^{-1}$ . In case 3,  $\dot{\epsilon}_{11} = -5000 \text{ s}^{-1}$ ,  $\dot{\epsilon}_{22} = \dot{\epsilon}_{33} = 500 \text{ s}^{-1}$ .



# **Atomic and Molecular Theory and Modeling**

The work in atomic and molecular theory and modeling described in the following articles encompasses a variety of topics that range from the oxidation chemistry of high explosive binders to quantum molecular dynamics calculations passing through work on the strength of composite materials.

# Aging of the Binder in Plastic-Bonded Explosive PBX-9501 and Free-Radical Oxidation

Joel D. Kress, T-12

The plastic-bonded explosive PBX-9501 is 94.9% by weight HMX explosive, 2.5% nitroplasticizer (NP), 2.5% Estane<sup>®</sup> 5703, and 0.1% Irganox 1010 stabilizer (antioxidant). Estane 5703 is a poly(ester urethane) segmented copolymer, which serves as a glue that binds the HMX crystals together to allow for machining of high-precision high-explosive parts. The urethane units segregate (phase separate) into “hard” domains that act as physical crosslinks between the “soft” polyester domains. (See Fig. 1.) The NP softens the Estane, and together this markedly decreases the mechanical sensitivity of the PBX. However, the Estane slowly degrades in time, and the corresponding effect on the mechanical properties of the PBX can potentially raise safety and reliability concerns.

Currently, we are studying the chemical mechanisms of degradation processes in PBX-9501 and developing kinetics models. Here we describe recent progress in the study of the free-radical oxidation of the urethane segments. Free-radical oxidation may be as important as hydrolysis [1], especially as the PBX dries out, since the storage environment has both low moisture and oxygen content. The nitroplasticizer (NP) is composed of a mixture of bis-2,2-dinitropropyl acetal (BDNPA) and formal (BDNPF). Various research efforts at Los Alamos and Pantex are showing that upon heating NP loses nitro ( $\text{NO}_2$ ) groups and produces oxidizing species that can degrade Estane (for example, the oxidation of the methylene bridge, left-hand dashed circle in Fig. 1). The oxidation of the Estane polymer can cause crosslinking (increased MW and gel formation), while

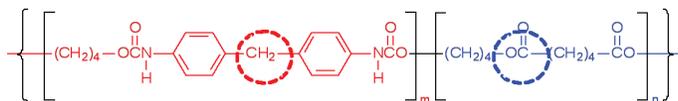
small molecule remnants of the reactions can be detected in gas analyses. Such large changes in MW can have drastic effects on the mechanical properties of the elastomer.

Results on the oxidative degradation of Estane are being obtained from the Constituent Aging Study (CAS). The goals [2] of the CAS are to artificially age the PBX-9501 constituents (HMX, Estane, NP, stabilizer) in 25 different combinations at accelerated rates, detect chemical reactions that take place, identify reaction products and possible degradation mechanisms, and provide data for lifetime prediction modeling. One of the most striking findings is that combinations containing NP (NP + Estane and NP + Estane + stabilizer) show a strong correlation between the total amount of gas products generated and significant changes in the molecular weight of the Estane.

Changes in number-average molecular weight ( $M_n$ ), weight-average molecular weight ( $M_w$ ) and polydispersity ( $MWD = M_w/M_n$ ) of the Estane were measured by gel permeation chromatography (GPC). Two different methods were used: (1) average molecular weight measurements relative to polystyrene standards with a differential refractive index (RI) detection, and (2) absolute molecular weight measurements using a multiangle light scattering (MALS) technique. To interpret the GPC measurements, Flory-Stockmayer (FS) theory of random crosslinking is used.

In another study to understand degradation mechanisms under simulated aging conditions, isotopically labeled model materials are used [3] to understand degradation mechanisms. Isotopic enrichment has been utilized to improve the sensitivity of spectroscopic techniques for

Fig. 1. Chemical structure of Estane<sup>®</sup>.



observing small quantities of degradation products. Detection of degradation products, one of which is the result of oxidation of the bridging methylene carbon of the MDI unit (red dashed circle around CH<sub>2</sub> in Fig. 1), was accomplished by 1-D and 2-D nuclear magnetic resonance (NMR) techniques. By quantifying the relative amounts of the oxidized product(s) to the starting material at different aging times and temperatures, the kinetics of the degradation process was determined.

In Table I we summarize our analyses of the kinetics of oxidation of Estane due to NP decomposition. Since differing amounts (masses) and ratios of Estane and NP were used in the various CAS samples, we attempt to place all of the CAS results on an equal footing by dividing (normalizing) the observed rates by the initial amount of NP (denoted [NP]<sub>0</sub>). For the rate of evolution, R<sub>NO<sub>x</sub></sub>, of NO<sub>x</sub> gasses [2] from eight different CAS samples the normalized rates R<sub>NO<sub>x</sub></sub>/[NP]<sub>0</sub> ranged between 5 × 10<sup>-4</sup>/week and 15 × 10<sup>-4</sup>/week. As derived from the isotopically labeled NMR experiments, the rate of carbonyl formation (R<sub>C=O</sub>) of the methylene bridge on the urethane unit gives a normalized rate of R<sub>C=O</sub>/[NP]<sub>0</sub> = 2.8 × 10<sup>-4</sup>/week. Finally, the crosslinking rate (R<sub>X-link</sub>) was calculated from GPC-MALS and GPC-RI data. (The crosslinking rates are displayed in Fig. 2.) This yielded normalized rates R<sub>X-link</sub>/[NP]<sub>0</sub> between 0.2 × 10<sup>-4</sup>/week and 0.5 × 10<sup>-4</sup>/week. The normalized NO<sub>x</sub> evolution rate is about 3 to 5 times faster than the normalized observed carbonyl formation rate. Also, the normalized observed crosslinking rate is about 6 to 13 times slower than the normalized observed carbonyl formation rate. Similar values (i.e., within an order of magnitude) of all of the observed rates normalized by the initial amount of NP may suggest that NP decomposition is the rate limiting step in both the carbonyl formation and Estane crosslinking in Estane/NP mixtures.

For more information contact Joel Kress at [jdk@lanl.gov](mailto:jdk@lanl.gov).

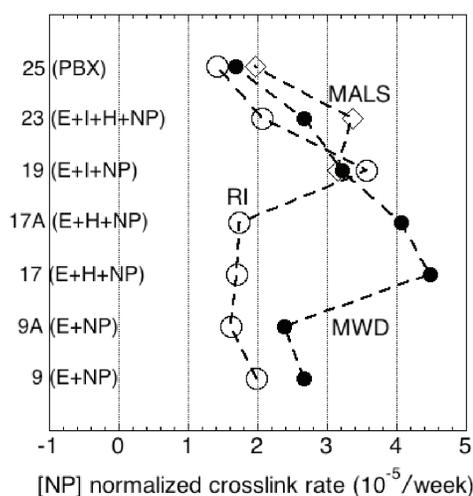
[1] M.R. Salazar, et al., *J. Polymer Sci.: Part A* **41**, 1136 (2003).  
 [2] D.K. Pauler, et al., *AIP Conf. Proceedings* **845**, 547 (2006).  
 [3] D.W. Wroblewski, et al., "Accelerated Aging and Characterization of a Plasticized Poly(Ester Urethane) Binder," *Polymer Preprints*, submitted.

#### Funding Acknowledgements

NNSA's Campaign 8, Enhanced Surveillance.

**Table I.** Observed rates at T = 64°C normalized by the initial amount of nitroplasticizer (NP), [NP]<sub>0</sub>.

	Normalized Rate (10 <sup>-4</sup> /week)	Experiment
R <sub>NO<sub>x</sub></sub> /[NP] <sub>0</sub>	5 to 15	CAS NO <sub>x</sub> gas evolution
R <sub>C=O</sub> /[NP] <sub>0</sub>	2.8	NMR MDI Oxidation
R <sub>X-link</sub> /[NP] <sub>0</sub>	0.20 to 0.34	CAS MW (GPC-MALS)
R <sub>X-link</sub> /[NP] <sub>0</sub>	0.16 to 0.36	CAS MW (GPC-RI)
R <sub>X-link</sub> /[NP] <sub>0</sub>	0.17 to 0.45	CAS MWD (GPC-RI)



**Fig. 2.** NP normalized crosslinking rate (R<sub>X-link</sub>/[NP]<sub>0</sub>) at T = 64°C for CAS Samples #9, #9A, #17, #17A, #19, #23, and #25. E = Estane, NP = nitroplasticizer, H = HMX, I = Irganox, PBX = PBX 9501. Open circles = rate determined from weight-averaged MW with RI detection. Closed circles = rate determined from MWD with RI detection. Open diamonds = rate determined from weight-averaged MW with MALS detection.

## Multi-Length Scale Modeling of Highly Filled Polymer Composites

David E. Hanson, Cynthia Reichhardt, T-12; Axinte Ionita, T-1

**H**ighly filled polymer composites, including plastic-bonded explosives such as PBX-9501, are a unique class of high-value materials. PBX-9501 consists of >90% HMX particles (filler) held together by a polymer matrix (binder). Currently, there are no predictive materials science models for these materials. We are developing robust, predictive constitutive models for these composites that span the molecular to the macroscopic length scales. Using existing and emerging models, we will determine how the physics and chemistry that define the mechanical behavior at each length scale can be propagated upward to the next level.

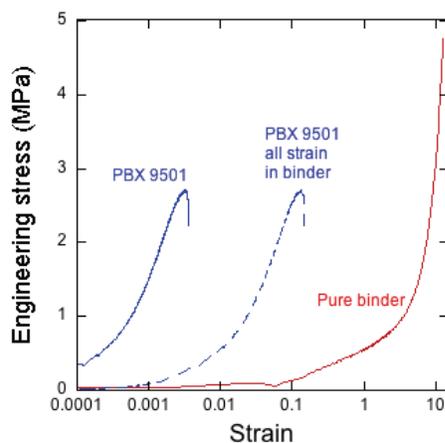
The mechanical behavior (e.g., stress-strain) of highly filled polymer composites is orders of magnitude different from the bulk behavior of its constituents. This is illustrated by compressive strain for PBX-9501, shown in Fig. 1. Over most of the strain, the stress for the pure binder material is about one tenth that in the composite material, but the strain-at-break is two orders of magnitude greater for the binder. The geometric constraints imposed by the binder adhering to the rigid filler particles dramatically change the mechanical response of the composite. Tensile experiments under cyclic deformation show that, even at low strains,

the material undergoes irreversible damage. The failure mechanisms responsible for this have their origins at the molecular level. We believe that when the local strain in the polymer network adjacent to a filler particle surface becomes sufficiently large, a polymer chain can detach from the surface, producing interfacial voids on the smallest filler particles. As the strain increases, the voids expand and propagate as cracks, ultimately causing macroscopic material failure.

It is important to model the material starting at the smallest length scales because it is here that changes due to long-term chemical aging (e.g., polymer chain scissions or oxidative cross linking) will appear. The challenge lies in developing physically faithful models for each length scale and coupling them together to achieve a computationally tractable protocol that can be implemented in Finite Element Method (FEM) simulation codes. Currently, these codes are not able to predict the effects of long-term chemical aging or thermo-mechanical insult unless all of the code parameters are explicitly fit to very specific experimental data (which are usually not available).

We represent the molecular physics occurring between the HMX substrate and the Estane with the polymer/node network model EPnet [1] (a unique asset of Los Alamos National Laboratory), using an estimate of the interfacial binding energy obtained from molecular dynamics (MD) simulations in conjunction with quantum chemistry calculations. The EPnet code treats polymer networks as an explicit collection of point-like nodes in 3 dimensions, connected by virtual chains characterized by a length distribution and a force model. The code supports void regions and will be modified to treat rigid spherical regions that are attached to the polymer network. Energy changes in the polymer network are computed by integrating the chain forces between

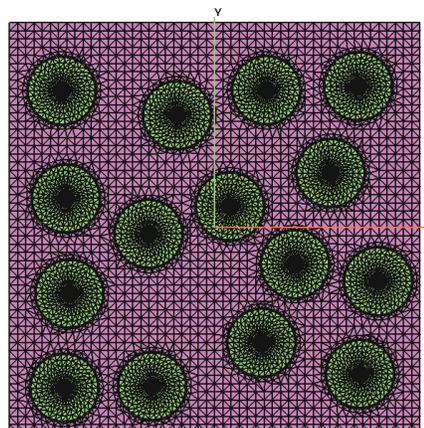
**Fig. 1.** Compressive strain for PBX-9501. Over most of the strain, the stress for the pure binder material is about one tenth that in the composite material, but the strain-at-break is two orders of magnitude greater for the binder. For the middle curve, the experimental strains for the PBX-9501 have been increased by a factor of about 40 to give the approximate strain in the binder material.



connected nodes and filler particle interfaces as the simulation cell undergoes strain. Parameters such as the chain forces will be obtained by fitting to tensile stress-strain experiments on material representative of the mesoscale region (small particles only). Void formation at the filler surface, generated in response to a tensile strain, will be studied in simulations of cells containing single spherical particles with periodic boundary conditions or single particles adjacent to a plane surface representing a large filler particle. Material failure at this length scale will occur when the void region (crack) bridges the simulation cell. A sufficient number of statistically independent network/filler particle cases will be run to determine the failure probability as a function of strain, parametric on other parameters (e.g., fill fraction).

This information will provide the input for the Probabilistic Crack Occurrence (PCO) statistical model, which is designed to connect the EPNet scale with the 2Scale FEM modeling technique. The PCO model is based on a material containing inclusions of two widely differing sizes. The small inclusions are assumed to be part of a background matrix, which is normally represented as a homogeneous substance by a standard FEM mesh. PCO is designed to reintroduce inhomogeneity in the background matrix in a controlled, statistical manner by providing a space-dependent criterion for the strain at which a crack should be opened in the FEM matrix based on the small scale properties of the material. The initial version of PCO is complete, and employs a statistical technique known as the weakest link approach. The small inclusions are assumed to follow a Poisson spatial distribution in the matrix. The solution of the model has the property that the matrix separating two large inclusions decreases in strength as the spacing between the large inclusions increases. The most likely locations for cracks to form will depend on the details of the random arrangement of large inclusions in a physically meaningful way: cracks are most likely to occur in regions with a relatively low density of large inclusions.

The largest scale of our model effort is the 2Scale-FEM analysis, which incorporates



**Fig. 2.** RVE mesh (LAGRIT) using a random spheres model distribution. Large grains are simulated by spheres and the “dirty binder” as a homogenized medium.

energy vs strain from EPNet as well as crack initiation criteria from PCO. In engineering applications using heterogeneous materials, a common approach is to start with a Representative Volume Element (RVE), the smallest domain that can be considered to have approximately constant average material properties. Here, we will determine the RVE response for a heterogeneous material (including cracking) through a Direct Numerical Simulation (DNS) using FEM in conjunction with the scheme described below. The geometry of the RVE can be constructed directly from micrographs as an idealized model in which large grains of HMX are modeled as random distribution of spheres, as shown in Fig. 2. Each constituent of the RVE is assigned a constitutive law determined from physics theories starting at the molecular level. To simulate cracking, a failure criterion for crack initiation and propagation is incorporated from the PCO analysis. The initial RVE (undamaged) is loaded until cracks occur. Cracks are propagated by stopping the FEM analysis and re-meshing the RVE with explicit crack boundaries between mesh points. The simulation is then restarted using the new geometry. The DNS at the RVE level offers the advantage of using a large variety of local microstructure of the RVE in conjunction with different types of loading conditions.

*For more information contact Cynthia Reichhardt at [cjrx@lanl.gov](mailto:cjrx@lanl.gov).*

[1] D.E. Hanson, *Polymer* **45** (3), 1055 (2004).

#### Funding Acknowledgements

NNSA’s Campaign 8, Enhanced Surveillance.

## Time-Reversible *ab initio* Molecular Dynamics

Anders M. N. Niklasson, T-1; C. J. Tymczak, Texas Southern University;  
Matt Challacombe, T-12

In classical molecular dynamics (MD) it is well known that computational schemes are improved significantly by imposing time-reversal symmetry on the integration of Newton's equations of motion. However, in *ab initio* Born-Oppenheimer MD, based on self-consistent field (SCF) theory, a time-reversible integration of the nuclear motion is problematic. In Born-Oppenheimer MD the electronic degrees of freedom, e.g., the density  $\rho(t)$ , is propagated by an extrapolation from previous time steps,  $\rho(t - n\delta t)$ . This is necessary to provide an accurate initial guess for the SCF optimization, which constrains the solution to the Born-Oppenheimer potential energy surface. A good initial guess often reduces the computational cost by an order of magnitude. The electron extrapolation from previous time steps combined with the SCF optimization constitutes an adiabatic propagation of the electronic degrees of freedom,

$$\rho(t + \delta t) = \text{SCF}[\rho(t), \rho(t - \delta t), \dots]. \quad (1)$$

Unfortunately, Eq. (1) is inconsistent with a time-reversible dynamics because of the nonlinearity and irreversibility of the SCF procedure. The remedy for this problem was previously to force the SCF optimization to a very high degree of accuracy, which is computationally expensive. At "exact" SCF convergence the optimized density is independent of the input density and there is therefore no longer a "propagation" of the density. However, a more efficient approach would be to keep the electron

propagation, but with a restored time-reversal symmetry. This was achieved with the introduction of time-reversible Born-Oppenheimer MD [1].

The principle of lossless time-reversible dual channel integration of the electronic degrees of freedom is illustrated in Fig. 1. The idea is to replace the irreversible adiabatic propagation of the electron density in Eq. (1) by a dual filter procedure with an additional auxiliary density channel denoted by a tilde. The scheme is perfectly reversible, despite the irreversibility of the SCF procedure.

Figure 2 illustrates the effect of replacing a lossy linear extrapolation scheme with a time-reversible lossless linear integration.

The key advantage with time-reversible Born-Oppenheimer MD is the improved global energy stability even under approximate SCF convergence. The computational cost can therefore be substantially reduced.

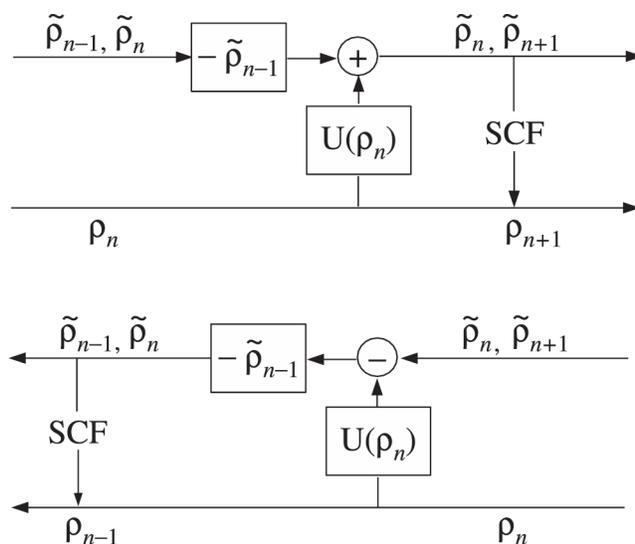
This work demonstrates a first step toward a new class of time-reversible and symplectic integrators that previously was limited only to classical MD. We hope that further development may lead to new standards for large-scale *ab initio* MD simulations.

*For further information contact Anders Niklasson at [amn@lanl.gov](mailto:amn@lanl.gov).*

[1] A.M.N. Niklasson, et al., *Phys. Rev. Lett.* **97**, 123001 (2006).

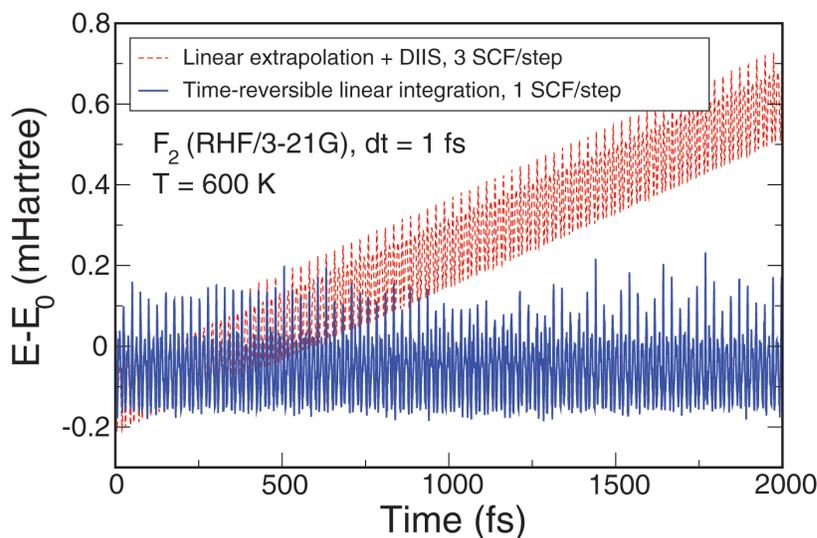
### Funding Acknowledgements

NNSA's Advanced Simulation and Computing (ASC), Materials and Physics Program; and Laboratory Directed Research and Development Program.



**Fig. 1.**

The principle of lossless dual-channel filter integration. The propagation is driven by the update filter  $U$ , where  $\tilde{\rho}_{n+1} = U[\rho_n] - \tilde{\rho}_{n-1}$  and  $\rho_n = \text{SCF}[\rho_n]$ . The filter scheme is perfectly reversible, either by going backwards and changing the  $\oplus$  sign to a  $-$  sign, or equivalently, by changing  $\delta t$  to  $-\delta t$ .



**Fig. 2.**

The energy fluctuations around the starting energy  $E_0$  as a function of time for a  $F_2$  molecule at a temperature  $T \approx 600\text{K}$  (Hartree-Fock theory with a Gaussian basis set RHF/3-21G and time steps  $\delta t = 1\text{ fs}$ ).

## Recent Developments in Modeling Atomic Processes in Dense Plasmas

James Colgan, Joseph Abdallah, Jr., T-4; Ben Streufert, University of Virginia;  
Christopher J. Fontes, X-1-NAD

**R**ecent experiments using intense laser pulses on thin targets have produced spectra in which it has been speculated that certain features are due to multiple ionization or recombination events [1]. To explore this possibility, a program of work has begun to add the rate coefficients for collisional double ionization and its inverse process, four-body recombination, to the collisional rate matrix computed within the Los Alamos plasma kinetics code ATOMIC.

First of all, the required collisional double ionization cross sections were obtained from semiempirical fits to experimental measurements, since there is no reliable theoretical method that can produce collisional double ionization cross sections in a timely manner for a wide range of atoms or atomic ions [2]. The semiempirical fits were designed to be reliable over a wide range of incident electron energies, and to be sufficiently versatile to describe cross sections from few-electron ions to larger atomic systems. These semiempirical fits were implemented in the Los Alamos multipurpose ionization code GIPPER [3]. Of course, in any rate matrix, one must also include the corresponding inverse processes so that ion populations may be properly obtained. In this case, the corresponding inverse process is four-body recombination, where three electrons interact with an ion such that two electrons recombine with the ion. We have derived the four-body recombination rates from detailed-balance considerations, for cases where the electron distributions are Maxwellian or non-Maxwellian.

The resulting collisional rate matrix is then solved as part of the Los Alamos plasma kinetics code ATOMIC [4]. We have examined emission spectra produced from solving the coupled rate equations, including the double ionization and four-body recombination rate coefficients, for an Ar plasma in which various fractions of hot electrons are present. We find that inclusion of these multiple-electron effects can make appreciable differences to the average ionization stage of the plasma and the resulting emission spectra at relatively high electron densities, for plasma conditions that contain appreciable fractions of hot electrons.

For example, in Fig. 1, we present emission spectra of Ar at various electron densities. The electron temperature is 120 eV with a 50% fraction of hot electrons, at various temperatures (10, 20, and 30 keV) as labeled in the heading. At this temperature and density, the dominant spectral peak is due to transitions in the He-like stage of Ar (at around 3.96 Å); peaks from lower (more electrons) ion stages with lower emissivities are also present. At lower electron densities the inclusion of the collisional double ionization and four-body recombination rates (labeled as CDI/4BR) make little difference to the emission spectra. However, as the electron density is increased to  $10^{23}/\text{cm}^3$ , the spectra including the CDI/4BR are quite different from that without CDI/4BR. The spectra at a density of  $10^{23}/\text{cm}^3$  show peaks corresponding to ions containing more electrons (even up to C-like and N-like Ar), indicating that the CDI/4BR allows the ions to recombine

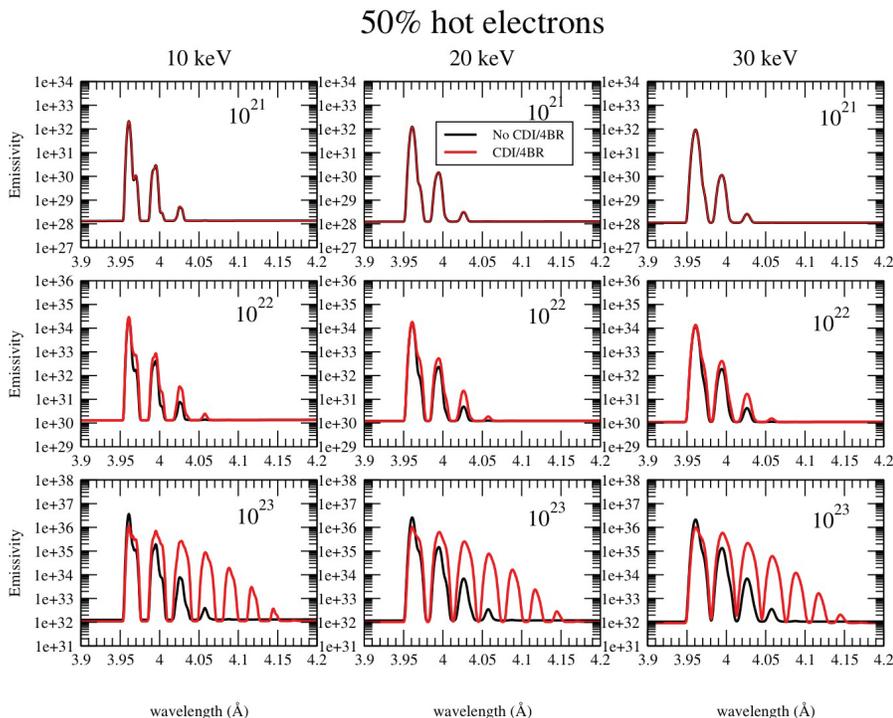
more quickly in this region. This effect is also seen in the average ionization stage, which is almost two units lower when CDI/4BR is included at a density of  $10^{23}/\text{cm}^3$ . Similar effects are also observed when lower fractions of hot electrons are included.

*For more information contact James Colgan at [jcolgan@lanl.gov](mailto:jcolgan@lanl.gov).*

- [1] R. Mancini, personal communication, 2006.
- [2] B. Streufert and J. Colgan, *Los Alamos Summer School Research Reports* 8 129, Los Alamos National Laboratory report LA-UR-06-5939 (2006).
- [3] B.J. Archer, et al., "Los Alamos Memorandum," Los Alamos National Laboratory report LA-UR-02-1526 (2002).
- [4] N.H. Magee, et al., "Atomic Processes in Plasmas," in *14th APS Topical Conference Proceedings*, (J.S. Cohen, S. Mazevet, and D.P. Kilcrease, Eds.), Santa Fe, NM, pp. 168–180 (2004).

### Funding Acknowledgements

NNSA's Advanced Simulation and Computing (ASC), Materials and Physics Program.



**Fig. 1.** Emission spectra of Ar at various electron densities.

## Quantum Molecular Dynamics Calculations of Opacities of Mixtures

Dan Horner, Lee Collins, T-4; Joel Kress, T-12

Opacities play a crucial role in characterizing a diverse set of phenomena including nuclear weapons, astrophysical objects, and laser plasmas. The associated environments generally consist of a mixture of species. Therefore, accurate mixture opacities become vital ingredients to many macroscopic modeling efforts, especially radiation-hydrodynamics. Most composite opacities employed in such models are determined from combination prescriptions based on independent atomic optical data. Such mixing rules ignore basic inter-atomic processes, which become especially important in the warm, dense matter regime. To examine the importance of such processes, we have systematically studied mixture opacities under a set of varying approximations from a fully quantum mechanical treatment of the mixture to simple atomic addition rules.

We determine the basic characteristics of the dense fluid (plasma) by quantum molecular dynamics (QMD) simulations, which provide a consistent set of static (equation-of-state), dynamical (diffusion/viscosity), and optical (opacities/conductivities) properties. All the electrons for the interacting species receive a full quantum mechanical treatment using sophisticated finite-temperature density functional theory (FTDFT). The electronic charge density for the entire system consistently relaxes to a completely integrated solution. All intra-atomic and inter-atomic interactions are represented. For the nuclei, a classical treatment suffices at the temperatures considered. The system thus evolves by a two-step process: 1) a

FTDFT calculation for a set configuration of atoms, and 2) movement of the nuclei according to the force determined. At a selected set of time steps, we perform calculations of the electrical conductivity. From this basic quantity, all other optical properties derive such as frequency-dependent absorption coefficients, indices of refraction, and reflectivities. By integrating the inverse of the absorption coefficient over the derivative of the Planck distribution, we produce the Rosseland Mean Opacity  $\kappa_R$ .

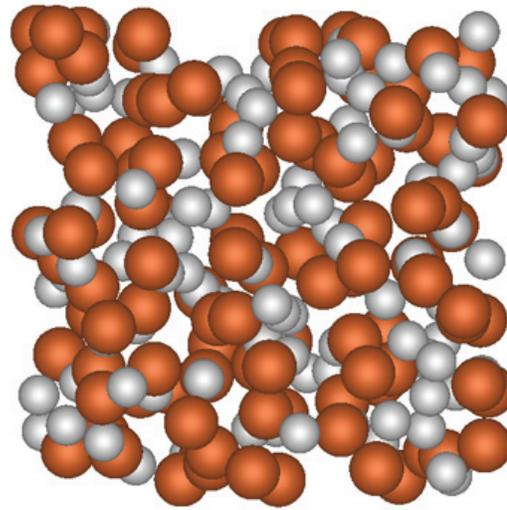
For the study of the opacities of mixtures, we have employed two levels of QMD calculations. In the first (QMD-mixture) we treat the entire composite as one system, allowing a complete intermixing of all the species. This gives the best possible value of the optical properties and forms the benchmark against which other schemes are judged. In the second procedure (QMD-species), we separately treat by QMD each of the pure species forming the mixture at the same temperature and appropriate number density. We then add the individual absorption coefficients and perform the integration over the Planck function to give a "mixed" opacity. This will serve as a best-case example for a mixing scheme based on combining the results of pure species calculations. Mixed opacities based on atomic results involve additional approximations.

As an example, we treat a LiH sample with equal fractions of the two species. In Fig. 1, we present a snapshot from the MD simulation that shows the complex intermingling of the species under conditions of high density ( $1.57 \text{ g/cm}^3$ )

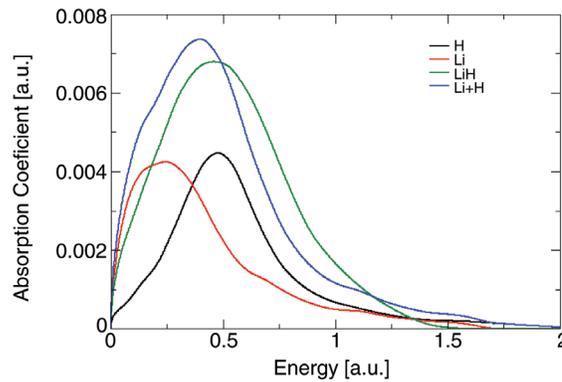
and low temperature (0.5 eV). Figure 2 displays the absorption coefficients calculated for the pure species H (black) and Li (red), for the sum of the two pure species (blue), and for the integrated mixture (green). Even for this best-case mixing procedure, the pure species sum departs both in magnitude and in peak position from the full LiH fluid case. In Table 1, we make comparisons of the Rosseland Mean Opacities. Even for this highly averaged property, significant differences arise between the mixture and the pure-species sums.

For more information contact Lee Collins at lac@lanl.gov.

**Funding Acknowledgements**  
 NNSA's Advanced Simulation and Computing (ASC), Materials and Physics Program.



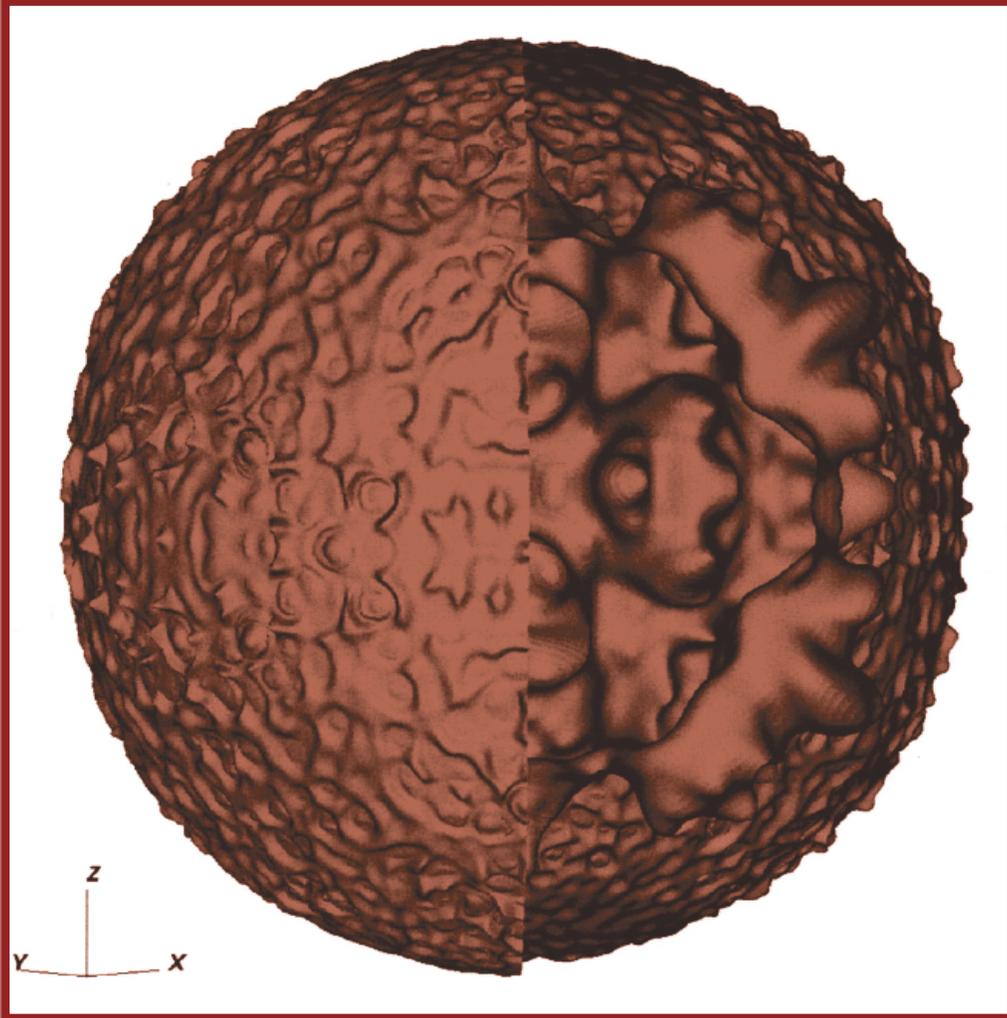
**Fig. 1.** Snapshot from a MD simulation of the LiH mixture.



**Fig. 2.** Absorption coefficients as a function of photon energy: a) pure H (black); b) pure Li (red); and c) sum of pure coefficients (H + Li: blue), mixture LiH (green).

**Table 1.** Rosseland Mean Opacities as a function of density and temperature for pure species (H and Li), for addition of pure species (H + Li), and for full mixture (LiH). The brackets indicate exponential power.

Density / T	$\kappa_R$ [1/cm]			
	H	Li	H+Li	LiH
1.58g/cc				
T = 0.5 eV	1.35[+5]	5.84[+5]	7.21[+5]	4.6[+5]
T = 1.0 eV	2.62[+5]	5.63[+5]	8.23[+5]	7.6[+5]
T = 2.0 eV	3.13[+5]	5.30[+5]	8.43[+5]	9.8[+5]
0.78 g/cc				
T = 0.5 eV	2.29[+4]	5.27[+5]	5.55[+5]	4.2[+5]
T = 1.0 eV	1.13[+5]	4.98[+5]	6.14[+5]	5.6[+5]
T = 2.0 eV	1.69[+5]	4.74[+5]	6.43[+5]	5.6[+5]



# Nuclear Physics

The work on nuclear physics highlighted in this section involves collaborations among researchers in Los Alamos Neutron Science Center (LANSCE), and Applied Physics (X) and Theoretical (T) divisions. The articles describe a gamut of applied work, such as the evaluation of cross sections for detector development and other applications, to somewhat longer-range projects such as the calculation of fission barriers.

## Iridium Cross-Section Evaluations

Patrick Talou, Toshihiko Kawano, Shannon Cowell, T-16; Ronald O. Nelson, Matthew Devlin, and Nik Fotiades, LANSCE-NS

A suite of activation detectors, also known as radiochemical detectors, was used during nuclear tests at the Nevada Test Site. Strategically loaded, these detectors permitted integral performance diagnostics including energy profiles of the neutron fluence as well as yields. The Iridium isotopes ( $^{191,193}\text{Ir}$ ) constitute a set of such detectors.

The usefulness of activation detectors as probes depends sensitively on accurate cross sections of neutron-induced reactions,  $(n,xn)$ . A vital component of modern reaction models is a detailed knowledge of the nuclear structure of detector nuclei. Until recently, there was little known about the nuclear excited states of Ir isotopes and precise cross-section evaluations were not possible. Recent experimental results, including those from GEANIE/LANSCE, have significantly improved our understanding of the Ir nuclear structure (see [1]). These improvements, together with recent advances in reaction models, afforded the opportunity to re-evaluate the relevant  $(n,xn)$  Ir cross sections.

Utilizing GNASH [2], a nuclear reaction code based on the Hauser-Feshbach theory of compound nuclei, new Ir  $(n,xn)$  cross sections have been obtained and incorporated into the ENDF libraries. Figure 1 gives an example of the new evaluations and compares the  $^{191}\text{Ir}(n,2n)^{190}\text{Ir}$  cross section to previous evaluations and experiment [3–6]. By comparing the 1999 calculation (red) by Chadwick and Hayes [7] with the modern evaluation (black), it is clear that

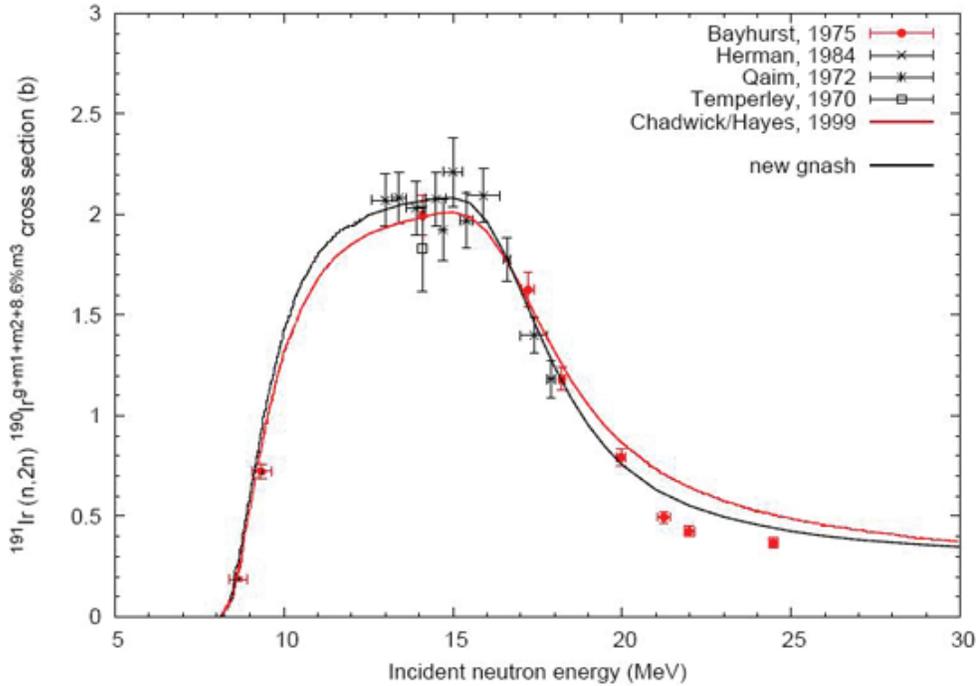
the current model better represents the high-energy data obtained by Bayhurst [3]. Though the improved nuclear structure is a vital component, the high-energy cross section is most affected by recent improvements in modeling of the pre-equilibrium nucleus.

Though experimental data is well reproduced, it is important for many applications to understand the uncertainties associated with model evaluations. GNASH incorporates several models, and detailing the uncertainties is difficult; methods are still under development. Currently, error estimates are obtained based on model parameter uncertainties and experimental data. An example is shown in Fig. 2 using the  $^{191}\text{Ir}(n,2n)^{190}\text{Ir}$  cross section (black). One-sigma “high” and “low” uncertainty quantifications (blue) are shown and in general agree with available experimental data.

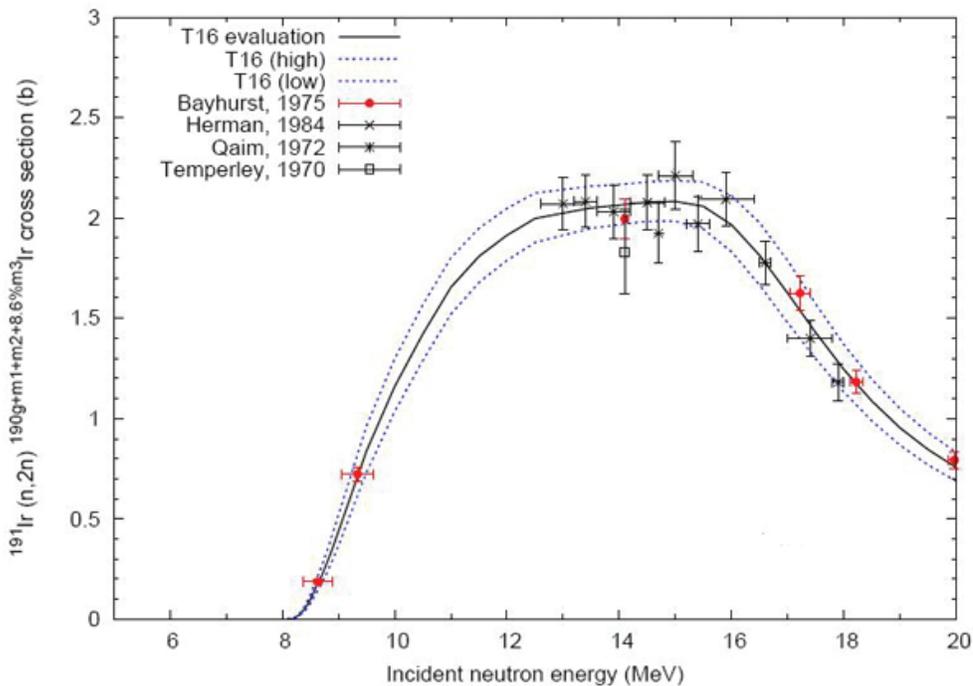
**For more information contact Shannon Cowell at [scowell@lanl.gov](mailto:scowell@lanl.gov).**

- [1] P. Talou, et al., “High-Spin Isomer states of Iridium and Gold” in this volume on p. 148
- [2] Proceedings of the Workshop: “Nuclear Reaction Data and Nuclear Reactors - Physics, Design, and Safety - Vol. 1,” International Centre for Theoretical Physics, Trieste, Italy, 15 April – 17 May, 1996, Ed. A. Gandini and G. Reffo (World Scientific, Singapore, 1999) pp. 227–404.
- [3] B.P. Bayhurst, et al., *Phys. Rev. C* **12**, 451 (1975).
- [4] M. Herman, et al., *Nuc. Phys. A* **430**, 69 (1984).
- [5] S.M. Qaim, *Nuc. Phys. A* **185**, 614 (1972).
- [6] J.K. Temperley and D.E. Barnes, BRL-1491 (1970).
- [7] M.B. Chadwick and A.C. Hayes “Iridium as a Radiochemical Detector for Neutron Spectrometry,” Los Alamos National Laboratory report LACP-99-226 (1999).

**Funding Acknowledgements**  
 NNSA's Advanced Computing and Simulation (ASC) Program.



**Fig. 1.**  
 An example of the new evaluations and comparison of the  $^{191}\text{Ir}(n,2n)^{190}\text{Ir}$  cross section to previous evaluations and experiment [3–6].



**Fig. 2.**  
 An example using the  $^{191}\text{Ir}(n,2n)^{190}\text{Ir}$  cross section (black). One-sigma "high" and "low" uncertainty quantifications (blue) are shown and in general agree with available experimental data.

## High-Spin Isomeric States of Iridium and Gold

Patrick Talou, Toshihiko Kawano, Shannon Cowell, T-16; Ronald O. Nelson, Matthew Devlin, and Nik Fotiades, LANSCE-NS

**I**ridium ( $^{191,193}\text{Ir}$ ) and Gold ( $^{197}\text{Au}$ ) isotopes served as radiochemical detectors during nuclear tests at the Nevada Test Site. A precise knowledge of the neutron-induced cross sections for the detector nuclei is crucial for extracting the energy components of the neutron fluence as well as yield information. Complete experimental data is not available for all relevant reactions and evaluated cross sections must be used. A vital component of modern reaction models is a detailed knowledge of the nuclear structure of the reaction nuclei. However, there was little known about the long-lived isomeric states of  $^{191,193}\text{Ir}$  and  $^{197}\text{Au}$  until recent experimental results from GEANIE/LANSCE. Figure 1 shows the details of the isomer states as well as the states that gamma transition to them. Experimental data for the production of these high-spin isomer states provide a unique opportunity to test and improve modern reaction models.

GNASH calculations [1], which utilize the Hauser-Feshbach (HF) compound nucleus theory, were performed to obtain gamma-ray production cross sections for the four dominant reactions feeding the  $11/2^-$  isomer of  $^{191,193}\text{Ir}$  and  $^{197}\text{Au}$ . In HF, a projectile incident on a target produces a compound nucleus (CN) that is long-lived enough to have reached equilibrium. From equilibrium, the CN decays via a series of gamma and particle emissions. GNASH calculates detailed level populations for each nuclear state at each step of the CN decay. As these populations determine all relevant gamma-ray and particle emission cross

sections, accurate nuclear structure information is vital.

Figure 2 shows gamma-ray production cross sections for the states feeding the isomer of  $^{197}\text{Au}$ . There is reasonable agreement between experiment (triangles) and theory (solid line). The cross sections from “high-spin” excited states, (594.4 keV and 358 keV transitions) are very well reproduced. The pre-equilibrium model used by GNASH significantly effects these transitions. By implementing the quantum mechanical theory of Feshbach-Kerman-Koonin, GNASH better predicts these high-spin transitions. The production cross sections from “low-spin” excited state, (174.8 keV and 538.6 keV) are often underestimated for incident neutron energies between 2–7 MeV. It is unclear if these discrepancies are primarily due to model deficiencies or assumptions made in the data analysis.

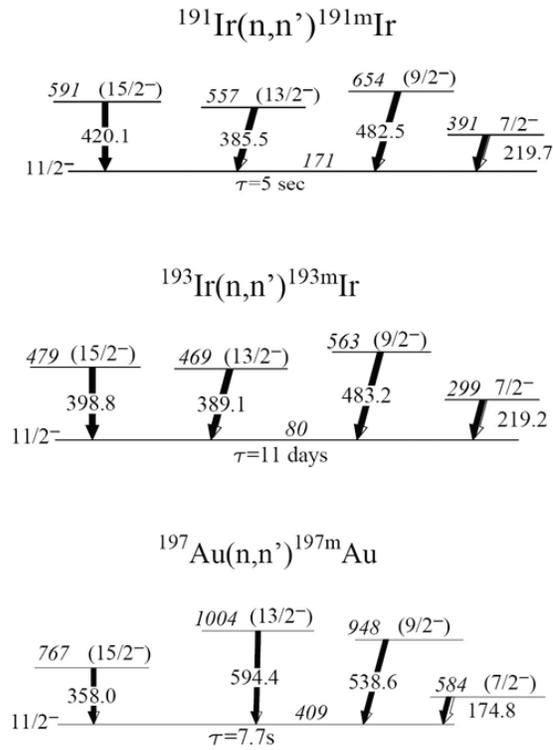
The improvements to evaluated neutron cross sections is discussed in another report, “Iridium Cross-Section Evaluations” [2].

*For more information contact Shannon Cowell at [scowell@lanl.gov](mailto:scowell@lanl.gov).*

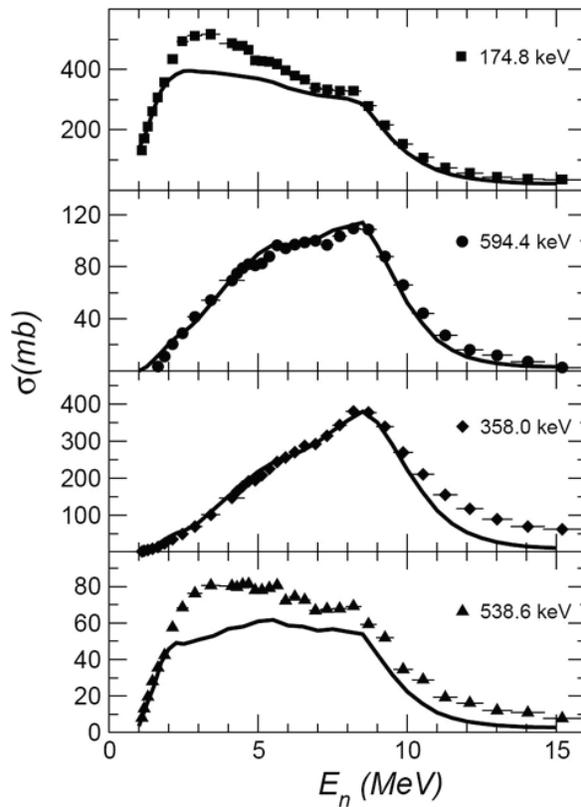
[1] Proceedings of the Workshop: “Nuclear Reaction Data and Nuclear Reactors - Physics, Design, and Safety, Vol. 1,” International Centre for Theoretical Physics, Trieste, Italy, 15 April – 17 May, 1996, Ed. A. Gandini and G. Reffo (World Scientific, Singapore, 1999) pp. 227–404.

[2] P. Talou, et al., “Iridium Cross-Section Evaluations,” in this volume on p. 146.

**Funding Acknowledgements**  
 NNSA's Advanced Computing and Simulation (ASC) Program.



**Fig. 1.** The details of the high-spin isomer states as well as the states that gamma transition to them. Experimental data for the production of these isomer states provide a unique opportunity to test and improve modern reaction models.



**Fig. 2.** Gamma-ray production cross sections for the states feeding the isomer of  $^{197}\text{Au}$ .

## Measurement and Analysis of the ${}^6\text{Li}(n,t){}^4\text{He}$ Cross Section

Gerald Hale, T-16; Matthew Devlin, LANSCE-NS

**T**he reaction  ${}^6\text{Li}(n,t){}^4\text{He}$  has important applications as a neutron flux monitor, and as a breeder of tritium. Surprisingly, although its cross section has been well measured in the region below a few hundred keV where it is used as a standard, it is not so well known at energies above one MeV. A high-priority measurement at the Weapons Neutron Research (WNR) facility has been to determine this cross section and its angular distributions in the MeV energy range.

Preliminary values of the measured angular distributions have been obtained recently from LANSCE-NS. The overall normalization of the data set is yet to be determined, but the relative information is almost final. This information has been included in a large, multichannel R-matrix analysis of reactions in the  ${}^7\text{Li}$  system in order to obtain better information about the  ${}^6\text{Li}(n,t)$  integrated cross section. The fits to some of the angular distribution data are shown in Fig. 1. The shape of the distribution around 2 MeV confirms the presence of the  $J^\pi=3/2^-$  resonance that had been obtained earlier, and indicates that a pronounced “shoulder” in the integrated cross section due to that resonance is more than 5% higher than obtained from the previous data. This can be seen in Fig. 2.

The angular distributions in Fig. 1 also indicate that additional resonances may contribute in the energy range above 2 MeV. We will continue the analysis as the final data become available, with the

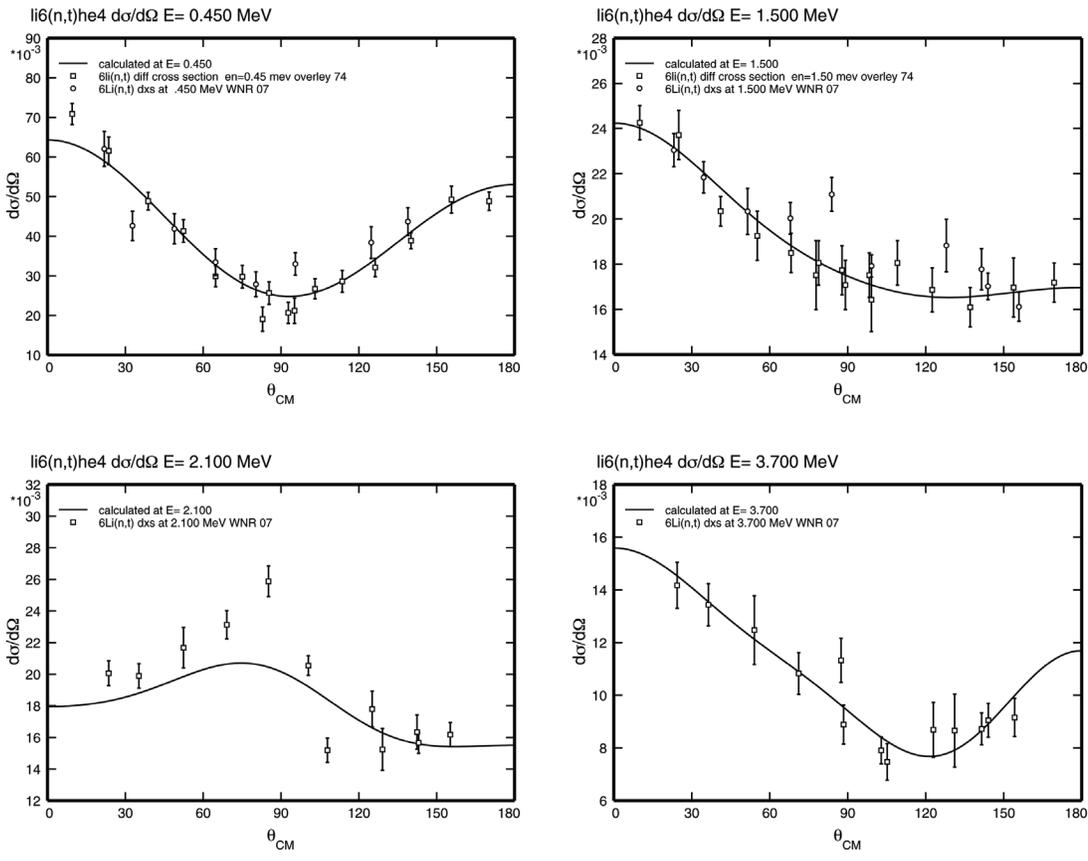
goal of providing a complete new evaluation of the  $n+{}^6\text{Li}$  cross sections, including covariances, for NW applications.

*For more information contact Gerald Hale at [ghale@lanl.gov](mailto:ghale@lanl.gov).*

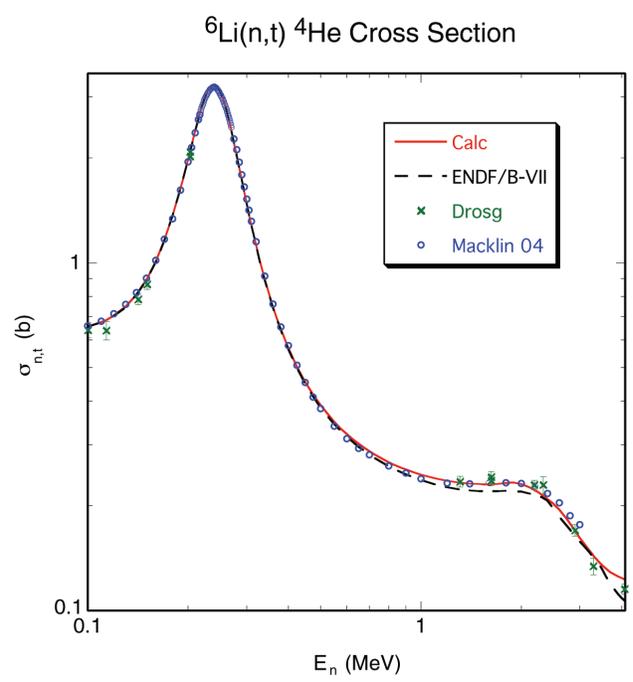
- [1] J.C. Overley, et al., *Nucl. Phys. A* **221**, 573 (1974).
- [2] M. Drosch, et al., *Nucl. Inst. Meth. B* **94**, 319 (1994).
- [3] R.L. Macklin, et al., *Nucl. Sci. Eng.* **71**, 205 (1979); ratio converted using revised  ${}^{235}\text{U}(n,f)$  cross sections (2004).

### Funding Acknowledgements

NNSA's Advanced Simulation and Computing (ASC), Advanced Applications Program, and Campaign 4, Reactivity and Compression.



**Fig. 1.** Measured and calculated angular distributions for the  ${}^6\text{Li}(n,t)$  reaction at energies between 0.45 and 3.7 MeV. The preliminary WNR measurement is shown with the data of Overlay et al. [1].



**Fig. 2.** Integrated cross sections for the  ${}^6\text{Li}(n,t)$  reaction. The solid red line gives the calculation from the current fit. The dashed black line is the result from the recent ENDF/B-VII (and IAEA standards) evaluation. The measurements shown are from Drogg [2] (green crosses) and from Macklin [3] (blue circles).

## Calculation of Charged-Particle Transport for NIF Mix Diagnostics

Anna Hayes, T-16; Gerard Jungman, T-6; Paul A. Bradley, X-2

The National Ignition Facility (NIF) at Lawrence Livermore National Laboratory is designed to demonstrate thermonuclear ignition of DT gas, with a net gain in energy, using the technology of indirect-drive inertial confinement fusion (ICF). Hydrodynamical instabilities result in mixing of the DT fuel with the outer shell material during the implosion of the capsule. We have been working on a new technique to diagnose the nature of mix occurring in ignited and partially failed NIF capsules.

Our mix diagnostic for NIF relies on the fact that energetic (knock-on) tritium nuclei from the fuel will interact with nuclei from the shell material, and these reactions are very sensitive to the amount and type of mix taking place. For example, triton reaction rates on the shell material can be increased by two orders of magnitude by mixing between the shell and the DT gas compared to a fiducial no-mixing case. These reactions produce high-energy  $\beta$ -emitters, and we presently have a Laboratory Directed Research and Development (LDRD) project to develop  $\beta$ -spectroscopy as a diagnostic for mix at NIF.

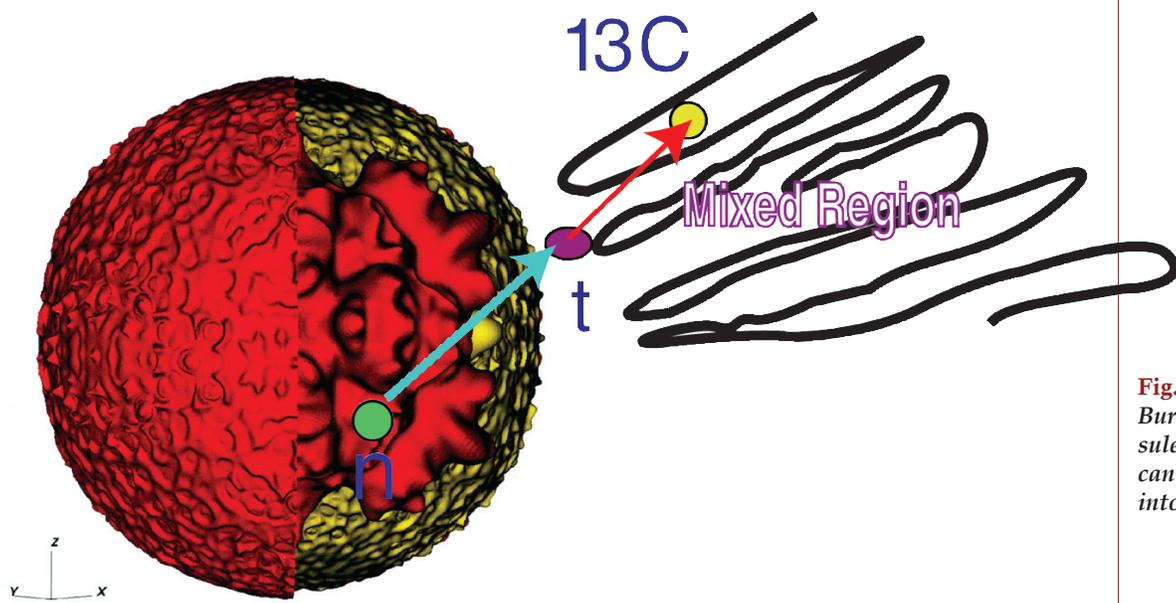
An ongoing experimental program at the Omega facility at Los Alamos National Laboratory has been fielded as a proof of principle for this idea. The present experiments use  $^{13}\text{C}$  as the capsule shell, although at NIF  $^9\text{Be}$  and/or plastic shells will be used. For these experiments the key reaction of interest is the  $^{13}\text{C}(t,\alpha)^{12}\text{B}(\beta^-)$  reaction. We ran a series of calculations for  $^{13}\text{C}$  ICF capsules using

the LASNEX code. The yields of these capsules are about  $10^{13}$ . We have developed the computational tools to determine accurately the production of knock-on tritons, their transport in the burning capsule, and subsequent yields for triton-induced diagnostic reactions (see Fig. 1). We coupled the output of LASNEX with our new charged-particle transport post-processor.

*For more information contact Anna Hayes at [anna\\_hayes@lanl.gov](mailto:anna_hayes@lanl.gov).*

### Funding Acknowledgements

Laboratory Directed Research and Development Program.



**Fig. 1.**  
*Burning DT NIF capsule: 14 MeV neutrons can knock the tritons into the  $^{13}\text{C}$  mix region.*

## Calculations of the Proliferation-Resistance of Recycled Fuel

Anna Hayes, Bill Wilson, T-16; Holly Trelue, X-1

One of the central concepts for the Global Nuclear Energy Partnership (GNEP) is the reduction of proliferation risks while expanding the global nuclear power industry. Achieving this goal requires an advanced reprocessing and recycling technology for the spent fuel. The present scheme being considered does not separate plutonium from other long-lived radioactive elements and thus requires detailed modeling to characterize the fuel.

In an effort to quantify the proliferation issues, we have been modeling the nuclear burn of these advanced fuels using the Los Alamos reactor burn code MonteBurns, which couples the MCNP transport code and the nuclear transmutation code CINDER'90. With this tool we can provide an accurate description of a broad class of fuels, including the reactor spatial and temporal power, the fuel composition, and the radiation and decay heating. We characterize the waste content and emission and examine several non-pure plutonium fuels in fast reactors to determine radiation signatures for advanced safeguards.

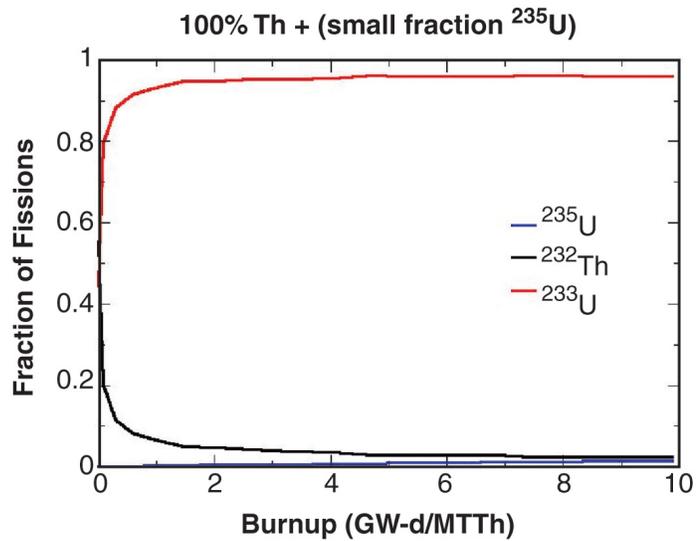
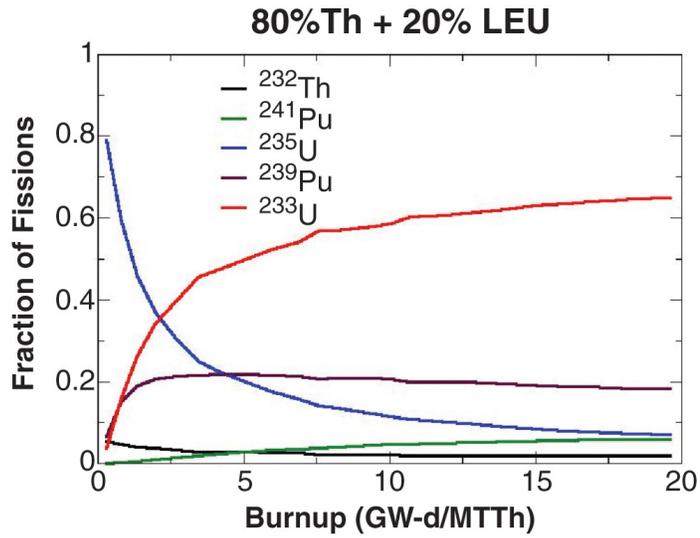
We have determined and verified the characteristic radiation signatures of irradiated fuel assemblies for Th-U fuels (see Fig. 1). These calculations provide detailed spectra of neutrons, gamma, betas, and charged particles emitted from the spent fuel as a function of cooling

and/or storage time. We have examined the proliferation index of these fuels and shown that unless 20% of lightly enriched uranium is included in the fuel, the proliferation index is violated.

*For more information contact Anna Hayes at [anna\\_hayes@lanl.gov](mailto:anna_hayes@lanl.gov).*

### **Funding Acknowledgements**

Department of Energy (DOE) Office of Research and Engineering (NA22) and Laboratory Directed Research and Development Program.



**Fig. 1.** Fuel burnup for the Th-U fuel. The curves show the fraction of the different isotopes burning as a function of time. (top) Shows the burn for normal Th-U fuel. (bottom) Shows the burn for the Th-U fuel that violates the proliferation index. Units in labels of x-axis = Giga Watt-day/metric ton of Th.

## Large-Scale Calculation of Fission Barrier Parameters for 5254 Nuclei with $171 \leq A \leq 330$

Peter Möller, Arnold J. Sierk, T-16

In previous *Theoretical Division Nuclear Weapons Highlights* issues, we have discussed calculations of fission potential-energy surfaces as functions of up to five different nuclear shape coordinates as the system evolves from a ground-state shape to two separated fission fragments. The five shape coordinates we consider beyond the second minimum in the fission barrier (fission isomer) are elongation, neck radius, spheroidal deformations of the emerging left and right fragments, and left-right mass asymmetry. For less deformed shapes between the spherical shape and the second minimum in the barrier we consider three shape coordinates in the Nilsson perturbed-spheroid  $\epsilon$  parameterization:  $\epsilon_2$  (quadrupole),  $\epsilon_4$  (hexadecapole), and  $\gamma$  (axial asymmetry). Important parameters of the calculated nuclear potential-energy surface, such as the local minima, saddle points between all pairs of minima, valleys as functions of elongation, and the height of the ridges that separate the valleys are extracted by similar immersion techniques to those used in geography. Some illustrative results for a few nuclei have been presented in *Nuclear Weapons Highlights* reports in previous years.

This year we have used our models to undertake production calculations so that we obtain detailed descriptions of the fission potential-energy for a large number of nuclei. Specifically we calculated 5-D potential energy surfaces for 5254 nuclei between the proton and neutron drip lines from  $A = 171$  to  $A = 330$ . For each of these nuclei the energy is calculated for 5,009,325 different nuclear shapes. In addition we studied the effect of triaxiality at the first peak in the barrier by calculating 3-D potential-energy surfaces for these same nuclei in the  $\epsilon$  parameterization. We have developed highly automated scripts that use our immersion codes to extract relevant barrier structure parameters from

the calculated potential-energy surfaces. At this stage we have extracted the heights of the first and second barrier peaks and the fission-isomeric state for all 5254 nuclei.

Figure 1 shows a comparison between calculated and experimental barrier parameters for a sequence of uranium nuclei. The experimental data are from a review by Madland [1]. Sometimes the results of several experiments are plotted for the same isotope. We have made such comparisons for isotope chains of all elements from Th to Es ( $Z = 90$  to  $Z = 99$ ). Such detailed knowledge about the barrier structure and other nuclear structure properties are necessary to model  $(n,f)$ ,  $(n,2n)$ , and many other cross sections. The potential-energy calculations took 30000 CPU hours and the subsequent analysis 20000 CPU hours (so far) on the T-16 cluster “nuclei.”

To model many astrophysical scenarios; for example to model the end of the rapid-neutron-capture process (r-process) in which many heavy elements are formed in stars, it is necessary to know fission-barrier heights of a large number of nuclei. When a neutron is captured in the r-process it is energetically possible for the nucleus to fission if the neutron binding energy of the compound system is larger than the fission-barrier height. We display in Fig. 2 the difference between the calculated fission-barrier height and the one-neutron separation energy. When this quantity is negative it is energetically possible for the nucleus to fission.

The duration of the large neutron flux that powers the r-process is thought to be of the order of 1 s. Thermonuclear explosions also generate large fluxes of neutrons, but these are of shorter duration, and neutron capture

in this environment is therefore referred to as prompt neutron capture. However, this process can be thought of as an approximation to the r-process on earth [2], at least over a limited range of nuclei. Several nuclear weapons tests between 1952 and 1969 involved studies of the prompt capture process. In Fig. 2, we have indicated where capture chains starting with  $^{232}\text{Th}$ ,  $^{238}\text{U}$ , and  $^{242}\text{Pu}$  terminate [2]. We note that the U and Pu chains end very near the dark green area, where nuclei would fission immediately upon neutron capture. Thus, our results are very consistent with the experimental observations, and a very encouraging sign that our models produce reliable fission-barrier parameters far from stability. Consequently, another exciting application is to use our calculated database in modeling the end of the r-process, which runs far from stability in regions of experimentally inaccessible nuclei where the neutron separation energy is 1–2 MeV or so. The Th decay chain does not reach the dark green area. However, it was observed [2] that this capture chain is severely blocked at  $A = 242$  and  $A = 244$  where the neutron capture cross section becomes very small. A more detailed interpretation of the observations on these capture chains requires more elaborate calculations of the reaction paths using realistic neutron spectra, which are not all thermal.

In the near future we plan to configure our immersion codes to establish when the calculated potential-energy surfaces exhibit more than one fission mode, the character (mass symmetric division or mass asymmetric division) of the different modes,

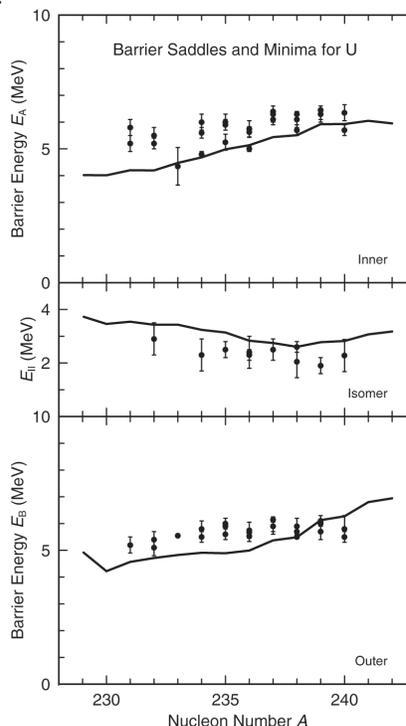
and the threshold saddle-point energies for these modes.

*For more information contact Peter Möller at [moller@lanl.gov](mailto:moller@lanl.gov).*

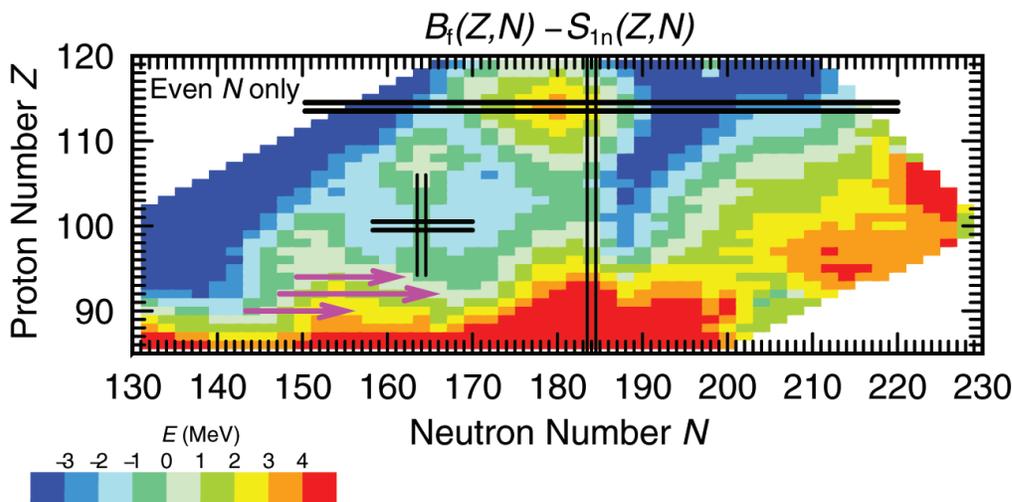
- [1] D.G. Madland, Los Alamos National Laboratory, personal communication, 2000.  
 [2] S.A. Becker, *Carnegie Observatories Astrophysics Series*, Vol. 4, A. McWilliam and M. Rauch, Eds. (Pasadena: Carnegie Observatories, <http://www.ociw.edu/ociw/symposia/series/symposium4/proceedings.html>).

### Funding Acknowledgements

NNSA's Advanced Simulation and Computing (ASC), Materials and Physics Program.



**Fig. 1.** Calculated height of the first  $E_A$  and second  $E_B$  peak in the fission-barrier and calculated energy  $E_{II}$  of the fission isomer for U isotopes, compared to experimental data where available. The calculated energies are all relative to the calculated energy of the ground-state minimum.



**Fig. 2.** Difference between calculated fission-barrier height and neutron separation energy. When this quantity becomes negative (dark green area) fission is energetically possible. We indicate in the figure the beginning and ending of these capture chains observed in weapons tests [2]. Two of the chains (U and Pu) end next to the dark green area where fission becomes energetically possible.

## Fission Fragment Excitation of Nuclear Isomers in Hot Dense Environments

Anna Hayes and Jim Friar, T-16; Gerard Jungman, T-6

**H**igh-energy fission fragments are born with kinetic energies of the order of 100 MeV, and, in principle, these fragments can cause nuclear reactions while moving through a sample of uranium. Typically, the fission fragments of  $^{235}\text{U}$  have nuclear charges  $Z \sim 40\text{-}50$ , which suggests that the barriers for Coulomb excitation are low enough to allow fast fragments to excite nuclear isomers. We have calculated the cross section for Coulomb excitation of the 26-minute, 77 eV isomer of  $^{235}\text{U}$ . For a typical fission fragment we found that the Coulomb excitation cross section is quite large and can be of the order of 10 barns.

**Fig. 1.**  
The cross section for Coulomb excitation of  $^{235}\text{U}$  by fast fission fragments is large. However, in high-density environments the fission fragments are stopped locally.

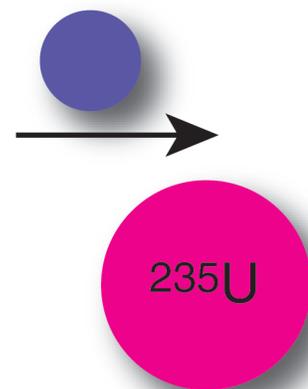
The probability that Coulomb excitation takes place depends on both the cross section and the mean free path of the fragment. We calculated the mean free path of fission fragments in extreme environments as a function of temperature and density (see Fig. 1). For very high densities we find that the fission fragments are stopped locally in the uranium (see Fig. 2). Thus Coulomb excitation of the isomer becomes very improbable, despite the large cross section for the process.

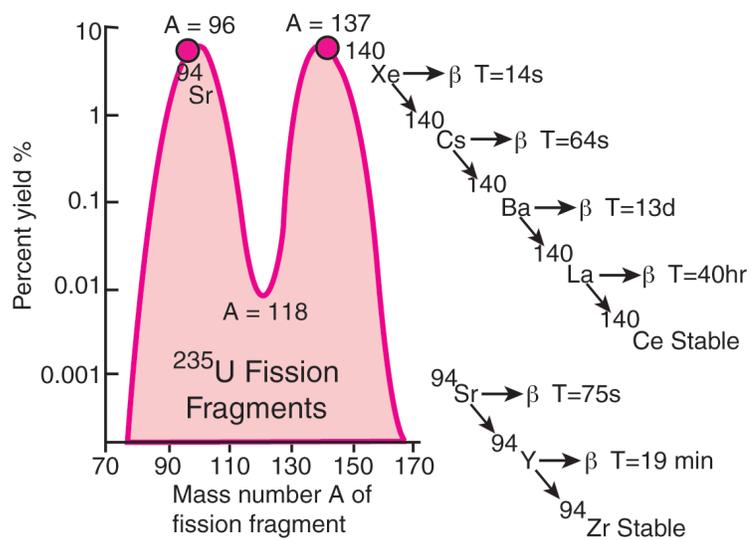
*For more information contact Anna Hayes at [anna\\_hayes@lanl.gov](mailto:anna_hayes@lanl.gov).*

### Funding Acknowledgements

NNSA's Campaign 4, Reactivity and Compression, and the Physics and Engineering Models Program.

K.E. = 100 MeV





**Fig. 2.**  
*Distribution of fission fragments for  $^{235}\text{U}$ .*

## Improved Neutron Fluence Spectra for Simulations of the Role of the Isomer of U-235

Anna Hayes, T-16; Gerard Jungman, T-6; Merri Wood-Schultz, X-2

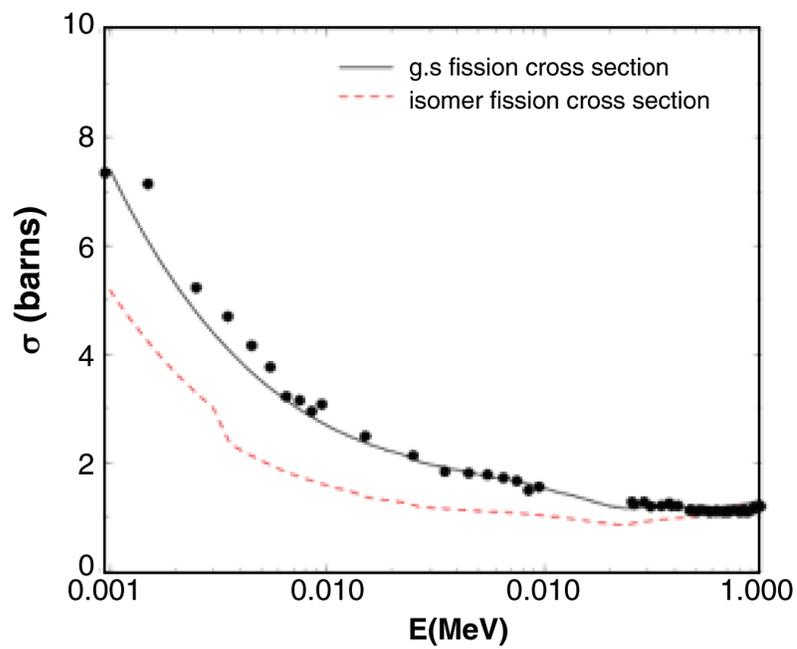
**T**he fission cross section for the isomer of  $^{235}\text{U}$  is predicted to be significantly lower than that for the ground state at neutron energies below 0.5 MeV. At these same neutron energies the neutron capture cross section is enhanced (see Fig. 1). An accurate determination of the role of this isomer in weapons simulations requires an accurate determination of the neutron fluence at these low energies.

For this purpose we carried out detailed 2-dimensional simulations for a dominantly  $^{235}\text{U}$  device using Monte Carlo transport of the neutronics. We found that the Monte Carlo treatment provided a significant improvement in the predicted low-energy neutron spectrum over other methods, mainly because of the improved treatment of neutron down-scattering processes. We determined the effect of including the isomer of  $^{235}\text{U}$  in the simulations on the yield and on the production of  $^{236}\text{U}$ .

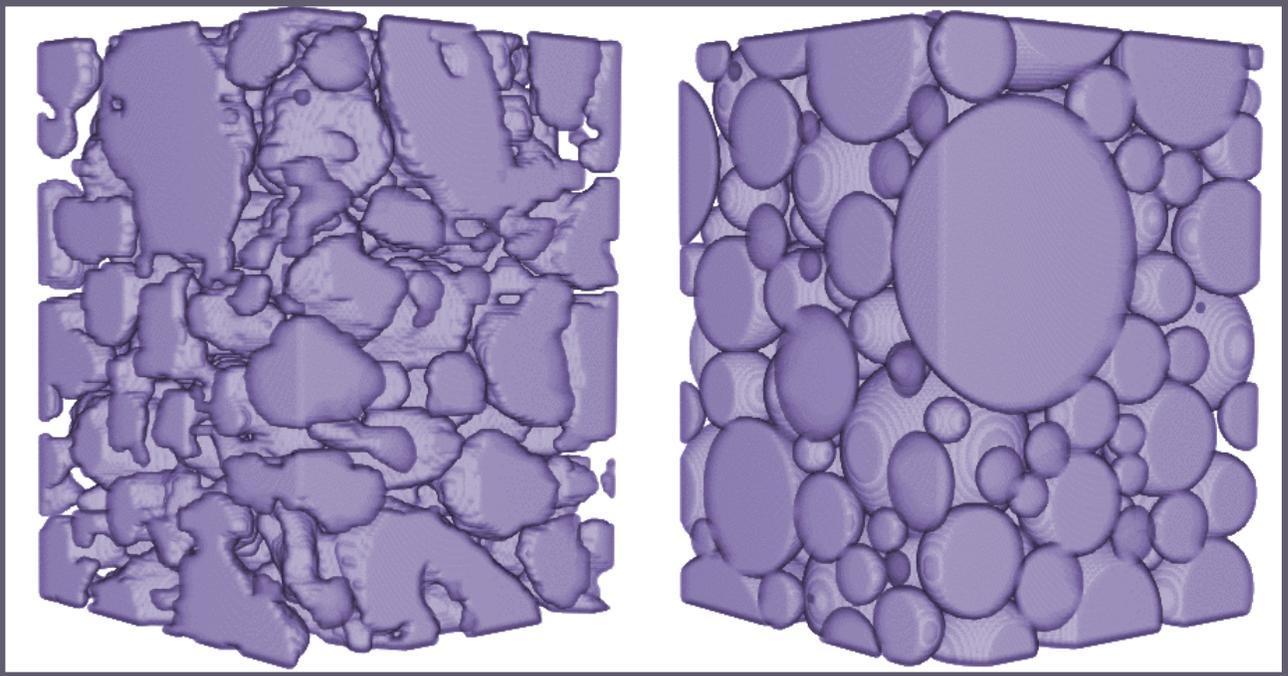
*For more information contact Anna Hayes at [anna\\_hayes@lanl.gov](mailto:anna_hayes@lanl.gov).*

### **Funding Acknowledgements**

NNSA's Campaign 4, Reactivity and Compression, and the Physics and Engineering Model Program, and Laboratory Directed Research and Development Program.



**Fig. 1.** Fission cross section for  $^{235}\text{U}$ . The black (line and dots) are for the ground state, and the red line is for the isomer.



# High Explosives and Organic Materials

Work in high explosives and organic materials has a long history in the nuclear weapons programs. In this issue we are highlighting work on the mechanical response of composite materials, where microstructure plays a crucial role, atomistic simulations of shock waves moving through polycrystalline materials, the structure of dislocations, foam mechanics, fluid dynamics, as well as the first molecular dynamics simulation with 320 billion atoms.

## Microstructure Effects on PBX Mechanical Response

Scott G. Bardenhagen, T-14

The mechanical response of plastic-bonded explosives (PBXs) is of interest in a variety of munitions and industrial applications. PBXs are composed of energetic grains embedded in a polymeric binder. The heterogeneity at this material scale localizes energy during deformation, resulting in damage nucleation sites and hot spots. These local events ultimately drive bulk material response, such as fracture/failure and deflagration/detonation. To develop predictive models, it is imperative to have a sound physical understanding of grain-scale material response and represent material-state heterogeneity in the models.

A fully 3-D characterization of material microstructure morphology may be obtained using x-ray microtomography. From a series of images collected as the sample is rotated, a 3-D image of the specimen is reconstructed. X-ray microtomography has been used to image the energetic material PBX-N109 [1]; an example grain morphology is depicted in Fig. 1.

Packings of spheres are common approximations to granular composite microstructures. These packings can be generated fairly easily, but they miss a distinguishing characteristic of PBXs, namely

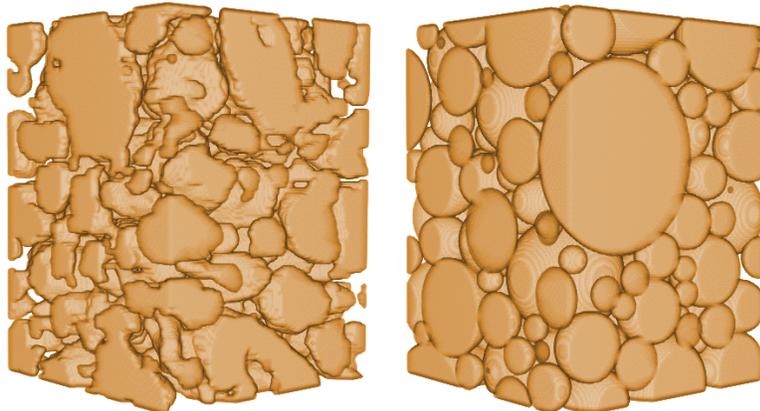
the irregularity of the grain shapes, and in particular the presence of corners that may cause regions of high local stress. An example is depicted in Fig. 1.

At Los Alamos National Laboratory, we performed computational experiments on various microstructures, and calculated bulk properties. These results are reported in Table 1 for both tomography data microstructures and spherical grains. “As-received” tomography data is labeled “raw.” Image processing techniques were used to soften features and remove pixel-scale noise in tomography data and this case is labeled “smoothed.” All cases give remarkably similar bulk properties.

It is concluded that grain shape may have very little effect on the initial bulk properties for an elastic, perfectly bonded mixture. Certainly, packed spheres provide an excellent surrogate for determining bulk properties. The grain volume fraction is the most important parameter by far. While somewhat surprising, the result is consistent with analyses using idealized microstructures. Good estimates of bulk properties (within 5%) may be obtained with surprisingly small RVEs (containing eight or fewer spheres).

Histograms were used to compare material state for different microstructures and assess the importance of grain shape to the distribution of strain-state statistics. Histograms of strain state for the cases in Table 1 are displayed in Fig. 2.

Fig. 1. Examples of grain morphologies used in computations. On the left is x-ray microtomography data for PBX-N109, on the right a packing of spheres.



**Table 1.** Comparison of Bulk Properties

Cases	Volume Fraction HMX	Bulk Modulus (GPa)	Shear Modulus (GPa)
Tomography (smoothed)	63.7%	9.61	3.79
Tomography (raw)	65.0%	9.65	3.80
Tomography (raw)	64.1%	9.61	3.79
Tomography (raw)	65.1%	9.65	3.80
Tomography (raw)	66.4%	9.66	3.80
Spheres	66.4%	9.67	3.80

These histograms are remarkably similar, again illustrating insensitivity to microstructural details for elastic, perfectly bonded mixtures. However, some differences may be seen in the main peak and the shoulder to its left. These features correspond to more compressive strain in the grains, and less in the binder, for packed-sphere microstructures. This difference is due to the microstructure. While in the tomography data each grain tends to be coated with dirty binder, in the case of the packed spheres, most spheres touch one another initially. This connected network of spheres provides an embedded load-bearing structure that carries slightly more load and shields the dirty binder.

The same data as Fig. 2 are replotted on a log scale in Fig. 3 to examine differences in the tails of the distributions. It is in the tails that the effect of irregular grain corners might be expected to become apparent, appearing as an increase in the frequency of large strains. The tails of the distributions for the packed-sphere microstructure are clearly shifted slightly toward larger strains than the tomography microstructures. This is precisely the opposite of what would be expected based on grain shape alone. The shift is again due to the connected substructure of contacting spheres, and strain concentrations associated with contact mechanics. The binder coating on the grains in the tomography data pads grain contacts and reduces strain concentrations.

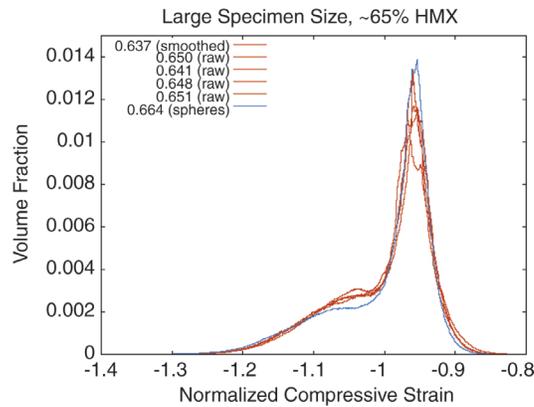
These considerations suggest enhanced sensitivity of the packed-sphere microstructure as one might expect simply from contact mechanics playing a large role. Some micrographs of PBXs suggest the larger grains are well coated with dirty binder in PBX-9501 [2], which should provide a desensitizing effect.

*For more information contact Scott Bardenhagen at bard@lanl.gov.*

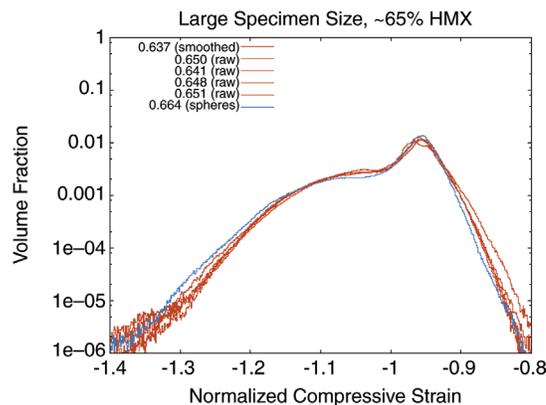
- [1] S. Lecume, et al., "Structure of Nitramines Crystal Defects Relation with Shock Sensitivity" in *Proceedings of the 35th International Conference of ICT* (Karlsruhe, Germany, 2004) pp. 2-1-2-4.
- [2] B.E. Clements and E.M. Mas, *Model. Sim. Mat. Sci. Eng.* **12**, pp. 1-15 (2004).

**Funding Acknowledgements**

NNSA's Advanced Simulation and Computing (ASC), Materials and Physics Program; and Campaign 2, Dynamic Materials Properties.



**Fig. 2.** Strain-state histograms.



**Fig. 3.** Strain-state histograms.

## Atomistic Simulations of Shock Waves in Polycrystalline Iron

Kai Kadau, T-14; Timothy C. Germann X-1-SMMP; Peter S. Lomdahl, T-14; Robert C. Albers, T-11; J.S. Wark and A. Higginbotham, University of Oxford; Brad L. Holian, X-1-SMMP

The propagation of shock waves through a polycrystalline iron sample is explored by large-scale atomistic simulations [1,2]. For large enough shock strengths the passage of the wave causes the body-centered-cubic (bcc) structure to transform into a close-packed structure with most structure being isotropic hexagonal-close-packed (hcp) and, depending on shock strength and grain orientation, some fraction of face-centered-cubic (fcc) structure [3]. The simulated shock state as represented by the Hugoniot is compared to experimental data (see Fig. 1). By calculating the extended x-ray absorption fine structure (EXAFS) directly from the atomic configurations obtained by our simulations [4], a comparison to recent experimental EXAFS measurements of nanosecond-laser shocks in polycrystalline iron shows that the experimental data are consistent with a phase transformation. However, the atomistically simulated EXAFS spectra also show that an experimental

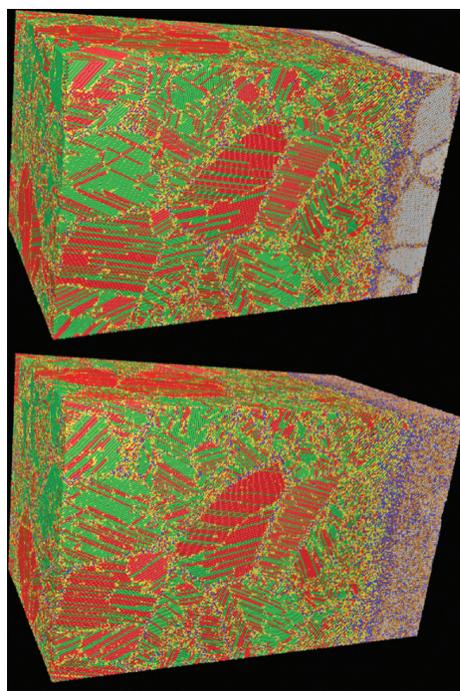
distinction between a product hcp or fcc phase is not possible based on the EXAFS spectra alone.

Figure 2 shows samples shocked with  $u_p = 0.906$  km/s, 14.6 ps after the impact. Top (bottom) sample has an initial temperature of 50K (300K) and reaches a temperature behind the shock front of 296K (622K), pressures of about 39 GPa, and volume compressions of 19%. The samples consist of 32 grains and about 30 M atoms confined in a 57.4 nm x 7.4 nm x 109.9 nm box. Color coding denotes the local neighborhood of each atom: gray: bcc; blue: uniaxially compressed bcc; yellow: grain-boundary; red: hcp; and green: fcc. Fluctuations in the bcc structure at 300 K (bottom) wash out the color-scheme analysis, and therefore make it hard to see the grain boundary structure between the bcc grains. The hcp/fcc ratio behind the shock front is on average 1.5 for this shock strength and increases with decreasing shock strength. The initial average grain size—as defined by the average caliper diameter—is 32.7 nm.

The present simulations of shock waves in polycrystalline iron demonstrate how atomistic modeling can enhance our understanding of ultrafast dynamical processes that take place in shock-induced phase transformations. Interpretations of experimental data can be scrutinized and checked with the simulation data set (see Fig. 3).

Specifically, atomistically simulated and experimental EXAFS spectra both support the conclusion that a phase transformation polycrystalline iron takes place under shock loading, yielding an almost isotropic close-packed product phase. However, the simulations also reveal the possibility of the product having some sizable fraction of metastable fcc product, as opposed to purely

**Fig. 1.** Experimental and simulated Hugoniots for polycrystalline iron. The full line is a linear fit to the experimental data [LANL data by Brown and McQueen], the dashed lines are fits to the three different parts of the simulated Hugoniots: elastic precursor, transformation wave, and the overdriven part. As a reference, the  $[001]_{\text{bcc}}$  single crystalline shock data are included.



hcp. This question cannot be resolved from the experimental EXAFS spectra alone, since the noise-levels for those ultrafast in-situ measurements are too large to differentiate between hcp and fcc. We have to emphasize that details of the interatomic potential change the hcp to fcc ratio in the simulations. However, all tested EAM potentials for iron show a sizable fraction of fcc in the product phase under shock loading. Since the hcp/fcc ratio within a grain decreases the more the shock direction deviates from the  $[001]_{\text{bcc}}$  direction with respect to the initial polycrystal, we speculate that there is a geometric component to this effect. The ABC stacking sequence along  $[111]_{\text{bcc}}$  makes it easier to transform under rapid compression in this direction into the ABC-stacked fcc structure instead of the AB-stacked hcp phase. For shocks along  $[001]_{\text{bcc}}$  the AB-stacked close-packed bcc planes transforming into the AB-stacked close-packed hcp planes more easily than into the ABC-stacked fcc planes. Fully resolved *ab-initio* molecular dynamics (MD) simulations would allow for a more reliable prediction of the hcp to fcc ratio of the product phase. However, such methods are presently far too computationally intensive for the large scales needed for a polycrystalline system. Our nonequilibrium MD simulations were on the order of some 10 ps, still too short to detect further relaxation processes that might occur on longer time scales. We leave the final clarification of the detailed structure and relaxation of the product phase to planned high-energy laser-based experiments that can access time scales on the order of nanoseconds.

For more information contact Kai Kadau at [kkadau@lanl.gov](mailto:kkadau@lanl.gov).

- [1] D.H. Kalantar, et al., *Phys. Rev. Lett.* **95**, 075502 (2005).
- [2] K. Kadau, et al., *Science* **296**, 1681 (2002).
- [3] K. Kadau, et al., *Phys. Rev. B* **72**, 064120 (2005).
- [4] J.J. Rehr, et al., *Phys. Rev. Lett.* **58**, 3397 (1992).

#### Funding Acknowledgments

NNSA's Advanced Simulation and Computing (ASC), and Laboratory Directed Research and Development.

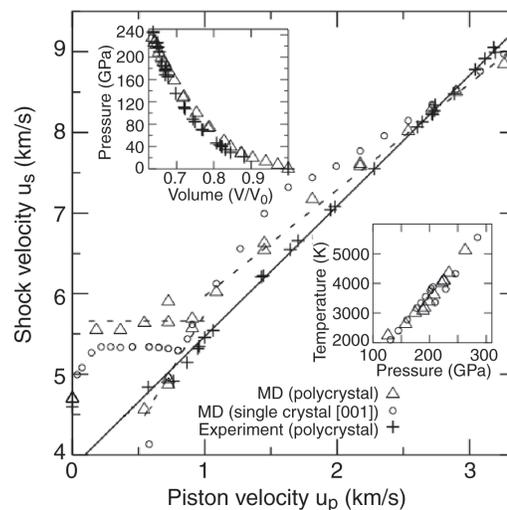


Fig. 2. Samples shocked with  $u_p = 0.906$  km/s, 14.6 ps after the impact. Please see the text for more details.

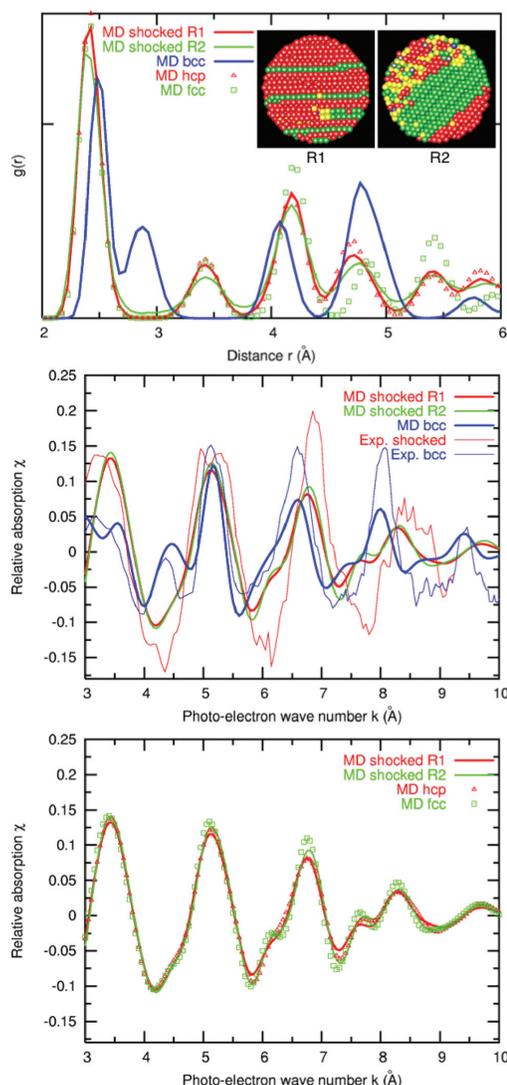


Fig. 3. Radial distribution functions and EXAFS signal of several different cutouts (each containing about 5000 atoms) from the sample shown in Fig. 2 (bottom): purely bcc, region R1 (hcp/fcc ratio 3.2) and R2 (hcp/fcc ratio 1.03). The radial distribution and the EXAFS signals of the regions R1 and R2 are compared to ideal hcp and fcc structures simulated under the same pressure and temperature conditions. A comparison to experimental in situ EXAFS data are shown in the middle panel.

# Molecular Dynamics Comes of Age: 320-Billion-Atom Simulation on BlueGene/L

Kai Kadau, Peter S. Lomdahl, T-14; Timothy C. Germann, X-1-SMMP

As the computational power is increasing, molecular dynamics (MD) simulations are becoming more important in physics, and science in general. We demonstrate weak and strong scaling of our MD code SPaSM on Lawrence Livermore National Laboratory's (LLNL's) BlueGene/L (BGL) architecture containing 131,072 IBM PowerPC440 processors. A maximum of 320 billion atoms have been simulated in double precision, which corresponds to a cubic piece of solid copper with an edge length of 1.56  $\mu\text{m}$ .

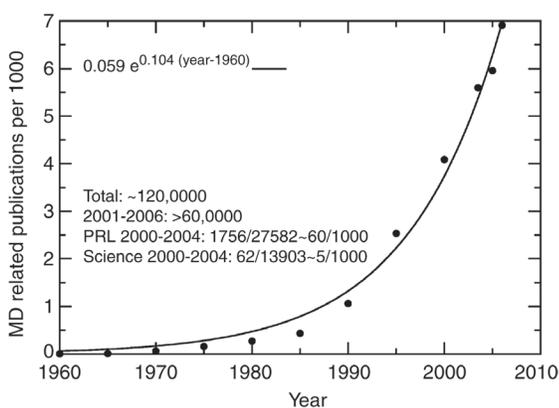
Since the invention of MD in 1957 by Berni Alder and Tom Wainwright [1], the computational power has increased from about 1000 floating point operations per second (Flops) of early vacuum tube machines, like the UNIVAC at LLNL or the MANIAC at Los Alamos National Laboratory, to 360 TFlops on architectures like the BlueGene/L at LLNL [2]. With this increase of computational power the early 100-particle hard sphere system has grown into more sophisticated particle interactions—like smooth pair and many

body potentials, as well as *ab initio* methods calculating the interatomic forces based on quantum mechanic principles—and the number of simulated particles has been increased to multiple billions. The importance of MD in science has evolved with this increase in capabilities as can be illustrated by the MD-related scientific literature fraction evolution (see Fig. 1). Today, 7 out of 1000 published scientific articles are related to MD; looking at a specific physics journal such as *Physical Review Letters*, this number is about 10 times larger, with the growth still in the exponential regime.

With the introduced SPaSM [3] algorithm, multiple short performance runs were carried out on the BGL system using 4096-131,072 CPUs and 1-320 billion particles arranged in a face-centered cubic lattice and interacting via a Lennard-Jones potential with  $r_{\text{cut}} = 2.5 \sigma$  and with  $r_{\text{cut}} = 5 \sigma$ . (See Fig. 2.)

The calculations were all performed in double precision, which means that each particle structure consisted of 88 bytes [position, velocity, and force vectors (24 bytes each), integer for type (4 bytes) and a tag (4 bytes), and a double to analyze/characterize the atoms (8 bytes)]. Hence, 320 billion atoms consume 26226 GB for the particle structure information alone. Additional memory is needed for the cell structure, various buffers, and the executable and operating system on each node. This additional memory only amounts to less than 20% of the whole memory consumption (LLNL's BGL has

**Fig. 1.** Evolution of the literature fraction of MD-related publications (i.e., molecular dynamics is a keyword or in the abstract). The literature search was performed using the following databases: BIOSIS, Engineering Index, Inspec, ISI Proceedings, ISI SciSearch, ISI Social SciSearch. An interdisciplinary journal such as *Science* follows this general trend, while specific physics journals such as *Physical Review Letters* have a significant larger fraction of MD-related publications. An exponential fit represents the general trend of the data.



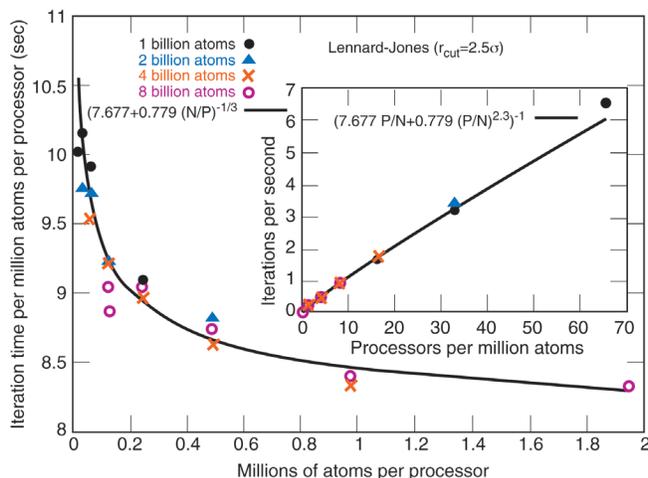
32768 GB main memory). These numbers suggest that SpaSM, with its efficient memory management, would be able to perform with almost a trillion atoms in single precision on ASC Purple, which has a main memory of 48832 GB.

Sophisticated use of the evergrowing computational power makes it possible to investigate challenges in physics and related science on the fundamental atomistic level. The 320 billion atoms simulated on BGL could represent a cubic piece of metal with an edge length well over a  $\mu\text{m}$  (Fig. 3). With one iteration taking only slightly more than 20 s, a physical simulation on the microscale in 3-D is now, in principle, possible in a matter of weeks. One of the biggest challenges to make these ultralarge-scale production runs feasible is the analysis and the input/output (I/O) of the enormous data produced on long-term storage, such as huge parallel file systems. The I/O is also very important for restart capabilities since time to failure of these enormous multiprocessor machines is usually only on the order of weeks if at all.

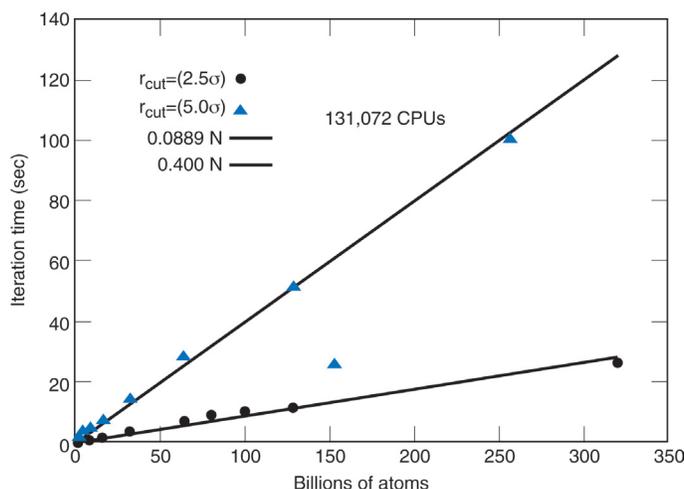
*For more information contact Kai Kadau at [kkadau@lanl.gov](mailto:kkadau@lanl.gov).*

[1] B. J. Alder and T. E. Wainwright, *J. Chem. Phys.* 27, 1208 (1957).  
 [2] T.C. Germann, et al., *Supercomputing '05*, SC05 IEEE ACM 1-59593-061-2/05/0011 (2005).  
 [3] D.M. Beazley and P.S. Lomdahl, *Parallel Computing* 20, 173–195 (1994).

**Funding Acknowledgments**  
 NNSA's Advanced Simulation and Computing (ASC).



**Fig. 2.** Strong scaling for 1–8 billion atoms distributed on 4096–65,536 processors of the IBM BGL machine at LLNL. For a small number of atoms per processor the communication/calculation ratio gets larger and the iteration time per atom per processor increases. The inset shows the speed-up as the number of processors per million atom increases.



**Fig. 3.** Iteration time for up to 320 billion atoms on 131,072 processors of the BGL architecture at LLNL. The achieved performance for 320 billion atoms simulated by a Lennard-Jones potential with a cutoff of  $2.5 \sigma$  was 27.2 TFlops, and 48.1 TFlops for 256 billion atoms interacting via a Lennard-Jones potential with a larger cutoff of  $5.0 \sigma$ .

# Shock-Front Broadening in Polycrystalline Materials

John L. Barber, Kai Kadau, T-14

Consider a polycrystalline solid comprising numerous grains, where each grain is composed of a different crystallographic orientation of some material. If a shock is induced at one side of the material and subsequently penetrates into the polycrystal, how does the initially sharp shock front grow in width with time? Meyers [1,2] has proposed a model in which the major contribution to the shock-front broadening comes from the differences in shock speed that distinct points on the shock front experience as they move through different grains. This model explicitly ignores other effects, such as scattering at grain boundaries, which may also influence the evolution of the shock-front width. The aim of this work is to explore the mathematical implications of the Meyers model, and to derive a useable analytical expression for the shock-front width.

Figure 1 shows a schematic representation of our polycrystal, along with an illustration of the coordinate system to be used. The shock front is assumed to originate on one face of the polycrystal, at  $y = 0$ , and penetrate into the bulk of the solid. Different points on the shock front propagate independently in straight lines in the  $+y$  direction, their speeds at any given location determined by the properties of the grain they are passing through. A particular sample ray arrives at a location  $y = x$  after a propagation time  $t$ . Since the geometry of the grains and the shock speeds therein are generally random, the relationship between  $x$  and  $t$  will be a random one. If the shock-front speed as a function of position for a sample ray is  $v(y)$ , then the propagation time to reach  $x$  is

$$t = \int_0^x dy \frac{1}{v(y)} = \frac{x}{c} + \int_0^x dy \left( \frac{1}{v(y)} - \frac{1}{c} \right), \quad (1)$$

where

$$c = \left\langle \frac{1}{v} \right\rangle^{-1} \quad (2)$$

is the reciprocal of the average inverse shock speed of the material, averaged over both the random grain geometry and the distribution of shock speeds within each grain. This is the overall rate at which the shock front will move through the crystal.

According to (Eq. 1), the propagation time can be separated into the sum of a deterministic part and a random part with mean zero. Using a version of the central limit theorem, it can be shown that in the limit of large  $x$ , the integral in (Eq. 1) reduces to a Gaussian random variable (RV), such that

$$t = \frac{x}{c} + \left( 2x \int_0^\infty dy \chi(y) \right)^{\frac{1}{2}} Z \quad (3)$$

Here  $\chi(y)$  is the spatial autocorrelation function of the integrand in (Eq. 1) and  $Z$  is a standard normal deviate.

Equation (3) gives the propagation time as a RV in terms of a fixed penetration distance  $x$ . By solving (Eq. 3) for  $x$ , we may obtain the random penetration distance as a function of fixed time  $t$ . The resulting relation can be manipulated to yield a relationship between the shock-front position  $x$  (i.e., the average value of  $x$ ) and the shock-front width (i.e., twice the standard deviation of  $x$ ):

$$\frac{\Delta x}{d} = \left( \frac{8c^2}{d} \int_0^\infty dy \chi(y) \right)^{\frac{1}{2}} \left( \frac{\bar{x}}{d} \right)^{\frac{1}{2}} \quad (4)$$

Here we have introduced  $d$ , a measure of the average grain diameter, in order to render (Eq. 4) scale invariant.

As it stands, (Eq. 4) is completely general, though unwieldy, since it requires complete knowledge of the autocorrelation function  $\chi(y)$ . We can obtain a more useable relation by making the additional assumption that the shock speed within each grain is

uncorrelated with the geometry of the grains themselves. In this case, it can be shown that (Eq. 4) reduces to

$$\frac{\Delta x}{d} = \left[ 4c^2 \chi(0) \frac{d^*}{d} \left( 1 + \frac{\sigma^2}{d^{*2}} \right) \right]^{\frac{1}{2}} \left( \frac{\bar{x}}{d} \right)^{\frac{1}{2}}, \quad (5)$$

where  $d^*$  and  $\sigma$  are known in stochastic geometry as the average chord length and standard deviation of the chord length, respectively [3]. Note that for the common special case of a Poisson Voronoi lattice in three dimensions, where  $d$  is taken as the mean caliper length [4], we have

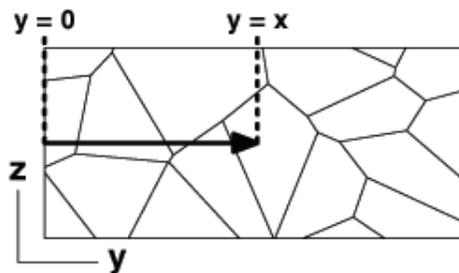
$$\frac{d^*}{d} \left( 1 + \frac{\sigma^2}{d^{*2}} \right) = 0.631. \quad (6)$$

The expression (Eq. 5), along with a choice for the distribution of shock speeds within the polycrystal, yields a testable prediction for the dependence of the shock-front width on shock position. By simulating the model directly, we can test the steps leading up to (Eq. 5). The results of several such simulations for various system widths  $L$  are shown in Fig. 2. The shock speed distribution and grain geometry were chosen to be consistent with a sample of polycrystalline iron at equilibrium in a body-centered cubic state, with a 3-D Poisson Voronoi geometry. We see that to within noise the simulation results are consistent with the theoretical prediction (Eq. 5). More realistic polycrystalline simulations using the embedded atom method [5,6] are currently being performed, in order to verify how well the assumptions and predictions of this model compare with a real physical system.

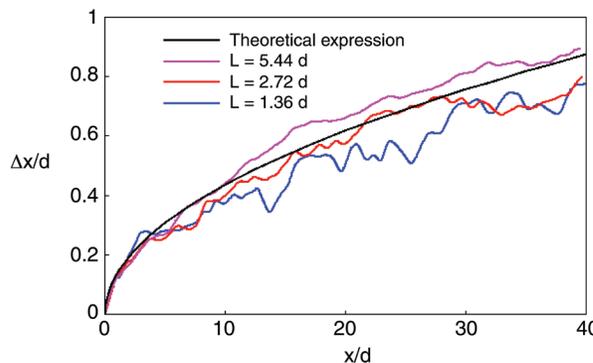
**For more information contact John L. Barber at [jlbarber@lanl.gov](mailto:jlbarber@lanl.gov).**

- [1] M.A. Meyers and M.S. Carvalho, *Mater. Sci. Eng.* **24**, 131 (1976).
- [2] M.A. Meyers, *Mater. Sci. Eng.* **30**, 99 (1977).
- [3] D. Stoyan, et al., *Stochastic Geometry and its Applications* (Wiley, 1995).
- [4] A. Okabe, et al., *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams* (Wiley, 1992).
- [5] K. Kadau, et al., *Science* **296**, 1681 (2002).
- [6] K. Kadau, et al., *Phys. Rev. B* **72**, 064120 (2005).

**Funding Acknowledgments**  
Laboratory Directed Research and Development Program.



**Fig. 1.**  
*A schematic representation of the coordinate system in our polycrystal. The arrow indicates the direction of shock propagation.*



**Fig. 2.**  
*Results for the evolution of the shock-front width, as determined by direct simulation of the ray model. Shown are curves for various values of the system width  $L$ .*

## New Shock-Induced Phase Transformation in RDX

Marc J. Cawkwell, Thomas D. Sewell, T-14

The response of crystalline solids to shock loading is complex and often involves extensive plastic deformation [1] and/or military phase transformations [2] in order to reduce the overall elastic strain energy in the material. A previously unknown shock-induced phase transformation has been identified in the secondary explosive RDX (cyclotrimethylene trinitramine) using large-scale molecular dynamics (MD) simulations. RDX is a molecular crystal that adopts an orthorhombic unit cell ( $\alpha$ -polymorph, spacegroup  $Pbca$ ) in ambient conditions [3].

The phase transformation was first observed in MD simulations of the collapse of cylindrical voids during shock loading normal to the (210) plane (see Fig. 1). The ordering of the molecular centers of mass was found noticeably different from that of the parent  $\alpha$  phase in volumes around the collapsing void. However, a coherent interface is found between parent and product phases. A sample of the new phase was “cut out” of the larger void collapse simulation block for further study and comparison with the  $\alpha$  phase.

Molecular dynamics simulations of the new phase showed that it also belongs to spacegroup  $Pbca$  but with relative translations of the molecules with respect to the parent structure (Figs. 2a and 2b). The molecular translations that give rise to the transformation result in a decrease in the  $b$  lattice parameter and an increase in the  $c$  lattice parameter by around 3.4% while the  $a$  lattice parameter remains almost unchanged. The transformation to the new phase results in an increase in volume by around 0.4%. Calculations of

the enthalpy difference between the two phases under hydrostatic compression shows that  $\alpha$ -RDX is more stable than the new phase by about 40 meV per molecule in the sampled pressure range of 1.5 GPa to 5.0 GPa. The new phase is unstable rather than metastable at hydrostatic pressures below 1.5 GPa.

Small volumes of the new phase were able to form in MD simulations of void collapse when the shock direction was normal to (210) because of compressive stresses exerted along [010] as the void collapses. The free surfaces at the void allow the expansion of the lattice parallel to [001] to be accommodated via extrusion of the material into the void. On the basis of these results, we were able to predict that this military phase transformation will take place in RDX single crystals shocked sufficiently strongly along [010]. The Hugoniot for shock loading along [010] was calculated using MD simulations (Fig. 3). Changes in gradient of the  $U_s - U_p$  Hugoniot correlate very closely with the onset of the phase transformation. Experimental studies are currently underway to evaluate these theoretical predictions.

*For more information contact Marc Cawkwell at [cawkwell@lanl.gov](mailto:cawkwell@lanl.gov).*

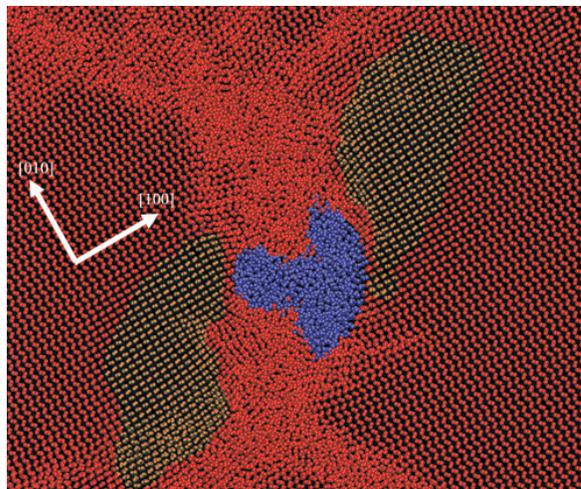
[1] B.L. Holian and P.S. Lomdahl, *Science* **280**, 2085 (1998).

[2] K. Kadau, et al., *Science* **296**, 1681 (2002).

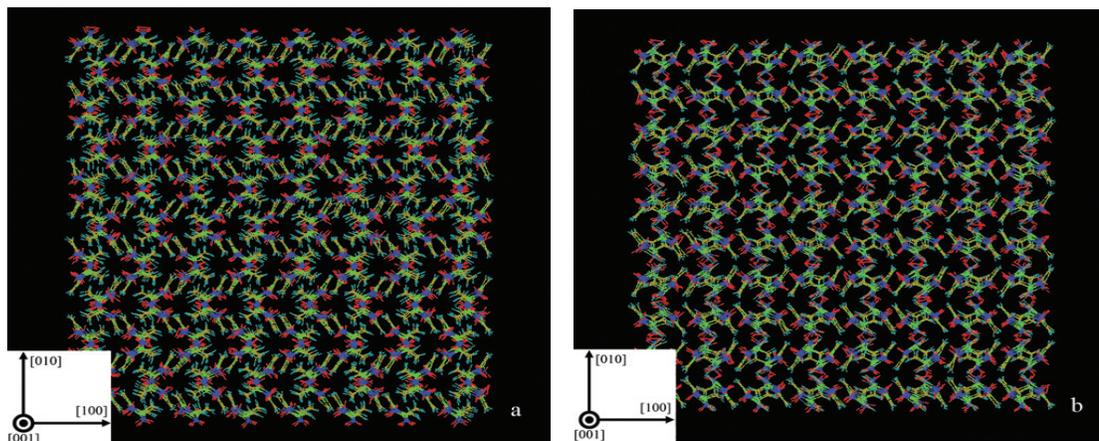
[3] C.S. Choi, and E. Prince, *Acta Cryst.* **B28**, 2857 (1972).

### Funding Acknowledgement

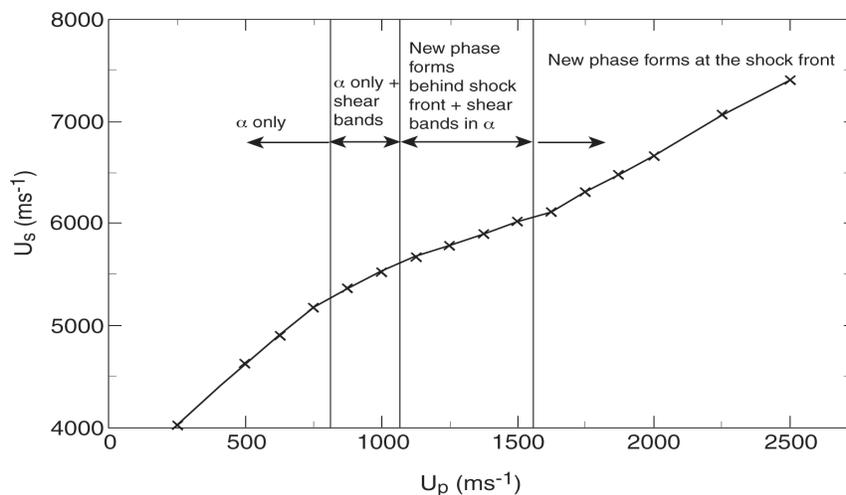
NNSA's Advanced Simulation and Computing (ASC), Materials and Physics Program.



**Fig. 1.** Snapshot of the molecular centers of mass during the collapse of a 200-Å diameter cylindrical void after shock loading normal to (210) at a particle velocity of  $630 \text{ ms}^{-1}$ . Color scheme: blue correspond to molecules that were on the surface of the void; yellow to molecules comprising the new shock-induced phase, and red to the  $\alpha$ -RDX structure and defects induced in the bulk during the simulation.



**Fig. 2.** Left—A projection along [001] of the molecular structure of  $4 \times 4$  unit cells of  $\alpha$ -RDX at a hydrostatic pressure of 1.5 GPa. Right—A projection along [001] of the molecular structure of  $4 \times 4$  unit cells of the predicted new phase of RDX at a hydrostatic pressure of 1.5 GPa.



**Fig. 3.** Hugoniot for shock loading parallel to [010] in single-crystal RDX calculated using MD simulations.

## The Core Structure of the [001] Screw Dislocation in $\alpha$ -RDX

Marc J. Cawkwell, Thomas D. Sewell, T-14

**P**lastic deformation in crystalline solids is mediated by the formation and glide of crystal defects known as dislocations [1]. Dislocations are line defects that possess a long-range strain field that is described by linear elasticity and a core region around their elastic center where the displacements of atoms or molecules from their lattice sites are sufficiently large that inelastic relaxations take place. The atomic scale structure of the dislocation core, i.e., how the Burgers vector of the dislocation is distributed spatially into the crystallographic planes in the zone of the dislocation line, has a significant influence on the mobility of the dislocation and the planes on which it can glide [2].

The secondary explosive RDX (cyclotrimethylene trinitramine) is a molecular crystal that adopts an orthorhombic crystal structure in ambient conditions ( $\alpha$ -polymorph) [3]. Despite its widespread use, a rigorous understanding of the mechanisms through which it deforms plastically has yet to be developed. Indentation studies have demonstrated that the primary slip system involves dislocations with Burgers vector [001] that glide on the (010) plane [4].

Atomistic studies of the core structure of the [001] screw dislocation in RDX have been performed in order to identify the slip planes in this material. Interatomic and intermolecular bonding in RDX is described using the non-reactive potential developed by Smith and Bharadwaj [5]. The simulation cell consisted of a cylinder three-Burgers-vectors thick and of diameter 460 Å. The

axis of the cylinder was parallel to [001] and periodic boundary conditions were applied in this direction. A perfect screw dislocation ( $\mathbf{b} = [001]$ ) was introduced into the center of the block by applying to the centers of mass of all molecules the displacement field calculated using the anisotropic elastic theory of dislocations [1,6]. The elastic constants used in the anisotropic elasticity calculations were determined at 300K from the Smith potential. The equilibrium structure of the dislocation at 300K was determined by performing molecular dynamics simulations while keeping all atoms in the outer 30 Å of the block fixed at their initial positions.

The core structure of the [001] screw dislocation is represented by a differential displacement map in Fig. 1. In this figure, relative displacements of the molecular centers of mass parallel to the dislocation line, i.e., the screw component of the Burgers vector, are plotted. It is evident that the displacements are confined to the (010) plane and that the core is planar. This result is in excellent agreement with experiment since such a configuration strongly implies that the glide of [001] dislocations on (010) planes will have the lowest attainable Peierls stress, i.e., the primary slip planes are (010). The edge components of the Burgers vector were found to be extremely small and there is no indication that the dislocation dissociates into partials. Atomistic calculations of the Peierls stress for the glide of the [001] screw dislocation in  $\alpha$ -RDX on (010), are currently underway in order to further resolve the experimental results in [4].

For more information contact Marc Cawkwell at [cawkwell@lanl.gov](mailto:cawkwell@lanl.gov).

- [1] J.P. Hirth and J. Lothe, *Theory of Dislocations* (Wiley-Interscience, New York, 1982).
- [2] V. Vitek, *Prog. Mat. Sci.* **36**, 1 (1992).
- [3] C.S. Choi and E. Prince, *Acta Cryst.* **B28**, 2857 (1972).
- [4] P.J. Halfpenny, et al., *J. Mater. Sci.* **19**, 1629 (1984).
- [5] G.D. Smith and R.K. Bharadwaj, *J. Phys. Chem. B* **109** 3570 (1999).
- [6] M.H. Yoo and B.T.M. Loh, "Numerical Calculation of Elastic Properties for Straight Dislocations in Anisotropic Crystals," Oak Ridge National Laboratory report ORNL-TM-3408 (1971).

### Funding Acknowledgements

NNSA's Advanced Simulation and Computing (ASC), Materials and Physics Program.

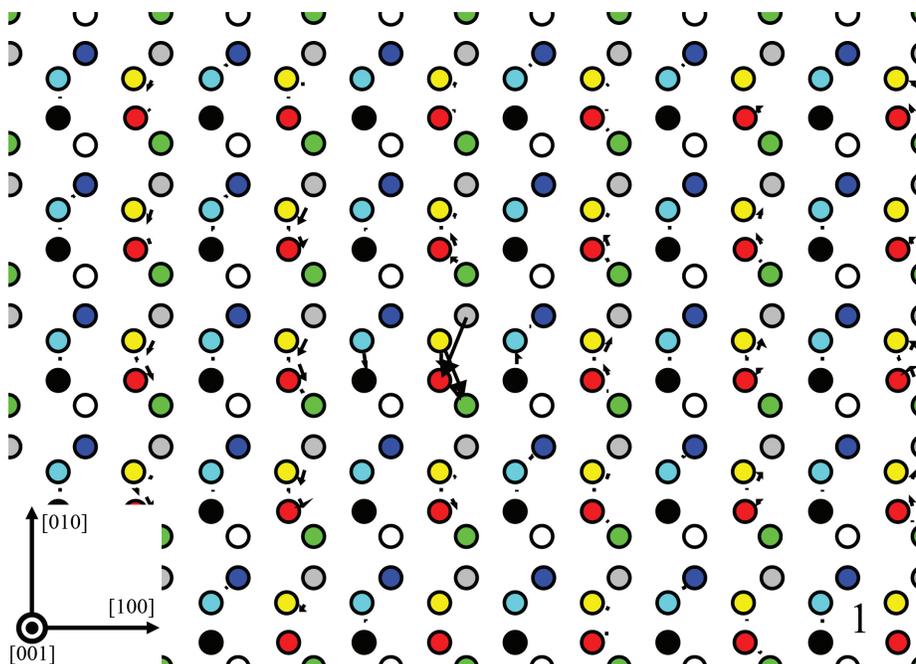


Fig. 1. A differential displacement map of the core of the [001] screw dislocation in  $\alpha$ -RDX. This map shows a projection perpendicular to the direction of the dislocation line, [001], of the centers of mass of molecules in one period of crystal. The colors differentiate molecular planes. The arrows represent relative displacements of centers of mass between neighboring molecules parallel to the dislocation line and their length is normalized such that it is equal to the separation of centers of mass when the relative displacement is equal to  $|b|/2$ .

## The Mechanics of LANL Foam Pads

Scott G. Bardenhagen, T-14; Brian M. Patterson, MST-7; Carl M. Cady, MST-8;  
Mathew W. Lewis, WT-1; Dana M. Dattelbaum, DE-9

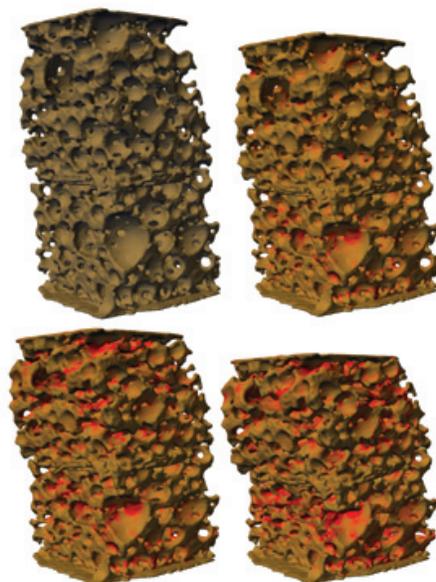
Cellular solids are ubiquitous in nature (e.g., wood, bone), and have been found increasingly valuable in meeting the material demands of an expanding suite of specialized engineering applications such as energy absorption and packaging (e.g., honeycombs, foams). These materials have unique characteristics relative to more common structural materials, including complex, irregular structure at the cellular scale, high strength to weight ratios, and routine operation over large deformation ranges. Recent research on polymeric foams is advancing modeling, simulation, and analysis capabilities applicable to Los Alamos National Laboratory (LANL) weapon system components in particular, and cellular solids more generally.

The character of foam material response is typically assessed via compression experiments on sufficient quantities of foam that the underlying foam cellular structure is largely unimportant.

These “bulk” measurements provide stress/strain curves that are useful in identifying common foam characteristics and categorizing response regimes. Examples may be seen in Fig. 2; S-shaped curves are typical of many types of foam. These data may also be used to calibrate engineering continuum models for bulk material response.

Bulk response data and models, while useful for many engineering tasks, are inherently inadequate for predicting the distributions and variability in material response evident in small foam specimens, or even cell-by-cell. Quantification of uncertainty in engineering systems requires as fundamental input material property variability. This is especially important in understanding the nature of damage and failure, which is typically controlled by the distribution of extreme events. The ultimate goals of this research on LANL foams are to characterize the distribution of states within a specimen, as well as the range of effective material property variability possible as the specimen size approaches the cell size.

Fig. 1.  
Foam structure as measured using microtomography (top left), and simulated compressed states, for a small column of foam.



Recent developments in particle-in-cell (PIC) methods indicate that these numerical techniques are suitable for precisely this class of problem [1]. Using foam structures determined using x-ray microtomography (similar to a CAT scan), quasi-static compression was simulated with results in agreement with experimental data in the literature. It was predicted that the full foam sample is an auxetic material at modest compressions, and that it becomes progressively more difficult to remove porosity, resulting in residual porosity even in “fully

densified" foam. This work was featured as an example of the contribution of numerical simulation and visualization in stockpile stewardship in the *Wall Street Journal* [2].

A similar study is now targeting LANL foam pads. This study is providing new estimates of foam "parent" material properties (the material properties of the individual foam struts). Because these are blown foams (manufactured using a curing process that evolves gas), it is impossible to manufacture the parent material in bulk for testing. One recourse is extrapolation from material testing results on ever more-dense foams, created by alterations in manufacturing that invariably affect parent material properties to an unknown extent. Another is extraction of structural elements (e.g., struts) for testing, which also has unresolved complications, in this case due to both the scale and geometry of the extracted parts. Numerical simulation provides an equally, if not more viable, alternative.

An in-house x-ray microtomography facility has recently been acquired and used to characterize LANL foam pads. These data and some of the first simulation results using these data may be seen in Fig. 1, where a column of foam in various states of compression is depicted.

Bulk stress strain curves for several foam pad samples are depicted in Fig. 2. The experimental data are colored by test lot, and cover a range of relative densities from 32–37%. Specimen variability within a test lot is fairly small. However, between lots it is larger, as illustrated by the highest relative density test lot stress strain curves (light blue) lying between lower relative density lots (purple, dark blue).

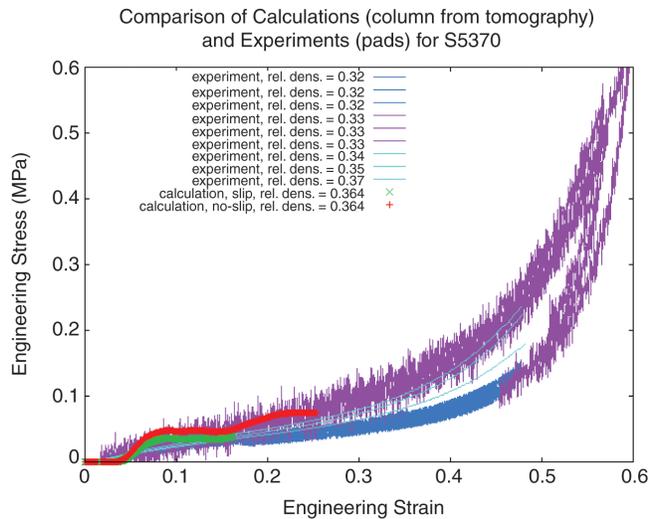


Fig. 2. Bulk stress-strain curves for LANL foam pads (blue, purple) and simulation results for the foam column (red, green).

Although the specimen in Fig. 1 is much smaller than the foam pad samples that generated the data in Fig. 2, its average response is expected to be similar. Simulation results for two different boundary conditions are depicted (red, green) and both are found to lie within the scatter of the experimental results. This fit required a substantial revision of the parent material properties from a previous estimate [3]. The shear modulus was reduced by a factor of three, and the bulk modulus increased by a factor of four. Future work includes determining the distributions of material states and effective material properties for input into uncertainty analyses.

*For more information contact Scott Bardenhagen at [bard@lanl.gov](mailto:bard@lanl.gov).*

- [1] A.D. Brydon, et al., *J. Mech. and Phys. Solids* 53, pp. 2638–60 (2005).
- [2] "Cold War Echo; U.S. Weighs Whether to Build Some New Nuclear Warheads," *The Wall Street Journal*, December 14, 2005.
- [3] M.W. Schraad and F.L. Harlow, *Intl. J. Mech. Sci.* 48, pp. 85–106 (2006).

### Funding Acknowledgements

NNSA's Advanced Simulation and Computing (ASC), Materials and Physics Program; and Campaign 2, Dynamic Materials Properties.

## Elastic Stress and Strain Distributions in Polymeric Materials

Panagiotis Maniadis, T-CNLS; Turab Lookman, Kim Ø. Rasmussen, T-11; Edward M. Kober, T-14

An understanding and modeling of the elastic properties of polymeric materials such as Estane®, which serves as the polymeric binder component in the high explosive PBX-9501, is an important issue for nuclear weapons. There are strong suggestions that aging effects in this composite arise from chemical and physical changes in the binder. Estane 5703 is a segmented block copolymer, the strength of which arises from the self-organization it undergoes.

Block copolymer melts in a microphase-separated state exhibit complex mechanical response and flow behavior that is fundamentally different from the behavior of chemically cross-linked materials such as rubber. Applied mechanical stress causes orientation phenomena, which in turn leads to the disruption and recombination of bonds into energetically favorable positions. These phenomena are thought to be responsible for the typical features of block copolymer materials such as high tensile strength, tear strength, and elongation. However, the interplay between self-assembly and elastic properties is sensitive to a wide range of variables, and a fundamental and comprehensive understanding has yet to be realized.

Our previously developed self-consistent field models can effectively and precisely determine the morphological structures emerging from the self-assembly of such segmented chains. In order to further enable this methodology to also characterize the elastic properties resulting from these self-assemblies, we have constructed a self-consistent field theory that in addition to the polymer density also incorporates the local stress and strain fields in the material. This development is based on previous work by Fredrickson [1] who recently developed a self-consistent field theory for homopolymer melts that simultaneously incorporates polymer density, stress, and strain fields. The strain is included at the microscopic level via a single chain expression due to Doi and Edwards [2]. The stress field arises as the conjugate to the strain field. We have extended this theory to self-assembling polymer systems, and we have initially computer-implemented this theory for the simplest such system, namely diblock copolymers.

Within this theory we are able to not only determine the morphological properties of the polymer density but also the corresponding internal stress distributions in the self-assembled structures. An example of this is

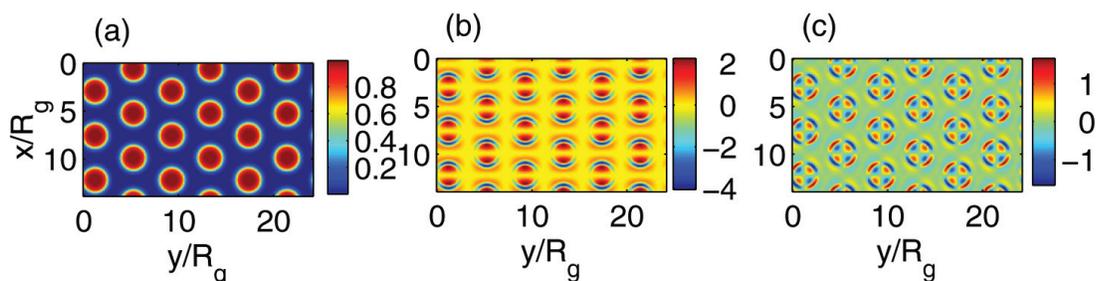
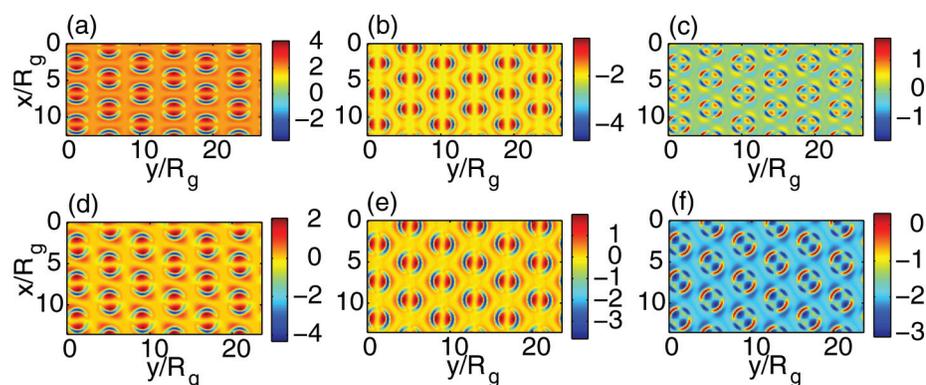


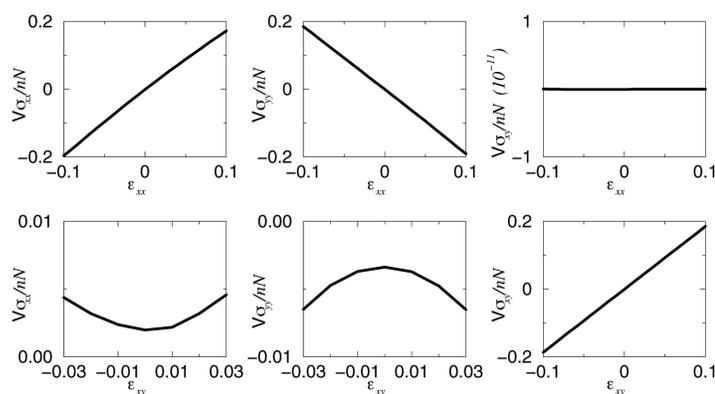
Fig. 1. Unstrained ( $\epsilon = 0$ ) equilibrium phases and stress tensor components for 30:70 AB diblock copolymer. The corresponding tensor components  $\sigma_{xx}$  and  $\sigma_{xy}$  are also shown.



**Fig. 2.** Components of the stress tensor for two different nonzero strain applied to the hexagonal phase ( $\text{tr } \epsilon = 0$  for both cases). For  $\epsilon_{xx} = -0.1$ ,  $\epsilon_{yy} = 0.1$ , and  $\epsilon_{ij} = 0$ , for  $i \neq j$   $\sigma_{xx}$  (a),  $\sigma_{yy}$  (b), and  $\sigma_{xy}$  (c) are shown. For  $\epsilon_{ij} = 0$ , for all  $i, j$  except  $\epsilon_{xy} = -0.1$  the components  $\sigma_{xx}$  (d),  $\sigma_{yy}$  (e), and  $\sigma_{xy}$  (f) are shown.

shown in Fig. 1 for the hexagonal phase of a diblock copolymer. In this particular case a block copolymer composed of just two different chemical species (say A and B) in the volume ratio 30:70 self-assembles into cylinders rich in the minority species (A), which are packed hexagonally in a matrix formed by the majority species (B). Together with a two-dimensional projection of polymer density we show a similar projection of the  $\sigma_{xx}$  and the  $\sigma_{xy}$  components of the stress tensor. Similarly, we can determine the remaining components of the stress tensor. Such stress distributions can also be determined in the presence of strain loadings (see Fig. 2), and closer scrutiny provides detailed understanding of the variation of stress along statistically typical polymer chains depending on their position in the overall structure of the self-assembled melt.

Our model further allows us to follow the behavior of the polymer melt in the presence of applied stress in greater detail as illustrated by Fig. 3 where we show the spatially averaged stress as a function of the external strain for the hexagonal phase. As expected it can be seen that a compression (expansion) of the melt, in the  $(i, j)$  directions, leads to



**Fig. 3.** The spatially averaged components of the stress tensor  $\sigma$  as a function of the external strain in the hexagonal phase.

a decrease (increase) of the spatially averaged stress in this direction (i.e. a change of  $\sigma_{ij}$ ). As the applied strain undergoes the relatively large variation from  $-0.1$  to  $0.1$  we see that the primary stress component changes linearly, while the more secondary elements display a nonlinear dependence for larger ( $|\epsilon_{ij}| > 0.01$ ) applied strain. From the slope of the lines shown in Fig. 3 in the small strain region, it is possible to calculate the corresponding components of the elasticity tensor  $K_{ij,kl}$ .

**For more information contact Kim Rasmussen at [kor@lanl.gov](mailto:kor@lanl.gov).**

- [1] G.H. Fredrickson, *J. Chem. Phys.* **117**, 6810 (2002).
- [2] M. Doi and S.F. Edwards, *The Theory of Polymer Dynamics* (Clarendon Press, Oxford, 1986).

#### Funding Acknowledgements

NNSA's Advanced Simulation and Computing (ASC), Physics and Engineering Model.

## Gravitational Effects in Reactive Flows: From Physics to Large Eddy Simulation Modeling

*Sergei Chumakov, T-CNLS; Natalia Vladimirova, University of Chicago*

**W**e use direct numerical simulations to investigate the effects of gravity in reacting flows. We consider two miscible fluids with small density difference; a reaction mechanism transforms one fluid into the other. We use a simple reaction model that is described by a single reaction progress variable, and the Boussinesq approximation for gravity. The combination of these two models allows us to focus on the interaction between gravity and reaction without interference of thermochemistry, compressibility, shock waves, density stratification, and other phenomena that make numerical simulations and interpretation of results much more complicated. The system is amenable to rigorous analysis [1] and yet is realistic enough for comparison with experiments, for instance, experiments with auto-catalytic reactions [2].

Here we describe the numerical setup that models a bubble of reaction products rising in reactant fluid. This setup has an astrophysical application. Several explosion scenarios of Type Ia supernovae involve a bubble of nuclear reaction products that rises towards the surface of the white dwarf star. Unlike regular hydrocarbon flames, the density change across the reaction front is very small and thus the Boussinesq approximation is justified. On the scale of the star the flame is very thin, and resolving the detailed reaction network at the flame scale is impractical; the simplified reaction model captures the most important effects.

The reaction at the surface of the bubble transforms “cold” material surrounding the bubble into “hot” reaction products. As the bubble grows in volume, its buoyancy increases. (See Fig.1.) When the bubble becomes large enough, the Rayleigh-Taylor instability (RTI) develops on the upper surface of the bubble. (See Fig. 2.) The instability stretches the reacting surface of the bubble and enhances the total reaction rate, further increasing the buoyancy.

The continuing efforts to follow the evolution of the bubble as a part of whole-star simulations [3–5] encounter the problem of tremendous separation of scales: millimeters (the flame front thickness) to thousands of kilometers (size of the star). All such simulations are initialized with relatively large bubbles and rely on assumptions about the state of such bubbles. Due to RTI, later bubble development is extremely sensitive to the subgrid (LES) model for reaction. Numerical simulations of small standalone bubbles are restricted to 2-D and axisymmetric bubbles [6,7]. Important 3-D effects, such as the development of the RTI on the top surface of the bubble, are not yet understood.

In our simulations, we study the effect of initial perturbations on the surface of a 3-D reacting bubble. The goal is to quantify the effect of the perturbations on the bubble growth rate and speed and how the induced flow influences the RTI on the surface of the bubble. The simulations are performed with the 3-D spectral code [8]; the 3-D

spectral elements code [9] and axisymmetric finite difference code [7] are used for cross-code comparison. All three codes are parallelized using MPI and perform well on distributed supercomputers.

The application of the DNS database is an *a priori* evaluation of the subgrid-scale models for Large Eddy Simulation (LES). The models in development include SGS scalar flux, SGS scalar dissipation, and SGS energy dissipation. An important extension of the existing set of models is the ability to predict the effect of SGS front stretching and curving on the gross reaction rate.

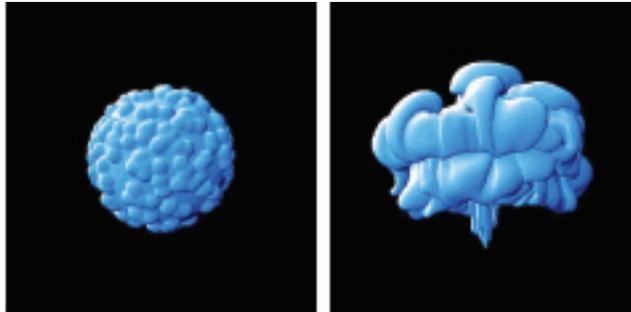


Fig. 1. Initially perturbed surface of the bubble (left) and the surface of the bubble at later time (right).

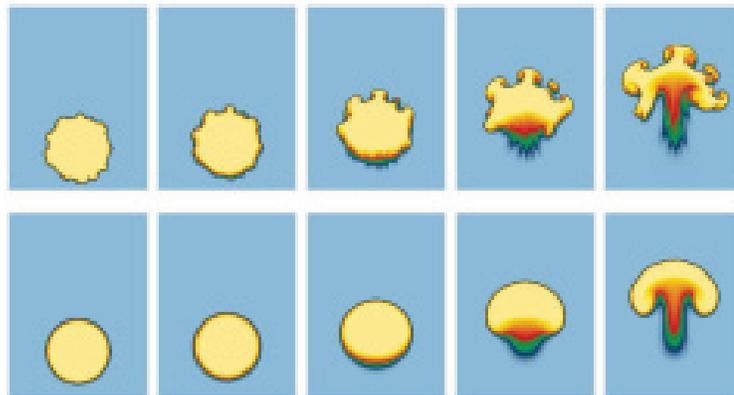


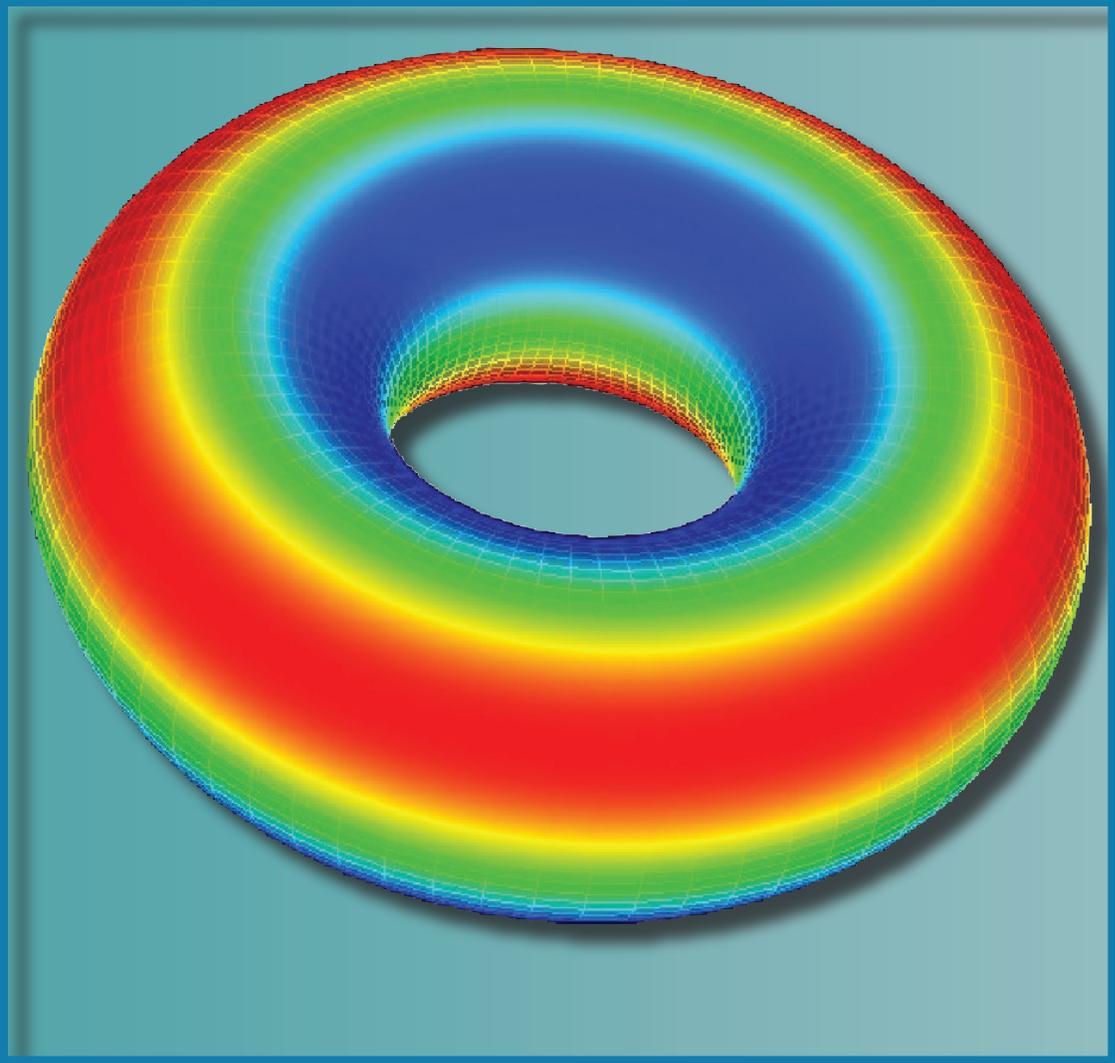
Fig. 2. Comparison of bubbles with initially perturbed and smooth surfaces; the images show the vertical slice of 3-D bubbles at equal time intervals.

For more information contact Sergei Chumakov at [chumakov@lanl.gov](mailto:chumakov@lanl.gov).

- [1] H. Berestycki, et al., *Analyse Non Lineaire* **23**, 407 (2006).
- [2] M.C. Rogers and S.W. Morris, *Phys. Rev. Lett.* **95**, 024505 (2005).
- [3] F.K. Roepke, et al., *arXiv:astro-ph/0609088* (2006).
- [4] T. Plewa, et al., *Astrophys. J.* **612**, L37 (2004).
- [5] A.M. Khokhlov, *Astrophys. J.* **449**, 695 (1995).
- [6] M. Zingale and L.J. Dursi, *arXiv:astro-ph/0610297* (2006).
- [7] N. Vladimirova, *Combustion Theory and Modelling* (2007), in press.
- [8] S.G. Chumakov, *J. Fluid Mech.* **562**, 405–414 (2006).
- [9] H.M. Tufo and P.F. Fischer, *J. Parallel and Distributed Comput.* **61**, 151 (2001).

### Funding Acknowledgements

NNSA's Campaign 4, Secondary Assessment Technology; and University of Chicago Center for Astrophysical Thermonuclear Flashes.



# Production Desktop, Supercomputing, and Support

The high-performance computing environment at Los Alamos National Laboratory (LANL) extends from the end-user's office computer to the supercomputer and storage systems to large-scale visualization facilities. These environments are complemented by software tools sets, application and input/output (I/O) libraries, and high-performance networks that are developed and supported by computer scientists, engineers, and technicians from the High-Performance Computing (HPC) and Computing, Telecommunications, and Networking (CTN) divisions. The intent is to provide an intuitive electronic workplace that enables simulation scientists to get the most out of the national assets in the three primary LANL computer centers: yellow (unclassified-protected), turquoise (unclassified-collaboration), and red (classified).

Adding to the complexity and challenge of these environments is the need for manipulating the huge data sets that are required to represent complex physical phenomena. This drives the need for advanced computer architectures, higher and higher speed networks, and a growing inventory of software support tools to mask the complexity of the computer environment from the end user.

The articles in this section focus on the "desktop to teraflops" supercomputing environment that is the primary modern simulation and modeling tool that scientists need and use to develop a first-principles understanding of the world and universe that we all share.

## Managing Software Tools in High-Performance Computing Environments

David R. Montoya, Georgia A. Pedicini, David O. Gunter, HPC-4

**P**ast high-performance computing (HPC) environments were based on a singular architecture with extensive support from the vendor. Examples at Los Alamos National Laboratory (LANL) include the Crays (1970s–1980s), Blue Mountain (1990s), and the Q clusters (2000). During the earlier days of supercomputers there had to be a close collaboration between the vendor and the implementation site personnel. Computers such as these could not be built with only local skills, and much of the software was custom-developed for the machine. Overall this was an expensive proposition. In those early days the Advanced Simulation and Computing (ASC) program provided an environment to build products that could support those new-generation machines. The ASC PathForward program, begun in 1998, helped tie universities, commercial vendors, and national laboratories into partnerships to develop the supporting software infrastructure in support those new machines. Capabilities and skill bases were created, which in turn fed new products into the HPC economy.

More recently, the open-source movement developed a process where software could be free and user maintained, allowing intellectual capital to grow based on the interest that could be generated in a given area or product. For user software, this movement fostered an explosion of products, driving both vendor-supported and open-source software products. Another development area is Linux, a low-cost operating system (OS). Ron Minnich (formerly of CCS Division) further developed this OS as BProc/Clustermatic, which emphasizes a lightweight kernel on compute nodes of a complex Linux cluster. This was the advent of what we call capacity HPC.

This phase of computing has moved into production capacity within the past 3–4 years, creating both opportunities

and challenges. Moore's Law, describing growth in hardware capability, continues to provide an indicator for new generations of hardware. Computational clusters are just reaching production quality when new hardware is already being procured for the next capacity cluster. Newer machines have different capabilities, may run newer versions of the Linux kernel, have different network fabric implementations, or have different parallel file systems. The result is a loss of consistency in the software and user environments. The challenge is to manage consistency to the best level possible. Tools to manage and plan for a diverse environment with different software requirements and configurations are essential.

### Environment

The details of each cluster that makes up the current heterogeneous HPC environment must be understood to grasp what impact its specific hardware has on the software environment. Table 1 describes the current inventory of HPC clusters at LANL. Footnotes to the table indicate differences among machines that are nominally of the same class.

From a software health perspective, attributes of interest are as follows:

- Number and type of segments describe the individual units that make up a portion of the larger machine. On some clusters, different segments provide different functionality and must be viewed as separate machines. Front-end and compile nodes, not reflected here, add another dimension of complexity. Taking these all into account, we monitor more than 50 machines.
- Processor and OS provide a base for the core software architecture of a machine. Subcharacteristics that have a large impact are kernel version, glibc version (the library that defines system calls), and interconnect. The variant of the OS

has a huge impact when using software products that are tightly coupled to the hardware. Many of our systems have a BProc implementation. This is a concern with software when it needs to tap into specific libraries with a known implementation (e.g., BProc/Myrinet and BProc/InfiniBand). Any tools that deal with process management and migration need to be rewritten to accommodate the BProc paradigm. Specifically, we had to modify MPI (message-passing interface) implementations and debuggers to run under BProc.

- Interconnect and storage systems come into play when we look at processor communication and I/O. Message Passing Interface (MPI) implementations, and software interacting with file systems, may need to be retuned for each implementation.

There are trade-offs when it comes to supporting software on current and emerging HPC platforms. As long as the growth curve is steep for architectures that provide increasing computational power, we need to invest in the infrastructure that must follow it to make it useable. This investment will have to grow as we move to the new accelerator and hybrid architectures for which there currently are no software tools.

*For more information contact David R. Montoya at [dmont@lanl.gov](mailto:dmont@lanl.gov).*

#### Funding Acknowledgements

This research was supported by the NNSA tri-Lab Advanced Simulation and Computing Program.

Secure Restricted Network (Red)									
Name (Program <sup>1</sup> )	Processor	OS	Segments	Nodes per Segment	CPUs: per Node / Total	Memory per Node / Total	Interconnect	Peak (GFlop/s)	Storage
QA / QB (ASC)	Alpha	Tru64	64 (Domains)	32	4 / 8192	8GB <sup>2</sup> / 22.4TB	Elan3 <sup>3</sup> -Dual	20,480	288 TB HP PFS
CA/CB/CC (ASC)	Alpha	Tru64	3x4 (Domains)	32	4 / 1536	4GB / 1.5TB	Elan3 <sup>3</sup> -Single	3,840	HP PFS
CX (ASC)	Alpha	Tru64	1 (Domain)	32	4 / 128	4GB / 128GB	Elan3 <sup>3</sup> -Single	320	HP CFS
Lightning (ASC)	AMD opteron	Linux BProc	13	255	2 <sup>4</sup> / 7,140	8GB <sup>5</sup> / 26.5TB	Myrinet/Lanai	30,600	160 TB Panasas
Unclassified Protected Network (Yellow)									
QSC (ASC/IC)	Alpha	Tru64	8 (Domains)	32	4 / 1024	16GB / 4TB	Elan3 <sup>3</sup> -Dual	2,560	26 TB HP PFS
Grendels (ASC)	Xeon	Linux BProc	1	126	2 / 252	2GB / 252GB	Myrinet	1.2	NFS
Flash (ASC)	AMD opteron	Linux BProc	5	255 <sup>6</sup>	2 / 1,906	8GB <sup>7</sup> / 8.58TB	Myrinet/Lanai	8,643	36 TB Panasas
Lambda (R)	PentiumIII	Linux	1	164	2 / 328	4GB / 656GB	Ethernet	918	NFS
Saguaro(R)	AMD opteron	Linux BProc	1	32	2 / 64	4GB / 128GB	Ethernet	307	NFS
Open Collaborative Network (Turquoise)									
Pink (IC)	Xeon	Linux BProc	1	958	2 / 1,916	2GB / 1.9TB	Myrinet Lanai	9,196	50 TB Panasas
Mauve (IC)	Itanium	Linux	1	64	4 / 256	16GB / 1TB	Numalink <sup>8</sup>	819	40 TB sgi
TLC (IC)	AMD opteron	Linux BProc	1	110	2 / 220	8GB / 880GB	Myrinet	880	50 TB Panasas
Coyote (IC)	AMD opteron	Linux BProc	5	258	2 / 2,550	8GB / 10.2TB	InfiniBand	13,485	50 TB Panasas

<sup>1</sup> Programs: IC=Institutional Computing, ASC=Advanced Simulation and Computing, R=Recharge  
<sup>2</sup> Exceptions: (qd26, qd27...qd31 and qd58, qd59...qd63) = 16GB per node; (qd22, qd23, qd54, qd55) = 32GB per node  
<sup>3</sup> Elan3 = Quadrics QsNet interconnect, either Dual rail or Single rail  
<sup>4</sup> Exception: lb-7 has "dual-core" CPUs, ie. 4 CPUs per node for a total of 1,020 processors.  
<sup>5</sup> Exceptions: (ll-1 and ll-2)=4GB per node, lb-1=16GB per node  
<sup>6</sup> Exceptions: flasha=300 nodes, flashd=127 nodes, flashdev=16 nodes  
<sup>7</sup> Exceptions: flashd=16GB per node, flashdev=4GB per node  
<sup>8</sup> Mauve is a Symmetric MultiProcessor with Direct Memory Access (DMA) shared among all its CPUs.

**Table 1.** Current High-Performance Compute Clusters managed by HPC Division (December 2006).

## Tripod Project as a Process

*Richard L. Graham, CCS-1; Josip Loncaric, HPC-5; Steven R. Shaw, HPC-3*

The National Nuclear Security Agency (NNSA) Advanced Simulation and Computing (ASC) program chartered the Tripod project, initially targeted at Linux capacity computing clusters, to develop a seamless software environment for use by the NNSA tri-lab community (Los Alamos, Livermore, and Sandia national laboratories). Los Alamos has made suggestions for viewing Tripod as an ongoing process, which were well received at the December 13, 2006, Tripod tri-Lab meeting at Livermore.

The rapid pace of change in the computer industry is legendary. To deal with this rapid change, Tripod needs to define an agile process. While the details have yet to be defined, the essential requirements for this process include the following:

- 1) Ongoing discussions and quarterly technology direction evaluation: Tripod should prepare a roadmap for high-performance computing (HPC) capacity systems based on market trends, re-evaluate this document quarterly, and update it as needed. This roadmap should propose a small number of strategies for top trends and encourage market competition within the chosen software interfaces. To evaluate these strategies under full production conditions, in time to guide decisions, “pathfinders” are needed. Once a path has been chosen, Tripod should educate the user community early about the necessary changes.
- 2) New capacity systems should follow an accelerated procurement

and deployment schedule to maximize their productivity. A rapid applications readiness process to stabilize new systems for production must not require a tri-Lab agreement on minor technical details. Each lab will have to comply with its local security requirements. Software choices should follow the Tripod roadmap in broad terms. The focus ought to be on standard interfaces, with equivalent functionality seen as acceptable. The Tripod roadmap should be reviewed and updated based on recent deployment experience so that the next deployment is made faster.

- 3) Old capacity systems in full production should be updated when justified, but those updates must be carefully managed to minimize disruption to users. It is essential that updates be governed by a cost/benefit analysis.

Los Alamos recommendations for Tripod policies are as follows:

- 1) New capacity systems should be consistent with the Tripod roadmap, with criteria divided into three tiers: (a) equivalent functionality based on public standards for software interfaces, (b) software robustness (consider dependencies and avoid single points of failure), and (c) specific Tripod software components used. For non-ASC users, other software stacks may be appropriate.
- 2) Upgrades of old capacity systems in stable production, which were Tripod compliant, should be governed by a

- cost/benefit analysis when Tripod software recommendations change.
- 3) Some capacity systems in full production should be used as “pathfinders” to evaluate benefits of alternative software stacks. Rapid evolution of HPC technology is expected to create radically different architectures within 5 to 10 years, including radical changes in Tripod software stack, yet we cannot get there without effort. Staying at the forefront of computing is critical to national security.

### Conclusion

Tripod should define a process that tracks the forefront of computing with agility, speed, and precision, and yet delivers reasonable stability to users. Our goals are user productivity at justifiable cost/benefit ratios and management of lifecycle costs. Therefore, Tripod should foster R&D in areas expected to improve productivity. We can use bidirectional leverage to reduce procurement and support costs: leveraging the mass market and contributing software as open source to a larger community of users to leverage their support. Tripod must also consider software component dependencies and reduce the impact of bugs.

Tripod policies should respect local security requirements, deliver a consistent tri-Lab user experience, yet allow for different types of capacity systems in production (“pathfinders”). Upgrades of old systems should be guided by cost/benefit-based decisions. Finally, robustness of the Tripod software stack requires avoidance of single points of failure wherever possible.

*For more information contact Josip Loncaric at [josip@lanl.gov](mailto:josip@lanl.gov).*

### Funding Acknowledgements

This research was supported by the NNSA tri-Lab Advanced Simulation and Computing Program.



[https://tripod.ca.sandia.gov/mediawiki/index.php/Main\\_Page](https://tripod.ca.sandia.gov/mediawiki/index.php/Main_Page)

## Quantifying Checkpoint Efficiency

John T. Daly, HPC-4

**A**dvanced Simulation and Computing (ASC) applications implement fault tolerance using a simple checkpoint rollback strategy that periodically saves a known good state of a calculation by writing a dump file. One would like to be able to quantify the effectiveness of this form of fault tolerance, particularly in light of decreasing application mean time to interrupt (MTTI) associated with high-performance computing (HPC) systems of increasing complexity and numbers of parts.

If *solve time* is defined as time spent performing productive computation, and input/output (I/O), and run time as solve time plus time spent in restart, rework, and writing dump files, then it follows that the *checkpoint efficiency* can be accurately quantified in terms of three parameters: *restart overhead* ( $R$ ), *dump time* ( $\delta$ ), *application MTTI* ( $M$ ), and *checkpoint interval* ( $t_c$ ), as follows [1]:

$$\text{Checkpoint Efficiency} = \frac{T_s}{T_r} = e^{-\frac{R}{M}} \left( \frac{\lambda - \frac{1}{2}\Delta^2}{e^\lambda - 1} \right)$$

$$\text{where } \lambda = \frac{t_c}{M}, \Delta = \sqrt{\frac{2\delta}{M}}.$$

By assuming that interrupts arrive randomly and are exponentially distributed [2], it is known that a solution for the optimum checkpoint interval exists in terms of dump time and application MTTI. It can be approximated to highest order when  $\Delta$  is small [3], or it can be expressed analytically in terms of the Lambert W-function and approximated quite accurately in terms of a perturbation series [4]. The checkpoint efficiency

as quantified above can be used to demonstrate that checkpoint schemes are not particularly sensitive to one's estimate of the application MTTI, as illustrated by Fig. 1. This property is important because in practice the application MTTI may be quite difficult to estimate accurately.

Since computing checkpoint efficiency in terms of an exact solution for the optimum checkpoint interval can be tedious, a simple approximation of  $\lambda$  that is accurate for both large and small  $\Delta$  would be desirable. Such an approximation may be constructed using the notion of asymptotic matching as follows [1]:

$$\begin{aligned} \lambda &\approx \Delta + \frac{1}{2}\Delta^2 \Rightarrow \text{Checkpoint Efficiency} \\ &\approx e^{-\frac{R}{M}} \left( \frac{\Delta}{e^{\Delta + \frac{1}{2}\Delta^2} - 1} \right). \end{aligned}$$

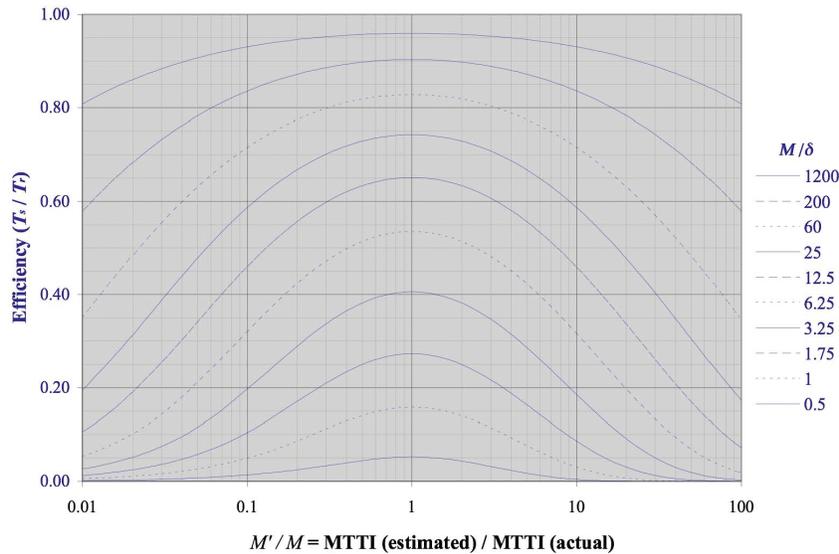
Figure 2 demonstrates the accuracy of this approximation. It also clearly illustrates that once a 90–95% efficiency has been attained, further improvements in the ratio of reliability to dump time are of limited value to a checkpointing application. This is an important design consideration for future HPC systems.

[1] J.T. Daly, NECDC Proceedings (2007).  
 [2] J. T. Daly, LA-UR-06-8089 (2006).  
 [3] J. W. Young, *Comm. of the ACM* **17**, 530 (1974).  
 [4] J. T. Daly, "A Higher Order Estimate of the Optimum Checkpoint Interval for Restart Dumps," *Future Generation Computer Systems* **22**, 300 (2006).

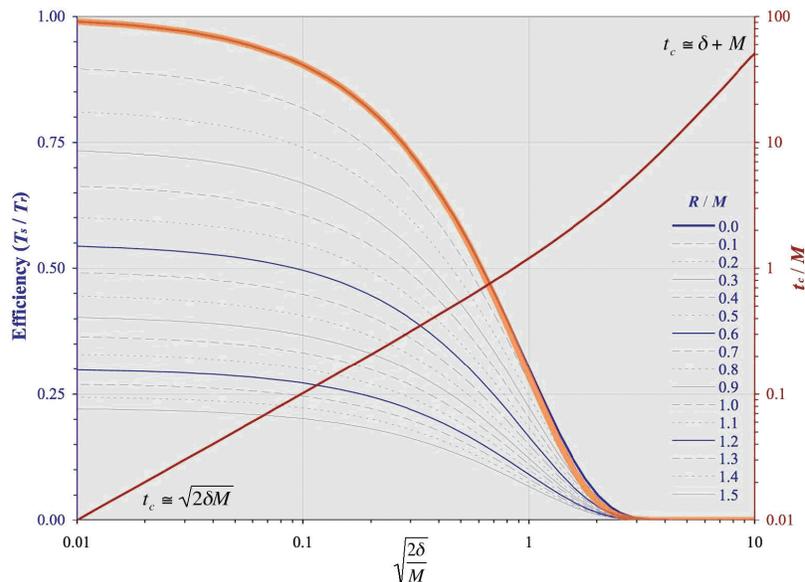
For more information contact John Daly at [jtd@lanl.gov](mailto:jtd@lanl.gov).

### Funding Acknowledgements

This research was supported by the NNSA tri-Lab Advanced Simulation and Computing Program.



**Fig. 1.** Checkpoint efficiency as a function of the accuracy of the estimated application MTTI based on the exact solution of the optimum checkpoint interval and assuming that restart overhead is negligible.



**Fig. 2.** Checkpoint efficiency is plotted in blue. It is based on an exact solution of the optimum checkpoint interval shown in red. The efficiency based on asymptotic matching is plotted in orange for  $R/M=0$ .

## Unified Data Model: A Library for Parallel I/O and Data Management

William W. Dai, Robbie Aulwes, HPC-4

**T**he Unified Data Model (UDM) is an original library for parallel input/output (I/O) and data management that we developed under the Advanced Simulation and Computing (ASC) program for code projects. The goal of the library is to provide sustainable, interoperable, efficient, scalable, and convenient tools for parallel I/O and data management for high-level data structures in applications, and to provide tools for the connection between applications, i.e., files generated from one application to be used in another application as inputs.

### Functionality for Array and Structured Mesh

The UDM library provides I/O tools for single and multidimensional arrays used in numerical simulations. It supports arrays with any number of dimensions. The minimum requirement to write an array in parallel is to set the size in each dimension on its own computer processor, and the starting address of data in memory. As a result, each processor contributes a part of a multidimensional array. The library also supports ghost cells, which are typically involved in simulations in parallel computer environments.

The capability for structured meshes in the library is very similar to the one for multidimensional arrays, but association between a structured mesh and a set of variables will be automatically built through the library. Variables may also be defined on different structured meshes, and they may be defined on nodes, faces, and mesh elements.

### Functionality for Unstructured Meshes and Variables

One of the important and powerful functions in the UDM library is the management of unstructured meshes and variables. The library supports a broad range of unstructured meshes, which include meshes with fixed shapes, arbitrary polygons, and arbitrary polyhedrons. A mesh element may be a zone, or face, or edge; i.e., zone-mesh, face-mesh, and edge-mesh. An edge-mesh may be 1-D, 2-D, or 3-D; and a face-mesh may be 2-D or 3-D. Mesh elements of a zone-mesh may be made directly from nodes, the elements may be made from edges, and the elements may also be made from faces, and the faces are made from either edges or nodes. The UDM library also supports ghost mesh elements, boundary faces, boundary edges, boundary nodes, slip faces, slip edges, slip nodes, and other options. The variables the library supports include node-, or edge-, or face-, or zone-variables, and the variables may be scalars, vectors, and tensors.

Although the UDM library covers a broad range of unstructured meshes, a user only has to set up one mesh definition and all other mesh definitions will be hidden from the user. For example, for an unstructured zone-mesh made from nodes, only a list of nodes for each element is needed, if the elements are of a fixed shape, such as prisms. If mesh elements are arbitrary polyhedrons made from nodes, two arrays are needed; one for the numbers of nodes for each element, and the other for the list of nodes for each element. Like the

capability for structured meshes, the association between a mesh and a set of variables is automatically built into the library.

### Functionality for Querying

A file written through the UDM library is self-descriptive. All the information in the file may be queried by calls to the library. For example, for a given file, users may find the number of arrays, meshes, and variables, the description of each array, mesh, and variable, and any association between meshes and variables. Through querying function calls in the library, meshes and variables may be viewed through graphics tools.

After a data object, such as array, or mesh, or variable, is written into a file, users may read any part of the data object, for example, as global IDs, or processor rank, or space domain. Figure 1 illustrates the capability for reading three parts of an unstructured mesh with 1.6 billion elements. The left image is a part read through a processor rank, the middle one is a part identified through global IDs, and the right image is read based on a space domain.

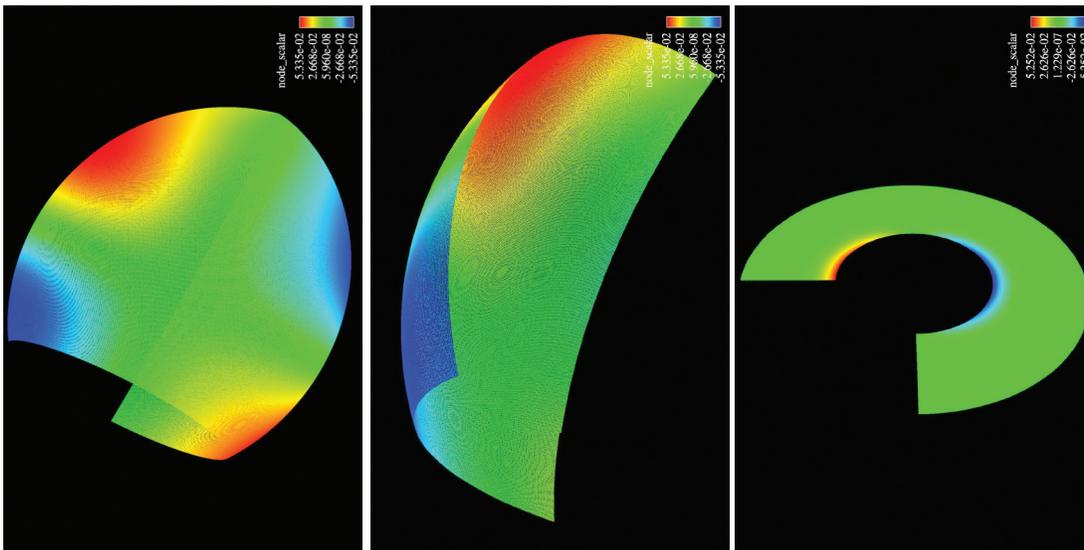
### I/O Performance

The UDM library is built on the top of message-passing interface in/out (MPI-IO). The performance of the library is approximately 95% that of the best MPI-IO performance. To get that performance, we first make sure that the data on each computer processor are contiguously written onto a disk file, and therefore there is no movement of data during writing. Second, the library collects all metadata, i.e., the description of data, and writes the collection only when a file is closed. Third, when reading a file, the library reads all the metadata together and reads it only once.

*For more information contact William W. Dai at [dai@lanl.gov](mailto:dai@lanl.gov).*

### Funding Acknowledgements

This research was supported by the NNSA tri-Lab Advanced Simulation and Computing Program.



**Fig. 1.** The three parts of an unstructured mesh and the variables defined on the three parts. The original full mesh contains 1.6 billion elements. Each part of the mesh is read from the original mesh through three different capabilities. The left image is read through a computer processor rank, the middle one is obtained through a set of global IDs, and the right image is obtained from a user-defined space domain.

## Applications Using the Unified Data Model

Rob Aulwes, William Dai, HPC-4

**T**he Unified Data Model (UDM) library was developed under the Advanced Simulation and Computing (ASC) program at Los Alamos National Laboratory to provide a high-performance parallel in/out (I/O) library for managing mesh connectivity and mesh data for simulation projects. The library includes a simple C application programming interface (API) that developers use to easily construct structured and unstructured meshes and their associated variables. The Unified Data Model supports multidimensional arrays and structured and unstructured meshes. Unstructured meshes can be constructed using fixed shapes or arbitrary polygons and polyhedrons. All mesh data are stored in a single file for easy file management, and the data file is platform-independent. The UDM library has been successfully integrated with the EnSight visualization tool [1] and the Shavano project [2].

### Integration with EnSight

Visualization software is an essential tool for analyzing meshes and mesh data. EnSight, a parallel visualization software developed by CEI, is the software commonly used at the Laboratory. It becomes a requirement with our customers that the mesh data written to files must be viewable in EnSight.

Part of the UDM distribution includes a shared library module that provides a plug-in for EnSight visualization. With the module, users are able to use EnSight to view a mesh and its variables that are stored in a UDM file. The plug-in supports 2-D and 3-D meshes.

EnSight provides a feature for handling distributed data, called server of servers (SoS). However, with earlier versions of EnSight, users were required to decompose an unstructured mesh manually or attempt to decompose the mesh using an unsupported partitioner. Consequently, they would have to decompose each time they used the server-of-servers (SoS) mode with a different number of servers. Since UDM stores mesh data in a single file, there is no partitioning required to view UDM data using the SoS mode.

### Integration with Shavano

The Shavano project is one of a number of ASC projects whose goal is to develop high-fidelity weapons simulations. The UDM library has been successfully integrated within the Shavano project to provide the ability to perform a parallel dump to a single file for restarting calculations. We added the ability to dump data to a UDM file for visualizing with EnSight. We are continuing our efforts to enhance data analysis using EnSight by including additional variable data.

We also demonstrated the ability to connect a legacy application to Shavano. We wrote output data from the legacy application to a UDM file, which was then imported into the Shavano project. Using the UDM format permitted the user to run calculations using the imported data with any number of CPUs. With the current process of running calculations, a user would need to generate a new set of files to run the same calculation with a different number

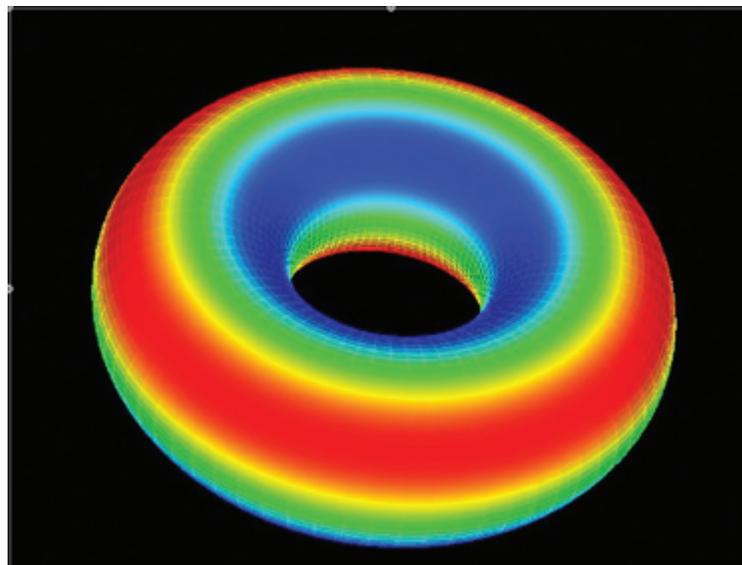
of CPUs. We are using UDM to eliminate this step by adding the capability to run with any number of CPUs using the same set of files. Figures 1 and 2 are 3-D simulations generated using UDM.

*For more information contact Rob Aulwes at [rta@lanl.gov](mailto:rta@lanl.gov).*

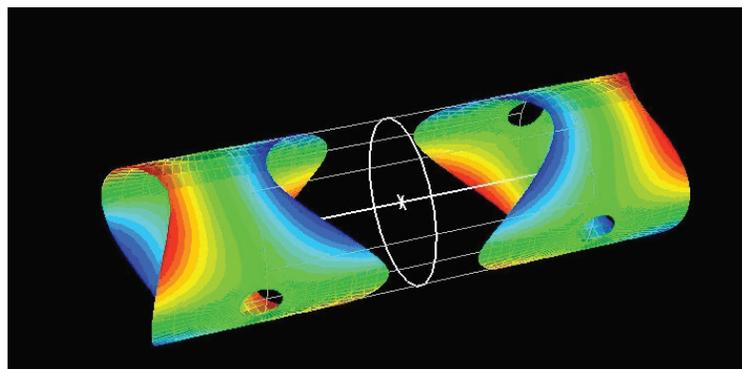
[1] EnSight Web site, <http://www.ensight.com>.  
[2] K. Roark, "Classified colloquium to explore the future of weapon simulation," LANL News Bulletin, [http://int.lanl.gov/news/index.php/fuseaction/nb.story/story\\_id/7520/nb\\_date/2005-11-16](http://int.lanl.gov/news/index.php/fuseaction/nb.story/story_id/7520/nb_date/2005-11-16).

### Funding Acknowledgements

This research was supported by the NNSA tri-Lab Advanced Simulation and Computing Program.



**Fig. 1.**  
*Torus generated using UDM.*



**Fig. 2.**  
*Intersection of torus with a cylinder.*

## Creating Future Network Architectures

Robert Martinez, Tim Merrigan, Denny Rice, Todd Bowman, CTN-5

Supercomputers were once built with tightly coupled CPUs (processors), data storage, and visualization hardware that communicated using vendor-specific interconnect hardware. Today's high-performance computing (HPC) platforms are architected with thousands of processors organized into clusters that communicate across multiple high-performance networks. The current HPC hardware paradigm is built on large numbers of commodity servers containing multiple processors operating in parallel, each processor operating on a small portion of the overall calculation.

Communications between hundreds or thousands of processors to achieve this parallelism is accomplished on multiple specialized networks optimized for their functions. Communications between processors within clusters to provide intraprocess message passing are interconnect networks that provide low-latency and high-bandwidth performance between processors.

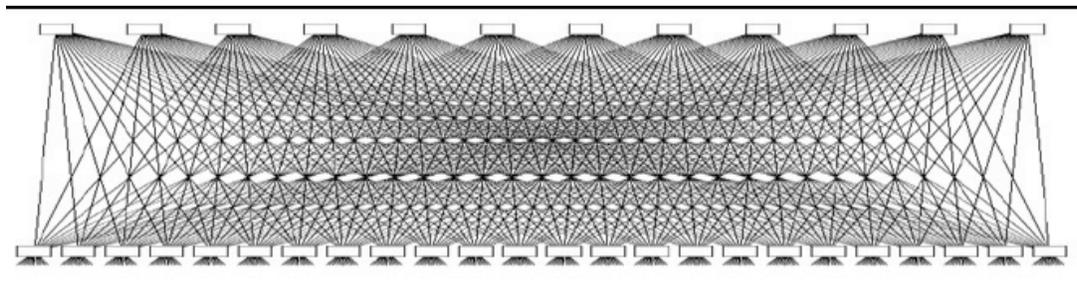
The message-passing networks or interconnect networks provide connections for internal communications between thousands of processors within the cluster as illustrated in Fig. 1.

Interconnect networks are primarily based on niche technologies designed specifically for cluster environments to provide low latency and high data rates, 10 and 20 Gigabits/s. Some interconnects are standards-based such as InfiniBand and Ethernet, while others are based on proprietary implementations such as Myrinet or Quadrics. CTN-5 is developing monitoring and troubleshooting tools to provide technicians with the ability to resolve InfiniBand issues.

A network drawing of an Interconnect network is shown in Fig. 1. Applications that are run or executed on HPC clusters produce data representing computational results that must be saved. The input/output (I/O) network provides high-performance access to the parallel storage file systems. The system administration of HPC clusters are performed through a management network built on commodity Ethernet.

The LANL Network Engineering group, CTN-5, provides network engineering design, installations, and operations support for the HPC clusters at Los Alamos National Laboratory. Specifically, CTN-5 and groups in HPC division work closely together in the parallel scalable backbone (PaScalBB) project to develop

**Fig. 1.**  
The Roadrunner cluster interconnect, used for message passing and within the cluster to move data to the I/O front ends or I/O nodes. Interconnects are built on the niche technologies like InfiniBand that will be used with Roadrunner.



future network architectures that will support the progressively higher performance clusters introduced into the LANL environment. CTN-5 is assuming support responsibilities for the various interconnect networks. During FY07, CTN-5 will be scaling up support for the InfiniBand interconnect technology used in the recently acquired clusters. CTN-5 staff provide 5x10 prime time and 7x24 on-call operations support.

The cluster I/O networks are based on the Gigabit and 10 Gigabit Ethernet technologies that CTN-5 supports on the Laboratory production networks. The Roadrunner cluster I/O network shown in Fig. 2 illustrates the relationship of I/O networks to the cluster and storage file systems. CTN-5 provides depth of knowledge in the design, installation, monitoring, and problem resolution for Ethernet-based networks.

### Conclusion

The current supercomputing paradigm is highly dependent on scalable high-performance network technologies and architectures. Future supercomputing platforms will require emerging network technologies to provide the required scalable performance.

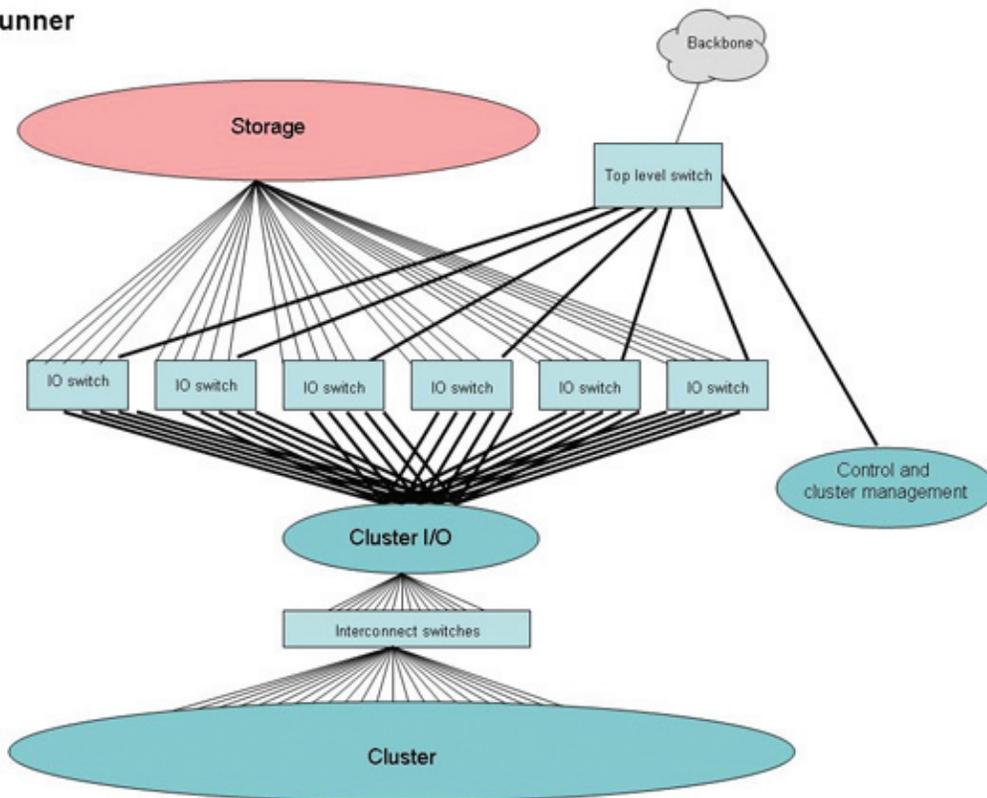
The challenge for HPC/CTN network engineers will be to apply these emerging technologies to provide the continually increasing computing demand by LANL researchers.

*For more information contact Robert Martinez at ram@lanl.gov.*

### Funding Acknowledgements

This research was supported by the NNSA tri-Lab Advanced Simulation and Computing Program.

### Roadrunner



**Fig. 2.** Roadrunner I/O network diagram illustrates the relationship between the computational cluster nodes, the interconnect, storage, and I/O switching fabric.

## The ASC Roadrunner System at LANL: Near-term capacity with a future hybrid accelerated option

Manuel Vigil, Roadrunner Project Manager, HPC-DO



**L**os Alamos National Laboratory (LANL), in partnership with International Business Machines (IBM), is in the early stages of developing and deploying the Roadrunner High Performance Computing system. LANL Divisions involved include High Performance Computing (HPC); Computer, Computational, and Statistical Sciences (CCS); Computing, Telecommunications, and Networking (CTN); Applied Physics (X); and Theoretical (T).



**Fig. 1.**  
*The Roadrunner base system has been installed at LANL.*

The primary goals for the Roadrunner system are as follows:

- Providing a large “capacity-mode” computing resource for LANL weapons simulations
- Implement and option for a petascale hybrid accelerated architecture capable of supporting the future LANL workload.
- Becoming a lead participant in the industry-wide path toward hybrid accelerated computing devices for HPC.

The Roadrunner Project has three phases. During Phase 1, already underway, LANL has acquired, installed, and deployed more than 81 teraFLOPS of a base capacity system to provide capacity computing cycles in the near term. Phase 2 is a technology refresh and assessment of the final system. Phase 3 consists of an option to install a hybrid computing architecture that has the potential for significant improvements in the price/performance curve to help meet future Advanced Simulation and Computing (ASC) computing requirements.

### Roadrunner Base System

The Roadrunner base system was delivered in September 2006. LANL completed acceptance testing of the system in December 2006. The Roadrunner base system has 14 connected units (~71 TeraFLOPS) for the classified computing environment and 2 connected units (~10 TeraFLOPS) for the unclassified computing environment. Each connected unit is 144 Opteron X64 processors from Advanced Micro Devices (AMD) connected with a high-speed InfiniBand 4X interconnect fabric. The system in the unclassified computing environment also includes initial testbeds of the IBM Cell Broadband Engine™ (Cell BE) processing elements. These testbeds are being used for initial applications and systems software work in advance of the optional hybrid system in Phase 3.

Following acceptance, the system has been undergoing a focused system integration effort for assimilation into the LANL classified computing environment and initial applications testing. This integration includes infrastructure planning and deployment of key network and I/O capabilities to accommodate Roadrunner in the Metropolis Center for Modeling and Simulation, and installing and testing the production software stack (compilers, debuggers, Message Passing Interface, I/O libraries, resource manager, etc.) required by LANL applications. System monitoring and performance tools and processes for tracking system reliability and usage statistics are also being integrated.

The system was transitioned to the classified computing environment following security accreditation in May 2007. Several key weapons applications are running on the base system in anticipation of meeting a Level-2 ASC Milestone by June 2007. The system will formally transition to production status in the summer of 2007. Roadrunner will more than double the enduring capacity computing available for nuclear weapons applications at LANL.

### Hybrid Computing Architecture – A New era in Scientific Simulation

Phase 2 is planning and assessing the hybrid architecture system targeted for deployment in Phase 3. The advanced hybrid architecture system will be both

Opteron processors and the IBM Cell Broadband Engine™ (Cell BE) processing elements. Roadrunner is the first supercomputer system to use this hybrid processor architecture.

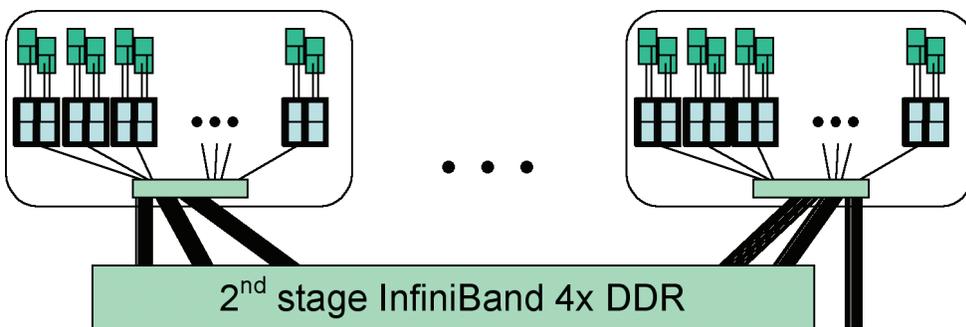
Los Alamos and IBM have been working together to improve the architecture of the Phase 3 system by improving performance both in the interconnect and between the hybrid nodes. The final classified system, planned for 2008, will have Opteron and IBM Cell Blades directly connected for sustained petaFLOP performance.

The final hybrid system will demonstrate a new paradigm for HPC, one in which accelerators are used to increase application performance rather than just adding more processors. Los Alamos and IBM are jointly developing a hybrid system programming model that will first be used in Roadrunner but will form the basis for future hybrid systems as well. Early initial performance on the Cell processors show great potential.

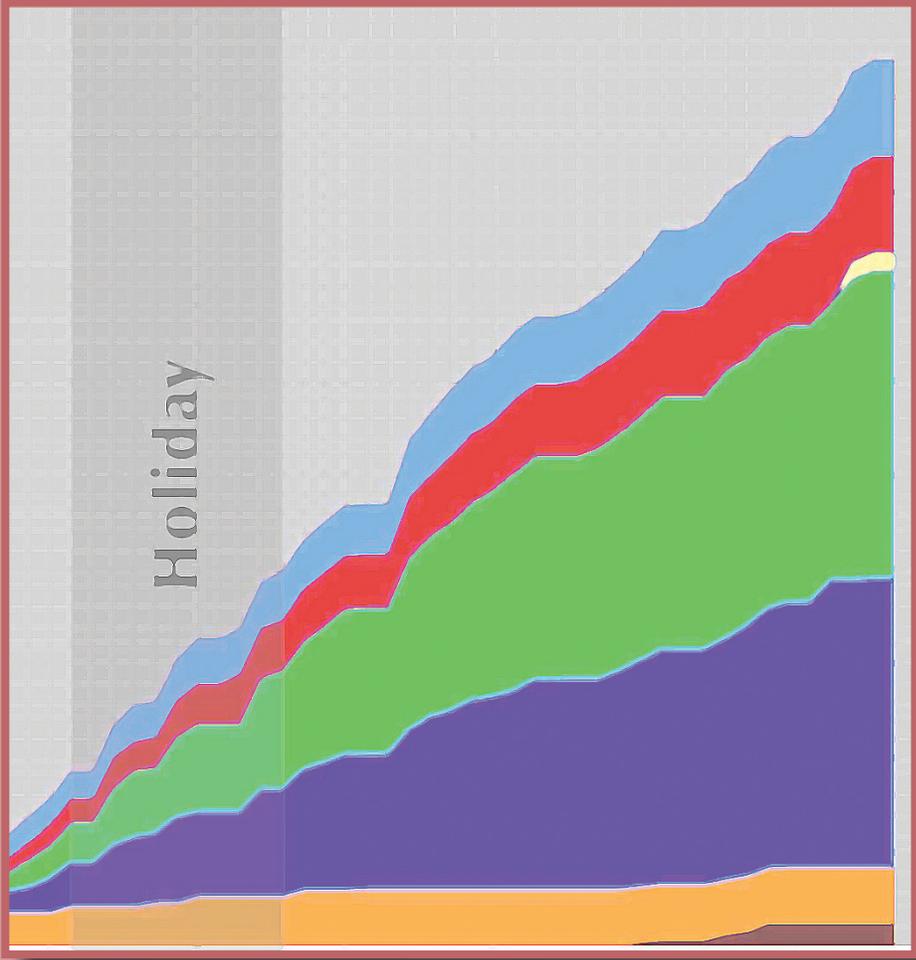
*For more information contact Manuel Vigil at [mbv@lanl.gov](mailto:mbv@lanl.gov).*

### Funding Acknowledgements

This research was supported by the NNSA tri-Lab Advanced Simulation and Computing Program.



**Fig. 2.**  
The Roadrunner hybrid architecture with Cell Blade accelerators.



# High-Performance Computing Systems Operations

The effective management of a large-scale, high-performance computer center requires a disciplined approach to system management, configuration management, and reliability and performance monitoring to enable hundreds of users to successfully use thousands of processors for their long-running parallel applications. In addition, a close collaboration with the user community is required to optimize the total effectiveness of the combined software and hardware systems. In this section we describe novel metrics to measure and optimize the overall application throughput. Approaches for monitoring both the health and configuration of overall system software components are described. Finally, this section includes a description of the New Mexico Supercomputing Challenge and its contributions to stimulating interest in high-performance computing in New Mexico's youth.

## Concurrent Application Testing Based on a Throughput Metric

John T. Daly, HPC-4

The Advanced Simulation and Computing (ASC) program has produced various metrics-based approaches for evaluating the effectiveness of a high-performance computing (HPC) system, however these typically reflect the system perspective [1]. What is lacking is a consistent set of accessible and intuitive metrics that accurately capture the application's perspective. By articulating the application throughput in terms of node performance, parallel scaling, checkpoint efficiency, and operational utilization on specified numbers of allocated nodes as follows, one can formulate requirements verification in terms of a new paradigm of concurrent application testing [2].

**Application Throughput = (Performance • Scaling • Efficiency • Availability) / Allocated Nodes**

Under this new paradigm, the performance and scaling data collected during system integration would be combined with efficiency and utilization data collected on the production platform to validate a system against specific application throughput requirements [2]. In this way, the bulk of the testing can be performed in concurrence with actual production workload.

By so doing, one increases the productive lifetime of the system by shifting weeks or even months of testing out of the integration phase. Typically, acceptance testing for a new system involves running hundreds of hours of simulated problems. By performing concurrent

application testing, the bulk of this data would be based on actual production jobs. This has the added advantage of increasing the certainty of the test results by basing them on data gathered over longer run times generated by a real workload. In particular, the efficiency and utilization are substantially governed by the application and system mean time to interrupt (MTTI) and mean time to repair (MTTR) [3]. Because these values are based on system events that occur randomly [4], measuring them over weeks or even months gives a better approximation of the expected values.

Figure 1 illustrates the concept of concurrent application testing using actual Red Storm (ASC supercomputer at Sandia National Laboratories) checkpoint efficiency results generated over a 5-month period from the run of a production problem on that system. The results demonstrate both the improvement of checkpoint efficiency over time, and the high level of agreement obtained between the model and actual data measured over significant time periods.

*For more information contact John Daly at [jtd@lanl.gov](mailto:jtd@lanl.gov).*

[1] J.R. Stearley, "Defining and Measuring

Supercomputer Reliability, Availability, and Serviceability (RAS)," <http://www.cs.sandia.gov/~jrstea/ras> (2005).

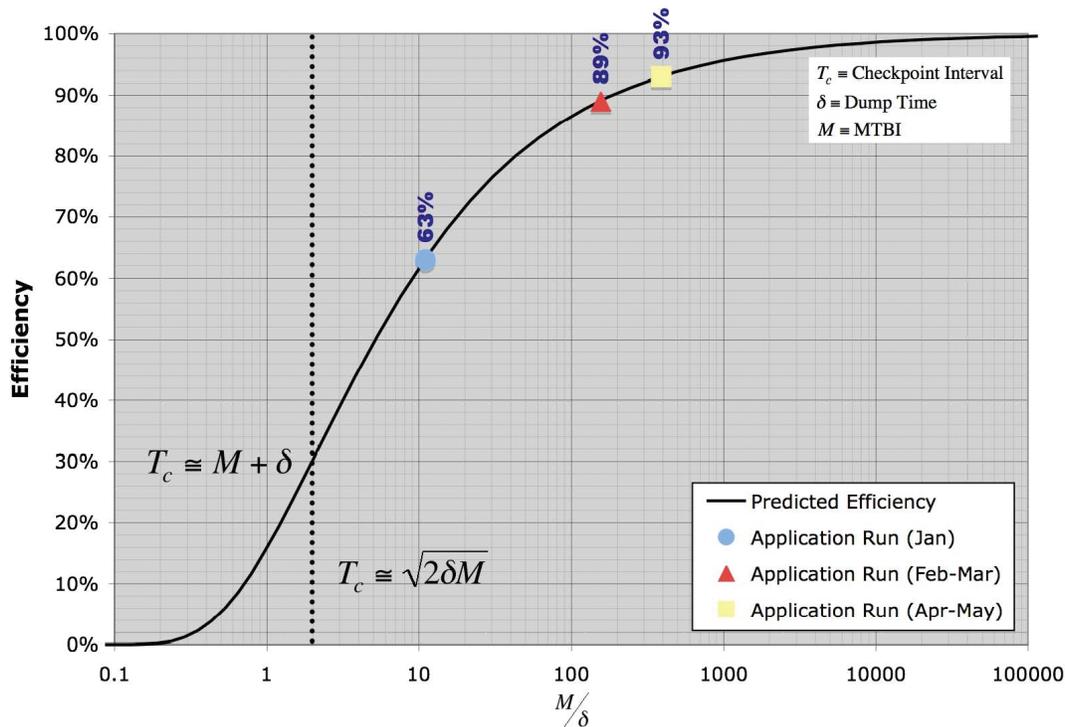
[2] J.T. Daly, "Methodology and Metrics for Quantifying Application Throughput," Proceedings of the Nuclear Explosives Code Developers' Conference (NECDC), in press.

[3] J.T. Daly, "Evaluating the Performance of a Checkpointing Application Given the Number and Types of Interrupts," Proceedings of the First Workshop on High Performance Computing Reliability Issues (HPCA-11), San Francisco, (2005).

4] J.T. Daly, "Failure Analysis From the Application's Point of View", Los Alamos Technical Report, LA-UR-06-8089.

### Funding Acknowledgements

This research was supported by the NNSA tri-Lab Advanced Simulation and Computing Program.



**Fig. 1.** Checkpoint efficiency as a function of dump time, application MTBI, and the restart overhead in blue. The optimum checkpoint interval is plotted in red to demonstrate its asymptotic behavior for large and small values of the nondimensional dump time.

# Facilitating High-Throughput ASC Calculations

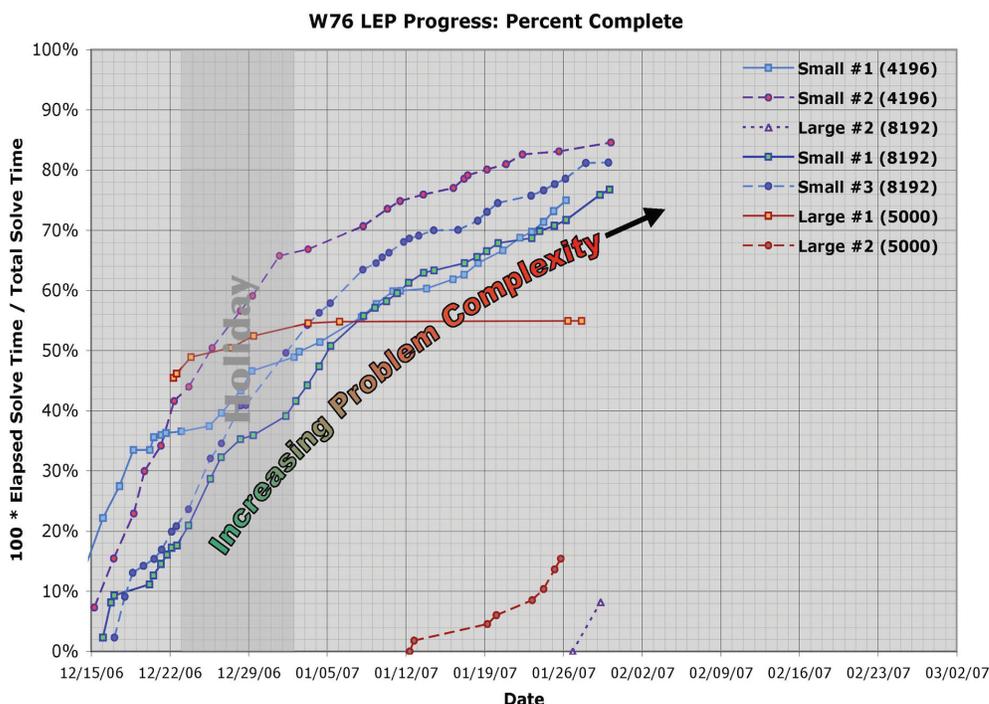
John T. Daly, HPC-4

The productivity of the application user is arguably one of the most significant metrics upon which high-performance computing (HPC) systems performance is evaluated for the Advanced Simulation and Computing (ASC) program. One important component of that productivity is the application throughput associated with the calculation. Application throughput for *capacity* calculations, characterized by large numbers of jobs on relatively few processors, is typically limited by the performance and scaling of an HPC system. However, in the case of capability calculations, characterized by relatively few jobs run on a significant fraction of the platform, many thousands of processors, the application throughput, or movement of data, is frequently limited by the reliability and availability of the system [1,2].

This paper discusses improvements in the throughput of HPC applications on capability systems realized by developing and implementing a unique *job\_suite* toolkit.

The *job\_suite* toolkit—which includes *job\_submit*, *job\_view*, *job\_list*, and *job\_kill*—was developed by the author specifically for the purpose of keeping these capability class calculations running and maximizing system utilization in the face of various types of interrupts. This is accomplished by a combination of job monitoring, to check that a running application is making progress, and automated restart, to get a stopped application running again after an interrupt. The *job\_suite* tools have run successfully on ASC platforms including Red Storm, Purple, BG/L, and Q. They can interact with jobs via any of the tri-Lab-supported schedulers

**Fig. 1.** Progress on the W76 LEP calculations in terms of operational utilization over time.



(i.e., LSF, LCRM, PBS, or Moab) or through a Posix interface on platforms without a scheduler.

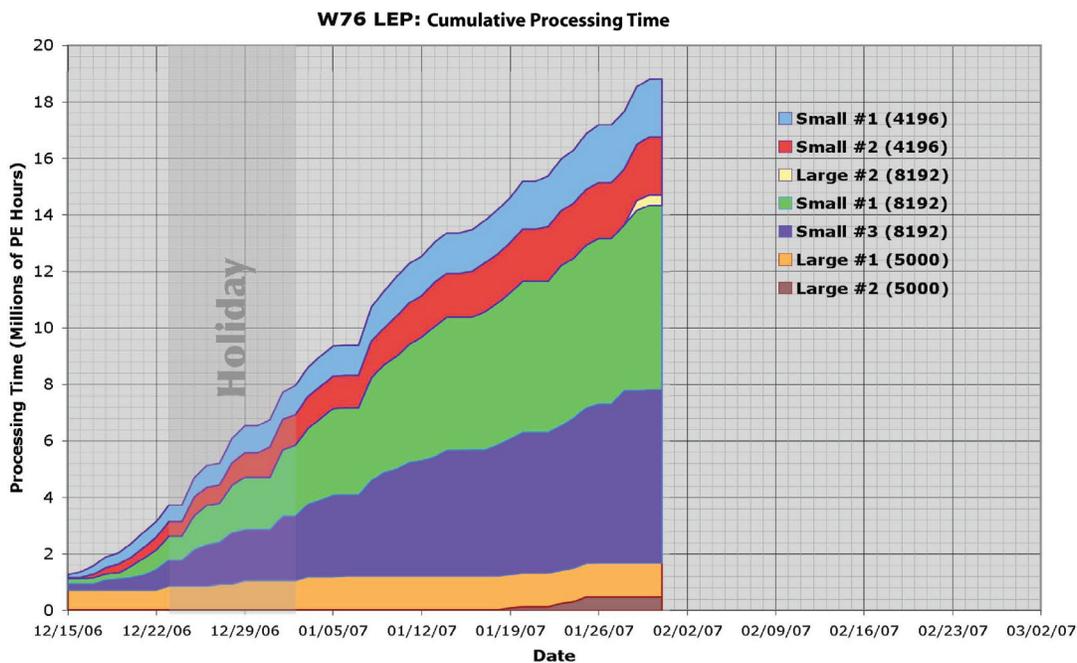
The *job\_suite* set of tools has been used to facilitate user productivity for other calculations [3], but the most compelling example of its effectiveness comes from the W76 LEP calculation run in tandem by a single user on Red Storm, Purple, and BG/L between December 2006 and February 2007. This calculation consisted of seven high-priority jobs, each utilizing 4096 to 8192 processors. It was consistently run on one-quarter or more of the available compute resources of all three platforms. The jobs typically sustained an aggregate of 11 interrupts per day. Figure 1 illustrates the operational utilization over time of the seven different jobs, and Fig. 2 depicts the wall-clock time accumulated by the jobs in processor hours. This level of throughput in face of frequent interrupts represents a substantial HPC capability for the weapons program.

[1] J.T. Daly, "Methodology and Metrics for Quantifying Application Throughput," Proceedings of the Nuclear Explosives Code Developers' Conference (NECDC), in press.  
 [2] J.T. Daly, "A Higher Order Estimate of the Optimum Checkpoint Interval for Restart Dumps," *Future Generation Computer Systems* **22**, 300 (2006).  
 [3] J.T. Daly, "CIS External Review: A Customer Perspective on Red Storm," Los Alamos Technical Report, LA-UR-06-3802.

*For more information contact John Daly at [jtd@lanl.gov](mailto:jtd@lanl.gov).*

**Funding Acknowledgements**

This research was supported by the NNSA tri-Lab Advanced Simulation and Computing Program.



**Fig. 2.** Progress on the W76 LEP calculations in terms of accumulated wall-clock hours of CPU-core time.

## Seventeen Years of the New Mexico Supercomputing Challenge

David Kratzer, HPC-3

In an effort to increase the future pool of talented scientists, the Los Alamos National Laboratory (LANL) Theory, Simulation, and Computation Directorate and the High-Performance Computing (HPC) Division support the New Mexico Supercomputing Challenge (Challenge), a program that offers a unique opportunity to students in New Mexico. Since 1990, teams of New Mexico high school students have worked for a school year to develop a HPC computational science project. More than 7,100 students and 1,200 teachers have participated in the program.

The Challenge has as its mission to improve students' understanding and use of technology by developing their skills in scientific inquiry, modeling, computing, communications, and team work. Each team of up to five students and a sponsoring teacher defines and works on a single computational project

of its own choosing. They submit a proposal at the beginning of the school year, prepare an interim report in December, present a preliminary report of their work in February, and submit a final written report in April. The teams come to LANL at the end of April for oral presentations of their projects and are given tours of the Laboratory. An awards ceremony takes place the next day where scholarships, savings bonds, trophies, and other awards are given to participating students, totaling more than \$90,000 in 2006 and 2007, in large part annual contributions of \$80,000 by the Computing, Computational, and Statistical Sciences (CCS) Division.

Challenge participants learn team work, time management skills, how to research a scientific topic, how to use a computer as a tool to model or simulate their project, and how to present their results orally, in written form, and by posting a web page.

**Fig. 1.**  
*Students receive training in Mathematica at the 2007 Challenge Kickoff Conference.*



Laboratory employees serve as team mentors, help judge the projects, and give talks and demonstrations to the Challenge participants when they come to Los Alamos for the year-end Expo and awards ceremony.

During the summers, a two-week institute for teachers is held to give the teachers professional development to better prepare them to sponsor Challenge teams.

*For more information, contact David H. Kratzer [dhk@lanl.gov](mailto:dhk@lanl.gov) or visit <http://www.challenge.nm.org>.*

### **Funding Acknowledgements**

Funding is provided by the State of New Mexico, businesses, New Mexico educational institutions, and LANL.



**Fig. 2.**  
*One of the LAHS Challenge teams (l. to r.) Ben Batha, Ryan Marcus, mentor Neale Pickett, Daniel Cox, and William Phillips.*

*Photo by S. King*

In a recent check, over 40 past Challenge participants are current LANL staff members and another 50 have been employed by LANL at some point. Neale Pickett, CTN-5, Network Engineering, is a past participant who remains involved in the program. In 1991, as a student at Albuquerque Academy, he and three teammates submitted a project on solar magnetic pole reversal, inspired by one of the students watching a PBS Nova episode on the subject. The model didn't work, but the simulation, done on the Connection Machine, went well, Neale recalls, and he and his team took first place in the very first Challenge.

Seventeen years later, with a degree from New Mexico Tech, Neale is a computer security programmer at LANL and mentor for the Los Alamos High School (LAHS) 2007 Supercomputing Challenge team. The four LAHS students are applying the traveling salesman problem to the Mars Rover route, computer networking efficiency, and nano-medicine delivery.

## Request Tracker (RT): A Tool for Software Change Control

Georgia A. Pedicini, David R. Montoya, HPC-4

The steps involved in any change-control process are relatively obvious: request, approve, schedule, perform, verify. The trick is to identify, and track, specific tasks associated with changes in a given environment. The high-performance computing (HPC) software support team has begun to use an open-source product called RT (Request Tracker) to track changes to the software products installed on production computers in the Laboratory Integrated Computing Network (ICN).

### RT capabilities

RT has built-in capabilities to manage task details at any level. Separate tickets can be generated for general tasks or detailed subtasks, and tickets can be grouped together by dependencies. For example, if a general ticket is defined for a software product update, then additional tickets can be entered for each machine or subset of machines. This allows incremental progress to be tracked more effectively.

Each ticket contains detailed history of all actions taken, from creation to email notifications, to each individual field update. If questions arise about how a change was managed, all the subtasks can be reviewed later. Having access to this level of detail is important for tracking the effectiveness of a change control process. It can also help improve the process for later requests.

RT allows actions to be linked to certain ticket updates. The most powerful action

is to send notifications to a defined set of users whenever a ticket status changes. This improves cross-team collaboration by making sure everyone is up to date on the next required task.

Tickets can be reviewed through standardized reports, such as all “open” tickets, but RT’s search tool also allows complex queries using any ticket fields (Fig. 1). These queries allow all users to quickly drill down to the most important information for that user, or in a specific context. Examples include all open tickets owned by a user, or all tickets containing this combination of terms.

RT’s flexible customizations are even more important than the built-in features. HPC has added custom “status” and “machine” fields. Responsibility for software updates changes hands several times between multiple subteams, and the custom status field allows for easier communication when this hand-off occurs. Due to complexities in the production ICN environment, a software change is generally staged across the environment. Separate tickets for a custom list of machines make it practical to keep track of progress.

### RT work in progress

One of the strengths of using the RT tool is the ability to customize it to fit the needs or adjust as better approaches are discovered. The custom status field has been updated to more accurately reflect the different steps of a software upgrade. Recent discussion clarified how these status values map to other stages of

software lifecycle, such as removing old versions. Notifications are being reduced from the built-in “daily” to a manageable “when something important happens,” where “important” will surely continue to be a topic for discussion. Tracking granularity is being changed to one ticket per production machine. As a result, we are experimenting with an open-source add-on and ticket cloning to make ticket creation easier.

The teams involved in software change control meet regularly to discuss experiences. We have worked through initial discrepancies in the mechanics of ticket handling. We continue to add customizations, gradually reflecting the reality of multiple ticket hand-offs and special requirements of the ICN. RT is helping these cross-organizational teams keep better control over software changes in a complex production environment.

*For more information contact Georgia Pedicini at gap@lanl.gov.*

### Funding Acknowledgements

This research was supported by the NNSA tri-Lab Advanced Simulation and Computing Program.

**Fig. 1.**  
*Query of open tickets in process.*

#	Subject	Status	Owner	Requestors	SW status	Yellow	Turquoise	Red	Product version
852	Install Intel 9.1 in Turquoise	open	markus	gap@lanl.gov	Pending		Pink-32 Compile Pink-32 Master TLC	Intel compiler	9.1
764	PGI compiler on Saguario needs license.dat and module files	open	Nobody	oog@lanl.gov	Tested	Saguaro Compile Saguaro Master		pgi	6.2-4
684	PGI 6.2-4 module files installed on Lightning	open	gap	oog@lanl.gov	Ready		Lightning-32 Compile Lightning-32 Master Lightning-64 Compile Lightning-64 Master	pgi	6.2-4
150	Totalview 7.2 / Red	open	gap	gap@lanl.gov	Installed		Lightning-32 Compile Lightning-32 Master Lightning-64 Compile Lightning-64 Master	TotalView	7.2
973	PGI 6.2-4 install Turquoise	open	gap	gap@lanl.gov	Installed		Coyote-64 Compile Coyote-64 Master	pgi	6.2-4

## Monitoring Software Health in a Heterogeneous HPC Environment

David R. Montoya, Georgia A. Pedicini, David O. Gunter, Jennis Pruett, Jennifer M. Frye, HPC-4

**H**igh-performance computing (HPC) clusters are often just reaching production quality when new hardware is arriving for the next capacity or capability cluster. Newer machines have different system, network, and input/output (I/O) capabilities, or may have a different architecture, as will be the case with the Roadrunner hybrid system. The result is a loss of consistency across generations of hardware and software. Software tools must be available to help manage and plan for a diverse environment with different software requirements and configurations.

In a constantly changing HPC environment, we need a global view of system health attributes. Automated tasks need to check the health of the systems. For example, what products are available or missing from all the relevant segments?

In our current environment, “segments” describe the individual units that make up a portion of the larger machine. On some clusters, different segments provide different functionality and must be viewed as separate machines. Front-end and compile nodes add another dimension of complexity. Take these all into account and we have more than 50 machines to monitor.

In the past few years we have been making plans to better support the changes in HPC environments. We have instituted projects to collect data from across clusters to support managing the

software. An initial tool allows for our software environment to be analyzed, assessing consistency across machines. Looking at the complete consolidated module listing (Fig. 1), there will be differences: certain products will not run with a specific processor, or at a kernel release, or library versions are not compatible with a specific release.

This display is a visual approach to the question of software availability across the production environment. Major differences are easily apparent. A nightly collection process also collects details about product versions. Further enhancements will add these details, making it easier to pinpoint minor discrepancies within clusters.

### System Health

The capability represented in Fig. 1 is just an initial step to a broader infrastructure to manage the health of the software environment. Two other questions need data: (1) whether the available products are in good working condition, and (2) evaluating usage patterns. Both efforts are underway. A regression-test suite is being added to the data collection framework. “Lite” tests can verify minimal product functionality; more complex tests can be developed to evaluate performance. We anticipate nightly regression runs reporting success or failure. For example, we could highlight if there were any failures for a given product or machine segment. Another capability would log details of the test results. Usage data collection is also being prototyped to collect usage

statistics for users and products. We are adapting the custom support service model to our heterogeneous HPC environment. Tools and data are needed to manage the software environment, to make decisions about which products provide value to the environment, and to manage product lifecycles. Initial steps have been taken, and enhancements are planned. With a focus on efficiency and value added, we are building a sustainable support capability.

*For more information, contact David Montoya at [dmont@lanl.gov](mailto:dmont@lanl.gov).*

### Funding Acknowledgements

This research was supported by the NNSA tri-Lab Advanced Simulation and Computing Program.

	pfe1	pink	tlc	cy-c1	cy-c2	cy-c3	mauve	cy-1	cy-2	cy-3	cy-4	cy-5
openmpi				X	X	X		X	X	X	X	X
papi	X	X										
pathscale				X	X	X						
perl	X	X										
pgi				X	X	X						
pgi32				X	X	X						
procmon	X	X										
purify	X	X										
purifyPlus	X	X										
python	X	X	X	X	X	X		X	X	X	X	X
sgi-mpt							X					
subversion	X	X		X	X	X		X	X	X	X	X
tcltk	X	X										
tkdiff	X	X										
totalview				X	X	X		X	X	X	X	X
ups	X	X										
unbrind				X	X	X		X	X	X	X	X

**Fig. 1.** A sample segment of software product distribution on the LANL Turquoise network. Tools are listed in the left column and machine segments along the top.



# Author Index

## A

Abdallah, Joseph Jr.....	140
Addressio, Frank L.....	130
ASC Telluride Project Team.....	90
Albers, Robert C.....	166
Alder, Berni.....	44
Anderson-Cook, Christine.....	80
Andrews, Malcolm J.....	36, 68
Aulwes, Robbie.....	190, 192

## B

Banerjee, Arindam.....	68
Barber, John L.....	44, 170
Bardenhagen, Scott G.....	164, 176
Barrett, Brian.....	2
Barker, Kevin J.....	4
Berndt, Markus.....	98
Bock, Nicolas.....	110
Boettcher, Stefan.....	12
Boettger, Jonathan.....	116
Bowman, Todd.....	194
Bradley, Paul A.....	152
Buksas, Michael W.....	16, 48, 58
Bulla, R.....	128
Burakovsky, Leonid.....	126

## C

Cady, Carl W.....	176
Castain, Ralph.....	2
Cawkwell, Marc J.....	172, 174
Challacombe, Matt.....	138
Chang, Jae H.....	62
Chen, Shao-Ping.....	116
Chertkov, Michael.....	72
Chisolm, Eric.....	110, 116
Chumakov, Sergei.....	180
Clements, Brad T.....	120, 122, 130
Colgan, James.....	140
Collins, Lee.....	142
Cowell, Shannon.....	146, 148
Crockett, Scott.....	110, 116, 124, 126

<b>D</b>	
Dai, William W.....	104, 190, 192
Daly, John T. ....	188, 200, 202
Dattelbaum, Dana M. ....	176
Davenport, Carolyn Conner .....	8, 24
Davis, Kei .....	4
DeBardleben, Nathan .....	30
Dendy, Edward.....	40
Densmore, Jeffery D.....	16, 48, 52, 56
Devlin, Matthew.....	146, 148, 150
Dimonte, Guy .....	36
Dreisigmeyer, David W. ....	102
DuBois, Andrew .....	8, 24
DuBois, David.....	8, 24
Dyadechko, Vadim.....	94

<b>E</b>	
Edgar, B. Kevin.....	104
Edwards, Jarrod.....	38
Evans, Thomas M.....	16, 48, 50, 52, 58

<b>F</b>	
Fontes, Christopher J. ....	140
Fotiades, Nik.....	146, 148
Francois, Marianne M.....	40, 42, 76
Friar, Jim .....	158
Frye, Jennifer M.....	208
Fryer, Chris.....	16, 58
Fugate, Michael L.....	74, 82
Fung, Jimmy.....	40, 100

<b>G</b>	
Garimella, Rao V.....	42
Gattiker, Jim .....	82
Germann, Timothy C.....	44, 166, 168
Goncalves, Bruno T.....	12
Graham, Richard L.....	186
Grove, John W.....	34
Graves, Todd.....	80
Greeff, Carl.....	110, 116, 124
Gunter, David O. ....	184, 208

<b>H</b>	
Hale, Gerald .....	150
Hamada, Mike .....	80
Hanson, David E .....	136
Hansson, Anders .....	12
Hayes-Sterbenz, Anna C. ....	152, 154, 158, 160
Higdon, Dave.....	82
Higginbotham, A.....	166

Hodson, Steve.....	10
Hogden, John.....	18
Holian, Brad L. ....	44, 166
Honnell, Kevin.....	116
Horner, Dan .....	142
Hungerford, Aimee.....	16, 58

**I**

Ionita, Axinte .....	136
Istrate, Gabriel .....	12

**J**

Johnson, J.D. ....	116
Joubert, Wayne .....	96
Jungman, Gerard.....	152, 158, 160

**K**

Kadau, Kai.....	44, 166, 168, 170
Kanzleiter, Randy.....	58
Kasiviswanathan, Shiva P. ....	12
Kawano, Toshihiko .....	146, 148
Kelley, Timothy M.....	16, 54
Kenamond, Mark A. ....	40
Kerbyson, Darren J.....	4
Ketelsen, Christian.....	98
Kirchner, S. ....	128
Kober, Edward M.....	178
Kothe, Douglas B.....	76
Kratzer, David.....	204
Kress, Joel D. ....	134, 142

**L**

Lally, Bryan .....	96
Lewis, Mathew W.....	176
Li, Shengtai.....	84
Lipnikov, Konstantin .....	100
Livescu, Daniel .....	66
Lomdahl, Peter S. ....	44, 166, 168
Loncaric, Josip.....	186
Lookman, Turab .....	178
Loubere, Raphael .....	42
Lowrie, Robert B.....	38, 40
Lushnikov, Pavel M. ....	70

**M**

Maniadis, Panagiotis .....	178
Martinez, Robert.....	194
Masser, Thomas .....	34
Merrigan, Tim.....	194
Modl, David.....	26, 28

Möller, Peter .....	156
Monroe, Laura .....	24, 26
Montoya, David R.....	184, 206, 208
Mosher, Scott.....	16, 50, 54, 56
Moulton, J. David.....	98, 100

## **N**

Nelson, Ronald O.....	146, 148
Niklasson. Anders M.N.....	138
Nunez, James .....	10

## **O**

## **P**

Patterson, Brian M.....	176
Pedicini, Georgia A. ....	184, 206, 208
Peery, Travis .....	110
Pernice, Michael .....	96
Peterson, Darrell.....	58
Philip, Bobby.....	96
Pimentel, David.....	116
Plohr, JeeYeon N.....	130
Poole, Stephen .....	8, 24
Percus, Allon G.....	12
Pruett, Jennis.....	208
Pugmire, David .....	24, 26, 88

## **Q**

## **R**

Ramaprabhu, Praveen.....	36
Rasmussen, Bryan .....	106
Rasmussen, Craig E. ....	14
Rasmussen, Kim Ø.....	178
Rauenzahn, Rick M.....	38
Reichhardt, Cynthia.....	136
Reidys, Christian .....	12
Rice, Denny .....	194
Rider, William J.....	84
Ristorcelli, J.R. Jr. ....	66
Rosa, Massimiliano .....	62
Rose, Harvey A.....	70
Rudin, Sven P.....	112, 114, 116, 124, 126
Runnels, Scott .....	100
Rust, William .....	30

## **S**

Sancho, Jose Carlos .....	4
Schofield, Samuel P.....	42

Sewell, Thomas D.....	172, 174
Shashkov, Mikhail .....	94
Shaw, Steven R.....	186
Shewmaker, Andrew .....	28
Shipman, Galen .....	2
Si, Qimiao .....	128
Sicilian, James M. ....	76
Sierk, Arnold J.....	156
Sievers, Cindy.....	20
Sims, Benjamin .....	80
Staley, Martin .....	22
Streufert, Ben .....	140
Swaminarayan, Sriram.....	6

## T

T-1 Equation of State (EOS) Team.....	118
Talou, Patrick .....	146, 148
Trellue, Holly .....	154
Turner, John.....	6
Tymczak, C. J.....	138

## U

Urbatsch, Todd .....	16, 58, 60
----------------------	------------

## V

Vander Weil, Scott .....	80
Vigil, Manuel.....	196
Vladimirova, Natalia .....	72, 180

## W

Wallace, Duane .....	116
Wallstrom, Timothy .....	74
Wark, J. S.....	166
Warsa, James S. ....	62
Watson, Greg.....	14
Williams, Brian .....	82
Wills, John .....	116, 124
Wilson, Alyson .....	80
Wilson, Bill .....	154
Wilson, Greg.....	80
Wood-Schultz, Merri.....	160
Woodward, Paul R. ....	104

## X

## Y

## Z

Zhu, Jian-Xin.....	128
--------------------	-----

