LA-UR-12-24563
Approved for public release; distribution is unlimited.

Title: Disruptive Innovation in Numerical Hydrodynamics

Author(s): Waltz, Jacob I.

Intended for: This is an LDRD proposal that needs an LA-UR for distribution to potential postdoctoral candidates.

Disclaimer:
Los Alamos National Laboratory, an affirmative action/equal opportunity employer,is operated by the Los Alamos National Security, LLC for the National NuclearSecurity Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Departmentof Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

Disruptive Innovation in Numerical Hydrodynamics

PI: Waltz, Jacob, I; XCP-8; jwaltz@lanl.gov

Introduction

We propose the research and development of a high-fidelity hydrodynamic algorithm for tetrahedral meshes that will lead to a disruptive innovation [1] in the numerical modeling of Laboratory problems. Our proposed innovation has the potential to reduce turnaround time by orders of magnitude relative to Advanced Simulation and Computing (ASC) codes; reduce simulation set-up costs by millions of dollars per year; and effectively leverage Graphics Processing Unit (GPU) and future Exascale computing hardware. If successful, this work will lead to a dramatic leap forward in the Laboratory's quest for a predictive simulation capability.

Project Goals

Our strategic goal is to deliver a disruptive innovation in the numerical modeling of Laboratory problems. Since its inception, the ASC program has made remarkable progress in the development of Laboratory simulation tools. However, in recent years the rate of that progress has slowed and is now largely incremental in nature. We seek to break that pattern by developing a novel simulation methodology that combines algorithmic advances with effective utilization of emerging architectures, such as the Graphics Processing Unit (GPU). Taken together, these characteristics will enable a dramatic leap forward in predictive capability.

Our R&D goal is to develop a multi-material hydrodynamic algorithm for tetrahedral meshes that is applicable to high-strain-rate Laboratory problems and designed a priori for next-generation compute platforms. Accomplishing this goal requires the simultaneous solution of two problems. First, the historical difficulties associated with tetrahedral meshes when applied to Laboratory applications must be overcome if the benefits of such meshes are to be realized. Doing so will require the development of innovative algorithms as well as the validation of existing algorithms that are heretofore unproven for Laboratory problems. Second, because emerging architectures such as the GPU are expected to differ dramatically from the machines they will replace, the potential performance gains of existing methods are inherently limited due to having been designed for hardware concepts that are decades old. Our simulation methodology must be redesigned to best leverage this new class of hardware. The specific algorithm we seek to develop combines high-order, unsplit hydrodynamic methods with tetrahedral meshes and Adaptive Mesh Refinement (AMR). Based on preliminary studies we expect to achieve a 300-fold performance improvement relative to major ASC codes on current compute platforms and an additional 30-fold performance improvement on GPU-based compute platforms, while simultaneously improving both accuracy and ease of setup.

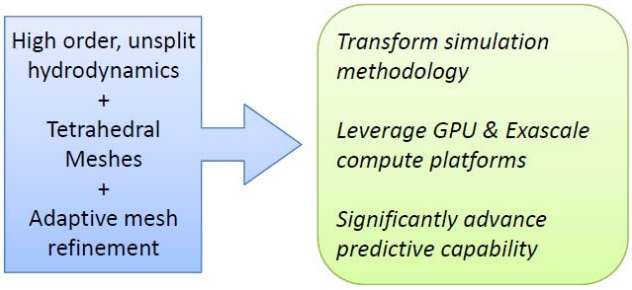

Figure 1: Our R&D and strategic goals

Our Verification and Validation (V&V) goal is to establish our methodology as a trusted, credible approach to Laboratory problems. Because our proposed methodology represents a significant departure from traditional approaches, rigorous and thorough testing is critical to both find-

ing flaws and building confidence. We intend to employ frequent verification testing throughout the R&D process using standard test problems, and to validate our methodology against experimental data from several high-strain-rate multi-material flows, which is the domain of interest.

Our technological goal is to establish a path forward for future codes on emerging architectures. Emerging architectures are expected to have low per-core memory and/or low memory bandwidth; a high potential FLOPs rate; and both threaded and vector components. Additionally, any Laboratory production system will not be a single compute node but rather a cluster of compute nodes. Therefore, message passing also will be required, but as part of a hierarchical approach that combines message passing with threading and vectorization. We intend to develop an approach that accounts for all of these factors and can achieve the overall fastest time to solution for a given compute platform. Because the optimal combination of these features will be platform-dependent, our approach will allow that combination to be selected at runtime.

Background and Statement of Problem

The ability to use tetrahedral meshes for Laboratory problems would provide several advantages over current hexahedral and arbitrary unstructured methods. First, tetrahedral meshes can be rapidly generated with commodity tools [2], which would reduce problem definition time while saving millions of dollars in setup costs. Second, hydrodynamic algorithms for tetrahedra can be highly optimized. Orders of magnitude performance improvements relative to ASC codes are possible [3-5]. Third, numerous algorithms have been developed for tetrahedral meshes in areas such as diffusion [6], transport [7,8], particle tracking [9], laser ray tracing [10], electromagnetics [11], and AMR [12,13]. The ability to leverage such algorithms would reduce code R&D costs.

Because of these recognized benefits the desire to use tetrahedral meshes for Laboratory applications is not new, and there have been previous attempts to do so. However, while algorithms based on such meshes have performed well for engineering fields such as Computational Fluid Dynamics (CFD) [14,15] and Computational Electromagnetics (CEM) [*e.g.* 11], these algorithms have not been successfully applied to high-strain-rate multi-material flows due to issues such as mesh stiffness, mesh imprinting, and spurious vorticity. Two examples from LANL are a series of Free Lagrange codes developed by Trease *et al* [16-18] in the late 1980s, and the CHAD Arbitrary Lagrangian-Eulerian (ALE) code [19-20], developed under the ASC Antero Project in the late 1990s and early 2000s (CHAD actually supported multiple zone types, but was tested with pure tetrahedral meshes on a number of problems [21]).

The Free Lagrange codes utilized a point-centered Finite Volume algorithm, which in principle can alleviate a problem known as mesh stiffness. As shown in Figure 2 for the 2D case, mesh stiffness occurs in Lagrange calculations when zones cannot deform isochorically, and it manifests as density errors that in turn lead to errors in pressure. The use of point-centered variables creates control volumes that are complex polyhedra and therefore generally not susceptible to mesh stiffness. However, because these codes differenced the point-centered pressure against a zone-centered pressure, the algorithm was stiff.

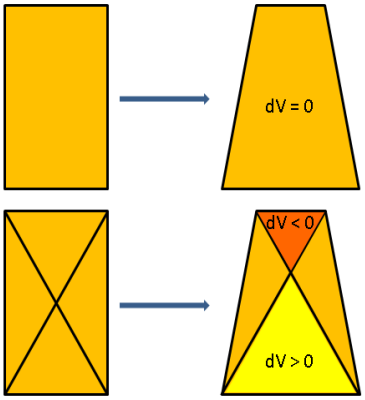

Figure 2: A quadrilateral can easily deform isochorically but triangles cannot. The result is density errors referred to as mesh stiffness.

CHAD mitigated mesh stiffness by combining unsplit time integration of the Lagrange and advection terms with a stencil that

only used point-centered pressures. Unsplit time integration approximates the face bending of an isochoric deformation by advecting excess material into or out of neighboring zones. However the CHAD algorithm had other flaws: the time integration method was inefficient in the high-speed flow regime; the mesh motion algorithm violated the Geometric Conservation Law, which caused mesh imprinting; and the discrete operators generated artificial vorticity.

Based on these observations, we posit that a hydrodynamic algorithm for tetrahedral meshes must minimize errors due to mesh stiffness, mesh imprinting, and artificial vorticity if it is to be successfully used for Laboratory applications. The development of such an algorithm constitutes the basic challenge of this proposal.

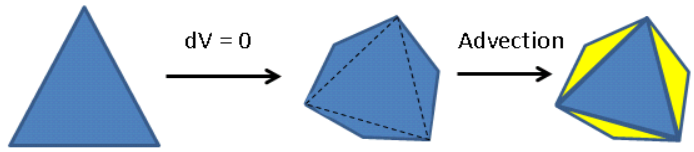

Figure 3: Unsplit time integration uses advection to approximate face bending. The highlighted portions are advected into neighboring zones simultaneously with the Lagrange step to avoid density errors.

However, the imminent deployment of GPU and emerging architectures creates additional constraints on physics algorithm development. The limited per-core memory of emerging architectures has two important consequences: first, algorithms with a high ratio of floating point operations to memory operations (referred to hereafter as *compute intensity*) are optimal; second, the traditional approach of reducing simulation errors through increased mesh resolution is no longer viable. While efforts are underway to deploy ASC algorithms on GPU hardware, whether or not these algorithms are appropriate for such architecture in the first place is unclear.

Either the development of a tetrahedral-based simulation methodology or the deployment of current algorithms on GPU hardware would represent a significant accomplishment. However, individually, neither would lead to the advance in predictive capability that we desire. Development of a new simulation methodology that did not account for emerging architectures might become a dead-end path on future hardware. Likewise, deployment of existing algorithms on emerging architectures will not fundamentally change the predictive capability of those algorithms. The combined challenge, therefore, is the problem we seek to solve: can we develop a hydrodynamic algorithm based on tetrahedral meshes that transforms simulation methodology *and* fully leverages emerging architecture, and in doing so deliver a disruptive jump in predictive capability.

Preliminary Studies

Existing technology base: Our proposed research code, CHICOMA, solves the compressible hydrodynamic equations on 3D tetrahedral meshes with AMR. CHICOMA is parallelized using a shared memory (threaded) approach and has been deployed on GPU hardware currently installed at LANL. CHICOMA has been demonstrated to solve 3D gas dynamics problems approximately 300-400 times faster than major ASC codes [4], which is the basis for our aforementioned performance goals. We also have demonstrated fast setup with the mesh generation tool CUBIT [2]. While the Finite Element (FE) hydrodynamic algorithm in CHICOMA is designed for engineering applications rather than Laboratory physics problems, the code can be used as a test bed for the proposed research. This approach will allow rapid implementation and testing in CPU, GPU, and hybrid environments, and eliminates the need to build a test code from scratch. The existing 3rd-order accurate time integration and AMR package can be used without modification.

Point-centered hydrodynamic algorithm: Using analytical tools developed by Burton and Morgan, we have derived a new point-centered unsplit ALE hydrodynamic algorithm [22] that addresses all of the sources of error listed previously. This algorithm was derived specifically for

implementation in CHICOMA and is compatible with both high-order spatial methods and AMR. While the scheme has yet to be implemented, it has been reviewed by local subject matter experts and is believed to be viable [23]; further details are discussed subsequently.

GPU optimization: Recent work [24] explored memory optimization techniques for GPU-based compute platforms. These techniques increase compute intensity by using graph theory methods to lower the cost of memory operations. Using CHICOMA as the test code, performance gains of 4-5 times were achieved for 3D gas dynamics problems. A hybrid CPU-GPU execution mode also was demonstrated. This work will serve as a reference baseline for GPU performance.

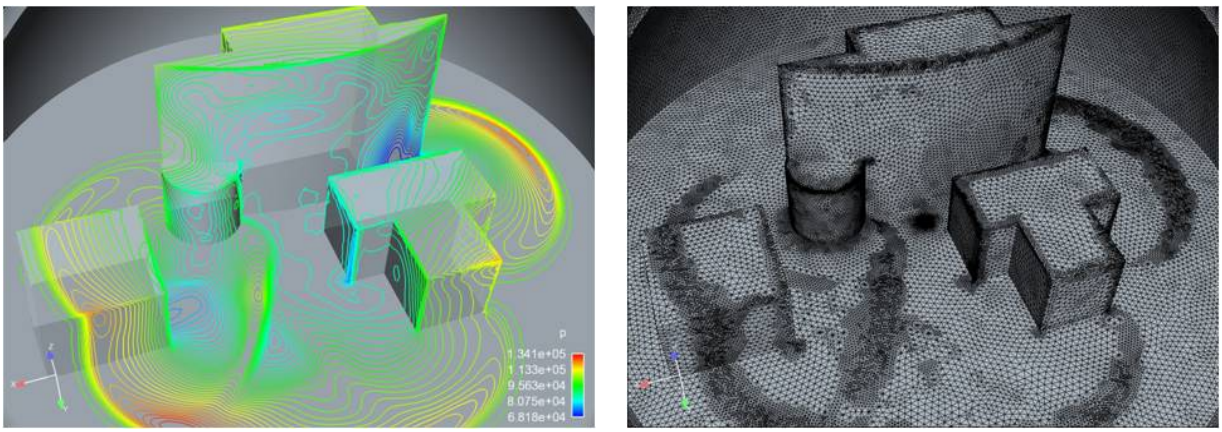

Figure 4: Surface pressure contours and adapted surface mesh from a 3D CHICOMA simulation of a ~ 0.1 ton air blast in a building complex using ~ 10^7 tetrahedra and AMR.

Proposed Innovation and Significance

We posit that a hydrodynamic algorithm suitable for Laboratory problems and emerging architectures can be developed for tetrahedral meshes. Our proposed innovation that will enable testing of that hypothesis is a hydrodynamic algorithm that combines the following:

- Vorticity compliance, to minimize rotational instabilities and spurious vorticity
- Unsplit high-order time integration, to minimize the deleterious effects of mesh stiffness, reduce operator splitting errors, and increase the allowable time step
- Point-centered variables, to minimize mesh imprinting due to flow misalignment
- High-order spatial methods, to minimize mesh imprinting and increase compute intensity
- AMR, to mitigate loss of accuracy at shocks
- Advanced hardware utilization models, to make effective use of emerging archtiectures

The notion of vorticity-compliant hydrodynamic methods is fairly recent. By vorticity-compliant we mean that the discrete equations satisfy a number of ancillary relations involving the curl and divergence operators. This feature mitigates the vorticity errors that were seen in the CHAD hydrodynamic algorithm and that are also present in current ASC hydrodynamic algorithms, none of which are vorticity-compliant. In Lagrange calculations these errors manifest as unphysical mesh motion such as hourglass and chevron modes. While a cell-centered scheme with vorticity-compliant properties was recently developed by Burton *et al* [25], we have been unable to locate in the relevant literature a point-centered scheme of the type we have derived.

Unsplit time integration, high-order time integration, and point-centered variables are common in CFD codes [26-29], but with few exceptions, have not been employed in Laboratory codes. Our work will therefore serve as a proving ground for these methods in Laboratory applications. As mentioned previously, unsplit time integration mitigates mesh stiffness but also reduces operator splitting errors [30]. High-order time integration allows for a time step that is closer to the Courant condition, which reduces the number of cycles and total runtime. Point-centering of variables reduces mesh imprinting errors by increasing the number of control volumes involved during advection. When that number is low, the results will be highly dependent on the alignment of those control volumes relative to the flow: poor alignment leads to errors referred to as mesh imprinting. When the number is high, the flow has more degrees of freedom and the negative effects of flow misalignment are mitigated. Point-centered schemes increase the number of control volumes involved because advection takes place between point-point pairs (*i.e.* along edges), as opposed to zone-zone pairs (faces) in a zone-centered scheme. On tetrahedral meshes, a zone has four faces whereas a point can easily have a dozen or more connected edges.

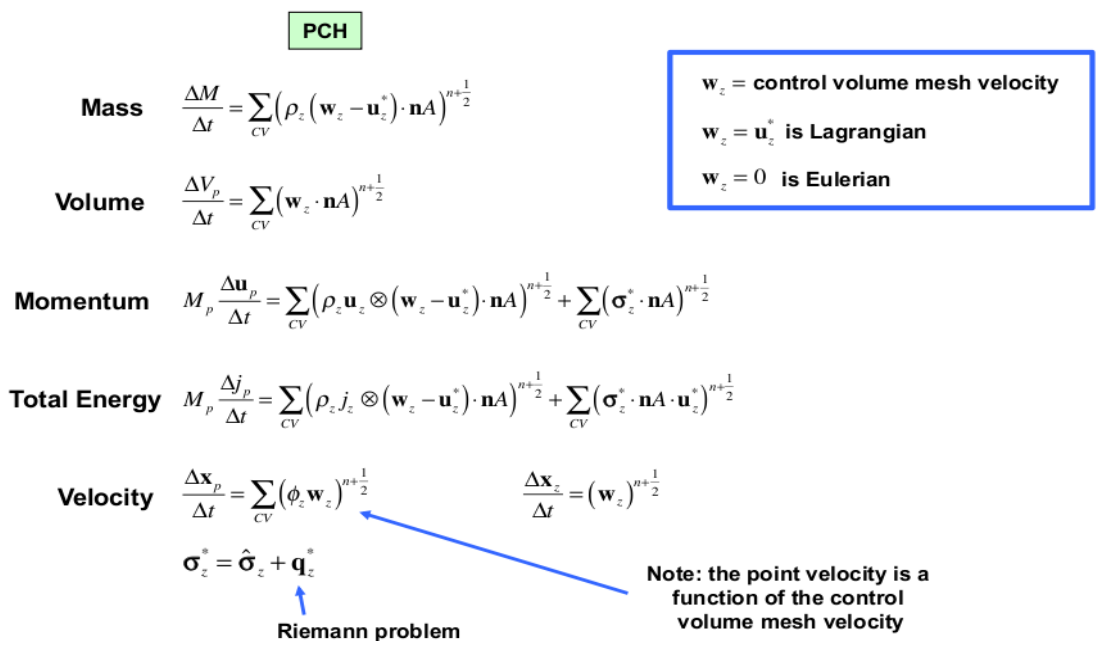

Figure 5: The discrete form of the ALE hydrodynamic equations in the point-centered, vorticity-compliant algorithm for 2nd-order unsplit time integration.

High-order spatial methods [31-34] are key to effective use of emerging architectures. Whereas 1st- and 2nd-order methods typically achieve speed-ups of 5-10 times on GPUs [24, 35], high-order methods achieve speed-ups of 25-50 times [36-37]. High-order methods allow simulation fidelity to be increased through enhancements in accuracy rather than resolution. This approach to error reduction is a departure from traditional Laboratory simulation methods, but one we believe is necessary given the expected memory characteristics of emerging architectures. An additional benefit of high-order methods is that they reduce mesh imprinting over traditional 1st- and 2nd-order methods [38]. We intend to employ reconstruction-based methods as opposed to higher-order FE methods; neither has been used in LANL hydrodynamic codes.

A reality of numerical hydrodynamics, however, is that all monotone methods are 1st-order accurate at shocks. Increased resolution is therefore still necessary near shocks to maintain solution

fidelity. AMR is an obvious means by which the extent of that resolution can be minimized. AMR is a defining capability of the ASC code RAGE [39] and has been used in LLNL multi-physics codes [40]. The use of high-order methods away from shocks and AMR near shocks also is well-established [41]. However, the effective use of AMR with emerging architectures remains somewhat of a research problem. Because AMR packages tend to be dominated by integer and memory operations, they generally are not amenable to GPU acceleration. We will develop a new AMR strategy that combines aspects of the standard patch-based and cell-by-cell approaches and also makes use of vectorized space-partitioning data structures [42].

Our hardware utilization paradigm is that a code should be flexible enough to allow whatever mode of execution leads to the fastest overall time to solution on a given compute platform. This again represents a departure from ASC codes, which have focused primarily on message passing. We will explore a variety of issues, including optimal memory access, threading across multiple GPUs, concurrent CPU-GPU execution, and hierarchical parallelism using message passing, threading, and GPUs. We expect the optimal combination of these features to be platform-dependent, and our utilization model will allow that combination to be selected at runtime.

Technical Impact

This work will have a number of technical impacts. First, as previously mentioned, is the development of a point-centered, vorticity-compliant hydrodynamics algorithm. Our research will evolve this nascent approach to numerical hydrodynamics both within the context of Laboratory applications as well as within the broader scientific community. Second is the validation of numerical techniques. Laboratory production codes are, out of necessity, based on trusted methods. By expanding the range of trusted methods our research will improve algorithmic diversity. Third is the development of effective strategies for emerging architectures such as GPUs. Our research will demonstrate a path forward for production codes on future compute platforms.

Mission Impact

Current Mission: We expect our research to lead to orders of magnitude increases in fidelity for capability-class problems, and orders of magnitude reductions in runtime for capacity-class problems. The impacts of these improvements will include faster responses to programmatic questions; increased population sizes for Uncertainty Quantification and other sensitivity studies; greater detail in discovery-scale simulations; and an enhanced ability to model realistic 3D features. Our work will also enable the use of commodity mesh generation software, which could lead to cost savings of several million dollars per year [43] and reduce problem setup time. Lastly, our work will allow future compute platforms to be used earlier and more effectively.

Future Mission: The jump in simulation capability that results from our research will enable the solution of entirely new classes of problems and therefore has the potential to significantly expand the scope of the Laboratory's simulation tools. New application areas might include design of blast mitigation structures for urban environments, energetic disablement calculations of Improvised Explosive Devices, anti-personnel and anti-structural analysis, and high-resolution studies of mix and ignition in Inertial Confinement Fusion targets.

R&D Methods and Anticipated Results

R&D Methods

The work associated with this proposal can be roughly broken down into five broad topical areas: ALE hydrodynamics, multi-material hydrodynamics, hardware utilization, AMR, and valida-

tion. Our overall R&D methodology will be test-driven, so that we can identify and correct deficiencies as early as possible. The specific test cases listed in what follows are generally considered sufficient to identify most potential issues in multi-material hydrodynamic codes. Likewise our proposed validation suite includes representative problems in the high-strain-rate regime. The complete problem set effectively constitutes the requirements for this project.

ALE hydrodynamics: We will initially implement and verify our proposed hydrodynamic algorithm for the single-material case. We have assembled a suite of hydrodynamic test problems that will be used to iteratively assess and improve the algorithm, starting in the Eulerian limit and then incorporating mesh motion. The problem suite, which tests basic fluid dynamics and shock physics capabilities, includes the following:

- Riemann problem [44]
- Sedov blast wave [45]
- Noh convergent shock [46]
- Gudderly convergent wave [44]
- Saltzmann piston [25]
- Taylor-Green vortex [44]
- Coggeshall compression [47]
- Adiabatic release [48]
- Shock diffraction over a step [49]
- Mach 3 flow over a step [50]

Definitions for several of these problems already exist for the FE hydrodynamic algorithm in CHICOMA. We will generate FE definitions for the remaining problems and then migrate the full set of problems to the ALE algorithm. Comparisons with the FE algorithm will be used to assess the new algorithm: if the ALE algorithm produces inferior results, we will consider abandoning the ALE approach and attempting to modify the FE algorithm for Laboratory problems. The consequences of this alternative outcome will be discussed subsequently.

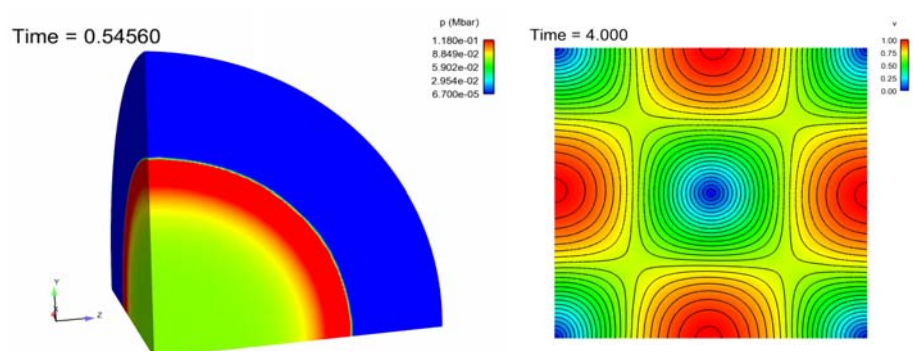

Figure 6: CHICOMA simulations of the Sedov (left) and Taylor-Green (right) test problems with the FE hydrodynamic algorithm.

Multi-material hydrodynamics: We will implement a multi-material treatment in the FE algorithm in parallel with the development of the ALE algorithm. If the ALE algorithm is successfully developed for the single-material case, the multi-material procedures developed in the FE algorithm will be migrated; otherwise we will move forward with the multi-material FE hydrodynamic algorithm. Our multi-material test suite includes multi-material variants of the Riemann problem along with the problems defined in [51]:

- Triple point problem
- Single-mode Rayleigh-Taylor instability
- Shock-bubble interaction

As with the previous set of test problems, input definitions will first be generated for the FE hydrodynamic algorithm, and then migrated to the ALE algorithm. Results will be compared to published results from other codes. The multi-material work also will include an interface to SESAME equations of state, implementation of the JWL equation of state, and programmed High Explosive (HE) burn, all of which are easily verified with simple manufactured problems.

Once a multi-material treatment is successfully implemented and verified, we will address material strength, in particular the Johnson-Cook model [50]. Our strength test problems will include:

- Verney spherical collapse [53]
- Howell cylindrical collapse [54]
- Elastic-plastic piston [25]

Hardware utilization: This topical area includes three phases. First, we will further explore the use of graph theory methods to increase compute intensity. In particular, memory access patterns will be examined with performance analysis tools to identify potential optimizations. Timing studies will be used to iteratively assess proposed improvements, as was done with the initial GPU version of CHICOMA. This will be done within the FE hydrodynamic algorithm, but the methods will be readily applicable to the ALE algorithm when it comes online. When the ALE algorithm does come online we will repeat this analysis to identify and correct its specific performance deficiencies. Performance will be compared against the FE algorithm.

Second, we will implement message passing and explore hierarchical parallelism. While a message passing implementation is straightforward, effective approaches to hierarchical parallelism are expected to require iteration. We have several large test problems that will be used to explore different strategies. We also expect the optimal strategy to be platform dependent.

In the final phase we will investigate advanced GPU usage. This includes threading across multiple GPUs as well as concurrent CPU-GPU execution modes. Again, we will examine code performance with performance analysis tools to identify potential optimizations, and iteratively assess proposed improvements with our suite of test problems.

AMR: The AMR package in CHICOMA [12] will first be examined to determine if any of its sub-components can make effective use of the GPU. If this is in fact the case, we will develop accelerated versions of those sub-components and assess what, if any, performance gain can be achieved. A number of test problems involving AMR are available to use for timing and verification studies, including variations of the standard problems listed previously as well as more complex geometries such as the simulation shown in Figure 4.

Once this phase is complete we will begin investigation of appropriate strategies for the use of the AMR package. The cell-by-cell AMR strategy currently in CHICOMA uses truncation error estimates to modify the mesh around specific features of interest, such as shocks and contact discontinuities, and does so every few time steps. In contrast a patched-based approach modifies much larger regions of the mesh but does so less often. We will first implement a patched-based approach, and, after baselining the performance of both strategies, we will explore hybrid patched/cell-by-cell strategies to iteratively optimize overall time to solution. Space-partitioning data structures will be used to construct localized patches from the cell-by-cell error estimates.

Validation: The ability to reproduce experimental results will ultimately determine the success or failure of our methodology. Our validation suite includes a number of high-strain-rate, multiple material problems, all of which have experimental data available for comparison:

- Taylor cylinder impact [55-56]
- Explosively formed projectile [57-58]
- HE cylinder tests [*e.g.* 59]
- Tungsten rod penetration of a steel target [60]
- HE acceleration of a metal shell [61]

Work on each problem will commence as soon as the necessary packages in CHICOMA have been implemented and verified, so that feedback can be incorporated as early as possible.

Expected Results

The expected result of this research is a validated 3D multi-material hydrodynamic algorithm for tetrahedral meshes with AMR that is at least 3rd-order accurate in both space and time. The algorithm will make use of emerging architectures and hierarchical parallelism. We expect a 300-fold performance improvement relative to major ASC codes on current compute platforms and an additional 30-fold performance improvement on GPU-based compute platforms.

If we are unable to successfully develop our proposed ALE algorithm, we will pursue extensions to the current FE algorithm in CHICOMA. This alternative path will allow many of the benefits of tetrahedral meshes to still be realized, but with a somewhat less accurate method. Compared to the CHAD algorithm, we expect that the improved FE algorithm would not suffer from the same time integration issues and would be less susceptible to mesh imprinting. The improved FE algorithm would be susceptible to spurious vorticity, although we note that the level of spurious vorticity in the current FE algorithm is comparable to or less than what is present in the major ASC codes [62]. The improvements in performance and ease of setup would still be realized.

Project plan

Our initial project plan showing the major tasks and milestones is shown in Figure 7. Also highlighted are items expected to result in journal submissions. The milestones are as follows:

1. 09/2013: Demonstrate single-material ALE simulations, including convergence analysis, accuracy measurements, comparison to experiment, and timing studies
2. 09/2014: Demonstrate hierarchical parallelism, GPU performance gains, and GPU-based AMR simulations, including timing studies
3. 09/2015: Demonstrate multi-material ALE simulations with strength, including convergence analysis, comparison to experiment, and timing studies

The interim milestones will be used to assess progress and adjust expected outcomes as needed.

Data Management Plan

This project will employ standard Software Configuration Management (SCM) practices to manage the CHICOMA source code and all associated data such as test problem definitions and documentation, including revision control, regression testing, build management, and defect tracking. Standard tools available on various LANL internal networks will be used for these tasks.

Transition Plan

Our initial discussions with the ASC Program Office indicate that if our research is successful, this project could possibly continue under ASC funding with the goal of developing a production simulation tool. If such funding is not available, we will pursue transfer of our technologies into

existing ASC projects. We intend to interact closely with the ASC Program Office throughout the LDRD and, as it nears completion, determine the most practical and cost-effective solution. We also will use our existing contacts in Global Security, Emergency Response, and Counter-terrorism programs to explore building a customer base within PADGS and/or partner agencies.

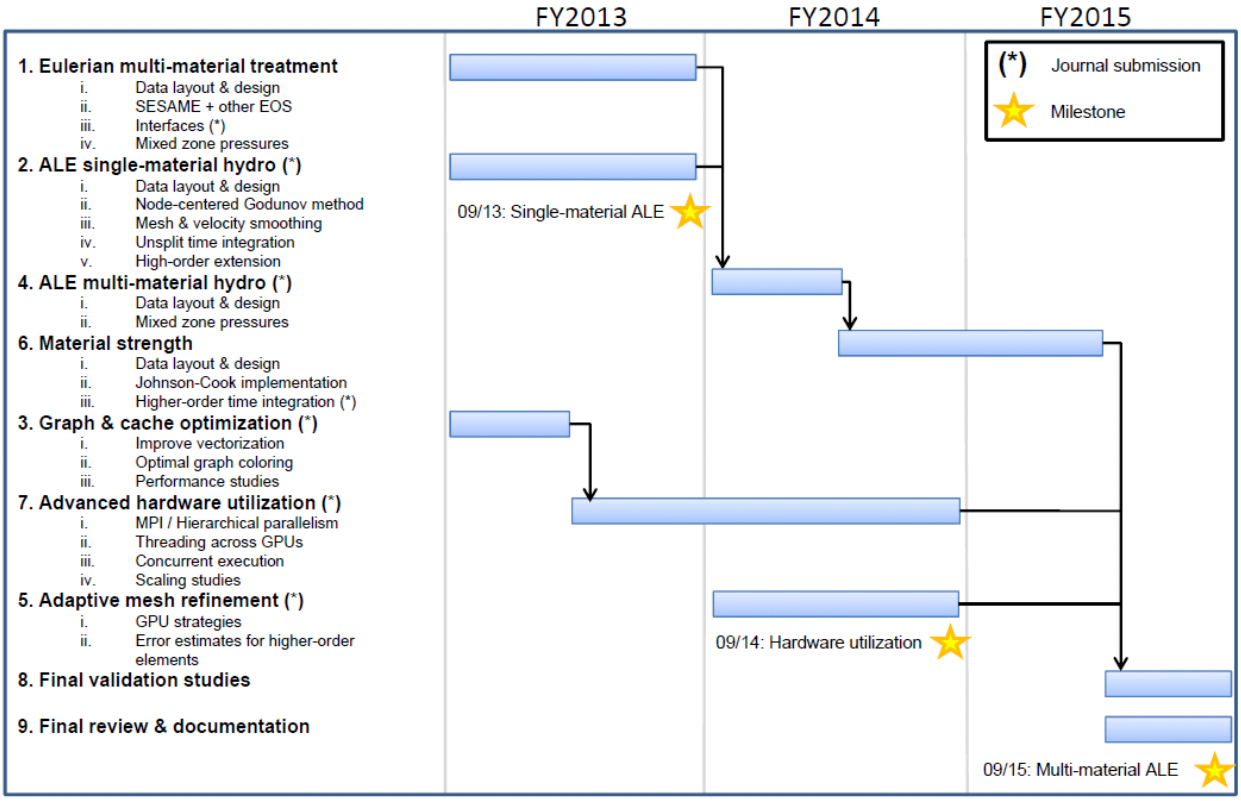

Figure 7: Initial project plan with major tasks and milestones

Budget Request

Our budget request corresponds to 3.5 FTEs of effort per year, for 3 years. This request derives from our team's experience in both the implementation of complex numerical algorithms and the management of scientific software development projects. We also have included travel budget of $25k per year, to include both domestic and international conferences, and a first-year expense of $50k for workstation purchases.

While funding the proposed research under ASC may seem natural, we believe doing so would actually increase risk. Thorough, rigorous testing is critical to establishing the validity of a new simulation methodology. However, the schedule pressures that exist within the programmatic environment can lead to premature deployment and compromised testing. This increases the likelihood that deficiencies will remain undetected until either the cost to fix them is unacceptably high or expectations are irreparably damaged. The aforementioned Free Lagrange codes suffered from such a fate. LDRD funding would shield the work from short-term programmatic pressures and thereby reduce the likelihood of this outcome.

Furthermore, the ASC program is strongly focused on near-term deliverables and improvements to existing simulation tools. While this emphasis may be appropriate, it creates an environment in which funding for high-risk/high-gain research could be difficult to justify on an ongoing basis, especially given current budget constraints. LDRD funding would eliminate this risk.

Citations

1. C. Christensen, *The Innovator's Dilemma*. Harper Business: New York (2000).
2. Sandia National Laboratories, "CUBIT 13.1 Online User's Manual", available from cubit.sandia.gov (2011).
3. F. Camelli, R. Löhner, J. Cebral and E.L. Mestreau, 'Timings of an unstructured-grid CFD code on common hardware platforms and compilers', AIAA Aerospace Sciences Meeting, AIAA-2008-0477 (2008).
4. J. Waltz, 'Spatial accuracy and performance of a mixed-order, explicit multi-stage method for unsteady flows', submitted to *Int. J. Num. Meth. Fluids* (2011).
5. R. Löhner, C. Yang, J. Cebral, O. Soto, F. Camelli, and J. Waltz, 'Improving the speed and accuracy of projection-type incompressible flow solvers', *Comp. Meth. Appl. Mech. Eng.* **195**, 3087 (2006).
6. F. Brezzi, K. Lipnikov, M. Shashkov, and V. Simoncini, 'A new discretization methodology for diffusion problems on generalized polyhedral meshes', *Comp. Meth. Appl. Mech. Eng.* **196**, 3682 (2007).
7. T.A. Wareing, J.M. McGhee, J.E. Morel, and S.D. Pautz, 'Discontinuous finite element S_N methods on three-dimensional unstructured grids', *Nuc. Sci. Eng.* **138**, 256 (2001).
8. J.S. Warsa and J.E. Morel, 'An S_N spatial discretization scheme for tetrahedral meshes', *Nuc. Sci. Eng.* **151**, 157 (2005).
9. G.D. Martin, E. Loth, and D. Lankford, 'Particle host cell determination in unstructured grids', *Comput. Fluids* **38**, 101 (2009).
10. T.B. Kaiser, 'Laser ray tracing and power deposition on an unstructured three-dimensional grid', *Phys. Rev. E* **61**, 895 (2000).
11. S. Nicaise, 'Edge elements on anisotropic meshes and approximation of the Maxwell equations', *SIAM J. Numer. Anal.* **39**, 784 (2001).
12. J. Waltz, 'Parallel adaptive refinement for unsteady flow calculations on 3D unstructured grids', *Int. J. Num. Meth. Fluids* **46**, 37 (2004).
13. R. Löhner and J.D. Baum, 'Adaptive *h*-refinement on 3D unstructured grids for transient problems', *Int. J. Num. Meth. Fluids* **14**, 1407 (1992).
14. J. Waltz, 'Microfluidics simulation using adaptive unstructured grids', *Int. J. Num. Meth. Fluids* **46**, 939 (2004).
15. R. Löhner, *Applied Computational Fluid Dynamics Techniques: An Introduction Based on Finite Element Methods (2nd ed).* Wiley & Sons: West Sussex (2008).
16. H.E. Trease, 'Three-dimensional Free Lagrangian hydrodynamics', LANL report no. LA-UR-85-2479 (1985).
17. D.A. Mandell and H.E. Trease, 'Parallel processing a real code – a case history', LANL report no. LA-UR-88-1836 (1988).
18. M.S. Sahota and H.E. Trease, 'A three-dimensional Free-Lagrange code for multimaterial flow simulations', LANL report no. LA-UR-90-3383 (1990).
19. P.J. O'Rourke and M.S. Sahota, 'A variable explicit/implicit numerical method for calculating advection on unstructured meshes', *J. Comput. Phys.* **143**, 312 (1998).
20. B.G. DeVolder, M.S. Sahota, and M.C. Cline, 'Verification of the CHAD hydrodynamics code', LANL report no. LA-UR-00-6039 (2000).
21. M.S. Sahota and M.C. Cline, 'Verification and validation of CHAD hydrocode', LANL report no. LA-UR-02-0868 (2002).
22. N. Morgan and J. Waltz, 'A new point-centered hydrodynamic scheme', LANL report no. LA-UR-12-01683 (2012).
23. D. Burton, personal communication (2012).

24. J. Waltz, 'Performance of a 3D unstructured mesh compressible flow solver on NVIDIA Fermi-class GPU hardware', submitted to *Int. J. Num. Meth. Fluids* (2012). Also LANL report no. LA-UR-12-01319 (2012).
25. D.E. Burton, T.C. Carney, N.R. Morgan, S.R. Runnels, S.K. Sambasivan, and M.J. Shashkov, 'A cell-centered Lagrangian Godunov-like method for solid dynamics', submitted to *Comput. Fluids* (2011).
26. R. Löhner and C. Yang, 'Improved ALE mesh velocities for moving bodies', *Comm. Num. Meth. Eng.* **4**, 599 (1996).
27. C. Farhat, P. Geuzine, and C. Grandmont, 'The discrete geometric conservation law and the nonlinear stability of ALE schemes for the solution of flow problems with unstructured grids', *J. Comput. Phys.* **174**, 669 (2001).
28. H. Luo, J.D. Baum, and R. Löhner, 'On the computation of multi-material flows using ALE formulation', *J. Comput. Phys.* **194**, 304 (2004).
29. D.S. Kershaw, M.K. Prasad, M.J. Shaw, and J.L. Milovich, '3D unstructured mesh ALE hydrodynamics with the upwind discontinuous finite element method', *Comput. Meth. Appl. Mech. Eng.* **158**, 81 (1998).
30. J. Waltz, 'Operator splitting and time accuracy in Lagrange plus remap solution methods', submitted to *J. Comput. Phys.* (2011).
31. V.A. Titarev and D. Drikakis, 'Uniformly high-order schemes on arbitrary unstructured meshes for advection-diffusion equations', *Comp. Fluids* **46**, 467 (2011).
32. P. Castonguay, P.E. Vincent, and A. Jameson, 'A new class of high-order energy stable flux reconstruction schemes for conservation laws on triangular grids', *J. Sci. Comp.* DOI: 10.1007/s10915-011-9490-6 (2011)
33. V-T. Nguyen, 'An arbitrary Lagrangian-Eulerian discontinuous Galerkin method for simulations of flows over variable geometries', *J. Fluids Struct.***26**, 312 (2010)
34. D.J. Mavriplis and C.R. Nastase, 'On the geometric conservation law for high-order discontinuous Galerkin discretizations on dynamically deforming meshes', *J. Comput. Phys.* **230**, 4285 (2011).
35. A. Corrigan, F.F. Camelli, R. Löhner, and J. Wallin, 'Running unstructured grid-based CFD solvers on modern graphics hardware', *Int. J. Num. Meth. Fluids* **66**, 221 (2011).
36. A. Klockner, T. Warburton, J. Bridge, and J.S. Hesthaven, 'Nodal Discontinuous Galerkin methods on graphics processors', *J. Comutp. Phys* **228**, 7863 (2009)..
37. P. Castonguay, D.M. Williams, P.E. Vincent, M. Lopez, and A. Jameson, 'On the development of a high-order multi-GPU enabled, compressible viscous flow solver for mixed unstructured grids', AIAA Computational Fluid Dynamics Conference, AIAA-2011-3229 (2011).
38. V.A. Dobrev, T.E. Ellis, T.V. Kolev, and R.N. Rieben, 'Curvilinear finite elements for Lagrangian hydrodynamics', *Int. J. Num. Meth. Fluids* **65**, 1295 (2011).
39. M. Gittings, R.Weaver, M. Clover, T. Betlach, N. Byrne, R. Coker, E. Dendy, R. Hueckstaedt, K. New, W.R. Oakes, D. Ranta, and R. Stefan, 'The RAGE radiation-hydrodynamic code', *Comput. Sci. Discov.* **1**, 63 (2008).
40. A.E. Kongies, R.W. Anderson, P. Wang, B.T.N. Gunney, R. Becker, D.C. Eder, B.J. BacGowan, and M.B. Schneider, 'Modeling NIF experimental designs with adaptive mesh refinement and Lagrangian hydrodynamics', *J. Phys. IV* **133**, 587 (2006).
41. N. Burgess and D.J. Mavriplis, 'An *hp*-adaptive Discontinuous Galerkin solver for aerodynamic flows on mixed-element meshes', AIAA Aerospace Sciences Meeting, AIAA-2011-0490 (2011).
42. J. Waltz, G.L. Page, S.D. Milder, J. Wallin, and A. Antunes, 'A performance comparison of tree data structures for N-body simulation', *J. Comput. Phys.* **178**, 1 (2002).
43. R. Webster, personal communication (2011).
44. J.R. Kamm, J.S. Brock, S.T. Brandon, D.L. Cotrell, B. Johnson, P. Knupp, W.J. Rider, T.G. Trucano, and V.G. Weirs, 'Enhanced verification test suite for physics simulation codes', LANL report no. LA-14379 (2008).

45. L.I. Sedov, *Similarity and Dimensional Methods in Mechanics.* Academic Press: New York (1959).
46. W.F. Noh, 'Errors for calculations of strong shocks using an artificial viscosity and an artificial heat-flux', *J. Comput. Phys.* **72**, 78 (1987).
47. S.V Coggeshall, 'Analytic solutions of hydrodynamics equations', *Phys. Fluids A* **3**, 757 (1991).
48. G. Bazan and R.N. Rieben, 'Adiabatic release test problem description', LLNL report no. LLNL-PRES-463883 (2012).
49. R. Hivier, 'Computation of shock wave diffraction at a ninety degrees convex edge', *Shock Waves* **1**, 89 (1991).
50. P. Woodward and P. Collela, 'The numerical simulation of two-dimensional fluid flow with strong shocks', *J. Comput. Phys.* **54**, 115 (1984).
51. R. Loubere, P-H. Maire, M. Shashkov, J. Breil, and S. Galera, 'ReALE: A reconnection-based arbitrary-Lagrangian-Eulerian method', *J. Comput. Phys.* **229**, 4724 (2010).
52. G.R. Johnson and W.H. Cook, 'A constitutive model and data for metals subjected to large strains, high strain rates, and high temperatures', Proc. 7th International Symposium on Ballistics (1983).
53. D. Verney, 'Évaluation de la limite élastique du cuivreet del' uranium par des expériences d'implosion', in *Behavior of Dense Media under High Dynamic Pressures.* Gordon and Breach: New York (1968).
54. B.P Howell and G.J. Ball, 'A Free-Lagrange augmented Godunov method for the simulation of elastic-plastic solids', *J. Comput. Phys.* **175**, 128 (2002).
55. G.I. Taylor, 'The use of flat-ended projectiles for determining dynamic yield stress, theoretical considerations', *Proc. R. Soc. London*, Series A, **194**, 1038 (1948).
56. M.L Wilkins and M.W Guinan, 'Impact of cylinders on a rigid boundary', *J. Appl. Phys.* **44**, 1200 (1973).
57. P.J. Maudlin, J.F. Bingert, S.R. Bingert, S-R Chen, and G.T. Gray III, 'Final report on the Tantalum K26 test bed series experimentation, characterization, and modeling simulation', LANL report no. LA-CP-10-01748 (2010). *[Export Controlled]*.
58. J. Waltz, 'Simulation of an explosively formed projectile with FLAG', LANL report no. LA-CP-00613 (2011). *[Export Controlled]*.
59. M.L. Rightley, 'Cylinder test calculations with the ASCI Shavano project code', LANL report no. LA-UR-01-6827 (2001).
60. C.E. Anderson, V. Hohler, J.D. Walker, and A.J. Stilp, 'Time-resolved penetration of long rods into steel targets', *Int. J. Impact Eng.* **16**, 1 (1995).
61. H.U. Freund, W. Geiger, and G. Honcia, 'Acceleration of spherical metal shells by high explosives: detonation velocity along the shell surface and efficiency of energy transfer', Proc. 6th International Symposium on Detonation (1976).
62. T. Masser, 'RAGE results for Sedov problem', LANL report no. LA-UR-11-04043 (2011).

Appendix

Dr. Jacob Waltz (PI, XCP-8) received a Ph.D. in Computational Sciences from George Mason University, specializing in Computational Fluid Dynamics, and joined LANL in 2002. He has over 15 years experience in the development of large-scale simulation codes. His experience includes Eulerian, Lagrangian, and ALE hydrodynamics methods; unstructured grid methods; turbulent and multi-fluid mix; AMR; parallel and high performance computing; multi-physics coupling; and code validation. He is a former Project Leader of the ASC Shavano / LAP Project and is the principal author of CHICOMA. Jacob is intimately familiar with the requirements and V&V of Laboratory multi-material hydrodynamic codes, and has strong relationships with application physicists. Most recently his work has involved multi-physics code validation and modeling in support of various Global Security programs, along with a small amount of research in the areas of Eulerian unstructured grid codes on GPU hardware and time integration for Lagrange-plus-remap methods. As Principal Investigator and chief code architect he will be involved in all aspects of the project. He will be supported at the 75% level.

Dr. Thomas Canfield (T-3) has over 30 years experience developing efficient numerical algorithms for scientific and engineering applications, including fluid and solid dynamics and structural analysis, on a broad range of computer architectures. Tom is a staff member in the T-3 Fluid Dynamics and Solid Mechanics Group, where he has played a key strategic role in the implementation and V&V of material models in ASC codes. He has made critical contributions to a number of high-level ASC milestones over the last decade, and has worked on both the ASC Shavano / LAP project and the now-defunct ASC Antero Project. Tom also has held positions in industry, academia, and at both Argonne and Sandia National Laboratories. His education includes a PhD in Applied Mechanics from the University of Illinois at Chicago. He is expected to play a lead role in all aspects of multi-material hydrodynamics, including theory, design, implementation, and V&V. He will be supported at the 50% level.

Dr. Nathaniel Morgan (XCP-8) has been a staff member at LANL since 2005, with several years of experience in weapon physics, weapon analysis, and validation. He is the former Primary Point of Contact for the W76 SLBM warhead, and continues to be active in several Global Security programs. For the last two years he has worked on the development of Lagrangian Cell-Centered Hydrodynamic methods under an LDRD-DR project, and has made key contributions in the development of vorticity-compliant operators, error analysis of current hydrodynamic methods, and physics-based artificial viscosity schemes. He is expected to play a lead role in hydrodynamic algorithm development and V&V. Nathaniel received a Ph.D. in Mechanical Engineering from Georgia Tech University and will be supported at the 75% level.

Larry (Dean) Risinger (CCS-7) received an M.S. in Mathematics from New Mexico State University and joined LANL as a staff member in 2000. His experience includes scientific programming, cyber-security, high-performance computing, and fault-tolerant message passing. He is currently working on an ASC Co-Design mini-app based on CHICOMA to explore implementation challenges on modern architectures. He also supports scientific programming for a Global Security project. Dean is expected to play a lead role in all aspects of hardware utilization and performance optimization and will be supported at the 50% level.

Dr. Shiv Sambasivan (CCS-2) received a Ph.D. in Mechanical Engineering from the University of Iowa and has been a post-doctoral research associate at LANL since 2010. He has 10 years of experience modeling high-speed multi-material compressible and incompressible flows, includ-

ing unstructured grid methods, AMR, level set methods, and material strength. He is currently the lead code developer for Lagrangian hydrodynamics methods in the Hydra multi-physics toolkit. He also has worked on the development of Cell-Centered Hydrodynamic methods under an LDRD-DR project, with key contributions related to the treatment of material strength during ALE. Shiv is expected to participate extensively in multi-material hydrodynamics and AMR activities, and will be supported at the 50% level.

Dr. John Wolhbier (CCS-2) received a Ph.D. in Electrical and Computer Engineering from the University of Wisconsin. As a staff member in CCS-2 he focuses on bridging the gap between applications and computer science. He has done considerable development of continuum dynamics plasma physics models in the RAGE code under the ASC Crestone / EAP Project. He also has invested a large amount of effort in programming a multi-material hydrodynamics algorithm for Roadrunner. Presently he is leading a team that is using OpenCL within the RAGE code to implement key algorithms for execution on modern architectures. The work on RAGE together with the work on Roadrunner has made him acutely aware of the challenges that will be faced for multi-physics application codes on modern architectures. John has strong relationships with both application physicists and computer science researchers, and has demonstrated the ability to successfully combine the latest ideas in computer science into applications code. John will be responsible for the investigation of approaches for algorithm implementations, including the comparison of OpenACC and OpenCL paradigms. He also is expected to contribute to AMR development and V&V. John will be supported at the 50% level.