Multi-Material ALE with AMR for Modeling Hot Plasmas and Cold Fragmenting Materials

Alice KONIGES[1], Nathan MASTERS[2], Aaron FISHER[2], David EDER[2],
Wangyi LIU[1], Robert ANDERSON[2], David BENSON[3], Andrea BERTOZZI[4]

[1]Lawrence Berkeley National Laboratory, 1 Cyclotron Rd. Berkeley, CA 94720, USA
[2]Lawrence Livermore National Laboratory, P. O. Box 808, Livermore, CA 94550, USA
[3]Jacobs School of Engineering, 9500 Gilman Dr., UCSD, La Jolla, CA 92093, USA
[4]Department of Mathematics, 520 Portola Plaza, UCLA, Los Angeles, CA 90095, USA

Abstract We have developed a new 3D multi-physics multi-material code, ALE-AMR, which combines Arbitrary Lagrangian Eulerian (ALE) hydrodynamics with Adaptive Mesh Refinement (AMR) to connect the continuum to the microstructural regimes. The code is unique in its ability to model hot radiating plasmas and cold fragmenting solids. New numerical techniques were developed for many of the physics packages to work efficiently on a dynamically moving and adapting mesh. We use interface reconstruction based on volume fractions of the material components within mixed zones and reconstruct interfaces as needed. This interface reconstruction model is also used for void coalescence and fragmentation. A flexible strength/failure framework allows for pluggable material models, which may require material history arrays to determine the level of accumulated damage or the evolving yield stress in J2 plasticity models. For some applications laser rays are propagating through a virtual composite mesh consisting of the finest resolution representation of the modeled space. A new 2nd order accurate diffusion solver has been implemented for the thermal conduction and radiation transport packages. One application area is the modeling of laser/target effects including debris/shrapnel generation. Other application areas include warm dense matter, EUV lithography, and material wall interactions for fusion devices.

Keywords: multi-physics simulations, arbitrary lagrangian eulerian hydrodynamics, adaptive mesh refinement, fragmentation, laser ray tracing, NIF

PACS: *02.70.−c, 46.15.−x, 52.50.Jm, 52.65.−y*

DOI: 10.1088/1009-0630/17/2/05

(Some figures may appear in colour only in the online journal)

1 Introduction

Applications in high energy density physics, such as modeling the laser interaction with matter, are traditionally modeled with an ALE (Arbitrary Lagrangian Eulerian) methodology, due to Lagrangian formulation's ability to deal with large displacements and the mesh-untangling capabilities of the ALE method. There is a variety of inertial confinement fusion (ICF) codes such as CHIC, Lasnex, Hydra, and FCI2 [1−4]. In a traditional ALE method, the pure Lagrangian time-advance phase is followed by a remap step (that may or may not occur at every Lagrangian time step). In the presence of vorticity, which can be physical or anomalous, or other small-scale variations in velocity, remesh and remap are necessary to maintain a viable computational solution. In other words, the remesh/remap stage serves to correct meshing errors, e.g., *bowties*, that can arise from both unphysical and physical effects in the material motion. A poor mesh can cause prohibitively small time steps and hamper the code's ability to progress forward in time. This has become the standard method used to model laser-plasma-solid interactions in the fluid regime.

Our work on combining ALE and Adaptive Mesh Refinement (AMR) stemmed from the need to use traditional ALE methods—because of physics and modeling constraints—but, the critical factor in our simulations is the need for much larger model domains than those used in traditionally small target simulations. In a traditional small target simulation, the computational domain is restricted to a small region surrounding the volume heated by the laser or ion source. Our need for modeling certain phenomena, such as debris and shrapnel formation and damage on laser facilities, as well as general target dynamics of laser and accelerator fields, motived our requirement to model very large domains. Additionally, the lack of symmetry in the physical effects following initial laser- or ion beam-target interaction requires a three-dimensional capability. Thus a simple uniformly or geometrically zoned Lagrangian mesh with relaxation cannot encompass the entire simulation domain of interest. Success with AMR treatments of flow situations such as those pioneered by Colella [5] motivated us to consider a combined ALE with AMR method. The methodology

encapsulated in our ALE-AMR Framework, builds off the previous work in incorporating ALE and AMR for gas-dynamics [6] using the Structured Adaptive Mesh Refinement Application Interface (SAMRAI) [7]. Initial work involved adapting typical AMR solution techniques, based on cell-centered discretizations, to staggered mesh discretizations using a nodal description of element position and velocity, but cell-centered descriptions of density and thermodynamic quantities [6].

Significant additions to this early methodology and framework have resulted in the development of a new fully featured 3D multi-physics code that contains sophisticated material modeling capabilities in addition to the usual radiation/hydrodynamics packages often contained in ALE modeling codes meant for the simulation of laser-driven fusion energy targets. The addition of multi-material and the requisite interface reconstruction capabilities provides the flexibility to relax the need to force material boundaries to correspond to zone boundaries. This capability is very important in modeling very complex targets such as those used at the National Ignition Facility (NIF), which contains the world's most energetic laser. The AMR capability is also used dynamically to provide higher resolution in regions with steep gradients during simulations. The implementation of an anisotropic stress tensor and support for numerous strength/failure models provides the flexibility to model a wide range of materials using anything from analytical flow-stress models to detailed single-crystal plasticity models [3]. Laser rays are propagated through a virtual composite mesh consisting of the finest resolution representation of the modeled space with inverse bremsstrahlung absorption. The energy deposited by the laser can create hot plasmas with subsequent energy transfer occurring by thermal conduction or radiation. A new 2nd order accurate diffusion solver has been implemented for the thermal conduction and radiation transport packages.

In the sections that follow, we give an introduction to Lagrangian dynamics and the basic ALE+AMR method in the context of the laser-driven and related simulations. We then describe in some detail how the introduction of stress and strain terms is incorporated into the basic formulation. One rather unique feature of our code method compared with traditional ICF codes is the ability to store material history data to facilitate use of engineering-based failure models. We describe our approach to multi-material issues using volume fractions. Our heat conduction, radiation transport, and laser/ion packages are described followed by a brief discussion of our work on surface tension models and our approach to visualization. We conclude with verification, validation, and sample simulations that show the capabilities of the code.

2 Lagrangian dynamics

The code is based on a staggered grid Lagrangian formulation with position and velocity being nodal variables and density, internal energy, temperature, pressure, strain, and stress being zonal (cell centered) variables. The thermal conduction and radiation transport equations are solved using the diffusion approximation with temperature and radiation energy being additional nodal variables. The plasma/fluid equations in a Lagrangian formulation (in vector and indicial notation $i, j, k = 1, 2, 3$) are:

$$\frac{\mathrm{D}\rho}{\mathrm{D}t} = -\rho\nabla\cdot\vec{U} = -\rho U_{i,i} \ , \tag{1}$$

$$\rho\frac{\mathrm{D}\vec{U}}{\mathrm{D}t} = \frac{1}{\rho}\nabla\cdot\boldsymbol{\sigma} = \frac{1}{\rho}\sigma_{ij,j} \ , \tag{2}$$

$$\frac{\mathrm{D}e}{\mathrm{D}t} = \frac{1}{\rho}V\boldsymbol{s}:\dot{\boldsymbol{\varepsilon}} - P\dot{V} = \frac{1}{\rho}V\left(s_{ij}\dot{\varepsilon}_{ij}\right) - P\dot{V} \ , \tag{3}$$

where

$$\frac{\mathrm{D}}{\mathrm{D}t} = \frac{\partial}{\partial t} + \vec{U}\cdot\nabla \ ,$$

is the substantial derivative, ρ is the density, $\vec{U} = (u, v, w)$ is the material velocity, t is time, $\boldsymbol{\sigma}$ is the total stress tensor, P is the pressure, e is the internal energy, V is the relative volume ($\rho V = \rho_0$ where ρ_0 is the reference density), $\boldsymbol{s}$ is the deviatoric stress, and $\dot{\boldsymbol{\epsilon}}$ is the strain rate tensor defined as

$$\dot{\epsilon}_{ij} = \frac{1}{2}\left(\frac{\partial U_i}{\partial x_j} + \frac{\partial U_j}{\partial x_i}\right) \ . \tag{4}$$

ALE-AMR uses a modified HEMP discretization [8–10] in evaluating the Lagrange update of the field variables and the strain rate tensor is evaluated using the numerical method proposed by Flanagan and Belytschko [11].

The deviatoric stresses are defined as

$$s_{ij} = \sigma_{ij} + P\delta_{ij} \ , \tag{5}$$

where δ is the Kronecker delta and P is pressure. The pressure is determined by the equation of state (EOS), which returns pressure as a function of density and either energy or temperature. A simple example of an EOS for solid material is a linear EOS where the hydrostatic pressure is computed from the material's bulk modulus K and the ratio of actual density to the reference density. ALE-AMR provides access to a variety of analytic and tabular EOS, as well as user defined EOS through Python functions.

The stress deviators are usually determined by constitutive relations. A simple example of a constitutive relation is Hooke's Law for isotropic elastic materials, which gives a linear relationship between deviatoric stress and deviatoric strain as

$$s_{ij} = 2\mu\epsilon_{ij} \ , \tag{6}$$

where μ is the shear modulus.

Unlike Refs. [1-4], we do not use separate electron and ion temperatures, because in our problems of interest (NIF debris and shrapnel, for example) equilibration is permitted. A possible extension of the code would be to include such effects with a two-temperature model.

An ALE method with an explicit time-marching Lagrange step is used to integrate the system. Following a traditional ALE approach, in the first phase, a Lagrange step advances the flow-field through a physical

time step. The optional second phase involves a modification of the grid and a remapping (interpolation) of the solution to the new grid. If the new grid is identical to the original grid prior to the Lagrange step, calculation resembles a fully Eulerian method, and the term Eulerian is often used to describe the mesh. Often in real ALE calculations the solution is mapped to a "relaxed" grid that has been smoothed in some manner. This procedure is used to alleviate the mesh tangling problems inherent in the Lagrangian methods. For the mesh relaxation algorithm we use a modified Laplace iteration, that has its origins in the work of Winslow [12]. The subsequent interpolation procedure is formulated as an advection problem, and is discussed in detail in Anderson [6]. When ALE is mixed with AMR, both the ALE and the AMR schemes can work independently to give zonal modification and potentially improve accuracy and/or runtime to solution. In practice, these two methods can also work against each other, whereby special ALE zoning techniques can be wiped out by an AMR scheme that is pushing the calculation towards refined zones. Some of the techniques we apply in applications of the code include limiting the refinement to less than full three dimension or preserving mesh feathering in the refinement scheme.

3 ALE with AMR methodology

In this section, we briefly review a basic method of combining ALE with AMR that was proven successful for gas dynamics [6]. Other ALE + AMR schemes have been proposed [13,14]. Our choice here is based on using the automatic parallelism of the underlying AMR library as well as the ability to retain compatibility with ALE schemes using a similar structured grid and Lagrangian integration method. For example, because the grid mapping is straight forward, partial solutions obtained using advanced physics models available in other structured ALE-only codes can be input directly into ALE-AMR simulations. Our calculations are computed on a structured mesh (quadrilaterals in two dimensions and hexahedra in three dimensions). The Lagrange step is a time-centered predictor-corrector discretization, and a staggered spatial discretization [15,16,8]. The scheme preferentially employs a monotonic artificial viscosity, q, due to Christensen [17], and a kinematic hourglass filter [18]. The Lagrange scheme conserves mass and momentum, and energy in 1D. However, in multidimensions, it does not conserve energy exactly. This scheme was chosen by Anderson et al., for advantages in the AMR remap step [6]. A Lagrange predictor step calculates the acceleration of each node, which is computed from a discretized control volume. Nodal acceleration is computed from $\vec{F} = \tilde{m}\vec{a}$, where $\tilde{m}$ is a nodal mass, and $\vec{F}$ is a sum of net force contributions (pressure, tractions, and artificial viscosity q acting on the faces of the control volume associated with the node in the HEMP discretization).

Nodal positions and velocities are advanced, and followed by an update of pressure, stresses and energy in each zone. A corrector step then repeats this process using time-centered values. A standard remap step is applied optionally after the Lagrange step. In an ALE method this allows one to 'arbitrarily' remap the grid to either a smoothed grid, or even back to the original grid rendering the calculation as functionally Eulerian.

We have available an equipotential method for defining smoothed grids based on early works [19,20]. Depending on the physics being modeled, it is advantageous to run problems with different combinations of ALE, e.g., fully Lagrangian with AMR refinement—later relaxing to ALE+AMR—or utilizing the full remapped Eulerian treatment. One real advantage of our code is that depending on the problem of interest, the user can intervene and pick a method appropriate for the detailed solution. Some problems are best with Lagranian alone for accuracy, some require full Eulerian to avoid tangling, and quite often we find that we need to add the AMR option to obtain sufficient zoning detail for the calculation to be effective and efficient. The AMR machinery also gives us a means to *shape-in* grids that are more appropriate for the problem than either a typical Lagrangian-optimized code mesh or a fully Eulerian AMR approach. Often at the beginning of the calculation, we already put in extra zoning in the areas of interest to the problem rather than just having the AMR package do refinement in area of changing conditions. Having the low resolution grid far from the simulation portion of interest effectively gives us an *air mesh* around our entire problem. This method avoids some of the traditional tangling problems associated with fully Lagrangian calculations. In section 12, we discuss some of the trade-offs for running the code in various modes and how the solution can vary with the Langrangian, ALE, or ALE with AMR modes.

The remap step consists of mapping the solution from the original Lagrange grid to the relaxed grid or the Eulerian grid. A convenient way to handle this remapping problem is formulating it as an advection equation, and often this step is simply referred to as advection. A staggered grid advection scheme based on Pember's work [21] is implemented. During the advection step ρ, volume fractions, and stress deviators (and associated strength quantities) are remapped on the original mesh as volume weighted quantities, followed by the momentum–computed and advected on a node centered dual-mesh. The total energy, E—formed from the cell-centered internal energy e and an average of the kinetic energy of the surrounding nodes—is then advected (cell-centered). The updated internal energy is then constructed from the remapped E and the updated kinetic energy obtained from the momentum update. In order to minimize some of the non-physical behavior that can be introduced by staggered mesh advection schemes under certain circumstances, we have implemented a Repair step, as proposed by Margolin et al [22].

In the AMR methodology a hierarchical grid structure is employed which changes dynamically in time and is composed of logically rectangular, grid "patches" of varying resolution. The collection of grid patches at a given resolution is referred to as a level, notated as L_θ, where θ is a level number counting from 0 at the

coarsest level. A grid hierarchy is constructed so at any refined level (L_θ, $\theta > 1$) there will be r^d fine grid cells precisely covering the coarse grid cell at the next coarser level ($L_{\theta-1}$), where r is called the refinement ratio and d is the problem dimensionality (1-, 2-, or 3-D). The solution is defined and updated on all cells, including coarse cells which underlay cells of finer resolution. However, during subsequent operations the solution on refined zones will be "coarsened down" to replace the solution on the underlying coarse cells.

Key to the ALE-AMR method is the construction of interlevel solution transfer operators. Properties of the algorithm are defined by choosing interlevel operators that maintain various properties to be consistent with combining the ALE and the AMR techniques. The choices of these operators are described in detail by Anderson [6], and we cite here the relevant implications for our ALE-AMR code. The main point is to restrict refinement to odd refinement ratios, because this preserves inversion and an exclusive $r : 1$ correspondence between fine and coarse node data locations. This allows for a conservative distribution of a quantity from the coarse mesh to its corresponding fine mesh stencil and the associated inversion with a simple summation. Solution refinement begins with linear mesh interpolation in each coordinate direction. The solution scalars are then interpolated onto the finer mesh using a series of linear interpolations utilizing either volume or mass coordinates in each coordinate direction.

Interpolation is used to map the coarse mesh onto the finer mesh, either to fill ghost zones or when a region is identified as needing refinement, using a constant refinement ratio r, but all constructions trivially extend to a vector of (odd-valued) ratios to allow for different ratios in each coordinate direction. The mapping conditions provided in Anderson, et al. [6], preserve constant fields properly, and monotonicity and oscillations are further addressed by a van Leer Limiter technique. The sequence of interpolations is to first interpolate density using cell volumes, compute the resulting fine mesh cell and nodal masses, and then use those bases to compute interpolated energies and velocities. An intermediate step is used during the coarsening process to allow the operator to be both constant field preserving and conservative. This intermediate step constructs a temporary remapped solution on a fine mesh but this mesh is not used in the time advance.

4 SAMRAI: structured adaptive mesh refinement application infrastructure

The ALE-AMR code currently uses SAMRAI [7], a structured grid adaptive mesh library, for the AMR parts of the algorithm. The SAMRAI framework is a C++ library, and the application code was developed using both C++ and FORTRAN 90, with FORTRAN 90 being reserved for performance of critical inner loop constructs. One of the primary benefits to the ALE-AMR code is that writing it on top of a parallel mesh refinement library hides a good majority of the parallel implementation from the physics development, and most of the physics code needs only be concerned with updating the solution on individual grid patches. Additionally, as changes are made to improve the scalability of the library, as long as the interface method to the library remains the same, these optimizations are automatically attained. Studies of the scalability of ALE-AMR with the SAMRAI library are given by Koniges, et al [23,24]. Encapsulation of the basic parallelism of the code to be within a library disjoint from the physics is critical to advancing the code to the next level of scalability as modern architectures advance. Very simply put, it allows us to routinely integrate new programming paradigms that are developed for the advanced architectures without adjusting our code that resides on top of the parallel library.

5 Stress tensor implementation

Addition of the full stress tensor terms for elastic-plastic flow to the basic Lagrangian equations in 1D, 2D, and 3D dimensions follows a straight-forward implementation of the equations given by Wilkins [8]. The text spells out the finite difference terms in the equations of motion. (We note that there a couple of minor typographical errors in the equations.) Wilkins also gives a series of test problems that we have used to benchmark the implementation. The ALE-AMR code does not include sliding interfaces. This inclusion would add significant complexity to the AMR method and is not generally required for our problems of interest.

6 Failure models

It is important that modeling tools allow for insertion of different physics models, as appropriate, for a given physical process. Fig. 1 shows a diagrammatic example of the plug-and-play features that allow a user to insert an appropriate model for both material strength and material failure in ALE-AMR. In a similar way, users can use different equation-of-state models, etc [25].

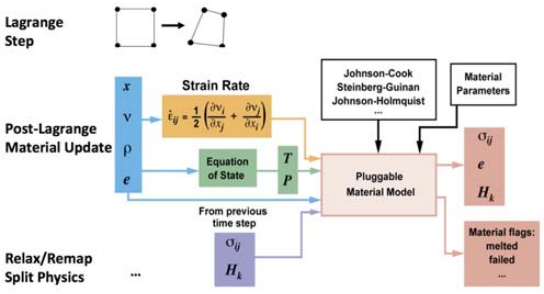

Fig.1 Diagrammatic framework of ALE-AMR code hydrodynamics, illustrating modularity and flexibility

In order to include material failure, the material strength model implementation is similar to using a standard equations of state except that the strength

models need to track some number of history variables. We use a C++ singleton design pattern to create a map between the region numbers and the material models applied to those regions. When a material update is required, the singleton runs through all of the cells, and calls the correct material models for each cell. Since different strength models have different numbers of history variables, the cells (and in the case of mixed material zones, the sub-cells) must be allowed to have different depths with regard to material history storage. The use of the material failure model enables us to denote a failed cell or zone, but does not permit the formation of fragments. We use a method called void insertion along with a material reconstruction algorithm (discussed in the next section) to create fragments in the code [26]. This implementation and structure is also crucial in implementing the surface tension algorithms.

7 Multi-material interface reconstruction

In almost all ALE-AMR simulations there is a need to have multiple materials in a zone. This can occur at the start of a simulation because it is often not practical to have all material boundaries for a complex target coincide with zone boundaries. In addition, unless the hydrodynamics are run only Lagrangian (very rarely done), the ALE formulation requires that material interfaces be tracked. ALE-AMR uses a Volume of Fluid (VOF) method and the CALE93 (Tipton) algorithm [27] for advection. During a simulation the volume fraction of each material in a mixed zone is stored and the actual interface is only reconstructed as needed, e.g., during mesh refinement. Volume fractions in mixed zones are used for weighting the pressures, densities, stresses, etc. for the different materials to obtain composite quantities. As discussed above, a special void material is introduced during failure and has an associated volume fraction as well. Fig. 2 shows how the code models fragment formation beginning from void coalesce.

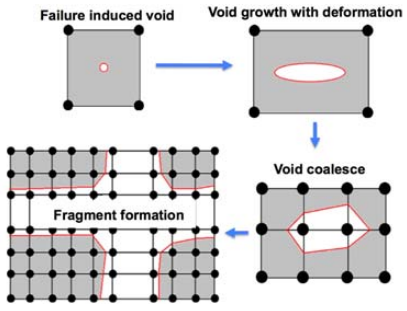

Fig.2 Fragment formation. When a material in a zone reaches a failure criteria, a small volume of void is inserted. The void volume will grow during deformation and using the multi-material framework, void volumes in different zones will be connected, resulting in void coalesce. Finally continued deformation results in fragment formation

During refinement we use volume fractions of neighboring zones to determine the orientation of each interface and the actual location determined by the volume fraction in the zone. Calculating the volume on either side of an interface in a 3D zone requires some effort. The hexahedral zones are bounded on six sides by doubly-ruled surfaces. When such a zone is truncated by a plane (see Fig. 3), we use the divergence theorem to transform the volume integral to a surface integral [28]. Making use of parametric coordinates as shown in the figure, it is possible to use sets of analytic expressions to solve the surface integral. Our approach is similar to that proposed by Kothe, et al [29].

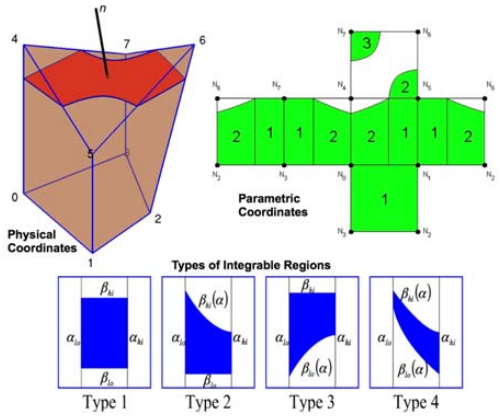

Fig.3 A truncated hexahedron associated with an interface reconstruction based on volume fractions (upper-left) with nodes shown in physical coordinates. The same nodes are given in parametric coordinates in the upper-right image. The intersection of the truncating plane with the six faces is also given in this image in addition to the corresponding type of integrable region. Of the potential 4 types shown, only 3 types are present for this cut

The CALE93 algorithm [27], provides a means to determine which materials in a multi-material zone are advected through a given volume flux. The following is a brief summary of the method. A normalized slope is calculated at each face for each (donor) element (see Fig. 4). This is done for each material in the donor element. The approach considers three potential flow patterns: Series flow: Materials are moved in order in which they are stacked up relative to the face; Parallel Flow: The components are moved simultaneously; Corner Flow: Treated as series flow until a critical volume fraction is achieved, at which point it becomes parallel flow. Since material can leave from multiple faces, the code checks that the total amount of material does not exceed the material current in the donor element.

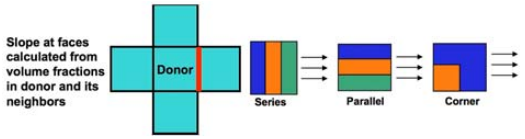

Fig.4 Flow pattern. During advection with multi-material the volume fractions in the donor zone and the neighbor zones are used to determine the flow pattern

Currently we implement a split advection scheme, where each logical direction is treated separately (the order rotated to avoid imprinting) to provide transport across zone corners without introducing inconsistencies between mixed and clean zone advection if single stage, e.g., Corner Transport Upwind [5], advection formulations are used.

8 Conduction and radiation

The radiation transport and heat conduction models in ALE-AMR are based on the diffusion approximation. The diffusion solver uses the Finite Element Method (FEM) which requires the data to be on a single level composite mesh. Details of the approach are given in a companion paper [30] and we only summarize the approach here. For each SAMRAI patch all zones at the finest level are assigned unique global ID that is used to create the composite mesh. A complex family of finite elements are required to treat all possible combinations of coarse-fine boundaries. The diffusion equation is discretized using the standard Galerkin approach using a set of quadrature rules to approximate integrals and construct matrices. The resulting linear system is solved using the HYPRE library [31], which shows 2nd order convergence. Heat conduction and radiation transport equations are time evolved implicitly. After the nodal energies and temperatures are evolved the zone based material temperatures and energies are updated by a mapping method that limits artificial diffusion.

9 Laser ray trace and ion package

Laser energy deposition and beam evolution uses the geometrical-optics approximation to calculate the trajectory of laser rays with deposition by inverse-bremsstrahlung [32]. In this approximation, ray propagation is completely determined by the electron-density gradient, which is approximated as linear within each zone resulting in parabolic ray trajectories, which must be tracked to identify exit points. Furthermore, as the electron density field is not continuous at zone interfaces, rays crossing into a neighboring zone experience an effective force normal to the zone interface, which discontinuously changes the component of their velocity in that direction (Snell's law). In 3D, quartic equations are solved to determine the exit face and point on the hexahedral zones. When there are multiple solutions the exit face with the shortest exit time is used. We use a virtual composite mesh consisting of the finest representation of the physical space using the AMR hierarchy [33]. The composite mesh is never explicitly constructed. When a ray enters a zone, it is determined if a finer representation exists, i.e., if the ray has crossed coarse-fine interface, and, if so, the fine destination index is found. Energy deposited by all rays is coarsened down to underlying patches at the end of the ray tracing step.

We process rays in batches that cascade to neighboring patches/processors, which evens the load on processors that are participating in ray tracing. However, there can be load balancing issues when there are large number of patches/processors with very few, if any, rays. Various approaches to improve load balancing are being evaluated. For many problems of interest, the duration of the laser or ion beams pulses is short compared to the total simulation time. The laser package can also be used to model ion beams, in which case, these beams are not refracted by density gradients (or zone interfaces) and therefore have linear trajectories. The energy of the ions are deposited along the trajectory using standard stopping formulas [34].

10 Surface tension models

One approach to modeling surface tension is through energy minimization where the surface energy is the surface tension coefficient times the surface area. Long-wavelength perturbations of a cylinder can decrease surface area while preserving volume, which is sometimes referred to as the Rayleigh instability. Nonlinear effects cause a given wavelength disturbance to subdivide into main and satellite droplets. If there is spatial variation of the surface tension coefficient Marangoni stress can be important. In addition to energy minimization, other approaches to model surface tension include both sharp and diffuse interface models. The interface is calculated explicitly for the case of sharp interface models, while the interface is inferred though the use of a phase field variable, e.g., concentration or density, in diffuse interface models. Volume fractions can also be used to calculate curvatures of material interfaces and subsequent surface tension forces. One application area where surface tension can play a role is EUV lithography droplet dynamics. We are exploring various approaches with the ALE-AMR code.

11 Visualization

The primary visualization tool used is VisIt, which can handle very large data sets, and is available for free under a BSD license [35]. VisIt was developed by the Department of Energy over the last decade with a number of features added to handle AMR data structures. VisIt can create material interface boundaries using volume fractions, however the algorithm in VisIt is close, but not identical to the algorithm in our code. Since the state of the materials is shown, this is not a problem in interpretation. VisIt also supports C++, Python and Java interfaces, and a variety of data formats including silo, hdf5, and VTK. ALE-AMR takes advantage of the ability to specify user defined functions operating on field variables to eliminate the need to store derived quantities (e.g., momentum), which may be used in post-processing, e.g., using VisIt to identify connected components (or "fragments") and computing the mass, and average velocity of these via the intermediary derived momentum.

12 Verification, validation, and sample simulations

In this section we present simulations exercising many of the features of the ALE-AMR code. We emphasize that our code was not developed to create a more accurate hydrodynamic modeling framework for shock dynamics. Rather it was developed as a means to simulate problems in failure and fragmentation in the high energy density regime that were simply not possible with traditional ALE codes because of the computational limitations.

We start with a traditional accuracy study based on hydrodynamics, and we then show more typical advanced simulations that we have used the code for in real experimental modeling. While our code was primarily developed to include advanced features such as a fragmentation model and a surface tension model, we can still consider some of the traditional hydrodynamics problems to illustrate various modes at which the code can be run (Lagrangian, Eulerian, and AMR) and to notice some of the features that distinguish the methods of running the code.

12.1 Sedov blast simulations

The Sedov Blast problem was independently solved by Sedov [37], Taylor [38], and von Neumann [39] and consists of a finite amount of energy released at the initial time at the origin of the domain that is filled with a polytropic gas. The problem is self-similar and Kamm [40] presents a concise algorithm for obtaining solutions for 1D (planar), 2D (Cylindrical), and 3D (Spherical) blasts with constant and power-law initial density profiles. In this section, the code found in the report by Timmes et al. [36] is used to evaluate these solutions for comparsion to the results of 2D ALE-AMR simulations of the cylindrical and spherical Sedov blasts run in Lagrangian, Arbitrary Lagrangian-Eulerian, and Eulerian modes with different levels of adaptive mesh refinement.

The initial blast energies, $E_{\rm blast}$, suggested by Kamm are used: $E_{\rm blast} = 0.311357$ and 0.851072 ergs for the cylindrical and spherical cases, respectively, with these values so chosen that for a uniform initial density of $\rho_0 = 1$ g/cm^3 and a $\gamma = 1.4$ polytropic gas the shock will be at $r = 0.75$ cm and 1.0 cm at the time $t = 1$ s.

The intent of these comparisons is to explore the impact of adaptive mesh refinement on the solution rather than a convergence study. As such, a relatively coarse mesh of 135×135 zones is used as the finest representation of the 1.35 cm×1.35 cm domain. This provides an initial mesh resolution of 0.01 cm, consistent with some of the results presented by Kamm. AMR will be used here such that each successive finer level refines the underlying zones by a ratio of 3 in each logical direction, ending with the finest level. Therefore for the current mesh, the limiting case is four levels of refinement with a coarsest mesh of 5×5 zones.

The blast energy is uniformly deposited as a specific energy into the four zones adjacent to the origin, which is the source of some of the deviation from the analytic solutions, particularly closer to the origin where density approaches zero and the specific internal energy is therefore unbounded. Pseudocolor plots of these simulations with one and four levels of refinement are presented in Figs. 5 and 6 for the cylindrical and spherical cases, respectively.

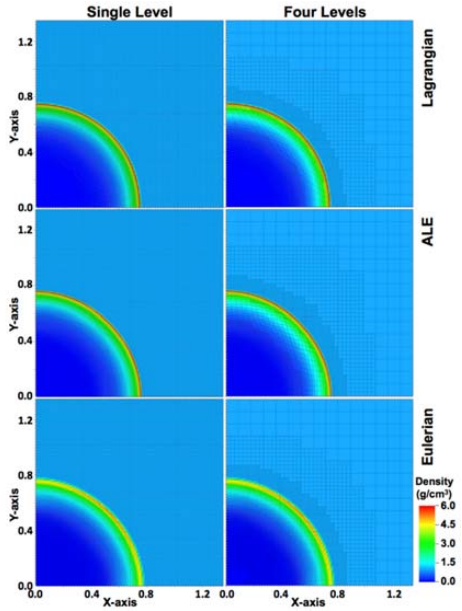

Fig.5 Simulation results for the cylindrical blast problem with one and four levels of refinement

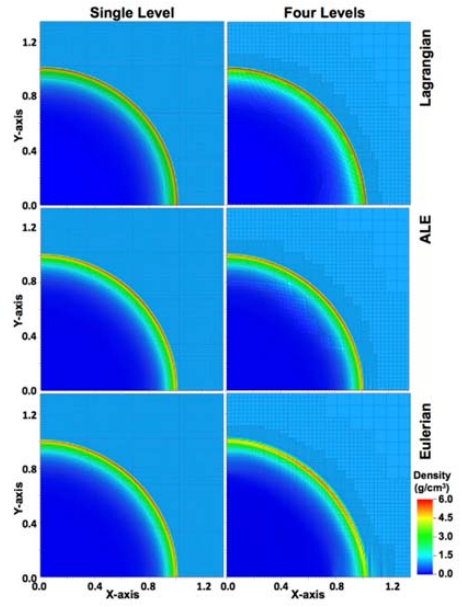

Fig.6 Simulation results for the spherical blast problem with one and four levels of refinement

Qualitatively, the shocks in all of these simulations have reached approximately the correct location at the specified time with reasonable symmetry. The blast has

traversed a large portion of the domain and a significant fraction of the mesh is refined at this point such that the coarsest level is entirely covered by the next finer level of refinement. In both the cylindrical and spherical simulations, the Lagrangian and ALE simulations have coarsened portions of the mesh behind the shock front, whereas the Eulerian simulation, remains fully refined behind the shock. The refinement leading the shock front can also be observed in the undisturbed regions of the mesh with coarser mesh farther from the shock front.

Since moving meshes are involved in the Lagrangian and ALE simulations whereas AMR has the potential to remove portions of the finest mesh from the solution, we extracted lineouts at 45° to the X-axis (Z-axis) to compare with the reference solutions. The reference solution is evaluated over a uniform 1.2 cm domain with a uniform 960 zone mesh and linearly interpolated as needed to compare with the Lagrangian and ALE simulations. Figs. 7-9 present lineouts from these simulations for zonal densities, material speeds, specific internal energy, and pressure. The observed errors are commensurate with those reported by Kamm for 1D RAGE simulations [40]. We observe that the Lagrangian and ALE simulations propagate the shock slightly faster than the analytic case and the Eulerian simulations are somewhat slower. However, as the current mesh is quite coarse, these discrepancies may be due to inadequate meshing. In addition to the reasonably good agreement with the analytic solution for all of the quantities near the shock, we observe that AMR does little to perturb the solution near the shock front, but does introduce some small anomalies, particularly in the post-shock region.

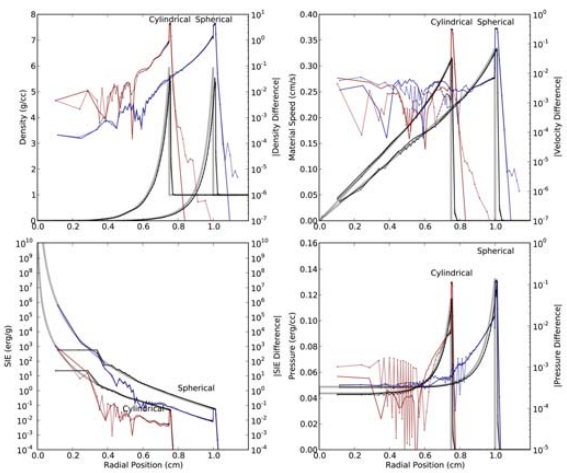

Fig.7 Lineouts from Lagrangian ALE-AMR solution to cylindrical and spherical Sedov blast problems. Gray line: Analytic solutions [36], Black lines: lineout from ALE-AMR simulations (at 45°), Red and Blue lines: absolute value of difference between ALE-AMR and analytic solutions for cylindrical and spherical blasts, respectively. Solid lines are for uniform mesh refinement, dashed lines represent four level AMR mesh

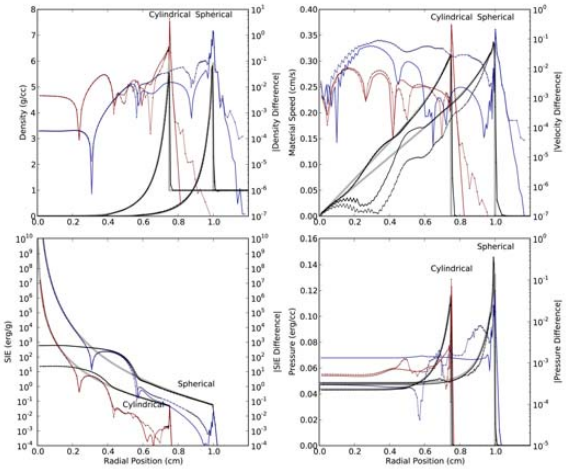

Fig.8 Lineouts from ALE-AMR solution (ALE Mode) to cylindrical and spherical Sedov blast problems. Gray line: Analytic solutions [36], Black lines: lineout from ALE-AMR simulations (at 45°), Red and Blue lines: absolute value of difference between ALE-AMR and analytic solutions for cylindrical and spherical blasts, respectively. Solid lines are for uniform mesh refinement, dashed lines represent four level AMR mesh

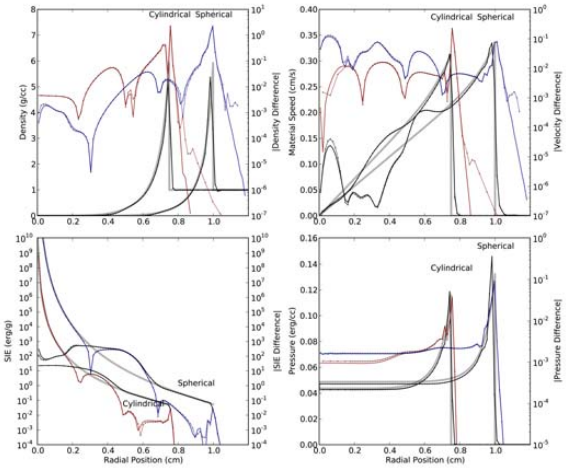

Fig.9 Lineouts from ALE-AMR solution (Eulerian Mode) to cylindrical and spherical Sedov blast problems. Gray line: Analytic solutions [36], Black lines: lineout from ALE-AMR simulations (at 45°), Red and Blue lines: absolute value of difference between ALE-AMR and analytic solutions for cylindrical and spherical blasts, respectively. Solid lines are for uniform mesh refinement, dashed lines represent four level AMR mesh

12.2 Droplet dynamics with AMR

The use of AMR can significantly reduce the computational cost of simulations. However, it is important to verify that the results are essentially the same with and without AMR. As an example, we show results of a recent simulation of EUV lithography droplet dynamics. The goal is to produce an efficient source of 13 nm emission for etching fine features in Si chips. One approach

is to heat and ionize Sn droplets with a laser and focus the resulting line emission with a multilayer mirror [41]. To improve efficiency a pre-pulse can be used to flatten the droplet and increase its size in the direction perpendicular to the laser axis. This is followed by the main pulse that ionizes the Sn that produces EUV emission. In our simulations with ALE-AMR we focus on the fluid dynamics associated with the pre-pulse. We show results of a simulation where an initial spherical droplet of Sn is heated by a CO_2 laser. As material is ablated by the laser, the remaining molten material is accelerated in the direction opposite the heating laser. The left image in Fig. 10 shows density variations within the droplet and shape changes 300 ns after the start of a $\approx$ 100 ns pulse from a simulation with single level of refinement. The initial spherical droplet has started to flatten and there are significant variations in density within the distorted droplet. The laser is aligned with bottom (Z) axis coming from the left and the droplet is accelerated to the right with radial expansion as well. This RZ (cylindrical) simulation includes the effects of surface tension. In the center image we show the results of using AMR with three levels of refinement that reduced the run time by a factor of 3. The finest level is the same as what was used for the image on the left. Comparing the two figures show that we achieve the same shape and density variations within the droplet with a significant reduction in run time. In the right image we show the AMR mesh and one can see that the regions of finest refinement are located along the edges of the droplet to resolve rapidly changing density.

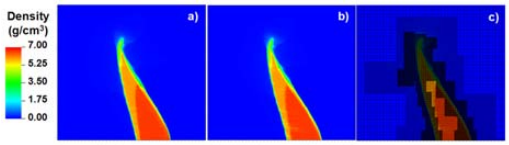

Fig.10 Droplet of molten Sn flattened by laser with uniform fine mesh (a), with an AMR mesh (b), and we show the 3 levels of refinement (c)

12.3 Validation and experimental guiding

The failure/fragmentation model has been validated against experiments such as an electromagnetically driven expanding ring where the measured number of Al fragments as a function of maximum radial velocity was compared to ALE-AMR simulation results (see Fig. 1 in Ref. [42]). The ALE-AMR code has also been used to guide experiments such as the one shown in Fig. 11 where the code prediction of fragment velocity of 750 m/s was used to set the gate times for the probe beam [25]. The in-flight fragments are seen in the upper right image. Fragments were also caught in aerogel foam as shown in the lower right image.

12.4 Need for air mesh and benefit of AMR

The development of our code was based on a recognition that traditional ALE codes were not able to model

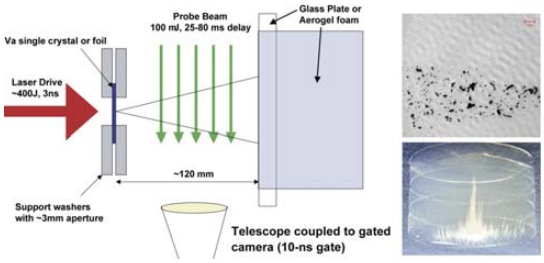

Fig.11 Laser-driven fragmentation experiments at the Janus laser at Lawrence Livermore National Laboratory (LLNL) using a probe beam and captured by glass plate or Aerogel foam

the full physics needed to study the fragmentation and failure models required for design of optics and diagnostics of laser facilities such as NIF. In order to study the very late time dynamics of laser or ion beam heated targets one generally has to add mesh that surrounds the target. This surrounding mesh is often referred to as an air mesh. This methodology is critical to move simulations to late time and effectively allows the code to proceed without too much ALE mesh tangling. It is often the case that the resolution needed away from the target is much less than what is required for the target components. Using standard ICF codes meshing procedures it is not possible to have appropriate resolution in the regions of interest for the fragmentation mechanisms while retaining a three-dimensional problem that is tractable on current computation platforms. This is seen in the left image of Fig. 12 where high resolution mesh lines extend into the mesh surrounding the target. In fact, this figure from a traditional ALE hydro code mesh does not even show the full mesh that was required for the simulation. If the full mesh was shown, the figure would just be a black image due to all of the mesh lines. Modeling experiments with laser ablation using Lagrangian hydrodynamics provides an accurate method to track the ablation front. The need to model such experiments out to late time using an air mesh was one of the driving reasons to develop the multi-physics ALE-AMR code. With AMR one can have the appropriate resolution in all parts of the simulation. An example of an ALE-AMR simulation with an air mesh is shown in the right image of Fig. 12.

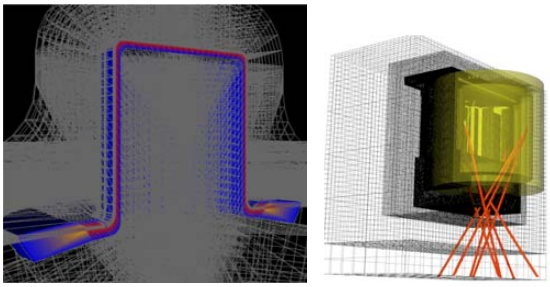

Fig.12 Left image: Example of air-mesh zoning of an ICF halfraum using a traditional ALE only Hydra code. Right image: Lasers entering an ICF halfraum with ALE-AMR

12.5 Simulations of debris and shrapnel generation

The ALE-AMR code was used as the primary tool for predicting and mitigating damage to NIF optics and diagnostics from debris and shrapnel during the National Ignition Campaign (2008-2012). These studies were crucial in the design of the target systems, the necessary features of debris shields, and even the materials used for the targets and ancillary structures such as cooling rings. Several examples of NIF ICF targets and the resulting predictions from ALE-AMR simulations were given in a recent paper by Eder, et al [42]. Here we discuss one ICF target simulation and one example of using high energy X-rays from a NIF target for exposure of samples to show the capabilities of the code and how it can be used. The predictions of fragment sizes and velocities have been validated with observations on a wide range of experiments. Current applications of the code include ion accelerator targets for warm dense matter experiments [43], EUV lithography target modeling and material wall interactions for fusion devices [44]. We note that the code itself is a framework for a variety of models and provides automatic parallelization. For example, Lui, et al. used the code as a framework for testing various surface tension models [45].

When simulations of laser or ion experiments include the entire target, as compared to localized regions near the energy deposition, extreme variations in densities and temperatures can exist from hot radiating plasmas to cold fragmenting solids [28]. The spatial and temporal scales for these full-target simulations can also be very large. The multi-physics code ALE-AMR was developed for such challenging simulations where we have used up to six levels of refinement, which roughly corresponds to a volume ratio of approximately 10^7 to 1 for the largest to smallest zones. The majority of ICF targets are cryogenically cooled with a capsule located inside a high-Z (Au or U) hohlraum liner having a wall thickness of $\sim$30 μm. The hohlraum liner is held by an outer Al structure having an average wall thickness of $\sim$300 μm [46]. Surrounding the Al casing there are two 500 μm thick Si cooling rings with slots and connecting tabs. There are 16 tabs on each ring. To model this target we often use a 22.5° wedge and model the upper half (1/32 of the target) as shown in the left image of Fig. 13. We show results of a simulation for a very low energy (14 kJ) experiment on the right image of Fig. 13 approximately 210 ns after the laser pulse. Of particular interest for potential damage to optics and diagnostics is the Al casing for which we use Johnson-Cook strength and failure models for Al 6061-T6. Damage in an element is accumulated until it reaches a value of 1.0, at which point failure occurs. Optionally failure may also occur if the principle stress exceeds a specified spall strength. The simulations was extended out to 750 ns where approximately 1/3 of the total target mass is vaporized and the molten/solid fragment velocities are ballistic.

In another experiment on NIF a sample holder was designed to hold 8 samples for radiation evaluation experiments. To block low-energy X-rays from reaching the samples, filters were placed in front of the samples. During initial fielding of this holder all 8 samples were damaged and no data were obtained. (See lower center image in Fig. 14 showing the holder following the shot.) The ALE-AMR team was asked to determine the reason for the damage and redesign the sample holder. The total X-ray fluence striking the filters is approximately 6 J/cm^2. In addition there is a debris wind from the vaporized target that strikes the filters later that is traveling at approximately 30 km/s. On the left side of Fig. 14 we give results of an ALE-AMR simulation with only X-ray loading that shows that the filter is launched with sufficient velocity to damage the sample. Simulations were used to determine the optimum angle of tilt to direct the filter away from the sample as shown in the upper right image in Fig. 14. The filter is moving at approximately 300 m/s and has not moved a significant distance when the debris wind arrives. Therefore, despite the X-ray load, the filter still blocks the majority of the target debris, protecting the sample. The redesign of the holder is shown in the lower right image with slots taken out of the cylinders holding the

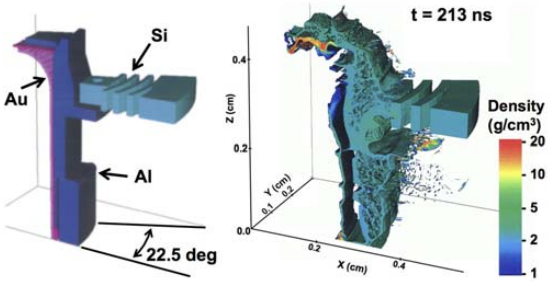

Fig.13 Three dimensional simulation of a portion of a ICF target consisting of an Au hohlraum inside an Al casing with Si cooling rings attached. The target material having a density greater than 1 g/cm^3 is shown 213 ns after being heated by 14 kJ of laser energy

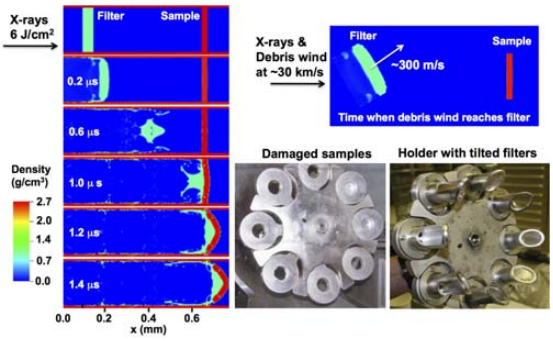

Fig.14 Simulation on left showing that X-ray loading on a filter can cause it to be launched towards sample, which explained reason for damaged samples shown in the center-bottom image. The simulation on the top right showing a tilted filter that does not damage sample, yet still shields sample from debris wind. Simulations were used to redesign the sample holder shown in the lower-right image

samples to allow the filter material to exit. All subsequent experiments using this holder produced valuable data with no samples being damaged [47]. ALE-AMR was also used to redesign targets and diagnostic components for numerous other experiments on NIF to minimize impacts associated with debris and shrapnel.

13 Conclusions

A new 3D multi-physics multi-material code, ALE-AMR, has been developed that can model hot radiating plasmas and cold fragmenting solids. Multi-material dynamics and interface reconstructions are determined using volume fractions. The multi-material mechanism is also used for void coalescence and fragmentation. In general, the failure models use history variables that are tracked for each material. A virtual composite mesh consisting of the finest resolution representation of the mesh is used for laser ray tracing with inverse-bremsstrahlung energy deposition. The same package is used for ion beam trajectories and deposition. Thermal conduction and radiation transport packages use a new 2nd order accurate diffusion solver developed for ALE codes using AMR. The ALE-AMR can be used to model a wide variety of problems including debris/shrapnel generation at large laser facilities, EUV lithography droplet dynamics, warm-dense-matter experiments using ion beams and material wall interactions for fusion devices.

Acknowledgments

We acknowledge the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science, U. S. Department of Energy under Contract No. DE-AC02-05CH11231. Work by LBNL under DE-AC02-05CH11231 was supported by the Director, Office of Science of the U. S. Department of Energy and the Petascale Initiative in Computational Science and Engineering. Work by LLNL was performed under the auspices of the U. S. Department of Energy by Lawrence Livermore National Security, LLC, Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. UCLA and LLNL acknowledge UC Lab Fees Research Grant 09-LR-04-116741-BERA. Work by UCLA also funded by NSF grant DMS-1312543.

References

1 Breil J, Galera S, Maire P-H. 2011, Computers and Fluids, 46: 161
2 Zimmerman G B, Kruer W L, et al. 1975, Comments on Plasma Physics and Controlled Fusion, 2: 51
3 Marinak M M, Haan S W, Dittrich T R, et al. Physics of Plasmas, 5: 1125
4 Holstein P A, Badonneau D, Bowen C, et al. 2004, Nuclear Fusion, 44: S177
5 Colella P. 1990, Journal of Computational Physics, 87: 171
6 Anderson R W, Elliott N S, Pember R B. 2004, Journal of Computational Physics, 199: 598
7 Hornung R and Kohn S. 2002, Concurrency: Practice and Experience, 14: 347
8 Wilkins M L. 1999, Computer Simulation of Dynamic Phenomena. Springer, Verlag
9 Wilkins M L. 1980, Journal of Computational Physics, 36: 281
10 Pember R B, Anderson R W. 2000, A Comparison of Staggered-Mesh Lagrange Plus Remap and Cell- Centered Direct Eulerian Godunov Schemes for Eulerian Shock Hydrodynamics. UCRL-JC-139820, Lawrence Livermore National Laboratory
11 Flanagan D P, Belytschko T. 1981, Int. J. Numer. Methods Eng., 17: 679
12 Winslow A. 1967, J. Comput. Phys., 1: 149
13 Morrell J M, Sweby P K, Barlow A. 2007, Journal of Computational Physics, 226: 1152
14 Morrell J M. 2011, Computers & Fluids, 46: 375
15 Wilkins M L. 1964, Meth. Comp. Phys., 3: 211
16 Sharp R W. 1978, HEMP Advection Model. UCID-17809, Lawrence Livermore National Laboratory
17 Christensen R. 1990, Godunov methods on a staggered mesh an improved artificial viscosity. UCRL-JC-105269, Lawrence Livermore National Laboratory
18 Margolin L, Pyun J. 1987, A method for treating hourglass patterns. LA-UR-87-439, Los Alamos National Laboratory
19 Winslow A. 1963, Equipotential zoning of two dimensional meshes. UCRL-7312, Lawrence Livermore National Laboratory
20 Tipton R. 1992, Grid optimization by equipotential relaxation. Lawrence Livermore National Laboratory
21 Pember R B, Anderson R W. 2001, Comparison of direct Eulerian Godunov and Lagrange plus remap, artificial viscosity schemes. Proc. 15th American Institute of Aeronautics and Astronautics Computational Fluid Dynamics Conference, Anaheim, CA, AIAA 2001-2644: 1
22 Margolin L G, Shashkov M. 2004, Computer Methods in Applied Mechanics and Engineering, 193: 4139
23 Koniges A E, Gunney B, Anderson R, et al. 2008, Advances in Parallel Computing, 15: 697
24 Koniges A, Preissl R, Kim J, et al. 2009, Application acceleration on current and future Cray platforms Proceedings. Cray User Group Meeting, Edinburgh, UK
25 Koniges A E, Masters N D, Fisher A C, et al. 2010, Journal of Physics: Conference Series, 244: 032019
26 Fisher A, Masters N, Dixit P, et al. 2008, Journal of Physics: Conference Series, 112: 022027
27 Jun B-II. 2011, A Simple Advection Scheme for Material Interface. UCRL-JC-139912, Lawrence Livermore National Laboratory
28 Koniges A, Debonnel C, Andrew J, et al. 2008, Journal of Physics: Conference Series, 112: 032072
29 Kothe D B, Williams M W, Lam K L, et al. 1999, A second-order accurate, linearity-preserving volume tracking algorithm for free surface flows on 3-D unstructured meshes. Proceedings of 3rd ASME/JSME Joint Fluids Engineering Conference (San Francisco, CA: ASME), FEDSM99-7109: 1
30 Fisher A C, Bailey D S, Kaiser T B, et al. 2015, Plasma Science and Technology, 17: 109
31 Falgout R D, Yang U M. 2002, Lecture Notes in Computer Science, 2331: 632

32 Kaiser T B. 2000, Physical Review E, 61: 895

33 Masters N D, Kaiser T B, Anderson R W, et al. 2010, Journal of Physics: Conference Series, 244: 032022

34 Atzeni S, Meyer-Ter-Vehn J. 2004, The Physics of Inertial Fusion: Beam Plasma Interaction, Hydrodynamics, Hot Dense Matter. Oxford University Press

35 Childs H, Brugger E, Whitlock B, et al. 2012, High Performance Visualization Enabling Extreme-Scale Scientific Insight. Editor E W Bethel and H Childs and C Hansen. Chapman & Hall, CRC Computational Science, p.357

36 Timmes F X, Gisler G, Hrbek G M. 2005, Automated Analyses of the Tri-Lab Verification Test Suite on Uniform and Adaptive Grids for Code Project A. Tech. Rep. LA-UR-05-6865, Los Alamos National Laboratory, p.116

37 Sedov L I. 1959, Similarity and Dimensional Methods in Mechanics. Academic Press, New York

38 Taylor G I. 1950, Proc. Roy. Soc. London A, 201: 159

39 Bethe Hans A, von Neumann J, Magee J L, et al. 1947, Blast Wave. Tech. Rep. LA-2000, Los Alamos National Laboratory, p.304

40 Kamm J R. 2000, Evaluation of the Sedov–von Neumann–Taylor Blast Wave Solution. Tech. Rep. LA-UR-00-6055, Los Alamos National Laboratory, p.46

41 Fomenkov I V, LaFontaine B, Brown D, et al. 2012, J. Micro/Nanolith. MEMS MOEMS, 11: 021110

42 Eder D C, Fisher A C, Koniges A E, et al. 2013, Nucl. Fusion, 53: 113037

43 Koniges A E, Liu W, Barnard J, et al. 2013, EPJ Web of Conferences, 59: 09006

44 Brooks J N, Hassanein A, Koniges A, et al. 2014, Contributions to Plasma Physics, 54: 329

45 Liu W, Bertozzi A, Kolokolnikov T. 2012, Comm. Math. Sci., 10: 387

46 Atherton L J. 2008, Journal of Physics: Conference Series, 112: 032063

47 Fournier K B, Celeste J, Rekow V, et al. 2010, Review of Scientific Instruments, 81: 075113

(Manuscript received 15 June 2014)
(Manuscript accepted 21 November 2014)
E-mail address of A. KONIGES: aekoniges@lbl.gov