

Three-dimensional Hybrid Continuum-Atomistic Simulations For Multiscale Hydrodynamics

H. S. Wijesinghe

Massachusetts Institute of Technology,
Cambridge, MA

R. D. Hornung

Lawrence Livermore National Laboratory,
Livermore, CA

A. L. Garcia

San Jose State University,
San Jose, CA

N. G. Hadjiconstantinou

Massachusetts Institute of Technology,
Cambridge, MA

We present an adaptive mesh and algorithmic refinement (AMAR) scheme for modeling multi-scale hydrodynamics. The AMAR approach extends standard conservative adaptive mesh refinement (AMR) algorithms by providing a robust flux-based method for coupling an atomistic fluid representation to a continuum model. The atomistic model is applied locally in regions where the continuum description is invalid or inaccurate, such as near strong flow gradients and at fluid interfaces, or when the continuum grid is refined to the molecular scale. The need for such "hybrid" methods arises from the fact that hydrodynamics modeled by continuum representations are often under-resolved or inaccurate while solutions generated using molecular resolution globally are not feasible. In the implementation described herein, Direct Simulation Monte Carlo (DSMC) provides an atomistic description of the flow and the compressible two-fluid Euler equations serve as our continuum-scale model. The AMR methodology provides local grid refinement while the algorithm refinement feature allows the transition to DSMC where needed. The continuum and atomistic representations are coupled by matching fluxes at the continuum-atomistic interfaces and by proper averaging and interpolation of data between scales. Our AMAR application code is implemented in C++ and is built upon the SAMRAI (Structured Adaptive Mesh Refinement Application Infrastructure) framework developed at Lawrence Livermore National Laboratory. SAMRAI provides the parallel adaptive gridding algorithm and enables the coupling between the continuum and atomistic methods.

[DOI: 10.1115/1.1792275]

1 Introduction and Background

Multiscale simulation of complex physical systems has received increasing attention in recent years. The primary challenge lies in resolving physical phenomena occurring over a broad range of spatial and temporal scales. Often, this challenge cannot be met by conventional, single-method formulations. Efficient multiscale formulations respond to this challenge by limiting the use of an expensive high-resolution model (e.g., atomistic) to the regions in which it is needed, while using a simpler, less expensive method, in the rest of the computational domain. Such a hybrid approach allows effective use of each method in different regions of the problem (e.g., interior and exterior of a shock wave).

Numerous hybrid methods have been proposed and demonstrated for solids [1–3], liquids [4–7], and gases [8–16]. The purpose of this paper is to illustrate one methodology for integrating the concepts of adaptive mesh refinement and hybrid methods. While our continuum-atomistic approach using adaptive mesh refinement uses an unsplit Godunov scheme for the continuum Euler equations (discussed in Section 2), the coupling methodology is in no way limited to a particular continuum solution method, provided the latter is a time-explicit conservative formulation [17]. This will be particularly important if the hybrid method needs to be used for the simulation of small-scale flows where Navier-Stokes capabilities are required. In fact, perhaps the most important message from this work is that powerful coupling methods can be developed from already existing continuum-continuum matching techniques [5,18]. This is possible because coupling can only be achieved in regions where both descriptions are equivalent, that is, where the continuum description is valid. This approach is in fact preferred because it typically results in coupling

methods that (1) can be rigorously shown to converge and (2) place very few restrictions on the continuum and atomistic constituents other than accommodating the coupling recipe.

In what follows, we will first describe the continuum and atomistic methods and their coupling. Next, the challenging question of choosing reliable grid and algorithm refinement criteria to track fluid interfaces is discussed. Finally, we present numerical results from several test cases and compare them with theory and other simulations to verify our approach.

2 Adaptive Mesh and Algorithmic Refinement

This section describes the Adaptive Mesh and Algorithm Refinement (AMAR) methodology in which a continuum algorithm is replaced by a particle algorithm at the finest grid scale in a hierarchical adaptive grid refinement (AMR) setting. Details of the general AMAR scheme have been presented elsewhere [14]. We will summarize them here for completeness and note differences from previous work.

2.1 AMR Algorithm for Continuum Hydrodynamics. In the AMAR implementation described herein, we employ a structured AMR grid hierarchy on which we solve the compressible, two-species Euler equations on every grid level except the finest. On the finest level, the solution is represented by the Direct Simulation Monte Carlo (DSMC) method. Note that AMAR uses the same adaptive meshing and time integration algorithms developed for continuum modeling of shock hydrodynamics [19,20].

Consider the two-fluid Euler equations in conservative integral form

$$\frac{d}{dt} \int_{\Omega} \mathbf{U} dV + \oint_{\partial\Omega} \mathbf{F} \cdot \hat{n} dS = 0, \quad (1)$$

where

Contributed by the Fluids Engineering Division for publication in the JOURNAL OF FLUIDS ENGINEERING. Manuscript received by the Fluids Engineering Division August 21, 2003; revised manuscript received March 29, 2004. Associate Editor: S. Balachandar.

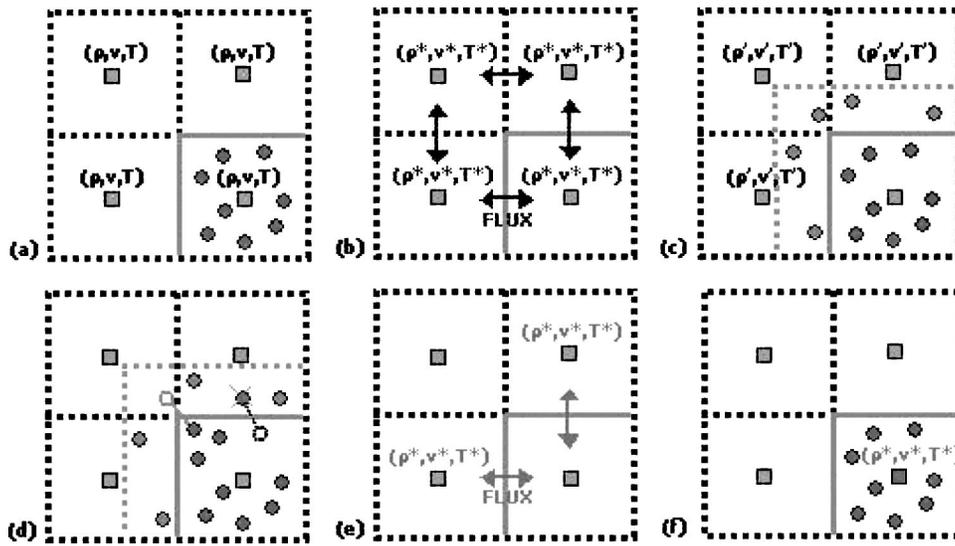


Fig. 1 Outline of AMAR hybrid: (a) Beginning of a time step; (b) advance the continuum grid; (c) create buffer particles; (d) advance DSMC particles; (e) refluxing; (f) reset overlying continuum grid.

$$\mathbf{U} = \begin{pmatrix} \rho \\ p_x \\ p_y \\ p_z \\ e \\ \rho c \end{pmatrix}; \quad \mathbf{F}^x = \begin{pmatrix} \rho u_x \\ \rho u_x^2 + P \\ \rho u_x u_y \\ \rho u_x u_z \\ (e + P)u_x \\ \rho c u_x \end{pmatrix} \quad (2)$$

Here, we provide only the x -direction component of the flux terms; other directions are similar. We assume a two-species gas with the mass concentrations of the two species being c and $(1 - c)$; generalization to more species is straightforward. Discrete time integration is achieved by using a finite volume approximation to Eq. 1. This yields a conservative finite difference expression with \mathbf{U}_{ijk}^n appearing as a cell-centered quantity at each time level and $\mathbf{F}_{i+1/2,j,k}^{x,n+1/2}$ located at faces between cells at half-time levels. We use a second-order version of an unsplit Godunov scheme to approximate the fluxes [21–23].

Time stepping on an AMR grid hierarchy involves interleaving time steps on individual levels [20]. Each level has its own spatial grid resolution and timestep (typically constrained by a CFL condition). The key to achieving a conservative AMR algorithm is to define a discretization for Eq. 1 that holds on every region of the grid hierarchy. In particular, the discrete cell volume integrals of \mathbf{U} and the discrete cell face integrals of \mathbf{F} must match on the locally-refined AMR grid. Thus, integration of a level involves two steps: solution advance and solution synchronization with other levels. Synchronizing the solution across levels assumes that fine grid values are more accurate than coarse grid values. So, coarse values of \mathbf{U} are replaced by suitable cell volume averages of finer \mathbf{U} data where levels overlap, and discrete fine flux integrals replace coarse fluxes at coarse-fine grid boundaries. Although the solution is computed differently in overlapping cells on different levels as each level is advanced initially, the synchronization procedure enforces conservation over the entire AMR grid hierarchy.

2.2 Atomistic Algorithm. Due to our interest in gas flows [24–28], the atomistic algorithm we use is the direct simulation Monte Carlo (DSMC) method [29]. In DSMC, the state of the system is given by positions and velocities of particles, $\{\mathbf{r}_\alpha, \mathbf{v}_\alpha\}$. The system evolves in time using the following two step approach. First, particles are moved without interaction; that is, their positions are updated to $\mathbf{r}_\alpha + \mathbf{v}_\alpha \Delta t_p$, where Δt_p is a DSMC

timestep. Appropriate boundary conditions are applied to particles that reach the boundary of the DSMC domain. Second, after all particles have moved, a given number are randomly selected for collisions. Rather than exactly calculate successive collisions, as in molecular dynamics [30], the DSMC method generates collisions stochastically with scattering rates and post-collision velocity distributions determined from the kinetic theory of a dilute gas. Accuracy of the splitting of streaming and collisions requires the time step Δt_p to be a fraction of the mean collision time τ_m for a particle [31,32]. DSMC has been rigorously shown to provide accurate solutions to the Boltzmann equation in the limit as the number of particles becomes large and the DSMC collision cell size and time step become small [33].

Although DSMC is orders of magnitude faster than molecular dynamics for simulation of gases, it is orders of magnitude slower than continuum algorithms for solving partial differential equations of hydrodynamics in the Navier-Stokes limit. Thus, only flow regions that require molecular resolution are treated by DSMC in the AMAR approach.

2.3 Continuum-Atomistic Coupling. During time integration of continuum grid levels, fluxes computed at each cell face are used to advance the solution \mathbf{U} (Fig. 1(b)). Continuum values are advanced using a time increment Δt_c appropriate for each level, including those that overlay the DSMC region. When the particle level is integrated, it is advanced to the new time on the finest continuum level using a sequence of particle time steps, Δt_p . The relative magnitude of Δt_p to the finest continuum grid Δt_c depends on the finest continuum grid spacing Δx (typically a few λ) and the particle mean collision time.

Euler solution information is passed to the particles via buffer cells surrounding the DSMC region. At the beginning of each DSMC integration step, particles are created in the buffer cells using the continuum hydrodynamic values (ρ, \mathbf{u}, T) and their gradients (Fig. 1(c)). Since the continuum solution is advanced first, these values are time interpolated between continuum time steps for the sequence of DSMC time steps needed to reach the new continuum solution time. DSMC buffer cells are one mean free path wide; thus, the time step Δt_p is constrained so that it is extremely improbable that a particle will travel further than one mean free path in a single time step. The particle velocities are drawn from an appropriate distribution for the continuum solver, such as the Chapman-Enskog distribution [34].

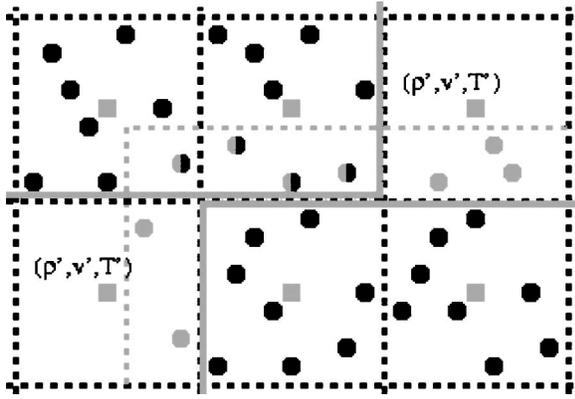


Fig. 2 Multiple DSMC regions are coupled by copying particles from one DSMC region (upper left) to the buffer region of an adjacent DSMC region (lower right). After copying, regions are integrated independently over the same time increment.

During each DSMC time integration step, all particles are moved, including those in the buffer regions (Fig. 1(d)). A particle that crosses the interface between continuum and DSMC regions will eventually contribute to the flux at the corresponding continuum cell face during the synchronization of the DSMC level with the finest continuum level. After moving particles, those residing in buffer regions are discarded. Then, collisions among the remaining particles are evaluated and new particle velocities are computed.

After the DSMC region has advanced over an entire continuum grid timestep, the continuum and DSMC solutions are synchronized in a manner analogous to the AMR level synchronization process described earlier. First, the continuum values in each cell overlaying the DSMC region interior are set to the conservative averages of data from the particles within the continuum grid cell region (Fig. 1(f)). Second, the continuum solution in cells adjacent to the DSMC region is recomputed using a “refluxing” process (Fig. 1(e)). That is, a flux correction is computed using a space and time integral of particle flux data,

$$\delta\mathcal{F} = -A\mathbf{F}^{n+1/2} + \sum_{\text{particles}} \mathcal{F}_p. \quad (3)$$

The sum represents the flux of the conserved quantities carried by particles passing through the continuum cell face during the DSMC updates. Finally,

$$\mathbf{U}^{n+1} = \overline{\mathbf{U}^{n+1}} + \frac{\Delta t_c \delta\mathcal{F}}{\Delta x \Delta y \Delta z} \quad (4)$$

is used to update the conserved quantities on the continuum grid where $\overline{\mathbf{U}^{n+1}}$ is the coarse grid solution before computing the flux correction.

In summary, the coupling between the continuum and DSMC methods is performed in three operations. First, continuum solution values are interpolated to create particles in DSMC buffer cells before each DSMC step. Second, conserved quantities in each continuum cell overlaying the DSMC region are replaced by averages over particles in the same region. Third, fluxes recorded when particles cross the DSMC interface are used to correct the continuum solution in cells adjacent to the DSMC region. This coupling procedure makes the DSMC region appear as any other level in the AMR grid hierarchy.

Multiple DSMC parallelepiped regions (i.e., *patches*) are coupled by copying particles from patch interiors to buffer regions of adjacent DSMC patches (see Fig. 2). That is, particles in the interior of one patch supply boundary values (by acting as a reservoir) for adjacent particle patches. After copying particles into buffer regions, each DSMC patch may be integrated *independently*,

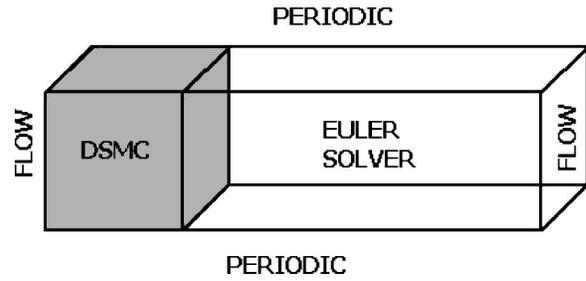


Fig. 3 3D AMAR computational domain for investigation of tolerance parameter variation with number of particles in DSMC cells

dently, in the same fashion that different patches in a conventional AMR problems are treated after exchanging boundary data.

2.4 Refinement Criteria. Criteria for refining the mesh and transitioning from continuum to a particle scheme is a significant research topic and is generally problem specific. Standard AMR methods assume that the differential equations of continuum hydrodynamics are valid at all length scales in the computation and grid refinement is often based on ad hoc notions (e.g., refine around steep gradients) or analytical error estimation techniques involving the continuum differential equations (e.g., Richardson extrapolation). In contrast, hybrid methods apply computational models matched to the flow properties at each physical scale. While results presented later demonstrate that we can effectively focus DSMC locally to capture shocks and diffusion fronts, many research issues remain.

The AMAR algorithm can refine the grid and algorithm based on any of a number of possible criteria. However, we have found that for single species flows, refinement based on density gradients is reliable. Tracking concentration gradients or concentration values within some interval is effective for multispecies flows involving concentration interfaces. Parameters for transitioning to DSMC are based on the continuum breakdown parameter method proposed by Bird [35], i.e., refinement is triggered by spatial gradients exceeding continuum tolerances. The gradient detector formula that we use is a variation of a sharp discontinuity detector by Trangenstein and Pember [36].

Due to spontaneous stochastic fluctuations in DSMC computations, it is important to track gradients in a manner that does not allow the fluctuations to trigger unnecessary refinement and excessively large DSMC regions. Let us consider the gas density as an example. Using the fact that for an ideal gas under equilibrium conditions the number of particles in a given volume is Poisson distributed, it can be shown that the standard deviation in the normalized density gradient is [37],

$$\sqrt{\left\langle \left(\frac{d\rho/dx}{\rho} \right)^2 \right\rangle} \approx \sqrt{\left\langle \left(\frac{N_{i+1} - N_i}{\Delta x \langle N_i \rangle} \right)^2 \right\rangle} = \frac{\sqrt{2}}{\Delta x \sqrt{\langle N_{\text{cell}} \rangle}}, \quad (5)$$

where N_{cell} is the number of particles in a cell where macroscopic properties are defined. The fluid density fluctuation can only be reduced by increasing the number of DSMC simulation particles. This has consequences for the use of density gradient tolerances R_ρ used for AMAR. In general, such tolerances must be based on the number of particles used for the atomistic domain since the spatial gradients of density on the *coarse* grid which is fluctuating (as shown below) are used to decide whether refinement will take place. In particular, in our version of density-gradient-based refinement, refinement occurs in regions where the nondimensionalized density gradients are above the R_ρ threshold, i.e.,

$$R_\rho < \frac{2\lambda}{\rho} \left| \frac{d\rho}{dx} \right| \quad (6)$$

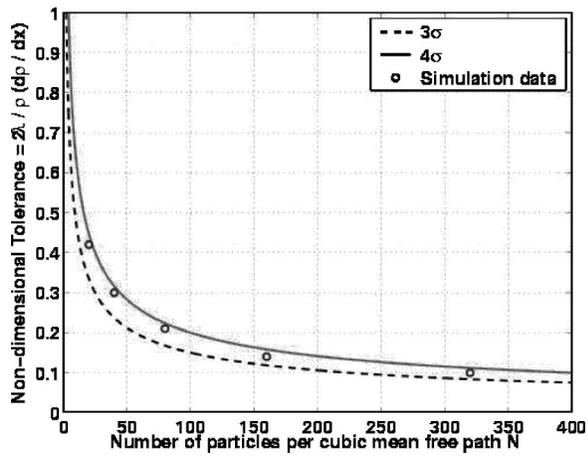


Fig. 4 Variation of density gradient tolerance with number of DSMC particles N . Here we use N because in our implementation $N_{\text{cell}}=N$.

Note that the density gradient is calculated in the continuum mesh on a coarser AMR hierarchy level than the DSMC level.

To determine the minimum value of R_ρ required to prevent growth of the atomistic region, simulations were conducted using the domain geometry shown in Fig. 3 for a range of N_{cell} . Grid refinement occurs during a “trigger” event where the density fluctuations exceed R_ρ and the atomistic subdomain grows in the axial direction by a single continuum cell width. The value of R_ρ that yields a 5–10% trigger rate (i.e., between 5–10 trigger events per 100 iterations) is plotted in Fig. 4 as a function of N_{cell} . In what follows we outline how theoretical predictions bounding these numerical results, shown as solid lines in Fig. 4, can be obtained.

For the geometry considered in this test problem, each continuum cell consists of 8 DSMC cells and hence effectively the contribution of $8 \times N_{\text{cell}}$ particles is averaged to determine the density gradient between continuum cells. If we use Eq. (5) for continuum cells, we obtain

$$\sigma = \sqrt{\left\langle \left(\frac{d\rho/dx}{\rho} \right)^2 \right\rangle_c} \approx \frac{1}{2\Delta x \sqrt{\langle N_{\text{cell}} \rangle}} \quad (7)$$

Note that we are assuming that the fluctuation of the continuum cells across from the atomistic-continuum interface is approximately the same as that in the atomistic region. This was shown to

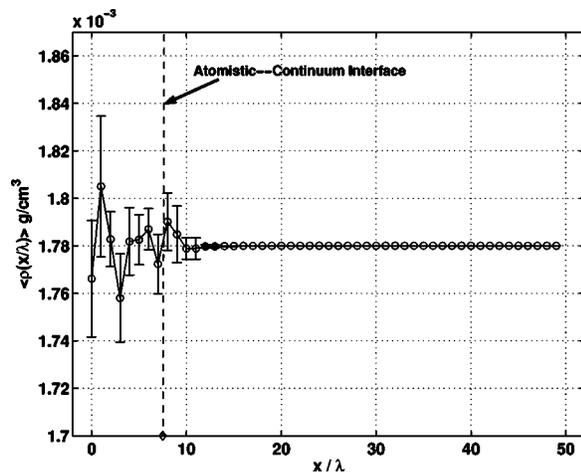


Fig. 5 Average density for stationary fluid Euler-DSMC hybrid simulation with $N_{\text{cell}}=80$. Error bars give one standard deviation over 10 samples.

be the case for the diffusion equation and a random walk model in Ref. [38], and is verified here for the Euler-DSMC system as shown in Fig. 5. This allows the use of Eq. (5) that was derived assuming 2 atomistic cells. Note that the observed trigger event is a composite of a large number of probable density gradient fluctuations that could exceed R_ρ ; gradients across all possible nearest neighbor cells, next-to-nearest neighbor cells and diagonally-nearest neighbor cells are all individually evaluated by the refinement routines and checked against R_ρ . For a 10% trigger rate (or equivalent probability of trigger) the probability of an individual cell-pair having a density fluctuation exceeding R_ρ can be estimated as $O(0.1/100)$ by observing that,

1. since the trigger event is rare, probabilities can be taken as additive,
2. for the geometry considered, there are ≈ 300 nearest neighbor, next-nearest neighbor and diagonal cells that can trigger refinement and
3. the rapid decay of the Gaussian distribution ensures the decreasing probability ($O(0.1/100) \sim O(0.001)$) of a single event does not significantly alter the corresponding confidence interval and thus an exact enumeration of all possible trigger pairs with correct weighting factors is not necessary.

For example, our probability estimate at $O(0.001)$ suggests that our confidence interval is $3\sigma-4\sigma$. This is verified in Fig. 4. Larger trigger rates can be achieved by reducing R_ρ . Curves shown in Fig. 4 help prototype tolerance criteria using a small number of particles prior to running larger simulations.

2.5 Euler-DSMC Code Implementation. The Euler-DSMC AMR code discussed here is composed of elements from the SAMRAI object-oriented framework, developed at Lawrence Livermore National Laboratory, and numerical routines specific to the application. SAMRAI provides a general, flexible software toolbox for developing multiphysics AMR applications and supports general parallel data management capabilities, including particle representations, on an AMR grid hierarchy.

The organization of major algorithmic parts in the hybrid Euler-DSMC code is similar to that of an Euler-only AMR code. However, the hybrid code requires a new level integrator that coordinates DSMC and Euler operations on different hierarchy levels. The new integrator, developed for this project, was constructed from elements in SAMRAI. It is interesting to note that all classes appearing in an Euler-only application are used without modification in the hybrid code. Also, the DSMC data structures and numerical routines, developed previous to incorporation in the hybrid application, were introduced without significant modification.

Recall that the Euler continuum model and DSMC particle model are vastly different numerical approaches. The Euler model represents compressible fluid flow as a deterministic system of partial differential equations containing a few grid-based variables. DSMC approximates the Boltzmann equation using a representative, stochastic sampling of a collection of particles whose state and motion are essentially gridless. The DSMC data structures and numerical routines are insulated from SAMRAI abstractions by a “wrapper” interface class. This class serves two important functions. First, it acts as a translator between SAMRAI patch data and the DSMC particle structures. Second, it allows the particles to be manipulated on a distributed parallel machine by SAMRAI. More importantly, the *serial* DSMC routines were coupled to the SAMRAI parallel communication framework without changing the particle structures or routines or recompiling SAMRAI library code. Additional details describing how this is done appear in Ref. [17].

3 Code and Algorithm Verification Tests

This section describes a number of test problems used to verify the AMAR hybrid formulation. Tolerance parameters used for grid refinement are also described. The single-species tests use gaseous

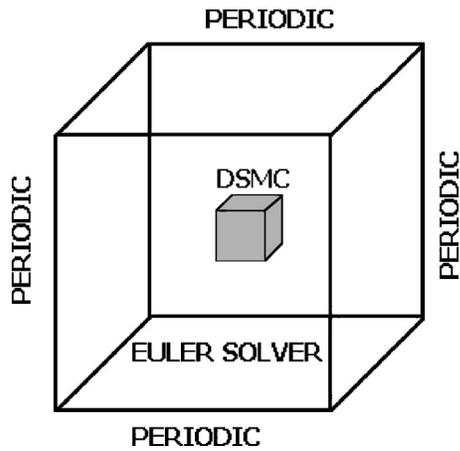


Fig. 6 Computational domain for uniform field test

Argon (molecular mass $m = 6.63 \times 10^{-23}$ g, hard sphere diameter $d = 3.66 \times 10^{-8}$ cm) at atmospheric conditions ($P = 1.013 \times 10^5$ Pa and average temperature $\bar{T} = 273$ K). This choice is for convenience reasons. The hard sphere diameter for Argon is well-known to reproduce equilibrium and nonequilibrium properties accurately. Argon is also used in a large number of DSMC studies and thus a substantial literature base of simulation and experimental results exist for verification purposes.

3.1 Thermodynamic Equilibrium. We start with the thermodynamic equilibrium test which serves to illustrate that due to the underlying $Re \rightarrow \infty$ ($Pe \rightarrow \infty$) assumption associated with our continuum solver, care needs to be taken in not using this model in nonconvectively dominated situations. The computational domain, illustrated in Fig. 6, consists of a cubic DSMC region of size $4\lambda \times 4\lambda \times 4\lambda$ embedded in the center of an Euler continuum grid of size $32\lambda \times 32\lambda \times 32\lambda$. The DSMC computation uses 800 particles per λ^3 . The continuum cell size is 2λ , while the DSMC grid size is λ . Periodic boundary conditions are applied at each face of the cubic domain. This geometry is the same as in Ref. [14] for comparison purposes. The density and temperature field were initialized at $\rho = 1.78 \times 10^{-3}$ g/cm³ and $T = 273$ K and the initial velocity field was set to zero in all directions. Adaptive refinement (both mesh and algorithmic) was disabled in order to illustrate the following effect. Figure 7 shows that the number of particles in the atomistic subdomain slowly increases in time. This phenomenon can be explained as follows. Although the initial

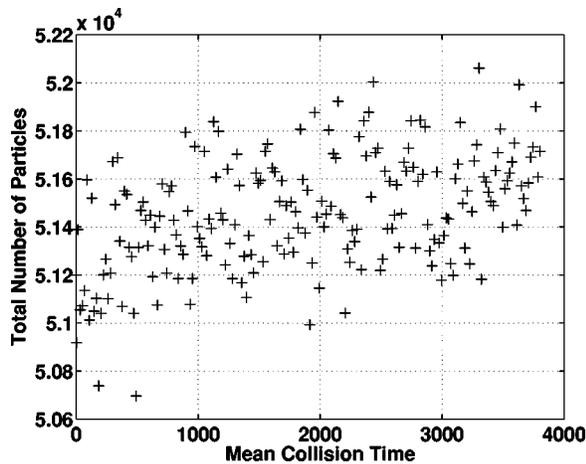


Fig. 7 Particle increase in the DSMC domain resulting from net heat flux transfer from the DSMC to the Euler region

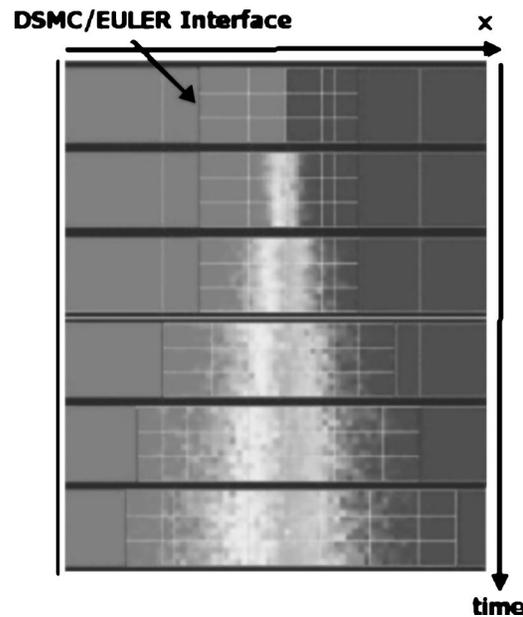


Fig. 8 Computational domain for self-diffusion interface tracked adaptively. The borders of DSMC patches are indicated by the boxes near the middle of the domain. The Euler model is applied in the remainder of the domain.

conditions are uniform, the statistical nature of atomistic solution generates fluctuations that transfer heat to the continuum subdomain. Since the Euler model possesses no mechanism to transfer thermal energy back to the atomistic domain, the result is an energy increase in the continuum subdomain and a corresponding energy decrease in the atomistic subdomain as total energy is conserved. This, in turn, produces an increase in density in the atomistic subdomain so as to maintain mechanical equilibrium (i.e., constant pressure). Hence the total number of particles in the atomistic subdomain increases. This phenomenon has been observed in other Euler/DSMC hybrid schemes [14] and has been shown to be suppressed when diffusive transport (Navier-Stokes terms) are included. Recall, however, that in the $Re \rightarrow \infty$ ($Pe \rightarrow \infty$) limit appropriate for Euler calculations, the diffusive timescale associated with molecular transport (the cause of this phenomenon) is much (infinitely) longer than the convective timescale which dominates the flow. As a result, this phenomenon should have no effect in the Euler limit. In some sense, it is a manifestation of the fact that equilibrium is not a convectively dominated situation appropriate to an Euler calculation. This is supported by simple tests we have performed which show that in the presence of a uniform flow this effect is reduced. Additionally, the excellent agreement between our results and benchmark Euler flows presented in the following sections shows that this phenomenon has no effect on solutions in the Euler limit.

3.2 Concentration Diffusion. Concentration diffusion tests were conducted to verify the ability of AMAR to track the spreading of an interface between two gases. The Euler model contains no diffusion terms (except for artificial numerical diffusion) so spreading of the interface is governed by physics modeled by the DSMC routines.

The diffusion coefficient for two gases modeled as hard spheres can be approximated as [39],

$$D_{12} = \frac{3}{16} \frac{\sqrt{2\pi k^3 T^3 / \mathcal{M}}}{P \pi d^2} = \frac{3}{8} \frac{1}{n d^2} \sqrt{\frac{kT}{2\pi \mathcal{M}}} \quad (8)$$

where $\mathcal{M} = (1/m_1 + 1/m_2)^{-1} = m_1 m_2 / (m_1 + m_2)$ is the reduced mass, and $d = (d_1 + d_2)/2$ is the average atomic diameter.

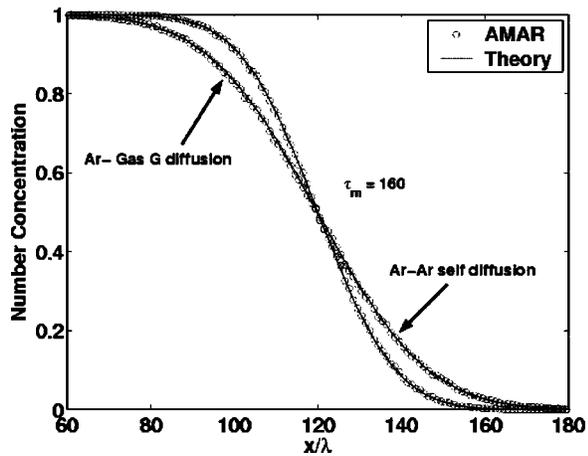


Fig. 9 Comparison of profiles obtained simulating diffusion with AMAR with theoretical diffusion profiles. Both self-diffusion and two-species diffusion are shown. Note λ refers to the Ar-Ar mean free path. The mean collision time τ_m is also associated with the Ar-Ar system.

A simple concentration diffusion test involves Argon gas on either side of an interface “colored” differently, i.e., the gases are essentially isotopes with negligible mass differences that undergo self-diffusion. The self-diffusion coefficient for Argon at standard temperature and pressure is $D_{11}=0.14 \text{ cm}^2/\text{s}$. Figure 8 shows the evolution of the atomistic-continuum computational domain for this self-diffusion test. Initially the gases are separated by a discontinuous interface corresponding to a step function profile for the gas concentration. The gradient of the corresponding mass concentration is used to place the atomistic region at the gas interface at the initial time. Subsequently, the mixing region is tracked using a “mass-concentration-value” refinement criteria which triggers below 0.001 and above 0.999. These values ensure that negligible concentration gradients exist across the atomistic-continuum interface.

The hybrid concentration profiles for the self-diffusion case are compared with theoretical profiles in Fig. 9. Also shown is the concentration profile for a test case using Argon and a fictitious gas G with hard-sphere diameter $d_2=1.516 \times 10^{-8} \text{ cm}$ such that

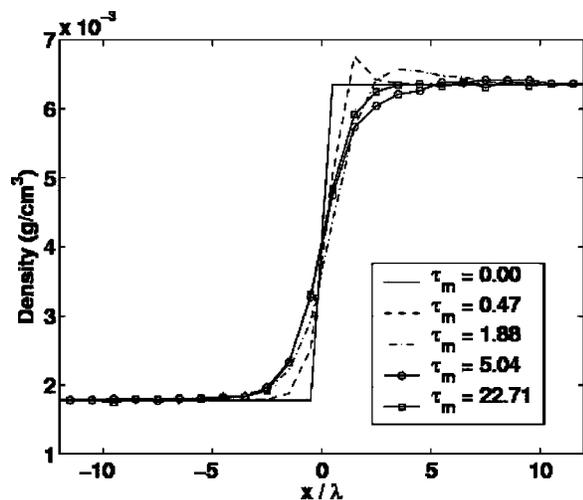


Fig. 10 Argon gas density profile evolution to equilibrium. τ_m is the mean collision time.

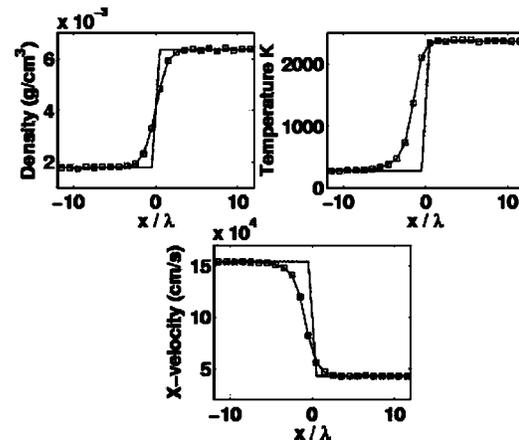


Fig. 11 Equilibrium shock wave profiles for density, temperature and velocity in a stationary Argon shock. The solid line connects the Rankine-Hugoniot jump values through a sharp jump centered on the AMAR shock location as determined by the AMAR density profile. The solid-square line is the AMAR result. Note that the AMAR temperature jump “leads” the AMAR density jump as documented in Ref. [29]. The agreement between AMAR and Rankine-Hugoniot jump values is excellent.

the diffusion coefficient is twice the self-diffusion coefficient; that is, $D_{12}=2 \times D_{11}=0.28 \text{ cm}^2/\text{s}$. The simulation results show excellent agreement with theory in both cases.

3.3 Stationary Shock Waves. Stationary shock simulations were conducted for both a single gas and a binary gas mixture.

3.3.1 Single Gas Stationary Shock Wave. For the single gas case, a shock wave is initialized using the discontinuous step profile given by the Rankine-Hugoniot conditions for Argon gas with an upstream Mach number of 5.0. The shock wave density, temperature and velocity ratios for this Mach number are 3.57, 8.68 and 0.28, respectively. A density gradient tolerance parameter $R_\rho=0.2$ was used to detect and refine the continuum grid across the shock. This value for R_ρ creates a stable $\pm 10\lambda$ atomistic region ahead of and behind the shock front.

The time evolution of the density profile is shown in Fig. 10. The initial step profile gradually transitions to a smoother curved profile within 20 mean collision times. The final equilibrium profiles for the pressure, density, velocity and temperature are shown in Fig. 11. The hybrid solution matches the Rankine-Hugoniot jump values in the far field while resolving the flow discontinuity at the shock front. Also note the temperature jump is shown to “lead” the density jump as documented in Ref. [29]. Note that since the initial density gradient in the streamwise direction is essentially infinite, the shock region will be refined for any setting of R_ρ . As the profile becomes smoother however, the value of R_ρ is critical to ensure the shock front remains tagged for refinement while preventing excessive atomistic subdomain growth due to spurious statistical fluctuations. This is achieved successfully with the $R_\rho=0.2$ setting.

3.3.2 Binary Gas Stationary Shock Wave. The binary gas shock simulation was conducted using a mixture of Helium and Xenon gases with number densities of 97% and 3% respectively. The hard sphere mass and diameter chosen to model Helium and Xenon were $m_1=6.65 \times 10^{-24} \text{ g}$, $m_2=2.18 \times 10^{-22} \text{ g}$ and $d_1=2.28 \times 10^{-8} \text{ cm}$, $d_2=5.18 \times 10^{-8} \text{ cm}$, respectively. The upstream flow Mach number was set to 3.89 with a temperature of 300 K and reference mass density of $1.07 \times 10^{-7} \text{ g/cm}^3$. These flow conditions were chosen to allow for convenient comparison with literature results. The corresponding Rankine-Hugoniot relations for the shock density, temperature and velocity ratios are 3.34, 5.59, and 0.3, respectively. Tolerance parameters were not

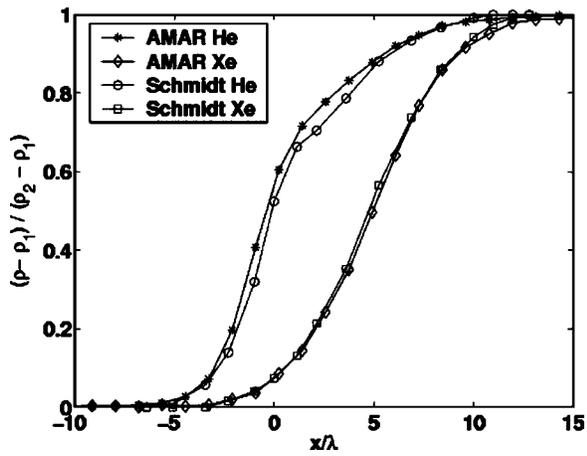


Fig. 12 Comparison of He-Xe binary gas shock wave equilibrium profiles computed with AMAR (blue lines) and with DSMC alone (red lines). The mixture mean free path $\lambda = 0.46$ mm for this test.

used for this test and instead the refinement region was user specified to extend 15λ ahead of and 35λ behind the shock front.

Similar to the single gas shock case, the binary gas shock profile also transitions from an initial discontinuous step profile to a smoother equilibrium profile. A comparison of the equilibrium hybrid density profiles with a fully DSMC simulation performed by Schmidt and Worner [40] are shown in Fig. 12. Good agreement can be seen for both Helium and Xenon density profiles.

3.4 Moving Shock Wave. Adaptive feature-tracking of the AMAR hybrid scheme is further verified using a $M=5$ moving shock passing through a stationary Argon gas. Figure 13 shows the atomistic subdomain dynamically tracking the passage of the shock front. Similar to the case of a stationary $M=5$ shock wave a density gradient tolerance $R_p=0.2$ was found sufficient to extend the atomistic subdomain $\pm 10\lambda$ about the shock front.

The density profile of the moving shock is shown in Fig. 14. Good comparison is seen with the analytical result. Note the hybrid profile does not produce spurious postshock oscillations well-known to plague continuum-only schemes [41,42]. Conventional shock capturing techniques for the Euler equations require artificial viscosity and enhanced smoothing techniques to reduce oscillations that cannot often be eliminated entirely. The AMAR hybrid

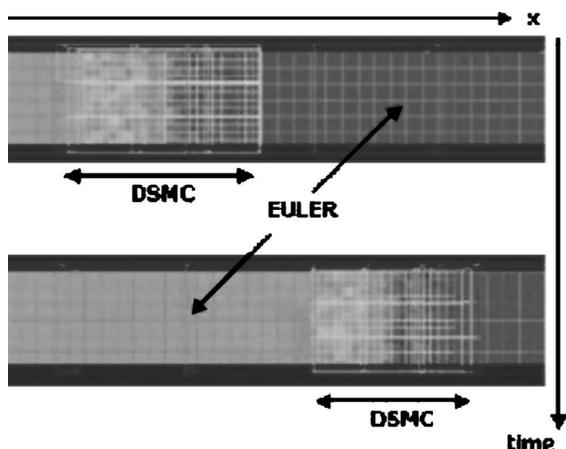


Fig. 13 A moving Mach 5 shock wave through Argon. The AMAR algorithm tracks the shock by adaptively moving the DSMC region with the shock front.

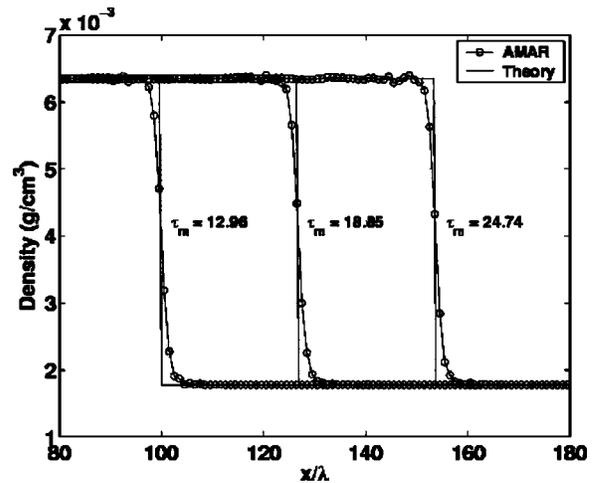


Fig. 14 A moving Mach 5 shock wave through Argon. The AMAR profile (red dots) is compared with the analytical time evolution of the initial discontinuity (blue lines). τ_m is the mean collision time.

scheme allows for accurate, adaptive and stable resolution of shock fronts without the need for artificial numerical constructs.

3.5 Richtmyer-Meshkov Instability. The Richtmyer-Meshkov instability (RMI) (Meshkov [43,44], Richtmyer [45]) is generated when the interface between two fluids is impulsively accelerated by a shock wave. The shock impulse causes perturbations on the interface to grow in size which in turn creates a mixing layer between the two gases [46]. In this paper RMI serves as a test problem from which to assess both the adaptive feature tracking and multispecies capability of the AMAR hybrid scheme. For the simulation considered here adaptive mesh refinement is considered only for tracking the shock front, and for computational convenience, only a small system was considered.

The computational domain illustrated in Fig. 15 was used for the RMI simulation. Argon gas and a fictitious gas B with hard sphere mass and diameter $m_B = 1.326 \times 10^{-22}$ g, $d_B = 3.66 \times 10^{-8}$ cm were chosen for the test. In order to reduce diffusion between the gases the cross collision diameter $d = (d_1 + d_2)/2$ was increased by a factor of 4. The density ratio across the gas-gas interface was initialized to $\rho_2/\rho_1 = 1.5$ (which corresponds to an Atwood number $At = ((\rho_2/\rho_1) - 1)/((\rho_2/\rho_1) + 1) = 0.2$). This

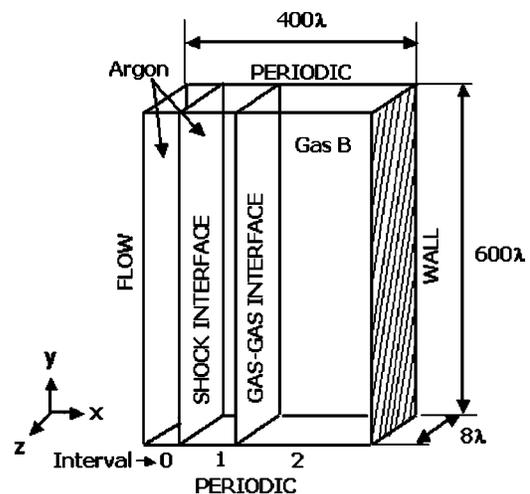


Fig. 15 Computational domain for Richtmyer-Meshkov instability simulation

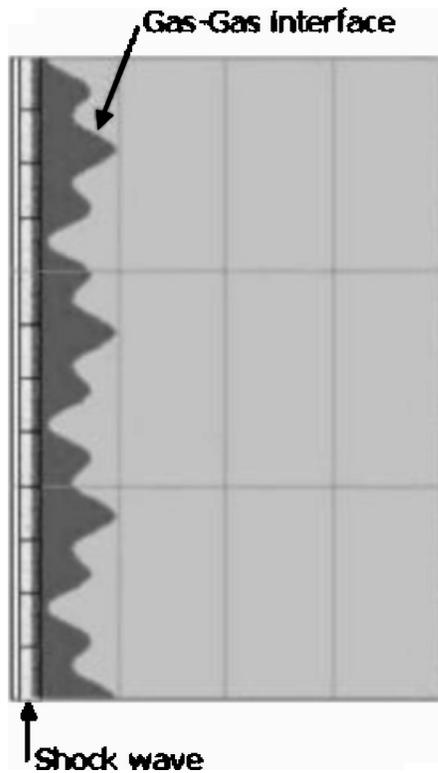


Fig. 16 Richtmyer-Meshkov instability simulation, time $t = 1.3\tau_m$ where τ_m is the Argon-Argon mean collision time. The shock wave is ahead of the gas-gas interface.

simulation utilizes 20 DSMC particles per cubic mean free path of Argon at the reference density. The lighter Argon gas occupies the left hand side of the interface while the heavier gas B occupies the right hand side. The shock wave is initialized upstream of the

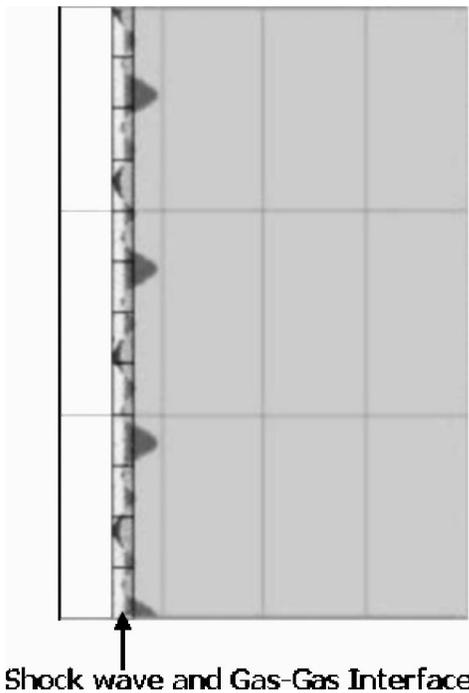


Fig. 17 Richtmyer-Meshkov instability simulation, time $t = 26.0\tau_m$ where τ_m is the Argon-Argon mean collision time. The shock wave intercepts the gas-gas interface.

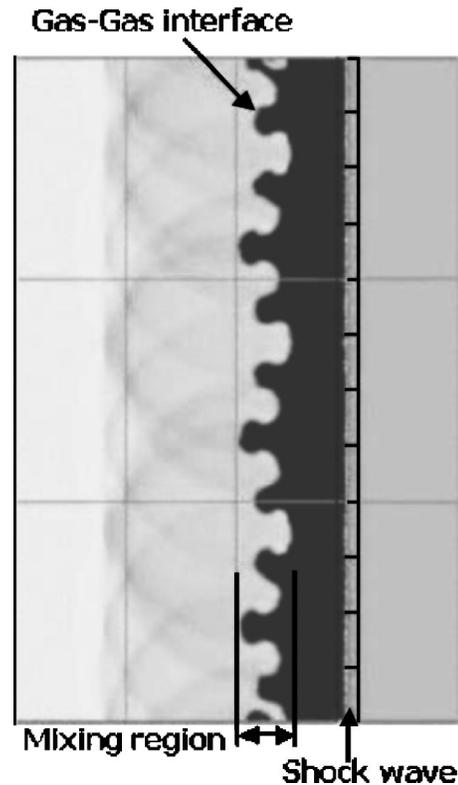


Fig. 18 Richtmyer-Meshkov instability simulation, time $t = 170.1\tau_m$ where τ_m is the Argon-Argon mean collision time. The shock wave has passed the gas-gas interface.

gas-gas interface and propagates at Mach number $M = 4.0$ through Argon gas. Classical Navier-Stokes methods have been unable to perform *rigorous* simulation at these Mach numbers [29] which are typical of experimental conditions [47]. The gas-gas interface has an initial sinusoidal profile with wavelengths 171λ and 57λ superimposed. Peak to peak amplitudes vary between 8λ and 32λ . Both the wavelength and disturbance amplitude are typical of real experimental conditions [47].

On interception with the interface, a reflected shock returns upstream and a transmitted shock continues through gas B. On reflection with the right hand wall, the transmitted shock returns past the interface and leaves the domain through the left face. A density gradient tolerance $R_\rho = 0.6$ ensures the atomistic subdomain is localized about the shock wave region only, while the gas-gas interface remains nonrefined. Figures 16, 17 and 18 show the propagation of the shock wave. Note that Fig. 18 also shows a reflected shock wave traveling upstream in addition to the transmitted shock through the gas-gas interface. The choice of R_ρ allows only for the transmitted shock wave to be refined in this case. This hybrid approach achieves a computational savings of $O(10)$ over a fully atomistic solution (assuming that the continuum solution cost is negligible compared to the DSMC solution cost the savings for this problem is more than a factor of 20). Our simulations show that through a judicious choice of refinement criteria and the development of a theory to quantify the effect of fluctuations, reliable fully adaptive mesh and algorithm refinement algorithms are possible. Larger calculations which will allow for comparison with experimental data and other simulation codes (for an example see Ref. [47]) are planned in the future.

4 Concluding Remarks

We have described an adaptive mesh and algorithmic refinement (AMAR) scheme for modeling multiscale, multispecies gas

dynamics and have demonstrated its effectiveness in a hybrid Euler-DSMC code via a series of computational test cases. The atomistic model is applied locally in regions where the continuum description is invalid or inaccurate, such as near strong flow gradients and at fluid interfaces. In particular, accurate and stable solutions for concentration diffusion, single and binary gas stationary shocks waves and moving shock waves have been obtained. These tests suggest that the AMAR code can perform highly-resolved DSMC calculations in a more cost effective manner than a DSMC-only code. An evaluation of code performance and scaling studies to larger problems will be the subject of future work.

While this effort shows the promise of the approach, much work remains to address research issues associated with such hybrid methods. In future work we plan to rigorously address issues related to accuracy and robustness of particle-continuum hybrids, criteria for adaptively switching between models, and effects of statistical fluctuations in particle schemes on the stability of continuum methods to which they may be coupled. In addition, we expect that future adoption of a Navier-Stokes model in the same algorithmic framework will yield a very useful tool for the design of micro and nano scale devices involving *compressible* flow features in which local continuum breakdown occurs.

Finally, the range of scales and dynamic nature of multi-algorithm hybrids that may be applied to a variety of important physical problems emphasizes the need for efficient computational approaches for large-scale parallel computing platforms. The AMR grid hierarchy paradigm offers many advantages to algorithm development and parallel code implementation, including the ability to manage both field data and particles in a single grid system while allowing workload and data for each method to be distributed in parallel independently of one other. We are currently working on more effective dynamic load balancing and data distribution algorithms for our code to increase the scale of the problems we can simulate. These results will be reported in the near future.

Acknowledgments

The authors wish to thank B. Alder, F. Alexander, J. Bell, S. Smith, and A. Wissink for helpful discussions. This work was performed under the auspices of the U.S. Department of Energy by University of California Lawrence Livermore National Laboratory under contract number W-7405-Eng-48.

Nomenclature

A	= signed area of a grid cell face
c	= mass concentration of reference fluid species
D	= diffusion coefficient
d	= atomic diameter
e	= total energy density
$\delta\mathcal{F}$	= refluxing correction
\mathbf{F}	= flux vector
k	= Boltzmann's constant
\mathcal{M}	= reduced mass
m	= atomic mass
N	= number of particles per cubic mean free path
N_{cell}	= number of particles per unit
n	= number density
\hat{n}	= surface normal
P	= pressure
$Pe = Re Pr$	= Peclet number
Pr	= gas Prandtl number
$\mathbf{p} = (p_x, p_y, p_z)$	= vector of momentum densities
Re	= Reynolds number
\mathbf{r}	= particle position vector
S	= surface element

T = temperature

Δt = time step

t = time

\mathbf{U} = Euler solution state vector

$\mathbf{u} = (u_x, u_y, u_z)$ = fluid velocity vector

\mathbf{v} = particle velocity vector

$\Delta x, \Delta y, \Delta z$ = grid spacing

Greek

λ = particle mean free path

Ω = fluid volume element

$\partial\Omega$ = boundary of fluid element

ρ = mass density

σ = standard deviation

Subscripts/Superscripts

0 = mean

1,2 = first, second species

α = particle index

c = continuum

i, j, k = coordinate indices

max = maximum

p = particle

n = temporal index

x, y, z = spatial coordinates

References

- [1] Abraham, F. F., Broughton, J. Q., Bernstein, N., and Kaxiras, E., 1998, "Spanning the continuum to quantum length scales in a dynamic simulation of brittle fracture," *Europhys. Lett.*, **44**, pp. 783–787.
- [2] Rudd, R. E., and Broughton, J. Q., 2000, "Concurrent coupling of length scales in solid state systems," *Phys. Status Solidi B*, **217**, pp. 251–291.
- [3] Shenoy, V. B., Miller, R., Tadmor, E. B., Rodney, D., Phillips, R., and Ortiz, M., 1999, "An Adaptive Finite Element Approach to Atomic-Scale Mechanics—The Quasicontinuum Method," *J. Mech. Phys. Solids*, **47**, pp. 611–642.
- [4] O'Connell, S. T., and Thompson, P. A., 1995, "Molecular dynamics-continuum hybrid computations: A tool for studying complex fluid flows," *Phys. Rev. E*, **52**, pp. R5792–R5795.
- [5] Hadjiconstantinou, N. G., 1999, "Hybrid Atomistic-Continuum Formulations and the Moving Contact-Line Problem," *J. Comput. Phys.*, **154**, pp. 245–265.
- [6] Li, J., Liao, D., and Yip, S., 1998, "Coupling continuum to molecular-dynamics simulation: Reflecting particle method and the field estimator," *Phys. Rev. E*, **57**, pp. 7259–7267.
- [7] Flekkoy, E. G., Wagner, G., and Feder, J., 2000, "Hybrid model for combined particle and continuum dynamics," *Europhys. Lett.*, **52**, pp. 271–276.
- [8] Wadsworth, D. C., and Erwin, D. A., 1990, "One-Dimensional Hybrid Continuum/Particle Simulation Approach for Rarefied Hypersonic Flows," *AIAA Paper 90-1690*.
- [9] Hash, D., and Hassan, H., 1996, "A Decoupled DSMC/Navier-Stokes Analysis of a Transitional Flow Experiment," *AIAA Paper 96-0353*.
- [10] Bourgat, J., Le Tallec, P., and Tidriri, M., 1996, "Coupling Boltzmann and Navier-Stokes Equations by Friction," *J. Comput. Phys.*, **127**, pp. 227–245.
- [11] Alder, B. J., 1997, "Highly discretized dynamics," *Physica A*, **240**, pp. 193–195.
- [12] Le Tallec, P., and Mallinger, F., 1997, "Coupling Boltzmann and Navier-Stokes Equations by Half Fluxes," *J. Comput. Phys.*, **136**, pp. 51–67.
- [13] Tiwari, S., and Klar, A., 1998, "Coupling of the Boltzmann and Euler equations with adaptive domain decomposition procedure," *J. Comput. Phys.*, **144**, pp. 710–726.
- [14] Garcia, A. L., Bell, J., Crutchfield, W. Y., and Alder, B. J., 1999, "Adaptive Mesh and Algorithm Refinement using Direct Simulation Monte Carlo," *J. Comput. Phys.*, **154**, pp. 134–155.
- [15] Aktas, O., and Aluru, N. R., 2002, "A Combined Continuum/DSMC Technique for Multiscale Analysis of Microfluidic Filters," *J. Comput. Phys.*, **178**, pp. 342–372.
- [16] Roveda, R., Goldstein, D. B., and Varghese, P. L., 2000, "Hybrid Euler/direct simulation Monte Carlo calculation of unsteady slit flow," *J. Spacecr. Rockets*, **37**(6), pp. 753–760.
- [17] Hornung, R. D., and Kohn, S. R., 2002, "Managing Application Complexity in the SAMRAI Object-Oriented Framework," *Concurrency and Computation: Practice and Experience*, **14**, pp. 347–368.
- [18] Quarteroni, A., 1999, *Domain decomposition methods for partial differential equations*, Oxford, New York; Clarendon Press, Oxford, New York.
- [19] Berger, M., and Olinger, J., 1984, "Adaptive Mesh Refinement for Hyperbolic Partial Differential Equations," *J. Comput. Phys.*, **53**, pp. 484–512.
- [20] Berger, M., and Colella, P., 1989, "Local Adaptive Mesh Refinement for Shock Hydrodynamics," *J. Comput. Phys.*, **82**, pp. 64–84.
- [21] Colella, P., 1985, "A Direct Eulerian MUSCL Scheme for Gas Dynamics," *SIAM (Soc. Ind. Appl. Math.) J. Sci. Stat. Comput.*, **6**, pp. 104–117.

- [22] Colella, P., and Glaz, H. M., 1985, "Efficient Solution Algorithms for the Riemann Problem for Real Gases," *J. Comput. Phys.*, **59**, pp. 264–289.
- [23] Saltzman, J., 1994, "An Unsplit 3D Upwind Method for Hyperbolic Conservation Laws," *J. Comput. Phys.*, **115**, pp. 153–167.
- [24] Hadjiconstantinou, N. G., and Simek, O., 2002, "Constant-Wall-Temperature Nusselt Number in Micro and Nano-Channels," *J. Heat Transfer*, **124**, pp. 356–364.
- [25] Hadjiconstantinou, N. G., 2002, "Sound wave propagation in transition-regime micro- and nanochannels," *Phys. Fluids*, **14**, pp. 802–809.
- [26] Hadjiconstantinou, N. G., 2003, "Comment on Cercignani's second-order slip coefficient," *Phys. Fluids*, **15**, pp. 2352–2354.
- [27] Hadjiconstantinou, N. G., and Simek, O., 2003, "Sound propagation at small scales under continuum and non-continuum transport," *J. Fluid Mech.*, **488**, pp. 399–408.
- [28] Zheng, Y., Garcia, A. L., and Alder, B. J., 2002, "Comparison of kinetic theory and hydrodynamics for Poiseuille Flow," *J. Stat. Phys.*, **109**, pp. 495–505.
- [29] Bird, G. A., 1994, *Molecular Gas Dynamics and the Direct Simulation of Gas Flows*, Clarendon, Oxford.
- [30] Allen, M. P., and Tildesley, D. J., 1987, *Computer Simulation of Liquids*, Clarendon, Oxford.
- [31] Garcia, A. L., and Wagner, W., 2000, "Time step truncation error in direct simulation Monte Carlo," *Phys. Fluids*, **12**, pp. 2621–2633.
- [32] Hadjiconstantinou, N. G., 2000, "Analysis of Discretization in the Direct Simulation Monte Carlo," *Phys. Fluids*, **12**, pp. 2634–2638.
- [33] Wagner, W., 1992, "A Convergence Proof for Bird's Direct Simulation Monte Carlo Method for the Boltzmann Equation," *J. Stat. Phys.*, **66**, pp. 1011–1044.
- [34] Garcia, A. L., and Alder, B. J., 1998, "Generation of the Chapman-Enskog Distribution," *J. Comput. Phys.*, **140**, pp. 66–70.
- [35] Bird, G. A., 1970, "Breakdown of Translational and Rotational Equilibrium in Gaseous Expansions," *Am. Inst. Aeronaut. Astronaut. J.*, **8**, p. 1998.
- [36] Trangenstein, J. A., and Pember, R. B., 1992, "Numerical Algorithms for Strong Discontinuities in Elastic-Plastic Solids," *J. Comput. Phys.*, **103**, pp. 63–89.
- [37] Hadjiconstantinou, N. G., Garcia, A. L., Bazant, M. Z., and He, G., 2003, "Statistical error in particle simulations of Hydrodynamic Phenomena," *J. Comput. Phys.*, **187**, pp. 274–297.
- [38] Alexander, F., Garcia, A. L., and Tartakovsky, D., 2002, "Algorithm Refinement for Stochastic Partial Differential Equations: I. Linear Diffusion," *J. Comput. Phys.*, **182**(1), pp. 47–66.
- [39] Hirschfelder, J. O., Curtiss, C. F., and Bird, B., 1964, *Molecular theory of gases and liquids*, Wiley, New York.
- [40] Schmidt, B., and Worner, M., 1983, "Problems with the Computation of the Shock Structure in Binary Gas Mixtures Using the Direct Simulation Monte Carlo Method," *Acta Mech.*, **1–4**, pp. 59–55.
- [41] Arora, M., and Roe, P. L., 1997, "On Postshock Oscillations Due to Shock Capturing Schemes in Unsteady Flows," *J. Comput. Phys.*, **130**, pp. 25–40.
- [42] Woodward, P. R., and Colella, P., 1984, "The Numerical Simulation of Two-dimensional Fluid Flow with Strong Shocks," *J. Comput. Phys.*, **54**, pp. 115–173.
- [43] Meshkov, E. E., 1969, "Instability of the Interface of two Gases Accelerated by a Shock Wave," *Fluid Dyn.*, **43**(5), pp. 101–104.
- [44] Meshkov, E. E., 1970, "Instability of a Shock Wave Accelerated Interface between two Gases," *NASA Tech. Trans.*, **F-13074**.
- [45] Richtmyer, R. D., 1960, "Taylor Instability in Shock Acceleration of Compressible Fluids," *Commun. Pure Appl. Math.*, **13**, pp. 297–319.
- [46] Brouillette, M., 2002, "The Richtmyer-Meshkov Instability," *Ann. Rev. Fluid Mech.*, **34**, pp. 445–468.
- [47] Holmes, R. L., Dimonte, G., Fryxell, B., Gittings, M. L., Grove, J. W., Schneider, M., Sharp, D. H., Velikovich, A. L., Weaver, R. P., and Zhang, Q., 1999, "Richtmyer-Meshkov Instability Growth: Experiment, Simulation and Theory," *J. Fluid Mech.*, **389**, pp. 55–79.