

A Multiscale Particle Approach for Continuum/Rarefied Flow Simulation

Jonathan M. Burt¹ and Iain D. Boyd²
*Department of Aerospace Engineering
University of Michigan, Ann Arbor, MI 48109*

A hybrid particle scheme is presented for the simulation of compressible gas flows involving both continuum regions and rarefied regions with strong translational nonequilibrium. The direct simulation Monte Carlo (DSMC) method is applied in rarefied regions, while remaining portions of the flowfield are simulated using a DSMC-based low diffusion particle method for inviscid flow simulation. The hybrid scheme is suitable for either steady-state or unsteady flow problems, and can simulate gas mixtures comprising an arbitrary number of species. Numerical procedures are described for strongly coupled two-way information transfer between continuum and rarefied regions, and additional procedures are outlined for the determination of continuum breakdown. The hybrid scheme is evaluated through a comparison with DSMC simulation results for a Mach 6 flow of N₂ over a cylinder, and very good overall agreement is observed.

I. Introduction

IN recent years, a number of new techniques have been developed for the simulation of compressible gas flows involving a wide range of Knudsen number regimes.¹⁻¹¹ In the flows of interest, some flowfield regions are within the low Knudsen number range for which the Euler or Navier-Stokes equations are valid, while other regions are characterized by strong thermal nonequilibrium and require a simulation method based on gas kinetic theory. Most of these simulation schemes are of the “hybrid” type, where continuum CFD methods are used in near-equilibrium regions, while considerably more expensive alternate methods are applied in nonequilibrium regions where the continuum equations fail. In such nonequilibrium regions, the shape of the local velocity distribution is not known a priori. The velocity distribution must instead be determined either by numerically solving the Boltzmann equation (or a simplified approximation of this equation) in discretized velocity space,⁹ or by applying a Monte Carlo method to a collection of particles which represent random points in velocity space.

The most mature and most frequently used simulation method for highly nonequilibrium gas flows is the direct simulation Monte Carlo (DSMC) method.¹² In DSMC, a large number of particles, each representing a large collection of atoms or molecules, are moved according to assigned velocities through a computational grid. During a simulation time step, some fraction of particles in each cell are selected to participate in binary collisions, and the velocity of each colliding particle is modified in a manner consistent with the collision term in the Boltzmann equation. The popularity of DSMC may be attributed to a number of potential advantages over other simulation methods for highly nonequilibrium gas flows, including advantages related to efficiency, memory usage, implementation of additional physics models, and parallelization.

One disadvantage however to DSMC is large statistical scatter. DSMC tends to suffer from a statistical noise problem common to Monte Carlo methods, where the scatter in cell-averaged bulk flow properties scales with the inverse square root of the number of particles per cell. This leads to difficulties when using DSMC as part of a hybrid approach involving continuum CFD methods, as information transfer from DSMC to the continuum flow solver across rarefied/continuum boundaries tends to require a much lower level of scatter in DSMC cell properties than is typically found through instantaneous averaging among all particles in a cell. Accurate simulation therefore requires some form of scatter reduction, either by using a very large number of particles per cell,¹ by utilizing computationally expensive “ghost cells” near rarefied/continuum boundaries,²⁻³ or through a weak coupling between DSMC and CFD domains where scatter is reduced by averaging DSMC cell properties over a large number of time steps.⁵⁻⁷ While successful hybrid simulations have been demonstrated using all of these approaches, each has major

¹ Post-doctoral research fellow, AIAA member.

² Professor, AIAA associate fellow.

drawbacks related to computational efficiency, complexity or the ability to simulate unsteady flows, and all require implementation and synchronization of two separate methods for continuum and rarefied flow simulation.

Ideally, a hybrid simulation technique utilizing DSMC would allow for strong two-way coupled information transfer across rarefied/continuum boundaries during each time step, and would avoid the complexity associated with using two independent methods for different flowfield regions. With these goals in mind, some authors have proposed hybrid simulation approaches which use DSMC type particles throughout the simulation domain.^{10-11,13-14} In these “all particle” hybrid algorithms, macroscopic flow properties are calculated in both rarefied and continuum regions by averaging quantities among representative particles, and DSMC collision calculations are replaced in continuum regions by procedures which involve redistributing particle velocities under an assumption of local thermal equilibrium. Particle velocities are redistributed either through the use of a DSMC collision limiter scheme,¹³⁻¹⁴ or through direct resampling from a Maxwellian distribution at the cell-averaged temperature and bulk velocity.¹⁰⁻¹¹ While such hybrid techniques are far easier to implement and provide a simple solution to the interface scatter problem in hybrid CFD-DSMC algorithms, these techniques are prone to large errors associated with numerical diffusion.

In simulations of continuum flows near the low Knudsen number limit, the ratio of the local mean free path to the computational cell size typically approaches zero. Random molecular motion in this type of flow is completely suppressed on length scales comparable to the cell size, and on these macroscopic scales individual molecules tend to move along smooth trajectories following the gas streamlines. In contrast, representative particles in DSMC-based equilibrium particle methods exhibit random motion on length scales comparable to the cell size, so that the cell size takes on properties of the local mean free path and becomes a scaling factor for numerical transport coefficients. Because the influence of collisions is neglected in fluxes based on particle motion between adjacent cells, the ratio of numerical transport coefficients (viscosity, thermal conductivity, mass diffusivity) to physical transport coefficients tends to scale with the ratio of the mean free path to the cell size.¹⁵⁻¹⁶ This ratio typically approaches zero due to practical considerations of computational expense, so that effects of numerical diffusion become prominent and large errors may appear in simulation results.

One promising alternative to existing equilibrium particle methods, for use in a hybrid algorithm with DSMC, is a DSMC-based particle method recently developed by the authors.¹⁷ In this method, intended for compressible gas flows near the low Knudsen number limit, a large number of representative particles are tracked through the grid in such a way that every particle maintains a constant relative position within a network of Lagrangian cells. Each Lagrangian cell is coincident with a cell in the fixed Eulerian grid at the beginning of each time step, but moves and deforms over the time step interval according to local bulk gas properties. Particles follow the macroscopic motion of Lagrangian cells, and move along trajectories which closely approximate the gas streamlines. Random particle motion associated with thermal energy is therefore greatly suppressed. As a result, numerical diffusion errors, as well as effects of statistical scatter, are generally far smaller than in other DSMC-based equilibrium particle methods.

Simulation procedures in the new particle method, termed the low diffusion (LD) method here for convenience, are presented in a recent paper.¹⁷ These procedures are closely based on the standard DSMC method, and are relatively easy to implement in an existing DSMC algorithm. This ease of implementation, along with considerable reductions in both numerical diffusion and scatter effects relative to other equilibrium particle methods, make the LD method a promising choice for inclusion in a hybrid code with DSMC. One such implementation of a hybrid LD-DSMC code is proposed here.

In the following sections, various components of the hybrid LD-DSMC algorithm are described, and results are presented for simulations of a two-dimensional test problem used to evaluate the algorithm. First, basic steps in the LD simulation procedures are outlined, and modifications to these procedures are proposed to increase stability and allow potentially larger time steps to be used. Next, a set of procedures are outlined for strongly coupled information exchange between LD and DSMC regions, and routines are described for automatically and dynamically allocating grid cells to either region. A test case is then examined, involving a Mach 6 flow of N_2 over a cylinder, with a global Knudsen number of 0.01. To assess overall accuracy of the hybrid LD-DSMC code, hybrid simulation results are compared with results from a DSMC simulation of the same flow. Finally, conclusions are presented and ideas for future work are proposed.

II. Simulation Procedures in the Low Diffusion Particle Method

Numerical procedures in the original LD method for inviscid flow simulation are described in detail in a previous paper.¹⁷ A brief summary of the basic steps is provided here for reference. First note that, as in DSMC, all particles in an LD simulation carry information for a position, a velocity used to update the position, and a species

identification number. However, in this case each particle is also assigned a temperature and a second velocity, termed the “bulk particle velocity”, which is used to allocate momentum among all particles in each cell. The following routines are performed during each simulation time step, in place of DSMC collision calculations:

1. The density, bulk velocity and characteristic thermal speed are evaluated in each cell as functions of cell-averaged particle properties during the current time step, and these quantities are stored in the cell data structure.
2. A velocity is calculated for the Lagrangian cell face corresponding to each cell face in the fixed Eulerian grid. Lagrangian faces are modeled as massless specularly reflecting walls, with cells on either side approximated as semi-infinite reservoirs in local thermodynamic equilibrium. An expression derived from kinetic theory is solved iteratively using the secant method to find the velocity of each face.
3. Bulk velocity and thermal speed values assigned to each cell are updated, based on the contribution of momentum and energy transfer across all corresponding Lagrangian faces over the time step interval Δt .
4. All particles in each cell are assigned new bulk velocity and temperature values based on cell quantities.
5. Velocities used for particle movement are updated, in such a way that all particles maintain a constant relative position in a Lagrangian cell over the time step interval.

Particle movement and time-averaged sampling procedures are then performed as in the DSMC method. A few additional modifications to standard DSMC procedures are required, particularly in the generation of new particles at inflow boundaries, and are described in a previous paper.¹⁷

While time step limitations in the LD method are generally governed by the same CFL condition used in explicit CFD methods, the scatter inherent in any particle method involving the division of conserved quantities (mass, momentum and energy) into discrete packets may produce fluctuations in the local stability criterion for CFL numbers near one. This is particularly true for simulations involving a very small number of particles per cell. To avoid a potentially unstable solution without requiring a reduction in the time step size, a subcycling procedure may be used here. In the subcycling procedure, steps 2 and 3 above are sequentially repeated some set number N of times during each time step, and a fraction $\Delta t/N$ of the time step interval is used in evaluations of momentum and energy transfer in step 3. For a reasonable number of particles per cell (between 20 and 40) the subcycling procedure is found to contribute roughly 20% to the total simulation time for a conservative N value of 4. Note that this procedure creates only a minimal improvement in overall accuracy, so is generally unnecessary unless instabilities are present. To reduce simulation time, no subcycling steps are performed in the LD method calculations described in this paper.

III. Information Transfer at LD-DSMC Boundaries

In the proposed hybrid algorithm, a simulated flowfield is periodically divided into continuum regions assigned to the LD method and nonequilibrium regions assigned to DSMC. Procedures related to the automatic determination of continuum breakdown are described in the next section. While effective assessment of continuum breakdown is a considerable challenge in the development of any hybrid scheme for continuum/rarefied flow simulation, accurate strongly coupled two-way transfer of flowfield information across DSMC domain boundaries may be an even greater challenge. As discussed in the introduction, the inherent scatter in DSMC makes a strongly coupled hybrid implementation difficult, and necessitates some procedure for scatter reduction. Moreover, accurate and efficient conversion of relevant flowfield quantities between DSMC and the continuum flow solver may be complex and require considerable code development. These problems are largely avoided in hybrid schemes involving DSMC-based equilibrium particle methods, as information transfer is performed simply and straightforwardly through two-way particle transport between neighboring DSMC and continuum cells. Corresponding information transfer procedures in the proposed hybrid LD-DSMC scheme are somewhat more complicated, due to differences between LD and DSMC particle properties. However, the use of similar particle methods in both rarefied and continuum regions allows for relatively simple two-way information transfer, and provides a basis for a strongly coupled hybrid algorithm suitable for application to either steady state or unsteady flow problems.

In the LD-DSMC hybrid simulation procedures, all cells in the fixed computational grid are periodically assigned one of four integer values, corresponding to the DSMC domain, the LD domain, or one of two buffer regions. These buffer regions, termed regions A and B, are used for information exchange between the DSMC and LD methods. Each buffer region is two cells thick, and is positioned along the boundary between DSMC and LD domains, such that a line drawn from a DSMC cell to a nearby LD cell will pass through at least two cells in buffer region A, then

through two cells in region B. Standard DSMC collision calculations are performed in all cells within region A, while all cells in region B are included in LD method calculations.

In step 2 of the LD simulation procedures listed above, a velocity is determined for the Lagrangian cell face corresponding to each fixed Eulerian cell face. This velocity is a function of the density, bulk velocity and characteristic thermal speed in the cells on either side of the face. Lagrangian face velocities are calculated for all faces of each cell in buffer region B, including any face bordering a cell in region A. Cell properties must therefore be evaluated for cells in region A, as given in step 1 above, to properly determine Lagrangian face velocities for any face between regions A and B. While the determination of cell properties in region A could be carried out in a similar manner to step 1 in the LD calculations, the large scatter associated with instantaneous cell-averaged properties in DSMC may introduce considerable unphysical fluctuations in flowfield properties computed using the LD method. To reduce these fluctuations without requiring an excessive number of particles per cell within buffer region A, a sub-relaxation procedure²¹ is used here.

As a first step in the sub-relaxation procedure, the instantaneous density ρ_{cell}^* and mass-averaged bulk velocity $\mathbf{U}_{\text{cell}}^*$ are calculated in each cell within region A, using all particles in the cell during the current time step. These calculations are performed during each time step, in place of step 1 in the LD procedures. Note however that in region A the cell-based bulk velocity is a function of velocities used for DSMC particle movement, not bulk particle velocities used to assign momentum in the LD method. Next, an instantaneous cell-averaged temperature T_{cell}^* is calculated:

$$T_{\text{cell}}^* = \frac{1}{(3 + \langle \zeta \rangle) Ru} \left[2 \langle E_{\text{int}} \rangle + \left(\frac{N_p}{N_p - 1} \right) \langle MW \rangle \left(\langle \mathbf{U} \cdot \mathbf{U} \rangle_{\text{m}} - \langle \mathbf{U} \rangle_{\text{m}} \cdot \langle \mathbf{U} \rangle_{\text{m}} \right) \right] \quad (1)$$

Here Ru is the universal gas constant, N_p is the number of particles in the cell, and \mathbf{U} , MW , ζ , and E_{int} are respectively the velocity, molecular weight, number of internal (rotational and vibrational) degrees of freedom, and internal energy (per kmol) for an individual particle. The operator $\langle \rangle$ denotes an unweighted average among all particles in the cell, while $\langle \rangle_{\text{m}}$ denotes a mass-weighted average such that $\langle \mathbf{U} \rangle_{\text{m}} = \mathbf{U}_{\text{cell}}^*$.

Once T_{cell}^* is known for a given cell, the reciprocal β_{cell}^* of the most probable thermal speed for an equilibrium gas is calculated from the instantaneous cell temperature.

$$\beta_{\text{cell}}^* = \sqrt{\frac{\langle MW \rangle}{2 Ru T_{\text{cell}}^*}} \quad (2)$$

Finally, density, velocity and thermal speed scale values stored in the cell data structure are updated as

$$\begin{aligned} \rho_{\text{cell}} &= \theta \rho_{\text{cell}}^* + (1 - \theta) \rho_{\text{cell}}^{n-1} \\ \mathbf{U}_{\text{cell}} &= \theta \mathbf{U}_{\text{cell}}^* + (1 - \theta) \mathbf{U}_{\text{cell}}^{n-1} \\ \beta_{\text{cell}} &= \theta \beta_{\text{cell}}^* + (1 - \theta) \beta_{\text{cell}}^{n-1} \end{aligned} \quad (3)$$

where the superscript $n-1$ denotes a value at the previous time step and θ is a relaxation factor. The density, velocity and thermal speed scale calculated in Eqs. (3) are then used to determine the Lagrangian face velocity, as described in a previous paper,¹⁷ for any cell face in buffer region B bounded on one side by a cell in buffer region A. Note that larger values of θ correspond to a greater weight for cell quantities during the current time step, which reduces the time lag in averaged values but potentially increases scatter associated with instantaneous fluctuations in cell quantities. For a reasonable balance between these two effects, a θ value of 0.01 is used here. This may be assumed to accurately represent flowfield properties in an unsteady flow so long as characteristic transient time scales are much greater than $1/\theta = 100$ time steps. Larger θ values may be required in simulations of highly unsteady flows.

During each time step, following either DSMC collision routines or the LD method calculations listed above, all particles in both LD and DSMC domains are moved using the same set of procedures. In these movement procedures, the position vector for each particle is updated by the product of the particle velocity \mathbf{U} and the time step

interval Δt . Immediately before the movement routines, every particle in both buffer regions A and B is cloned, so that two particles are located at each occupied position in any buffer cell. In buffer region A, where all particles are initially of the DSMC type, each newly generated clone is treated as an LD type particle, and is assigned a new velocity used for movement procedures and a bulk velocity used to allocate momentum. Both velocities are equal to the cell velocity \mathbf{U}_{cell} as calculated in Eqs. (3). Each clone in region A is also assigned a temperature T_{cell} , which is computed as a function of β_{cell} and the cell-averaged molecular weight $\langle \text{MW} \rangle$.

$$T_{\text{cell}} = \frac{\langle \text{MW} \rangle}{2Ru(\beta_{\text{cell}})^2} \quad (4)$$

In buffer region B, where all particles are initially of the LD type, newly generated clones are regarded as DSMC particles, and are assigned values consistent with an equilibrium velocity distribution at the cell-averaged bulk velocity and temperature. The assigned temperature of each clone is set to zero, and velocity and internal energy values are determined through probabilistic sampling procedures based on expressions in Appendix C of Bird.¹² Given two random numbers R_1 and R_2 in the range (0,1] a clone particle velocity component U_i is calculated as

$$U_i = \mathbf{U}_{\text{cell}} \cdot \mathbf{n}_i + \frac{\sin(2\pi R_1)}{\beta_{\text{cell}}} \sqrt{-\frac{\langle \text{MW} \rangle}{\text{MW}} \ln R_2} \quad (5)$$

where \mathbf{U}_{cell} and β_{cell} are the instantaneous cell-averaged bulk velocity and thermal speed scale as calculated in step 1 of the LD simulation procedures, \mathbf{n}_i is the unit vector in the i direction, $\langle \text{MW} \rangle$ is the cell-averaged molecular weight, and MW is the molecular weight for the assigned particle species. (Note in Eq. (5) that, for a mixture, the thermal speed scale used in the particle velocity distribution is a function of the species molecular weight and the cell temperature. The cell temperature is in turn a function of the cell-averaged molecular weight and the thermal speed scale β_{cell} stored in the cell data structure.) For a diatomic gas species with no vibrational excitation, the corresponding internal energy is computed as

$$E_{\text{int}} = -\frac{\langle \text{MW} \rangle}{(\beta_{\text{cell}})^2} \ln R_3 \quad (6)$$

where R_3 is another random number in (0,1].

Immediately following particle movement procedures during each time step, any LD type particle remaining in buffer region A is removed from the simulation, and any DSMC type particle in region B is also removed. (LD and DSMC particles may be distinguished because only LD type particles have a nonzero temperature.) Thus, only the newly generated LD particles which have moved into cells assigned to LD cells or buffer region B are retained, and any LD particles which pass from region B into a cell in region A are destroyed. Likewise, only DSMC type clone particles which move into DSMC cells or buffer region A are retained, and DSMC particles which pass from region A into region B are destroyed as well. While these procedures only enforce mass, momentum and energy conservation in an average sense, they provide a relatively simple and strongly coupled means of exchanging information between rarefied and continuum flowfield regions in a hybrid simulation.

IV. Continuum Breakdown and Cell Assignment

Proper allocation of cells to LD or DSMC domains depends on accurate determination of continuum breakdown. This determination is itself an area of current research, and existing techniques for assessing continuum breakdown have found only moderate success. In general, such techniques are either intended only for limited types of flows and certain modes of nonequilibrium,¹⁸⁻¹⁹ or are very computationally expensive and inappropriate for use in simulation methods which include near-equilibrium assumptions.²⁰ For the hybrid algorithm presented here, a relatively simple and conservative technique is therefore proposed.

First, a gradient length local Knudsen number¹⁸ Kn_{GLL} is periodically (once every several hundred time steps) evaluated for each cell in the grid. The nondimensional parameter Kn_{GLL} is calculated as the ratio of the local mean free path to a length scale based on the density gradient, which in turn is given by the ratio of the gas density to the

magnitude of the density gradient.¹⁹ To reduce scatter and avoid spurious determination of continuum breakdown, cell density values ρ are averaged over a large number of time steps using a similar sub-relaxation procedure to that employed in Eqs. (3). Wang and Boyd¹⁹ found that representative hypersonic flows may be assumed sufficiently close to equilibrium to justify near-equilibrium assumptions when $\max\{\text{Kn}_{\text{GLL}}\} < 0.05$, where gradient length Knudsen numbers based on density, temperature and bulk velocity are all considered. As the LD method assumes an inviscid gas at the low Knudsen number limit, and as only the Kn_{GLL} value based on gas density is considered here, we use a more conservative cutoff value of 0.02 to evaluate continuum breakdown.

It should be emphasized that any continuum breakdown parameter based on gradients in macroscopic flow properties is at best marginally accurate in determining the general suitability of continuum flow simulation techniques. The density-based parameter used here is particularly problematic near the downstream edge of strong shocks, and in low density subsonic regions where considerable translational and translational-rotational nonequilibrium may exist. While the first problem may be addressed (as described below) by extending the DSMC domain by some set number of cells in all directions, the second problem requires a modification to the condition for continuum breakdown. In this case the DSMC method should be used in regions where the mean free path λ is large, regardless of the magnitude of any macroscopic gradients. A modified continuum breakdown condition

$$\max\left\{\frac{\lambda}{\rho}|\nabla\rho|, \frac{\lambda}{L}\right\} > 0.02 \quad (7)$$

is therefore used here, where L is a global characteristic length scale based on the boundary geometry. While there is little physical basis behind this modification, and the use of the length scale L is to some extent arbitrary, this does provide a simple and effective means of assigning to DSMC any region of sufficiently low density for DSMC calculations to be reasonably efficient.

In the continuum breakdown calculations performed periodically during a simulation, all cells are initially assigned to either LD or DSMC regions based on (7). Next, any DSMC cell within two cells of a LD cell is itself reassigned as an LD cell. This is done sequentially, by reassigning all LD cells which border a DSMC cell, then repeating the procedure an additional time. Any LD cell within ten cells of a DSMC cell is then converted to a DSMC cell. These reassignment procedures are performed with two intentions: First, the DSMC domain is extended on average eight cells beyond the boundary at which continuum breakdown is initially assumed. This reduces the probability that significant nonequilibrium occurs beyond the boundary, and allows some room for movement of this boundary in an unsteady flow. Second, any very small region initially assigned to either the LD or DSMC methods is removed. As a result, the boundary between LD and DSMC domains is considerably smoother and less prone to scatter effects.

As a next step in the allocation of cells to different regions, any LD cell within two cells of a DSMC cell is assigned to buffer region A. Any LD cell within two cells of this buffer region is then assigned to buffer region B. Finally, all LD type particles in cells newly assigned to either DSMC or region A are converted to DSMC type particles, using the probabilistic sampling procedures described above, and all DSMC type particles cells within the LD region or buffer region B are converted to LD type particles.

V. Evaluation of the Hybrid Algorithm

All hybrid simulation procedures outlined above have been implemented in a modified version of the DSMC code MONACO.²² The resulting hybrid LD-DSMC code allows for parallel simulations of two-dimensional planar or axisymmetric flows, and is intended for use with either non-Cartesian structured grids or mixed structured/unstructured grids employing quadrilateral cells in continuum regions. Features include adaptive parallel domain decomposition and non-Cartesian subcells for nearest neighbor collision selection in DSMC.

As a preliminary test case to demonstrate accuracy and proper implementation of the hybrid procedures, we consider a hypersonic blunt body flow problem previously used in hybrid CFD-DSMC simulations of Schwartztruber and Boyd.⁶ This case involves a flow of N_2 over a cylinder, with a freestream Mach number of 6.0 and a global Knudsen number of 0.01. The freestream gas temperature is 217.45 K and the freestream number density is $1.6 \times 10^{21} \text{ m}^{-3}$, which roughly corresponds to atmospheric conditions at an altitude of 70 km. To minimize errors associated with the lack of viscous terms in the LD method, a specularly reflecting boundary condition is used at the cylinder surface. Also, to reduce errors related to the lack of distinction in the LD method between thermal energy in translational and internal modes, the DSMC rotational collision number (i.e. the probability of rotational-translational energy exchange during a collision) is set to one, and vibrational energy excitation is neglected.

A hybrid LD-DSMC simulation is run in parallel on eight AMD Opteron processors, using a structured grid with 45,000 cells adapted so that the cell size is no larger than twice the local mean free path. DSMC subcells are used to avoid the selection of collision pairs separated by more than one-half the mean free path. A constant time step interval of 5×10^{-8} s is used, in line with standard DSMC guidelines,¹² and the simulation is run for 80,000 time steps. The last 50,000 time steps are used for sampling following convergence to steady state conditions. Approximately 4.2 million particles are tracked through the grid at steady state, and the total simulation time is about 5.9 hours. For comparison, a second simulation is run using only DSMC, and requires about 5.1 hours on the same number of processors.

Note that, while the hybrid simulation is around 16% more expensive than the DSMC simulation, the focus of the comparison here is on assessing the accuracy of the hybrid algorithm. By running hybrid and DSMC simulations using identical grids and input values, we can evaluate the accuracy of the hybrid scheme independent of any approximations related to cell size, time step or other numerical parameters. The LD method allows for cell sizes and time step intervals potentially several orders of magnitude larger than DSMC, as similar criteria are used to determine sufficient time step and grid refinement in the LD method as in an explicit finite volume solver for the compressible Euler equations.¹⁷ Thus, an ideal comparison of computational efficiency between the hybrid scheme and DSMC would require the use of automatic grid adaptation procedures during a hybrid simulation, so that cells in the LD domain would be far larger than those in the DSMC domain. Such a comparison is out of the scope of the present work, but the reader is referred to a previous paper¹⁷ for a discussion of computation efficiency in the LD method.

The grid boundary geometry is shown in Fig. 1, along with streamlines and boundaries of the LD and DSMC domains at the end of the sampling period in the hybrid simulation. As expected, a detached bow shock is captured in the DSMC domain, and the low density wake is also assigned to DSMC, while both freestream and high density post-shock regions are assigned to the LD method. Blue areas along the boundary between LD and DSMC domains denote cells within buffer region A, and neighboring orange areas indicate cells within buffer region B.

Figure 2 shows contours of bulk velocity magnitude for both the hybrid and DSMC simulations. Iso-contour lines for the hybrid simulation are colored according to the contour value, while lines for the DSMC simulation are black. Very good overall agreement is observed between the two results. In particular, the shock standoff distance and shock angles are nearly identical, and there is almost no discernable difference in the contour lines within the post-shock forebody region and the wake. However, noticeable disagreement is found in the post-shock region far from the cylinder, where the hybrid simulation slightly underestimates the velocity magnitude, and near the point of maximum thickness at the cylinder surface, where there is a small overshoot in the hybrid simulation result. While some significant differences are observed here, the level of disagreement in local velocity values is generally very small. The largest such difference occurs in the interior of the shock, where a roughly 0.7% overshoot in the shock standoff distance from the hybrid simulation corresponds to a maximum 8.2% difference in velocities, due to the extremely large gradients in this very narrow region. Excluding the interior of the shock, the maximum velocity difference anywhere in the flowfield is approximately 0.8%, so the overall level of agreement may be considered good.

Pressure contours from both simulations are shown in Fig. 3. The overall agreement here is excellent, although a few noticeable differences are found in the wake, along the line of maximum pressure just downstream of the shock, and at the outflow boundary on the right side of the grid. Still, most contour lines for the two simulations are nearly indistinguishable within the resolution of the figure.

Contours of gas density are shown in Fig. 4. The level of agreement observed in the figure is relatively good, although inferior to that found in the pressure contours of Fig. 3. The density is noticeably greater in the hybrid simulation results within the post-shock forebody region and in the wake region far downstream of the cylinder. The maximum difference in local density between the two simulations is approximately 1.0% in the post-shock forebody region, and about 0.3% in the wake.

Figure 5 shows contours of translational temperature. Note that only a single temperature is used in LD method calculations, while translational and rotational temperatures may differ in regions assigned to DSMC. Very good agreement is found near the cylinder surface, just downstream of the shock, and throughout most of the wake region. Note in particular the nearly identical temperature spikes in the shock, which are associated with finite rate collision processes and a rotational temperature lag in DSMC shock structure calculations. Some temperature overshoot is observed however in hybrid simulation contours within the post-shock region far from the cylinder. Excluding the interior of the shock, the maximum local temperature difference between the two simulations is found in this farfield post-shock region, and is approximately 2.1%.

To demonstrate the presence of rotational-translational nonequilibrium, a contour plot for rotational temperature is included as Fig. 6. In comparing Figs. 5 and 6, some significant differences are observed. These differences are

most prevalent in regions of low density or large gradients which are assigned to DSMC in the hybrid simulation. In particular, the post-shock spike in translational temperature is absent in the rotational temperature contours, and the local temperature maximum along the centerline in the wake is considerably lower and further downstream for rotational temperature. A close examination of DSMC simulation contour lines in Figs. 5 and 6 also shows significant differences between translational and rotational temperatures within continuum regions assigned to the LD method in the hybrid simulation. (These temperature differences are reduced, but not eliminated, by employing a rotational collision number of one.) As these differences are neglected in LD method calculations, some disagreement is expected here between DSMC and hybrid simulation results. Despite any errors associated with rotational-translational nonequilibrium in continuum regions, good overall agreement is found between the two results shown in Fig. 6. The maximum rotational temperature difference outside of the shock is about 1.5%, and occurs in the post-shock region far from the cylinder.

Profiles of Mach number and pressure along the stagnation streamline are shown in Fig. 7. The shaded portion of the figure denotes the region surrounding the shock which is assigned to DSMC in the hybrid simulation. This DSMC region is found, as expected, to successfully capture areas of large gradients within the shock. However, some disagreement is observed between Mach number profiles from the two simulations near the upstream edge of the shock. The discrepancy is likely due to the lack of viscous transport in the LD method, and could be removed by either reducing the cutoff value for continuum breakdown in (7) or by extending the DSMC domain by a larger number of cells during the cell assignment procedures in the hybrid simulation. Aside from this difference, excellent agreement is observed between the two Mach number curves in Fig. 7. As mentioned in the discussion of Fig. 2, a small overshoot in the shock standoff distance occurs in the hybrid simulation, which leads to some noticeable difference in the Mach number profiles within the shock. Near perfect agreement (within the level of DSMC statistical scatter) is found in Mach number values downstream of the shock. Similar results are shown for the pressure curves, with the large gradient region within the shock appearing slightly further upstream in the hybrid simulation, and with excellent agreement in the post-shock region.

Figure 8 shows profiles of density, translational temperature and rotational temperature along the stagnation streamline. While relatively good overall agreement is observed here for all three quantities, the disparities between hybrid and DSMC results are significantly higher than those shown in Fig. 7. All three hybrid simulation curves in Fig. 8 display a shock positioned slightly upstream of that in the DSMC simulation. A small (approximately 1.3%) unphysical increase in density is observed at the LD-DSMC interface along the downstream edge of the shock, with a corresponding drop in translational temperature. While there are several different possible causes for this error, the most probable cause is associated with homogeneous flow approximations in the generation of LD type clone particles in cells within buffer region A. As described above, these clone particles are assigned velocities used for movement procedures which are equal to a cell-averaged bulk velocity. The use of identical velocities for all LD type clones in each cell provides a very simple and efficient means of transferring information to the LD domain, but neglects the influence of local density gradients. In regions of strong compression where LD cells appear downstream of DSMC cells, the use of a cell-averaged velocity for clone particle movement results in overly large velocities for clone particles near the LD-DSMC interface, and tends to overestimate the mass flux between neighboring LD and DSMC cells. An unphysical density increase appears as a result of the overestimate in mass flux, as is observed in the figure. Note that this error is found only at the downstream edge of the shock in the forebody region near the stagnation streamline. Further outward from the centerline the post-shock density gradient is far smaller, so any errors due to such homogeneous flow assumptions become negligible.

In addition to approximations used in the generation of clone particle properties, several other potential error sources likely contribute to the disagreement between DSMC and hybrid simulation results observed in Figs. 2 through 8. Other possible error sources include a lack of exact mass, momentum and energy conservation in LD-DSMC coupling procedures, which may lead to random walk effects among various flow quantities. As discussed above, discrepancies between rotational and translational temperature in DSMC calculations within continuum regions may lead to additional disagreement with hybrid simulation results. One major difference between the DSMC and hybrid simulations is the lack of consideration for viscous transport in the LD method, which likely accounts for some of the disparity found in continuum flow regions. This seems like the most probable cause for differences in velocity contours shown in Fig. 2, particularly near the cylinder surface where large transverse velocity gradients exist.

VI. Conclusions and Future Work

A hybrid particle scheme for the simulation of multiscale compressible gas flows has been presented. The DSMC method is used in rarefied regions, and a new DSMC-based low diffusion particle method for inviscid flow

simulation is employed in continuum regions. The use of similar particle methods in both rarefied and continuum regions allows for a comparatively simple algorithm to simulate flows involving a wide range of Knudsen number regimes. Strong two-way coupling between the two methods permits simulation of either steady state or unsteady flows, and numerical procedures are generalized for use with polyatomic gases or mixtures involving a number of different species. The proposed simulation framework is considerably simpler to implement, given only an existing DSMC code, than coupled CFD-DSMC hybrid approaches, and avoids the numerical diffusion problems inherent in existing hybrid schemes which utilize DSMC-based particle methods in continuum regions.

A two dimensional hypersonic flow of N_2 around a cylinder was used as a test case to evaluate the hybrid LD-DSMC algorithm, through a comparison of DSMC and hybrid simulation results. This flow problem includes a number of complex phenomena, and may be considered a particularly challenging case for use in assessing hybrid simulation accuracy. Phenomena observed here include development of a strong detached shock, subsonic regions within both rarefied and continuum domains, strong expansion and rarefied compression, a translational temperature spike and other trends associated with internal shock structure. While some noticeable differences were observed between results from DSMC and hybrid simulations, the overall level of agreement was very good. Shock locations were nearly identical for the two simulations, and excluding the interior of the shock, all flowfield quantities considered were found to differ by no more than about 2%.

The work presented here is intended to represent progress toward a more general simulation scheme for high speed, low density flows. Additional modeling capabilities should be incorporated into the numerical procedures to extend applicability or increase accuracy, and further evaluation of the method is necessary to properly assess simulation efficiency. One area of current work is the implementation of viscous transport in the LD method, for simulating near-equilibrium viscous flows within the validity range of the compressible Navier-Stokes equations. The kinetic theory basis of the LD method allows for simple implementation of wall velocity and temperature slip models, in a similar manner to kinetic flux vector splitting.²³ The addition of viscous flow capabilities also permit the LD method to be applied to low Knudsen number regimes where viscous effects are important but DSMC is prohibitively expensive.

Another intended future capability of this work is the simulation of continuum flows involving nonequilibrium distributions in internal energy modes. As rotational and (particularly) vibrational energy accommodation processes tend to occur over time scales much larger than the mean collision time, there is some potential in a variety of flow problems for these energy modes to experience significant nonequilibrium even when the local Knudsen number is very small. The assignment of internal energy values to LD particles should allow for the same capabilities for arbitrary rotational and quantized vibrational energy distributions currently available in DSMC, while also permitting the use of similarly large time steps and cell sizes as explicit CFD methods.

Other intended areas of future work include the characterization of hybrid code efficiency, using much larger cells in continuum regions, and implementation of procedures for automatic grid adaptation. We also plan to evaluate hybrid simulation accuracy for additional types of flows. Of particular interest for this project are unsteady flow problems related to high altitude rocket exhaust flows and plume impingement.

Acknowledgments

Financial support for this work was provided by Spectral Sciences, Inc. through contract W9113M-06-C-0122, a Missile Defense Agency (MDA) Small Business Innovative Research (SBIR) Phase II award. The authors gratefully acknowledge Jason Cline and Matt Braunstein for their oversight of this work.

References

- ¹Wijesinghe, H. S., Hornung, R. D., Garcia, A. L., and Hadjiconstantinou, N. G., "Three-dimensional Hybrid Continuum-Atomistic Simulations for Multiscale Hydrodynamics," *Journal of Fluids Engineering*, Vol. 126, 2004, pp. 768-777.
- ²Roveda, R., Goldstein, D. B., and Varghese, P. L., "Hybrid Euler/Particle Approach for Continuum/Rarefied Flows," *Journal of Spacecraft and Rockets*, Vol. 35, No. 3, 1998, pp. 258-265.
- ³Roveda, R., Goldstein, D. B., and Varghese, P. L., "Hybrid Euler/Direct Simulation Monte Carlo Calculation of Unsteady Slit Flow," *Journal of Spacecraft and Rockets*, Vol. 37, No. 6, 2000, pp. 753-760.
- ⁴Wang, W., and Boyd, I. D., "Hybrid DSMC-CFD Simulations of Hypersonic Flow Over Sharp and Blunted Bodies," AIAA Paper 2003-3644, 2003.
- ⁵Schwartzentruber, T. E., and Boyd, I. D., "A Hybrid Particle-Continuum Method Applied to Shock Waves," *Journal of Computational Physics*, Vol. 215, 2006, pp. 402-416.
- ⁶Schwartzentruber, T. E., and Boyd, I. D., "A Modular Particle-Continuum Numerical Method for Hypersonic Non-equilibrium Gas Flows," *Journal of Computational Physics*, 2007 (article in press).

⁷Wadsworth, D. C., and Erwin, D. A., "Two-Dimensional Hybrid Continuum/Particle Approach for Rarefied Flows," AIAA Paper 92-2975, 1992.

⁸Garcia, A. L., Bell, J. B., Crutchfield, W. Y., and Alder, B. J., "Adaptive Mesh and Algorithm Refinement Using Direct Simulation Monte Carlo," *Journal of Computational Physics*, Vol. 154, 1999, pp. 134-155.

⁹Kolobov, V. I., Arslanbekov, R. R., Aristov, V. V., Frolova, A. A., and Zabelok, S. A., "Unified Solver for Rarefied and Continuum Flows with Adaptive Mesh and Algorithm Refinement," *Journal of Computational Physics*, Vol. 223, 2007, pp. 589-608.

¹⁰Tiwari, S., and Klar, A., "An Adaptive Domain Decomposition Procedure for the Boltzmann and Euler Equations," *Journal of Computational and Applied Mathematics*, Vol. 90, 1998, pp. 223-237.

¹¹Macrossan, M. N., "A Particle-Only Hybrid Method for Near-Continuum Flows," *Rarefied Gas Dynamics: 22nd International Symposium*, American Institute of Physics, 2001, pp. 388-395.

¹²Bird, G. A., *Molecular Gas Dynamics and the Direct Simulation of Gas Flows*, Clarendon Press, Oxford, 1994.

¹³Bartel, T. J., Sterk, T. M., Payne, J. L., and Preppernau, B., "DSMC Simulation of Nozzle Expansion Flow Fields," AIAA Paper 94-2047, 1994.

¹⁴Titov, E. V., Zeifman, M. I., and Levin, D. A., "Application of the Kinetic and Continuum Techniques to the Multi-Scale Flows in MEMS Devices," AIAA Paper 2005-1399, 2005.

¹⁵Breuer, K. S., Piekos, E. S., and Gonzales, D. A., "DSMC Simulations of Continuum Flows," AIAA Paper 95-2088, 1995.

¹⁶Macrossan, M. N., "The Equilibrium Flux Method for the Calculation of Flows with Non-equilibrium Chemical Reactions," *Journal of Computational Physics*, Vol. 80, 1989, pp. 204-231.

¹⁷Burt, J. M., and Boyd, I. D., "A Low Diffusion Particle Method for Simulating Compressible Inviscid Flows," AIAA Paper 2007-4259, 2007.

¹⁸Boyd, I. D., Chen, G., and Candler, G. V., "Predicting Failure of the Continuum Fluid Equations in Transitional Hypersonic Flows," *Physics of Fluids*, Vol. 7, No. 1, 1995, pp. 210-219.

¹⁹Wang, W., and Boyd, I. D., "Predicting Continuum Breakdown in Hypersonic Viscous Flows," *Physics of Fluids*, Vol. 15, No. 1, 2003, pp. 91-100.

²⁰Camberos, J. A., Schrock, C. R., McMullan, R. J., and Branam, R. D., "Development of Continuum Onset Criteria with Direct Simulation Monte-Carlo Using Boltzmann's H-Theorem: Review and Vision," AIAA Paper 2006-2942, 2006.

²¹Sun, Q., and Boyd, I. D., "Evaluation of Macroscopic Properties in the Direct Simulation Monte Carlo Method," *Journal of Thermophysics and Heat Transfer*, Vol. 19, No. 3, 2005, pp. 329-335.

²²Dietrich, S., and Boyd, I. D., "Scalar and Parallel Optimized Implementation of the Direct Simulation Monte Carlo Method," *Journal of Computational Physics*, Vol. 126, 1996, pp. 328-342.

²³Chou, S. Y., and Baganoff, D., "Kinetic Flux-Vector Splitting for the Navier-Stokes Equations," *Journal of Computational Physics*, Vol. 130, 1997, pp. 217-230.

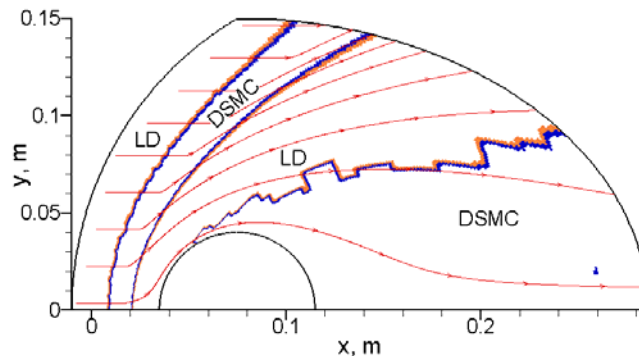


Figure 1. Grid geometry, streamlines and boundaries between LD and DSMC regions in the hybrid simulation.




Figure 2. Contours of bulk velocity magnitude from DSMC and hybrid simulations. DSMC contour lines are shown in black. Values are in m/s.




Figure 3. Pressure contours from DSMC and hybrid simulations. Values are in Pa.




Figure 4. Contours of gas density. Values are in kg/m^3 .




Figure 5. Contours of translational temperature. Values are in K.




Figure 6. Contours of rotational temperature. Values are in K.




Figure 7. Profiles of Mach number and pressure along the stagnation streamline.

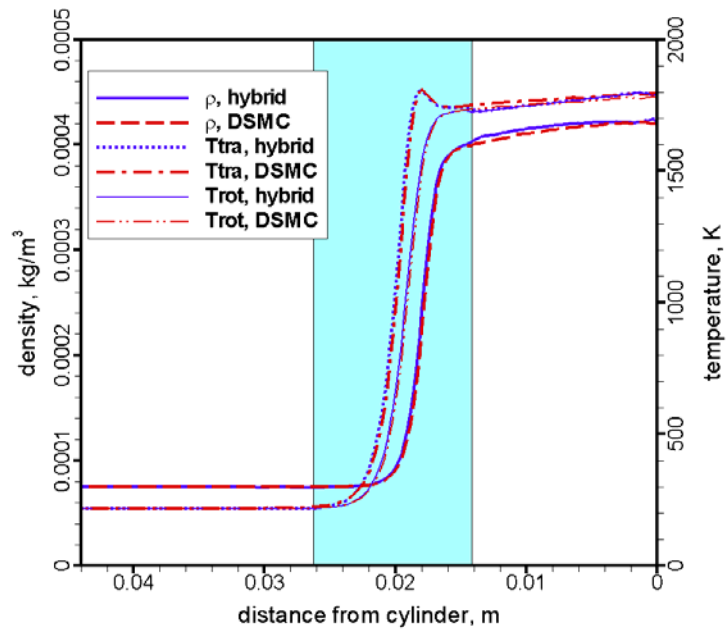


Figure 8. Density, translational temperature and rotational temperature profiles along the stagnation streamline.