Journal of Computational Physics 228 (2009) 460–475

Contents lists available at ScienceDirect

Journal of Computational Physics

journal homepage: www.elsevier.com/locate/jcp

A hybrid particle approach for continuum and rarefied flow simulation

Jonathan M. Burt *, Iain D. Boyd

Department of Aerospace Engineering, University of Michigan, 1320 Beal Ave., Ann Arbor, MI 48109, United States

ARTICLE INFO

Article history:
Received 4 June 2008
Received in revised form 18 September 2008
Accepted 24 September 2008
Available online 10 October 2008

Keywords:
Monte Carlo methods (65C05)
Particle methods (76M28)
Direct simulation Monte Carlo

ABSTRACT

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 N_2 over a cylinder, and good overall agreement is observed. Large potential efficiency gains (over three orders of magnitude) are estimated for the hybrid algorithm relative to DSMC in a simple example involving a rarefied expansion flow through a small nozzle into a vacuum chamber.

© 2008 Elsevier Inc. All rights reserved.

1. 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.

1.1. Hybrid CFD–DSMC approaches

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.

* Corresponding author. Tel.: +1 734 255 0464.
E-mail address: burtj@umich.edu (J.M. Burt).

0021-9991/$ - see front matter © 2008 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcp.2008.09.022

One disadvantage to DSMC, however, 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 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.

1.2. Hybrid simulations using equilibrium particle methods

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 than coupled CFD–DSMC algorithms and provide a simple solution to the interface scatter problem in CFD–DSMC algorithms, these all-particle hybrid techniques are prone to large errors associated with numerical diffusion.

Large numerical diffusion errors in all-particle hybrid simulations result from an extreme sensitivity to computational cell size in the DSMC-based equilibrium particle methods employed within continuum regions. This sensitivity to cell size may be understood through the following explanation: in simulations of continuum flows near the low Knudsen number limit, the ratio of the local mean free path to the cell size typically approaches zero. Random molecular motion in this type of flow is almost completely suppressed on length scales comparable to the cell size, and on these macroscopic scales individual molecules tend to move along smooth trajectories which approximately follow 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 and 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.

1.3. The low diffusion particle method

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 simulating 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.

1.4. Outline of the paper

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 sim-

ulation 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 is 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 representative hypersonic flow problem 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. An additional test case is then used to demonstrate applicability and efficiency of the hybrid scheme for a rarefied expansion flow. Finally, conclusions are presented and ideas for future work are proposed.

2. Simulation procedures in the low diffusion particle method

While 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 an LD method simulation 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 procedures are performed during each time step in an LD method simulation, 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 Lagrangian 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.

3. 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 and the additional momentum and energy transport through Lagrangian cell faces in the LD method. 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 at least two cells in region B. The relative locations of these buffer regions are illustrated in Fig. 1. Standard DSMC collision calculations are performed in all cells within region A, while all cells in region B are included in LD method calculations.

3.1. Determination of flow properties in buffer cells

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 procedure [21] is used here.

As a first step in the sub-relaxation procedure, the instantaneous density ρ^*_{cell} and mass-averaged bulk velocity U^*_{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 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:

$$T^*_{\text{cell}} = \frac{1}{(3+\langle\zeta_i\rangle)R_u}\left[2\langle E_i\rangle + \left(\frac{N_p}{N_p-1}\right)\langle \text{MW}_i\rangle(\langle U_i\cdot U_i\rangle_m - \langle U_i\rangle_m\cdot\langle U_i\rangle_m)\right] \tag{1}$$

Here R_u is the universal gas constant, N_p is the number of particles in the cell, i is the particle index, and $\mathbf{U}_i$, MW_i, ζ_i, and E_i 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 $\langle\ \rangle_m$ denotes a mass-weighted average such that $\langle U\rangle_m = U^*_{\text{cell}}$.

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.

$$\beta^*_{\text{cell}} = \sqrt{\frac{\langle \text{MW}_i\rangle}{2R_uT^*_{\text{cell}}}} \tag{2}$$

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

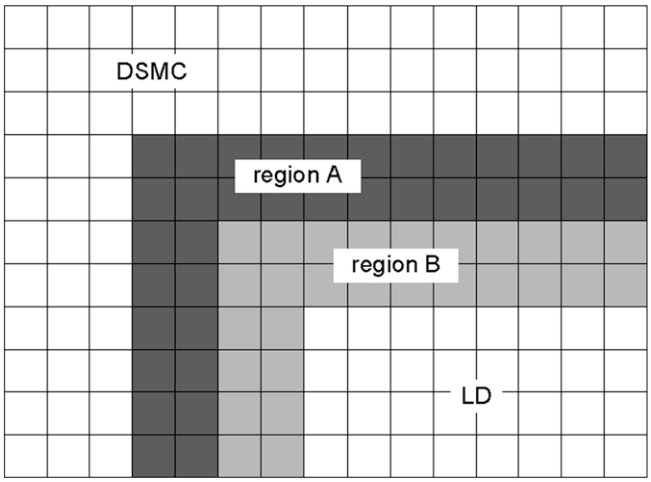

Fig. 1. Location of buffer regions in a hybrid simulation.

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

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 Eq. (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.

3.2. Particle cloning in buffer cells

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 $\mathbf{U}_i$ 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 U_{cell} as calculated in Eq. (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 $\langle \text{MW}_i \rangle$.

$$T_{\text{cell}} = \frac{\langle \text{MW}_i \rangle}{2R_u(\beta_{\text{cell}})^2} \tag{4}$$

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 Ref. [12]. Given two random numbers R_1 and R_2 in the range (0,1] a clone particle velocity component $U_{i,j}$ in the j direction for a particle with index i is calculated as

$$U_{i,j} = \mathbf{U}_{\text{cell}} \cdot \mathbf{n}_j + \frac{\sin(2\pi R_1)}{\beta_{\text{cell}}}\sqrt{-\frac{\langle \text{MW}_k \rangle}{\text{MW}_i}\ln R_2} \tag{5}$$

where $\mathbf{U}_{\text{cell}}$ and β_{cell} are the instantaneous cell-averaged bulk velocity and thermal speed scale as calculated in step 1 of the LD simulation procedures, $\mathbf{n}_j$ is the unit vector in the j direction, $\langle \text{MW}_k \rangle$ is the cell-averaged molecular weight, and MW_i 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

$$E_i = -\frac{\langle \text{MW}_k \rangle}{(\beta_{\text{cell}})^2}\ln R_3 \tag{6}$$

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

3.3. Particle removal in buffer cells

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. It should be noted that these procedures only enforce mass, momentum and energy conservation in an ensemble-averaged sense; exact conservation is enforced only with $\theta = 1$ as the number of particles per cell approaches infinity. However, errors associated with a lack of exact conservation should generally be small, and these procedures do provide a relatively simple and strongly coupled means of exchanging information between rarefied and continuum flowfield regions in a hybrid simulation.

4. 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]. A comparative analysis of different breakdown criteria is out of the scope of the present study, and the reader is referred to an extensive body of existing research on this topic [11,12,14,18–20]. For the hybrid algorithm presented here, a relatively simple and conservative technique is therefore proposed.

4.1. Determination of continuum breakdown

First, a gradient length local Knudsen number [18] 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 Eq. (3). Wang and Boyd [19] found that representative hypersonic flows may be assumed sufficiently close to equilibrium to justify near-equilibrium assumptions when $\max\{Kn_{GLL}\} < 0.05$, 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

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

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.

4.2. Cell assignment procedures

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 an 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.

5. Evaluation of the hybrid algorithm for a hypersonic blunt body flow

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.

5.1. Simulation setup

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 Schwartzentruber 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 1.6×10^{21} 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 (the reciprocal of 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 h. For comparison, a second simulation is run using only DSMC, and requires about 5.1 h on the same number of processors.

Note that, while the hybrid LD–DSMC 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 for this flow would require the use of automatic grid adaptation procedures during the hybrid simulation, so that cells in the LD domain would be far larger than those in the DSMC domain. (Some form of grid adaptation is required for maximum efficiency using *any* cell-based scheme which can accurately resolve strong shocks while employing efficient continuum simulation techniques in near-equilibrium regions.) Such a comparison is out of the scope of the present work, but a clear demonstration of efficiency gains relative to DSMC is provided below for a rarefied expansion flow. The reader is referred to a previous paper [17] for further discussion of computational efficiency in the LD method.

5.2. Hybrid simulation results and comparison with DSMC

The grid boundary geometry is shown in Fig. 2, 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 domain. Dark gray areas along the boundary between LD and DSMC domains denote cells within buffer region A, and neighboring light gray areas indicate cells within buffer region B.

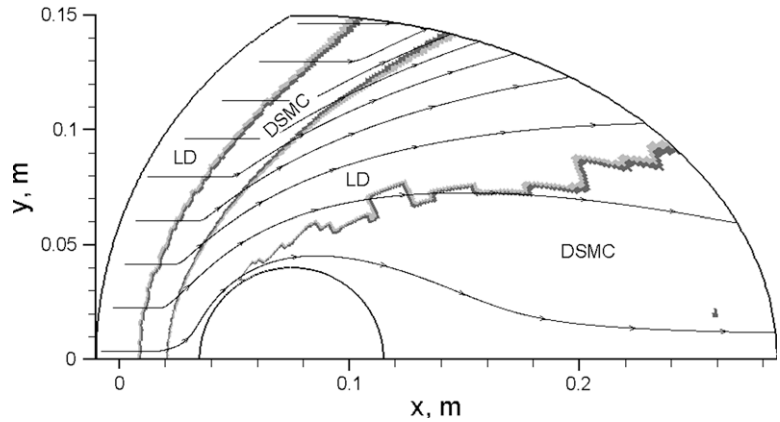

Fig. 2. Grid geometry, streamlines and boundaries between LD and DSMC regions in the hybrid simulation.

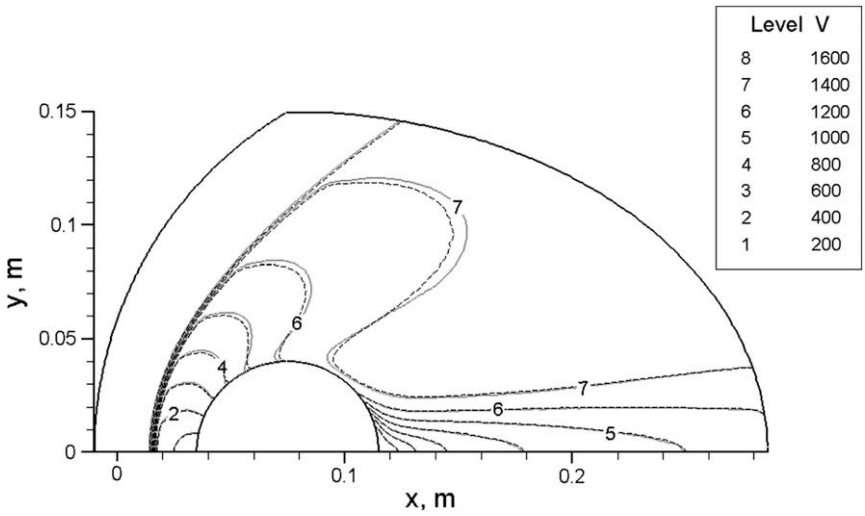

Fig. 3. Contours of bulk velocity magnitude from DSMC and hybrid simulations. Dashed contour lines are for DSMC. Values are in m/s.

Fig. 3 shows contours of bulk velocity magnitude for both the hybrid and DSMC simulations. Iso-contour lines for the hybrid simulation are shown in gray, while black dashed lines indicate results from the DSMC simulation. 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 within 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. 4. 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.

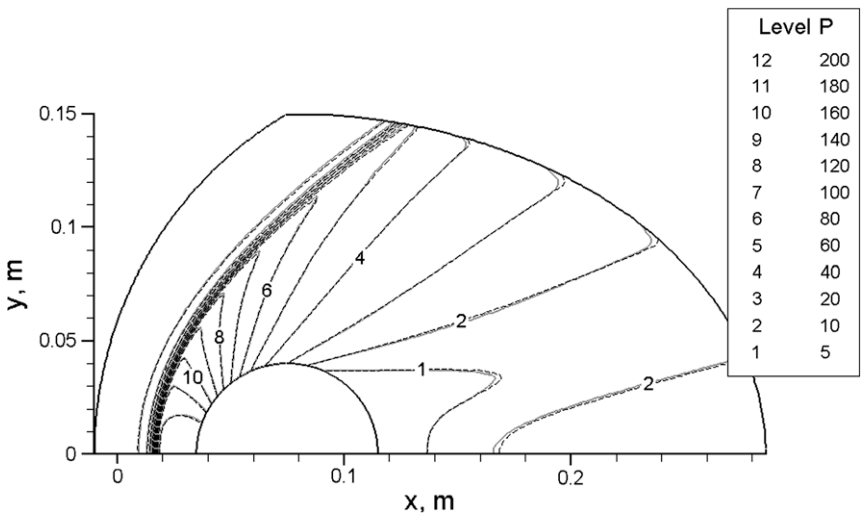

Fig. 4. Pressure contours from DSMC and hybrid simulations. Values are in Pa.

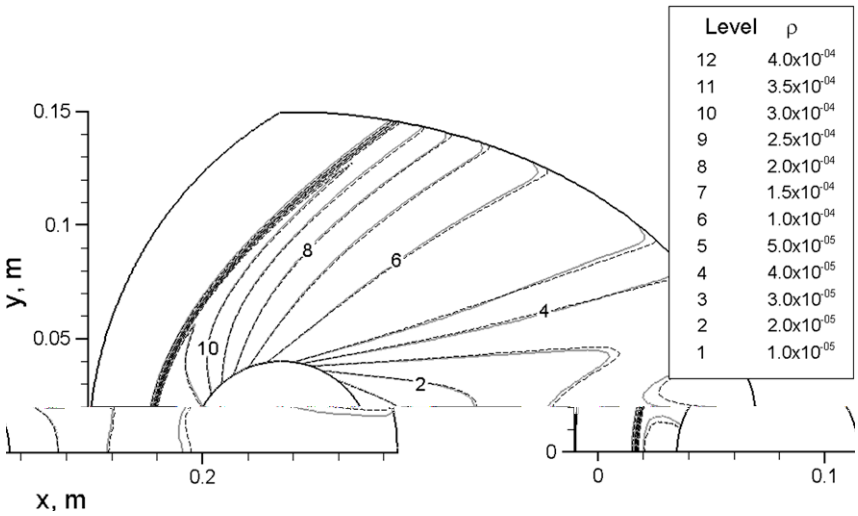

Fig. 5. Contours of gas density. Values are in kg/m^3.

Contours of gas density are shown in Fig. 5. The level of agreement observed in the figure is relatively good, although inferior to that found in the pressure contours of Fig. 4. The density is noticeably higher 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.

Fig. 6 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. 7. In comparing Figs. 6 and 7, 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. 6 and 7 also shows significant differences between translational and rota-

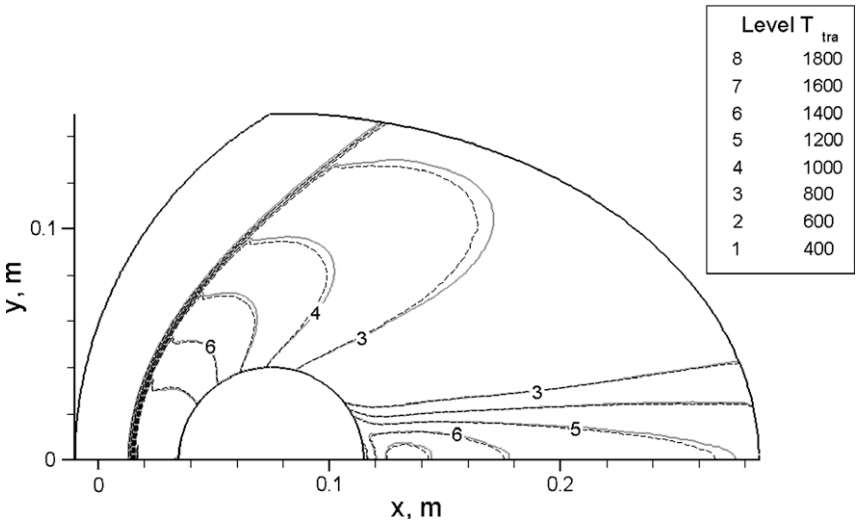

Fig. 6. Contours of translational temperature. Values are in K.

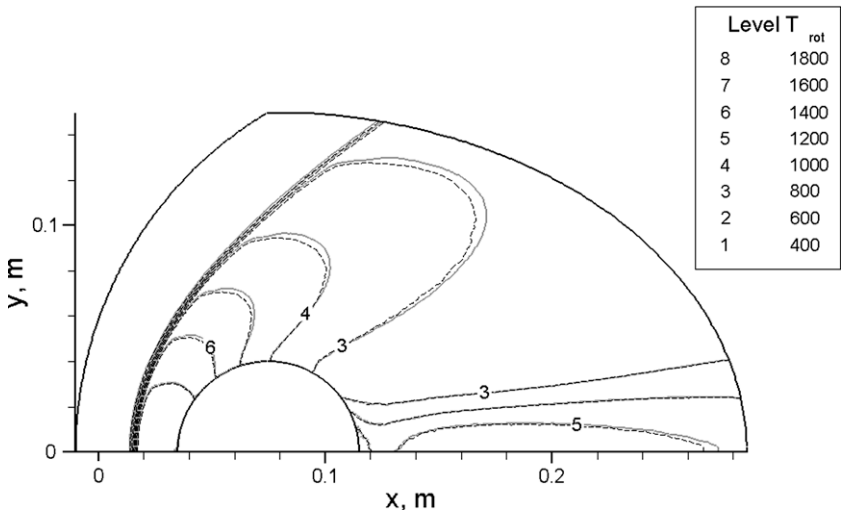

Fig. 7. Contours of rotational temperature. Values are in K.

tional temperatures within continuum regions assigned to the LD method in the hybrid simulation. (The 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. 7. 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. 8. 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

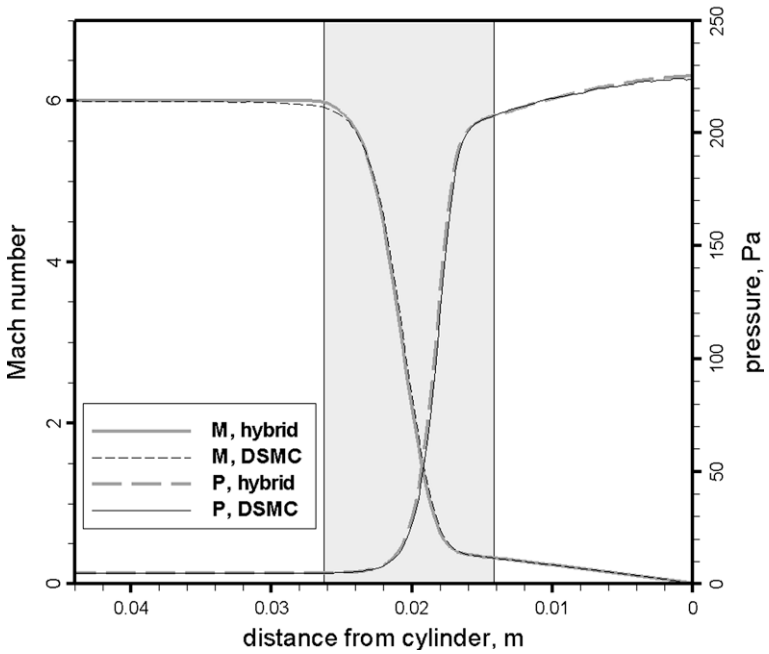

Fig. 8. Profiles of Mach number and pressure along the stagnation streamline.

continuum breakdown in Eq. (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. 8. As mentioned in the discussion of Fig. 3, 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.

Fig. 9 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. 8. All three hybrid simulation curves in Fig. 9 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 flow 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. 3–9. Other possible error sources include the fact that mass, momentum and energy are conserved in LD-DSMC coupling procedures only in an ensemble-averaged sense, as mentioned in Section 3.3. This may lead to random walk effects among various flow quantities, particularly when the local Mach number is very small along the boundary between buffer regions A and B. 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. 3, particularly near the cylinder surface where large transverse velocity gradients exist.

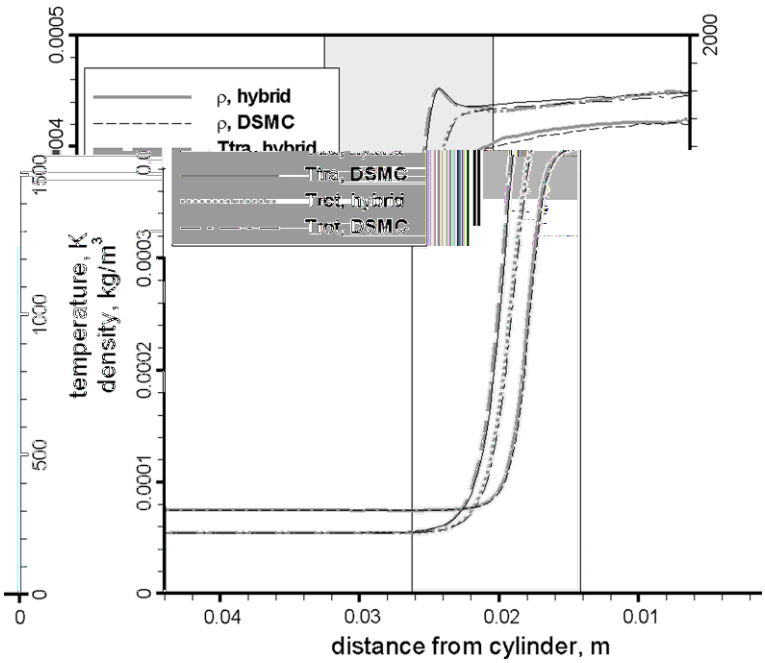

Fig. 9. Density, translational temperature and rotational temperature profiles along the stagnation streamline.

6. Application to a rarefied expansion flow

While the test case examined in the previous section was used effectively to evaluate the accuracy of the proposed hybrid particle scheme for a complex hypersonic flow problem, the hybrid simulation required more CPU hours than a DSMC simulation of the same flow and provided a relatively poor demonstration of hybrid simulation efficiency. As discussed above, the main potential advantage of the hybrid scheme relative to DSMC is that, within continuum regions assigned to the LD method in a hybrid LD–DSMC simulation, the cell size and time step can be far larger than the mean free path and mean collision time respectively. This property was demonstrated in a previous paper [17], and should allow for simulation speed-up relative to DSMC of potentially several orders of magnitude if the cell dimensions and time step size are properly rescaled in continuum regions. Because the location of continuum regions in a hypersonic blunt body flow is not generally known a priori, some automatic grid adaptation procedure must be built into the hybrid simulation code to take advantage of the less stringent cell size requirements in the LD method. It should be emphasized that any cell-based scheme intended for efficient and accurate simulation of very low Knudsen number flows with regions of strong translational nonequilibrium (such as hypersonic wakes and strong shocks) requires some form of either automatic or manual grid adaptation for maximum efficiency. The use of much larger cells in continuum regions is the primary source of efficiency gains for all hybrid schemes which employ DSMC in nonequilibrium regions.

While implementation of an automatic grid adaptation procedure is out of scope of the present study, the potential efficiency benefits of the hybrid scheme can be more easily demonstrated using another type of flow, where the location of continuum regions may be more accurately estimated during the grid generation process before simulation startup. One such flow is a nozzle/plume expansion flow in which gas is expelled through a small convergent–divergent nozzle into a near-vacuum. In this case, flowfield regions near and upstream of the nozzle throat are within the low Knudsen number regime where near-equilibrium assumptions underlying the LD method are valid, while regions further downstream of the throat may include considerable nonequilibrium effects, and therefore require a cell size (suitable for DSMC) which is much smaller relative to the mean free path.

Although this type of flow is well suited for demonstrating hybrid scheme efficiency in relation to DSMC, the mean free path within the convergent nozzle region may be sufficiently small to make any DSMC simulation of the entire flowfield prohibitively expensive. This means that a direct comparison of hybrid simulation results with DSMC results, as performed in the previous section, is not generally possible for this type of flow. However, some limited comparison can be made between hybrid LD–DSMC results and results from a CFD simulation for the compressible Euler equations.

6.1. Simulation setup

As a representative nozzle/plume expansion flow test case to demonstrate hybrid simulation efficiency, we choose a flow based on that of Boyd et al. [23]. Here N_2 is exhausted through a convergent–divergent nozzle then into a vacuum chamber, where the throat diameter is 3.18 mm, the nozzle divergence half-angle is 20°, and the nozzle area ratio is 100. The simulated flowfield geometry is shown in Fig. 10. Gas flow parameters include a stagnation temperature of 699 K, a stagnation pressure of 6400 Pa, a mass flow rate of 6.8×10^{-5} kg/s, and an ambient temperature and pressure of 300 K and 10^{-2} Pa respectively in the vacuum chamber. Along the inflow boundary upstream of the nozzle throat, the bulk velocity is 5.8 m/s, the Mach number is 0.01 and the number density is 6.63×10^{23} m^{-3}. As in simulations described in the previous section, all particle-wall collisions in the DSMC domain involve specular reflection, the VHS collision model is employed in DSMC calculations, the DSMC rotational collision number is set to one, and subcells of length one-half mean free path are used in the DSMC collision partner selection procedure. The nozzle throat diameter is used as a characteristic global length scale in the periodic determination of continuum breakdown using Eq. (7), and the simulation domain is initialized at startup with the subsonic nozzle inflow conditions upstream of the throat and with the vacuum chamber ambient conditions elsewhere in the flowfield.

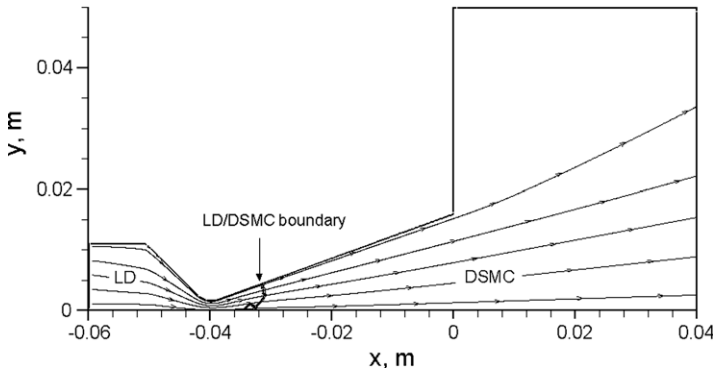

Fig. 10. Grid geometry, streamlines and location of LD and DSMC regions in the hybrid nozzle/plume flow simulation.

An axisymmetric hybrid LD–DSMC simulation is performed using a structured grid with 38,200 cells. All cells in regions of expected continuum breakdown are refined to less than two local mean free paths, and cells in continuum regions are sufficiently refined to assure grid independence. The simulation employs a uniform global time step interval of 10^{-8} s. This interval was estimated to meet DSMC time step requirements in regions expected to be assigned to DSMC, and to meet a standard CFL criterion in continuum regions assigned to the LD method. Calculations are performed for a startup period of 100,000 time steps, followed by a sampling period of another 100,000 time steps once steady state conditions have been reached.

6.2. Hybrid simulation efficiency relative to DSMC

The simulation is run in parallel on 12 processors for a total time of about 28.6 h. At steady state roughly 5.9 million particles are simultaneously tracked through the grid, with about 57% of cells located in regions assigned to the DSMC domain. The boundary between continuum (LD) and nonequilibrium (DSMC) regions at steady state is shown in Fig. 10, and is located approximately 9 mm downstream of the nozzle throat. As expected, most of the divergent nozzle region and all of the plume region is assigned to DSMC, while the entire convergent nozzle region upstream of the throat is simulated using the LD method. (A small nearly isolated area of DSMC cells is found in Fig. 10 just upstream of the main LD/DSMC boundary, and likely appears as a result of scatter along the axis in density gradients used to determine continuum breakdown. The uneven nature of the LD/DSMC boundary here is similar to that observed near the wake in Fig. 2 for the hypersonic cylinder flow simulation, and in both cases there seems to be no significant affect on simulation results.) Generally expected trends are also observed in the streamlines shown in Fig. 10, with converging streamlines upstream of the throat and an increase in streamline divergence angles downstream of the nozzle exit.

As noted above, a DSMC simulation of the entire flowfield would be prohibitively expensive, due to the extremely small mean free path and mean collision time in regions near and upstream of the nozzle throat. While no DSMC simulation of this flow could therefore be run for comparison, we can estimate the relative cost of a DSMC simulation through a series of rough approximations: First, the average ratio of the local mean free path to the cell size is observed to be around 0.015 in regions assigned to the LD method in the hybrid simulation. As a ratio of at least 0.5 is desired for a DSMC simulation, cell dimensions in continuum regions would need to be reduced on average by a factor of about 30. This means that both the number of cells and the number of particles in continuum regions would increase by a factor of roughly 10^3, following the relation between cell length and cell volume. In order to meet standard DSMC guidelines for the time step interval in the regions of highest density, the required time step would need to be reduced to around 4×10^{-9} s, which would again more than double the simulation time. As the computational expense of the LD method is comparable to DSMC or other DSMC-based methods when measured as the calculation time per particle per time step [17], it seems reasonable to assume that a full DSMC simulation of the nozzle/plume configuration considered here would be on the order of 2×10^3 times more expensive than the hybrid LD–DSMC simulation. Such a DSMC simulation would therefore require about 7×10^5 CPU hours, and would take over 6 years to run on the 12 processors employed in the hybrid simulation.

6.3. CFD simulation of the nozzle flow

Although a DSMC simulation of the entire flowfield requires computational resources far beyond those available to the authors, some comparison with hybrid simulation results can be performed by running an inviscid CFD calculation for the flow within the nozzle. A simulation of the nozzle flow is run using the LeMANS CFD code [24], with point-implicit time integration and a modified form of Steger–Warming flux vector splitting used to solve the axisymmetric compressible Euler equations. Calculations are performed on a structured grid similar to that used in the hybrid LD–DSMC simulation, and grid independence is confirmed through additional simulations which employ a more refined grid. The total CFD simulation time is about 7.9 h when run in parallel on four processors.

It should be emphasized that the CFD simulation cannot be assumed to accurately characterize flow properties in nonequilibrium regions downstream of the continuum breakdown boundary shown in Fig. 10. Moreover, the CFD simulation domain does not even include the plume region beyond the nozzle exit, due to CFD stability problems associated with the extremely low densities and rapid expansion through much of this region. Still, a comparison between CFD and hybrid simulation results within the nozzle provides a useful reference for evaluating hybrid simulation accuracy in regions where the CFD results can be considered valid. Such a comparison can also be used to demonstrate the importance of nonequilibrium effects through much of the divergent nozzle region where CFD and hybrid simulation results may differ significantly.

Note that a comparison of hybrid LD–DSMC simulation results with results from an uncoupled set of CFD and DSMC simulations, as in the viscous flow simulations in Ref. [23], would likely provide a better assessment of hybrid simulation accuracy for this flow. However, this uncoupled approach should not be viewed as a comparable alternative to coupled hybrid schemes such as that presented here, and the use of multiple simulations involving different methods has some significant limitations: first, an uncoupled CFD–DSMC approach will only work on steady state flow problems where the flow normal to the continuum breakdown boundary is supersonic. This contrasts with the LD–DSMC scheme, which includes no steady state assumptions and which allows two-way information transfer along subsonic boundaries, as demonstrated for the hypersonic cylinder flow in the previous section. Second, an uncoupled CFD–DSMC approach requires the user to run two separate

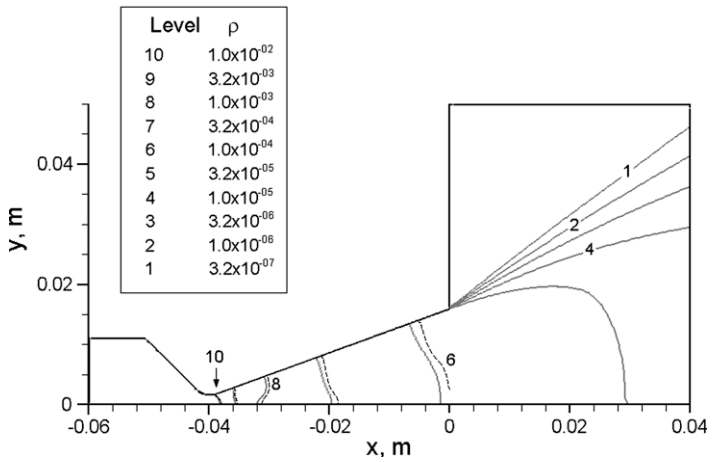

Fig. 11. Contours of density. Solid lines indicate contours for the hybrid simulation, and dashed lines are for the CFD nozzle flow simulation. Values are in kg/m^3.

simulations using two different codes, and to define a separate DSMC simulation domain based on some post-processor evaluation of continuum breakdown in the CFD results.

6.4. Hybrid simulation results and comparison with CFD

Contours of density are shown in Fig. 11 for the hybrid LD–DSMC simulation and for the CFD nozzle flow simulation. Dashed black lines represent CFD contours. One important characteristic observed in the figure is the enormous range of density values; the density varies by almost five orders of magnitude between the nozzle throat and the strong expansion region around the nozzle lip. In both hybrid results and CFD results, there is a general increase in density toward the axis within the nozzle, due to an increased expansion rate associated with the divergence of streamlines near the outward-sloping nozzle wall. One notable exception to this trend is at contour level 8, where results from both simulations show a reduction in density at the wall. This is likely related to the decrease in slope for streamlines downstream of the high curvature region around the throat.

Excellent agreement is found between results from the two simulations in regions upstream of the continuum breakdown boundary shown in Fig. 10. As expected, the level of agreement in density contours becomes progressively worse with downstream distance beyond the continuum breakdown boundary, with a relative difference of about 6.3% at the nozzle exit. Most of this difference can be attributed to the lack of consideration for translational nonequilibrium effects in the CFD simulation, although the contribution of other effects such as rotational temperature lag in the DSMC calculations may also be significant.

Contours of Mach number for the hybrid and CFD simulations are shown in Fig. 12. As in Fig. 11, very good agreement is observed between the two sets of results upstream of the continuum breakdown boundary, while significant disagreement is found further downstream within the divergent nozzle region. A sonic line is shown at the nozzle throat, and an additional sonic line is shown within the centered expansion fan which extends outward from the nozzle lip. The low Mach numbers observed within the plume region far from the axis may be attributed to the low density ambient gas in the simulated vac-

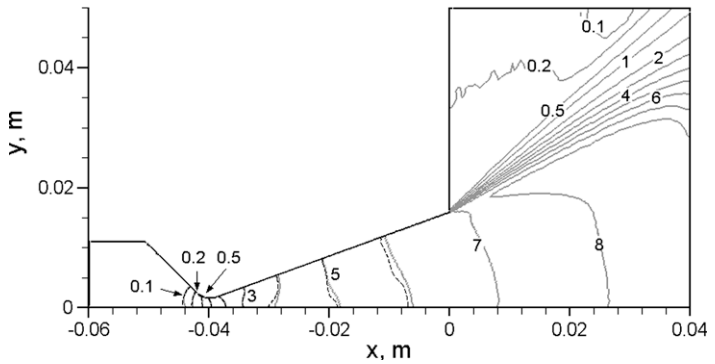

Fig. 12. Mach number contours. Dashed lines are for the CFD nozzle flow simulation.

uum chamber. Considerable scatter is found in these low Mach number contour lines, and results from small particle populations as well as a tendency for increased scatter in DSMC bulk velocity values within subsonic regions.

7. Conclusions

A hybrid particle scheme for the simulation of continuum and rarefied 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, requiring 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%. While there is no theoretical relation between the level of simulation accuracy and the cutoff value for continuum breakdown, it is interesting to note that the 2% figure in observed errors is consistent with the cutoff value of 0.02 used in Eq. (7).

To demonstrate the range of applicable flow problems and to estimate potential efficiency gains relative to DSMC, an additional hybrid simulation was run for a rarefied nozzle/plume expansion flow. A series of approximations were used to estimate the hybrid simulation expense as around three orders of magnitude smaller than that of an equivalent DSMC simulation, due primarily to differences in cell size requirements within the convergent nozzle region. A CFD nozzle flow simulation for the compressible Euler equations was run for comparison with hybrid simulation results. As expected, excellent agreement was observed between the two simulations upstream of a continuum breakdown boundary within the divergent nozzle region, while noticable disagreement was found downstream of this boundary.

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. In hybrid LD–DSMC simulations, the addition of viscous flow capabilities should permit the LD method to be applied to continuum regions where viscous effects are important but DSMC is prohibitively expensive. Other proposed extensions to the present work include the further characterization of hybrid code efficiency, using much larger cells in continuum regions, and evaluation of hybrid simulation accuracy for additional types of flows. Unsteady flow problems related to high altitude rocket exhaust flows and plume impingement should be particularly well suited to the proposed hybrid LD–DSMC scheme.

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. The authors would also like to thank Matt McKeown at the University of Michigan for his help with the CFD simulations that were performed as part of this study.

References

[1] H.S. Wijesinghe, R.D. Hornung, A.L. Garcia, N.G. Hadjiconstantinou, Three-dimensional hybrid continuum-atomistic simulations for multiscale hydrodynamics, J. Fluids Eng. 126 (2004) 768–777.
[2] R. Roveda, D.B. Goldstein, P.L. Varghese, Hybrid Euler/particle approach for continuum/rarefied flows, J. Spacecraft Rockets 35 (3) (1998) 258–265.
[3] R. Roveda, D.B. Goldstein, P.L. Varghese, Hybrid Euler/direct simulation Monte Carlo calculation of unsteady slit flow, J. Spacecraft Rockets 37 (6) (2000) 753–760.
[4] W. Wang, I.D. Boyd, Hybrid DSMC–CFD simulations of hypersonic flow over sharp and blunted bodies, AIAA Paper 2003-3644, 2003.
[5] T.E. Schwartzentruber, I.D. Boyd, A hybrid particle-continuum method applied to shock waves, J. Comput. Phys. 215 (2006) 402–416.
[6] T.E. Schwartzentruber, I.D. Boyd, A modular particle-continuum numerical method for hypersonic non-equilibrium gas flows, J. Comput. Phys. 225 (2007) 1159–1174.
[7] D.C. Wadsworth, D.A. Erwin, Two-dimensional hybrid continuum/particle approach for rarefied flows, AIAA Paper 92-2975, 1992.
[8] A.L. Garcia, J.B. Bell, W.Y. Crutchfield, B.J. Alder, Adaptive mesh and algorithm refinement using direct simulation Monte Carlo, J. Comput. Phys. 154 (1999) 134–155.

[9] V.I. Kolobov, R.R. Arslanbekov, V.V. Aristov, A.A. Frolova, S.A. Zabelok, Unified solver for rarefied and continuum flows with adaptive mesh and algorithm refinement, J. Comput. Phys. 223 (2007) 589–608.
[10] S. Tiwari, A. Klar, An adaptive domain decomposition procedure for the Boltzmann and Euler equations, J. Comput. Appl. Math. 90 (1998) 223–237.
[11] M.N. Macrossan, A particle-only hybrid method for near-continuum flows, in: Proceedings of the 22nd International Symposium on Rarefied Gas Dynamics, American Institute of Physics, 2001, pp. 388–395.
[12] G.A. Bird, Molecular Gas Dynamics and the Direct Simulation of Gas Flows, Oxford University Press, New York, 1994.
[13] T.J. Bartel, T.M. Sterk, J.L. Payne, B. Preppernau, DSMC simulation of nozzle expansion flow fields, AIAA Paper 94-2047, 1994.
[14] E.V. Titov, M.I. Zeifman, D.A. Levin, Application of the kinetic and continuum techniques to the multi-scale flows in MEMS devices, AIAA Paper 2005-1399, 2005.
[15] K.S. Breuer, E.S. Piekos, D.A. Gonzales, DSMC simulations of continuum flows, AIAA Paper 95-2088, 1995.
[16] M.N. Macrossan, The equilibrium flux method for the calculation of flows with non-equilibrium chemical reactions, J. Comput. Phys. 80 (1989) 204–231.
[17] J.M. Burt, I.D. Boyd, A low diffusion particle method for simulating compressible inviscid flows, J. Comput. Phys. 227 (2008) 4653–4670.
[18] I.D. Boyd, G. Chen, G.V. Candler, Predicting failure of the continuum fluid equations in transitional hypersonic flows, Phys. Fluids 7 (1) (1995) 210–219.
[19] W. Wang, I.D. Boyd, Predicting continuum breakdown in hypersonic viscous flows, Phys. Fluids 15 (1) (2003) 91–100.
[20] J.A. Camberos, C.R. Schrock, R.J. McMullan, R.D. Branam, Development of continuum onset criteria with direct simulation Monte Carlo using Boltzmann's H-theorem: review and vision, AIAA Paper 2006-2942, 2006.
[21] Q. Sun, I.D. Boyd, Evaluation of macroscopic properties in the direct simulation Monte Carlo method, J. Thermophys. Heat Transfer 19 (3) (2005) 329–335.
[22] S. Dietrich, I.D. Boyd, Scalar and parallel optimized implementation of the direct simulation Monte Carlo method, J. Comput. Phys. 126 (1996) 328–342.
[23] I.D. Boyd, P.F. Penko, D.L. Meissner, K.J. DeWitt, Experimental and numerical investigations of low-density nozzle and plume flows of nitrogen, AIAA J. 30 (10) (1992) 2453–2461.
[24] L.C. Scalabrin, I.D. Boyd, Development of an unstructured Navier–Stokes solver for hypersonic nonequilibrium aerothermodynamics, AIAA Paper 2005-5203, 2005.