

# A hybrid particle-continuum method applied to shock waves

T.E. Schwartzentruber<sup>\*</sup>, I.D. Boyd

*University of Michigan, Department of Aerospace Engineering, Francois-Xavier Bagnoud Building, 1320 Beal Avenue,  
Ann Arbor, MI 48109, United States*

Received 22 April 2005; received in revised form 10 October 2005; accepted 31 October 2005  
Available online 7 December 2005

## Abstract

A hybrid numerical scheme designed for hypersonic non-equilibrium flows is presented which solves the Navier–Stokes equations in regions of near-equilibrium and uses the direct simulation Monte Carlo method where the flow is in non-equilibrium. Detailed analysis of each stage of the hybrid cycle illustrates the difficulty in defining physically correct DSMC boundary conditions in regards to both macroscopic state and velocity distribution. However, results also show that DSMC boundary conditions have little effect on a previously initialized interior particle domain. A sub-relaxation technique capable of determining macroscopic, hydrodynamic properties in a DSMC simulation is used to determine low-scatter boundary conditions for the NS domain. Particle and continuum domains adapt during the hybrid simulation through application of a continuum breakdown parameter based on the gradient-length Knudsen number. The hybrid code reproduces experimental results and full DSMC simulations in half the time for a large range of 1D shock waves in argon and diatomic nitrogen gas.

© 2005 Elsevier Inc. All rights reserved.

*Topic classification:* 65C05, 82C80

*Keywords:* Direct simulation Monte Carlo; Hybrid particle-continuum; Hypersonics; Non-equilibrium; Re-entry vehicles

## 1. Introduction

Computation of the aerothermodynamics of hypersonic re-entry vehicles along their entire trajectory involves continuum conditions at low altitudes and rarefied, or *non-equilibrium*, conditions at high altitudes. Well-established simulation methods already exist for each of these flow regimes. For example, the continuum Navier–Stokes (NS) equations may be solved using algorithms from Computational Fluid Dynamics (CFD) and the particle based direct simulation Monte Carlo method (DSMC) may be used for the non-equilibrium flows. A hybrid method that blends the CFD and DSMC techniques is an attractive approach for simulation of flows involving a mixture of both continuum and non-equilibrium flow regimes. This is the situation at intermediate altitudes where within a mostly continuum flow, there may be local regions of non-equilibrium

<sup>\*</sup> Corresponding author. Tel.: +734 994 4924.

*E-mail addresses:* [schwartt@engin.umich.edu](mailto:schwartt@engin.umich.edu) (T.E. Schwartzentruber), [iainboyd@engin.umich.edu](mailto:iainboyd@engin.umich.edu) (I.D. Boyd).

flow generated by both the rapid expansion behind a re-entering capsule as well as by strong gradients in shock waves and boundary layers.

For dilute gases, such as the earth’s atmosphere, the most popular numerical method for simulating non-equilibrium flow is the DSMC method developed by Bird [1]. Here, the trajectories of a large number of simulated particles are followed simultaneously through a grid of computational cells. For each iteration, particles are first moved along their trajectories without colliding, after which particles residing within the same cell are randomly selected for a collision process. A major limitation of DSMC is that in order for this collision process to be physically accurate, the cell size must be on the order of the mean free path ( $\lambda$ ) while containing at least 20 simulated particles per cell [1]. As a result, 2D and 3D DSMC simulations can require prohibitively high numbers of computational cells and therefore simulated particles, especially in regions where  $\lambda$  is very small. However, it is precisely in these regions where the continuum approximation is valid and the NS equations can be solved without the same restriction on cell size. This is the motivation behind developing a hybrid CFD-DSMC numerical scheme for hypersonic non-equilibrium flows.

### 2. Hybrid scheme considerations

Two major problems that must be addressed in such a hybrid scheme include the details of how particle and continuum regions transmit information across their interface and determining *where* this interface should be placed. Typically, particle and continuum regions are determined by applying a continuum breakdown parameter to the flow field. This study uses the gradient-length Knudsen number

$$Kn_{GLQ} = \frac{\lambda}{Q_{\text{free-stream}}} |\nabla Q|, \tag{1}$$

where  $Q$  represents the parameter of interest, such as density ( $\rho$ ), velocity ( $V$ ), or temperature ( $T$ ). The actual continuum breakdown parameter is then the maximum of these, that is:

$$Kn_{GL} = \max(Kn_{GL\rho}, Kn_{GLV}, Kn_{GLT}). \tag{2}$$

It has been shown for flows representative of hypersonic re-entry problems [2], that in regions of the flow field where  $Kn_{GL} < 0.05$ , the discrepancy between a NS and DSMC solution is less than 5%. Thus, these regions could be solved using a continuum solver with little error.

The transfer of information between particle and continuum domains is typically handled using one of the two methods depicted in Fig. 1. Flux-based coupling, Fig. 1(a), involves calculating the fluxes of mass, momentum, and energy at the interface according to the particle cell ( $F_p$ ), and according to the continuum cell ( $F_c$ ).  $F_p$  is calculated by tracking individual simulation particles as they cross the interface whereas  $F_c$  is extrapolated using macroscopic gradients in the continuum domain; a standard procedure in a continuum

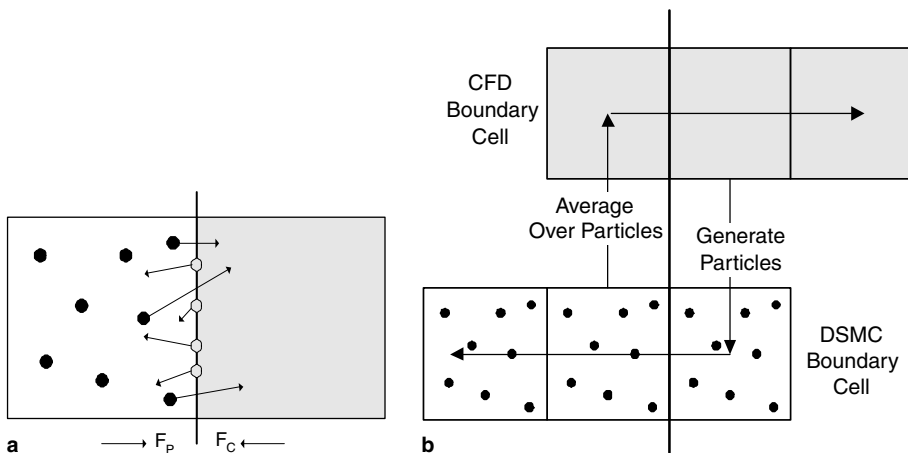


Fig. 1. Typical hybrid coupling procedures.

solver. However, in general  $F_P \neq F_C$  and a modified flux must somehow be calculated such that the transfer of information across the interface remains conservative. This is analogous to the conservative flux functions used widely in pure continuum solvers. The modified flux is not only used to update the continuum solution, but is also used to create a distribution of particles on the interface which are then allowed to propagate into the particle simulation. State-based coupling, Fig. 1(b), temporarily averages particle information to obtain a macroscopic state on one side of the interface and at the same time generates a distribution of particles from a macroscopic state on the other side of the interface. In this way, the flux into both the continuum domain and particle domain is handled exclusively by its respective solver and is thus inherently conservative. In regards to the frequency of information transfer, a “decoupled” hybrid simulation will refer to a simulation in which both particle and continuum domains are solved to steady-state before transferring information across their interface. A “loosely” coupled approach increases the frequency of information exchange, whereas a “strongly” coupled approach exchanges information at every time-step. With respect to the averaging of particle information, the statistical scatter involved in determining  $F_P$  in Fig. 1(a) is much higher than that associated with averaging particle information to obtain a macroscopic state as seen in Fig. 1(b). In fact the statistical error in the flux quantities can be shown [3] to scale as

$$E_{\text{flux}} \sim \frac{E_{\text{state}}}{Kn}, \quad (3)$$

where  $Kn$  is the Knudsen number which at the interface (near-equilibrium conditions) will typically be close to 0.01. The error in such an average is also known to scale with the inverse square-root of the number of samples. This implies that the number of samples required to reduce  $E_{\text{flux}} \approx E_{\text{state}}$  will scale as  $1/Kn^2$ ; a significant disadvantage of flux-based coupling.

In order to reduce the statistical scatter when averaging DSMC particle information, Wang and Boyd used the Information Preservation (IP) scheme [4] which preserves microscopic and macroscopic information for each DSMC simulation particle. Statistical scatter was then eliminated by using state-based coupling and averaging over the IP-particle information. Although successful for certain 2D flows [5], when applied to 1D normal shock waves, it was found that the IP scheme produced an incorrect post-shock state and a shock wave that was too thin. A new formulation for the IP energy flux [6] was able to remedy these problems somewhat, however, at large computational expense.

Other researchers have attempted to directly couple a continuum solver to a particle solver. Ref. [7] presents a discussion of hybrid method considerations as well as a summary of published work on such hybrid schemes. For example, Hash and Hassan present both a decoupled DSMC-NS hybrid simulation of a hypersonic blunted cone [8] and a loosely coupled simulation of Couette flow [9] using flux-based coupling. For the study of Couette flow, the authors analyze different conservative hybrid flux formulations, concluding that use of the Marshak condition [10] is most accurate and efficient. These studies also maintain a fixed particle continuum interface set by a conservative value of  $Kn_{GL} = 0.005$ . Wadsworth and Erwin developed a strongly coupled, flux-based, hybrid DSMC-NS scheme and applied it to both 1D shocks [11] and 2D rarefied slit flow [12]. In these studies, a Maxwellian distribution was used to generate simulation particles at the interface and the domain boundaries remained fixed. Other studies have been carried out where the domain boundaries are re-evaluated during the simulation. Le Tallec and Mallinger use criterion based on the residual of the Grad 13-moment equations to adaptively position the interface and applied their scheme to 2D flow around an ellipse and a flat plate [13]. Roveda et al. strongly couple the Euler equations with DSMC for time-accurate moving 1D shock waves [14] and 2D unsteady slit flow [15]. Since the Euler equations are used, a Maxwellian distribution is sampled to generate simulation particles at the interface. The authors employ state-based coupling and effectively “clone” particles near the interface in order to reduce the statistical scatter transferred to the continuum domain. Finally, the domains are adaptively defined by first applying a smoothing function to the DSMC domain and then calculating continuum breakdown using a low value of  $Kn_{GL} = 0.005$ . Quite recently, a sophisticated 3D hybrid DSMC-Euler scheme has been proposed which uses flux-based coupling and embeds the particle solver in the finest level of an adaptive mesh and algorithmic refinement (AMAR) scheme [16]. In addition to these hybrid methods which attempt to couple *standard* CFD and DSMC methods, other researchers have developed new, single, methods designed to work over a large Knudsen number range. For example, Pareschi and Caglioli formulate a Time Relaxed Monte Carlo (TRMC) method and apply it to a

space homogeneous case [17]. Computational complexity and statistical variance are reduced because fewer particles are used to represent the solution, the remainder of which is represented by a continuous Maxwellian distribution. As equilibrium is approached, more and more of the solution is represented by the continuous Maxwellian and the method essentially becomes an Euler solver in this limit.

This paper analyzes in detail, each step in a hybrid cycle designed for steady-state hypersonic flows. Focusing on steady-state solutions allows for temporal averaging, which will be shown in upcoming sections to effectively control statistical fluctuations. The hybrid scheme strongly couples existing NS and DSMC solvers. The interface locations are adapted during the simulation based on the published recommendation [2] of  $Kn_{GL} = 0.05$ . State-based coupling takes advantage of the existing boundary procedures of both solvers to handle information transfer while introducing less statistical scatter than a flux-based scheme. The hybrid scheme is applied to 1D normal shock waves in argon and diatomic nitrogen gas. First, the magnitude of statistical scatter is analyzed as well as the accuracy of the velocity distributions generated at the interface and their effect on the particle domain. This information is then used to implement a simple, robust and efficient hybrid code. The code is tested for shock Mach numbers ranging from 1.5 to 10 where the user is required to change only the gas properties and free-stream conditions. The results are compared with experimental data and the hybrid code is profiled to analyze its numerical efficiency.

### 3. Numerical models

The hybrid CFD-DSMC numerical scheme studied in this paper, divides the flow field into continuum regions, where the CFD method is used, and particle regions where the DSMC method is used. This hybrid scheme uses standard CFD and DSMC solvers that have been extensively researched and validated in the literature. The major challenge that this research addresses is how to couple a probabilistic particle solver to a deterministic continuum solver. The CFD code solves the two-dimensional, laminar, Navier–Stokes equations using an explicit, second-order, finite volume algorithm. The inviscid fluxes are calculated using modified Steger–Warming flux-vector splitting [18] and the viscous derivatives are calculated using centered differencing. The particle regions are solved using MONACO [19], a general, cell-based implementation of the DSMC method which statistically simulates the Boltzmann equation. The viscosity model used in the NS solver is

$$\mu = \mu_{\text{ref}} \left( \frac{T}{T_{\text{ref}}} \right)^{0.75}, \quad (4)$$

where  $\mu_{\text{ref}} = 2.13 \times 10^{-5}$  and  $1.67 \times 10^{-5}$  N s/m<sup>2</sup> are the reference viscosities of argon and nitrogen respectively, at  $T_{\text{ref}} = 273$  K. The parameters used in the DSMC collision model (variable hard-sphere) are consistent with this viscosity model. For diatomic nitrogen, MONACO employs the variable rotational energy exchange probability model of Boyd [20] where the reference temperature for rotational energy exchange is specified as 91.5 K and the maximum rotational collision number as 18.1. Energy transfer to vibrational modes is not considered. In this study, the hybrid code is applied to one-dimensional normal shock waves in argon and diatomic nitrogen gas. Results are compared with an existing data-set of detailed experimental measurements for these flows obtained by Alsmeyer [21]. Measurements were made using an electron beam absorption technique and consist of density profiles over a range of shock Mach numbers. The grid employed for *all* hybrid, full DSMC and full NS simulations presented in this paper has 400 uniform cells in the flow direction and five cells in the transverse direction. Although computed on a 2D grid, the resulting flow fields are strictly 1D and are only presented for the middle (3rd) row of cells. The grid spans approximately 40 upstream mean free paths, specified by Alsmeyer as  $\lambda = 1.098 \times 10^{-3}$  m. The free-stream conditions are  $T_1 = 300$  K and the number density  $n_1 = 1.6095 \times 10^{21}$  particles/m<sup>3</sup>. This results in free-stream density values of  $\rho_1 = 1.069 \times 10^{-4}$  kg/m<sup>3</sup> and  $\rho_1 = 7.480 \times 10^{-5}$  kg/m<sup>3</sup> for argon and nitrogen, respectively. The velocity is then adjusted to achieve the desired flow Mach number. The simulation particle weight, equal to the ratio of real to simulated particles, is set to  $5.0 \times 10^{19}$  for all simulations in this study. This results in the number of simulation particles per cell ranging from 30 upstream of the shock to 50–100 downstream of the shock for Mach numbers ranging from 1.55 to 10, respectively. The DSMC time-step,  $\Delta t_{\text{DSMC}}$  is specified such that  $V_{\text{free-stream}} \times \Delta t_{\text{DSMC}} = \frac{1}{5} \Delta x$ , and the NS time-step is matched to this value. To maintain consistency for efficiency comparisons, a full DSMC simulation *begins* with particles initialized by a Chapman–Enskog

distribution based on a NS solution, similar to the initial particle regions in a hybrid simulation. The full DSMC simulation is considered to have reached steady-state when the total number of simulation particles no longer varies appreciably. In contrast, a hybrid simulation is considered to have reached steady state only after the interfaces have adapted to their final positions *and* the total number of particles has reached steady state. In order to obtain the same level of statistical scatter in the final solutions, both full DSMC and hybrid simulations are then sampled for 3000 time-steps after reaching steady state.

#### 4. Accuracy of DSMC boundary conditions

The procedure for imposing boundary conditions on the DSMC domain was depicted earlier in Fig. 1(b). For each time-step, all simulation particles in these boundary cells are first deleted and then re-generated based on current NS information. The number of new particles is determined directly from the NS cell density and are randomly distributed within the DSMC boundary cell volume. The velocities of these newly generated particles are sampled from a Chapman–Enskog distribution [22] based on the local macroscopic state and gradient, known from the NS solver. The boundary cells then become an extension of the DSMC domain for one standard DSMC cycle.

As previously mentioned, one of the important tasks in a hybrid scheme is determining in what regions a particle formulation is necessary. The current hybrid implementation begins with a full NS solution and applies the breakdown parameter (Eq. (2)) to determine the initial particle and continuum domains. Particles are then generated in all cells of the particle domain using the Chapman–Enskog distribution, based on this initial NS solution. It is well known that the NS equations predict a shock that is too thin. As shown in Fig. 2(a), even the application of a very conservative breakdown parameter may result in the initial DSMC domain being too narrow. Focusing only on the DSMC region in Fig. 2(a), the hybrid simulation will begin with an incorrect DSMC solution that is enclosed by boundary cells fixed at a physically incorrect *state*. This raises a fundamental question, namely, how will the DSMC solution respond to such ill-posed boundary conditions?

To answer this question, for a Mach 5 shock wave in argon gas, three DSMC domains are determined by applying three values of the continuum breakdown parameter,  $Kn_{GL} = 0.01, 0.1$  and  $1.0$ . This results in initial DSMC domains that are too narrow, where hybrid interfaces now lie inside the physically correct shock profile as seen in Fig. 2(a). The DSMC solver is then run until the solution no longer varies noticeably ( $\approx 5000$  time-steps) while keeping the boundary conditions fixed. The resulting macroscopic velocity profiles in the narrow DSMC computational domains are shown in Fig. 2(b). It is clear that although the boundary conditions

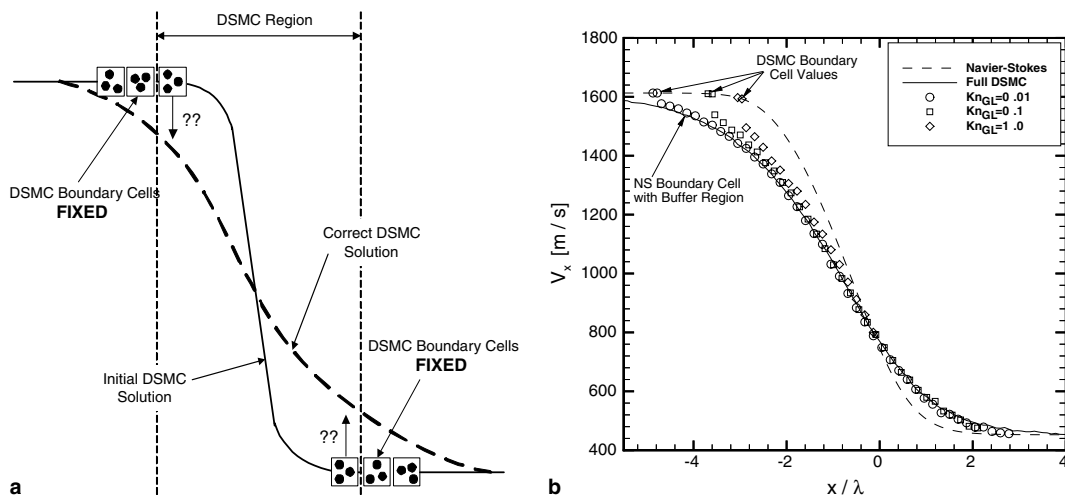


Fig. 2. Imposing incorrect DSMC boundary conditions. (a) Typical velocity profile showing initial DSMC domain. (b) Effect of ill-posed DSMC boundary conditions.

are fixed on the incorrect NS solution, the entire DSMC solution (including the cells immediately adjacent to the incorrect boundary cells) proceed towards the correct solution. Thus, these ill-posed boundary conditions seem to have little effect on the interior domain. This may be a result of the fact that relatively few simulation particles enter via the boundary cells compared to the vast majority of particles that were initially created to fill the interior of the DSMC domain. This behavior, verified to be independent of grid density and Mach number, is an important aspect of the hybrid cycle. Clearly, if interior cells were strongly influenced by these ill-posed boundary conditions then this incorrect state would be reflected back into the NS domain. A further way to reduce this problem would be to generate the NS boundary condition from a DSMC cell located slightly inside the interface as opposed to the cell immediately touching the interface. Such a “buffer” would not allow any local effects of ill-posed DSMC boundary conditions to be transferred back into the continuum domain.

Ultimately, a DSMC simulation requires boundary cells not only representing the correct macroscopic state, but whose particles represent the physically correct velocity distribution at that location in the flow field. Most DSMC simulations apply boundary conditions in uniform flow regions where a Maxwellian distribution is physically correct. A hybrid scheme attempts to place boundary conditions in regions of near-equilibrium flow and therefore requires at least a Chapman–Enskog velocity distribution. However, even if the DSMC boundary cells are successfully set to the correct macroscopic state, the generated Chapman–Enskog distribution is not necessarily the correct velocity distribution. This situation is portrayed in Fig. 3 where the flow gradient used to generate the distribution comes from the NS solution, which in this case has a much steeper gradient than the true physical gradient at that point in the shock wave. Fig. 4 shows the true velocity distribution (calculated from a full DSMC simulation) as well as the Chapman–Enskog distributions generated using the steep NS gradient and the true physical gradient for both pre and post-shock interfaces. These plots are a result of the Mach 5 simulation using  $Kn_{GL} = 0.01$ , whose interface locations can be seen from Fig. 2(b). The post-shock interface (see Fig. 4(a)) is seen to be in a region of near-equilibrium and thus all distributions agree well. However, at the pre-shock interface (see Fig. 4(b)), the Chapman–Enskog distribution based on the NS state and gradient is a very poor DSMC boundary condition. Even if this distribution is generated using the true physical gradient (which is unknown during the simulation), it still does not reproduce the *true* distribution. Therefore, until the interfaces are correctly located in regions of near-equilibrium, the DSMC boundary conditions may contain both an un-physical state and velocity distribution. Fortunately, these ill-posed boundary conditions were shown not to have a large influence on the interior DSMC solution.




Fig. 3. Generation of velocity distributions in boundary cells.






Fig. 4. Velocity distributions in DSMC boundary cells. (a) Post-shock boundary cell. (b) Pre-shock boundary cell.

### 5. Reducing scatter in NS boundary conditions

Now that it has been determined that despite initially, ill-posed boundary conditions, the DSMC region evolves towards the correct solution in time, this information must be transferred into the continuum domain. For example, Fig. 2(b) shows the variation in macroscopic velocity calculated in the DSMC region for a Mach 5 shock wave with boundaries set using  $Kn_{GL} = 0.01$ . As suggested at the end of the previous section, a small overlap region is created such that the NS boundary cell is placed 5 cells inside the interface. Fig. 2(b) shows that in the pre-shock NS boundary cell, the velocity decreases from 1614 to 1550 m/s (roughly 4%) during approximately 5000 DSMC time-steps. It is desirable to represent this changing DSMC information as a time varying NS boundary condition. As depicted earlier in Fig. 1, this is achieved by averaging over all particles in a DSMC cell to obtain macroscopic properties at each time-step. Figs. 5 and 6 display the success of various averaging procedures in the pre-shock NS boundary cell compared to the raw DSMC data for velocity and temperature, respectively. Here the raw DSMC velocity is simply the average of all particle velocities in the boundary cell at a given time-step. Clearly, this simple spatial average is subject to large statistical fluctuations which completely overwhelm the physical variation of both velocity and temperature. In the pre-shock NS

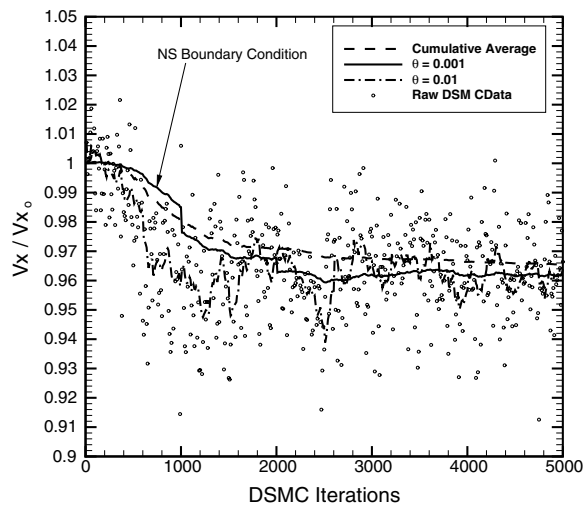


Fig. 5. Reduction of DSMC velocity scatter for NS boundary condition.

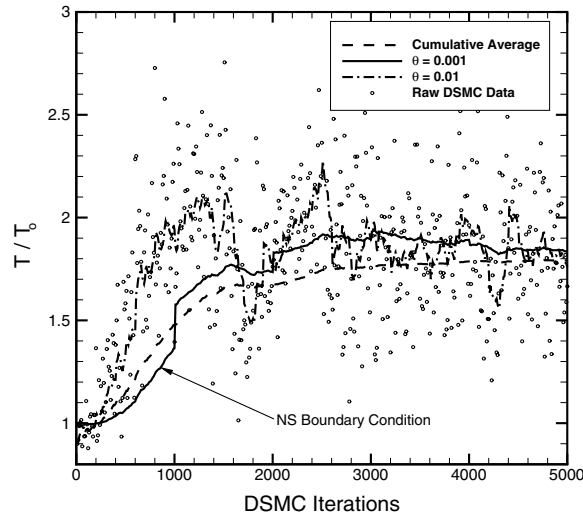


Fig. 6. Reduction of DSMC temperature scatter for NS boundary condition.

boundary cell, the velocity fluctuates by as much as 8%, and the temperature fluctuates by as much as 100%. Although not shown, the statistical variations in density, velocity and temperature in the post-shock NS boundary cell are roughly 30%, 50% and 20%, respectively. The macroscopic temperature in a DSMC cell at a given time-step (the raw data) is calculated by

$$T = \frac{m}{3k} \left( \langle V_x^2 + V_y^2 + V_z^2 \rangle - \langle V_x \rangle^2 - \langle V_y \rangle^2 - \langle V_z \rangle^2 \right), \quad (5)$$

where  $k$  is the Boltzmann constant and  $m$  is the mass of the particle. It has been observed that sampling from a small number of DSMC particles (less than 100) consistently depresses the value of the translational temperature. Roveda et al. [14] correct for this, however, in the current simulations this discrepancy is less than 2% and Eq. (5) is used exclusively. Figs. 5 and 6 also plot the cumulative time-average of the raw DSMC data. The inclusion of the history of the flow properties is seen to reduce the scatter almost completely while at the same time introducing a time lag. Recently, a method for evaluating instantaneous macroscopic hydrodynamic properties from a DSMC simulation has been proposed. The sub-relaxation technique by Sun and Boyd [23] includes the cumulative history using a relaxation factor and in addition, allows for “old history” to be removed from the average when it begins to pollute, or lag, the true variation. The sub-relaxation formula for the temporal average  $\bar{A}_j$  is:

$$\bar{A}_j = (1 - \theta)\bar{A}_{j-1} + \theta A_j \quad (6)$$

and the correction, which removes the “old history” is given by:

$$\bar{A}'_j = \bar{A}_j + \frac{(1 - \theta)^{j-i}}{1 - (1 - \theta)^{j-i}} (\bar{A}_j - \bar{A}'_i), \quad (7)$$

where  $j$  is the current time-step,  $i$  is the previous time-step at which a correction was made, and  $\bar{A}'_i$  is the temporal average recorded at this previous time-step  $i$ . As detailed in Ref. [23], Eq. (7) effectively removes the history before time-step  $i$ . In order to maintain stability, this correction is performed only when the coefficient of this correction is between zero and one. This condition is approximately satisfied when  $j = \frac{1}{\theta} + i$ . Note that the amount of scatter associated with a given  $\theta$  value is the same as the scatter resulting from averaging over  $\frac{1}{\theta}$  time-steps [23]. Figs. 5 and 6 show the success of this averaging technique applied to the DSMC velocity and temperature variations, respectively. Although when using  $\theta = 0.01$  the average follows the DSMC variation with almost no lag, the scatter is far too large to use this average as a NS boundary condition. Clearly, large scatter in a NS boundary condition will induce large numerical waves in the continuum domain which will likely cause instability and failure of the hybrid code. However, a time lag is more acceptable, considering



this hybrid code is being developed for steady-state flows. By decreasing  $\theta$  to 0.001, the scatter is reduced significantly with an acceptable lag. Notice that the sub-relaxation technique reaches the steady state value at approximately 2500 time-steps while the cumulative average still has not reached it at 5000 time-steps. Additionally, in contrast to the cumulative average, the sub-relaxation average behaves smoothly at early time-steps since only a small weight of  $\theta$  is applied to each successive raw DSMC measurement. For these reasons, the sub-relaxation technique will be used in the current hybrid code to determine the NS boundary condition at every DSMC time-step, using a value of  $\theta = 0.001$ .

## 6. Proposed hybrid cycle

In previous sections, the methods used to obtain boundary conditions for both the NS and DSMC domains were outlined. The NS boundary condition is able to follow a changing DSMC solution with manageable lag and low scatter. The DSMC boundary condition was shown to be valid only in regions of near-equilibrium flow and thus it is desirable to locate the proper interface location quickly. Note however, results also concluded that the DSMC boundary condition has little effect on the interior DSMC solution. This section will implement all of the above results and recommendations into a strongly coupled hybrid cycle. The proposed hybrid cycle involves the following steps:

- (1) Obtain an initial solution to the NS equations.
- (2) Generate DSMC/NS domains by applying the  $Kn_{GL}$  continuum breakdown parameter to NS regions only (initially, the entire hybrid domain is a NS region). Generate new simulation particles as needed using the Chapman–Enskog velocity distribution.
- (3) Create particles in the DSMC boundary cells from the current NS solution using the Chapman–Enskog velocity distribution.
- (4) Advance the DSMC domain by one time-step,  $\Delta t_{DSMC}$ .
- (5) Set the NS boundary conditions using the sub-relaxation technique applied to the updated DSMC particles.
- (6) Advance the NS domain in time by  $\Delta t_{DSMC}$ .
- (7) Return to step (3) -OR- Every  $N_{Br}$  time-steps, return to step (2)
- (8) If step (2) no longer modifies the DSMC/NS domains, then fix the domains. Begin cumulative DSMC sampling and now generate the NS boundary condition based on this cumulative average. Repeat steps (3) to (6) until the degree of scatter in the DSMC solution, and the NS residual, fall below threshold values.

Fig. 7 exhibits how this hybrid cycle proceeds to solve a Mach 6.5 shock wave. In order to clearly visualize the progression of the hybrid solution, the breakdown parameter is applied only every  $N_{Br} = 2000$  time-steps and the recommended [2] value of  $Kn_{GL} = 0.05$  is used. All variables are normalized as,

$$\rho_n = \frac{\rho - \rho_1}{\rho_2 - \rho_1}, \quad Vx_n = \frac{Vx - Vx_2}{Vx_1 - Vx_2}, \quad T_n = \frac{T - T_1}{T_2 - T_1} \quad (8)$$

with the spatial dimension normalized by the mean free path of the gas upstream of the shock. Only temperature profiles are presented, as they involve the largest variations in both the DSMC and NS solutions, especially upstream of the shock wave. Initially, the entire hybrid solution is a NS domain set to a NS solution of the shock wave problem. Immediately, the breakdown parameter is applied resulting in the initial DSMC and NS domains depicted by vertical lines in Fig. 7(a). Here, the hybrid solver has already cycled for 2000 time-steps and as a result, the hybrid DSMC solution has moved significantly towards the full DSMC solution. Recall, this important behavior was observed previously in Fig. 2(b). The hybrid NS solution is seen to have followed this changing DSMC solution smoothly with some lag. In fact, for the first 2000 time-steps, the exact NS boundary condition in the pre-shock region was shown previously in Fig. 6. The breakdown parameter is then applied to the hybrid NS portion of the solution seen in Fig. 7(a). Since the gradient in the NS region has now been extended further in each direction, the new DSMC domain (shown in Fig. 7(b)) is now larger. Here, the hybrid solver has cycled for another 2000 time-steps with its new domains and the DSMC solution has

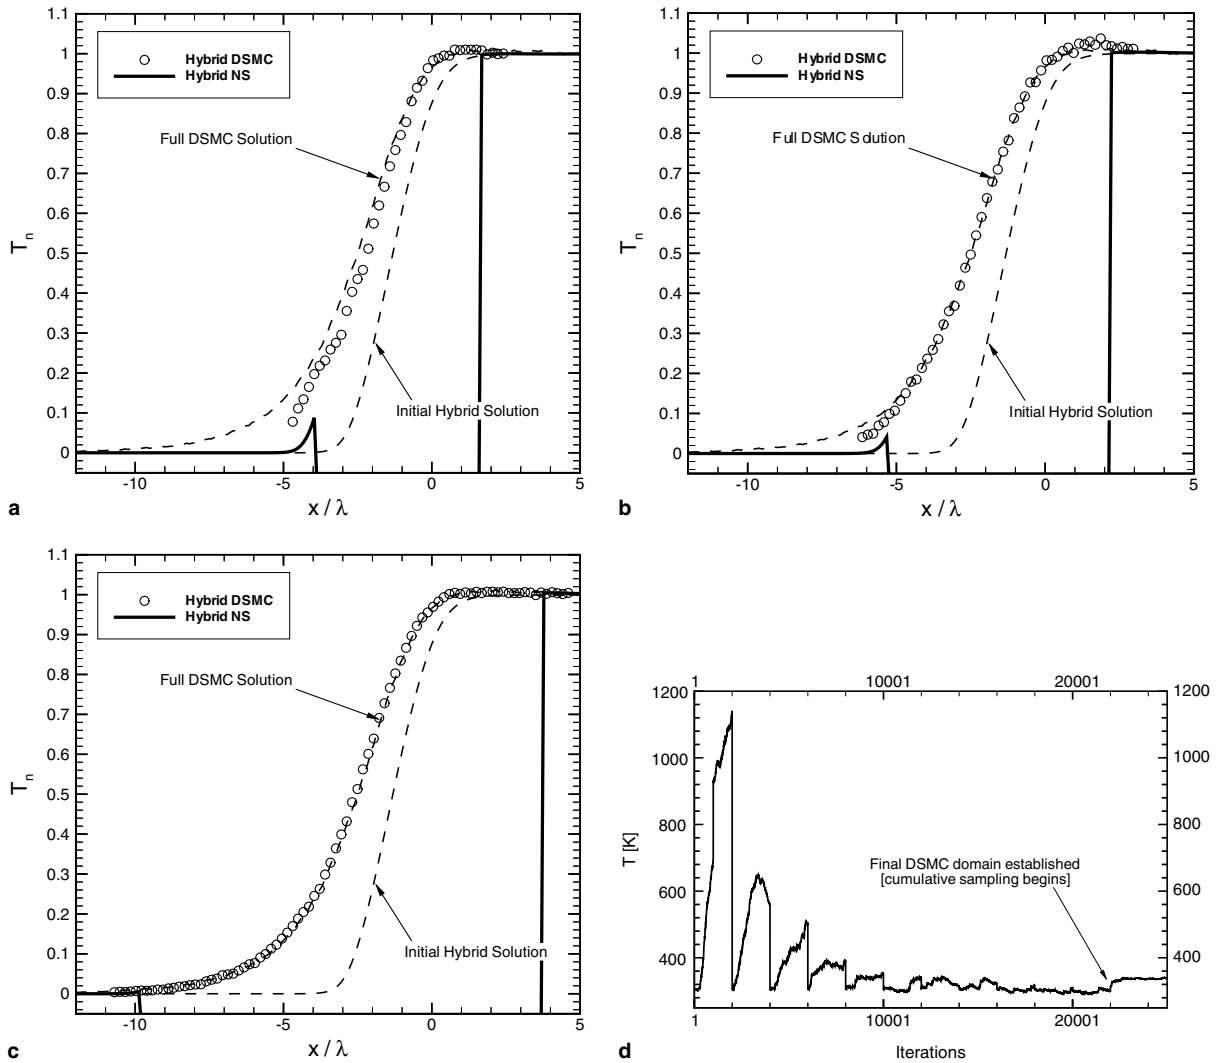


Fig. 7. Progression of a hybrid simulation for a Mach 6.5 shock. (a) After 2000 time-steps. (b) After 4000 time-steps. (c) Final hybrid solution. (d) Temperature variation in pre-shock NS boundary cell.

progressed even further towards the correct solution. The DSMC domain continues to expand every 2000 time-steps, until at 22,000 time-steps application of the breakdown parameter no longer results in a change in the DSMC or NS domain size. This signifies that the DSMC solution is no longer changing in time and has thus reached steady state. At this point the domain sizes are fixed and cumulative DSMC sampling begins in order to reduce the scatter in the hybrid DSMC solution to an arbitrarily low value. Since the DSMC solution is now in steady-state, the sub-relaxation averaging is abandoned and the NS boundary condition is set using the cumulative DSMC average. The hybrid cycle is run for 3000 more time-steps with the finalized interfaces in order to reduce the scatter in the final solution which is shown in Fig. 7(c). Clearly the hybrid solution reproduces exactly the full DSMC solution. Notice that the DSMC portion of the solution transitions very smoothly to the NS portion and that the two solutions agree very well in the buffer region. This signifies that the choice of  $Kn_{GL} = 0.05$  does in fact place the interface in a region of near-equilibrium that can be modeled successfully with the NS equations.

The actual boundary condition seen by the pre-shock NS solution is plotted in Fig. 7(d). Although some of the variations are quite large, they occur over a large number of time-steps. Since the NS boundary condition is updated at every DSMC time-step (strongly coupled), these large variations present no problem to the NS

solver. It should be noted that the discontinuous variations (every 2000 time-steps) do not occur in the *same* NS cell. Recall that every 2000 time-steps, the interfaces and thus the NS boundary cells are relocated. Clearly, there is still scatter present in the NS boundary condition, but its magnitude is far less than the scatter in the raw DSMC data presented in the previous section. This low scatter can be completely eliminated by switching to cumulative DSMC sampling after the final domains are established (in this case after 22,000 time-steps).

## 7. 1D normal shock results for Argon and Nitrogen

In the previous section the number of time-steps between application of the breakdown parameter was set at a large value,  $N_{Br} = 2000$ , in order to *clearly* demonstrate the stages involved in the hybrid cycle. However, one of the previous conclusions drawn is that the correct interfaces should be established as quickly as possible. This may be achieved by applying the breakdown parameter more often. However,  $N_{Br}$  must remain large enough to allow the DSMC solution to change before the breakdown parameter is applied again, otherwise this would waste computational time. For all cases presented in this section, an optimal value of  $N_{Br} = 500$  time-steps was determined to find the correct interfaces in the least amount of time. In fact, on average, the final interfaces are located in 4000 time-steps. It is very important to note that this is approximately the same number of time-steps (only slightly more) that a full DSMC simulation requires to reach steady state. Thus, for these cases, the correct interfaces are effectively determined *while* the DSMC portion is proceeding towards steady state. The remaining hybrid cycle parameters are the same for every case and have already been given above. Using these parameters, the hybrid code is able to produce all of the results in this section by simply changing the gas properties and free-stream velocity.

All hybrid results will be compared with full DSMC simulations, full NS solutions and with experimental data from Alsmeyer [21]. Alsmeyer collected detailed density profiles inside shock waves at various Mach numbers between 1.2 and 9.0. Figs. 8(a),(b) and (c), (d) show the hybrid results in argon gas for Mach numbers of 3.38 and 9.0, respectively. Fig. 8(a) shows the density profiles calculated from full DSMC, full NS and hybrid solvers compared with experimental data for a Mach number of 3.38. Fig. 8(b) depicts profiles of the normalized density gradient for the various methods. The gradient at  $\rho_n = 0.5$  is also referred to as the *reciprocal shock thickness* and will be used later. For this case, both Figs. 8(a) and (b) verify that the hybrid solution reproduces exactly the full DSMC solution. Additionally, the full DSMC solution is seen to accurately predict the experimental results. Similarly, the results for the Mach 9.0 case, depicted in Figs. 8(c) and (d), show the hybrid solution reproducing exactly the full DSMC solution which again agrees very well with experimental data. For this high Mach number case, the final DSMC domain is very large. In fact, the final interfaces extend beyond any significant density gradient (especially upstream of the shock). The reason for this is that the temperature gradient is known to precede the density gradient in high-speed shock waves. Although Alsmeyer did not measure temperature, this phenomenon was seen earlier in Fig. 7(c). Since the breakdown parameter detects the maximum of density, velocity and temperature gradients (see Eq. (2)), the temperature gradient seen in Fig. 7(c) will push the interface upstream. Similar results are presented for nitrogen simulations at shock Mach numbers of 2.0 and 6.1 in Figs. 8(e) and (f), respectively. For the nitrogen cases, full DSMC and full NS results are not plotted, however for Mach numbers ranging from 1.5 to 10.0, the hybrid results are verified to match full DSMC solutions and result again in thicker shock waves than those obtained by full NS solutions. Experimental results exist only for the density profiles and the hybrid results (solid line) shown in Figs. 8(e) and 8(f) agree well with these experimental measurements. Since nitrogen is a diatomic molecule, the DSMC region now contains both a translational temperature ( $T_{TRA}$ ) and a rotational temperature ( $T_{ROT}$ ). No vibrational degrees of freedom are included in the DSMC simulations. It is interesting to note that at the interfaces the gas is close to thermal equilibrium, that is, there is only a small difference between  $T_{TRA}$  and  $T_{ROT}$ . Thus although  $Kn_{GL}$  does not directly account for thermal non-equilibrium, it still positions the interfaces in regions quite close to thermal equilibrium. However, the gas is not exactly in thermal equilibrium and future simulations may either reduce the continuum cutoff parameter below  $Kn_{GL} = 0.05$  or possibly incorporate thermal non-equilibrium into the definition of continuum breakdown.

Full NS, DSMC and hybrid simulations were run for 13 Mach numbers between 1.2 and 10. The reciprocal shock thickness predicted by these methods for argon gas was calculated and compiled into one graph for comparison with experimental data over the entire Mach number range. This result is shown in Fig. 9(a).

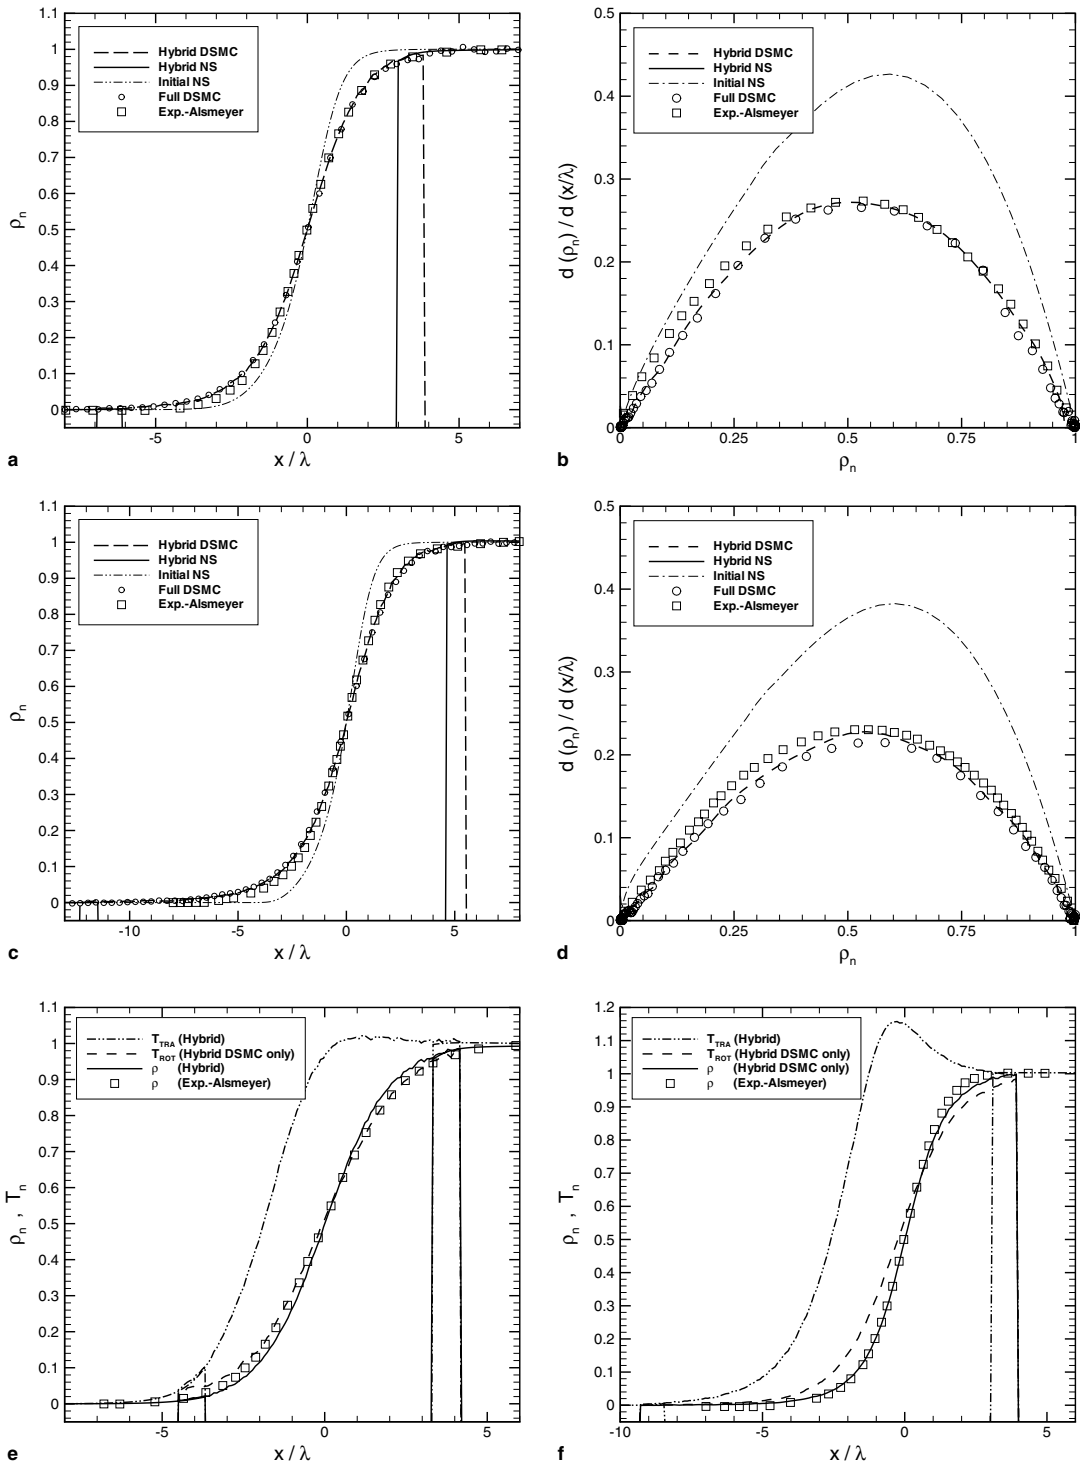


Fig. 8. Hybrid code results for shock waves of various Mach number. (a)  $\rho$ , Mach 3.38 (argon). (b) Reciprocal shock thickness, Mach 3.38. (c)  $\rho$ , Mach 9 (argon). (d) Reciprocal shock thickness, Mach 9. (e)  $\rho$  and  $T$ , Mach 2.0 (nitrogen). (f)  $\rho$  and  $T$ , Mach 6.1 (nitrogen).

Clearly, the NS solver predicts a shock wave that is far too thin, while the DSMC solver, using the viscosity law in Eq. (4), agrees with the experimental results very well. In addition, the hybrid solver is seen to reproduce the DSMC results almost exactly for each Mach number. Finally, it should be noted how all three solvers

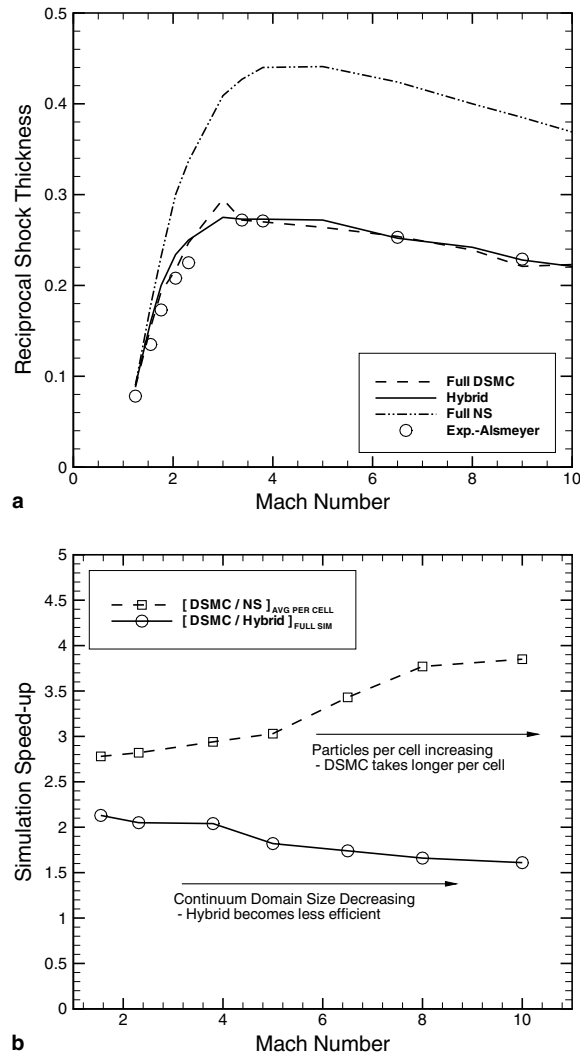


Fig. 9. Accuracy and efficiency of the hybrid code for 1D normal shocks in argon. (a) Reciprocal shock thickness over entire Mach number range. (b) Hybrid efficiency for 1D normal shock waves.

predict the same reciprocal shock thickness when the Mach number is small. In fact, for a Mach 1.244 case, application of the continuum breakdown parameter on the initial NS solution results in no DSMC domain at all. Thus, the hybrid solver immediately returns the full NS solution as the final solution. The NS solution for this case does indeed agree both with experiment and a full DSMC simulation and thus lends further support for using  $Kn_{GL} = 0.05$  as the cutoff for continuum breakdown.

Computational profiling of the hybrid code reveals an average speed-up factor of 2 compared with a full DSMC simulation for 1D shock waves in argon. Clearly, both the size of the overall hybrid domain as well as the average number of particles per DSMC cell will greatly influence the efficiency of the hybrid code and are indeed chosen somewhat arbitrarily for this study. The efficiency comparisons are thus only included for completeness and cannot be readily extended to other physical problems or multi-dimensional simulations. Having said this, Fig. 9(b) displays the increase in efficiency gained by the hybrid code for this study. It is evident that as the Mach number increases, the density in the post-shock region increases as well. Since, on average, there are now more particles per DSMC cell, the ratio of time spent by the DSMC solver for one cell to that spent by the NS solver increases with Mach number. For this reason, the hybrid solver would be expected to increase in efficiency as Mach number increased. However, at the same time, the hybrid particle domain gets

larger as Mach number increases. This reduces the efficiency of the hybrid solver as fewer cells may be simulated as continuum cells. The overall result is that the efficiency degrades only slightly as Mach number increases and a full DSMC simulation remains roughly 2 times slower than a hybrid simulation. According to the profiling, the DSMC solver takes only 3–4 times longer per cell than the NS solver. Thus, even if the entire domain could be represented as continuum cells, the maximum speed-up would be a factor of 3–4 times. It is very important to realize, however, that the increase in efficiency gained by a hybrid solver ultimately comes from simulating far fewer total cells. In this study, the continuum mesh maintained the same density as the particle mesh for simplicity. The true advantage of a hybrid code, as mentioned in the introduction, is that the continuum mesh is able to be far coarser than the particle mesh. In multiple dimensions such a large decrease in the number of cells simulated will result in much larger computational savings. Although this research strongly couples an explicit NS solver to the DSMC method, future research will investigate the benefits of using an implicit NS solver. This will enable significant time-scale decoupling (discussed in Ref. [7]) and therefore allow the continuum cells to be updated far less frequently than the particle cells. The combination of simulating fewer cells for fewer time-steps in large continuum portions of the flow field is expected to greatly increase the computational efficiency of such a hybrid scheme.

## 8. Conclusions

This study presented a detailed analysis of each stage of a hybrid numerical scheme which strongly couples a NS solver to a DSMC solver in order to efficiently model high speed non-equilibrium flows. During the hybrid simulation, due to the inherent difficulty in determining physically valid particle and continuum regions without having access to the final physical solution, the state represented in DSMC boundary cells may be un-physical. In addition, before the interface adapts to a region of near-equilibrium, the velocity distribution in the DSMC boundary cell may be un-physical as well. However, detailed analysis showed that this boundary condition has only a small effect on the interior of the DSMC domain when it has been initialized with particles that correspond to a macroscopic NS solution. Also, including a small overlap or “buffer” region to separate DSMC and NS boundary cells reduced the effect of ill-posed DSMC boundary conditions and inhibited the reflection of this error back into the NS domain. For such steady-state problems, state-based coupling using a sub-relaxation technique to include recent history was found to be an efficient method of controlling statistical scatter when generating NS boundary conditions from particle information. Using a continuum breakdown parameter of  $Kn_{GL} = 0.05$  resulted in less than a 5% discrepancy between continuum and particle solutions; thus re-enforcing previous results. A hybrid code based on these findings was able to reproduce experimental and full DSMC simulation results in half the time for 1D argon and nitrogen shock waves over a large Mach number range. The hybrid code is simple, robust, and only requires the user to modify the gas properties and free-stream conditions before each simulation.

## Acknowledgment

This work is sponsored by the Space Vehicle Transportation Institute, under NASA Grant NCC3-989 with joint sponsorship from the Department of Defense.

## References

- [1] G.A. Bird, *Molecular Gas Dynamics and the Direct Simulation of Gas Flows*, Oxford University Press, New York, 1994.
- [2] W.-L. Wang, I.D. Boyd, Predicting continuum breakdown in hypersonic viscous flows, *Physics of Fluids* 15 (2003) 91–100.
- [3] N.G. Hadjiconstantinou, A.L. Garcia, M.Z. Bazant, G. He, Statistical error in particle simulations of hydrodynamic phenomena, *Journal of Computational Physics* 187 (2003) 274–297.
- [4] Q. Sun, I.D. Boyd, A direct simulation method for subsonic, micro-scale gas flows, *Journal of Computational Physics* 179 (Jul) (2002) 400–425.
- [5] W.-L. Wang, I.D. Boyd, Hybrid DSMC-CFD simulations of hypersonic flow over sharp and blunted bodies, *AIAA Paper 03-3644*, 2003, Presented at the 36th AIAA Thermophysics Conference, Orlando, FL.
- [6] W.-L. Wang, I.D. Boyd, A new energy flux model in the DSMC-IP method for nonequilibrium flows, *AIAA Paper 03-3774*, 2003, Presented at the 36th AIAA Thermophysics Conference, Orlando, FL.



- [7] H.S. Wijesinghe, N.G. Hadjiconstantinou, A discussion of hybrid atomistic-continuum methods for multiscale hydrodynamics, *International Journal for Multiscale Computational Engineering* 2 (2004).
- [8] D.B. Hash, H.A. Hassan, A decoupled DSMC/Navier–Stokes analysis of a transitional flow experiment, AIAA Paper 96-0353, January 1996, Presented at the 34th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV.
- [9] D.B. Hash, H.A. Hassan, Assessment of schemes for coupling Monte Carlo and Navier–Stokes solution methods, *Journal of Thermophysics and Heat Transfer* 10 (1996) 242–249.
- [10] F. Golse, Applications of the Boltzmann equation within the context of upper atmosphere vehicle aerodynamics, *Computer Methods in Applied Mechanics and Engineering* 75 (1989) 299–316.
- [11] D.C. Wadsworth, D.A. Erwin, One-dimensional hybrid continuum/particle simulation approach for rarefied hypersonic flows, AIAA Paper 90-1690, June 1990, Presented at the 5th AIAA/ASME Joint Thermophysics and Heat Transfer Conference, Seattle, WA.
- [12] D.C. Wadsworth, D.A. Erwin, One-dimensional hybrid continuum/particle approach for rarefied flows, AIAA Paper 92-2975, July 1992, Presented at the 23rd AIAA Plasmadynamics and Lasers Conference, Nashville, TN.
- [13] P. Le Tallec, F. Mallinger, Coupling Boltzmann and Navier–Stokes equations by half fluxes, *Journal of Computational Physics* 136 (1997) 51–67.
- [14] R. Roveda, D.B. Goldstein, P.L. Varghese, Hybrid Euler/particle approach for continuum/rarefied flows, *Journal of Spacecraft and Rockets* 35 (3) (1998) 258–265.
- [15] R. Roveda, D.B. Goldstein, P.L. Varghese, Hybrid Euler/direct simulation Monte Carlo calculation of unsteady slit flow, *Journal of Spacecraft and Rockets* 37 (6) (2000) 753–760.
- [16] H.S. Wijesinghe, R.D. Hornung, A.L. Garcia, N.G. Hadjiconstantinou, Three-dimensional hybrid continuum-atomistic simulations for multiscale hydrodynamics, *Journal of Fluids Engineering* 126 (2004) 768–777.
- [17] L. Pareschi, R.E. Caflisch, An implicit Monte Carlo method for rarefied gas dynamics, *Journal of Computational Physics* 154 (1999) 90–116.
- [18] R.W. MacCormack, G.V. Candler, The solution of the Navier–Stokes equations using Gauss–Seidel line relaxation, *Computers and Fluids* 17 (1989) 135–150.
- [19] S. Dietrich, I.D. Boyd, Scalar and parallel optimized implementation of the direct simulation Monte Carlo method, *Journal of Computational Physics* 126 (1996) 328–342.
- [20] I.D. Boyd, Analysis of rotational nonequilibrium in standing shock waves of nitrogen, *AIAA Journal* 28 (1990) 1997–1999.
- [21] H. Alsmeyer, Density profiles in argon and nitrogen shock waves measured by the absorption of an electron beam, *Journal of Fluid Mechanics* 74 (1976) 497–523.
- [22] A.L. Garcia, B.J. Alder, Generation of the Chapman–Enskog distribution, *Journal of Computational Physics* 140 (1998) 66–70.
- [23] Q. Sun, I.D. Boyd, Evaluation of macroscopic properties in the direct simulation Monte Carlo method, *Journal of Thermophysics and Heat Transfer* 19 (3) (2005) 329–335.