

Hybrid Particle–Continuum Simulations of Hypersonic Flow over a Hollow-Cylinder-Flare Geometry

Thomas E. Schwartzenruber*
University of Minnesota, Minneapolis, Minnesota 55455
Leonardo C. Scalabrin†
ESI Group North America, Huntsville, Alabama 35806
and
Iain D. Boyd‡
University of Michigan, Ann Arbor, Michigan 48109

DOI: 10.2514/1.36681

A modular particle–continuum numerical method is used to simulate steady-state hypersonic flow over a hollow-cylinder-flare geometry. The resulting flowfield involves a mixture of rarefied nonequilibrium flow and high-density continuum flow. The hybrid particle–continuum method loosely couples direct simulation Monte Carlo and Navier–Stokes methods, which operate in different regions, use different mesh densities, and are updated using different-sized time steps. Hybrid numerical results are compared with full particle and full continuum simulations as well as with experimental data. The hybrid particle–continuum simulations are demonstrated to reproduce experimental and full particle simulation results for surface and flowfield properties including velocity slip, temperature jump, thermal nonequilibrium, heating rates, and pressure distributions with high accuracy. The hybrid method, which uses particle simulation next to the surface, is also shown to predict accurate heating rates even when a highly dissipative numerical scheme is used for the continuum solver. For this particular flow, a hybrid simulation is obtained with modest computational savings over full particle simulation.

Nomenclature

Br_{cutoff}	=	cutoff value of the continuum-breakdown parameter	V	=	bulk Y velocity of gas, m/s
C_p	=	coefficient of pressure, $(p - p_\infty)/(0.5\rho_\infty U_\infty^2)$	$ V $	=	velocity magnitude, $\sqrt{U^2 + V^2}$, m/s
d	=	diameter of the molecule, m	X, Y	=	Cartesian coordinates, m
Kn	=	Knudsen number, λ/l_c	x_{sep}	=	X coordinate of flow separation on cylinder surface, m
Kn_{GL}	=	gradient-length Knudsen number	Δt_{ss}	=	minimum simulation time before sampling can begin
k	=	Boltzmann constant, 1.38×10^{-23} , J/K	λ	=	mean free path, m
l_c	=	gradient length, $Q/ \nabla Q $	μ	=	viscosity, (Ns)/m ²
M	=	Mach number	ρ	=	density, kg/m ³
m	=	mass of molecule, kg	τ_c	=	mean free time, s
n_{overlap}	=	number of overlap cells where both particle and continuum methods are used	ω	=	power-law temperature dependence of μ
p	=	pressure, N/m ²	<i>Subscripts</i>		
Q	=	macroscopic property	ref	=	reference value of the parameter
q	=	heat transfer per unit area, W/cm ²	∞	=	freestream value of the parameter
R	=	cylinder or flare radius, m			
St	=	coefficient of heat transfer (Stanton number), $q/(0.5\rho_\infty U_\infty^3)$			
T	=	continuum combined translational–rotational temperature, K			
T_{rot}	=	rotational temperature, K			
T_{tra}	=	translational temperature, K			
T_{vib}	=	vibrational temperature, K			
T_{wall}	=	wall temperature, K			
U	=	bulk X velocity of gas, m/s			

Presented as Paper 3892 at the 39th AIAA Thermophysics Conference, Miami, FL, 25–28 June 2007; received 15 January 2008; revision received 7 April 2008; accepted for publication 8 April 2008. Copyright © 2008 by Thomas E. Schwartzenruber, Leonardo C. Scalabrin, and Iain D. Boyd. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0001-1452/08 \$10.00 in correspondence with the CCC.

*Assistant Professor, Department of Aerospace Engineering and Mechanics; schwartz@aem.umn.edu. Member AIAA

†Software Development Engineer; lsa@esi-group-na.com.

‡Professor, Department of Aerospace Engineering; iainboyd@umich.edu. Associate Fellow AIAA.

I. Introduction

HYPERSONIC flow over a hollow-cylinder-flare geometry generates interactions between shock waves and boundary layers, resulting in a complex and highly nonlinear flowfield. Such hypersonic interaction flows occur in the inlets of supersonic and hypersonic propulsion devices and in the vicinity of control surfaces on hypersonic vehicles. In such flows, shock–shock interactions often produce reflected shock waves that impinge on the boundary layer next to the vehicle surface. The interaction between a strong shock wave and a boundary layer may then cause the flow to separate and form a region of high-pressure recirculating gas next to the surface. High-speed flow hitting such a recirculating region significantly alters shock structures and interaction regions, which in turn influence the extent of flow separation. Peak aerothermal loads are observed at the location of shock impingement on the vehicle surface, and therefore accurate prediction of this phenomenon is important in the design of a hypersonic vehicle. In addition, the heating and frictional loads measured within a recirculating flow region are significantly different from when the flow is purely attached. Such differences may reduce the effectiveness of a control surface, and thus accurate prediction of these phenomena is important during the design of hypersonic vehicles. For most

blunt-body flows, the strong bow shock wave does not interact with the vehicle surface. However, various blunt-body deceleration techniques that use supersonic parachutes or ballutes involve highly complex flowfields in which many regions of shock–shock and shock–boundary-layer interactions are present [1,2].

For hypersonic interaction flows at high altitudes, the local mean free path between gas–particle collisions may approach the characteristic lengths associated with the high-gradient regions. Under such conditions, a continuum description of the gas may no longer be accurate and a particle or kinetic description of the gas may be required. The direct simulation Monte Carlo (DSMC) particle method [3] is currently the most popular method for high-speed nonequilibrium flows. DSMC directly simulates a gas flow by tracking a representative number of simulated molecules through a computational mesh, which collide with each other as well as with the vehicle surface. A limitation of the DSMC method is that it becomes computationally expensive in the continuum regime due to correspondingly small molecular spatial and temporal scales, which must be resolved to the order of the mean free path λ and the mean free time τ_c , respectively. However, under such continuum conditions, the Navier–Stokes (NS) equations accurately model the flow and can be solved efficiently. For this reason, various researchers have proposed hybrid numerical methods that adaptively reposition particle and continuum computational domains within a single hybrid simulation and couple particle and continuum regions by transferring information across the interface [4–10]. A more detailed overview of the various methods and comparison with the hybrid numerical algorithm used in this paper can be found in [11]. In this study, the modular particle–continuum (MPC) algorithm [11] is used to simulate a steady-state hypersonic interaction flow. The MPC method was initially developed and tested for 1-D normal shock waves [12] as well as for hypersonic flow over a 2-D-cylinder geometry [13]. The MPC algorithm loosely couples DSMC and NS regions, which have different mesh densities and are updated using different time steps. Such an approach enables both spatial- and temporal-scale decoupling while lending itself to a modular implementation [14], which uses existing state-of-the-art DSMC and NS codes (unmodified) within the hybrid code. For flow over a 2-D cylinder, the MPC method has been shown to reproduce the flowfields, surface properties, and local velocity distributions predicted by full DSMC simulations 1.4 to 3.3 times faster while requiring less memory [13,14].

In this paper, the MPC method is used to simulate hypersonic flow over a hollow-cylinder-flare geometry. This axisymmetric, highly nonlinear flow provides a significant challenge for the MPC algorithm, compared with the previous MPC research and other cited hybrid research. High-quality experimental data are available for this flow problem and have been used extensively to validate both DSMC and NS simulations in the literature. A detailed description of the problem and a comparison of experimental results with full NS and full DSMC simulations are presented in Sec. II. A brief overview of the MPC numerical cycle is presented by showing the progression of a hybrid simulation in Sec. III. The accuracy of the final MPC solution with respect to flowfield features, velocity slip, temperature jump, thermal nonequilibrium, and surface properties is investigated in Sec. IV. The computational efficiency gained over full DSMC simulation is also discussed in Sec. IV, and conclusions are presented in Sec. V.

II. Full Navier–Stokes and DSMC Simulation

Particle regions are simulated using MONACO [15], a general cell-based implementation of the DSMC method [3]. The variable-hard-sphere collision model is employed, which results in the following macroscopic viscosity model [3]:

$$\mu = \mu_{ref} \left(\frac{T}{T_{ref}} \right)^\omega, \quad \mu_{ref} = \frac{15\sqrt{\pi mkT_{ref}}}{2\pi d_{ref}^2(5-2\omega)(7-2\omega)} \quad (1)$$

All numerical results presented in this paper are for diatomic nitrogen with a reference diameter of $d_{ref} = 4.17 \times 10^{-10}$ m at a reference

temperature $T_{ref} = 273$ K. The power-law exponent ω is set equal to 0.75, m is the mass of an N_2 molecule, and k is the Boltzmann constant. MONACO employs the variable rotational-energy-exchange probability model of Boyd [16], where the reference temperature for rotational energy exchange is specified as 91.5 K and the maximum rotational collision number is 18.1. Energy transfer to vibrational modes is not considered.

Continuum regions are simulated using the LeMANS code [17]. For the results of this paper, it is assumed that rotational and translational energy modes can be described by a single temperature T in continuum regions. The vibrational energy mode is not considered. The resulting governing equations are the well-known 2-D laminar compressible NS equations. The viscosity in the NS solver is modeled using Eq. (1) to exactly match the viscosity model used in DSMC. LeMANS solves this set of equations using a finite volume formulation. Unless otherwise noted, the inviscid fluxes between the mesh volumes are discretized using a modified form of the Steger–Warming flux vector-splitting [18,19], which is less dissipative than the original form. The modified form is thus adequate to calculate boundary layers, and the scheme switches back to the original form of Steger–Warming near shock waves, where added numerical dissipation is needed for stability. The viscous terms are calculated using the values of properties at the cell centers and at the nodes. The time integration is performed using a point-implicit method. Finally, for continuum-NS simulation, no-slip conditions are applied to both velocity and temperature on all surfaces. Specific details of the continuum numerical method are contained in [17].

The hollow-cylinder-flare geometry is detailed at the top of Fig. 1. The viscous interaction at the sharp leading edge generates a boundary layer and weak shock wave, which then interact with a stronger shock generated by the flow reattachment downstream of the flare junction. This shock–shock interaction, which occurs near the surface of the flare after the junction, in turn affects the size of the recirculating region. This strongly coupled and highly nonlinear flow provides a challenging case for the validation of both NS and DSMC methods. The flow conditions investigated are those of run 11 performed in the Large Energy National Shock (LENS) facility at the Calspan—University of Buffalo Research Center (CUBRC). The experimental results were made public after a blind code-validation exercise at the 2001 AIAA Aerospace Sciences Meeting in Reno, NV [20,21]. During the conference, invited numerical solutions using both NS and DSMC methods were presented by various researchers, including Candler et al. [22], Gnoffo [23], Moss [24], Kato and Tannehill [25], Boyd and Wang [26], and Roy et al. [27]. For the run-11 conditions, Navier–Stokes simulations [22] predicted a larger




Fig. 1 Hollow-cylinder-flare geometry (dimensions in mm) and variation of λ .




Fig. 2 Grid convergence studies for the hollow-cylinder-flare geometry.

recirculation region than measured experimentally. The same simulations also highlighted the sensitivity of the solution to grid resolution, numerical flux function, and limiters used in the NS code. Further investigation by Candler et al. [28] determined that the freestream conditions had a significant amount of thermal nonequilibrium due to frozen vibrational relaxation during the rapid nozzle expansion in the LENS facility. As a result, new revised freestream conditions were defined for this case. The current study uses the following revised flow conditions for CUBRC run 11 with a hollow-cylinder flare: $M_\infty = 12.4$, $U_\infty = 2484.1$ m/s, $\rho_\infty = 5.566 \times 10^{-4}$ kg/m³, $T_\infty = 95.6$ K, $T_{\infty \text{vib}} = 2486.7$ K (not used), and $T_{\text{wall}} = 297.2$ K.

It should be noted that the vibrational energy mode is assumed to be frozen throughout the flow, and the vibrational temperature is therefore not used in any NS, DSMC, or hybrid simulations presented in this paper. The LeMANS code, as already described, is used to solve the NS equations for the hollow-cylinder-flare geometry detailed in Fig. 1, with the preceding flow conditions. Solutions are obtained on three mesh densities of 550×125 , 1100×250 , and 2200×500 cells, corresponding to mesh 1, mesh 2, and mesh 3, respectively. Results for the coefficient of pressure [$C_p = (p - p_\infty)/(0.5\rho_\infty u_\infty^2)$] are displayed in Fig. 2a to verify grid independence and compare the NS solution with experimental results. As seen in Fig. 2a, the solutions on mesh 2 and mesh 3 are virtually identical and the NS equations predict a recirculation region significantly larger than measured experimentally. As a result, the flare shock forms further upstream, which causes the strong interaction region, and thus the peak surface pressure, to occur further downstream than observed experimentally. It will be shown later that the NS equations (using no-slip boundary conditions) also significantly overpredict surface heating and shear stress at the sharp leading edge of the hollow cylinder. The surface pressure distribution predicted by LeMANS in Fig. 2a is in excellent agreement with simulations performed by Candler et al. [28] with respect to the precise location of flow separation, the peak surface pressure, and the overprediction of surface pressure along the flare.

In the blind code-validation exercise, the DSMC method well predicted leading-edge surface properties and the size of the separation region [21]. The largest discrepancy with experimental results was found downstream of the flare junction, where λ becomes small. This can be seen at the bottom of Fig. 1, in which the hard-sphere value of λ is seen to decrease by an order of magnitude after the flare junction. For this reason, fine resolution (cell size and time step) is required for the flare region in a DSMC simulation. Subsequent DSMC simulations [29] have improved the resolution

and produced excellent results, compared with the experimental data. The MONACO code, as already described, is used to calculate the flowfield and surface properties for the preceding run-11 flow conditions using both mesh 2 and mesh 3. The least resolved region in mesh 2 and mesh 3 is at the surface along the flare, where cell dimensions are 4λ and 2λ , respectively, in the direction normal to the wall. If necessary, cells are further subdivided into collision subcells, which are less than 1λ in dimension to satisfy the resolution requirements for the DSMC method. The mesh-3 simulation uses approximately 4 times as many simulated particles as the mesh-2 simulation and both contain at least 15 particles per cell. The flow is axisymmetric, where cell volumes in DSMC become larger as distance from the axis increases and therefore contain more particles. However, because the entire computational domain is located away from the axis of symmetry, the effect is minimal and no cell weighting is used for DSMC simulation of the hollow-cylinder flare. Constant DSMC time steps of 5×10^{-9} s are used, which are verified to be less than $0.6 \times \tau_c$ everywhere. In addition to cell size and time-step restrictions, a DSMC simulation must be allowed to reach steady state before sampling of the solution begins. To demonstrate this, Fig. 2b shows the MONACO solution on mesh 2 sampled at various times. A close-up view of the separation region clearly demonstrates that the size of the separation region increases significantly between the sampling periods of 0.5–0.75 ms and 1.5–1.75 ms and remains relatively constant thereafter. The precise location of flow separation is best determined as the location on the surface where the shear stress becomes negative. Although profiles of shear stress are not shown here (and were not measured experimentally), the point of separation x_{sep} predicted by MONACO for each sampling interval is listed in Table 1. By noting the percentage change in x_{sep} per millisecond, it is evident that the size of the separation region has effectively reached steady state after 1.5 ms. DSMC simulations of the same case performed by Moss and Bird [29] draw the same conclusion and predict a level of agreement with experimental data similar to the

Table 1 DSMC prediction for the location of flow separation and approach to steady state

Sampling interval, ms	x_{sep} , m	% change, 1/ms
0.5–0.75 ms	0.09316	—
1.5–1.75 ms	0.08936	4.0%
2.0–2.25 ms	0.08856	1.8%
3.0–3.25 ms	0.08806	0.6%

MONACO results in Fig. 2b. Compared with the NS results in Fig. 2a, DSMC is seen to better predict the extent of flow separation for the run-11 conditions. Finally, comparison of the solutions obtained on mesh 2 and mesh 3 (for the same sampling interval) demonstrates that mesh 2 provides adequate resolution for a DSMC simulation. Although the mesh-3 solution predicts a slightly higher surface pressure inside the separated region, the location of flow separation and the remainder of the profile agree very well with the mesh-2 solution. It is interesting to note that for the run-11 conditions, the mesh density required to obtain a grid-independent NS solution (mesh 2) is the same as that required by a corresponding DSMC simulation.

The importance of adequate resolution for collision cells is portrayed in Fig. 3. Here, the mesh-2 MONACO simulation is repeated without using collision subcells. Because collision cells next to the flare surface are now up to 4λ in the direction normal to the surface, substantial numerical diffusion is introduced into the boundary layer, causing an overprediction of 30% in the peak heating rate. It should be noted that previous MPC results, presented in [30], did not use this subcell feature within the hybrid code and thus overpredicted the heat transfer by a similar margin. For the results presented in this paper, the MPC method uses the identical subcell algorithm used by MONACO for full DSMC simulations.

III. Modular Particle–Continuum Numerical Method

A. Problem Setup

The MPC numerical method used in this paper is completely detailed in [11] and is now briefly outlined. The MPC method begins with a NS solution obtained on a mesh designed to solve the NS equations. An estimate of local continuum breakdown is provided by the gradient-length Knudsen number [31]:

$$Kn_{GL} = \max\left(\frac{\lambda}{Q} |\nabla Q|\right) \quad (2)$$

where λ is calculated consistent with the variable-hard-sphere model and the maximum is taken over the flow quantities Q of interest: density ρ , temperature T , and velocity magnitude $|V|$. Cells in which $Kn_{GL} > Br_{cutoff}$ are labeled as DSMC cells and the remainder are labeled as NS cells, thus defining an interface between the two regions. Previous studies have recommended [31] and validated [11–13] a cutoff value of $Br_{cutoff} = 0.05$. A slightly more conservative value of $Br_{cutoff} = 0.03$ is used for the MPC simulations presented in this paper. To create an overlap region in which both methods are used, the DSMC regions are then extended further into the continuum region by a number ($n_{overlap}$) of cells. For the




Fig. 3 Effect of collision subcells on the predicted heating-rate profile.

hollow-cylinder-flare simulation, a value of $n_{overlap} = 4$ is used. Extended DSMC regions are further surrounded by one additional DSMC boundary cell, and NS regions are surrounded by one additional NS boundary cell. Now that all regions and boundaries have been initialized, the mesh inside DSMC regions is refined to the local value of λ . This procedure is described in [14], which also uses flow gradient information from the NS solution to set more stringent refinement in the direction of flow gradients and to relax the refinement in other directions. The initial particle region (including the overlap) for the hollow-cylinder flare is shown at the top of Fig. 4, and a portion of the mesh is depicted at the bottom of Fig. 4. Initially, particles are generated in all DSMC cells (using a Chapman–Enskog velocity distribution [32]) such that they are consistent with the NS solution. Before each DSMC simulation time step, all particles within DSMC boundary cells are first deleted and then regenerated to maintain the desired boundary conditions, consistent with the corresponding NS solution. As the DSMC regions iterate and particles collide with each other and with surfaces, the DSMC solution proceeds away from the initial continuum-NS solution toward the correct nonequilibrium solution. As this progression and relaxation occurs, DSMC and NS regions are loosely coupled and interface locations are periodically updated throughout the simulation. Before the NS portions of the hybrid numerical cycle, the macroscopic state variables are set in each NS boundary cell using a subrelaxation average [33] of the molecular properties in each corresponding DSMC cell. As an example, the final steady-state interface locations for the hollow-cylinder-flare simulation are shown at the top of Fig. 4. Finally, because the creation of particle and continuum regions as well as the mesh refinement are automatic in an MPC simulation, the user time required to set up an MPC simulation involves creating the NS mesh and is therefore identical to the setup time required for a full NS simulation of the problem.

B. Numerical Cycle and Progression of the Hybrid Solution

The loosely coupled numerical cycle employed in the MPC method [11] can be summarized as follows:

- 1) Using Kn_{GL} , set up the initial DSMC and NS domains based on an initial NS solution and refine DSMC regions to λ . Generate particles throughout the entire DSMC domain.
- 2) Allow DSMC regions to progress and relax with the current boundary conditions, while adaptively repositioning the interfaces (using the overlap region) without using the NS solver at all.
- 3) After the DSMC solution and interfaces stop changing, use the current DSMC solution to set the NS boundary conditions. Significantly converge the NS region. If the new NS solution changes




Fig. 4 Interface locations (top) and mesh refinement (bottom) for the hollow-cylinder flare.

the state in any DSMC boundary cell, then return to 2. Else, if the new NS solution does not change the state in any DSMC boundary cell, then continue to 4.

4) Because the NS region is converged *and* the DSMC solution is no longer changing, the interfaces will no longer change and steady state has been reached. Lock the interfaces, remove the overlap regions, and cycle both the DSMC and NS methods (coupling occasionally) until the DSMC scatter and NS residual fall below threshold values.

The hybrid numerical cycle used by Wu et al. [10] is similar to the preceding MPC cycle and adds further evidence that a loosely coupled approach is well suited for high-speed steady-state flows. One difference is that whereas Wu et al. use a cumulative (sampled) average inside time-varying DSMC regions, the MPC method uses a subrelaxation average [33] to track macroscopic variations with low scatter in DSMC regions. Tracking variations in DSMC regions with the subrelaxation average provides very useful information to the MPC cycle regarding when to transfer information and allows the interfaces to adapt *while* DSMC regions are progressing toward steady state [12]. A second important difference lies in step 2 of the MPC cycle: allowing the interfaces to stop moving before any information is transferred. This ensures that interfaces are indeed located in near-continuum flow as specified by the continuum-breakdown parameter Kn_{GL} before any information is transferred into the continuum region. This has been shown to be especially important in the vicinity of strong shock waves [11].

Figure 5 shows MPC solutions at various stages during step 2 of the preceding MPC cycle for the hollow-cylinder-flare case. In Fig. 5, the translational temperature profile along cut C2 (shown previously in Fig. 4) is plotted at four intervals: t_0 , t_1 , t_2 , and t_3 . Because the initial boundary-layer development (at the leading edge) is enclosed entirely by a DSMC region (see the top of Fig. 4), both temperature jump and velocity slip at the surface are captured. As a result, the weak leading-edge shock predicted by DSMC lies closer to the surface than predicted by the NS equations employing no-slip conditions. As DSMC regions are allowed to iterate (from time t_0 through t_3), the solution progresses significantly toward the final correct solution predicted by a full MONACO (DSMC) simulation, including velocity slip, temperature jump, and thermal non-equilibrium. As this progression occurs, the interface locations are updated, which (as seen in Fig. 4) actually follow the shock movement toward the lower position predicted by full DSMC. The movement of the hybrid interfaces is entirely facilitated by the overlap region, which allows for local differences between DSMC and NS solutions to develop and be detected by the continuum-breakdown parameter [11]. A close-up look at the shock region in

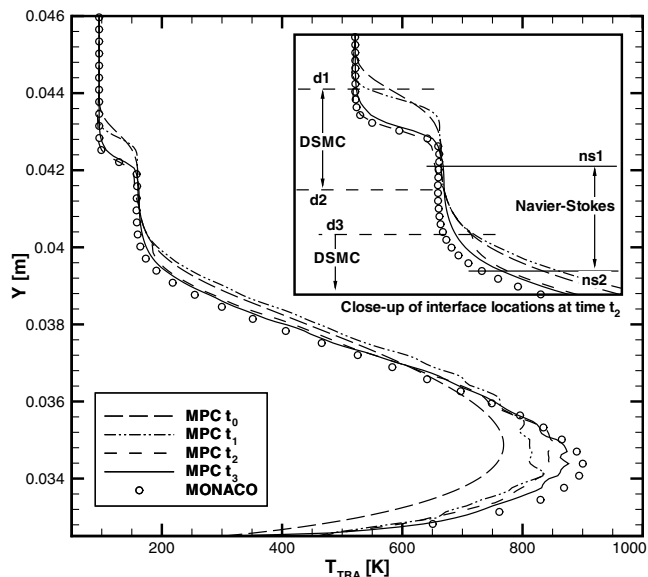


Fig. 5 Progression and relaxation of DSMC regions within the MPC cycle.

Fig. 5 shows three overlapping regions: DSMC inside the shock (lines $d1-d2$), NS in the continuum region between the shock and boundary layer (lines $ns1-ns2$), and DSMC in the boundary layer (below line $d3$). The MPC solution is initialized to the NS solution (the profile at time t_0) and after 2000 DSMC time steps has progressed to the profile labeled as t_1 , which has not clearly progressed in the proper direction at this point. However, in step 2 of the MPC cycle, as long as the DSMC regions and interfaces continue to change, no information is transferred to the NS regions, and DSMC iterations continue. After 4000 more DSMC time steps, the solution reaches the profile labeled as time t_2 . At this point, variations inside the DSMC regions have ceased and repeated implementation of step 2 no longer results in movement of the interfaces. The MPC cycle has reached step 3 and information is now transferred to the NS regions. It is important to notice the role of the overlap region between lines $d3$ and $ns2$. Although the t_2 profile still uses the (incorrect) initial NS solution for conditions at line $d3$, within this overlap region, the MPC solution recovers and by line $ns2$ is very close to the final DSMC result. Thus, during step 3 of the MPC cycle, when information is transferred back to the NS regions, the new NS boundary condition is taken on line $ns2$. After convergence of the NS regions, the profile labeled t_3 is obtained, which shows significant progression toward the correct DSMC profile. At this point, step 3 is complete and the cycle returns to step 2 with new (and more accurate) DSMC boundary conditions supplied from the NS regions at lines $d1$, $d2$, and $d3$.

IV. MPC Simulation Results

Previous research focusing on 1-D normal shock waves [12] and hypersonic flow over a 2-D cylinder [11,13] demonstrated the accuracy of the MPC method for computing flowfield data, velocity distributions, and surface properties. The hypersonic interaction flow studied in this paper provides a significant challenge for the MPC algorithm, compared with the previous MPC research. The results of both the previous and current section demonstrate further essential capabilities of the MPC method. Such capabilities include the ability to adapt interfaces to moving flow features, information transfer that significantly alters the solution in hybrid NS regions, and the ability to handle axisymmetric flows.

The current implementation of the MPC method is restricted to serial (single CPU) processes. Because of computational time and memory restrictions, the initial NS solution used is that obtained on mesh 1. Particle regions are further refined to the local value of λ , which results in a mesh density slightly higher than that of mesh 2 for regions in which the DSMC method is used. A portion of this hybrid mesh is shown in Fig. 4. As seen in Fig. 2a, the NS solution on mesh 1 is very close to the grid-converged result on mesh 2, except that the separation region is slightly smaller. The size of the separation region on mesh 1 is still sufficiently larger than that predicted by DSMC and thus still provides a challenge for the MPC method. Beginning the MPC simulation using the NS solution on mesh 1 also allows testing of the mesh refinement procedures for the hollow-cylinder-flare problem.

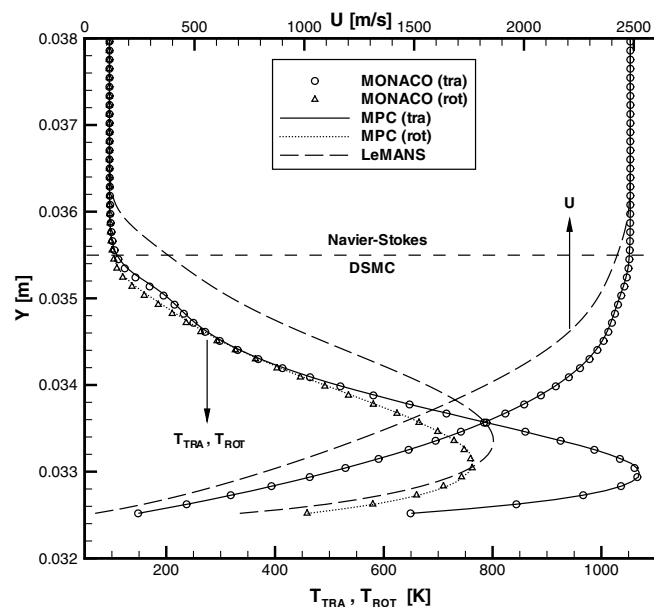
A. Flowfield Results

To provide detailed comparisons between full NS, DSMC, and MPC simulations, flow properties are extracted along various cuts normal to the cylinder-flare surface. The results along cuts C1, C2, C4, and C5 (from Fig. 4) are shown in Figs. 6a–6d, respectively. Near the leading edge, the breakdown parameter sets up a DSMC region that completely envelops the initial boundary-layer growth and leading-edge shock. As seen in Fig. 6a, the shock-boundary-layer height predicted at location C1 by full DSMC is lower than that predicted by the initial NS solution. In addition, DSMC is seen to predict significant velocity slip, temperature jump, and thermal nonequilibrium at this location. The MPC method is seen to very accurately reproduce full DSMC results near the leading edge and, at the same time, has lowered the interface location from its initial position of approximately $Y = 0.0364$ m to a final position of $Y = 0.0355$ m. At location C2, the shock and boundary layer


become more distinct, and Fig. 6b shows that the MPC method now solves the NS equations in a continuum region between the shock and boundary layer. Again, the MPC method captures the new shock location, as well as the thermal nonequilibrium and temperature jump predicted by full DSMC, very accurately. The MPC profile in Fig. 6b also demonstrates that the changing DSMC regions have successfully set new NS boundary conditions that significantly shift the continuum solution between the shock and boundary layer. The progression of the MPC solution at cut C2 was detailed previously in Sec. III. It is important to note that the flow is in thermal equilibrium in the continuum region between the shock and boundary layer and is adequately modeled by the NS equations. This adds further credibility to the use of the gradient-length Knudsen number [31] to detect continuum breakdown. Figure 6c depicts the temperature profiles at location C4, which lies after boundary-layer reattachment, just downstream of the shock–shock-interaction region. Here, DSMC predicts the translational temperature increase (due to the strong shock) to occur further upstream (or higher) than predicted by the NS equations. In addition, DSMC shows a high degree of thermal nonequilibrium inside the shock region. Both of these effects are

captured by the MPC method. However, it should also be noted that other than inside the shock, full DSMC and full NS solutions agree quite well at this location, especially close to the surface. This indicates that at this point on the flare, the flow may lie in the continuum regime. Finally, Fig. 6d depicts the temperature profiles well downstream of the shock–shock interaction, where a strong shock is distinctly separated from the boundary layer. Here, the MPC method solves the NS equations in the large continuum region between the shock and boundary layer. As seen in Fig. 6d, the flow is in thermal equilibrium in this region and the NS equations accurately model a large portion of the flowfield.


Location C3 cuts through the recirculating flow region, which is evident in Fig. 7, in which the x velocity U is negative. Here, DSMC is seen to predict a smaller recirculation region than the initial NS solution. Similar to a full DSMC simulation, an MPC simulation also requires significant time for the recirculating flow to transition from the initial NS solution and reach a steady state. However, because the MPC interface envelops the entire recirculation region, the interface locations do not move at all during this lengthy time. This causes a problem for the MPC cycle outlined in the previous section, because




a) Temperature and velocity profiles along C1



b) Temperature profiles along C2



c) Temperature profiles along C4



d) Temperature profiles along C5

Fig. 6 Full DSMC, NS, and MPC flowfield results at various locations along the hollow-cylinder flare.

little or no interface movement signifies steady state, at which point the MPC cycle wants to lock the interfaces and begin sampling. To allow the MPC cycle to handle slow recirculating flows, a user-defined parameter Δt_{ss} is introduced that specifies the minimum number of iterations (or minimum physical time) before sampling begins. This parameter is analogous to the user-defined parameter found in most DSMC implementations that specifies when sampling is to begin. The MPC solution along cut C3 is plotted for various values of Δt_{ss} , and Fig. 7 shows that a time of 0.3 ms is sufficient for the MPC simulation to reach steady state. This is significantly less than the time required for a full DSMC simulation to reach steady state (see Fig. 2b) and is entirely attributed to the fact that an MPC simulation *begins* with particles initialized to match the full NS solution.

B. Surface Properties

The coefficient of pressure C_p predicted by full DSMC, NS, and MPC simulations is plotted in Fig. 8a. The size of the recirculation region is successfully reduced by the MPC method and agrees well

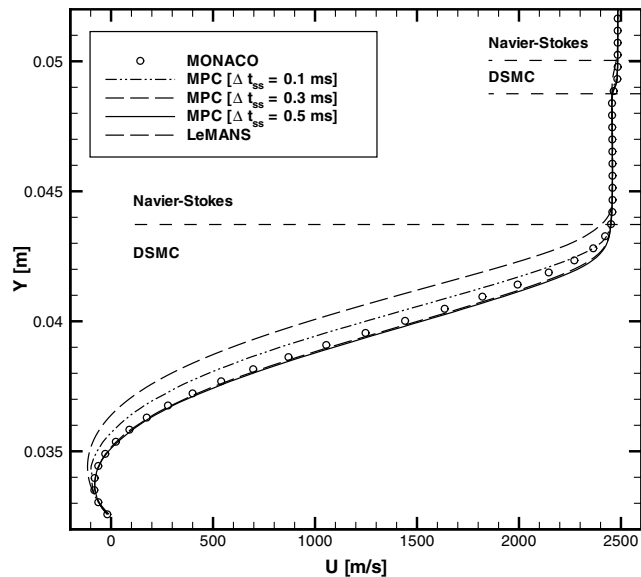




Fig. 7 Velocity profile along C3.



a) Coefficient of pressure



b) Coefficient of heat transfer (Stanton number)

Fig. 8 DSMC, NS, and MPC surface properties for the hollow-cylinder flare.

with that predicted by full DSMC. The coefficient of heat transfer, also referred to as the Stanton number [$St = q / (0.5\rho_\infty u_\infty^3)$], where q represents the amount of energy transferred to the surface per unit time per unit area, is plotted in Fig. 8b for each simulation method. Here, the MPC method is seen to improve the prediction of heat transfer inside the recirculation region, compared with the NS prediction. Because the extent of flow recirculation is reduced into better agreement with experimental data, the location of peak heating predicted by the MPC method is also shifted upstream into excellent agreement with full DSMC and into better agreement with experimental measurement when compared with the NS solution. Furthermore, the heating rate near the leading edge of the hollow cylinder is shown in Fig. 9. Here, the NS equations predict an infinite heating rate due to the no-slip boundary condition, whereas a full DSMC simulation predicts a finite heating rate because DSMC naturally allows for velocity slip. As seen in Fig. 9, the leading edge is a region of high heating and the MPC method is able to reproduce the finite heating rate predicted by full DSMC simulation. This is expected, because the MPC method is shown in Fig. 6a to accurately capture velocity slip, temperature jump, and thermal nonequilibrium near the leading edge.

An interesting result is obtained when the *original* Steger–Warming flux vector-splitting approach [18,19] is used exclusively within LeMANS and therefore within the MPC method in continuum regions. Figure 10 demonstrates that the numerical dissipation inherent in the original Steger–Warming flux function results in a significant (27%) overprediction in the peak heating rate resulting from a full NS (LeMANS) simulation. Because the MPC method uses the DSMC method along the entire cylinder-flare surface, the excess numerical dissipation is contained within low-gradient continuum regions. As seen in Fig. 10, MPC simulation using the highly dissipative original Steger–Warming flux function still reproduces the heating rates predicted by full DSMC simulation and measured experimentally.

C. Computational Efficiency

As mentioned earlier, each MPC simulation begins with a NS solution on a mesh designed for the NS equations. Mesh refinement within DSMC regions during an MPC simulation is automatic. Thus, the setup time required for an MPC simulation is identical to that required for a full NS solution. For these reasons, the setup time required for DSMC, NS, and MPC simulations is not included in the timing comparisons. In addition, the full DSMC simulations described in this paper are run in parallel on a cluster of CPUs,




Fig. 9 DSMC, NS, and MPC leading-edge heating rates.




Fig. 10 DSMC, NS, and MPC heating rates using the original Steger-Warming flux function.

whereas all NS and MPC simulations were performed in serial on a single CPU. Thus, the parallel efficiency of the DSMC code becomes a second factor that is included approximately into the timing comparisons that follow.

The time required for a full DSMC simulation is taken as the time required to reach steady state plus the time required to sample the solution. The time required by an MPC simulation consists of the time to obtain the initial NS solution, plus the time required to reach step 4 in Sec. III.B (MPC steady state), plus the time required to sample and converge DSMC and NS regions (step 4). To remain consistent, each MPC solution is sampled for the same physical time and number of time steps as the corresponding full DSMC solution. The performance parameters of most interest are listed in Table 2 for each case and are now briefly described. The most practical

parameters include the ratio of MPC to full DSMC simulation times (the *speedup* factor) and the ratio of *memory* (RAM) requirements. The largest contributing factor to the speedup is the fact that the time required by a DSMC simulation scales directly with the number of particles used in the simulation. Therefore, if all hybrid NS operations (initial NS solution and NS updates) within the MPC cycle take a negligible time when compared with that spent simulating hybrid DSMC regions, then the speedup factor achieved by the MPC method will scale directly with the number of particles eliminated (replaced with a continuum description) by the MPC simulation. This ratio of *particles* used by the MPC method to that used by DSMC is also included in Table 2. To show the relative cost of hybrid NS operations, the ratio of time required for *initial NS* simulation compared with the full DSMC simulation and the ratio of the *time step* used in NS regions compared with that used in DSMC regions are included in Table 2. Finally, the time required for initial NS simulation plus the time required for the MPC simulation to reach steady state is often less than that required by a full DSMC simulation to reach steady state. Because both methods are sampled for the same time, an MPC simulation often requires fewer total iterations, which also contributes to the speedup factor. In other words, not only does an MPC simulation require fewer particles, but due to the efficiency of the initial NS solution, it may also reach steady state faster than full DSMC. The ratio of iterations required for the MPC method to reach *steady state* (after obtaining the initial NS solution) to that required by full DSMC simulation is also included in Table 2 for each case.

MPC simulation of the hollow-cylinder flare is seen to use 53% of the particles and 80% of the memory required by a full DSMC simulation. In addition, despite reaching steady state much faster than full DSMC, the MPC method still requires 71% of the time needed to obtain a full DSMC solution (a speedup factor of 1.4). This signifies that the time required for the initial NS solution and NS operations within the MPC cycle are not negligible when compared with the time spent simulating DSMC regions. This is not surprising because, as mentioned in Sec. II, NS simulation of the hollow-cylinder-flare problem (run-11 conditions) requires a mesh resolution similar to that of DSMC and the size of particle and continuum regions are roughly equal. It should be noted that a hybrid simulation of a much larger continuum flow over an entire vehicle in which a shock-interaction region is embedded would result in a much larger speedup factor for the MPC method, compared with full DSMC simulation. Finally, larger speedup factors would result from MPC simulation of hypersonic blunt-body flows that involve extreme variations in λ between compressed forebody and rarefied wake regions. Such DSMC simulations become computationally demanding, especially for three-dimensional flows.

V. Conclusions

1) In comparison with previous MPC research, application of the method to the hollow-cylinder-flare problem demonstrates new essential capabilities that include the ability of MPC interfaces to track moving flow features and the ability of the MPC cycle to significantly alter NS boundary conditions that then shift the solution in NS regions of the simulation. An important result is that these capabilities are demonstrated for a loosely coupled approach that is able to maintain spatial- and temporal-scale decoupling throughout the simulation of a highly nonlinear flowfield. Additional capabilities involve variation in mesh density between continuum and particle regions and the ability to handle axisymmetric flows.

2) For flow over the hollow-cylinder-flare geometry, an MPC simulation initialized with a NS solution is able to transition away from the no-slip NS solution and reproduce the velocity slip,

Table 2 Computational efficiency of the MPC method

Hollow-cyl. flare	Speedup	Memory	Particles	Initial NS	Time step	Steady state	Sampling
DSMC	100%	100%	100%	N/A	1	100%	100%
MPC	71% (1.4 \times)	80%	53% (1.9 \times)	4%	28	20%	100%

temperature jump, thermal nonequilibrium, and surface properties predicted by full DSMC near the leading edge and along the cylinder surface. MPC simulation also successfully reduces the oversized separation region predicted by the NS equations to the size predicted by DSMC, which agrees more closely with experimental results.

3) Because the MPC method uses kinetic (DSMC) simulation in high-gradient regions, the use of a highly dissipative numerical flux function within the continuum solver does not affect the accuracy of predicted surface properties. This conclusion lends credibility to the concept of algorithm refinement, where a more physically based (kinetic) method is used to simulate nonequilibrium regions instead of attempting to locally enhance the continuum solution method.

4) The MPC method is able to produce highly accurate results approximately 1.4 times faster than full DSMC simulation while using 80% of the memory for the conditions under consideration. The fine resolution required to solve the NS equations for the hollow-cylinder-flare problem combined with the relative size of the nonequilibrium flow region limit the speedup factor achieved by the MPC method. If this hypersonic interaction flow were embedded as part of a much larger continuum flow about an entire vehicle, the MPC method would have the potential to offer significant computational savings over full DSMC simulation.

Acknowledgments

This work, performed at the University of Michigan, is sponsored by the Space Vehicle Transportation Institute under NASA grant NCC3-989 with joint sponsorship from the U.S. Department of Defense and from the U.S. Air Force Office of Scientific Research grant FA9550-05-1-0115. This work is also supported by the Francois-Xavier Bagnoud Foundation.

References

- [1] Moss, J. N., "Direct Simulation Monte Carlo Simulations of Ballute Aerothermodynamics Under Hypersonic Rarefied Conditions," *Journal of Spacecraft and Rockets*, Vol. 44, No. 2, 2007, pp. 289–298. doi:10.2514/1.22706
- [2] Gnoffo, P. A., Buck, G., Moss, J. N., Nielsen, E., Berger, K., Jones, W. T., and Rudavsky, R., "Aerothermodynamic Analyses of Towed Ballutes," 9th AIAA/ASME Joint Thermophysics and Heat Transfer Conference, San Francisco, AIAA Paper 06-3771, June 2006.
- [3] Bird, G. A., *Molecular Gas Dynamics and the Direct Simulation of Gas Flows*, Oxford Univ. Press, New York, 1994.
- [4] Hash, D. B., and Hassan, H. A., "Assessment of Schemes for Coupling Monte Carlo and Navier-Stokes Solution Methods," *Journal of Thermophysics and Heat Transfer*, Vol. 10, No. 2, 1996, pp. 242–249.
- [5] Hash, D. B., and Hassan, H. A., "A Decoupled DSMC/Navier-Stokes Analysis of a Transitional Flow Experiment," 34th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, AIAA Paper 96-0353, Jan. 1996.
- [6] Roveda, R., Goldstein, D. B., and Varghese, P. L., "Hybrid Euler/Particle Approach for Continuum/Rarefied Flows," *Journal of Spacecraft and Rockets*, Vol. 35, No. 3, 1998, pp. 258–265.
- [7] Roveda, R., Goldstein, D. B., and Varghese, P. L., "Hybrid Euler/Direct Simulation Monte Carlo Calculation of Unsteady Slit Flow," *Journal of Spacecraft and Rockets*, Vol. 37, No. 6, 2000, pp. 753–760.
- [8] Wang, W. L., and Boyd, I. D., "Hybrid DSMC-CFD Simulations of Hypersonic Flow over Sharp and Blunted Bodies," 36th AIAA Thermophysics Conference, Orlando, FL, AIAA Paper 03-3644, 2003.
- [9] Wijesinghe, H. S., Hornung, R. D., Garcia, A. L., and Hadjiconstantinou, N. G., "Three-Dimensional Hybrid Continuum-Atomistic Simulations for Multiscale Hydrodynamics," *Journal of Fluids Engineering*, Vol. 126, No. 5, 2004, pp. 768–777. doi:10.1115/1.1792275
- [10] Wu, J. S., Lian, Y. Y., Cheng, G., Koomullil, R. P., and Tseng, K. C., "Development and Verification of a Coupled DSMC-NS Scheme Using Unstructured Mesh," *Journal of Computational Physics*, Vol. 219, No. 2, 2006, pp. 579–607. doi:10.1016/j.jcp.2006.04.013
- [11] Schwartzentruber, T. E., Scalabrin, L., and Boyd, I. D., "A Modular Particle-Continuum Numerical Method for Hypersonic Nonequilibrium Gas Flows," *Journal of Computational Physics*, Vol. 225, No. 1, 2007, pp. 1159–1174. doi:10.1016/j.jcp.2007.01.022
- [12] Schwartzentruber, T. E., and Boyd, I. D., "A Hybrid Particle-Continuum Method Applied to Shock Waves," *Journal of Computational Physics*, Vol. 215, No. 2, 2006, pp. 402–416. doi:10.1016/j.jcp.2005.10.023
- [13] Schwartzentruber, T. E., Scalabrin, L. C., and Boyd, I. D., "Hybrid Particle-Continuum Simulations of Nonequilibrium Hypersonic Blunt-Body Flowfields," *Journal of Thermophysics and Heat Transfer*, Vol. 22, No. 1, 2008, pp. 29–37. doi:10.2514/1.30216
- [14] Schwartzentruber, T. E., Scalabrin, L. C., and Boyd, I. D., "Modular Implementation of a Hybrid DSMC-NS Solver for Hypersonic Non-Equilibrium Flows," 45th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, AIAA Paper 07-613, Jan. 2007.
- [15] Dietrich, S., and Boyd, I. D., "Scalar and Parallel Optimized Implementation of the Direct Simulation Monte Carlo Method," *Journal of Computational Physics*, Vol. 126, No. 2, 1996, pp. 328–342. doi:10.1006/jcph.1996.0141
- [16] Boyd, I. D., "Analysis of Rotational Nonequilibrium in Standing Shock Waves of Nitrogen," *AIAA Journal*, Vol. 28, No. 11, 1990, pp. 1997–1999.
- [17] Scalabrin, L. C., and Boyd, I. D., "Numerical Simulation of Weakly Ionized Hypersonic Flow for Reentry Configurations," 9th AIAA/ASME Joint Thermophysics and Heat Transfer Conference, San Francisco, AIAA Paper 06-3773, 2006.
- [18] MacCormack, R. W., and Candler, G. V., "The Solution of the Navier-Stokes Equations Using Gauss-Seidel Line Relaxation," *Computers and Fluids*, Vol. 17, No. 1, 1989, pp. 135–150. doi:10.1016/0045-7930(89)90012-1
- [19] Druguet, M. C., Candler, G. V., and Nompelis, I., "Effect of Numerics on Navier-Stokes Computations of Hypersonic Double-Cone Flows," *AIAA Journal*, Vol. 43, No. 3, 2005, pp. 616–623. doi:10.2514/1.6190
- [20] Holden, M. S., and Wadhams, T. P., "Code Validation Study of Laminar Shock/Boundary Layer and Shock/Shock Interactions in Hypersonic Flow, Part A: Experimental Measurements," AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, AIAA Paper 01-1031, Jan. 2001.
- [21] Harvey, J. K., Holden, M. S., and Wadhams, T. P., "Code Validation Study of Laminar Shock/Boundary Layer and Shock/Shock Interactions in Hypersonic Flow, Part B: Comparison with Navier-Stokes and DSMC Solutions," AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, AIAA Paper 01-1031, Jan. 2001.
- [22] Candler, G. V., Nompelis, I., and Druguet, M. C., "Navier-Stokes Predictions of Hypersonic Double-Cone and Cylinder-flare Flow Fields," 39th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, AIAA Paper 01-1024, Jan. 2001.
- [23] Gnoffo, P. A., "CFD Validation Studies for Hypersonic Flow Prediction," 39th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, AIAA Paper 01-1025, Jan. 2001.
- [24] Moss, J. N., "DSMC Computations for Regions of Shock/Shock and Shock/Boundary Layer Interaction," 39th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, AIAA Paper 01-1027, Jan. 2001.
- [25] Kato, H., and Tannehill, J. C., "Computation of Hypersonic Laminar Separated Flows using an Iterated PNS Algorithm," 39th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, AIAA Paper 01-1028, Jan. 2001.
- [26] Boyd, I. D., and Wang, W.-L., "Monte Carlo Computations of Hypersonic Interacting Flows," 39th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, AIAA Paper 01-1029, Jan. 2001.
- [27] Roy, C. J., Gallis, M. A., and Payne, J. L., "DSMC and Navier-Stokes Predictions for Hypersonic Laminar Interacting Flows," 39th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, AIAA Paper 01-1030, Jan. 2001.
- [28] Candler, G. V., Nompelis, I., Druguet, M. C., Holden, M. S., Wadhams, T. P., Boyd, I. D., and Wang, W. L., "CFD Validation for Hypersonic Flight: Hypersonic Double-Cone Flow Simulations," 40th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, AIAA Paper 02-0581, Jan. 2002.
- [29] Moss, J. N., and Bird, G. A., "Direct Simulation Monte Carlo Simulations of Hypersonic Flows with Shock Interactions," *AIAA Journal*, Vol. 43, No. 12, 2005, pp. 2565–2573. doi:10.2514/1.12532
- [30] Schwartzentruber, T. E., Scalabrin, L. C., and Boyd, I. D., "Hybrid Particle-Continuum Simulations of Low Knudsen Number Hypersonic Flows," 39th AIAA Thermophysics Conference, Miami, FL, AIAA Paper 07-3829, June 2007.
- [31] Boyd, I. D., Chen, G., and Candler, G. V., "Predicting Failure of the Continuum Fluid Equations in Transitional Hypersonic Flows," *Physics of Fluids*, Vol. 7, No. 1, 1995, pp. 210–219.

- [32] Garcia, A. L., and Alder, B. J., "Generation of the Chapman–Enskog Distribution," *Journal of Computational Physics*, Vol. 140, 1998, pp. 66–70.
doi:10.1006/jcph.1998.5889
- [33] Sun, Q., and Boyd, I. D., "Evaluation of Macroscopic Properties in the Direct Simulation Monte Carlo Method," *Journal of Thermophysics and Heat Transfer*, Vol. 19, No. 3, 2005, pp. 329–335.
doi:10.2514/1.12542